IDC
Analyze the Future

WHITE PAPER

Enabling Customer Choice: Sun Mainstreams Server Virtualization

Sponsored by: Sun Microsystems

John Humphreys

July 2006

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

EXECUTIVE SUMMARY

The rapid spread of software-based server virtualization on x86 platforms over the past few years is making virtualization an important solution for hardware resource sharing and server consolidation. As virtualization is being used more and more to increase utilization within production workloads, server vendors are responding with platforms, management technologies, and solutions that incorporate and leverage the underlying virtualization technologies. Sun Microsystems is taking a three-step approach to bringing more business benefits of deploying virtualization to enterprises that use x86 servers. One, Sun is emphasizing choice in how companies virtualize. Two, Sun is designing an impressive range of x86 platforms with the enterprise-class reliability, availability, and serviceability (RAS) features the company is known for with its SPARC systems. Three, Sun is incorporating physical and virtual machine management into its N1 System Manager so that customers have a single pane of glass when managing their infrastructures. Sun's customer-driven approach for how companies organize and manage their virtual environments, along with compelling RAS features, will drive virtualization as a standard for the enterprise.

INTRODUCTION

Virtualization in the form of partitioning a physical server into multiple virtual servers or virtual machines is a long-established and widely accepted solution for hardware resource sharing on large, scale-up server platform technologies. For example, Sun's Dynamic Domains have been available on the company's high-end SPARC platforms for many years. In such environments, virtualization has evolved over many years to become a stable and robust technology for improving hardware utilization and efficiently processing production workloads, including high-value, mission-critical applications. Virtualization has evolved over the years from basic resource partitioning, to a more software-based partitioning strategy, to partitioning within the operating system (OS) as customers take advantage of the performance benefits of hardware and processor improvements.

With the advent of software-based virtualization solutions from vendors such as VMware, Microsoft, and the Xen open source project, virtualization has emerged during the past few years as an increasingly important resource-sharing strategy on scale-out x86 hardware platforms. The number of servers supporting virtual machine capabilities is expected to grow rapidly in the future. Today, virtualization on x86 platforms is based largely on software implementations, such as virtual machine monitors (VMMs)

or software hypervisors, which can partition physical servers into one or more "logical servers" or virtual machines. This software virtualization on x86 systems has become a practical reality due also to the greatly increased server capacity and RAS features now available on these platforms. Each virtual machine can contain an image of multiple operating systems, such as Solaris, Windows, or Linux, which in turn can support standard applications.

According to IDC, as shown in Figure 1, worldwide shipments of server hardware platforms with virtualization capabilities installed are currently forecast to grow rapidly from 172,000 units in 2004 to 1,209,000 units in 2009, a compound annual growth rate (CAGR) of 47.7%. The number of virtualized system images running in partitions is expected to grow from 778,000 in 2004 to 5,107,000 in 2009, a CAGR of 45.7%. In the mix of virtualized server platforms, Windows and Linux images are expected to increase at CAGRs of 51.4% and 55.3%, respectively, whereas Unix images are forecast to increase at a lower CAGR of 34.7%.

FIGURE 1

The Rise of Virtualization

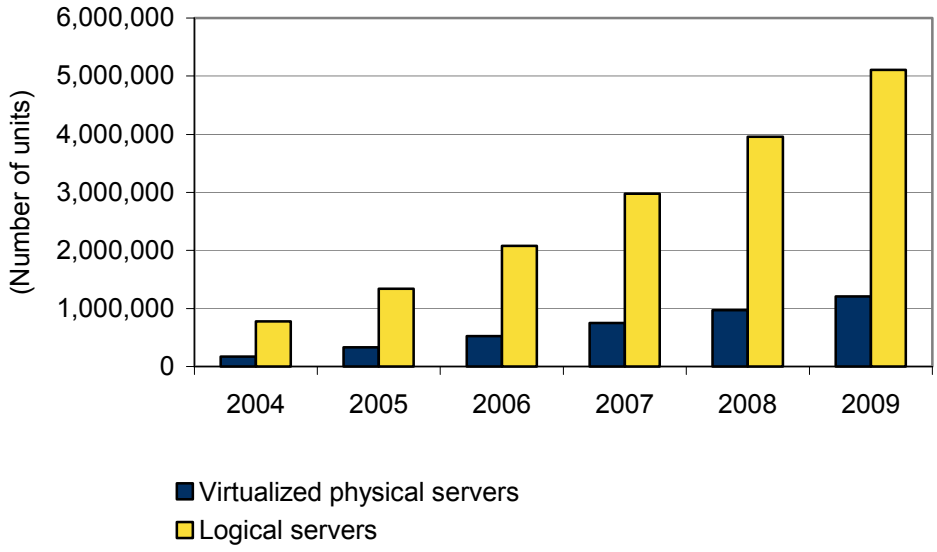

Source: IDC, 2006

Why Virtualize Systems?: The Business Benefits

There is growing evidence that virtualization is moving beyond early adopters and gaining acceptance with more mainstream customers as the market sees the benefits and many uses of virtualization technologies.

At the highest level, IDC has found two major benefits of the virtualization of server systems in partitioning and the impact on operational costs. Partitioning allows users to increase the server utilization of the current systems in their environments. Data suggests that x86 system utilization averages between 5% and 15% in most situations, which means that 85% to 95% of the capacity is idle. This information, combined with the fact that many customers assign one server per application, underscores the reality that many customers have dramatically overprovisioned their hardware resources. Virtualization is a means to isolate multiple applications running on a single host from each other so that utilization can be increased and capital costs used efficiently. This is one of the major reasons that most customers today initially started to virtualize their servers.

The other high-level benefit has been the impact on operational costs as customers manage virtual machines. By virtualizing the server, IT administrators are effectively encapsulating the application and decoupling it from the underlying hardware. By separating the application from the underlying hardware, IT administrators are able to more easily manage, provision, restart, and migrate applications on a shared pool of server hardware. This ability begins to create the IT computing utility that both users and industry have been articulating over the past few years. Perhaps even more important is the dramatic impact on operational costs. Customers report massive increases in server-to-administrator ratios, which significantly drive down the cost of managing large IT infrastructures. Because organizations can grow their infrastructures without adding staff members, they can offer more services much faster. This situation ultimately leads to competitive advantages for the business.

In fact, it is not uncommon to find users who use virtualization as the standard method for deploying new applications. Many administrators require users to have a solid business case for not having their applications hosted on virtual machines, as virtual machines can be much easier and quicker to deploy than a physical host. This stance is becoming increasingly common as companies look to drive down costs and increase utilization on x86 systems.

In addition, IT is finding that virtualization is a means to speed change in the datacenter. This need for more responsive IT is also driving the adoption of virtualization technologies. Companies using virtualization technologies report that the time to deliver a new server within their organizations has been reduced from days or weeks to just a few hours or, in some cases, only minutes.

This ability to respond quickly to changing needs in turn is driving how companies view virtualization technology. A recent IDC survey showed that over 50% of the workloads running on virtual machines are workloads that have been traditionally categorized as mission critical. This move into the enterprise applications space also means changes in the types of systems being virtualized and how those systems are being configured. Users are four times as likely to use blades for virtualization than other platforms because blade technologies are largely considered an enterprise platform that can be used in consolidation efforts. Additionally, the configurations of the systems themselves have become much richer, with a higher attach rate of processors and significantly larger memory as more and more applications are loaded onto a physical host.

Server Virtualization Use and Business Cases

Because of the benefits of virtualization, including the capital and operational cost benefits and the ease of managing applications, users are developing an increasingly broad set of usage scenarios for applying virtualization technologies in their environments. These major scenarios are as follows:

- **Hosting legacy applications.** Virtualization allows users to run applications of older operating systems that are no longer supported and that can no longer run on today's servers. The business case for legacy applications is that customers benefit from the increased performance of newer hardware without the cost of having to port the application off unsupported legacy operating environments.

- **Test and development.** An increasingly popular use is to run pilot solutions in different virtual partitions. This approach takes advantage of all the benefits of server consolidation (reducing the number of "nonproductive" servers) and allows increased agility through accelerated application rollout. New applications can be validated inside partitions before being rolled out across the datacenter, increasing the robustness of the solution. Additionally, developers can easily compile and test their applications under multiple operating systems using a single server, without multiple reboots and high investments. ISV applications can often be released on multiple operating systems at the same time. The business benefits are higher system utilization, thus requiring less hardware, and a faster time to market for new software versions.

- **Server consolidation.** The ability to run multiple partitions with each partition includes virtual BIOS, OS, and applications and provides sharing of physical hardware resources with complete software isolation. It leads to reduced hardware costs (though software costs will be similar), easier administration and simplified manageability, as well as reduced power consumption, cabling, and costly datacenter real estate.

- **Hardware migrations and upgrades.** Virtualization allows a staged migration of OS and applications onto new hardware. By validating the solutions in virtual partitions, IT shops reduce the possibility of disruption of service due to migration. When migrating to a new architecture, staging the migration in virtualized servers allows robust validation before a widespread rollout. The ROI with migration comes from the level of automation and reduced IT staff time needed to move applications across platforms.

- **Business continuity.** Most IT shops today deploy some form of failover, usually involving replicated servers. Because many failures are associated with software, users are recognizing that replicating servers in virtual partitions on the same machine provides the benefits of higher application availability from built-in virtualization failover at a reduced hardware cost.

- **Capacity planning.** Virtual partitions can be sized and resized as required. This capability offers great flexibility to IT shops, which can configure large compute resources when required (e.g., end of the quarter), then scale down as needed. By allocating compute resources as needed, IT shops can optimize the overall utilization of the server. In the future, the provisioning of compute resources will become even more dynamic, providing the ability to add and remove resources as required.

- **Load balancing.** Load balancing refers to policy-based tools that monitor utilization and virtual machines that are moved as needed to balance peak capacity with headroom. The benefit of load balancing is that users can more effectively pool their infrastructures and share a centralized managed pool across multiple applications. This ability to balance loads and provision applications dynamically helps further reduce hardware costs and increase service levels and application availability.

This ever-expanding portfolio of usage for virtualization is what makes the technology so compelling to end users and, in turn, is why Sun, its partners, and leading virtualization vendors are working together to create a robust platform that encompasses the hardware, software, and management layers of the solution. This collaborative approach is critical if the challenges of virtualization are to be overcome and the full potential of the technology is to be realized.

The Importance of Choice

In essence, virtualization is a hardware resource–sharing strategy. The core advantage of implementing virtual machines is the ability to drive higher utilizations on physical servers by sharing the use of hardware resources across several workloads, moving away from the "one application per server" approach that often characterized the use of nonvirtualized servers in the past. Another advantage is the ability to isolate one workload from another, in terms of memory, data, and hard drive contents, avoiding such problems as driver conflicts when two applications are running on the same server.

As the operational and business benefits delivered by virtualization have been realized, implementing virtualization has become a major initiative for organizations. There is growing recognition among organizations of the need for diversity in how they virtualize or decouple the application stack from the underlying hardware. Having the right tools for the job is as critical for virtualizing as it is for choosing a hardware platform or vendor.

Sun has responded by offering its customers choices at the hardware, operating system, and virtualization layers. In this manner, customers can continue to choose not only the application that best meets their needs but also the virtualization software, operating system, and hardware that will support those needs.

From a hardware partitioning perspective, Sun continues to offer customers Dynamic System Domains in its SPARC line of servers. This technology has long been used in support of running multiple mission-critical applications on the same host, yet with a high degree of electrical isolation. Hardware partitioning allows the user to take down (for maintenance or upgrades) one partition while ensuring the other partitions on the machine continue to operate. This approach to partitioning is a staple in the RISC Unix world and continues to advance in terms of the granularity of the partitions and the degree of flexibility with which the customer can manage them.

In terms of software-based or hypervisor partitioning, Sun x64 systems work with all the major implementations, most notably those from VMware (Infrastructure Server, ESX Server, and VMware Server [formerly GSX Server]), Microsoft (Microsoft Virtual Server), and Xen. In these software-only implementations, the software monitor provides a layer of abstraction (the virtual machine or logical server) between the base hardware (the physical server) and the operating system (Solaris, Windows, Linux, and many other OSes) that allows each operating system running in a virtual machine to operate as if it has control of the underlying hardware. The VMM allows multiple virtual machines to exist concurrently on the same hardware platform, providing the ability to run multiple instances of different operating systems and applications on a single physical server.

Within its Solaris 10 Operating System, Sun has developed Container technology that allows users to run multiple applications on the same OS instance while isolating applications and any faults that may occur. While the user can partition Solaris into two or more instances, a "global zone" exists for administration, helping customers address concerns around virtual machine or OS sprawl. Each type of virtualization has benefits and challenges. Hypervisor technologies allow users to run different operating environments on the same machine, but the technology does little to consolidate the number of OS instances an organization must manage. Conversely, while OS virtualization allows users to consolidate OS instances, the applications must be supported on the underlying operating environment. The key commonality across all types of virtualization is that organizations are now running more applications on fewer hardware instances, and hence the scalability and the RAS of the hardware become even more critical.

HARDWARE CONSIDERATIONS

To meet demand for increased capacity for smaller applications and Web services solutions, IT organizations historically have preferred to scale out an infrastructure of blade servers or rack-optimized servers in small increments. These applications typically aggregate their workloads across multiple similar servers to achieve the work of a single, much larger server. In some cases, the application does the aggregation; in others, the network infrastructure does the aggregation transparently. For larger and more critical applications, IT organizations favored scaling up to midrange and larger SMP servers. Large, critical applications typically run on a single server engineered to a higher degree of RAS.

One of the benefits of x86 systems has been dramatically reduced cost of acquisition. Today, about 90% of systems sold are based on x86 chips. The most common practice is to deploy a server with only a single application to avoid system faults or application contention issues. This practice has led to underutilization of server resources, massive power and cooling costs in the datacenter, and an explosion in the number of servers an organization must manage. By drawing on the best from both its volume and value systems, Sun has been able to develop volume x64 systems that offer the reliability of scale up and the performance and price of scale out. This combination of features makes the Sun modular computing concepts very complementary to the adoption and use of virtualized systems.

Modular Computing at Sun

Sun has recognized the need for enterprise features in its x64 systems and has developed a three-step design process. The first step is decoupling the server system into independent modular components. Sun's six modules are *compute* (i.e., processor technology), *input/output* (I/O), *power* (i.e., electrical power), *storage*, *cooling*, and *system management*. The second step in the modular computing approach is developing best-of-breed solutions for each module. The criteria for judging best of breed are enterprise-class requirements for capacity, I/O and connectivity, and RAS. The third step is combining modules to create a server well suited to enterprise-class workloads. IDC believes that Sun's modular approach provides a strong framework for forecasting evolution in the volume server market. Technology change is afoot that affects all six modules. Perhaps the best place to start is with processor technology and then follow with the other modules:

- With the advent of multicore processors, single-server x64 systems are approaching the capabilities previously associated with midrange SMP servers. Especially for processors running the x86 instruction set, servers support all major operating environments — Unix, Linux, and Windows, in particular. From an enterprise requirements perspective, the size of a workload that can run on a single blade is reaching enterprise class.

- Improved processor throughput has an immediate effect on I/O and network connectivity requirements. System bandwidth must be increased to avoid creating a choke point for server system performance. In-chassis I/O aggregation is needed to interface with different interconnect fabrics (e.g., Gigabit Ethernet [GbE], Fibre Channel, InfiniBand).

- Increasing processor capability and compute density also have an immediate effect on electrical power and cooling requirements. The server must also be engineered to offer efficient and robust power and cooling. These two components will need to be redundant and replaceable without causing server disruption to meet enterprise RAS requirements.

- Buyers of enterprise IT systems are keenly aware of the long-term costs of system maintenance. Hence, system management will be a critical feature and its quality will be assessed by many customers who have recently consolidated workloads on fewer servers precisely to gain greater control over server resources at a lower cost. When managing the IT budget, enterprise IT professionals are fully aware of the importance of downstream operating expenses in addition to up-front capital expenses.

The modular computing approach is a reminder that system components must be balanced to provide optimal throughput. Additional processing capability can easily be defeated by I/O bottlenecks. For example, inadequate cooling leads to thermal decay and shorter, less reliable server life. In addition, systems that require individual system administration will raise the total cost of ownership.

Sun's clean-sheet approach to modular computing allows customers to leverage the company's holistic system design — which includes computing, storage, I/O, cooling, power, and system management. Sun has designed its products to work at the datacenter level to ensure that industry-standard systems and full management capabilities — hardware, software, power, and cooling — are at the forefront of its modular computing design and permeate the Sun product line. Key design points for Sun include:

- **Use of industry-standard components.** This includes standard processors, memory, and storage as well as leveraging standards such as PCI where needed in the design. The use of industry standards not only helps drive economies of scale but also enables common design points across the modular product line.

- **Ease of integration into installed systems** with a key design point of simplified scalability. Expansion of Sun Fire X4600 and Sun Blade 8000 involves simple board adds that provide greater scalability.

- **More exacting levels of manageability and RAS.** This includes designing hot swappability into nearly all the components, making field upgrades easy and straightforward, and placing system environmentals (heat and power) in the center of the design process.

- **Performance and flexibility.** Not only is Sun using industry-standard performance processors with AMD, but it also will leverage the hardware virtualization assist capabilities that AMD is including this year. These capabilities will work to enhance software virtualization performance and security.

- **Consistent I/O throughout the datacenter.** Sun has long been known for its balanced approach to systems. These principles are central to the modular line as Sun has worked to ensure I/O is not a bottleneck in its modular and blade server line. This focus on I/O performance also ensures headroom for future board and blade upgrades.

- **Manageability.** In the Sun modular design, systems are more than the sum of the hardware components. Sun has made integrated manageability a central theme across its entire server line. The company has taken a nested approach that starts with common on-board service processors (Integrated Lights Out Management [ILOM] service processor) and extends to the infrastructure with N1 System Manager. At the application and service level, N1 System Manager nests with N1 Service Provisioning System to cover management from the most granular component to the broadest datacenter service.

Sun's portfolio of products for modular computing draws from across the company's server product lines. The company's modular computing initiative utilizes best-of-breed products from across its midrange and volume server product lines. It enables customers to deploy systems that include Sun Fire SMPs, Sun Fire x64 rack-optimized servers, and server blades in varying numbers and in varying architectures, depending on the exact needs of each customer.

This holistic approach to modular computing extends beyond just the hardware. Sun has extended the concept of modularity into the software space, specifically focusing on modularity in the operating system with Solaris 10 Containers, in the management layer with N1 System Manager and N1 Service Provisioning System, and in the application server space with Java Enterprise Systems and Java System Suites.

Sun Fire X4600

As part of the portfolio of modular systems, Sun is extending the Sun Fire line with the addition of X4600. This product blurs the line between scale out and scale up in that customers can configure the system to have between 4 and 16 Opteron processor cores. This flexibility allows the system not only to cover a great deal of enterprise workloads in the datacenter but also to deliver a great deal of investment protection from both capacity planning and use case perspectives.

The system is a 4U box configurable with up to two eight-processor boards. It works with both single-core and dual-core Opteron processors so that it can have up to 16 64-bit processors. The system includes 32 DIMM slots that hold up to 128GB of memory and has 20Gbps of bidirectional I/O.

In terms of enterprise features, many components are hot swappable, including the small form factor 2.5in. disk drives that are configurable with RAID 0 and 1. The system also includes power supplies and fans.

Sun Fire X4600 also comes standard with four 1GB Ethernet ports, six PCI slots, and the ILOM service processor, which integrates with N1 System Manager and is N1 Service Provisioning System ready.

Sun Blade 8000

The capabilities and advancements in each generation of blade products have generated more customer adoption of blade technology. Initially blades were almost exclusively an "edge of the datacenter" play. With the second generation of blade servers, customers began putting more application-tier workloads on the form factor. By continuing to enrich the feature set and performance of its blade products, Sun is squarely focused on taking blades to more mission-critical and robust datacenter workloads. The Sun Blade 8000 offers enterprise-class RAS as well as hot-pluggable and redundant I/O aggregation, fans, power supply units (PSUs), disks, and ExpressModules. Additionally, the I/O capabilities with the blade line scale up as the number of cores deployed increases. Each blade can hold four Opteron dual-core processors and offers 64GB of memory and two hot-swappable disk drives for up to 320GB of storage per blade.

The modular I/O is a big differentiator for Sun. The Sun Blade 8000 can reach a chassis I/O bandwidth of 2Tb across 10 blades with up to 80 cores — approximately 192Gbps to each blade. This bandwidth is achieved using standard PCI Express links out of the blade, across the passive backplane, and to networking modules in the back of the chassis. Each blade has up to six PCI Express links per blade. Two x8 links are used to connect to an ExpressModule that provides dedicated I/O to each blade. The remaining four links can be used to aggregate I/O across multiple blades in a Network ExpressModule with up to four modules per chassis.

This combination of I/O allows customers to dedicate I/O and share different fabrics including GbE, Fibre Channel, and InfiniBand within a chassis. It also ensures that customers will have I/O headroom as communication needs expand. Both the ExpressModule and Network ExpressModule are hot pluggable and field replaceable without having to enter the chassis.

Additionally, each blade is equipped with Sun's ILOM service processor. This processor allows the blades to be managed in the same way one would manage a standalone rack-optimized Sun server, ensuring that best practices and administrator training do not need to change simply because the form factor has changed.

By designing the platform with a five-year life cycle in mind, Sun has enabled customers to upgrade modules independently. For example, power and cooling modules can be upgraded to keep up with power draws that increase over time and threaten to limit the relevance of the platform.

Common Management Tools

The common management tools and framework across the x64 line of systems form a key value proposition for Sun. This line includes not just the Sun Fire X4600 and blade products but also the Sun Fire X2100, X4100, X4200, and X4500 systems.

The foundation of this framework is the ILOM service processor. This standard feature provides users with remote and secure KVM capabilities via Ethernet or serial port access. A standard Web browser GUI and a command line interface leverage standards such as SNMP and IPMI 2.0 for the collection of important system health and performance information.

In addition, customers can use the N1 System Manager for infrastructure management across the entire Sun server portfolio. This software package allows users to manage groups of servers, provision operating system patches and firmware updates remotely, monitor hardware and operating system events, as well as have lights-out and role-based access to the systems.

Data from N1 System Manager also feeds into the N1 Service Provisioning System so that administrators can focus on managing not just infrastructure but, more importantly, the services IT provides to end users. Increasingly, service management is becoming the focus of large organizations; therefore, automation in the management of services and the underlying infrastructure is a key feature. The N1 Service Provisioning System is the vehicle for Sun to enable its customers to begin to deliver automation at all levels in the technology solution stack.

CHALLENGES/OPPORTUNITIES

Sun and other vendors must continue to engage customers in industrywide educational discussions that articulate how to map IT requirements to the different virtualization, operating environment, and server platform choices available. Having a broad set of choices is the first step in full coverage of customer requirements.

In addition to choice, vendors must be able to offer the tools to manage both virtual and physical infrastructures. Tool proliferation is one of the many issues that inhibit streamlining IT. Sun is working to deliver more complete virtualization management within its N1 System Manager and N1 Service Provisioning System products. This work will need to continue to develop for virtualization to become even more pervasive.

Customers and vendors must not overlook the importance of robust, scalable, and balanced systems in a virtual world. As more applications are placed on fewer servers, enterprise features and design become increasingly important. Sun needs to continue to draw technologies and capabilities from its high-end systems and incorporate them in the volume system line if it is to continue to differentiate itself in the hardware layer.

CONCLUSION

Virtualization is shaping up to be one of the major trends that impact the server market as well as customer datacenters. The technology has already found a broad set of usages, ranging from improved hardware utilization in test and development, to application life-cycle extension, to high availability and disaster recovery.

As the technology becomes more mainstream and pervasive, it is even more critical that customers not only have choices in terms of how they virtualize their systems but also that the systems have the scalability, RAS, and manageability to ensure that increased utilization does not come at the expense of application availability or business performance.

Sun's investments in Containers, Dynamic Domains, and partnerships with key software virtualization vendors show that the company is making choice in virtualization a top priority. Recent investments in blade and modular server platforms indicate that Sun is also working to provide customers with access to enterprise features in its volume system line of products. With the integration Sun is doing in N1 System Manager and N1 Service Provisioning System, the company can target a single pane of glass from which users will be able to manage physical and virtual systems from the infrastructure through the services provided on those systems.