

Deploying Microsoft Exchange Server 2007 mailbox roles on VMware Infrastructure 3 using HP ProLiant servers and HP StorageWorks



Executive summary.....	2
Target audience	2
Introduction.....	2
Disclaimer	3
Testing overview.....	3
Lab configuration	4
Overview on virtual machine benchmarking.....	7
Phase one: I/O stress testing using Jetstress	7
Phase two: mailbox simulation using LoadGen	9
500 building block test.....	11
1000 building block test.....	13
2000 building block test.....	15
Summary	17
Appendix A	19
Appendix B – vendor support.....	20
VMware support statements.....	20
For more information.....	21

Executive summary

Historically, some IT groups have been reluctant to virtualize applications such as Microsoft®'s Exchange Server 2003 (Exchange 2003). The performance demands Exchange 2003 placed on the server's hardware left little processor headroom for additional server consolidation. A large Exchange mailbox server could fully utilize a four-processor server. However, with the introduction of dual-core and quad-core processors, server hardware performance can now scale significantly higher. Furthermore, architectural changes in Microsoft Exchange Server 2007 have provided performance improvements that also reduce the performance cost of virtualization.

For this white paper, a series of mailbox building blocks were tested using Exchange Server 2007 Enterprise Edition (Exchange 2007). The building blocks were used to measure performance of a virtualized mailbox server role using VMware Infrastructure 3 (VI3)¹. Performance was measured for a mailbox server configured as guest virtual machines (VMs) with 500, 1000 and 2000 mailboxes.

Test results indicate that the performance of each of the building block configurations was acceptable in terms of CPU and I/O performance. Even though virtualizing the mailbox server role added a significant increase in CPU utilization, there was sufficient headroom available within each of the Exchange Server 2007 building block configurations.

Target audience

The information contained in this white paper is intended for solutions architects, engineers, and project managers involved in the planning and virtualization of Microsoft Exchange Server 2007. You should be familiar with Exchange Server 2007 terminology and best practices as well as deployment and support of VMware Infrastructure 3 solutions.

Introduction

Using the building block approach to Exchange analysis serves two purposes. First, multiple measurements of various Exchange loads and CPU configurations provide trending data that assist in the understanding of the load characteristics. Second, CPU usage of heavily loaded, differently-sized building blocks suggests configurations when scaling out an Exchange deployment to multiple VMs on a single machine. With accurate VM sizing data, smart decisions in the tradeoff between highly-flexible, easily-migrated small VMs and minimal-maintenance large VMs can be made. For more information refer to the "VMware ESX Server: A comprehensive guide on how ESX virtualizes HP ProLiant Servers" white paper, located at <http://h71019.www7.hp.com/ActiveAnswers/cache/394553-0-0-225-121.html>.

To establish a baseline, the building blocks were initially tested using the physical server, and then repeated as a virtualized server. The Exchange Server 2007 building block testing used Microsoft's recommended heavy mailbox profile guidelines for system memory per user for database cache and for the number of mailbox users per CPU core necessary to support each of the building block configurations.

Testing was performed using the HP ProLiant DL580 G4 with a total of eight CPU cores and 32GB of system memory. For more information on the ProLiant DL580 G4 server, go to <http://h10010.www1.hp.com/wwpc/us/en/en/WF05a/15351-3328412-241644-3328422-1137825.html>.

¹ <http://www.vmware.com/products/vi/>

For even more performance headroom, the recently announced HP ProLiant DL580 G5 server offers quad-core processors, which doubles the number of CPU cores to 16 and support for up to 128GB of system memory. For more information on the ProLiant DL580 G5 server, go to <http://h10010.www1.hp.com/wwpc/us/en/en/WF05a/15351-15351-3328412-241644-3328422-3454575.html>.

Testing was conducted in two phases; the first phase was to measure the storage I/O performance. This was done using the Microsoft Exchange Server Jetstress tool (Jetstress). Testing was performed on the physical server and then repeated using the guest VM, to compare I/O storage performance between the two environments. The second phase of testing used the Microsoft Exchange Load Generator tool (LoadGen) to simulate messaging client traffic and was used to measure performance of the fully configured Exchange 2007 mailbox server, both as a physical server and as guest VM.

Performance was measured in two key areas, CPU utilization and storage I/O performance, which are typically the most demanded resources when virtualizing an Exchange server. To provide an accurate measurement of virtual machine workloads, CPU performance metrics were collected using esxtop². esxtop is a VMware performance data collection tool, used to measure server performance metrics.

Test results were then validated using the LoadGen pass/fail criteria to validate server performance metrics. Appendix A lists the performance counters' pass/fail criteria and expected messages sent and received per user per day.

Disclaimer

Microsoft does not support its software running in conjunction with non-Microsoft hardware virtualization software, except for Microsoft customers who have a Premier-level support agreement. For Premier-level customers, Microsoft will use commercially reasonable efforts to investigate potential issues with Microsoft software running in conjunction with non-Microsoft hardware virtualization software. For Microsoft customers who do not have a Premier-level support agreement, Microsoft will require the issue to be reproduced independently from the non-Microsoft hardware virtualization software. Please refer to Microsoft's Support statement for more details <http://www.support.microsoft.com/kb/897615/>. For addition support details refer to [Appendix B – vendor support](#).

Testing overview

The test matrix was broken down into two discrete phases:

1. Storage I/O stress testing using Jetstress
2. Mailbox simulation using LoadGen

Initial testing was performed to validate storage performance, for both the physical server and guest VM. Jetstress³ is specifically designed to exercise the storage subsystem. Jetstress simulates the I/O patterns used by the Exchange Jet database engine. As a best practice, Jetstress should be used to perform I/O performance testing to validate the storage configuration before deploying any Exchange 2007 mailbox server role.

After the initial storage configuration was validated using Jetstress, the Exchange 2007 topology was built and tested using LoadGen⁴. Unlike Jetstress, LoadGen requires the complete Exchange topology, including the Active Directory domain, the client access server (CAS), the hub transport server (HUB)

² http://www.vmware.com/pdf/esx2_using_esxtop.pdf

³ <http://www.microsoft.com/downloads/details.aspx?FamilyId=73DFE056-0900-4DBB-B14A-0932338CECAC&displaylang=en>

⁴ <http://www.microsoft.com/downloads/details.aspx?FamilyId=0FDB6F14-1E42-4165-BB17-96C83916C3EC&displaylang=en>

and mailbox server roles to be installed. LoadGen is used to simulate some of the key functions of the Microsoft Office Outlook 2007 client to exercise the mailbox server.

Testing of the mailbox building block configurations was initially performed on the physical server (non-virtual) and then repeated as a virtual machine. To assist direct comparison, both the physical and virtual machine used the identical hardware configuration including the number of processors, memory size and storage subsystem.

Lab configuration

ProLiant DL580 G4 (server under test):

- Four Dual-Core 64-bit Intel® Xeon® 7130M Processor (3.20 GHz, 150 Watts, 800MHz FSB)
- 32GB PC2-3200 Registered Memory (DDR2-400) with Advanced ECC functionality
- HP StorageWorks FC2242SR Dual Channel 4 Gb PCI-e HBA
- Two 72GB 10k Small form factor serial attached SCSI (SAS) drives for host operating system (OS)
- Two 72GB 10k SAS drives for guest OS
- RAID1 disk arrays for host OS disk and guest OS disk
- HP Smart Array P400 SAS controller with battery backed write cache (BBWC)
- BBWC set to 50/50 read/write
- Integrated NC371i Multifunction Gigabit Network Adapter, 1 Gb network
- Intel Virtualization Technology (Intel VT) enabled in BIOS
- Hyper-Threading enabled

HP StorageWorks 8000 Enterprise Virtual Array (EVA8000) storage array:

- 168 36GB 15K fibre channel disks
- Two disk groups were created: one for the mailbox databases and a second for the transaction logs

Local network:

- 1 Gb single switched network (no V-LAN)

Software configuration (physical server and guest OS):

- Microsoft Windows® Server 2003 Enterprise x64 Edition SP1
- Multiprocessor hardware abstraction layer (HAL) (used for all test configurations)
- Hotfixes KB898060, KB904639 and KB918980
- Exchange Server 2007 Enterprise Edition build 685.24
- Emulex HBA driver 7.1.207 (physical server only)
- Network driver 2.8.13 (physical server only)
- VMware tools (for VM's guest OS only)
- Video hardware acceleration (for guest OS only)
- LSI Logic SCSI driver (for guest OS only)

Tuning settings:

- Boot.ini (physical server only)
 - /NumProc will be used to define the number of processors
 - /MaxMem will be used to define the amount of memory to be used

- VM settings (guest VM OS only)
 - Number of processors (SMP) defined per VM
 - Memory defined per VM
- HBA settings (physical server only)
 - Queue depth set to 32 (default setting)
 - Queue target set to per LUN (default setting)
- Network:
 - 1 Gb full duplex
- Multipath I/O (MPIO):
 - Shortest queue service time (default setting)
- Windows disk administration:
 - Basic disk partitions
 - Default allocation size (OS disks and Log)
 - Sector aligned database logical unit numbers (LUN) using 64k allocation units and 64 as offset

VMware ESX Server 3.0.1 (host installation):

- Partitions options:
 - Installed host on SAS 72GB disk (RAID1)
 - Used recommended disk selection
 - Used recommended partition settings
 - Used recommended boot specification, install on master boot record (MBR)
 - Exchange database and logs disks were mapped as RAW devices
- Network:
 - NetXtreme II 5706
 - Default network for VM
 - Single v-network


Support servers:

- Active Directory server (physical)
- CAS and HUB transport server (physical)

LoadGen configuration:

- Office Outlook 2007 online mode using the heavy profile
- Build 8.01.0154 (web release)

Figure 1. Lab layout



Overview on virtual machine benchmarking

The performance results that follow demonstrate several artifacts of virtual system performance that may be unfamiliar to those used to physical system benchmarking. All of these have to do with the manner in which CPU cycles are used and analyzed and are worth description.

ESX Server executes code running on a virtual machine's virtual CPU (vCPU) in a linear manner. That is, code running on a single vCPU is not ever concurrently run on multiple physical CPUs. However, it is the prerogative of the hypervisor's scheduler to move that code between different cores or sockets for its own goals. The result of in-host CPU migrations is a distribution of vCPU work across multiple physical CPUs (pCPUs), which is done to increase total system efficiency.

Virtual device drivers running on the host can and are often implemented to use multiple threads. In addition, other host activities (such as the web interface, console OS time keeping, etc.) utilize non-zero CPU resources that are scheduled across the available pCPUs. In short, scheduler distribution of single-threaded vCPU work, multi-threaded driver execution, and single- or multi-threaded non-VM processes can produce a distribution of work across many physical processors.

In comparing virtual platform performance to the non-virtual (native) counterpart, we have chosen to include all server functions in the "cost of virtualization." While this is certain to include non-essential process processor usage, it is also guaranteed to be a pessimistic bound on the maximum usage on the virtual system. CPU usage on the host will then be reported as a percentage of all available pCPUs. In the native measurements, the operating system was restricted to match the number of vCPUs in the VM, which means that normalization of pCPU utilization is required. One manner for doing this is to report "Megacycles," which calculate work per available resources. In this manner, utilization across systems with a differing number of virtual and physical CPUs can be determined.

A final word on CPU utilization: in this experiment, the Hyper-Threading feature was available on the processors. Hyper-Threading provides the appearance of two cores for each single core on the system. These two presented cores (called logical cores) provide an incremental benefit over the single real core (called a physical core). While the value of Hyper-Threading can be debated, it is beyond debate that two logical cores provide more processing power than one physical core but less than two physical cores. Because of this non-equal processing power, CPU usage across multiple logical cores on a Hyper-Threading system is non-linear. Use of 100% of one physical core indeed represents completion of work twice as fast as 50% of one physical core. But use of 100% of two logical cores may only translate into use of as little as 110% of a single physical core.

Ultimately, the presence of non-critical hypervisor activities, Hyper-Threading in the CPU, and unequal CPU counts in the physical and virtual setups makes direct comparison difficult. This paper is not meant to serve as an analysis of the overhead due to virtualization but rather as a guide to the amount of utilization of this ProLiant DL580 G4 server by various building blocks.

Phase one: I/O stress testing using Jetstress

Phase one testing used Jetstress to validate the storage subsystem performance. The tests were initially run on the physical server and then repeated on the VM. The results were then compared to measure differences in I/O performance between the physical server and the virtual machine.

Setup of the physical server Jetstress test:

- Boot.ini was modified to /NumProc=4 and /MaxMem=12288
- Jetstress was configured with four storage groups, each with one database
- The test ran with 64 threads per database (defined by auto tuning)

Setup of VM Jetstress test:

- Memory: 12,288MB
- CPU: 4
- SCSI Controller: LSI Logic
- Hard disk 1: VMware Virtual Machine Disk Format (VMDK) (OS)
- Hard disk 2: Raw device mapping (RDM) (database (DB))
- Hard disk 3: RDM (Log)
- Hard disk 4: RDM (backup)
- Jetstress was configured with four storage groups, each with one database
- The test ran with 60 threads per database (defined by auto tuning)

EVA8000 configuration:

- Database disk group 40 disks using VRAID1
- Logs disk group 8 disks using VRAID1

Jetstress was configured with four storage groups (SGs) using the following system parameters:

Log buffers	9000
Minimum database cache	128.0MB
Maximum database cache	1024.0MB
Insert operations	25%
Delete operations	10%
Replace operations	50%
Read operations	15%
Lazy commits	80%

The test was configured to run the two-hour performance test (random I/O) and the thread workload was defined by auto tune in both physical and VM testing.

Measured areas:

- Achieved IOPS (I/O operations per second) on the random I/O test
- Disk read and write latencies on the random I/O test

Data will be evaluated using the following criteria:

- Testing will measure IOPS for random workload
- Pass/Fail criteria will be based on the Jetstress test constraints
- Maintain read and write latency 20 ms or less

The Jetstress results in Table 1 indicate that the VM was capable of providing storage performance on parity with the similarly configured physical server. In fact, IOPS achieved by the VM were about 10% higher than the physical test. However, the disk write latency was at the 20 ms maximum.

Table 1. Jetstress I/O test comparison

Test	Number of Disks	IOPS	Observed CPUs	Memory	Avg. Disk sec/read	Avg. Disk sec/write	Number of threads	Number of SGs/DBs	IOPS per disk	CPU utilization
Physical	40	4499	4	12GB	0.012	0.008	65	8	112	21%
VM	40	5034	8	12GB	0.010	0.020	60	8	125	7.79%

As stated above, direct comparisons of CPU usage are not possible given the data presented in Table 1. However, given that on half as many Hyper-Threaded processors, the physical setup was using a little more than double the virtual setup's processors one can conclude that the CPU overhead due to virtualization is fairly small in this experiment.

Table 2. esxtop Physical CPU\Processor Time

Counter	Average
Physical CPU\0\%Processor Time	15
Physical CPU\1\%Processor Time	0.064
Physical CPU\2\%Processor Time	0.088
Physical CPU\3\%Processor Time	0.215
Physical CPU\4\%Processor Time	13
Physical CPU\5\%Processor Time	11
Physical CPU\6\%Processor Time	11
Physical CPU\7\%Processor Time	12
Avg. % Processor time Total	7.79

Phase two: mailbox simulation using LoadGen

Phase two testing consisted of a series of load generation tests to measure performance of the mailbox server role in a fully configured Exchange 2007 environment. In a manner similar to the process used in phase one, comparison testing was performed initially on the physical server and then in the VM. Testing used the Office Outlook 2007 online mode with the heavy profile to simulate a mailbox user profile. Phase two testing required a complete Exchange 2007 environment, including an Active Directory server, CAS, and HUB in addition to the mailbox server. To provide a control for the test process, only the mailbox server role was virtualized; each of the other server roles were deployed as physical servers.

OS setup (physical and guest):

- Windows Server 2003 Enterprise x64 Edition SP1
- Multiprocessor HAL
- Hotfixes KB898060, KB904639 and KB918980

Exchange Server 2007 setup:

- Build 685.24 Enterprise Edition
- Dedicated HUB and CAS server
- Dedicated Active Directory domain controller
- Mailbox physical server configured per the build block
- Mailbox server in VM configured per the build block
- One storage group/database per 500 users
- Content indexing off

LoadGen setup:

- Build 8.0.1.0154
- Test duration 12 hours test run, four-hour ramp up and eight-hour steady state
- Office Outlook 2007 online mode using the heavy action profile

EVA8000 storage configuration:

- EVA storage was configured with two disk groups
- Database disk group was defined for each building block (as noted in each test case)
- Log disk group remained fixed with eight disks (minimum disk group size) for all build blocks
- VDIs for databases and Logs were configured using VRAID1 (disk mirroring)

Database LUNs:

- Database LUN was sector aligned with Microsoft's Diskpart tool using a 64 KB offset
- Database LUN was formatted with 64 KB allocation units

The three building blocks tested were sized at 500, 1000 and 2000 user mailboxes. Testing followed Microsoft's CPU core and memory guidelines for the heavy mailbox user, recommending one CPU core per 500 users and 5MB of system memory per mailbox plus an additional 2GB of memory for the operating system and the Exchange 2007 application.

Measured areas:

- CPU and memory metrics in the physical server test using Perfmon
- CPU and memory metrics with in VM using esxtop
- Exchange performance counters including the MExchangeS Mailbox counter 'Messages Queued for Submission', and the MExchangeS counters 'RPC Average Latency' and 'RPC Requests'.
- Disk performance counters, including the 'Disk Transfer/Sec', the calculated IOPS per users, the 'Avg. Disk sec/Read' and 'Avg. Disk sec/write'

Data will be evaluated using the following criteria:

- CPU performance comparison on the given build blocks
- Disk performance to compare storage I/O
- Pass/Fail criteria will be based on the LoadGen metrics (see Appendix A)

500 building block test

For the 500 user building block comparison testing, the guidelines recommend a single CPU core and 2GB + 500 x 5MB = 4.5GB of system memory.

Physical server configuration:

- Boot.ini file was modified using Numproc=1 and MaxMem=4608

Virtual machine configuration:

- Memory: 4608MB
- CPU: 1
- SCSI controller: LSI Logic
- Hard disk 1: VMware ESX Server File System (VMFS) (OS)
- Hard disk 2: RDM (DB)
- Hard disk 3: RDM (Log)
- Hard disk 4: RDM (backup)

Exchange:

- One storage group with one database
- Initialized with 500 mailboxes using the heavy profile

EVA8000 storage configuration:

- Database disk group was configured with 12 disks using VRAID1
- Log disk group was configured with 8 disks using VRAID1

The results of the 500 user physical and virtual tests listed in Table 3 indicate acceptable disk performance and latencies and also success in terms of pass/fail criteria and expected messages sent and received per user per day (see Appendix A). Performance metrics for storage show low disk latencies below 10 ms for sec/read and less than 5 ms for sec/write. For Exchange, the three key counters remain well below the required thresholds.

Table 3. 500 user LoadGen test

Counter	Physical	Virtual
Disk Transfers/sec	238	215
IOPS/user	0.48	0.42
Avg. Disk sec/Read	0.006	0.006
Avg. Disk sec/Write	0.002	0.001
Disk Read Bytes/sec	2,228,655	1,742,788
Disk Write Bytes/sec	409,681	467,798

Counter	Physical	Virtual
MSExchangeIS Mailbox\messages queued for submission	0	0
MSExchangeIS counter\RPC Average Latency	10 ms	9 ms
MSExchangeIS counter\RPC Requests	1	0

As stated above, direct comparison of the physical server CPU utilization with the VM CPU utilization is not an accurate method for comparison. For the 500 user test, for the physical server, the CPU averaged 37% CPU utilization using a single processor, while the VI3 server CPU utilization is only 4.04%. However, for the virtual test, the CPU utilization is averaged across all eight processors in the server, not just the single CPU assigned to the VM. In Table 4 the individual processor times for each CPU on the VI3 server are listed along with the processor time average across all eight CPUs.

Table 4. esxtop CPU processor time for guest VM with 500 users

Counter	Average
Physical CPU\0\%Processor Time	2
Physical CPU\1\%Processor Time	0.104
Physical CPU\2\%Processor Time	0.088
Physical CPU\3\%Processor Time	0.05
Physical CPU\4\%Processor Time	0.05
Physical CPU\5\%Processor Time	0.059
Physical CPU\6\%Processor Time	28
Physical CPU\7\%Processor Time	2
Avg. % Processor time Total	4.04

Since the processor utilization between the two test environments is not derived by the same means, it would be inaccurate to compare the 37% CPU utilization on physical server against the 4.04% measured (average across the eight CPUs) on the VI3 server. A better method to compare the CPU consumption between different test environments is to calculate the megacycles per user. Typically, the megacycles per user calculation is used to compare various client workloads and their respective CPU consumption. However, it can also be used to compare identical workloads against different processors configurations, such as dual-core and quad-core, or in this case between physical and virtual processors.

The formula of megacycles per user is:

$$((\text{CPU utilization} * \text{CPU speed (MHz)} * \text{number of CPUs}) / \text{number of users under test})$$

So, for the physical server test, the calculation is:

$$((0.37 * 3200 * 1) / 500) = 2.36 \text{ megacycles per user}$$

For the virtual server, the calculation is:

$$((0.04 * 3200 * 8) / 500) = 2.04 \text{ megacycles per user}$$

The test resulted in 15% delta in megacycles between the physical server and virtual machine indicating relatively low virtualization overhead with 500 mailbox user building block configuration.

1000 building block test

For the 1000 user building block comparison, the testing guidelines recommend two CPU cores and 2GB + 1000 x 5MB = 7GB of system memory.

Physical server configuration:

- Boot.ini was modified to /NumProc=2 and /MaxMem=7168

Virtual machine configuration:

- Memory: 7168MB
- CPU: 2
- SCSI controller: LSI Logic
- Hard disk 1: VMFS (OS)
- Hard disk 2: RDM (DB)
- Hard disk 3: RDM (Log)
- Hard disk 4: RDM (backup)

Exchange:

- Two storage groups, each with one database
- Initialized with 1000 mailboxes using the heavy profile
- 500 mailboxes per database

EVA8000 storage configuration:

- Database disk group was configured with 22 disks using VRAID1
- Log disk group was configured with 8 disks using VRAID1

The results of the 1000 user physical and virtual tests listed in Table 5 indicate acceptable disk performance and latencies and also success in terms of pass/fail criteria and message sent and received per user per day (see Appendix A). Performance metrics for storage subsystems show low disk latencies below 10 ms sec/read and below 5 ms sec/write. For Exchange, the three key counters remain well below the require thresholds.

Table 5. 1000 user LoadGen test

Counter	Physical	Virtual
Database Cache Size (MB)	5,127	5,657
Disk Transfers/sec	431	358
IOPS/user	0.43	0.36
Avg. Disk sec/Read	0.006	0.005
Avg. Disk sec/Write	0.001	0.001
Disk Read Bytes/sec	4,272, 573	4,245,423
Disk Write Bytes/sec	824,586	785,161
MSExchangeIS Mailbox\messages queued for submission	1	0
MSExchangeIS counter\RPC Average Latency	14 ms	14 ms
MSExchangeIS counter\RPC Requests	1	1

For the 1000 user test, the physical server consumed 27% of the CPU resources across the two processors, while the V13 server consumed 10.67%. Table 6 lists the individual CPU processor times for each of the eight CPUs on the V13 server.

Table 6. esxtop CPU processor time for guest VM with 1000 users

Counter	Average
Physical CPU\0\%Processor Time	4
Physical CPU\1\%Processor Time	0.197
Physical CPU\2\%Processor Time	0.119
Physical CPU\3\%Processor Time	0.086
Physical CPU\4\%Processor Time	26
Physical CPU\5\%Processor Time	4
Physical CPU\6\%Processor Time	24
Physical CPU\7\%Processor Time	27

Counter	Average
Avg. % Processor time across the active CPUs	10.67

Using the megacycle comparison, the physical server resulted in 1.73 megacycles per user, compared with the 2.56 megacycles per user for the VM test. While this would appear to display a 48% increase in overhead between the physical and virtual server, as already mentioned it is difficult to do direct comparisons of processor utilization, especially with several processors running below 1% utilization (in the virtual test). Nonetheless, CPU resources on ProLiant DL580 V13 server are still considered very low at only 10.67%.

Another interesting observation was that the virtual test resulted in a nearly 20% reduction in disk IOPS when compared to the physical test. While both tests were configured with identical memory configurations, the virtual test database cache grew over 500MB larger (see Table 5).

2000 building block test

For the 2000 user building block comparison, the testing guidelines recommend four CPU cores and $2GB + 2000 \times 5MB = 12GB$ of system memory.

Physical server configuration:

- Boot.ini was modified to /NumProc=4 and /MaxMem=12288

Virtual machine configuration:

- Memory: 12288MB
- CPU: 4
- SCSI controller: LSI Logic
- Hard disk 1: VMFS (OS)
- Hard disk 2: RDM (DB)
- Hard disk 3: RDM (Log)
- Hard disk 4: RDM (backup)

Exchange:

- Four storage groups, each with one database
- Initialized with 2000 mailboxes using the heavy profile
- 500 mailboxes per database

EVA8000 storage configuration:

- Database disk group was configured with 40 disks using VRAID1
- Log disk group was configured with 8 disks using VRAID1

The results of the 2000 user physical and virtual tests listed in Table 7 indicate acceptable disk performance and latencies and also success in terms of pass/fail criteria (see Appendix A). Performance metrics for storage show low disk latencies below 10 ms sec/read and below 5 ms sec/write. For Exchange, the three key counters remain well below the require thresholds.

Table 7. 2000 user LoadGen test

Counter	Physical	Virtual
Database Cache Size (MB)	9,583	10,257
Disk Transfers/sec	923	698
IOPS / user	0.46	0.34
Avg. Disk sec/Read	0.005	0.005
Avg. Disk sec/Write	0.002	0.001
Disk Read Bytes/sec	9,715,212	8,290,040
Disk Write Bytes/sec	1,533,689	1,584,557
MSExchangeIS Mailbox\messages queued for submission	1	1
MSExchangeIS counter\RPC Average Latency	12 ms	15 ms
MSExchangeIS counter\RPC Requests	1	1

For the 2000 user test the physical server consumed 28% of the CPU resources across the four processors, while the VI3 server consumed 23%. Table 8 lists the individual processor time along with the average of 23% across all eight CPUs in the VI3 server.

Table 8. esxtop CPU processor time for guest VM with 2000 users

Counter	Average
Physical CPU\0\%Processor Time	6
Physical CPU\1\%Processor Time	0.526
Physical CPU\2\%Processor Time	0.081
Physical CPU\3\%Processor Time	0.12
Physical CPU\4\%Processor Time	46
Physical CPU\5\%Processor Time	46
Physical CPU\6\%Processor Time	42
Physical CPU\7\%Processor Time	45

Counter	Average
Avg. % Processor time across the active CPUs	23.21

Using the megacycle comparison, the physical server resulted in 1.79 megacycles per user, compared with the 2.94 megacycles per user for the VM test. In this comparison there appeared to be significant overhead of virtualization. However, as already discussed, it is difficult to do direct comparisons given that several of the processors were running at less than 1% utilization. Nevertheless, even with the added overhead, CPU resources on the VI3 are still considered low at only 23.21%.

As was observed in the previous test, the virtual test resulted in a 35% reduction in disk IOPS when compared to the physical test. While both tests were configured with identical memory configurations, the virtual test database cache grew over 674MB larger (see Table 7).

Summary

Organizations looking to consolidate their server infrastructure looked last to Exchange due in part to Exchange 2003 demands on the server's processors and I/O storage subsystems. This essentially left little room for potential server consolidation on older two and four processor servers. However, with the introduction of dual-core and quad-core processors, server hardware can now scale significantly higher. Furthermore, architectural changes in Microsoft Exchange Server 2007 have provided performance improvements that also reduce the performance cost of virtualization

The building blocks tested in this document demonstrate the performance impact or added overhead of moving the Exchange 2007 mailbox role to a virtualized environment. As previously discussed, it is difficult to do a direct processor utilization comparison between the physical server and virtual server tests. However, by measuring the CPU megacycles consumed per user, a more accurate picture of the impact of virtualizing the server can be gleaned. Using the megacycles per user calculation the overhead of virtualization can be compared between the physical server and virtual server. Table 9 summarizes the 500, 1000 and 2000 building block test results.

While the larger build blocks appeared to display significant virtualization overhead when compared to the physical server test, it is difficult for direct comparison, especially with the effects of Hyper-Threading and factoring processors with little processor utilization. However, measuring the overall CPU demands on the server, it is considered quite low. The 2000 virtualized mailbox build block averaged only 23% CPU utilization on the eight core server, providing additional head room for further consolidation.

Table 9. Building Block test summary

Type	Number of users	Processors	Processor utilization	Megacycles per user
Physical	500	1	37%	2.36
VM	500	1 VM / 8	4.04%	2.04
Physical	1000	2	27%	1.73
VM	1000	2 VM/ 8	10.67%	2.56
Physical	2000	4	28%	1.79
VM	2000	4 VM/8	23.21%	2.94

Appendix A

Table 10 lists key counters and their respective pass/fail thresholds, used to determine if the LoadGen test was considered successful.

Table 10. LoadGen pass/fail criteria

Object and Counter	Expected Value
Processor: % Processor Time (_Total)	Average less than 90
MSExchangeIS: RPC Requests	Average less than 50 and max of 100
MSExchangeIS: RPC Averaged Latency	Average less than 50 ms and max of 100 ms
Logical disk: Average Sec/Read	Less than 50 ms at all times.
Logical disk Average Sec/Write	Less than 50 ms at all times.
MSExchangeIS Mailbox: Message Queued for Submission	Average less than 250 and max of 1000

Table 11 lists the expected message workload using LoadGen's Office Outlook 2007 online mode heavy profile

Table 11. Message sent and received per user per day

Counter	Expected Value	Error threshold
Messages Submitted/User/Day:	18	+/- 5%
Messages Delivered/User/Day:	31	+/- 5%
Recipients/message ratio	1.70	+/- 5%

Appendix B – vendor support

VMware requires that Support and Subscription services (SnS) must be included with every VMware virtualization software sale.

For more information on HP's customer technical support options for VMware Virtualization software go to:

<http://h18004.www1.hp.com/products/servers/vmware/services.html>

VMware support statements

For more information on VMware's customer technical support options for Microsoft products running within VMware virtual machines go to:

http://www.vmware.com/files/jp/pdf/ms_support_statement.pdf

For more information

For more information on planning, deploying, or managing Microsoft Exchange Server on ProLiant servers and HP storage see:

www.hp.com/solutions/exchange

For more information on planning, deploying, or managing a virtual infrastructure on ProLiant servers see:

www.hp.com/go/vmware

For more information on HP ProLiant servers see:

www.hp.com/go/proliant

For more information on HP storage solutions see:

www.hp.com/go/storage

To help us improve our documents, please provide feedback at www.hp.com/solutions/feedback.

© 2007 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries. Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

4AA1-5890ENW, Revision 2, October 2007

