WHITE PAPER

SPARC Servers: An Effective Choice for Efficiency in the Datacenter

Sponsored by: Oracle

Jean S. Bozman Matthew Eastwood

April 2012

EXECUTIVE SUMMARY

Limited IT budgets and rising operational costs are combining to create a difficult situation for CXOs, IT managers, and datacenter managers worldwide. Efficiency is the top priority for these managers — and technologies and business practices are being tapped to improve efficiency wherever possible.

In fact, datacenter efficiency needs to be continuously optimized in terms of people, time, money, processes, and infrastructure — and the use of datacenter "real estate." To support those goals, IT is harnessing a number of technical approaches, including virtualization, workload consolidation, dense computing, new form factors, and efficient management practices that avoid server sprawl — including physical server sprawl and virtual machine (VM) sprawl (virtual servers).

Oracle's SPARC servers have been designed to support virtualization, workload consolidation, and comprehensive and highly granular management. These capabilities are built on a suite of technologies that have been available on the SPARC and Oracle Solaris platform since the 1990s.

But now, there is a new opportunity for IT managers to evaluate their datacenter efficiency and their ability to support a range of demanding workloads (for example, databases, business processing, and decision support) for the next generation of enterprise applications. Agility can also be added to their IT operations through more rapid application rollout and an easy means to move services around in a highly available, secure, and scalable environment. A new generation of SPARC servers, based on SPARC T-Series processors, provides a wider portfolio of hardware and software features that can be leveraged to dramatically improve efficiency and agility.

These SPARC systems run Oracle Solaris 11 — and can host applications from Oracle Solaris 10 — with all of these applications running inside Oracle Solaris Zones. Oracle Solaris Zones technology enables the combining of applications for improved server utilization, and it supports management with a high degree of controllability and granularity. This means that an organization can consolidate workloads from older servers and then rehost them on highly efficient SPARC T4–based servers for IT flexibility and reduced operational costs.

Business benefits associated with this technical change include more efficient placement of applications within the datacenter, fewer IT staff hours spent managing the older workloads, and improved uptime. This means that overall, fewer steps need to be taken by IT to manage the applications and to run them efficiently with respect to power/cooling and system administration.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

OVERVIEW: HIGHER SYSTEM UTILIZATION LEADS TO DATACENTER EFFICIENCY

Datacenter efficiency is a complex concept involving rightsizing and the continuous optimization of the infrastructure (IT and facilities), people, and processes within the datacenter. In terms of servers and people, one way of increasing efficiency is to increase system utilization. In recent years, this has been accomplished through virtualization and consolidation of systems. However, Unix-based systems have been reaching high levels of utilization for quite a while.

Figure 1 depicts IDC Business Value research data for customers' relative cost experiences with Unix and industry-standard servers. Customers reported that three-year costs per 100 users for acquisition, management, maintenance, and power and cooling of Unix servers were over 25% lower than those same costs for x86-based systems (a total cost of over $52,000 for Unix compared with a total cost of over $77,000 for industry-standard servers). The large win for Unix is in terms of management. The cost of managing a Unix system for 100 users over three years is $8,100 compared with $50,400 for an x86 server. This ease of management is partially due to the increased utilization and scalability of Unix compared with x86. IT managers can consolidate many workloads onto a single system with Unix and experience ease of management, increased availability, and lower overall TCO.

- **Fewer physical systems to manage.** Often, when system utilization is increased, fewer physical systems are needed to run the same workloads. This means less maintenance, less management, and fewer procurement hassles. For example, increasing system utilization by running multiple virtual servers on fewer physical server "footprints" creates multiple positive ripple effects across the datacenter. These downstream efficiencies include simplified management, reduced IT staff time, reduced complexity, and increased uptime.

- **Less personnel time spent managing physical systems.** By deploying fewer servers, system administrators dedicate less of their time to ongoing maintenance. This frees up resources for innovation and contributes back to the business unit in terms of creating new software, increasing integration, and being more responsive to business needs.

- **Decreased power and cooling costs.** Physical servers are the single largest draw of power and the single largest source of cooling costs in the datacenter. For IT organizations, having fewer systems means having fewer systems to power and cool, which results in lower utility bills through energy efficiency. Although many times this savings is not attributed back to the IT group's activities, it results in savings for the overall business. Reduced power consumption and cooling requirements also enable a costly datacenter buildout or retrofit to be delayed. By getting more out of what is already on the datacenter floor and, in some cases, consolidating the server footprint, IT managers will have more room for unexpected capacity needs.

FIGURE 1

Reported Three-Year Costs of Industry-Standard and Unix Environments per 100 Users

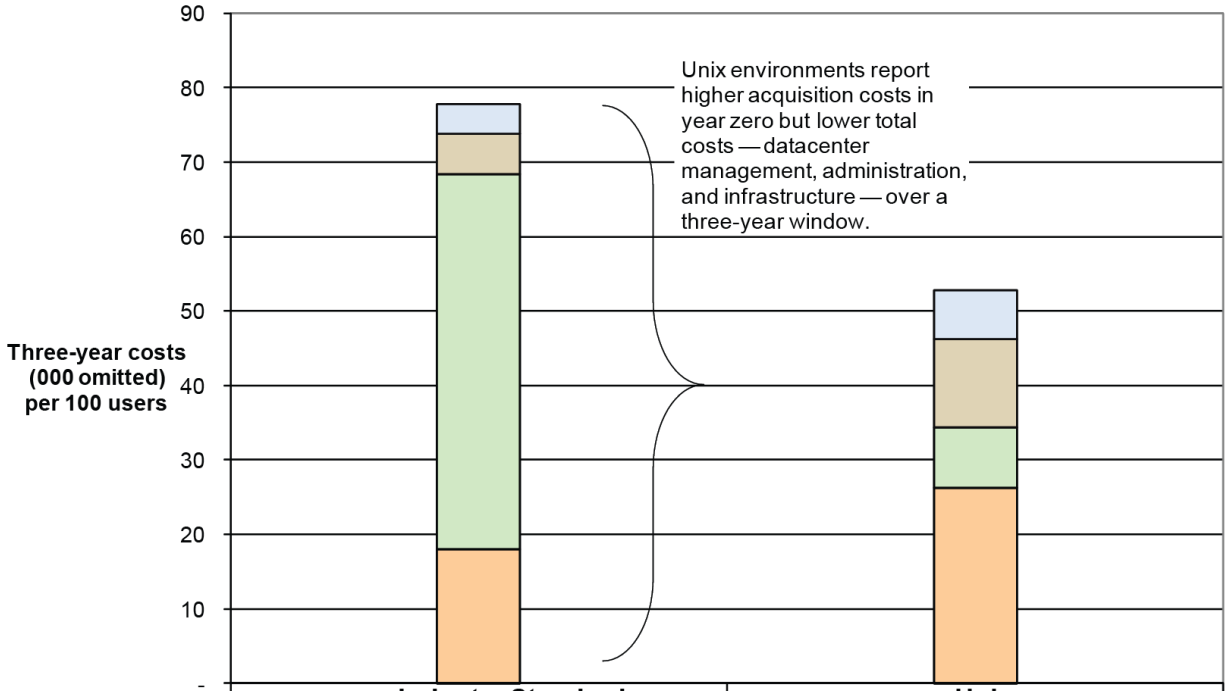

	Industry Standard	Unix
Power and Cool [2]	3.9	6.6
Maintain [2]	5.4	11.8
Manage [2]	50.4	8.1
Acquire [1]	18.0	26.2

[1] One-time acquisition costs occur in year zero.

[2] Average costs are for a three-year period.

- Cost information was collected over a series of engagements measuring overall datacenter costs.
- The term "industry standard" refers to x86 server platforms.

Source: IDC Business Value Research, 2011

Optimizing Virtualized Environments for Datacenter Efficiency

Many times, virtualization is deployed to achieve higher system utilization on Unix-based servers. It is a mature technology that has been deployed since the 1990s. However, server virtualization is creating a management conundrum for many IT organizations. The increase of partitions is difficult to control. Each of the zones, domains, or virtual machines needs to be managed, to some extent, as if it were a physical server. Ultimately, each of these logical systems runs on a physical server that needs adequate power, cooling, patching, monitoring, and management. Without that physical backbone, the logical systems cannot function.

Because of this increase in the number of logical servers, IT administrators must grapple with the consequences of virtual mobility and the inability to track the virtual back to the physical. Virtual machines are being moved around the datacenter and are being abstracted away from the IT hardware (the physical host, networking, and storage) and facilities (power and cooling). As a result, IT administrators spend most of their energy monitoring and managing the software rather than focusing on the underlying IT infrastructure hardware.

This rise in virtual infrastructure leads to increasing costs and confusion in managing virtualized environments. Figure 2 illustrates the climbing costs associated with management and administration of virtual servers through 2013. IDC predicts that, as of 2013, the cost to manage virtual servers will dwarf spending on new server hardware on an annual basis. Accordingly, datacenter managers are looking for new ways to manage their IT environments to combat this proliferation of virtualized servers.

However, IT managers can take a number of steps to ease this transition to a more fully virtualized world. Managing custom images is a task that is much more complicated than managing a set number of standard or "golden" images that can be replicated across the network for optimal performance results. Business units can achieve the benefits that virtualization brings — namely, flexibility in deployment and reduced operational costs — but they must be willing to change their long-held notions about capacity planning and upgrading through "swaps" of physical servers.

FIGURE 2

Worldwide Spending on Servers, Power and Cooling, and Management/Administration

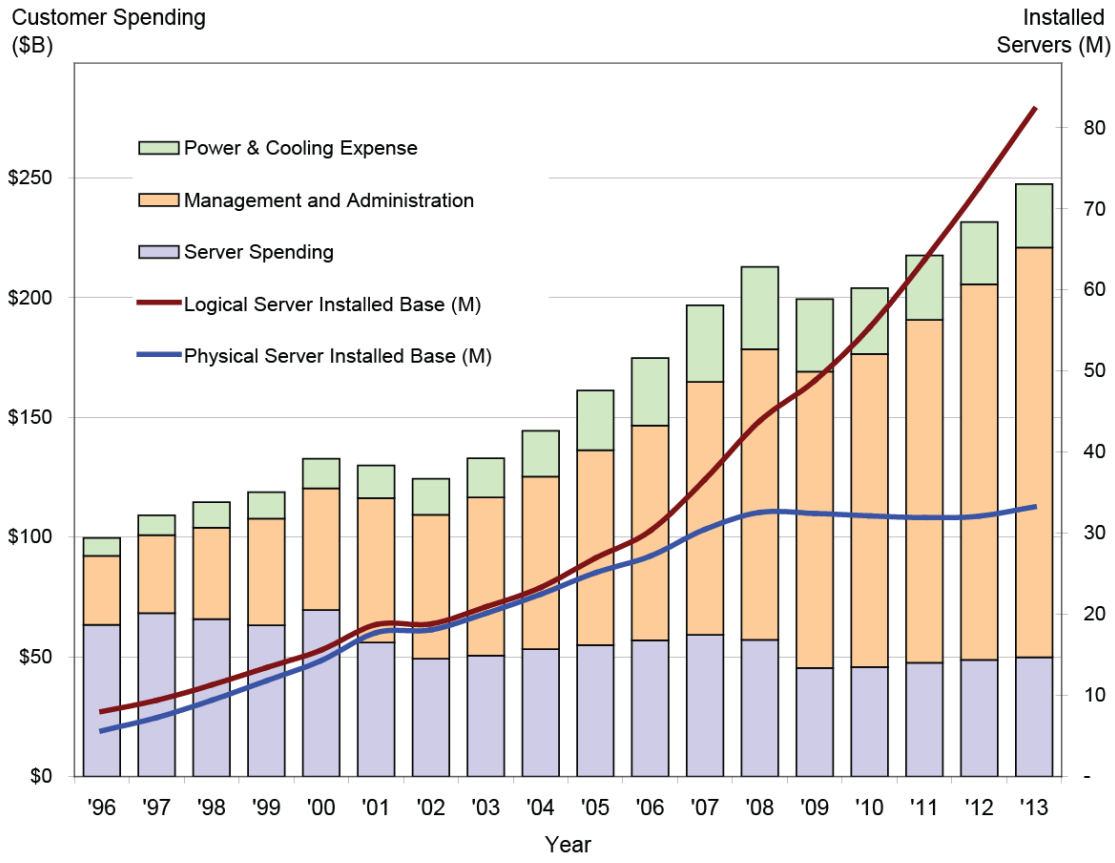

Source: IDC, 2012

Another step IT organizations should consider to save on management and administration costs is represented in Figure 1. Without a reduction in those costs, which have grown dramatically in recent years, the growth in operational costs (opex) will continue to outpace the gains (savings) made in reducing capex for acquired systems.

IDC Business Value customer research shows that large management savings and overall TCO savings can be achieved by using Unix servers to manage multiple system components, including CPUs (processors), operating systems, and associated devices. It is vital for system administrators, IT managers, and datacenter managers to consider the overall TCO, especially the often-overlooked management costs, before choosing server and virtualization solutions. IT organizations should take a holistic approach to datacenter efficiency, considering all the resources in the datacenter, particularly those that are not seen as hard costs, such as time and people costs.

Another challenge for IT is the prospect of increased downtime in the event that a highly virtualized physical server goes offline. A physical server hosting many VMs experiencing downtime is a larger problem for many businesses than a nonvirtualized system that is experiencing downtime. There are more workloads — 6.6 on average in 2010 (IDC's *Server Virtualization Multiclient Study*, November 2010) — on a virtualized physical server than on a nonvirtualized server. However, different types of server platforms can provide higher levels of reliability than others, making reliability, availability, and serviceability (RAS) features and overall availability more important in a virtualized environment than in a nonvirtualized environment.

The least efficient situation is one in which a system is offline for an extended period. During that period of downtime — which could be minutes, hours, or days — resources are consumed with abandon to get the system back up and running. People, money, infrastructure, and time are all factors to consider when calculating the resulting decrease in datacenter efficiency caused by unplanned downtime. To decrease the likelihood that this situation will occur, especially in a highly utilized, virtualized scenario, datacenter managers — during the procurement cycle — increasingly must consider RAS and operating system support for servers. Having high levels of RAS at the hardware layer, along with operating system support for those RAS features, is key to enabling efficiency in the datacenter and avoiding the worst-case scenario — one in which a system goes down and is inaccessible to end users.

Virtualization and Oracle SPARC Servers

Oracle's Sun SPARC servers have supported virtualization for more than 15 years, offering a mature and well-integrated approach to virtualization technology. This means that virtualization is well established on the SPARC platform rather than a new feature that is not yet fully developed. It is ready for use with today's production applications and databases, ensuring that the benefits of virtualization and workload consolidation can be realized. Furthermore, virtualization is included in SPARC servers at no cost to the user. Datacenter managers are able to take advantage of virtualization on SPARC servers without having to worry about the cost of these technologies. SPARC servers support the following virtualization technologies:

- Dynamic Domains technology is a hardware partitioning technology. Available on the SPARC Enterprise M-Series servers, Dynamic Domains have no overhead yet provide full resource, security, fault, and serviceability isolation.
- Oracle VM Server for SPARC (formerly called LDoms) is a hypervisor-based virtualization technology that is available on SPARC T-Series servers and the SPARC SuperCluster T4-4. Oracle VM Server for SPARC supports secure live migration as the virtual machines move from server to server, making failover and availability easier to achieve and improving resource utilization. At this time, live migration is not available for the SPARC T4-4 server when used in the SPARC SuperCluster.
- Oracle Solaris Zones technology (formerly called Oracle Solaris Containers) is an OS partitioning solution that enables thousands of software-isolated partitions to be deployed in a single instance of Oracle Solaris. As a built-in feature of the Oracle Solaris operating system, Oracle Solaris Zones technology is available in Oracle Solaris running on SPARC and x86 systems.

All these SPARC virtualization technologies are provided at no additional cost, as part of the hardware acquisition. Dynamic Domains, Oracle Solaris Zones, and Oracle VM Server for SPARC are all considered "hard partitions" with regard to Oracle per-core software licensing. Customers running Oracle software on SPARC servers need to license only those cores in partitions that are running Oracle software. For deployment of mission-critical applications and services, an additional layer of security and high availability can be provided through the use of Oracle Solaris Cluster. Oracle Enterprise Manager Ops Center can manage all these technologies, including creating, provisioning, and monitoring SPARC partitions. The combination of SPARC, Oracle Solaris, and Oracle Enterprise Manager Ops Center provides a complete virtualization solution on SPARC servers.

Because datacenter managers are "putting many eggs in one basket," an analogy to workload consolidation that puts multiple workloads on the same hardware system,

they need to be more certain than ever that their infrastructure — including physical servers, networking, storage, and power/cooling — is up and running all the time through high availability. Therefore, datacenter managers who embrace virtualization for all its benefits need to have very effective plans and products to ensure that uptime is maintained. They need to mitigate the possible increased risk of having many "eggs in one basket." They need to ensure high availability and to prepare for the worst with realistic disaster recovery plans.

ORACLE'S APPROACH TO DATACENTER EFFICIENCY

The technologies found in Oracle's SPARC servers running Oracle Solaris 11 work at three main levels: hardware, software, and network. When these technologies are deployed together, they can improve operational efficiency in several dimensions, as follows:

- **SPARC servers.** Oracle's SPARC systems range from a single processor blade up to a 64-processor chassis system. Oracle's SPARC server product line now includes the SPARC SuperCluster T4-4, which is a SPARC-based engineered system, including servers, storage, and software, optimized to support a variety of workloads, including transactional workloads, database workloads, and general-purpose workloads. All SPARC servers offer "no-cost virtualization" solutions because the virtualization is embedded in the server itself. As examples, Oracle Solaris Zones technology is available on all SPARC-based Oracle Solaris servers; SPARC Enterprise M-Series servers come with Dynamic Domains partitioning; and SPARC T-Series servers and the SPARC SuperCluster T4-4 server include Oracle VM Server for SPARC.

- **Oracle Solaris.** Oracle Solaris has guaranteed binary compatibility from release to release, which means there is no need to rewrite applications for new operating system releases. Oracle Solaris 10 included native virtualization in the form of Oracle Solaris Zones (formerly called Oracle Solaris Containers), which is a "physical to virtual" (P2V) virtualization technique. This technology allows both Oracle Solaris 8 and 9 environments to be rehosted, making consolidation of applications from older, less energy-efficient systems very simple. It was followed by Oracle Solaris 11, which will preserve binary compatibility for existing Oracle Solaris workloads on SPARC hardware. In all, more than 10,000 ISV applications run on SPARC-based Oracle Solaris systems (even more, if custom applications are included).

- **Oracle Solaris Cluster.** The RAS capabilities of SPARC servers and Oracle Solaris — which combine to make a server resilient — are enhanced by the addition of Oracle Solaris Cluster software. The Oracle Auto Service Request for Sun Systems feature alerts system administrators about system conditions that have the potential to cause system failure, flagging those systems for immediate attention so that the possibility of unplanned downtime can be avoided.

- **Oracle Enterprise Manager Ops Center.** Oracle SPARC-based solutions offer the scale, reliability, and performance needed for mission-critical applications. The Oracle SPARC technologies offer an extremely high thread count and memory density in a small and ecofriendly form factor. Corporations do not wish to exclude those qualities from their cloud foundation. Combined with Oracle Enterprise Manager Ops Center, a SPARC private cloud solution offers a way to leverage existing infrastructure and opens the door to running business-critical applications in a more efficient self-service model. Oracle Enterprise Manager Ops Center drives the deployment, configuration, life cycle, and observability of the SPARC fleet of servers. The Enterprise Manager solution connects the infrastructure layer of the datacenter to the application stack. It offers business-driven metrics into operational telemetry to help customers find the most value in their SPARC investments. Oracle Enterprise Manager Ops Center is now included, at no additional cost, with any server that has an Oracle premier support contract.

SPARC Technologies Enable Flexibility

SPARC technologies help IT managers increase the energy efficiency of older applications by redeploying the applications on new hardware with up-to-date processors. This redeployment can be done without having to rewrite the applications, which translates into savings on developer resources and potential unplanned downtime costs.

SPARC servers have a number of technologies that help IT managers create modular and rightsized IT environments:

- **Binary compatibility.** Oracle's SPARC-based Oracle Solaris systems are binary compatible at the processor instruction set level and at the operating system API level, both from generation to generation and within a generation of SPARC systems. This means that applications that work on earlier generations of SPARC-based Oracle Solaris servers will also work on the latest SPARC-based Oracle Solaris technology. Customers do not need to rewrite applications within the SPARC environment, and older Oracle Solaris applications can be deployed on the latest SPARC systems. This has been an important benefit of deploying on SPARC-based Oracle Solaris systems because it preserves customers' previous investments in custom applications and in older ISV applications that are still in use.

- **Building blocks for upgradeability.** The SPARC Enterprise M-Series servers support the mixing of different speeds and generations of SPARC64 processors. This capability means that IT departments can purchase a larger server that is not fully populated with processors. In this way, the system can grow through the addition of the newest and fastest processors to meet changing business needs over time.

  Importantly, the newer SPARC64 processors coexist with the older SPARC64 processors, and they all run at their native speeds. Thus, SPARC Enterprise M-Series customers can benefit from the latest generation of technology and still use their older processors. This capability is unique in the industry to Oracle's SPARC Enterprise M-Series servers. The processors and memory on the SPARC Enterprise M8000 and SPARC Enterprise M9000 servers are all hot swappable, so more processors and memory can be added easily, with no downtime for applications.

Oracle SPARC Product Families

Oracle has three SPARC-based product lines: the SPARC T-Series, the SPARC Enterprise M-Series, and the new SPARC SuperCluster T4-4. Oracle's SPARC servers promote datacenter efficiency with enterprise-tested technologies. They improve efficiency with respect to scalability, virtualization, high availability, power, cooling, and IT staff time spent on deployment and maintenance. These savings occur at three levels: processor, operating system, and management software. When combined into a single system offering, these technologies improve resource utilization (compared with older technologies), increase processing densities, and supply reliability features that allow computing to proceed.

SPARC T-Series Servers

The SPARC T-Series is now based on the fourth-generation, eight-core SPARC T4 processors. From the beginning, the SPARC T-Series systems were designed to run "cool" — meaning that power and cooling costs could be reduced compared with the power/cooling costs for earlier generations of SPARC systems.

The SPARC T-Series servers use an innovative "system on a chip" architecture in which many formerly separate functions are now part of the processor. Processor cores, cache, security coprocessors, and I/O and networking capabilities are all part of the SPARC T4 processors. SPARC T-Series servers scale from one to four processors and have up to 1TB of RAM.

The SPARC T-Series servers were originally designed to host network-facing workloads, and they are often deployed in telecommunications, financial services, and service provider sites because they can be deployed in dense arrays to support highly parallelized workloads. In the years following the introduction of the SPARC T-Series servers, many of the SPARC T-Series workloads were primarily Web-based workloads. However, with the fourth-generation SPARC T4 processor, SPARC T-Series servers have been optimized to support a much wider range of workloads, including application, batch processing, single-thread, and database workloads. The result is that customers can run a mix of workloads on the SPARC T-Series systems.

SPARC SuperCluster T4-4

The SPARC SuperCluster T4-4 server is based on SPARC T4-4 compute nodes. It is an Oracle engineered system in which the compute nodes, storage nodes, networking, and software are tightly integrated to provide high performance, availability, and simplified deployment. The compute nodes are SPARC T4-4 servers, and the storage nodes are optimized for performance with Oracle Database 11g Release 2. SPARC SuperCluster machines have the same virtualization capabilities as standalone SPARC T4 servers, supporting Oracle Solaris Zones and Oracle VM Server for SPARC; they can run Oracle Solaris 10 and Oracle Solaris 11. The tight integration of storage, computing, networking, and software inside a SPARC SuperCluster engineered system can reduce application deployment time, which increases datacenter efficiency.

Acquiring the integrated system means that customers can avoid doing onsite system integration, saving money and time. In this way, they can bypass many time-consuming tasks associated with integrating servers, storage, networking, and software because this work has already been done. In addition, the single-pane-of-glass management window in Oracle Enterprise Manager greatly reduces the amount of time that a system administrator would spend to cover whole-rack management.

SPARC Enterprise M-Series Servers

The SPARC Enterprise M-Series consists of scalable servers based on SPARC64 processors. These servers are jointly manufactured by Oracle and Fujitsu Ltd. through a longstanding relationship between the two companies that began in the 1990s between Oracle's precursor, Sun Microsystems, and Fujitsu and has continued after Oracle's acquisition of Sun. These processors have been optimized for single-thread performance, which is why SPARC Enterprise M-Series systems are most often deployed to run highly demanding enterprise workloads.

Today, the SPARC Enterprise M-Series servers are known for their support of very large databases and line-of-business (LOB) workloads, such as OLTP, ERP, CRM, and other types of highly stateful applications. These applications benefit from running on scalable servers — with capacities ranging from 1 quad-core processor to 64 quad-core processors.

The SPARC Enterprise M-Series servers are also used to support highly scalable databases and data warehouses, along with business intelligence workloads. The larger SPARC Enterprise M-Series systems do not need to be shipped fully populated with SPARC processors; rather, they can be populated as needed. This flexibility in deployment decreases energy consumption and reduces capital expenditures that would otherwise be incurred by running processors that are not in active use.

Oracle Solaris Supports Datacenter Efficiency

Through key technologies for server virtualization and advanced system management, Oracle Solaris and SPARC servers improve datacenter efficiency.

These technologies — Oracle Solaris Zones, Dynamic Domains, Oracle VM Server for SPARC, Oracle Enterprise Manager Ops Center, and Oracle Solaris Cluster — cover both software and hardware and both people and system efficiency in the datacenter. They combine easily, are designed for investment protection, are scalable, and enable increased resource utilization in datacenters.

- **Cloud-ready virtualization for servers, storage, and networking.** Oracle Solaris 11 provides greater agility and flexibility in deploying cloud-ready datacenter infrastructure through its deep support for virtualization, which now includes built-in network virtualization. As a fully virtualized operating system, Oracle Solaris 11 maps to processor, network, and storage devices, allowing secure hosting of multiple virtualized workloads on a single system or engineered system.

- **Efficient use of datacenter "real estate."** New installations of Oracle Solaris 11 and associated software management capabilities further improve efficiency when deploying many hundreds of systems in a typical datacenter, introducing automation that reduces boot time and provisioning time. The portfolio of ISV applications for Oracle Solaris includes more than 10,000 software application titles. These applications and all their configuration details can run directly on Oracle Solaris 11, or — by using built-in virtualization capabilities — they can run entire Oracle Solaris 10 environments within Oracle Solaris Zones. Dozens, or even hundreds, of Oracle Solaris Zones can be hosted within a single Oracle Solaris 11 instance. Custom applications also benefit from binary compatibility, but it should be noted that additional performance and, therefore, greater overall efficiency might be gained by recompiling the applications with the latest Oracle Solaris Studio tools and, where necessary, recoding to take advantage of new operating system capabilities.

- **RAS features.** Building on the hardware RAS features in each individual SPARC system, high-availability software scales the solution out to include multiple SPARC-based systems, including multisite high availability and disaster recovery, in highly secure computing environments that isolate individual workloads. Due to isolation, a single application can be prevented from causing systemwide downtime. Oracle Solaris Predictive Self Healing protects uptime by identifying pending faults and taking steps that proactively avoid unplanned downtime. In cases of unplanned downtime, the self-healing capability brings systems back up quickly to ensure high service levels.

- **High degree of granularity, customization, and control.** Dozens of Oracle Solaris Zones can be deployed in one instance of the OS, which minimizes the effort required to upgrade all the Oracle Solaris Zones. An Oracle Solaris Zone can be configured with its own network and storage resources. Its management is delegated to the "owner" of the application that is hosted in the zone, supporting high levels of granular controllability for workloads. Oracle Solaris Zones are a good fit for consolidating workloads deployed across multiple servers, including custom applications, because of their binary compatibility.

- **Complete security and resource isolation.** Dynamic Domains are meant for consolidation on the SPARC Enterprise M-Series systems. They provide disaster recovery support for memory, processor, and I/O resources. Dynamic domains have several levels of isolation and have a granularity of one processor socket. They are ideal for hosting complex, mission-critical applications that need complete security and resource isolation.

- **Dynamically optimized workloads.** Oracle VM partitions are used to consolidate various workloads, such as applications and databases, onto SPARC T-Series servers and SPARC SuperCluster T4-4 servers. Except with the SPARC SuperCluster T4-4 server, Oracle VM Server for SPARC also supports live migration between physical hosts as an IT organization's deployment needs change over time.

- **Highly effective disaster recovery.** Oracle Solaris Cluster adds highly effective disaster recovery capabilities to these virtualization and management technologies. It even adds further reliability to Oracle Solaris and Oracle Real Application Clusters (Oracle RAC), delivering accurate failover between physical or virtual environments, across the datacenter, or at great distances — across continents — to ensure business continuity for large datacenters.

Oracle Solaris Zones

Oracle Solaris Zones technology is a software-based partitioning technology that enables thousands of partitions to be created on a single instance of Oracle Solaris. Oracle Solaris Zones are available on any server running Oracle Solaris, not just on SPARC servers. Each Oracle Solaris Zone can be thought of as a secure virtual machine with its own allocation of virtualized network and storage resources (each of which can be encrypted). At the same time, management of an Oracle Solaris Zone is totally isolated from management of the physical server in the global zone. Many attributes of Oracle Solaris Zones enable the setting of minimum and maximum resource thresholds (via parameters) to support the appropriate number of end users, applications, and processes. These attributes simplify the consolidation of workloads on SPARC servers and provide enormous flexibility by allowing dynamic reallocation of the processor and memory resources.

Customers have long valued the ability to bring their custom or packaged software applications forward to each new Oracle Solaris release — and to run those software environments on new SPARC platforms. In cases where Oracle's binary compatibility guarantee is not sufficient, application environments from Oracle Solaris 8 or Oracle Solaris 9 systems can be run in Oracle Solaris Zones on Oracle Solaris 10 systems. Similarly, this P2V migration can be used to rehost Oracle Solaris 10 environments onto systems running Oracle Solaris 11. This capability makes it possible to consolidate legacy applications on newer and more efficient SPARC servers, thereby improving the efficiency of the entire datacenter. Datacenter managers operating in an Oracle Solaris environment can upgrade to the latest release of Oracle Solaris without fear of leaving stranded applications behind.

Importantly, a high level of granular control is supported because each individual Oracle Solaris Zone can be managed separately. For example, one Oracle Solaris Zone can be "failed over" to another through the use of high-availability software, or Oracle Solaris Zones can be retired from use, with the contents of the zones replicated to other zones, as needed. Running all the Oracle Solaris Zones under a single copy of Oracle Solaris enhances the manageability of the software environment and enables system administrators in the datacenter to work more efficiently.

In a SPARC-based Oracle Solaris environment, datacenter managers can virtualize their computing environment by partitioning one instance of Oracle Solaris into several Oracle Solaris environments with separate security environments and applications on each. This makes maintenance of each instance much simpler than it would be if each workload were running on a separate physical server. Importantly, once this takes place, much of the upgrading activity happens at the Oracle Solaris "instance" level (within a copy of Oracle Solaris) rather than at the partition level. As a result, an organization can avoid the phenomenon of box-swapping that occurs when it is rightsizing the workloads that run on physical servers.

Dynamic Domains

SPARC Enterprise M-Series servers offer a hardware partitioning technology called Dynamic Domains. Each Dynamic Domain has its own operating system instance, resulting in the partitioning of SPARC Enterprise M-Series servers into multiple virtual servers. Dynamic Domains have no overhead, meaning that there is no reduction in application performance. Dynamic Domains can be as small as one processor and as large as 64 processors, providing a high degree of scalability. They support complete isolation of resources, service, faults, and security access — all of which enables a high degree of workload consolidation. Dynamic Domains can be resized while an application is running. Processor, memory, and I/O resources can be moved from one domain to another to meet changing workload requirements. Oracle Solaris Zones can be deployed inside Dynamic Domains to provide additional partition granularity.

Oracle VM Server for SPARC

SPARC T-Series servers and SPARC SuperCluster servers offer advanced hypervisor-based partitioning that comes standard with every system. The logical domains in Oracle VM Server for SPARC environments are simple to deploy and to manage because of the built-in hypervisor. In addition, administrators can use Oracle Solaris advanced technologies, including ZFS cloning and "snapshots" that decrease disk requirements. For the purpose of availability, these environments can be built with redundancy at both the network level and the disk level.

For improved datacenter efficiency, Oracle VM Server for SPARC enables dynamic workload optimization to meet changing business demands. This saves datacenter managers resources in terms of infrastructure, people, process, time, and money. For configurations other than SPARC SuperCluster configurations, workloads can be easily migrated between systems, when necessary, with little latency in overall performance. Datacenter efficiency is improved by making change easier and ongoing operations more efficient.

Oracle Solaris Cluster: High Availability and Disaster Recovery for Traditional and Virtualized Environments

Oracle Solaris Cluster supports and is pretested with Oracle's SPARC systems on all of the previously mentioned virtualization technologies.

Oracle Solaris Zones and Oracle VM Server for SPARC can be protected by Oracle Solaris Cluster from outages caused by hardware and software failure. Both the VM and the applications running inside the VM can be monitored and moved around within the physical cluster, whether this occurs for planned downtime (for example, planned repairs) or for unplanned downtime caused by outages (for example, network outage, power outage, or natural disaster).

Additional protection and isolation is built into an Oracle Solaris Zones cluster: It is possible to create virtual clusters that are based on Oracle Solaris Zones. These virtual clusters are created by combining Oracle Solaris Zones that are running on different physical servers within an Oracle Solaris Cluster; the independent operation and management of each of these zones effectively isolates the workloads, preserving business continuity.

Adding a dimension of disaster recovery capabilities, the Oracle Solaris Cluster Geographic Edition feature enables customers to fail over applications from a primary geographic site to a secondary site. This works across unlimited distances (working over LANs or WANs), ensuring continuity of business services in the event of a localized disaster within any given geography. Together with the virtualization technologies previously discussed, these capabilities provide protection for important workloads in a way that the use of virtualization alone cannot.

Oracle Enterprise Manager Ops Center: Achieving Management Efficiencies

Oracle Enterprise Manager Ops Center is Oracle's virtualization, energy, and systems management solution. It is tightly integrated with Oracle Enterprise Manager, which offers the unique ability to manage Oracle databases, middleware, and applications. Oracle Enterprise Manager Ops Center is designed to optimize datacenter efficiency by enabling the management of assets at enterprise scale. Key attributes of Oracle Enterprise Manager Ops Center include the discovery, provisioning, and monitoring of physical and virtual servers. For system administrators, Oracle Enterprise Manager Ops Center's "single pane of glass" capabilities lower the cost of owning multiple SPARC servers. In addition, Oracle Enterprise Manager Ops Center can manage multiple x86 operating systems, including Oracle Linux, Oracle VM, Oracle Solaris, and non-Oracle operating systems. This decreases the number of tools needed and gets system administrators closer to leveraging IT resources to solve business problems.

Energy Management with Oracle Enterprise Manager Ops Center

System managers can dynamically tune SPARC servers from performance mode to energy-saving mode to optimize their consumption of energy. Oracle Enterprise Manager Ops Center unlocks these features at the datacenter level.

Oracle Enterprise Manager Ops Center allows the server administrator to switch all SPARC T-Series servers from performance mode to the energy-saving elastic mode, enabling server administrators to rapidly manipulate the energy consumption of hundreds of servers. With these features, management can align server demand with changes in energy policy.

Additionally, Oracle Enterprise Manager Ops Center shows the connection between operating system performance and server energy consumption, providing the administrator with better information about performance and energy consumption.

Oracle Enterprise Manager Ops Center can also help identify the hot spots in the datacenter, showing trends in the in/out air temperatures and mapping those temperatures to fan RPM speeds. With this information, administrators can find and reposition servers that are working too hard to cool themselves due to rack placement.

This power and cooling analysis can decrease the probability of downtime due to inadequate use or misuse of datacenter resources.

Virtual IT Discovery with Oracle Enterprise Manager Ops Center

In a virtualized environment, Oracle Enterprise Manager Ops Center has the capability to find virtual servers in the datacenter and tie them, and the corresponding storage and network infrastructure, back to the physical server. The ability to visualize where the virtual servers are located — and how they are being utilized — provides system administrators with a powerful utility that helps them manage server utilization.

Oracle Enterprise Manager Ops Center allows for multiple SPARC servers to be placed into a server pool that operates under defined placement (preferred virtual host for new guests) and balancing (relative workload demands on hosts) policies. Based on utilization, Oracle Enterprise Manager Ops Center supports efficient mapping of the workloads across the hardware infrastructure by moving the virtual servers from one physical system to the next until the grouping is optimized for efficiency in terms of people, processes, infrastructure, and energy.

CHALLENGES AND OPPORTUNITIES

Challenges in the Marketplace

With the rapid growth of x86 server deployments, the benefits of SPARC server deployments are not as widely known and understood as they were in the 1990s, when SPARC servers were first ramping in volume in the datacenter. And yet, today, these systems are widely deployed in many industries, including financial services, telecommunications, government, healthcare, and retail, among other vertical markets. This is one reason why IT managers should become more familiar with this Oracle computing environment as part of a wider effort to manage systems more efficiently across the entire datacenter — including x86 and non-x86 systems.

At the same time, because of the competition in the marketplace within the Unix server segment, customers will need to evaluate the product features and capabilities that differentiate Oracle's SPARC-based servers from those of its top Unix system

competitors in terms of support for virtualization, manageability, and operational efficiency. It is also worth noting that customers can consider running Oracle Solaris on their SPARC and x86 servers so that Solaris applications can span both platforms.

Opportunities

Now that Oracle has completed its acquisition of Sun Microsystems, it is able to leverage Sun's deep inventory of hardware and software products and technologies in customer deployments of its products. Oracle has the opportunity to offer the benefits of rapid deployment and efficient operations that are associated with integrated technology stacks from Oracle. This includes Oracle's SPARC hardware and the Oracle Solaris operating system, along with the Oracle Database 11*g* product, Oracle Java-based middleware, Oracle applications (for example, Oracle Financials), and other Oracle software products and services.

Although Oracle has clearly described the advantages of its integrated technology stacks, the company should continue to demonstrate the IT benefits and business benefits of SPARC-based Oracle Solaris deployments through customer case studies, public demonstrations (for example, trade shows and Webcasts), and before/after comparisons of results following deployments of new SPARC-based Oracle Solaris systems. All these approaches should result in a raised awareness about these technologies at existing customer sites and those of prospective customers.

CONCLUSION

Oracle's SPARC servers promote datacenter efficiency in multiple dimensions — with an impact on energy efficiency, datacenter space utilization efficiency, and operational efficiency.

As the datacenter evolves to support more network-centric and Web-centric workloads — in addition to traditional enterprise workloads — achieving these efficiencies will become increasingly important to achieving overall efficiency goals for the business and to meeting IT budgetary constraints.

As this paper describes, SPARC-based systems address efficiency at the processor level, the network level, the systems level, the operating system and the virtualization level, and the management level. Taken together, these efficiency considerations have the capacity to impact IT budgets and to affect the IT time horizon for near-term and long-term planning and implementation of new systems.