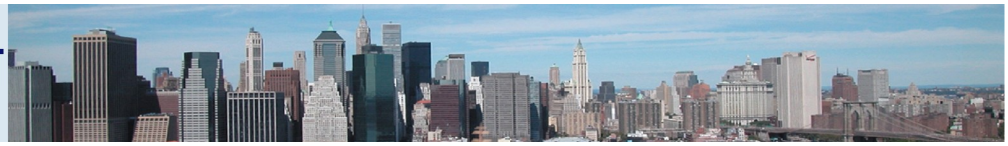

Sun Solaris Container Technology vs. IBM Logical Partitions and HP Virtual Partitions

Business Strategy Report

For

Sun Microsystems

Printed in the United States of America.

First Publication: December 2004

Produced by: Joe Barkan, Lead Analyst; Craig Norris, Editor; Barry Cohen, Editor-in-Chief

Table of Contents

Executive Summary

Edison Group has performed a comparative study of virtual server technologies from the three leading UNIX vendors: Sun, HP, and IBM.

Virtual server technology allows customers to partition a single physical server into more than one protected operating environment, with applications running in each virtual server essentially isolated from all of the others. This technology is designed to solve several long-standing problems, including underutilization of computing resources and the requirement to scale system administrative staff linearly with workload growth. With the release of Solaris 10, Sun has delivered state-of-the-art virtual server technology, with enhanced application isolation, security, and operational flexibility.

Edison Group set up a laboratory environment for testing and analyzing the features, functional richness, manageability, and flexibility of Sun's new Solaris Container virtual sever technology against HP and IBM's respective virtual server technologies (HP “vPars” and IBM “LPARs”).

The study results show that the Sun’s Solaris Containers technology does provide a more flexible and scalable virtual server environment than those provided by competitive technologies. Relative to the competition, the benefits of Sun’s product include:

- Increased hardware utilization: Solaris 10 users can run up to 4,000 containers on a single server; containers are not “locked” to any hardware, nor limited by hardware configuration. The number of virtual partitions allowed on IBM[1] and HP machines cannot exceed the number of installed processors.
- Fastest “time to production”: In our tests Solaris 10 provided the fastest “install and bring-up” times for new virtual servers; up to two times faster than the other vendors.
- High availability: Solaris 10 provided the fastest virtual server reboot times; containers can be rebooted or recovered from a failure up to 17 times faster than can the other vendors' virtual severs.

---

[1] The limitation of one partition per processor is for current IBM P4 series servers. The new P5 series of servers were not yet available at the time of this evaluation, so the new Micro-Partitioning feature could not be tested. According to IBM literature, Micro-Partitioning will enable up to ten virtual servers per processor. Sun Solaris 10 Containers will support as many as 4000 virtual servers.

- Granular security: Multiple root access privileges can be defined on a single server (versus one per server in the IBM and HP implementations), providing much more flexible security implementation, which provides an increased ability to consolidate disparate workloads on a single machine.
- Reduced operating costs: Solaris 10 supports globally managed resources where it makes sense; e.g., operating system image and storage, which reduce those administrative tasks by an n:1 basis (n= # of virtual servers on the physical server). Support of a single OS instance on the machine significantly reduces operating costs related to operating system patch management; IBM and HP require a separate OS instance for each virtual machine.
- Reduced hardware costs: Sun offers the unique ability to use their virtual server technology on the full range of Sun-branded servers and any x86-based server running Solaris 10 for x86. IBM and HP limit support for virtual servers to a subset of their proprietary Unix platforms.

We found that with appropriate expertise, all three virtual server implementations were relatively easy to configure and manage; however, the ability of Solaris 10 to host thousands of virtual servers on a single physical server, globally manage storage and I/O resources, and the implementation of a single OS instance across all of the virtual servers creates significant administrative efficiencies and, hence, cost savings. Our test results highlighted some significant functional differences as well.

Our findings lead us to believe that Sun has developed the most functionally flexible virtual server implementation, one that will serve well most enterprises looking to provide very high levels of application availability at the lowest hardware and administrative cost.

About This Report

This report documents the results of a head-to-head product comparison of the virtual server products from Sun, IBM, and HP. The study focuses on the relative strength of each vendor's virtual server implementation against using a dedicated server for each IT-implemented application or service. Its other objective is to reveal the comparative functional differences among the vendors' implementations and how those differences are reflected in the operating costs of the environments.

Common virtual server-related administrative tasks, from initial installation and virtual server partition configuration through load rebalancing, were performed for each vendor's virtual server implementation. The results were compared for their ease and speed of use. The purpose was to objectively measure (in quantitative and qualitative terms), the relative ease of everyday management (which relates to administrative costs) and functional flexibility (which relates to hardware efficiencies that can be realized through increased utilization).

Important Note: *Each vendor uses different terms and brand trademarks to name their virtual server technology. Sun's brand name for their technology is "Solaris 10 Containers Technology" and the individual virtual servers are referred to as "containers." Sun uses the name "Dynamic System Domains" to describe the hardware-based partitions supported on its higher-end servers. HP uses the term "vPar," short for Virtual Partition, to define its virtual servers. HP uses the term "nPar" to describe the electrically isolated hardware-based partitions supported on its higher-end servers. IBM uses the term "LPARs," short for Logical Partitions, to describe its virtual servers. In the test results and commentary that follows, we use the term "virtual server" generically, and we use each vendor's specific naming/branding conventions where appropriate.*

Who Should Read this Report?

The primary audiences for this report are technical reviewers of server technologies and products, data center managers, IT directors, CIOs, and other business decision makers who are interested in reading the results of third-party validated testing.

Methodology Overview

This study of virtual server technology, conducted by Edison Group, compares the functional flexibility of Sun's Solaris Containers to HP's vPars and IBM's LPARs, and assesses their relative merits and value in practical business application. It involves specifically applying a proprietary, general-purpose methodology developed by Edison Group for making product comparisons. The result is a summary definition of the strengths of each vendor's virtual server technologies as well as the benefits that may be realized relative to each other and to a non-consolidated server environment.

In the course of this study, the capabilities of Sun Solaris Containers, HP vPars, and IBM LPARs were compared using a set of tests to determine which of the three products accrues the greatest business benefits in terms of flexibility of implementation and system administrator efficiency.

Contents of this Report

The following is a brief overview of the sections contained in this document, to provide for quick reference.

- ***Rationale Behind this Report*** — a discussion of the reasons Edison Group engaged in this research.
- ***The Methodology Defined*** — a description of the tests used in the study.
- ***Test Results*** — results of each set of tests, providing summary findings and a discussion of their relevance to business operations.
- ***Conclusion*** — summarization of our findings.
- ***Appendices*** — details on the test platforms and clarification of architectural and terminology issues.

Rationale Behind This Comparison

Sun, IBM, and HP are the three leading commercial UNIX vendors. In the UNIX server market, they have been locked in a battle for customers over features, reliability, performance, and pricing for more than 10 years. All three of the vendors are now espousing a “utility computing” model, whereby customers will eventually pay for exactly the amount of computing resources required — similar to the way companies pay for electricity today. Certain key foundational services must be delivered on computing platforms to enable cost-effective utility computing. One of these key services is virtual server technology, which allows multiple server images to exist on a single physical server; this provides performance and security isolation, while allowing the flexibility required for meeting multiple customer Service Level Agreements (SLAs) on a single box. The net result is increased utilization and reduced hardware and management costs.

Historically, enterprises have dedicated individual servers to each application in a data center to provide both predictable performance and fault isolation. Using that approach, servers must be sized to accommodate peak workloads. This in turn leads to high IT costs due to over-provisioning of the hardware, high per-server administration costs, and the replicated infrastructure (e.g., power, networking) required in support of vast distributed data center architectures. Production servers are often replicated for testing, compounding the problem. Because many modern applications encompass multiple tiers of processing, single applications tend to spread across multiple servers, again increasing the cost of data center computing. When all is said and done, the average UNIX or Windows server runs at 10 to 15 percent of its capacity, wasting CPU cycles that could be used to perform useful work.

Consolidation of distributed servers has been the proposed solution to these issues for more than 10 years. Every relevant trade journal has published a multitude of articles recommending that users consolidate their servers. Nevertheless, application vendors have almost universally expected their applications each to run on their own server (or servers). Some vendors have required it as a condition of support. It is common knowledge that consolidating mixed workload types or multiple applications onto single servers in a UNIX environment can indeed cause resource conflicts as well as performance and reliability issues.

Virtual Server technology is designed specifically to let multiple workloads peacefully coexist on a single physical server and to allow enterprises to realize greater levels of capacity utilization and, hence, cost savings. While the concept of virtual partitions has been around in the UNIX market for several years, and early static partitioning implementations date back at least seven years, only now are comprehensive and extensively flexible technologies coming to market. The best virtual server technology will make UNIX and Windows servers act much more like traditional mainframes — by creating fully protected partitions that can shrink or grow as required by application loads.

Administrative costs savings are realized by the difference in administrator effort and the time it takes to bring up a new virtual server partition versus configuring a new server, and the reduction in administrative time that arises from simply having fewer physical machines and operating system images to manage. A reduction in the number of machines in a data center will also reduce network and storage component costs. Test and production environments can often run on the same server.

Downtime can also be minimized, as it is likely that rebooting a virtual partition will always be faster than rebooting a physical machine.

Enterprises evaluating virtual server implementations need to look at the levels of partitioning granularity available in terms of subdividing CPU horsepower, memory, and I/O bandwidth to determine if it is appropriate for their mix of applications. Dynamic workload rebalancing and full resource shifting capabilities are preferable over static partitions or “locked” hardware.

Edison Group developed a set of tests that would fairly compare the capabilities of the vendors' virtual server implementations and provide a guide for customers looking to reduce data center hardware and administrative costs.

Edison Group analysts performed hands-on testing, then documented, compiled, and analyzed the results.

The main thrust of this paper is to independently compare the capabilities that each vendor brings to the table in order to arrive at a conclusion regarding the efficiency and flexibility of the solutions and to elaborate on the benefits to enterprise customers.

The Test Methodology Defined

For purposes of this study, the methodology is defined as a product efficiency and flexibility evaluation process, whereby the three products in question are compared against a set of task-oriented objective and subjective metrics in order to derive an accurate set of analytical results.

The methodology employed to conduct this comparison consists of the following elements.

- **Test development:** A set of tests were devised to highlight any differences among the virtual server implementations and also to demonstrate some of the raw benefits of using virtual server technology (versus stand-alone servers).
- **Assembly of the test environment:** Servers from each of the three vendors were sourced and installed in a test lab. An isolated 100 Mb/s Ethernet network was used to connect the test servers to a workstation environment running the system consoles.
- **Test deployment:** Seasoned System Administrators performed the defined tests on each system. Test results were documented.

The tests used to perform this study were:

1. Compare the difference in time over rebooting a virtual server versus rebooting the physical server.
2. Compare the time and effort required for the installation of a physical server (single partition — no virtual servers) environment (new server installation and initialization) versus that of a new virtual server in a partitioned environment.
3. Simulate a fault that would bring down a virtual server. Model the resilience of the environment and time required to restart the virtual server.
4. Rebalance resources between virtual servers; i.e., change resource distribution. Evaluate time required to perform these tasks, and system flexibility.
5. Assign new resources to virtual servers; e.g., add additional CPU/memory to containers. Evaluate time required to perform these tasks, and system flexibility.

Also considered in this report are the differences among the vendors' virtual server implementations regarding security granularity: the ability to assign system administration privileges per container versus a requirement to assign global (system-wide) root privileges to administrators of each virtual server. We also considered administrative differences related to operating system patch management.

Test Considerations

Each of the following was considered during each test, where applicable:

- Speed of execution — the wall clock time it takes for the administrator of the system in question to kick off a job and the time to complete the job once it has been submitted.
- Steps — the number of administrator steps required to complete the tasks specified in each test (typically commands typed on the console command line).

Test Results

The table below summarizes the quantifiable test results. There were also qualified details related to these test results, which will be expressed later in this report.

Test #	*Test Detail*	*Time in Minutes*	*Time in Minutes*	*Time in Minutes*
		Sun	*HP*	*IBM*
1	Time to reboot physical server (cold boot)	13	6	12
	Time to reboot virtual server	<1	2	5
2	Time to bring up new physical server OS environment[1]	75	39	15
	Time to bring up new virtual server	14	31	15
3	Time to restart after failure of a virtual server partition	<1	3	5
4	Time required to rebalance virtual partitions within a server:			
	CPU	<1	<1[2]	<1
	Memory	N/A[3]	5[4]	<1
5	Time required to add new resources to a virtual server:			
	CPU	<1	<1[2]	<1
	Memory	N/A[3]	5[4]	<1

[1]Assumed that hardware was installed, configured, and ready for OS installation.

[2] In HP's vPar implementation, unbound CPUs can be added or moved almost instantly; CPUs that are "bound" to a partition cannot be moved without a system reboot.

[3] In the first release of Solaris 10, all system memory is pooled to be used by any container partition. "Update 1" will allow hard partitioning of memory.

[4]Requires reboot of the affected virtual servers.

Test Results Detail

Test 1: Rebooting

Compare the difference in time over rebooting a virtual server versus rebooting the physical server.

The intent of this test was to compare the length of time it takes to reboot a physical server with the length of time it take to reboot a virtual server on the same platform. Actual time to reboot any server is dependent on the hardware configuration and the amount of storage that is mounted.

The premise behind this test was to show how a well-implemented virtual server environment could reduce system/application downtime by decreasing reboot speeds, and to compare the time required to reboot virtual partitions on each vendors' platform.

SUN

The cold boot of the Sun server took 13 minutes. A warm boot of the Solaris 10 OS (which bypasses the hardware system tests) took three minutes. Rebooting a container took 18 seconds.

HP

A cold boot of the HP server took six minutes. A reboot of a vPar environment using the HP-UX 11.i v. 1.0 OS took two minutes.

IBM

A cold boot of the IBM server took 12 minutes. A reboot of an IBM LPAR took five minutes.

Edison Group Findings

The virtual servers from all three vendors rebooted significantly faster than a corresponding reboot of their entire server, demonstrating a key benefit to virtual server technology. Sun showed the fastest virtual server reboot time. Sun's ability to restart a working partition more than 10 times faster than HP and 17 times faster than IBM would serve to best minimize downtime, and would be particularly effective in highly consolidated environments and those needing a great deal of processing flexibility.

Test 2: Installation and Initialization

Compare the time and effort required for the installation and initialization of a physical server environment (single partition – no virtual servers) with that of a new virtual server in a partitioned environment.

The intent of this test was to demonstrate the time and effort savings provided when a virtual server can be deployed in place of a new physical server. For the purpose of this test we did not consider the time and effort required to buy the server and to unpack and install the hardware. We started with the assumption that the server hardware was ready to be deployed, but that the OS needed to be installed before readying the server for application deployment. The OS for each server was installed over a network rather than from CDs.

Sun

It took 75 minutes to install and initialize Solaris 10 on the Sun server. By contrast, it only took 14 minutes to create a new virtual server, which included typing 10 commands at the command line and waiting for the necessary (replicated) files to be copied to the new virtual server.

HP

It took 39 minutes to install and initialize the HP-UX OS on the HP server using HP's “Advanced Mode” (preset system parameters: hostname, time zone, IP address, file system and swap sizes, and default selection of file sets). This process consisted of six steps. Creating a new vPar on the HP server took 31 minutes. Installation of an HP vPar only differs from installing a fresh OS on a new box in that the reboot during the vPar installation does not include hardware tests. This accounts for the eight minute difference between the two.

IBM

Installation and initialization of a full system and LPARs are the same on AIX 5.2. Each LPAR gets a complete and separate AIX image. The process took 15 minutes in both cases and consisted of five discrete steps.

Edison Group Findings

The Sun implementation displayed the greatest time differential between installation and initialization of a full system and installation and initialization of a new virtual server. This is due to the fact that, unlike Solaris 10 Containers, the HP and IBM virtual server implementations maintain completely separate and distinct OS images for each virtual server partition. There is an advantage to the HP and IBM approaches. Each IBM LPAR and HP vPar can run a different version of the operating system (as long as that OS version supports the vendor's virtual server technology), or a different patch level of the operating system. This is useful for testing OS-level or patch-level upgrade compatibility with applications one virtual server at a time. Sun provides the same functionality via its Dynamic System Domain technology.

Comparing the deltas between the three systems for a single additional virtual partition does not demonstrate a significant advantage for any one system. However, virtual partition installation and initialization times are cumulative. For installations requiring a significant number of virtual partitions, an implementation on HP can take a significantly longer time compared to the other two platforms. A Solaris 10 implementation will provide an advantage over one from IBM if the maximum number of partitions possible for any comparable class of machine is taken into account as, similar to HP, the IBM server cannot have more virtual servers than it does CPUs.

Test 3: Resilience

Simulate a fault that would bring down a virtual server. Model the resilience of the environment and time required to restart the virtual server.

The intent of this test was to determine how long it would take for a virtual server partition to recover from a failure and to see how gracefully the system reacted to and recovered from the failure of a single partition.

Sun

To simulate the fault on the Sun server we killed a single virtual server process from outside the container (in the Global container). We then issued the *zoneadm -z zone1 boot* command. The killed container restarted in 13 seconds. The failure of the container did not affect the other running containers. While we did not use it for this test, Sun has a product called the "Predictive Self Healing Solaris Service Manager" that recognizes if a registered process or application has stopped, and restarts it automatically. This could be used to monitor the availability of any or all virtual servers, and to restart them after a failure.

HP

To simulate a virtual server failure and test the recovery times on a HP vPar, the *vparreboot–t* command was used from another running vPar. The *vparreboot–t* causes memory dump and reboot of the vPar. The default setting for memory dump was used (only selected memory pages 256 MB). The vPar (virtual server) rebooted in three minutes. The failure of the vPar did not affect the other running vPars.

IBM

To simulate a virtual server (LPAR) failure on the IBM server, we simulated an LPAR lockup and performed an LPAR restart. The restart of the LPAR took five minutes. The failure of the LPAR had no effect on the other running LPARs.

Edison Group Findings

All three systems recovered from our simulated virtual server failure without difficulty, and all three recovered a virtual server in much less time than it would have taken to reboot a physical server, proving the value of virtual server technology for high-availability scenarios. The results of this test tracked the results of our virtual server reboot test, with the Sun system providing the fastest recovery; seven times faster than HP and 17 times faster than IBM. However, when planning for high-availability, enterprises must consider other factors such as detection time (how long it takes a system or systems management tool to recognize a failure and initiate a restart) and application reload time, and should also consider the availability of clustering techniques and products to speed recovery in particular circumstances.

Test 4: Rebalance

Rebalance resources between virtual servers; i.e., change resource distribution. Evaluate time required to perform these tasks, and system flexibility.

The rationale behind this test was to determine both the capabilities and flexibility of each vendor's virtual server implementation as it relates to moving virtual server resources from one partition/container to another. We also judged the complexity (or lack thereof) of performing these tasks on each platform. For each system, we moved processor resources and (separately) memory between partitions/containers. We also looked at each platform's ability to move I/O resources, and we will discuss capabilities and restrictions below.

Sun

Once Dynamic Resource Pools have been configured (a 10 minute one-time process) CPUs dedicated to one processor pool (containing a single or multiple containers) can be moved to other virtual servers with a single command, and the movement of the resource is essentially instantaneous (the CPU’s instruction cache must clear before the move). CPUs can also be moved between partitions automatically based on defined parameters related to workload requirements, goal seeking, and current performance. Solaris 10 Containers also permit the administrator to define CPU and sub-CPU allocations per container.

While Solaris can limit memory allocation within a processor pool, the first release of Solaris 10 Containers does not support memory limits per container when multiple containers are running on top of one processor pool. Sun plans to support that feature in the first update of Solaris 10. Once that becomes available, a single command will move memory from one container to another. If the memory is in use before the move, a short period is required to swap the data out of memory before the resource is moved. Otherwise the movement is essentially instantaneous. In the Sun implementation, I/O and disk are global resources. Global container (full system) administrators can bind network cards to any or all containers, specifying different IP addresses to the same card for each virtual server. These assignments can be quickly changed at the command line. Disk resources can be assigned to containers via a volume mount.

HP

In the HP Virtual Server (vPar) implementation, two types of CPUs are defined. The first type is "Bound CPUs." Every vPar must have at least one bound CPU. Bound CPUs cannot be moved to another vPar. The other type of CPU is "Unbound CPUs." These can be moved across vPars. In our tests, the HP system was able to move CPUs from one vPar to another almost instantaneously (once any queued instructions have completed). On a more granular level, HP provides a workload manager that allows administrators to optimize and rebalance processor utilization across multiple processes running on a single virtual server. HP does not support moving memory or I/O between vPars without rebooting the affected vPars. Three steps are required — shut down both vPars, move the resources, and boot the partitions. In our tests that process took five minutes.

IBM

It was simple to move both CPUs and memory between LPARs on the IBM system. During our tests, our administrator determined that the fastest method of managing the IBM LPARs was via IBM's Graphical User Interface (GUI). We used that GUI to migrate the resources. Completing the necessary tasks on the GUI requires only a few seconds. The migration of resources happens immediately after the resources are cleared of instructions and data (< 1 minute). IBM's LPARs implementation does not support migration of I/O resources without restarting the affected LPARs. In our tests that process took eight minutes. Note that each IBM LPAR requires a minimum configuration of one CPU, one gigabyte of memory, and one PCI slot, limiting granularity and flexibility.

Edison Group Findings

The results of this test highlighted the flexibility and performance advantages of Sun's virtual server implementation. Simply put, HP and IBM's virtual server implementations are more closely tied to the hardware — in HP's case, moving both memory and I/O require rebooting the virtual servers and, in IBM's case, reboots are required to move I/O capacity. While rebooting virtual servers is quicker than rebooting entire systems, required reboots are unacceptable in production environments. HP and IBM customers should ensure that they can dedicate enough memory (in HP's case) and I/O (in both HP's and IBM's case) to each virtual server in order to meet peak demands. In both HP and IBM's implementations, certain resources must be dedicated to each virtual server.

In the first release of Solaris 10, Sun is missing the ability to lock memory to specific containers (and hence move memory between containers), but we expect this to be remedied in the first Solaris 10 service update (available within 90 days of the Solaris 10 release). This upgrade will function similarly to Sun's ability to move CPU resources. It should be noted that Sun, IBM, and HP have workload manager software that allows an administrator to limit processor and memory and I/O utilization for specific processes or applications running *within* a single virtual server. But while workload managers can shift CPU resources, they do not provide process, fault, or security isolation.

Solaris 10 also demonstrates a significant advantage related to storage administration. In Sun's implementation, file system and volume management occurs in the Global container, and is performed once for the all containers. By contrast, in the HP and IBM implementations, since each virtual server runs its own instance of the operating system, these tasks must be performed separately and repeatedly for each virtual server.

Test 5: Assigning Resources

Assign new resources to virtual servers; e.g., add additional CPU/Memory to containers. Evaluate time required to perform these tasks, and system flexibility.

The rationale behind this test was to evaluate the degree of simplicity involved in adding to existing partitions new (previously unused) hardware resources within the machine.

Sun

Adding CPU resources to a container was a single-command operation with near-instantaneous response. Moving I/O consisted of binding a NIC to a particular container and assigning an IP address — a two-step process. Adding a disk volume to a container requires mounting the volume on the virtual server — a single command. Again, in the current Solaris 10 version, memory is shared among the virtual servers.

HP

Unbound CPU resources can be moved immediately with one command in the HP-UX environment. However, adding memory or I/O (network or disk) requires a reboot of the vPar, which took five minutes in our tests.

IBM

As in our rebalance tests, the IBM LPAR implementation allows for simple and near instant movement of both CPU and memory resources. Moving I/O (network or disk) within the IBM system requires a reboot of the LPAR, which took seven minutes in our tests.

Edison Group Findings

As expected, the results of this test did not vary significantly from those of the rebalance test (Test 4); this highlights the operational flexibility and high-availability characteristics of the Solaris 10 system.

Virtual Server Root Security Considerations

As part of our testing we examined each vendor's implementation of virtual server root security. Beyond the ability to segregate and isolate workloads, we believe that — from an administrator perspective — it is important that virtual server implementations provide security isolation. Security isolation among virtual servers provides an extra layer of protection against downtime and additional flexibility, since multiple root administrators can be defined on a single server. For example, two IT administrators from different departments can both manage *their* server without worrying about affecting each other's application environments. Individual virtual servers could be rebooted with no chance of accidentally rebooting others — or the entire machine. In our tests, we found that only Sun provides this level of security granularity.

Sun

Solaris Containers provide application and process isolation on top of a single instance of the Solaris operating system. Each container has its own root that has no power over anything in any of the other containers. Users or applications in one container cannot see or affect applications in any other container. Network isolation is also provided. Even though multiple containers may share a single NIC, each container can only see its own traffic.

A containerized server also has a "Global container" that is comparable to a normal Solaris instance. The Global container has access to the entire system; hence the Global container administrator also has access to the entire system, and has the ability to create new containers as desired. Each container (other than the Global container), has its own separate root access, so individual administrators can be assigned to each container — and will only see the resources assigned to that container, the applications and processes running in that container, and the users defined in the namespace of that container. Users in different containers can have the same user ID without conflict.

HP

HP implements a single root user across all vPars in a single nPar (electrically isolated, fixed hardware partition). This root user has full administrative access to all virtual servers running on the nPar. HP recommends that users concerned about administrative security isolation use nPars instead of vPars.

IBM

Similar to HP, IBM implements a single root user for the entire system. No root security isolation exists.

Patch Management Considerations

An additional, important benefit to Sun's single operating system that became clear when we looked into the security features of each vendors' technology relates to operating system patch management. Sun's single OS instance method of virtual server deployment means that only one copy of the OS needs to be managed and patched. The competing technologies require that each virtual server instance has its own copy of the operating system — any upgrade or security patch must be applied individually to each virtual server. There is a real administrative workload savings related to patch management of *n:1* (where n= the number of defined virtual servers). A single OS patch installation could take more than one hour of administrator time and often requires the reboot of the server; multiplied out by hundreds of virtual servers managed, the costs savings will be significant.

Conclusions

Our tests revealed that the virtual server technologies from all three vendors are (given the appropriate expertise) relatively easy to install, configure, and manage. We found no significant differences in terms of ease of use from an operational management perspective — within the limits of each product's capabilities.

We did, however, find a major architectural difference between Sun's approach to virtual servers and those of HP and IBM. The easiest way to understand this difference is to note that all of the Sun Solaris Containers running on a server run in virtualized partitions on a *single instance* of the Solaris 10 operating system. In both the HP and IBM implementations, each virtual server running on the machine loads its own instance of the operating system. Sun has taken hardware or firmware-based partitioning approaches one better by providing a higher layer of software abstraction, which creates the required isolation for highly flexible, easily implemented and managed virtual servers while reducing relative administrative costs

The primary advantage revealed in this study of HP and IBM's hardware-oriented, multiple OS-instance approaches relative to Solaris Containers is their ability to run different versions of their respective operating systems, or different OS patch levels, simultaneously on multiple virtual servers running on a single machine. This is especially useful for OS upgrade testing prior to moving an application into production on the new OS level. Sun provides the same functionality via its Dynamic System Domains.

Our testing revealed some of the disadvantages of HP's and IBM's approach: longer restart times and less flexible resource rebalancing capabilities — for example, some resources fixed to virtual servers, virtual server reboots required to move certain resources, hardware-based limits to maximum number of virtual servers per machine. Relative to the competitive vendors' approaches, Sun's sub-processor allocation technology allows for many times the number of virtual servers to be installed and run on a single server. IBM and HP's workload manager software provides sub-CPU allocation capabilities among applications but, unlike Sun's Solaris Containers, they do not provide security, namespace, and fault isolation.

With Solaris 10 Container Technology, Sun has added security, namespace, and fault isolation in addition to the full resource containment available in Solaris 9. Fully isolated containers can run on a single-processor system all the way up through Sun's 72-processor E25K. The technology is available for both Sun SPARC-based servers and x86 architecture servers running Solaris x86. IBM's LPARs and HP's vPars operate only on a high-end subset of their respective UNIX server product lines. Sun's technology can be used to run two or more instances of the same application or database (e.g., Oracle) on a single system, to easily create and tear down multiple test environments; this can be used, for example, to run dozens of web servers each with its own isolated Port 80, or to create hundreds of virtual servers on a single server for a community of university computer science students while adding very little overhead

Sun has also simplified the acquisition and deployment of virtual servers by making the technology an integrated part of Solaris 10, not a bolt-on or extra-cost package. Also, unlike competitive technologies, it does not require the purchase of any additional hardware — for example, a management workstation.

Customers evaluating virtual server technologies should note these significant architectural and feature differences, match them to application requirements, and determine the cost savings that can be achieved through each vendors' virtual server technologies.

Solaris Containers in the Real World

The tests performed by the Edison Group were generic in nature — that is, the results can be applied to all customers in all industries. However, our research has established several industry-related scenarios where we believe that the highly flexible virtual sever technology provided by Sun's Solaris 10 will provide a considerable degree of benefit. While these examples are specific, many of the attributes and benefits we describe below will apply to other industries.

Financial — The financial industry is undergoing a wave of rapid consolidation. Consolidation of data centers and applications is required after every merger or acquisition if expected cost synergies are to be achieved. Solaris Containers help in several ways:

- With Solaris 10's ability to host multiple applications on a single server without requiring administrators to dedicate CPUs to each application, applications can be migrated from one data center to another without migration of the underlying server hardware.
- New applications (acquired in the M&A transaction) can be rapidly brought up and tested in parallel with production systems in different data centers without the need to purchase new hardware on which to run them.
- With the granular security provided by Solaris Containers, user accounts and data can be rapidly migrated to alternate data centers and consolidated, sharing hardware resources while preserving separate administrative control and guaranteed performance levels to each group.
- The increased system utilization that can safely be achieved with Solaris Containers reduces the total number of required servers after a merger; once applications and/or user accounts are consolidated, recovered servers can be sold on the secondary market.
- The planning and implementation of data center and systems consolidations takes long enough that the situations and needs often change before they are finished. Solaris Containers offer a degree of versatility that has enabled greater flexibility for these organizations.

Telco industry — The rise of new telecommunications services, IP telephony, and G3 Mobile telecommunications technologies require significant server resources — which must scale with the expected rapid increase in customer demand. Using Solaris Containers, Telco customers will benefit from:

- Reduced server hardware costs: Telcos will be able to safely combine application services for multiple customers on a single server — increasing utilization while maintaining necessary SLAs around performance, availability, and — perhaps most importantly — security.
- Reduced data center floor space, power, and cooling requirements via the secure server consolidation provided by Solaris Containers of hundreds or thousands of applications or customers.
- Rapid application deployments as new applications/services are purchased on a per-customer basis.
- Fast, automated, flexible resource rebalancing across virtual servers to meet peak service demand.
- Administrative cost savings as fewer OS images (one per physical server) need to be managed and patched.

Security — Access control is a key component of all commercial and governmental security strategies. The extremely granular administrative rights capabilities provided by Solaris Containers, along with strong fault isolation, are of great benefit where security concerns are paramount.

- Government agencies can benefit from the ability of Solaris Containers to reduce costs via consolidation while maintaining necessary security isolation.

Manufacturing — Manufacturing companies are moving to reduce costs by simplifying and streamlining supply chains and by off-shoring where possible. Achieving both of these goals simultaneously requires flexible system capabilities — including the ability to rapidly shift and rebalance workloads while achieving SLAs and maintaining 24x7 availability. Solaris Containers allows for:

- Greater system utilization can be achieved using Solaris Containers, reducing hardware costs, OS images, and related system administration costs.
- Flexible application deployment. With Solaris Containers, new OS images and/or applications can be brought up in seconds.

- Solaris Containers enable rapid workload-to-hardware resource rebalancing as required by off-hours processing or workforce time-container changes.
- New supply chain applications can be tested on production systems without putting those systems at risk of downtime.

In summary, the net result of our study is that we have determined the advantages of Sun's Solaris 10 Containers over its competitors include:

- The ability to run up to 4,000 containers on a single server; containers are not "locked" to any hardware, nor limited by hardware configuration. The number of virtual partitions allowed on IBM and HP machines cannot exceed the number of installed processors. Sun's implementation best allows administrators to achieve very high server utilization levels while meeting or exceeding customer SLAs.
- The fastest "install and bring-up" times for new virtual servers; up to two times faster than the other vendors.
- The fastest virtual server reboot times; containers can be rebooted or recovered from a failure up to 17 times faster than the other vendors. This significantly reduces administrative costs and "time-to-production."
- Multiple root access privileges can be defined on a single server (versus one per server in the IBM and HP implementations), providing much more flexible security implementation, which creates an increased ability to consolidate workloads.
- Reduced operating costs due to support of globally managed resources where it makes sense, e.g., operating system image and storage, which reduce those administrative tasks by an n:1 basis (n= # of virtual servers on the physical server).
- Reduced operating costs related to operating system patch management; one OS image is used across all virtual servers on a Solaris system. IBM and HP require a separate OS instance for each virtual machine.
- Highly flexible CPU, memory, network, and storage resource sharing and rebalancing capabilities with the highest degree of granularity.
- The unique ability to use the virtual server technology on their full range of Sun-branded servers and any x86-based server running Solaris 10 for x86.

Our findings lead us to believe that Sun has developed the most advanced and functionally flexible virtual server implementation, one that will serve well enterprises looking to increase server utilization and simplify administration in an effort to reduce IT costs.

Appendix I - Test Platform Details

Below are the technical specifications of the lab servers used to conduct the tests in this report:

Sun:

Sun Fire 6800

1 Uniboard Domain 4 x 900 MHz CPUs

16 GB Mem

1 36 GB 7200 RPM hard disk

Solaris 10 build 69 (Publicly available as Solaris Express 10/04)

HP:

HP rp7410 server:

1 Cell Board (4 x 875 MHz processors)

4 GB memory

2x 36 GB 10 KRPM internal hard disks

2x 1 Gb LAN

64-bit HP-UX 11.i v. 1.0 (June 2004)

Mission-critical OE (default selection of filesets)

Standard patch Bundle Release June 2004

Virtual Partitions 3.01

IBM:

1 pSeries 650 6M 2-Rack Server

6 1.2 GHz PWR4+ Processors

6 GB SDRAM DIMM Memory

2 I/O Drawers

8 36.4 GB 10K Ultra SCSI hard disk

4 10/100/1000 Ethernet cards

2 Dual Ultra3 SCSI adapters

AIX version 5.2 ML3

IBM Hardware Management Console running on a Linux-based x86 server. This is a dedicated machine required to configure and manage the LPARS.