



white paper

Proving the Benefits of Deploying VMware ESX Server 3.0.1 on a Scalable Platform

Performance and Economic Advantages Using Unisys ES7000/one Enterprise Servers

Jerrold Buggert, Linda Krasinski

productivity

Implementation of virtual machine (VM) technologies is an emerging trend and initial, small-scale deployments have demonstrated their ability to both improve server utilization and reduce management costs. As such, the next logical step is to employ virtualization in a centralized and disciplined fashion within production environments. But, doing so requires a different mindset and class of hardware.

This white paper compares virtualized workload capacities between the scale-up Unisys ES7000/one Enterprise Server and a scale-out enterprise server approach – providing valuable insight to those organizations that are considering taking virtualization to an enterprise-class level.

Table of Contents

Executive Summary	4
Introduction	5
Test Environment	6
Scale-up Server Platform: 16-processor ES7000/one Enterprise Server	6
Scale-out Server Platform: 4-processor White Box Server	6
Additional Benchmark Environment Hardware Components	7
Virtual Machine Environment and Configuration	7
Benchmark Application Characteristics	8
Test Methodology	9
Test Results	10
Phase One Test Results: Maximum VMs per Server	10
Phase Two Test Results: Variable Workload Performance	10
Economic Implications	12
Approach	12
Results	12
Conclusion	15
About the Authors	15

Executive Summary

The high-level business benefits of virtualization, including reduced infrastructure outlays, increased responsiveness and flexibility, and consistent and predictable operations, are clear enough that the technique has become fairly commonplace in test and development environments. However, enterprises still struggle with doing it cost-effectively on any real grand scale, and as such, IT organizations have used it in a limited and opportunistic fashion – mostly migrating legacy and other generally small workloads to virtual instances.

Although these deployments are on smaller 2- and 4-processor servers, they have resulted in significant advantage compared to the one-application-per-server paradigm. In other words, there appears to be additional and tremendous opportunity within the data center to improve infrastructure by further and fully leveraging virtualization.

This initial success with regard to virtualization, combined with the increasing maturation of the virtualization layer and associated management capabilities, has made it an attractive prospect for enterprise deployments, as well. However, employing virtualization in a centralized and disciplined fashion in production environments requires a different mindset and class of hardware. Thus, there are many technical issues to consider when choosing a target server for enterprise-class virtualization. They include:

- The ability to host the largest workloads and support greater amounts of memory
- Additional capabilities that enhance the overall system resiliency and reduce or minimize the effects of a broader range of potential sources of error
- The agility to address workload shifts and creation of physical computing entities within a system
- Very high levels of redundancy to mitigate hardware failures

This white paper is intended to demonstrate an important dimension of optimized virtualization infrastructure and its results should be of significant relevance to organizations interested in deploying virtual machines (VMs) in a consistent, standardized, and optimized manner within enterprise environments. Specifically, we endeavor to compare virtualized workload capacities between the scale-up Unisys ES7000/one Enterprise Server and a scale-out enterprise server approach.

And our data shows that the Unisys ES7000/one Enterprise Server with 16 dual-core Intel® Xeon® processors running VMware® ESX Server™ 3.0.1 not only helps maintain better service quality, but also *improves total cost of ownership (TCO) by 33 percent*. It reduces costs for VMware ESX Server 3.0.1 licensing, client operating system licensing, power and cooling, platform/host management, and network and disk connectivity. The scale-up Unisys ES7000/one Enterprise Server can host virtual machines at a greater density (VMs per processor) than is possible on a system with four dual-core processors.

The conclusion is that the ES7000/one Enterprise Server and VMware ESX Server 3.0.1 are a very suitable combination for a large scale, enterprise-class virtualization solution and is more cost effective than the 2- and 4-processor servers typically used for VMware solution deployments.

Introduction

This technical paper is intended for IT professionals who are responsible for selecting an optimal VMware ESX Server 3.0.1 deployment solution or preparing and deploying an ES7000/one Enterprise Server hosting VMware ESX Server 3.0.1. As such, it is assumed that the reader has some knowledge of virtualization and VMware ESX Server 3.0.1.

Most systems do not run with a constant workload. And instead, the workload fluctuates throughout the day or business cycle. It is expected that the server platform hosting VMware ESX Server 3.0.1 can handle these fluctuations transparently and with minimal impact to end user quality of service. In this paper, however, Unisys demonstrates the scale-up ES7000/one Enterprise Server with VMware ESX Server 3.0.1 is better equipped to handle workload fluctuations than smaller, scale-out servers.

We also demonstrate what happens to a four dual-core¹ processor VMware ESX Server 3.0.1 host experiencing a peak period for work requests, when it can become overdriven to the point that transaction response time (RT) becomes unacceptable. In contrast, we show that the 16 dual-core processor ES7000/one Enterprise Server can handle such peaks in work requests with less than half the average transaction RT than the 4-processor host during this same peaking period. And this is all accomplished while non-peaking VMs on the ES7000/one Enterprise Server maintain an acceptable transaction RT.

For a given workload, the 16-processor ES7000/one Enterprise Server is managing 128 VMs or 8 VMs per processor versus the 4-processor host, which supports 12 VMs or 3 VMs per processor. This disparity represents nearly 11 times more VMs or nearly 3 times more VMs per processor. The capability to house more VMs per processor, while better managing fluctuating workloads, is a significant reason to consolidate VMs onto a fully configured ES7000/one Enterprise Server.

Unisys ES7000/one Enterprise Server

The ES7000/one Enterprise Server offers a flexible, single server solution that can be configured to meet enterprise workloads as they change. With the ability to grow from 4 to 32 Intel® Xeon® processors MP™, or from 4 to 32 Intel® Itanium® 2 processors, the ES7000/one Enterprise Server scales beyond conventional Intel® processor-based platforms while delivering improved price/performance.

The ES7000/one Enterprise Server also has the ability to create up to eight independent partitions that concurrently run different operating systems, making it an excellent solution for consolidating applications, migrating from expensive proprietary RISC systems, deploying business intelligence applications, and running large-scale databases.

Additional and more detailed product information is available on the Unisys website at <http://www.unisys.com>.

The benchmark test described in this paper reveals the potential of a fully-configured ES7000/one Enterprise Server hosting VMware ESX Server 3.0.1. Because the VMs and application in our test represent only one type of virtualized environment, Unisys engineering recommends that a customer's anticipated mix of virtual machines and applications is assessed using appropriate capacity planning tools prior to deployment.

¹ All references to processors in this paper are to dual-core processors.

Test Environment

The primary system under test was the ES7000/one Enterprise Server running VMware ESX Server 3.0.1-30353 Release Candidate. This system was evaluated in relationship to a server with four AMD processors.

Scale-up Server Platform: 16-processor ES7000/one Enterprise Server

The ES7000/one Enterprise Server was configured with 16 processors (32 cores, 32 logical processors) and 64GB of memory – the maximum processor and memory allowed with VMware ESX Server 3.0.1. Therefore, hyper-threading technology (HTT) was disabled on all 16 processors for the benchmark tests. (Note: VMware ESX Server 3.0.1 does support HTT.)

Figure 1 summarizes the ES7000/one Enterprise Server system components.

Figure 1:
ES7000/one Enterprise Server Profile

System Component	Description
Software	VMware ESX Server 3.0.1-30353 Release Candidate
Processors	16 dual-core Intel® Xeon® processors 7140M (3.4GHz), 16KB L1, 1MB L2 per core, 16MB L3, hyper-threading disabled
Memory	64GB RAM plus 48MB shared cache per four processors
Network Interface Cards (NIC)	6 Intel® PRO/1000 F Server Adapters 1 Intel® 8254NXX Gigabit Ethernet Controller (inbuilt, both ports used)
Host Bus Adapter (HBA)	Emulex LightPulse LP11000 HBA
Disk subsystem	EMC® CLARiiON® CX300 networked storage system, 45 disks, 72GB per disk, 15,000 RPM, 5 LUNs, RAID 0

The following changes were made in the BIOS to the Platform/OS Control section:

- Linux OS: YES
- IO Space Enable: ALL
- P64H2 1K I/O Option: YES
- Hot Addable Memory: Disable

Scale-out Server Platform: 4-processor White Box Server

The ES7000/one Enterprise Server performance was evaluated in relationship to a “white box” lab system assembled from off-the-shelf components, including four AMD processors. Figure 2 summarizes the white box server system components.

Figure 2:
White Box Server Profile

System Component	Description
Software	VMware ESX Server 3.0.1-30353 Release Candidate
Processors	4 Dual-core AMD Opteron 2.2-GHz processors, 64KB L1, 1MB L2 per core
Memory	16GB RAM
NICs	2 Intel® PRO/1000 XF Server Adapters 1 Broadcom BCM5704C adapters (inbuilt, both ports used)
HBA	Emulex LightPulse LP11000 HBA
Disk subsystem	Eurologic Systems fiber channel network storage system, 13 disks, 18GB per disk, 15,000 RPM, JBOD

Additional Benchmark Environment Hardware Components

The benchmark environment included a number of supplementary hardware components as outlined in Figure 3.

Figure 3:

Additional Hardware Components

System	Function	Hardware Description
Database server 1	Processes database queries from the applications running on the virtualized machines	White box lab server with four AMD Opteron processors, 2.4GHz, 8GB RAM, 1-Gb NIC, running Microsoft® Windows® Server 2003 Enterprise x64 Edition operating system
Database server 2	Processes database queries from the applications running on the virtualized machines	Unisys ES3020L server with two Intel® Xeon® processors, 3.2GHz, 12GB RAM, 1-Gb NIC, and running Windows Server 2003 Enterprise Edition operating system
Business logic emulator server	Runs three emulators – Payment Gateway, Purchase Order Verification, and Shipment Notification – which provide data to the application on the virtualized machines via web services requests	Unisys ES3020L server with two Intel® Xeon® processors, 3.2GHz, 12GB RAM, 1-Gb NIC, and running Windows Server 2003 Enterprise Edition operating system
Remote business emulator servers	Run the Emulated Businesses (EBs) that drive the benchmark workload	43 Dell OptiPlex GX280 PCs with one Intel® Xeon® processor, 3.2GHz, 512MB RAM, 1-Gb NIC (used at 100Mb speed), and running Microsoft Windows XP Professional with Service Pack (SP) 2 operating system

Virtual Machine Environment and Configuration

VMware ESX Server 3.0.1-30353 Release Candidate was used. No tunings were applied to the VMware ESX Server 3.0.1 software.

One template VM was created. On the ES7000/one Enterprise Server, the template VM was replicated 127 times; resulting in a total of 128 VMs. On the white box server, the template VM was replicated 11 times; resulting in a total of 12 VMs. Each VM had the following characteristics:

- OS: 32-bit Windows Server 2003 Standard Edition operating system
- VMware Tools: Installed
- Memory: 512MB
- Processors: 1 virtual processor
- Network: Two 1-Gb NICs; VMware Accelerated AMD PCNet Adapters
- Disk: 5GB
- Application: Refer to the following section

Benchmark Application Characteristics

Our testing used an application that portrays a retail distributor on the Internet with ordering capabilities and product-browsing scenarios. The application accepts incoming web service requests from other businesses or store fronts to place orders, view catalog items, make changes to the catalog, update or add customer information, and request the status of an existing order. The majority of requests generate purchase order activity with a small portion leading to additional item or catalog information inquiries.

The application environment is characterized by:

- Multiple on-line business sessions
- Commercially available application environment
- Use of XML documents and SOAP for data exchange
- Business-to-business application logic
- Reliable and durable messaging
- Dynamic web service response generation with database access and update
- Simultaneous execution of multiple transaction types that span a breadth of business functions
- Databases consisting of many tables with a wide variety of sizes, attributes, and relationships
- Transaction integrity (ACID properties)

The application was derived from the TPC-App benchmark. The business workload in this benchmark has a wide variety of transaction types and a variable distribution of transaction peaks and valleys. The TPC-App benchmark application was modified to better align with the typical profile we see within the Unisys customer and prospect base, particularly in the area of IO. For example, the amount of data read in a search and size of images transmitted over the network were reduced. Because of these modifications, the results of this testing are not comparable to published TPC-App benchmark results.

There are seven different types of web service requests (transactions) from the emulated businesses (EBs) to the application:

- New Customer
- Change Payment Method
- Create Order
- Order Status
- New Products
- Product Detail
- Change Item

Some of these requests result in the application sending web service requests to three external servers: a Payment Gateway emulator, a Purchase Order Verification emulator, and a Shipment Notification emulator. These servers are on one physical system in our test environment.

The application on each virtualized machine interacts with an external backend database server. We used two smaller database servers versus one large database server and statically assigned VMs/applications to one or the other database to balance the load.

The application runs in Microsoft .NET Framework Common Language Runtime (CLR) environment, uses Microsoft Message Queuing (MSMQ) service, and does not use the SSL protocol.

The workload referred to in this paper is generated from the web service requests (transactions) from EBs to the applications. When the number of EBs is increased, so is the workload. During benchmark testing the workload generated an average throughput per VM of 114KB/sec for disk and 405KB/sec for network. The average RT to web service requests is used as the performance metric for the results in this paper and is reviewed in more detail in the following section.

Test Methodology

To show the potential of the 16-processor ES7000/one Enterprise Server, a variable workload consisting of valleys with low utilization and peaks with high utilization was applied and measured against the ability of the 4-processor white box system to handle the same variations in workload. The key differences between the two systems are outlined in Figure 4.

One significant difference between the two systems is that the ES7000/one Enterprise Server can host nearly 11 times more VMs than the 4-processor server – and with only four times the number of processors. Refer to the Phase One testing results on this page for more details on how the number of VMs per server was determined.

Our testing method is to start with a workload that leaves ample headroom for increase. At this base (or valley),

the workload is equally divided across each VM. After a pre-defined elapsed time, the workload on some VMs is increased. For the 4-processor white box host, all 12 VMs experience an increased workload. For the 16-processor ES7000/one Enterprise Server, 12 of the 128 VMs experience the same increased workload (equivalent to the peak load of one 4-processor system). The workload on the other 116 VMs remains constant.

The test was scripted to mimic real-world, business cycle-driven conditions – reducing and increasing the workload to create valleys and peaks. The peaking workload is 14 times greater than the base (or valley) workload for the 12 peaking VMs.

The average transaction RT during the valleys and peaks are calculated for each server and compared. An average RT of less than 1 second is considered acceptable.

Figure 4:

Test System Comparison

Test System	Type of Processor	Number of Processors	Memory	VMs
ES7000/one Enterprise Server	Dual-core Intel® Xeon® processor 7140M 3.4GHz	16 processors, 32 cores	64GB	128 VMs
White box server	Dual-core AMD Opteron processor 2.2GHz	4 processors, 8 cores	16GB	12 VMs


Test Results

Testing was performed in two phases. During the first phase, the workload per VM was kept constant while the number of VMs was increased. The objective of this test was to determine the maximum number of VMs supported while maintaining RTs within the one-second specification. The second phase used the information gained from Phase One (maximum number of VMs per server) and studied the impact of workload peaks and valleys on RT. The results of both testing efforts are reviewed in the following sections.


Phase One Test Results: Maximum VMs per Server

During Phase One, tests were conducted to determine the maximum number of VMs the scale-up and scale-out systems could support when running the benchmark application – while keeping within the specified one-second RT.

The four-processor white box host was configured with 8, 12, or 16 VMs and the peak workload per VM was applied. Figure 5 shows the average RT exceeded 1 second with 16 VMs. 12 VMs, with less than 1 second RT, was chosen as the optimal number of VMs – 3 VMs per processor.



The 16-processor ES7000/one host was configured with 112, 120, or 128 VMs and the peaking workload was applied. Figure 6 shows the average RT remained constant – *and well below one second* – for all configurations. With 128 VMs being the maximum allowed with VMware ESX Server 3.0.1, the 16-processor host can manage 8 VMs per processor. Compared to 3 VMs per processor for the four-processor host, the 16-processor ES7000/one Enterprise Server supports nearly three times the number of VMs per processor.



Phase Two Test Results: Variable Workload Performance

This section presents Phase Two test results and discusses the performance of the 4-processor and 16-processor servers when presented with the same variable workload. As mentioned earlier in the Methodology section, such workloads are more representative of real-world conditions where transaction volumes are oftentimes unpredictable.

Not surprisingly, Phase Two tests showed that as workload increased, transaction RT also increased, and as the workload decreased, so did the RT.

Figure 7 charts the valley and peak workload applied to the 4-processor white box server. It also shows that at the peak workload, the RT skyrocketed to above 2 seconds as the host struggled to accommodate the volume.

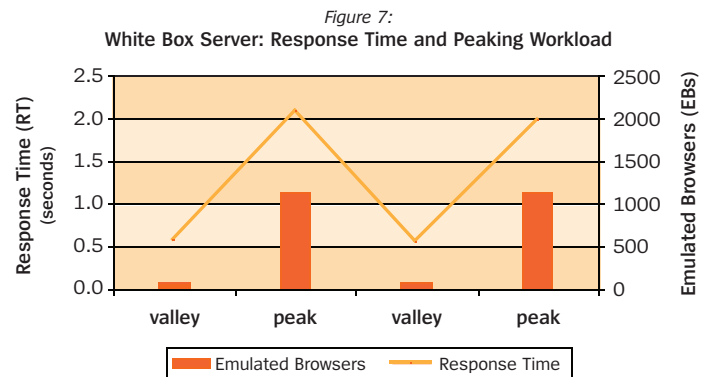
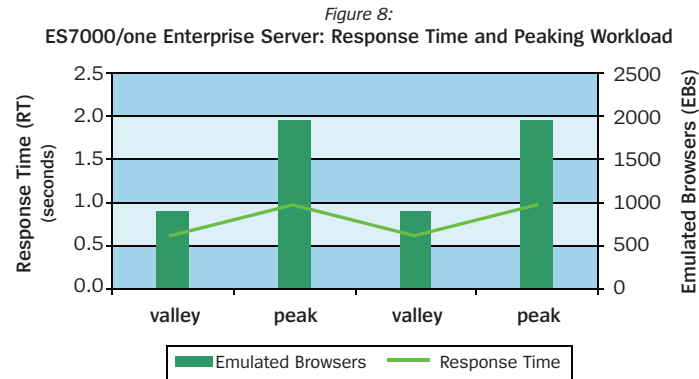
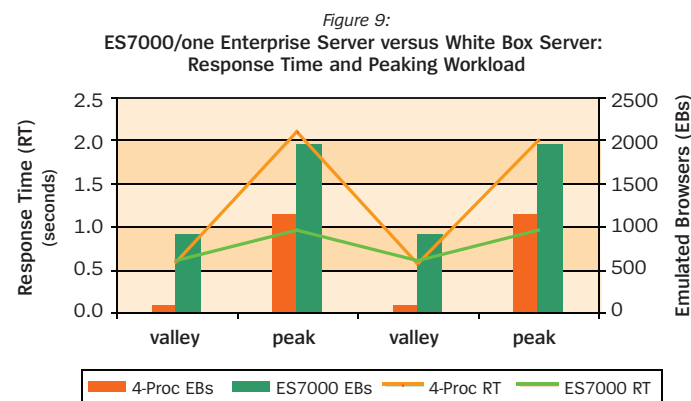


Figure 8 shows the valley and peak workload applied to the 16-processor ES7000/one host. The chart also shows that at the peak workload, this platform maintained the less than one second RT criteria.



The chart in Figure 9 combines the benchmark data from both servers into a single view. It clearly shows that, while the workload is peaking, the 16-processor host maintains less than half the RT of the 4-processor host. This is because the ES7000/one Enterprise Server offers better peak workload handling capabilities, including ramp up and ramp down of the workload. And, this is in addition to the fact that the 16-processor host maintained the workload of an additional 116 non-peaking VMs. Note that the ES7000/one valley – or base – workload is much greater than that of the 4-processor server (896 EBs versus 84 EBs).



Important Update

Since this benchmarking work was conducted in the Unisys labs, AMD released a 2.8-GHz Opteron processor with improved performance. In order to provide a fair and contemporary comparison between the scale-up and scale-out servers, a 33 percent performance adjustment has been made to the scale-out 4-processor server. This adjustment changes the virtual machine density from the three-VMs-per-processor level observed in the benchmark testing to four-VMs-per-processor. This adjustment is reflected in the economic analysis described in the following section.

Economic Implications

Compelling though the raw results of this testing are, however, technical considerations alone do not provide a complete basis for a decision on major IT infrastructure deployment models. For this reason, an economic analysis was undertaken that describes and compares the costs associated with the infrastructure required to host a significant number of workloads/applications in both a scale-up and scale-out fashion.

Approach

A total workload of 128 VMs was selected as a representative figure since this is the number that can be comfortably deployed on a single scale-up server. Per our benchmark testing, and making a subsequent performance adjustment as noted on page 11, eight servers, each with four processors, are required to host this same number of VMs using a scale-out model.

Six primary categories of capital and operating expense were evaluated and included in the analysis, all at retail list selling price to provide an objective basis for comparison:

- Platform hardware/maintenance cost
- Connectivity switching cost
- VMware license/support costs
- Operating system/support costs
- VMware host management costs
- Power and cooling costs.

A three-year time horizon was chosen since the implications of a given architecture decision generally have a significant impact on operational expenses. The specific factors affected by this time period assumption are hardware and software support/subscription, as well as host management. A 24x7 level of support was used since this is consistent with a data center, enterprise class operating model. As a simplifying assumption, the analysis was performed on a cash basis (no depreciation, annualized cash payments for support, or time value of money was considered).

Results

The results of this analysis (summarized in Figure 10 on page 13) are quite telling. *The total benefit of deploying on a scale-up architecture is 33 percent over a three-year period.* In other words, deploying in a scale-out virtual machine environment can cost nearly one-third as much as deploying on a scale-up model over a three-year period.

As mentioned, in order to host 128 VMs, one 16-processor server is required while eight 4-processor servers are required. This equates to 16 processors to host the defined workloads in the scale-up scenario compared to 32 processors in the scale-out scenario. Virtual machine density per processor on the scale-up server is therefore two times that of the scale-out architecture: eight compared to four. The majority of the economic benefit highlighted in this analysis derives from this fact: VMware ESX Server 3.0.1 cost, OS cost, and power and cooling are directly related to the number of processors. Host management and connectivity costs are more closely related to the number of VMware ESX Server 3.0.1 virtual hosts and physical servers respectively, which provides an even greater benefit to the scale-up model.

But, the one area where the scale-up model does not trump the scale-out is in platform cost. The reason for this is clear. These servers have been designed for the one-application-per-server model, while the scale-up ES7000/one Enterprise Server has been designed as a robust, scalable, reliable, and available platform for enterprise-class deployment of core and mission-critical workloads, such as databases and large business processing applications, as well as hosting a large number of workloads through virtual machine technology or other consolidation methodologies. (Refer to page 14 for more details.)

Figure 10:
Economic Analysis

	Scale-up Model	Scale-out Model
	16 processors	32 processors
	128 VMs	128 VMs
Base Platform (includes three-year maintenance and support cost)	204,715	178,544
Power/Cooling (three-year cost)	26,910	48,941
VMware Cost	80,500	161,000
Connectivity Costs	8,700	24,800
Management Costs - VMware ESX Server 3.0.1 Host Mgmt (three-year cost)	4,680	12,480
Microsoft Windows OS Requirement, Datacenter Edition	62,021	156,422
Total Three-year Cost	387,527	582,207
Three-year TCO/VM	3,028	4,548
Annualized Cost/VM	1,009	1,516
Percent TCO Savings	33.4%	

Enterprise-class Versus Standard High-volume Servers

What makes a server “enterprise class” and how does it differ from a standard high-volume server?

	Standard High-volume Server	ES7000/one Enterprise Server
Performance Features		
Scalability	2 or 4 processors	Up to 32 processors
Shared (processor external) cache to improve performance	No	48MB per four Intel® Xeon® processors, 64MB per four Intel® Itanium® 2 processors
Maximum memory	64GB (typical)	512GB
Directory-based memory to improve data access times	No	Yes
Manageability Features		
System monitoring	Yes	Yes
“Call home” in the event of HW and/or OS issues/failures	No	Yes
Error detection	Error-Correcting Code (ECC) memory	ECC memory, as well as full data path data integrity checking
Flexibility Features		
Partitioning	Only through virtualization, limited in size	Supports both virtual and hard partitions, which can be any size
Capacity on demand	No	Yes
Support for Xeon® and Itanium® processors in the same system	No	Yes
Availability Features		
N+1 power and cooling	Yes	Yes
Redundant system clock	No	Yes, optional
Redundant service processors	No	Yes, optional
Redundant cache	No, no cache option available	Yes
Intelligent memory subsystem to improve error handling	No	Yes

Conclusion

The ES7000/one Enterprise Server, configured with the maximum dual-core² processors and memory allowed by VMware ESX Server 3.0.1 is capable of handling peaks from variable workloads more efficiently than an aggregate of multiple, four-processor configurations. The capacity of more processors in a single host affords the VMs more flexibility, resulting in greater productivity and creation of the most efficient VMware ESX Server 3.0.1 host.

Our test showed that a 4-processor configuration did not have the ability to handle the workload peaks resulting from a variable workload within the criteria of less than one second average RT. In fact, the RT was over two seconds at the peak – more than double the acceptable criteria.

In contrast, our test showed that the 16-processor ES7000/one Enterprise Server, managing a large number of VMs, can accommodate the workload peaks from a variable workload within the criteria of less than one second average RT. In fact, the RT at the peak was less than half the RT of the 4-processor host at the peak. When there was an increase in the workload on some of the VMs, the corresponding demand for more processor cycles was easily absorbed.

The economic impact of this disparity is a 33 percent improvement in total cost of ownership that also ensures better service quality. These savings are due to reduced costs for VMware ESX Server 3.0.1 licensing, client operating system licensing, power and cooling, platform management, and network and disk connectivity.

As businesses and public sector organizations begin to broadly adopt virtualization for the heavy lifting of their core and mission-critical workloads, they are realizing that the underlying platform must be up to this challenge. In order to best take advantage of powerful virtual machine technology and drive cost per virtual instance to the lowest level, deployment on a scale-up platform, such as the ES7000/one Enterprise Server, is strongly indicated.

About the Authors

Jerrold Buggert

Jerrold Buggert has been managing performance and technology evaluation groups for the past 13 years. He is currently the director of the Unisys PDT Performance Group, which does characterization, analysis, modeling, and measurement of all the Unisys product lines. His group is responsible for providing modeling of hardware and software in development, generation of capacity planning data, proof points, and marketing performance data, including Transaction Processing Performance Council (TPC) benchmark results. The group is also involved in performance analysis and optimization of large customer environments, such as NASDAQ and Amadeus. He is the primary Unisys representative on the TPC; a member of the TPC steering committee, TPC Public Relations committee, and Pricing subcommittee; and the chairman of CPU definition workgroup. Previously, he was the chairman of the TPC-W Ecommerce benchmark development subcommittee and the chairman of the TPC.

Linda Krasinski

Linda Krasinski has been with Unisys for 28 years. Most recently, she has worked with Unisys ES7000 servers and Windows Server operating systems in the areas of benchmarking, measurement, system analysis, system performance, and networking. She also recently worked with SPECweb, Windows Terminal Server, and various virtualization technologies.

² Reminder, all references to processors in this white paper are to dual-core processors.

Contributors

Rick Freeman

Rick Freeman has been involved with performance analysis and measurement for nearly 30 years. Most recently, he has managed projects to publish industry-standard TPC and SPEC benchmarks, as well as to produce technical white papers and to characterize, analyze, and improve the performance of large-scale Intel® processor-based systems. He is a member of the TPC Technical Advisory Board and is Unisys Primary Representative to Standard Performance Evaluation Corporation (SPEC).

Steve Morris

Steve Morris has more than 30 years of experience in computing. Most recently Steve's work at Unisys has concentrated on load generation for the TPC-W, TPC-App, and TPC-E benchmarks. He has also been an active participant in the subcommittees that developed the specifications for the TPC-W and TPC-App benchmarks.

Alan Chan

Alan Chan has been with Unisys for eight years. He has made significant contributions to the development of TPC-W and TPC-App benchmarks. Most recently, he developed a test package to evaluate VMware ESX Server 3.0 and has been involved with VMware ESX Server operations.

For more information, contact your Unisys representative.

Or call Unisys today at:

1-800-874-8647, extension 765 (U.S. and Canada)

00-1-585-487-2430 ext 765 (Other countries)

In a hurry to learn more? Visit:

<http://www.unisys.com/tech/virtualization>

Specifications are subject to change without notice.

© 2007 Unisys Corporation.

All rights reserved.

Unisys is a registered trademark of Unisys Corporation. Microsoft and Windows are registered trademarks of Microsoft Corporation. Intel, Itanium, and Xeon are registered trademarks of Intel Corporation. VMware is a registered trademark and ESX Server is a trademark of VMware, Inc. TPC and TPC-App Benchmark are trademarks of the Transaction Processing Performance Council. EMC and CLARiiON are registered trademarks of EMC Corporation. Linux is a trademark of Linus Torvalds. All other brands and products referenced herein are acknowledged to be trademarks or registered trademarks of their respective holders.

