

PERFORMANCE STUDY

A Performance Comparison of Hypervisors



Contents

Introduction	1
Virtualization Approaches	1
Enterprise Virtualization Infrastructure	2
Test Methodology and Configuration	4
Guest Operating System.....	4
Test Workloads.....	4
Hardware configuration.....	5
Software Configuration.....	5
Virtual Machine Configuration.....	5
Test Results	6
SPECcpu2000 Integer	6
Passmark.....	8
Compile Workloads.....	10
Netperf.....	11
SPECjbb2005.....	13
Discussion	14
Single Virtual-CPU Tests.....	14
Virtual SMP Tests.....	14
Qualitative Comparison	15
Future Work	16
Conclusion	16
References	17
Appendix A: Test Configuration	18
Hardware Configuration.....	18
Client (for netperf tests).....	18
Native and Guest Operating System Configuration.....	18
Native:.....	18
VMware ESX Server 3.0.1.....	18
Xen 3.0.3-0.....	18
Hypervisor Configurations.....	19

Introduction

Virtualization has rapidly attained mainstream status in enterprise IT by delivering transformative cost savings as well as increased operational efficiency, flexibility and IT service levels. Intel and AMD have independently developed virtualization extensions to the x86 architecture called hardware virtualization. This and other recent hardware advances such as multicore processors are further fueling the adoption of virtualization.

While a full virtual service-oriented infrastructure is composed of a wide array of technologies that provide resource aggregation, management, availability and mobility¹, the foundational core of virtual infrastructure is the hypervisor.

This paper provides a quantitative and qualitative comparison of two virtualization hypervisors available for the x86 architecture — VMware ESX Server 3.0.1 and open-source Xen 3.0.3 — to validate their readiness for enterprise datacenters. A series of performance experiments was conducted on the latest shipping versions (at the time of this study in November 2006) for both hypervisors using Microsoft Windows as the guest operating system. This white paper discusses the results of these experiments. The discussion in this white paper should help both IT decision makers and end users to choose the right virtualization hypervisor for their datacenters.

The experimental results show that VMware ESX Server delivers the superior, production-ready performance and scalability needed to implement an efficient and responsive datacenter. Furthermore, while we had no problems exercising enterprise virtualization capabilities such as Virtual SMP and virtual machine scalability using the VMware ESX Server hypervisor, we were not successful in running similar tests with the Xen 3.0.3 hypervisor due to product failures.

Virtualization Approaches

The x86 architecture is the most popular computer architecture in enterprise datacenters today, hence virtual infrastructure for the x86 architecture has tremendous benefits. The two leading software virtualization approaches to date have been full virtualization and paravirtualization. AMD and Intel have recently introduced new processor instructions to assist virtualization software.

- The **full virtualization** approach allows datacenters to run an unmodified guest operating system, thus maintaining the existing investments in operating systems and applications and providing a nondisruptive migration to virtualized environments. VMware uses a combination of direct execution and binary translation techniques [1] to achieve full virtualization of an x86 system
- The **paravirtualization** approach modifies the guest operating system to eliminate the need for binary translation. Therefore it offers potential performance advantages for certain workloads but requires using specially modified operating system kernels [2]. The Xen open source project was designed initially to support paravirtualized operating systems. While it is possible to modify open source operating systems, such as Linux and OpenBSD, it is not possible to modify “closed” source operating systems such as Microsoft Windows. It is also not practical to modify older versions of open source operating systems that are already in use. As it turns out, Microsoft Windows is the most widely deployed operating system in enterprise datacenters. For such unmodified guest operating systems, a virtualization hypervisor must either adopt the full virtualization approach or rely on hardware virtualization in the processor architecture.

¹ Resource aggregation refers to the capability to pool, share, and throttle memory, processing power, network, and storage across server instances. Mobility refers to the capability to perform live migrations of running virtual machines from one physical server to another in response to availability requirements.

- The **hardware virtualization support** enabled by AMD-V and Intel VT technologies introduces virtualization in the x86 processor architecture itself. While first-generation hardware assist support includes CPU virtualization only, later generations are expected to include memory and I/O virtualization as well. The emergence of virtualization hardware assist reduces the need to paravirtualize guest operating systems. In fact, Xen vendors such as Virtual Iron have announced that they are supporting only full virtualization using AMD-V and Intel VT processors and are not supporting paravirtualization [9].

While an architectural comparison between these approaches is of interest to those trying to predict the long-term direction of virtualization technology, the advantage of any one approach for any single element of virtualization overhead may be outweighed by a variety of other datacenter requirements outlined in the “Datacenter Requirements” section of this paper. It is a combination of all three approaches that will ultimately help architect a successful virtual datacenter.

Enterprise Virtualization Infrastructure

Enterprise datacenters typically start by implementing virtualization as the basis for server consolidation and containment². Over time, IT staff tend to branch out in their use of virtualization, to the point where it becomes a standard part of the production datacenter infrastructure. While this standardization on virtual infrastructure provides tremendous value in improved resource utilization, superior manageability and flexibility, and increased application availability, these benefits cannot be achieved through the hypervisor alone. Enterprise virtualization is a broad IT initiative, of which basic server partitioning is just one facet.

² Consolidation is the process and result of shrinking the overall server footprint in a datacenter to a smaller number of virtualized servers. Containment is the process and result of containing the further proliferation of physical servers, beginning at a particular time, through virtualization.

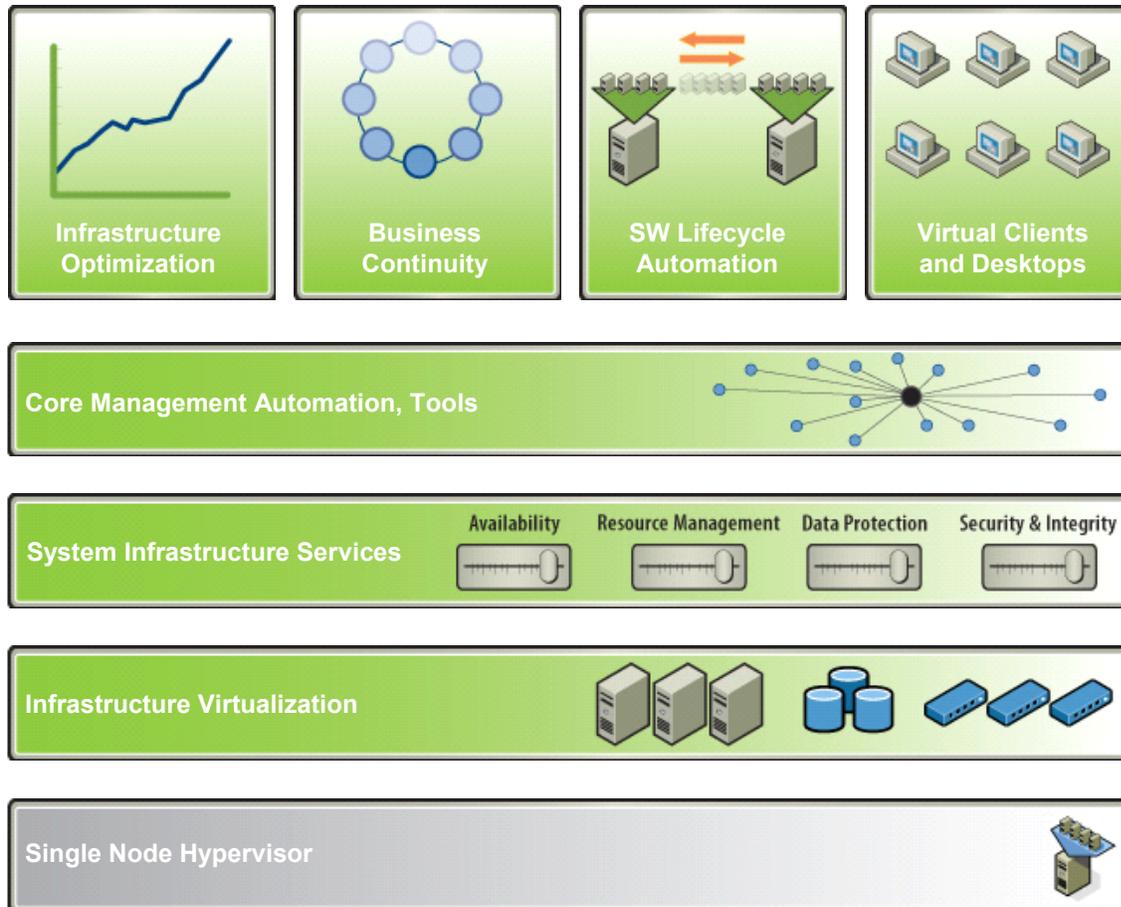


Figure 1 — Enterprise Virtualization Infrastructure

As illustrated in Figure 1, enterprise virtualization infrastructure consists of the following components:

- **Single-node hypervisor** to enable server partitioning capability.
- **Infrastructure virtualization** that virtualizes and aggregates industry standard servers and their attached network and storage into unified resource pools.
- A set of virtualization-based distributed **system infrastructure services** such as resource management to dynamically and intelligently optimize the available resources among virtual machines. High availability for better service levels, data protection for reliable and cost effective disaster recovery, and security and integrity to better protect existing infrastructure investments from typical datacenter vulnerabilities.
- A suite of **management automation technologies and tools** that provide virtualization-specific capabilities such as comprehensive system resource monitoring (of metrics such as CPU activity, disk access, memory utilization, and network bandwidth), automated provisioning, cloning, and workload migration support.
- A set of **end-to-end solutions** such as infrastructure optimization, business continuity, software lifecycle automation, and Virtual Desktop Infrastructure complete the virtual infrastructure.

Together, all these components in the enterprise virtualization infrastructure are required to successfully implement virtualization inside datacenters. The foundational element of the virtual infrastructure, however, is the hypervisor.

The next two sections provide a comparison of the operational characteristics of the VMware ESX Server 3.0.1 and Xen 3.0.3 hypervisors, with a specific focus on their respective performance characteristics.

Test Methodology and Configuration

To conduct a quantitative comparison of hypervisors, a few key decisions had to be made — the choice of the guest operating system and the workloads to use for the evaluation.

Guest Operating System

Microsoft Windows Server 2003 was selected as the guest operating system for these tests for several reasons. First, Microsoft Windows operating systems are the most widely deployed operating systems on x86 platforms. Second, typical enterprise customers run standard, off-the-shelf operating systems and software in their virtual machines to maintain compatibility and compliance with their support agreements. To run such unmodified guest operating systems with the Xen 3.0.3 hypervisor, one needs to use the latest generation of hardware that supports virtualization in the x86 processors.

The Linux community has adopted the “para-virt ops” API approach based on VMware’s proposed Virtual Machine Interface (VMI), a completely open hypervisor interface. The “para-virt ops” API together with back-end VMI support is scheduled to be included in the Linux kernel version 2.6.21. VMware has already announced plans to support these paravirtualized guest operating systems and we will revisit these tests for both unmodified and paravirtualized Linux guest operating systems at that time.

Test Workloads

A typical enterprise datacenter runs a mix of CPU-, memory-, and I/O-intensive applications. Hence the test workloads chosen for these experiments comprise several well-known standard benchmark tests, as listed below:

- The integer component of the SPECcpu2000 benchmark suite, available from SPEC® (Standard Performance Evaluation Corporation), was chosen to represent CPU-intensive applications [6].
- Passmark, a synthetic suite of benchmarks intended to isolate various aspects of workstation performance, was selected to represent desktop-oriented workloads [4].
- Netperf was used to simulate the network usage in a datacenter [3].
- The SPECjbb2005 benchmark suite from SPEC was used to represent the Java applications typically used in the datacenters [7].
- A compile workload — build SPECcpu2000 INT package — was also added to capture typical IT development and test usage in datacenters.

The objective of these experiments was to test the performance and scalability of the two virtualization hypervisors. The tests were performed using a configuration with a single virtual CPU. We attempted to repeat the single-virtual-CPU tests using virtual SMP configurations (for example, two virtual CPUs and four virtual CPUs), as well as to run scalability tests using multiple virtual machines. However, it was not possible to run the Xen 3.0.3 hypervisor either in multiple

virtual CPU configurations or using multiple virtual machines. More details have been provided later in the paper.

It must be noted here that any experimental test setup based solely upon resource-intensive workloads driving a physical system into and past saturation is not a likely customer scenario. In most production deployments, IT managers conduct detailed capacity planning and sizing exercises and the average utilization of the servers is kept within reasonable limits to allow for usage spikes and future capacity growth. The benchmark test suites are used in these experiments only to illustrate performance and scalability of the two virtualization hypervisors.

Hardware configuration

The system used to run all benchmark tests was an IBM X3500 server with two VT-enabled dual-core 3GHz Intel Woodcrest CPUs (total four cores). Although the test system had 5GB of RAM installed, it was booted with only 1GB of RAM for native tests. Additionally, the test system was configured with a dualport 1Gbps Ethernet adapter and two 146GB SAS disk drives. For native operating system tests, all data was captured using Windows Server 2003 Enterprise Edition R2 32-bit. Only the Netperf tests needed a client, which used either one or two Microsoft Windows 2000 clients. The client used was a Dell 1600SC server configured with two Pentium 4 2.4GHz processors and one 1Gbps network adapter card. All tests were controlled from within the virtual machine itself. As shown in Figure 2, both Netperf clients communicated with a single virtual machine.

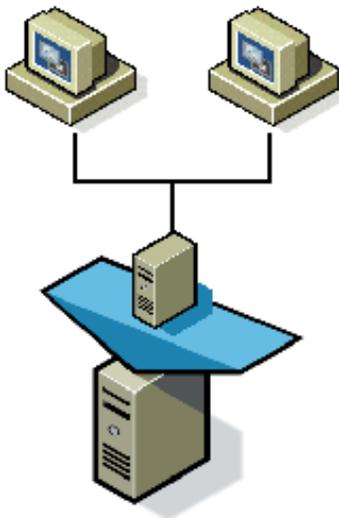


Figure 2 — Configuration for two-client Netperf test

Software Configuration

All the experiments described in this paper were run using ESX Server 3.0.1 GA release and Xen 3.0.3-0 release. Both were the latest shipping releases for the two virtualization hypervisors at the time of this testing in November 2006. We downloaded the Xen 3.0.3 version from University of Cambridge Computer Laboratory [8].

Virtual Machine Configuration

Each virtual machine was configured with one virtual CPU and 1GB of memory unless specifically noted. For the SPECjbb2005 tests, each virtual machine was configured with 1.6GB of memory and two or four virtual CPUs based on the test run. The Windows Server 2003 EE R2 32-bit operating

system was installed inside the virtual machine. The 32-bit version was chosen because it is still the most deployed operating system. We plan to run similar tests using the 64-bit version in the future.

No attempt was made to optimize the benchmark test results in any way. Default tools and settings were used in all cases. For SPECjbb2005 tests, BEA Systems' JRockit 5.0 R26.4.0-63 Java virtual machine (JVM) was used. The java virtual machine options for SPECjbb2005 tests were set to `-Xms960m -Xmx960m -Xgc:parallel -XXaggressive:opt -XXcompactratio8 -XXminblocksize16k`. For each test run, only the single active virtual machine was powered on, since idling virtual machines continue to consume a small amount of resources and can skew results.

Test Results

This section provides detailed results for each of the experiments. All results, unless specifically noted, have been normalized to native performance on a throughput basis to make it easier to illustrate the slowdown resulting from virtualization. Higher numbers indicate better performance, unless indicated otherwise.

For an enterprise datacenter, better performance implies significant benefits in several ways:

- Better performance across different workloads implies that more application types can be successfully deployed in production in virtual environments.
- Near-native performance also indicates that more virtual machines can be deployed on a single physical server, resulting in higher consolidation ratios. This can help even if an enterprise plans to standardize on virtual infrastructure for server consolidation alone.
- Finally, better performance can have a measurable impact on many of the costs, such as hardware, software, administration, support, system downtime, and user productivity, which influence the total cost of ownership (TCO).

Hence a hypervisor's performance can contribute considerably towards easing the initial cultural shift to adopting virtualization inside datacenters as well as transparently migrating end users to virtualized environments. However, as stated earlier, a hypervisor is just one component required for successfully implementing enterprise virtualization infrastructure.

SPECcpu2000 Integer

This benchmark comprises mostly user-level computation, hence we expect both virtualization hypervisors to score close to native. The results, as shown in Figure 3, show a slowdown over native ranging from 0–6 percent for the VMware ESX Server and from 1–12 percent slowdown for the Xen 3.0.3 hypervisor. Overall, Xen 3.0.3 shows twice the overhead of VMware ESX Server, an average slowdown of 6 percent compared to 3 percent.

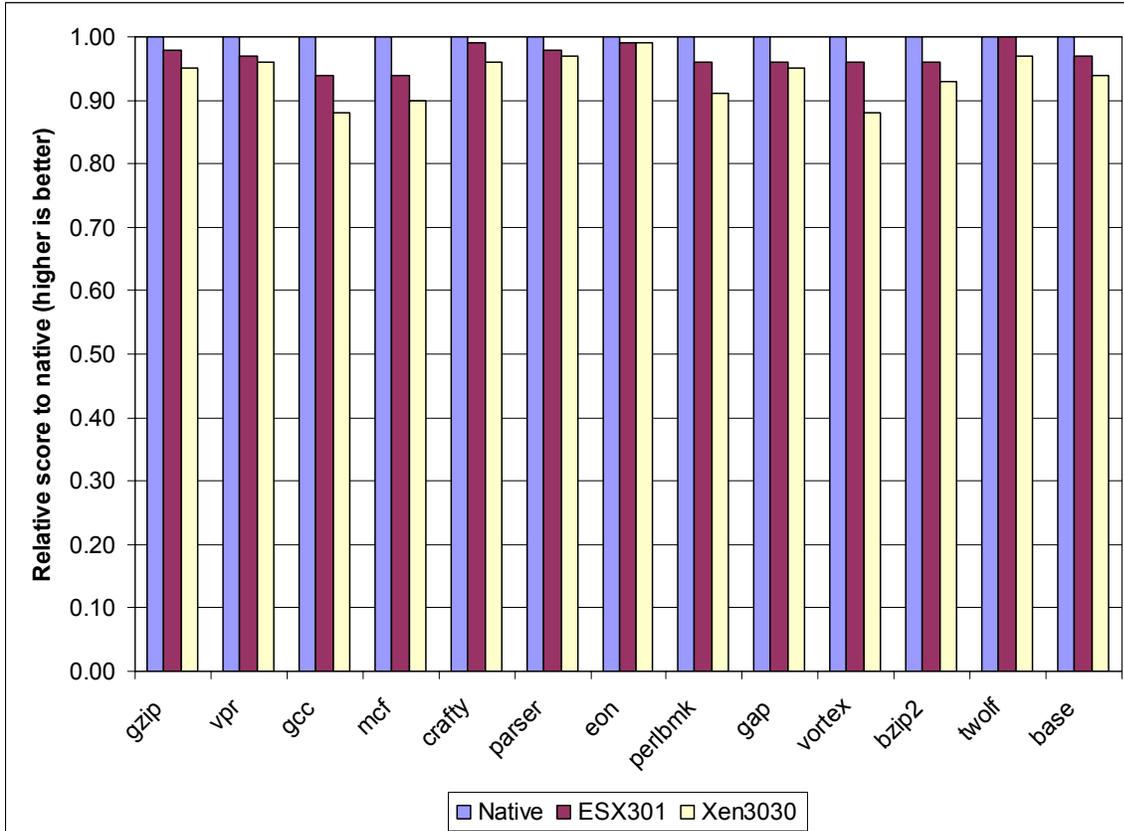


Figure 3 — SPECcpu INT 2000 results compared to native (higher values are better)

Passmark

Figure 4 shows the results obtained for CPU tests in the Passmark benchmark suite. The following CPUmark subtests were run during these experiments: IntMath, FPMath, MMX, SSE/3DNow, Compression, Encryption, ImageRotate, and StringSort. These tests comprise mostly user-level computation, hence we expect both virtualization hypervisors to score close to native. The results show a slowdown over native ranging from 4–18 percent for VMware ESX Server and from 6–41 percent overhead compared to native for the Xen 3.0.3 hypervisor. Overall, Xen 3.0.3-0 shows almost twice the overhead, an average slowdown of 17 percent compared to 9 percent for VMware ESX Server.

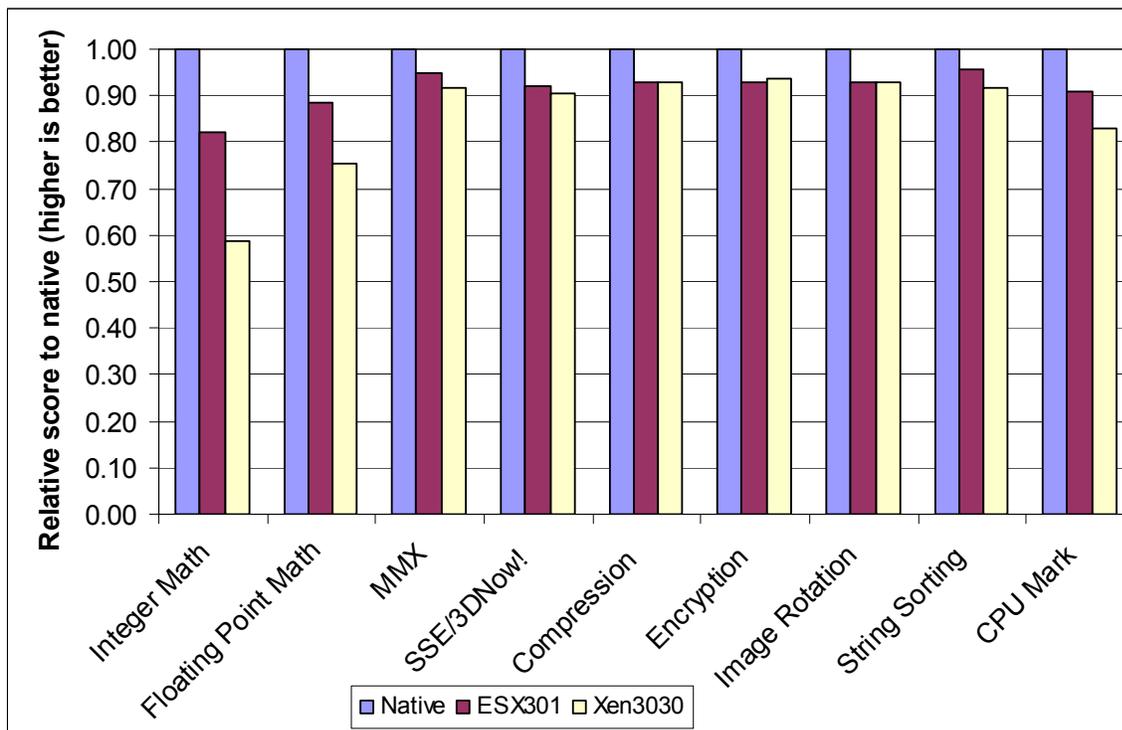


Figure 4 — Passmark – CPU results compared to native (higher values are better)

Both SPECcpu2000 Integer and Passmark – CPU tests demonstrate that VMware ESX Server can handle CPU intensive applications — such as database servers, application servers, file servers, terminal servers, and mail servers — in a typical enterprise datacenter more efficiently than the Xen hypervisor.

Figure 5 shows the results for Memory tests in the Passmark benchmark. The Memorymark subtests included: AllocateSmallBlock, ReadCached, ReadUncached, and Write. Both VMware ESX Server and Xen hypervisors demonstrate near native performance. The VMware ESX Server shows an average 2 percent overhead compared to native, while the Xen results show an average of 3 percent overhead compared to the native performance.

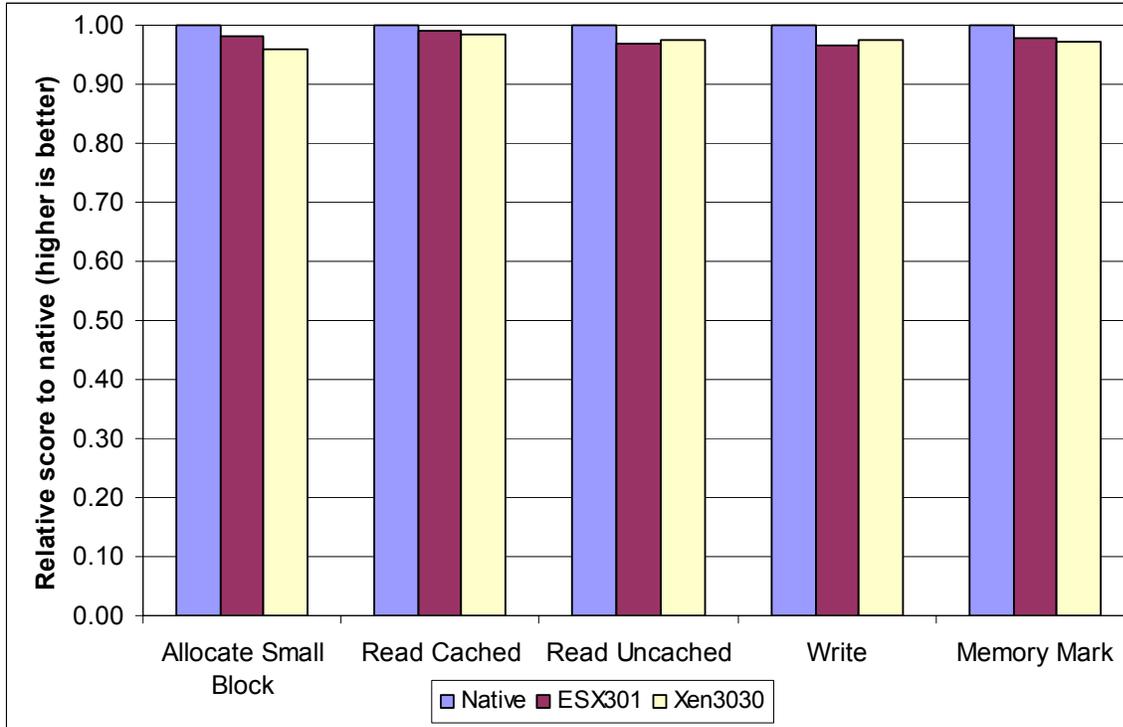


Figure 5 — Passmark - Memory results compared to native (higher values are better)

Compile Workloads

We also examined a compile workload during these experiments: build SPECcpu2000 INT package for Windows. The workload used Microsoft Visual C++ 2005 Express Edition compiler with Microsoft PSDK for Windows Server 2003 R2. VMware ESX Server 3.0.1 performed better than Xen 3.0.3. For the SPECcpu2000 Int compile job, the native test took 102 seconds, the VMware ESX Server test took 113 seconds (90 percent of native performance), and the Xen test took 149 seconds (68 percent of native). Figure 6 shows the relative throughput (inverse of elapsed time) for the compile workload as normalized to the throughput in the native environment. In previously published papers using a paravirtualized Linux guest under Xen, compile benchmarks show near-native performance. The current results demonstrate that such results do not carry over to fully virtualized guests using hardware-assisted virtualization.

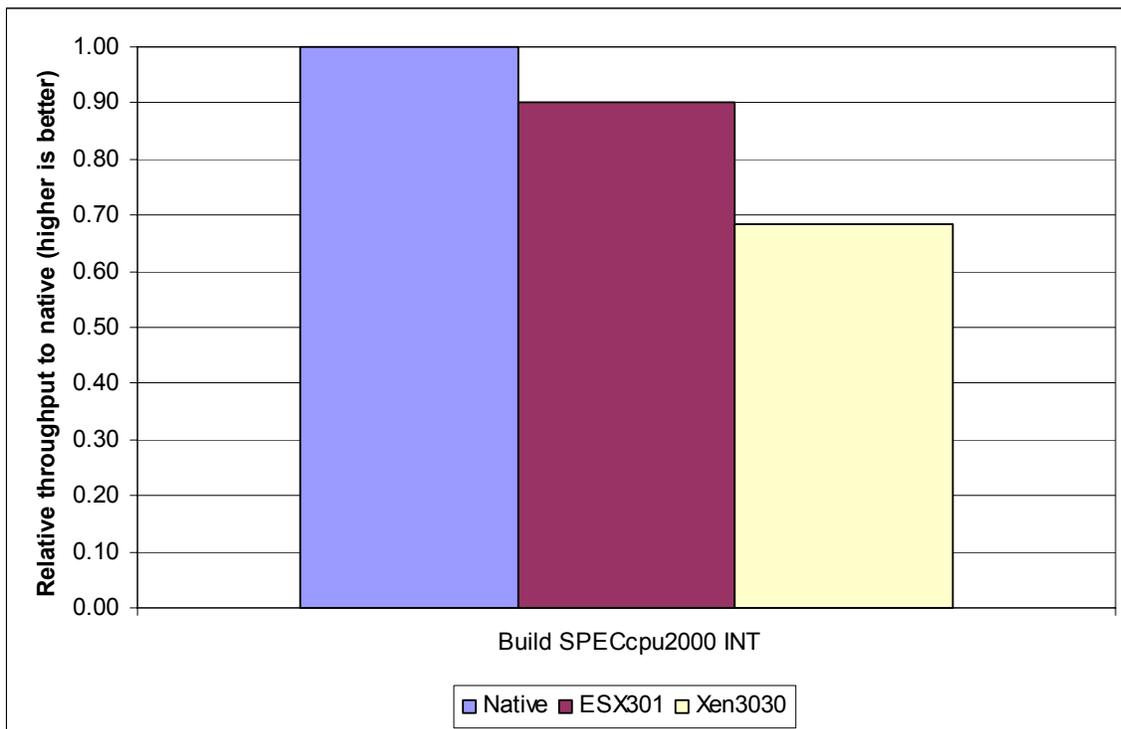


Figure 6 — Compile workload result compared to native (higher values are better)

Netperf

These experiments involved running single or multiple client processes communicating with a single uniprocessor virtual machine through a dedicated physical Ethernet adapter and port. All tests were based on the Netperf TCP_STREAM test. The MessageSize was set to 8192 bytes and the SocketSize was set to 65,536 bytes.

Figure 7 shows the Netperf results for send and receive tests for both one and two clients. VMware ESX Server delivers near native performance for both one- and two-client tests. The Xen hypervisor, on the other hand, is extremely slow, performing at only 3–6 percent of the native performance.

Figure 8 shows how the total throughput in Mb/sec scales from one client to two clients for both send and receive tests. This comparison of one-client tests to corresponding two-client tests shows that the native tests scale almost perfectly: both throughput and CPU utilization double. VMware ESX Server does very well, too: the throughput for two-client tests goes up 1.9–2 times compared to the one-client tests. Xen is almost CPU saturated for the one-client case, hence it does not get much scaling and even slows down for the send case.

The Netperf results prove that by using its direct I/O architecture together with the paravirtualized vmxnet network driver approach, VMware ESX Server can successfully virtualize network I/O intensive datacenter applications such as Web servers, file servers, and mail servers. The very poor network performance makes the Xen hypervisor less suitable for any such applications.

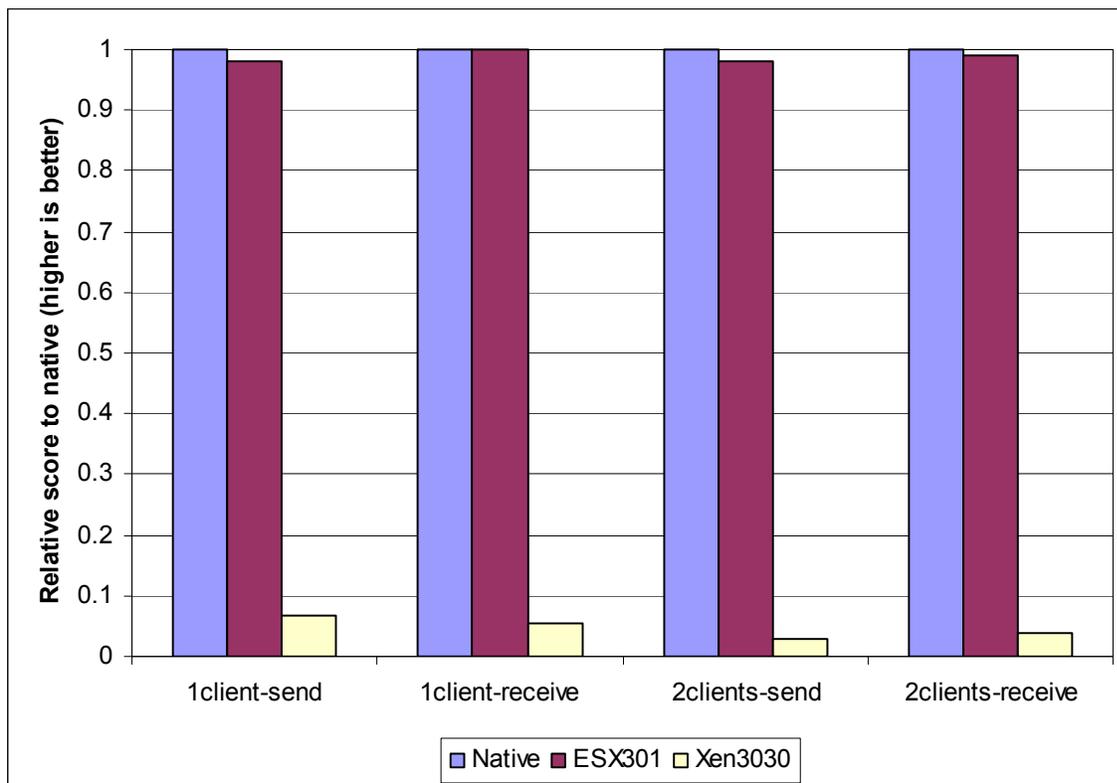


Figure 7 — Netperf results compared to native (higher values are better)

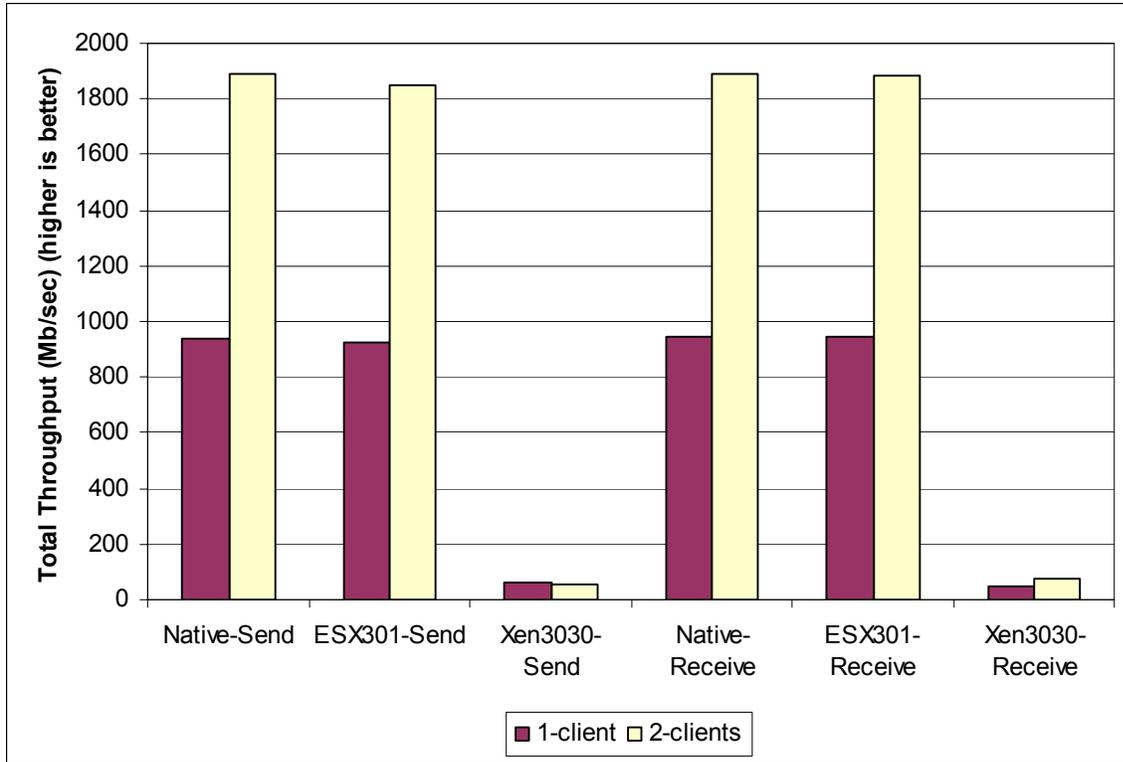


Figure 8 — Netperf throughput scalability results (higher values are better)

SPECjbb2005

The SPECjbb2005 benchmark tests server-side Java Virtual Machine (JVM) performance and does not do any network or disk I/O. No results could be obtained for Xen since it could not boot SMP Windows. Figure 9 therefore compares only the results obtained from VMware ESX Server and native tests. As shown, the VMware ESX Server performance is 91 percent of native using two virtual CPUs, and 88 percent of native when using four virtual CPUs. Since the JVM runs as a single user-level process, direct execution dominates SPECjbb2005's runtime within a virtual machine, just as with the SPECcpu200 integer tests.

Most enterprise applications, such as J2EE application servers, database servers, file servers and mail servers, rely on additional CPU resources to offer increased scalability. The results demonstrate that enterprise customers can deploy VMware ESX Server to scale these applications successfully in virtual environments as well. The Xen hypervisor, on the other hand, is not yet ready for such virtual SMP configurations. Furthermore, these experiments also prove that VMware ESX Server can virtualize enterprise Java applications without any show-stopping performance degradation.

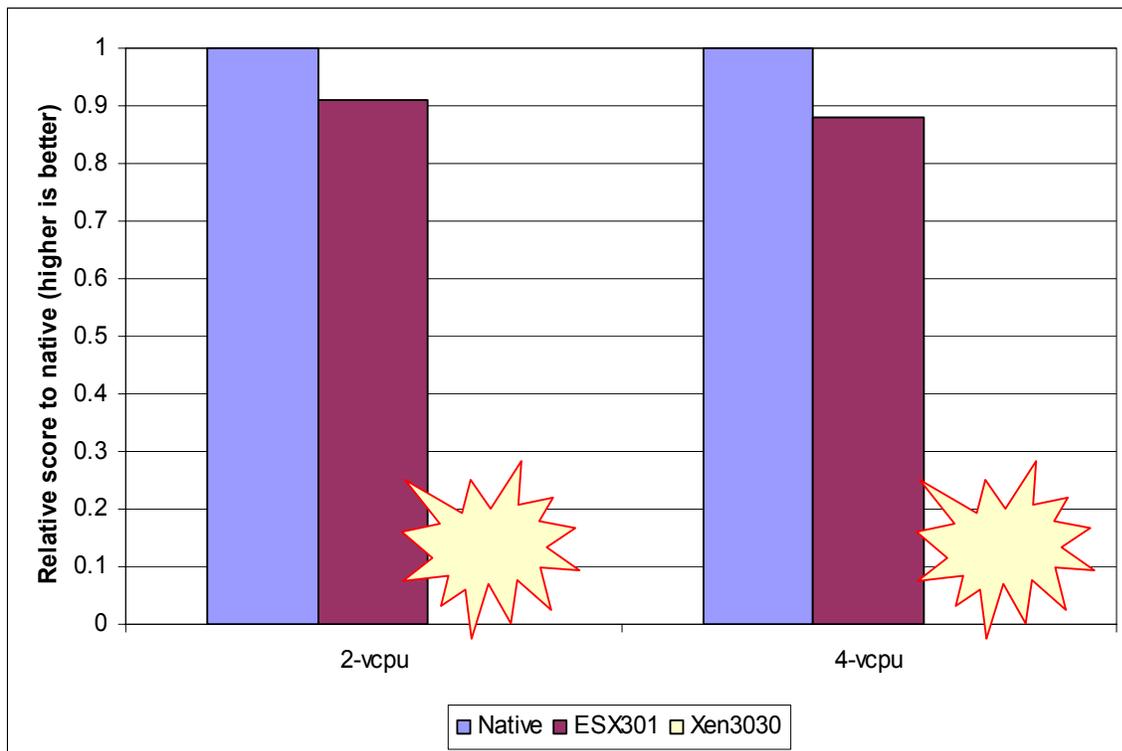


Figure 9 — SPECjbb2005 results compared to native (higher values are better)

Discussion

The objective of this evaluation was to validate the performance and scalability of VMware ESX Server and Xen hypervisors. Unlike the Xen hypervisor, the proven VMware ESX Server hypervisor successfully delivered a combination of performance and scalability requirements necessary for enterprise datacenters. The tests also highlighted several key issues that are important for successful datacenter deployments. Both VMware ESX Server and Xen claim to support virtual SMP configurations for guest virtual machines. Hence the initial plan called for repeating the single virtual CPU tests with configurations featuring two virtual CPUs and four virtual CPUs. Furthermore, the test plan called for testing the scalability of both virtualization hypervisors by running several virtual machines concurrently.

Single Virtual CPU Tests

For both SPECcpu2000 and Passmark –CPU tests, the Xen hypervisor showed on average twice the overhead compared to VMware ESX Server. For enterprise applications that are sensitive to CPU resources, this means that the Xen hypervisor can deliver much lower throughput than VMware ESX Server for the same CPU utilization. Furthermore, most enterprises start implementing virtualization to consolidate underutilized servers. These results imply that VMware ESX Server can support many more virtual machines per core compared to the Xen hypervisor.

The Netperf tests performed extremely poorly for the Xen hypervisor compared to VMware ESX Server. We believe that this happened because the Xen hypervisor lacks an open-source paravirtualized network driver for Windows, similar to the paravirtualized vmxnet driver provided by VMware ESX Server. The commercial versions of Xen are expected to offer paravirtualized network drivers similar to the vmxnet driver. However, such proprietary guest drivers will further add to forking of the open Xen source code and make it difficult for datacenter customers to migrate between various flavors of Xen. Furthermore, unlike the Xen hypervisor, these commercial and supported versions will not be free, and hence will change ROI and TCO calculations that were based on an open-source free offering.

Virtual SMP Tests

The tests for both virtual SMP configuration as well as virtual machine scalability could not be run due to issues with the Xen virtualization hypervisor. A two virtual CPU Windows guest could not be booted using the Xen hypervisor.

The virtual machine scalability tests could not be run because more than two uniprocessor Windows guests could not be booted using the Xen hypervisor. At this time, it is not known when this issue will be fixed and the tests can be tried again.

While Xen claims to support virtual SMP and virtual machine scalability, the results from these experiments demonstrate that enterprise customers should run their own tests to make sure such configurations actually work.

Qualitative Comparison

In addition to performance, customers have identified a number of operational characteristics that are crucial to successful enterprise deployments:

- Maturity of the hypervisor
- Reliability, availability, and serviceability (RAS)
- Scalability
- Management and automation
- Support and maintainability
- Security

While a detailed comparison of ESX Server and Xen is outside the scope of this document, it is worthwhile to provide a brief discussion of how the hypervisors deliver on these operational characteristics.

ESX Server is widely acclaimed for its rock-solid reliability, stability, and maturity. It is a third-generation product whose design and capabilities reflect more than five years of production deployment history at more than 20,000 enterprise customers worldwide. It has been proven to support a variety of operating systems and business-critical applications. VMware customers have reported continuous and ongoing uptime of more than 1,000 days with VMware ESX Server. Xen to date has gained hardly any real-life mileage in datacenter deployments running mission-critical applications that demand negligible downtime

ESX Server is designed to provide the highest level of reliability and availability through a number of powerful RAS capabilities. Examples include capabilities to automatically overcome partial hardware failures as well as full hardware failures through feature such as NIC teaming and bonding, storage mutipathing, etc. ESX Server features an integrated clustered volume manager that enables virtual machines and virtual machine metadata to be stored on enterprise SANs and shared across multiple ESX Server hosts. This architecture enables virtual machines to be restarted on any ESX Server hosts in the event of a server failure. Rigorous interoperability and certification with enterprise SAN as well as host-based replication ensures that virtual machines can be migrated to secondary datacenters in the event of a datacenter or storage failure. These industrial-strength RAS characteristics are yet to be developed for any other virtualization solution, much less proven in the marketplace. Adding to the reliability is rigorous end-to-end certification and interoperability testing of ESX Server with more than 250 systems, storage, and hardware devices that cover the vast majority of networking and storage equipment deployed in datacenters today.

The practical scalability (that is, the number of virtual machines per server in production scenarios) of a hypervisor is largely determined by the hypervisor's ability to make efficient use of system resources, especially system memory. With its unique advanced memory management capabilities such as page sharing and memory ballooning, ESX Server can effectively maximize consolidation ratios and deliver a superior ROI. In addition, the resource management capabilities enable ESX Server to reserve CPU , and I/O resources per guest to maximize scalability while ensuring that each virtual machine has sufficient resources to meet its SLA.

The distributed virtualization capabilities of VMware Infrastructure further enable RAS and scalability beyond the boundaries of a single physical server. Components such as VMware Distributed Resource Scheduler enable flexible allocation of capacity to seamlessly accommodate demand spikes without sacrificing SLA. VMware HA automates recovery from failures of hosts

running ESX Server. VMware Consolidated Backup frees production CPU cycles by offloading backup tasks to a centralized server.

Last, but not least, are the management capabilities required, including multihost management platform, virtual machine lifecycle management, performance management, image management, and APIs that enable enterprise management frameworks to manage virtual infrastructure.

The lack of such RAS, scalability, management, and distributed virtualization capabilities constrain Xen from being a viable, end-to-end enterprise virtualization infrastructure stack.

Future Work

The requirements for datacenter adoption are demanding and far beyond single-server virtualization. Future work under consideration includes more subjective tests covering a wider set of applications. The future tests will also include 64-bit guest operating systems, unmodified and paravirtualized Linux guest operating systems, as well as the scalability tests that could not be done during this round.

Conclusion

IT managers are increasingly looking at virtualization technology to lower IT costs through increased efficiency, flexibility, and responsiveness. As virtualization becomes more pervasive, it is critical that virtualization infrastructure can address the challenges and issues faced by an enterprise datacenter in the most efficient manner. Any virtualization infrastructure looking for mainstream adoption in datacenters should offer the best-of-breed combination of several important enterprise readiness capabilities such as maturity, ease of deployment, manageability and automation, support and maintainability, performance, scalability, reliability, availability and serviceability, and security. We found that VMware ESX Server is far better equipped to meet the demands of an enterprise datacenter than the Xen hypervisor. While Xen-based virtualization products have received much attention lately, customers should take a closer look at the enterprise readiness of those products. The series of tests conducted for this paper proves that VMware ESX Server delivers the production-ready performance and scalability needed to implement an efficient and responsive datacenter. Furthermore, we had no problems setting up and running virtual SMP and virtual machine scalability tests with the reliable and proven third-generation VMware ESX Server hypervisor. Despite several attempts, we were not successful in running similar tests with the Xen hypervisor.

References

1. Adams K. and Agesen O. A Comparison of Software and Hardware Techniques for x86 Virtualization. *ASPLOS* October 2006. http://www.vmware.com/pdf/asplos235_adams.pdf
2. Barnum P., Dragovic B., Fraser K., Hand S., Harris T., Ho A., Nuegebauer R., Pratt I., and Warfield A. Xen and the Art of Virtualization. *Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles*, January 2003
3. Netperf. <http://www.netperf.org/netperf/>
4. Passmark. <http://www.passmark.com/products/pt.htm>
5. Popek, G. J., and Goldberg, R. P. Formal requirements for virtualizable third generation architectures. *Commun. ACM* 17, 7 (1974), 412–421.
6. Standard Performance Evaluation Corporation (SPEC). <http://www.spec.org/cpu2000/>
7. Standard Performance Evaluation Corporation (SPEC). <http://www.spec.org/jbb2005/>
8. University of Cambridge Computer Laboratory
<http://www.cl.cam.ac.uk/research/srg/netos/xen/index.html>
9. Virtual Iron Virtualization Blog.
<http://www.virtualiron.com/fusetalk/blog/blogpost.cfm?threadid=10&catid=3>
10. VMware, Inc. white paper. Virtualization Overview. <http://www.vmware.com/pdf/virtualization.pdf>
11. VMware, Inc. white paper. Virtualization: Architectural Considerations and Other Evaluation Criteria. http://www.vmware.com/pdf/virtualization_considerations.pdf

Appendix A: Test Configuration

Hardware Configuration

Server

- IBM X3500 four-core Intel Woodcrest 3GHz
- 5GB memory (only 1 GB was used for both native and virtual tests)
- 2 73GB 15k RPM SAS drives
- NIC: Intel PRO/1000 MT Dual Port Server Adapter, 8254NXX Gigabit Ethernet Controller

Client (for netperf tests)

- Dell 1600SC, two-way P4 2.4GHz, Windows 2000 Professional
- NIC: Intel Pro/1000 MT Server Adapter
- Interrupt Moderation Rate (IMR): minimum

Native and Guest Operating System Configuration

- Windows Server 2003 Enterprise Edition R2, 1024MB, 1 virtual CPU, 1024x768 video, 32-bit

Native:

- ACPI Multiprocessor HAL, 10GB NTFS partition
- SCSI controller: IBM ServeRAID 8k/8k-I Controller
- Display adapter: ATI ES1000
- Network adapters: Intel PRO/1000 MT Dual Port Server Adapter
- IMR: minimum

VMware ESX Server 3.0.1

- ACPI Uniprocessor HAL, vmdk thick disk on vmdk partition, VMware Tools installed, vmxnet
- SCSI controller: LSI Logic PCI-X Ultra320 SCSI Host Adapter
- Display adapter: VMware SVGA II
- Network adapters: VMware Accelerated AMD PCNet Adapter, MaxTsoSegSize=3952, MinTsoSegCount=2, TsoEnable=1

Xen 3.0.3-0

- Standard PC HAL, 7.6GB file-backed file system on separate disk, qemu harddisk, guest config
- Network adapters: Realtek RTL8139 Family PCI Fast Ethernet NIC, Receive Buffer Size=64KB

Hypervisor Configurations

- ESX301: VMware ESX Server 3.0.1 GA release using BT monitor, 32-bit, esx.conf
- Xen3030: Xen 3.0.3-0 release, Intel VT, 32 bit.
Dom0: XenLinux 2.6.16.29, 32-bit, FC5 distribution, 192MB, kernel build config

Revision: 20070201 Item: PS-004-INF-01-002



VMware, Inc. 3145 Porter Drive Palo Alto CA 94304 USA Tel 650-475-5000 Fax 650-475-5001 www.vmware.com
© 2007 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156, 6,795,966, 6,880,022, 6,961,941, 6,961,806, 6,944,699, 7,069,413; 7,082,598 and 7,089,377; patents pending.
VMware, the VMware "boxes" logo and design, Virtual SMP and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

