WHITE PAPER

Microsoft Exchange Server 2007

Performance on VMware vSphere™ 4

vmware®

Table of Contents

Introduction 3

Executive Summary 3

Experiment Configuration and Methodology 3

Testbed Configuration 3

Test and Measurement Tools 5

Test Cases and Test Method 5

Experimental Results and Performance Analysis 6

Single Virtual Machine Performance 6

Exchange Latencies 6

Processor Utilization 7

Consolidated Performance 8

Latency Fairness in VMs 8

Processor Utilization 9

Best Practices 10

Conclusion 10

About the Author 10

Acknowledgements 10

Introduction

Increased adoption of VMware infrastructure and continued performance advances in hardware and software have driven the virtualization of Microsoft Exchange. The current version, Microsoft Exchange Server 2007, is a resource-hungry 64-bit application that has reduced storage demands as compared to previous versions. The increased use of memory and heavier CPU load requires efficient virtualization capabilities to meet customer demands. This paper will detail experiments that demonstrate the efficient performance and scalability of Exchange when running inside a virtual machine on a VMware vSphere™ 4 environment.

The key performance requirements for Exchange virtualization are efficiency of single instance virtualization, uniform and predictable performance in consolidation environments, and fair distribution of resources to multiple virtual machines (VMs). Acceptable single and multiple VM performance is required to maintain a positive user experience and assist virtualization and Exchange administrators with capacity prediction and maintenance. Resource distribution fairness allows administrators to identify performance bottlenecks in a predictable manner. With unfair resource distribution it can be difficult to find and isolate the source of user concerns.

This paper addresses all three of these performance characteristics:

- The performance implications of running the Exchange 2007 Mailbox Server role in a virtual environment versus on a physical system.
- The performance of virtualized Exchange 2007 Mailbox Server role when scaling up (adding vCPUs to a virtual mailbox instance), and scaling out (adding virtual mailbox instances to the host).
- Data showing fair virtual machine management, which provides similar service levels for multiple mailbox roles in a consolidated environment.

Some performance best practice recommendations for the virtualized Exchange environment are also provided.

Executive Summary

The tests in this paper use Microsoft's recommended Exchange sizing guidelines both in the virtualized and physical environments. The experiments compare the Exchange 2007 Mailbox Server role in virtual and physical environments. In both cases, the mailbox role is running on an HP ProLiant DL580 connected to an EMC CX3-40 fiber channel attached storage. All peripheral server roles (Active Directory, DNS, Hub Transport, and Client Access Server) are on a Dell PowerEdge 2950 physical system. The Microsoft Exchange Load Generator (LoadGen) Heavy User profile was used to stress the System Under Test (SUT).

Key observations and takeaways from the paper are that the VMware vSphere 4 can virtualize large Exchange deployments with only a fraction of the CPU resources available on modern servers. vSphere is shown to provide near-linear scale-up and scale-out capabilities, which allows administrators to deploy Exchange in configurations with multiple small mailbox instances or with fewer large ones. Latencies observed by the user are less than half the acceptable latencies as recommended by Microsoft.

Experiment Configuration and Methodology

The performance and sizing studies were done in VMware's internal labs. The purpose of the tests was to measure, analyze, and understand the performance of Exchange in both the physical and virtual environments. In the following sections, the test bed configuration used for the experiments is described in detail, and the test tools are discussed. Finally, a description of the experiments is presented.

Testbed Configuration

The HP ProLiant DL580 server is configured with two Quad-Core Intel Xeon X7350 processors and 64 GB of physical Memory. Exchange Server 2007 SP1 was installed on Windows Server 2008 in both physical and virtual environments. The virtual environment used VMware vSphere 4 to virtualize the Windows/Exchange environment.

Here is the detailed test bed configuration:

System Under Test (SUT)

- Server: HP ProLiant DL580
- Processors: 2 Quad-Core Intel Xeon 2.93 GHz X7350
- Memory: 64GB DDR-2 DIMM

- HBA: Emulex HBA 4Gb PCI-Express
- Virtualization Software: VMware vSphere 4.0
- Operating System: Microsoft Windows Server 2008 Datacenter Edition (64 bit)
- Application: Microsoft Exchange Server 2007 SP1 (64 bit)

Storage

- Storage: EMC CX3-40 with two storage processors
- Hard Drives: 120 146GB 15K RPM drives
- Disk configuration for Exchange: 104 disks for DB and 16 for logs.

Client

- Client: Dell PowerEdge 1950
- Processors: 1 Quad-Core Intel Xeon 2.66 GHz X5355
- Memory: 16GB DDR-2 DIMM
- Operating System: Microsoft Windows Server 2003 Enterprise Edition (64 bit)
- Application: Microsoft Exchange Load Generator

Server Hosting Active Directory, DNS, Exchange Hub and CAS Roles

- Server: Dell PowerEdge 2950
- Processors: 2 Quad-Core Intel Xeon 2.66 GHz X5355
- Memory: 32GB DDR-2 DIMM
- Operating System: Microsoft Windows Server 2008 Enterprise Edition (64 bit)
- Application: Microsoft Exchange Server 2007 SP1 (64 bit)

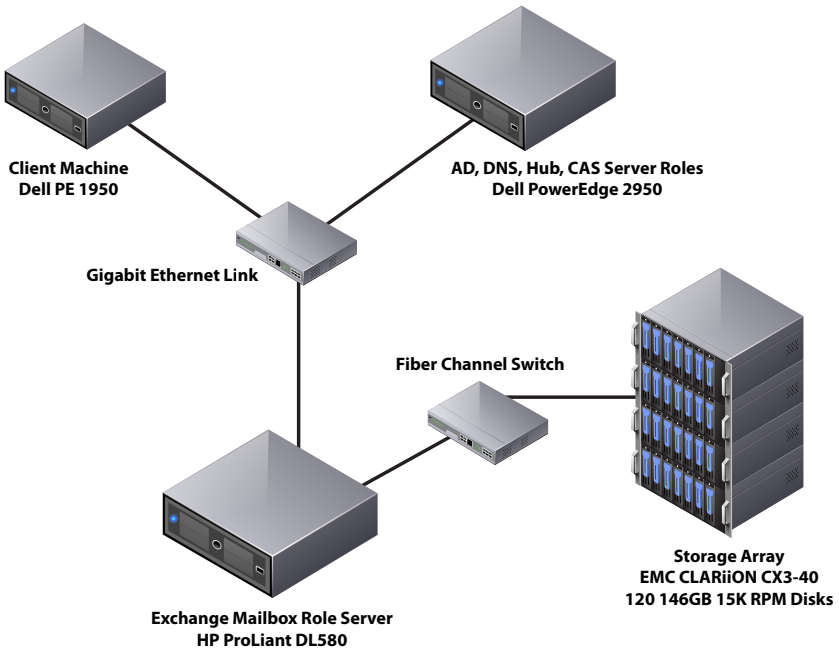

Test and Measurement Tools

Microsoft's Exchange Load Generator (LoadGen) runs on a client system and simulates the messaging load for both physical and virtual Exchange environments. LoadGen provides tools to measure system response for a given number of Exchange users. In this case LoadGen was configured to use Messaging Application Programming Interface (MAPI) clients using the LoadGen Heavy User profile. Results included a variety of Exchange transaction latencies during an eight hour run period, selected to represent a typical workday. Industry consensus exists in using the Send Mail 95th percentile transaction latency as an accurate measure of mailbox performance. The value of this result, which represents the maximum time needed to send an email for 95 percent of the transactions, should be reported below 500 ms to represent an acceptable user experience.

The LoadGen tries to closely model the normal daily email usage of real users, in order to provide an estimate of the number of users a system can support. While this workload is widely used to measure the performance of Exchange platforms, as with all benchmarks, the results may not match the specifics of your environment.

In this example, Microsoft Windows Perfmon (http://technet.microsoft.com/en-us/library/bb490957.aspx) was used to examine performance levels of the physical Exchange configuration. Perfmon was configured to log relevant CPU, memory, disk, network and system counters, as well as Exchange-specific counters.

For VMware ESX 4.0, esxtop was used to record both ESX Server and virtual machine related performance counters. Esxtop was configured to log CPU, memory, disk, network and system counters during the LoadGen runs.

Test Cases and Test Method

The two primary objectives in performing these experiments were:

1. Compare the performance and scalability of physical Exchange installations with their virtual counterparts, in a single virtual machine. For this, the following experiments were conducted:
   - Physical scale up experiments with 1, 2, 4, and 8 physical processors
   - Single VM scale up experiments with 1, 2, 4, and 8 vCPUs
2. Understand the scalability of the Virtual Machines (VMs), with respect to the number of VMs and virtual CPU count. The following experiments were conducted:
   - 2 vCPU VMs scale out with 1, 2, 4, 6, and 8 VMs
   - 4 vCPU VMs scale out with 1, 2 , 3, and 4 VMs

User and resource (processor and memory) sizing were based on Microsoft's recommendations (http://technet.microsoft.com/en-us/library/aa998874.aspx and http://technet.microsoft.com/en-us/library/bb738124.aspx). 500 Heavy users were configured per processor and sized memory to 2 GB plus 5 MB per mailbox. As an example, 2,000 Heavy users require 4 CPUs and 12 GB for the virtual machine or physical environment.

The table below shows the number of CPUs, memory size and number of Heavy users for each scenario:

Number of Virtual/ Physical CPUs	Memory Size (GB)	Total Number of Heavy Users
1	4.5	500
2	7	1,000
4	12	2,000
8	22	4,000

In the consolidated environments, two different virtual machine configurations were chosen, based on number of vCPUs, to demonstrate scale-out capability. Whenever the total number of vCPUs from all VMs exceeds eight, physical CPU resources of the system are overcommitted. The degree of overcommitment is 1.5x for 6x2 and 3x4 scenarios and goes up to 2x for 8x2 and 4x4 scenarios.

These configurations are detailed here:

VM Count and vCPUs per VM	Memory per VM (GB)	Total Memory in VMs (GB)	Total Number of Heavy Users
2 x 2	7	14	2,000
4 x 2	7	28	4,000
6 x 2	7	42	6,000
8 x 2	7	56	8,000
2 x 4	12	24	4,000
3 x 4	12	36	6,000
4 x 4	12	48	8,000

In each of the tested scenarios, the two most interesting metrics collected were the Exchange transaction latencies and system CPU utilization. Transaction latency directly predicts user experience and CPU utilization reflects efficiency. The tests were run for eight hours as per LoadGen guidelines, which included a ramp-up period that was not included in the reported performance metrics.

For the scale up comparison between physical and virtualized systems, apples-to-apples comparable configurations were maintained, i.e., the same processor count and configured memory. For the physical system, Windows Server 2008's 'bcdedit' tool was used to configure the number of processors and amount of memory for the specific test. In the virtual environment, the VM was configured with matching CPU count and amount of memory.

Experimental Results and Performance Analysis

Experiments are presented to show the performance of a single virtual machine and many VMs in a consolidation environment. Single virtual machine performance provides insight into the characteristics of the workload and consolidated performance shows ESX's capabilities in real production environments.

Single Virtual Machine Performance

For an apples-to-apples comparison, identical software stacks were configured in both the physical and virtual environments and conducted the corresponding scale up experiments moving from one to eight physical and virtual processors respectively. The results are discussed next.

Exchange Latencies

Among all of LoadGen's reported transaction latencies, Send Mail is the most prominent and its latency is considered a yard stick of user responsiveness. The 95th Percentile Send Mail latencies being at or below 500 ms is considered acceptable user responsiveness. Send Mail is therefore used to compare performance across the various configurations.

Figure 1. 95th Percentile Send Mail Latency (Physical vs. Virtual)

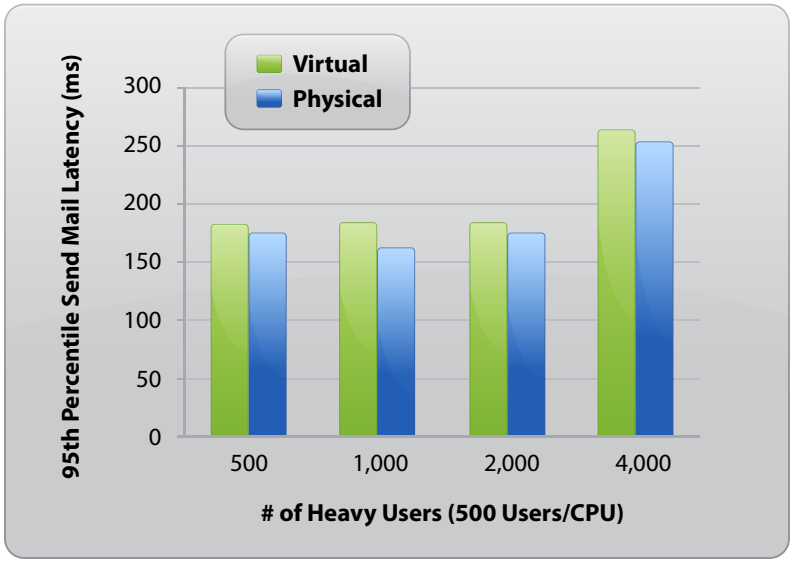

Figure 1 shows Send Mail time staying at a quarter of a second or less as mailbox count and number of CPUs are increased. In all cases, the 95th Percentile Send Mail latencies are far lower than the standard requirement of 500 ms and ranged between 163 ms (for 500 Heavy users) to 264 ms (in case of 4,000 Heavy users). The latencies of the virtual Exchange environments are very similar to the comparable physical environment. Latencies in the virtual environment average within five percent of those measured in the physical environment.

Processor Utilization

For each of the test cases, the overall ESX host CPU utilization has also been measured. An application when virtualized may have to pay some small CPU virtualization overhead cost. This CPU overhead in case of Exchange is negligible and supporting even 4,000 Heavy users consumes only a small fraction of the host CPU resources. This leaves plenty of room to consolidate many more Exchange VMs, to support more users, or run VMs hosting other enterprise applications.

Figure 2. Percentage of CPU utilization in virtual configurations

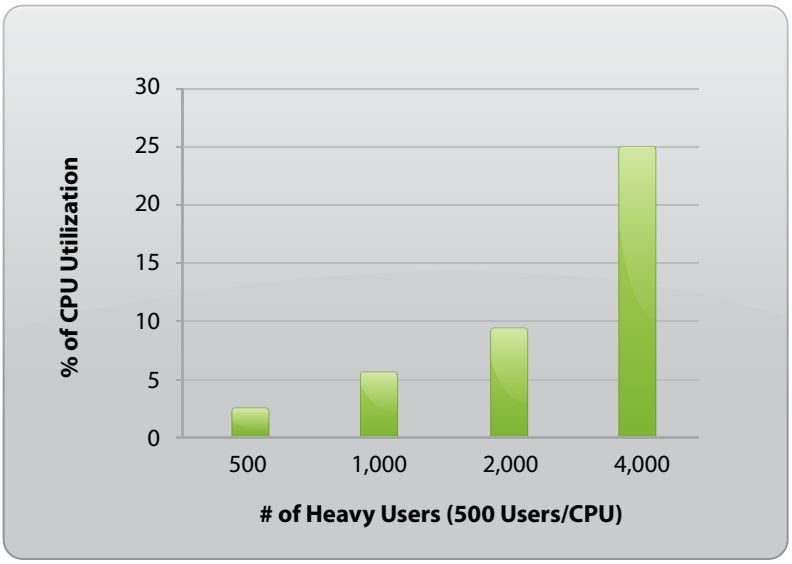

Figure 2 shows the ESX host CPU utilization is low in various configurations when increasing numbers of users. The eight vCPU VM case supporting 4,000 Heavy users consumes barely 25 percent of the host processing capacity.

Consolidated Performance

Now consider how effectively vSphere scales multiple Exchange VMs running on the same host. Consolidated environments increase load and stress ESX's storage and networking stacks along with its resource management (CPU, memory scheduler) modules.

The consolidation tests use multiple two and four vCPU VMs loaded with the identically configured Exchange VMs, following Microsoft's sizing recommendations described above. With both the 2-way and 4-way virtual machines, the Exchange environment was scaled out to demonstrate up to 8,000 Heavy users. In configurations supporting 8,000 Heavy users, the VMware ESX host's physical CPUs have been overcommitted to 16 vCPUs on only eight physical processors.

Exchange Latencies

Figure 3 below depicts the 95th Percentile Send Mail transaction latencies for the multiple 2 and 4 vCPU VMs running concurrently.

Figure 3. 95th Percentile Send Mail Latency (2 vCPU VM vs. 4 vCPU VM)

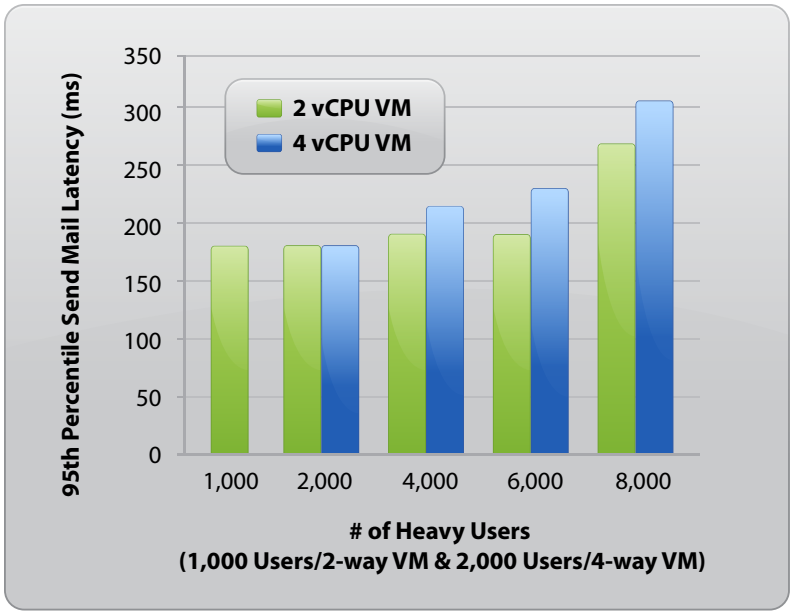

The 95th Percentile Send Mail latencies at 8,000 users were collected with twice as many virtual CPUs as physical cores. At this level of over-commitment, those latencies are:

- 273 ms with eight 2 vCPU VMs
- 304 ms with four 4 vCPU VMs

In both cases, the Send Mail latency is far below the 500 ms threshold recommendation. There is not a significant performance difference between the different configurations, giving Exchange administrators flexibility in sizing. Even in the largest environments, with 8,000 active Heavy user mailboxes and two times more virtual CPUs than physical cores, the heavily consolidated environment shows Send Mail latencies as low as 100 ms greater than the smallest configuration.

Latency Fairness in VMs

The performance of each VM with respect to the average of all instances is a critically important quality of consolidation software. With Exchange, this means comparing the Send Mail latency of each VM to the average of the scaled out virtual environment. The fair distribution of resources to each of the VMs provides uniform performance for all users in the environment. An unfair distribution of resources would lead to unpredictable results and make it difficult or impossible to diagnose performance problems. This section contains details evaluating ESX 4's ability to fairly distribute resources to Exchange VMs.

The RPC averaged latency counter (reported by Exchange by Windows Perfmon) signifies the amount of time taken by that particular Exchange instance to handle Outlook RPC requests (tasks such as Send Mail, Browsing, Calendar, etc). This counter was used as an estimation of each Exchange VM's transaction processing latency for requests received from LoadGen. Figure 4 shows the distribution of RPC latencies for each of the eight 2 vCPU VMs from the 2x overcommitted scenario.

Figure 4. Latency Fairness in VMs

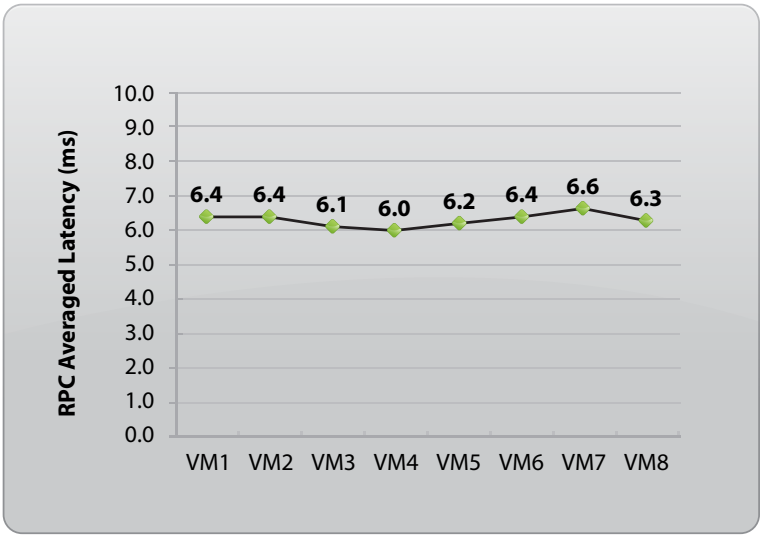

The figure shows no mailbox instance varies more than 0.3 ms from the consolidation-wide RPC latency average, 6.3 ms. With such similar response times, Exchange and virtual infrastructure administrators can be certain of fairly distributed resources, even performance across all users, and predictable and repeatable response time.

Processor Utilization

During the course of the scale out experiments, ESX host CPU consumption was also measured, as depicted in Figure 5.

Figure 5. Percentage of CPU utilization (2 vCPU VM vs. 4 vCPU VM)

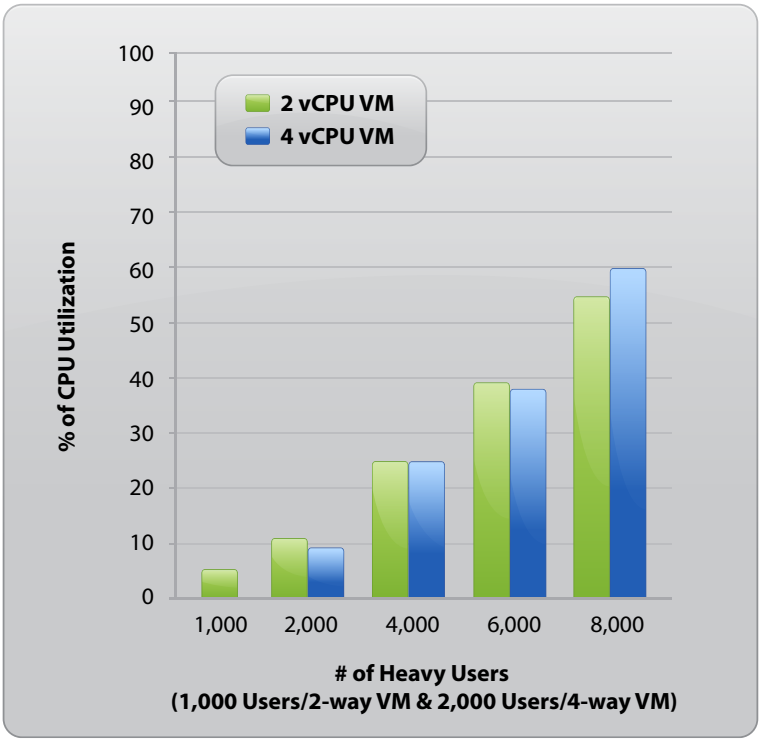

The CPU utilization increased linearly as the Exchange VMs were loaded with more users. The CPU cost of virtualizing the same mailbox count with either the 2 or the 4 vCPU VMs is similar. Even at 8,000 Heavy Exchange users, representing a medium to large business, the CPU utilization on a single host is below 60 percent. This leaves room for additional Exchange user growth or other application consolidation.

Best Practices

This section describes some performance best practice recommendations for improving virtualized Exchange performance:

- If your virtual machines were created in a previous version of ESX server, upgrade those virtual machines to Hardware Version 7 in order to enable the latest virtual machine features like 8 vCPUs. Hardware Version 7 is introduced in vSphere 4. You can use the VI Client to upgrade the Hardware Version of your virtual machines.
- Ensure that VMware tools within the VM are installed and up to date. The tools package includes performance enhancements that take advantage of features in newer versions of ESX.
- Align the VM's VMDK disks along the sector boundary alignment recommended by the storage vendor. The VMFS partitions were aligned automatically by using the VI Client. Detail information is available at http://www.vmware.com/pdf/esx3_partition_align.pdf.
- Make sure the system has enough memory to avoid ESX host swapping, else performance in the virtual machines is significantly reduced. Pay attention to the balloon driver's inflation, which induces swapping in the guest. Detail information is available at http://www.vmware.com/pdf/esx3_memory.pdf.
- Follow the same memory and CPU Best Practices guidelines recommended by Microsoft for physical environment in configuring your Virtual Machine CPU and Memory resources as well. For above and other Microsoft recommended Exchange 2007 best practices for Server and Storage design, please visit http://technet.microsoft.com/en-us/library/bb738142.aspx.

Conclusion

The results in this paper show that the VMware vSphere 4 platform has excellent performance and scalability when running Microsoft Exchange 2007 in virtual machines. It shows that Exchange Send Mail 95th Percentile latencies remain far lower than the standard requirement of 500 ms for both single VM and consolidated environments. In the CPU overcommitted scenarios, transaction latencies were maintained below 500 ms (273 ms supporting 8,000 Exchange Heavy users with eight 2 vCPU VMs). These low latencies reported for 95 percent of the mailboxes are the key to providing an excellent user experience in an Exchange environment.

In addition to the demonstrated low latencies, 8,000 Heavy Exchange users consume less than 60 percent of host CPU resources, leaving room for further user growth and/or further consolidation. ESX also ensures similar and consistent performance across all consolidated virtual machines. In the eight 2 vCPU configuration Exchange transactions across all of the eight VMs were serviced with similar response times.

Through concrete scale up and scale out data, vSphere 4 is shown to be an ideal virtualization platform for your small, medium, or large Exchange production deployments.

About the Author

Vincent Lin is a Performance Engineer at VMware. In this role, his primary focus is to evaluate and help improve the performance of VMware products in better supporting key Enterprise applications. Prior to VMware, Vincent was a Principal Performance Engineer at Oracle. Vincent received a Master of Science degree from the University of West Florida.

Acknowledgements

The author would like to thank the VMware development teams for their efforts in developing and optimizing vSphere. He would also like to thank Charles Windom for the Microsoft Exchange production deployment insights. Finally he would like to thank members of the VMware Performance team who have helped with the performance analysis and creation of this document: Jennifer Anderson, Nikhil Bhatia, Scott Drummonds, Seongbeom Kim, Aravind Pavuluri, and Kiran Tati.

VMware, Inc. 3401 Hillview Ave Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com
 VMW_09Q2_WP_VSPHERE_EXCHANGE_EN_P11_R1

vmware®