

Session WV21:

IBM Lotus Domino and IBM Lotus Sametime on VMware Infrastructure 3: From Planning to Successful Deployment

Ivan Dell'Era

Software Engineer, IBM

Purvi Trivedi

Software Engineer, IBM

Agenda

- ◊ **IBM support policy for VMware ESX**
- ◊ **Identify good candidates for virtualization on VI 3**
- ◊ **Lotus Domino benchmark results**
- ◊ **Lotus Sametime benchmark results**
- ◊ **Customer experience**

Introduction to IBM Lotus Domino and IBM Lotus Sametime

○ IBM Lotus® Domino®

- Messaging
- Groupware
- Workflow
- Application server

○ IBM Lotus Sametime®

- Security-rich Instant Messaging server for Enterprise
- Web Conference
- IBM Lotus Sametime Gateway to supported public Instant Messaging Providers
- Telephony Integration

IBM Support Statement for VMware ESX Server

- **IBM offers support for IBM products on VMware ESX and optional support for VMware ESX**
 - > <http://www.ibm.com/servers/eserver/xseries/vmware.html>
- **Lotus Domino products are officially supported on ESX**
 - > <http://www-1.ibm.com/support/docview.wss?rs=203&uid=swg21106182>
- **What to pay attention when using Lotus Domino on VMware in production environments**
 - > <http://www.ibm.com/support/docview.wss?rs=203&uid=swg21267837>
- **Other virtualization technologies and hypervisors are currently not supported for Lotus Domino based products**

Why Testing Is Important in a Virtualized Environment

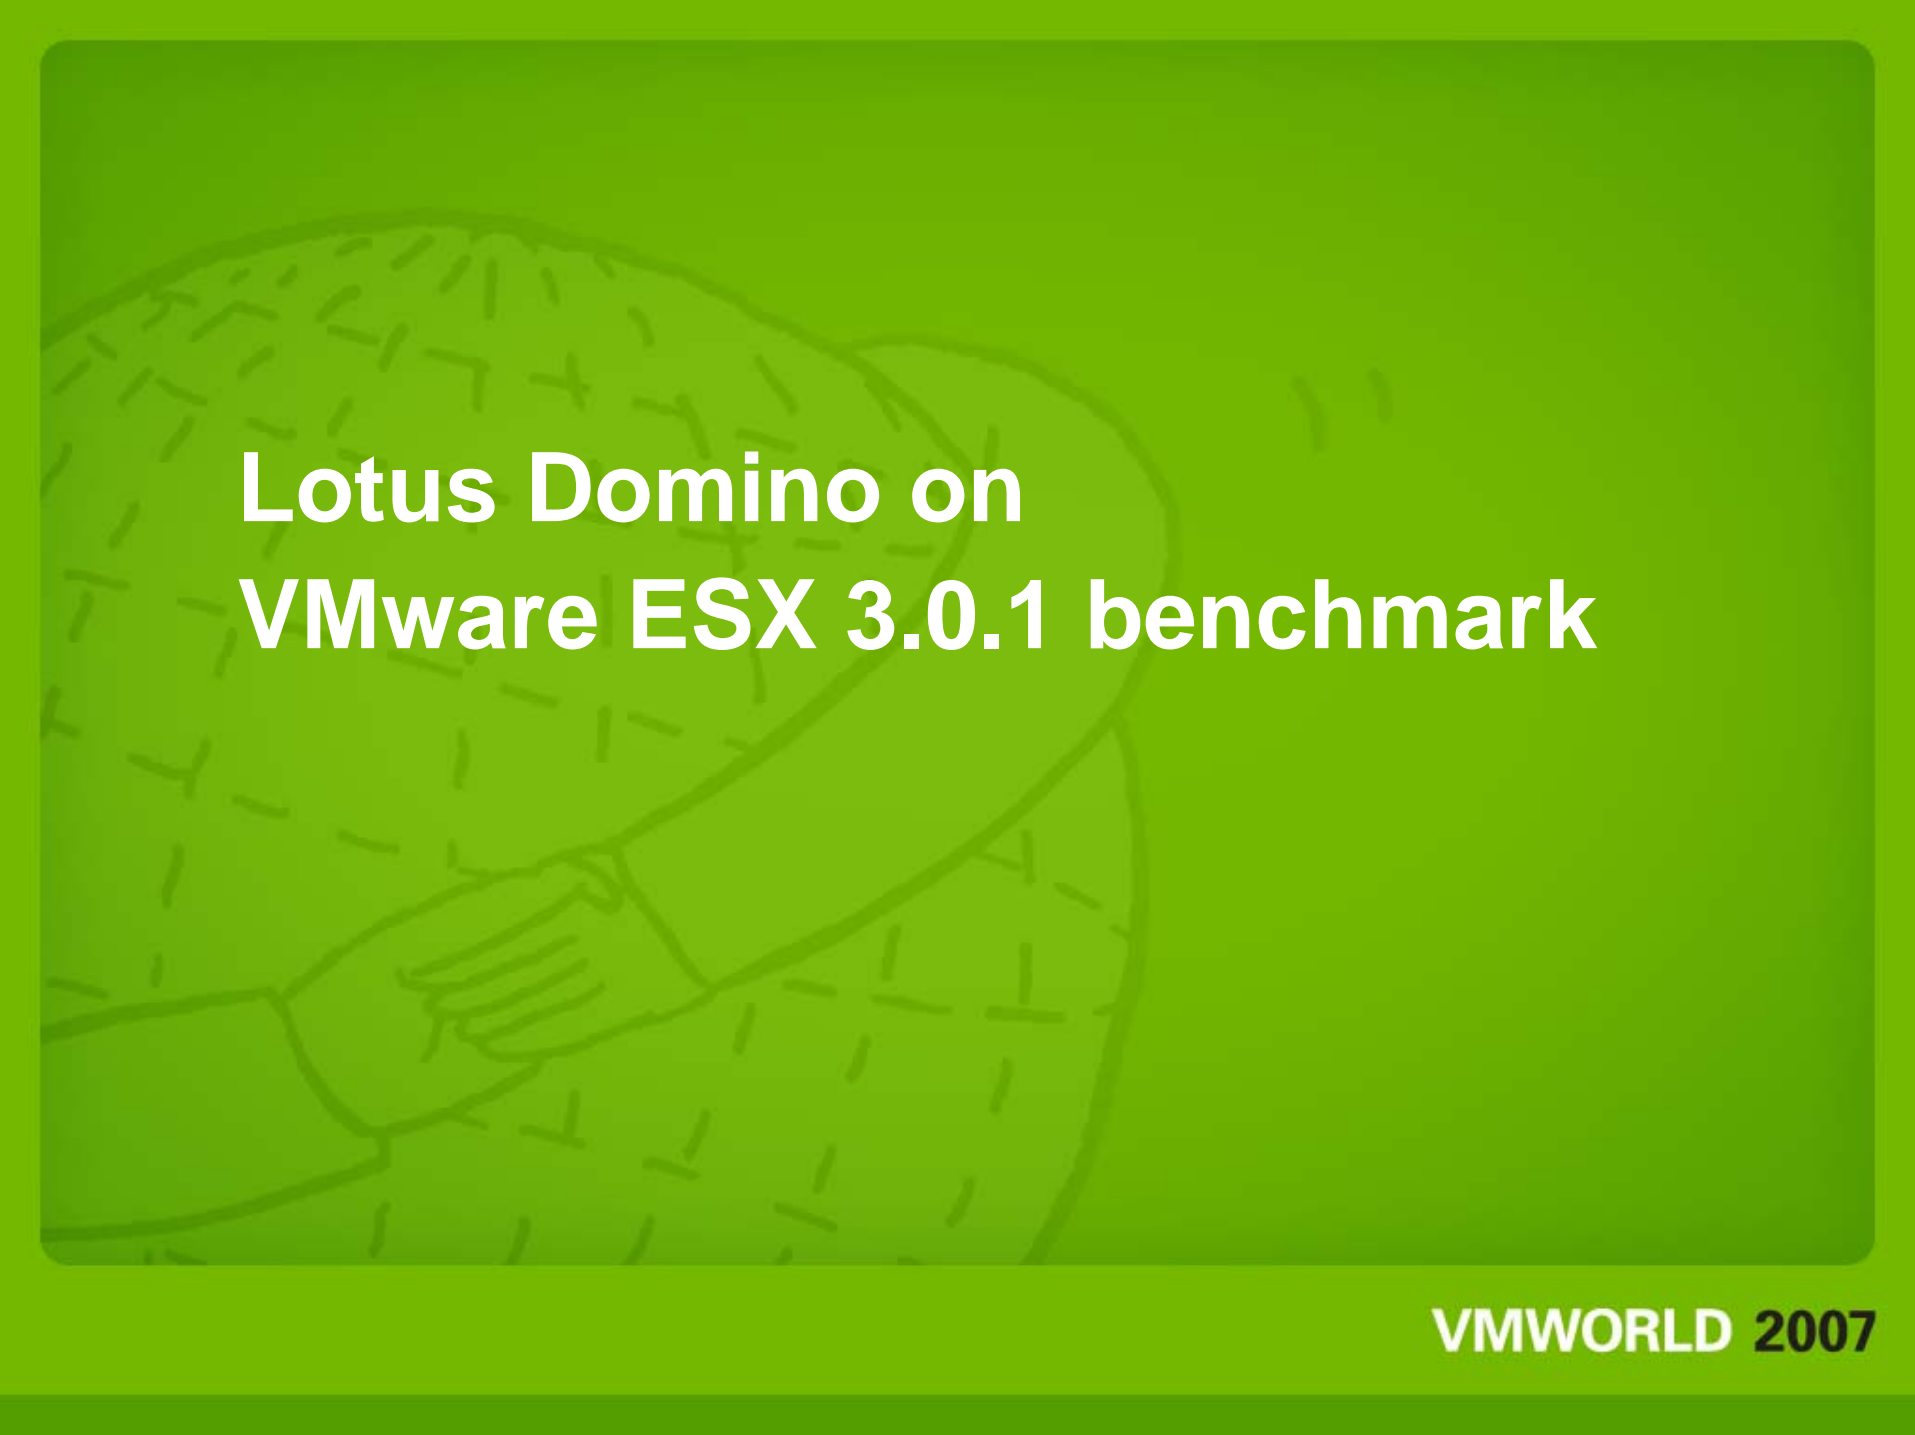
- **Virtualization is relatively new to the x86 platform**
- **Intel® and AMD provide virtualization support in their CPUs**
- **What about disk and network I/O?**
- **No “one size fits all” solution exists**
 - Differences in hardware, storage solutions, disk sizes, just to name a few
 - What other VMs are running on the same box?
 - Not all physical servers are good candidates for virtualization
- **Virtualization always works well in QE, test, development or pre-production**
 - Production may be a different story
- **This is why you need to test your virtualized environment**
 - Don't simply pull the plug and hope it works
- **If done correctly virtualization can work for you too**
 - Several US customers are currently testing Lotus Sametime on VMware ESX Server, some with over 40K concurrent Lotus Sametime users

Before We Start...

- **You need to find the answers to these questions:**
 - Why do you want to virtualize?
 - What is your goal?
 - Is it more important to save a few dollars or to prevent issues difficult to troubleshoot?
 - What do you have today, and where are you going?
 - Did you test your new environment?
 - How do you measure success?
- **If you don't, your users will, and they will let you know**

What Servers Are Good Candidates for Virtualization?

- There isn't a "one size fits all" rule
- **General rule: very busy, CPU saturated, intensive I/O servers may not be good candidates**
 - > Exception: Moving from older to newer and more powerful hardware may be possible as long the response time for end users doesn't degrade
- **Test, development, and branch servers are generally fine**



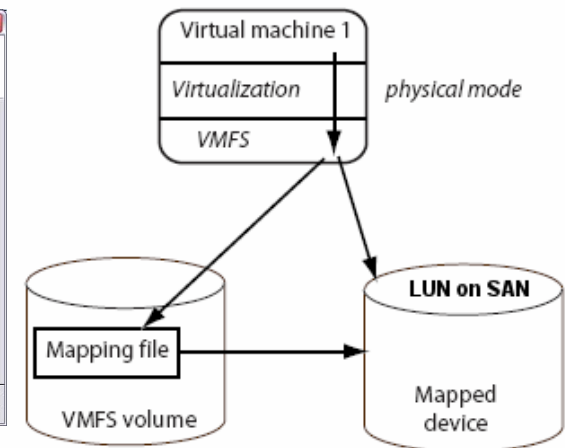
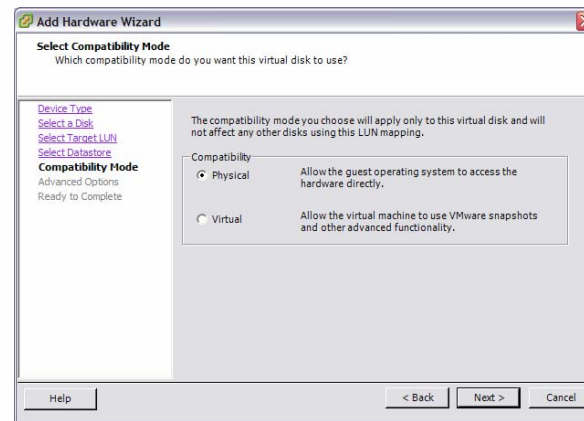
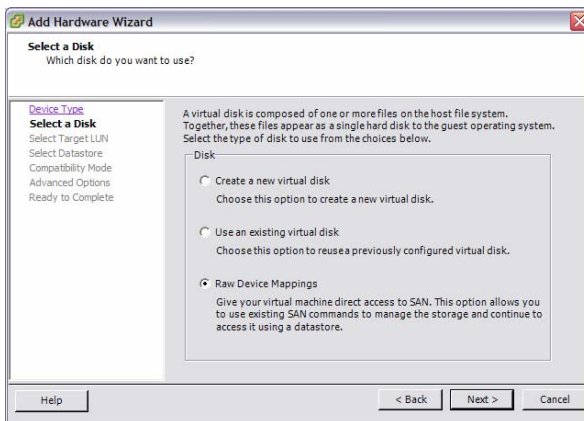
Lotus Domino on VMware ESX 3.0.1 benchmark

IBM Lotus Domino Benchmark goal

- **NRPC mail workload performed out of the box with minimal tuning of Lotus Domino 7.0.2FP2**
 - > 6,000 NRPC mail users simulated using a single driver
 - > Mail file size 87.5 MB, totaling 512 GB at the beginning of the test
 - > NSF_DbCache_Maxentries=6050
- **Verify scalability of Lotus Domino in VMware ESX environment**
 - > This benchmark focused on a single VM running on the server
- **Identify tuning, if required**
- **Verify the difference of performance of VMFS vs RDM**
- **Results considered acceptable with response time < 1 second**
- **Note: this is a simulated test under ideal conditions not from a real environment; your milage may (and will) vary**
 - > Typically you need to at least cut in half the benchmark results

What Is RDM?

- Physical Raw Device Mapping specifies minimal SCSI virtualization; in physical mode, the VMkernel passes all SCSI commands to the device
- See http://www.vmware.com/pdf/esx25_rawdevicemapping.pdf for more information
- In theory RDM provides better performance, higher throughput, and lower CPU utilization



Hardware

- **IBM System X™ 366 (about 3-year-old system)**
 - > 16 GB RAM
 - > 4 Intel Xeon® CPUs 3.66 GHz
 - > QLogic QLA2340
- **Intel Clovertown whitebox** (this is a server not yet supported by VMware – it is expected to be added to the Hardware Compatibility List in the near future)
 - > 16 GB RAM
 - > 2 QuadCore CPUs 2.33 GHz
 - > QLogic QLA2360
- **SAN DS4800**
 - > 8 GB cache
 - > 15K RPM
 - > 36 GB disks
 - > 35+ spindles used for Lotus Domino Data

IBM System X 366

OS: Microsoft® Windows® 2003 Enterprise Edition SP2 32 bit

o Physical test

- 4 CPUs
- HT off
- 6 GB RAM
- Boot disk & OS C: (local)
- Program Directory E: (SAN RAID-5)
- Data Directory F: (SAN RAID-5)
- Transaction logs G: (SAN RAID-1)
- Page file on E: configured for 4096-20000 MB
- /numproc=4 /PAE
/burnmemory=10240 in boot.ini

o VMware test

- 4 vCPU
- HT on – VMware recommendation
- 6 GB RAM
- Boot disk & OS C: (SAN RAID-5)
- Program Directory C: (SAN RAID-5)
- Data Directory F: (SAN RAID-5)
- Transaction logs G: (SAN RAID-1)
- Page file on C: configured for 4096-20000 MB
- Single VM running on VMware ESX 3.0.1

IBM System X 366 Test Results

1,000-4,000 Users

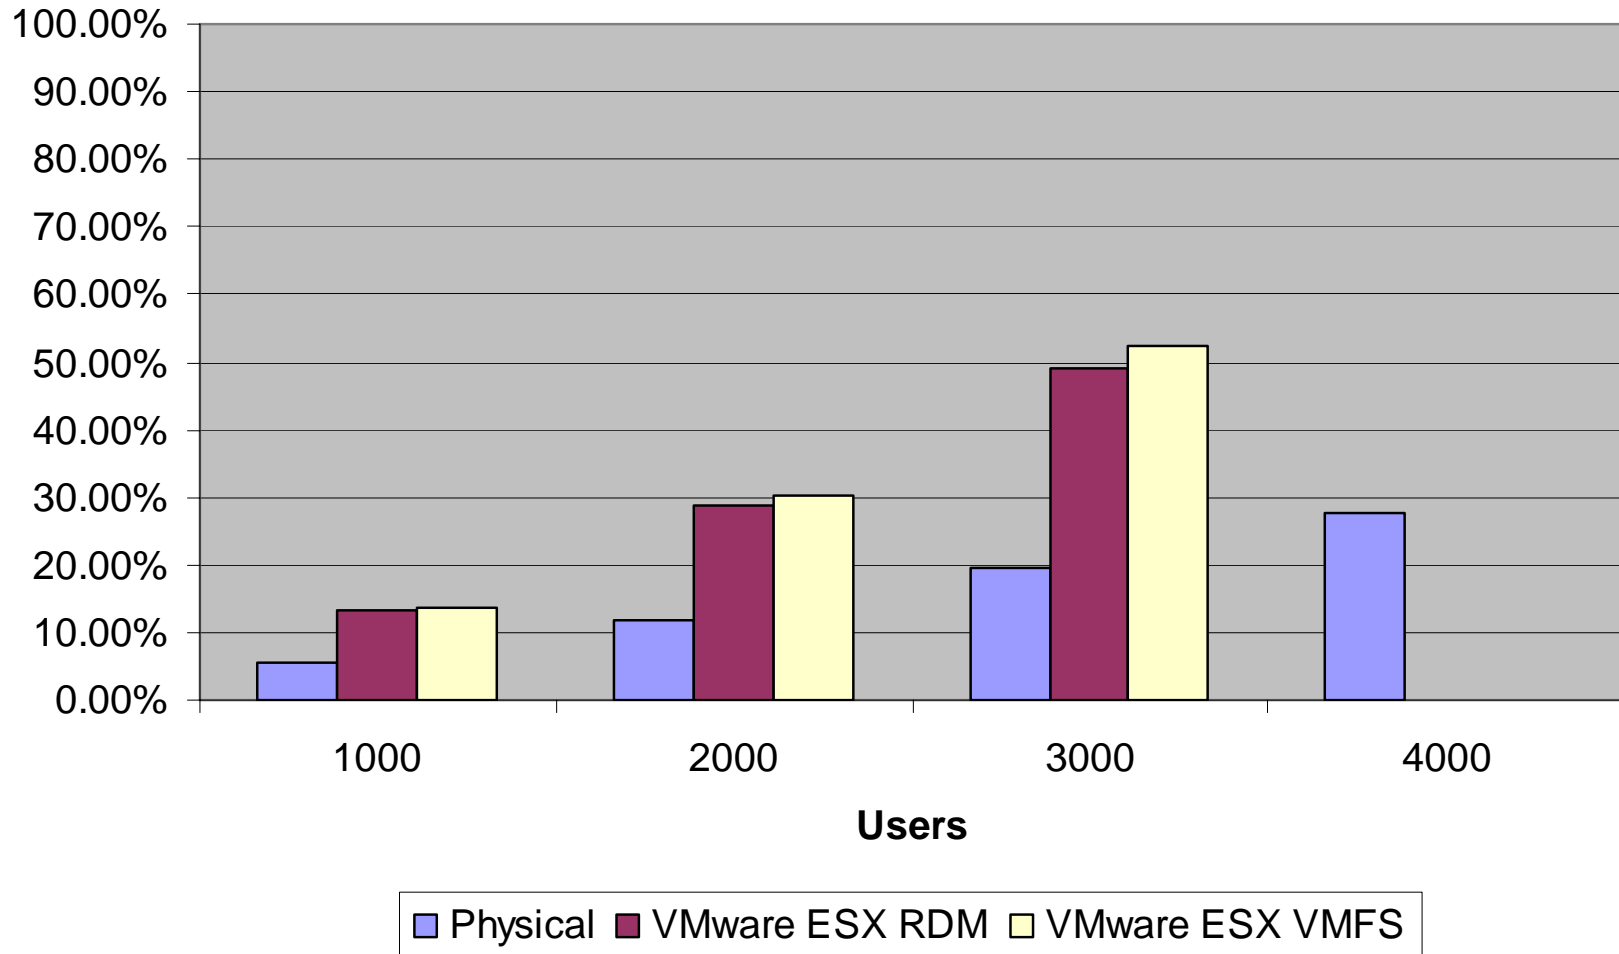
x366 CPU utilization	1000 users	2000 users	3000 users	4000 users
Physical	5.50%	11.80%	19.50%	27.70%
VMware ESX RDM	13.40%	28.70%	49.20%	
VMware ESX VMFS	13.70%	30.20%	52.40%	

x366 response time (ms)	1000 users	2000 users	3000 users	4000 users
Physical	56	99	155	696
VMware ESX RDM	50	91	180	
VMware ESX VMFS	927 (*)	87	164	

(*) We suspect this result value to be incorrect, we didn't have time to repeat this run on time for this presentation

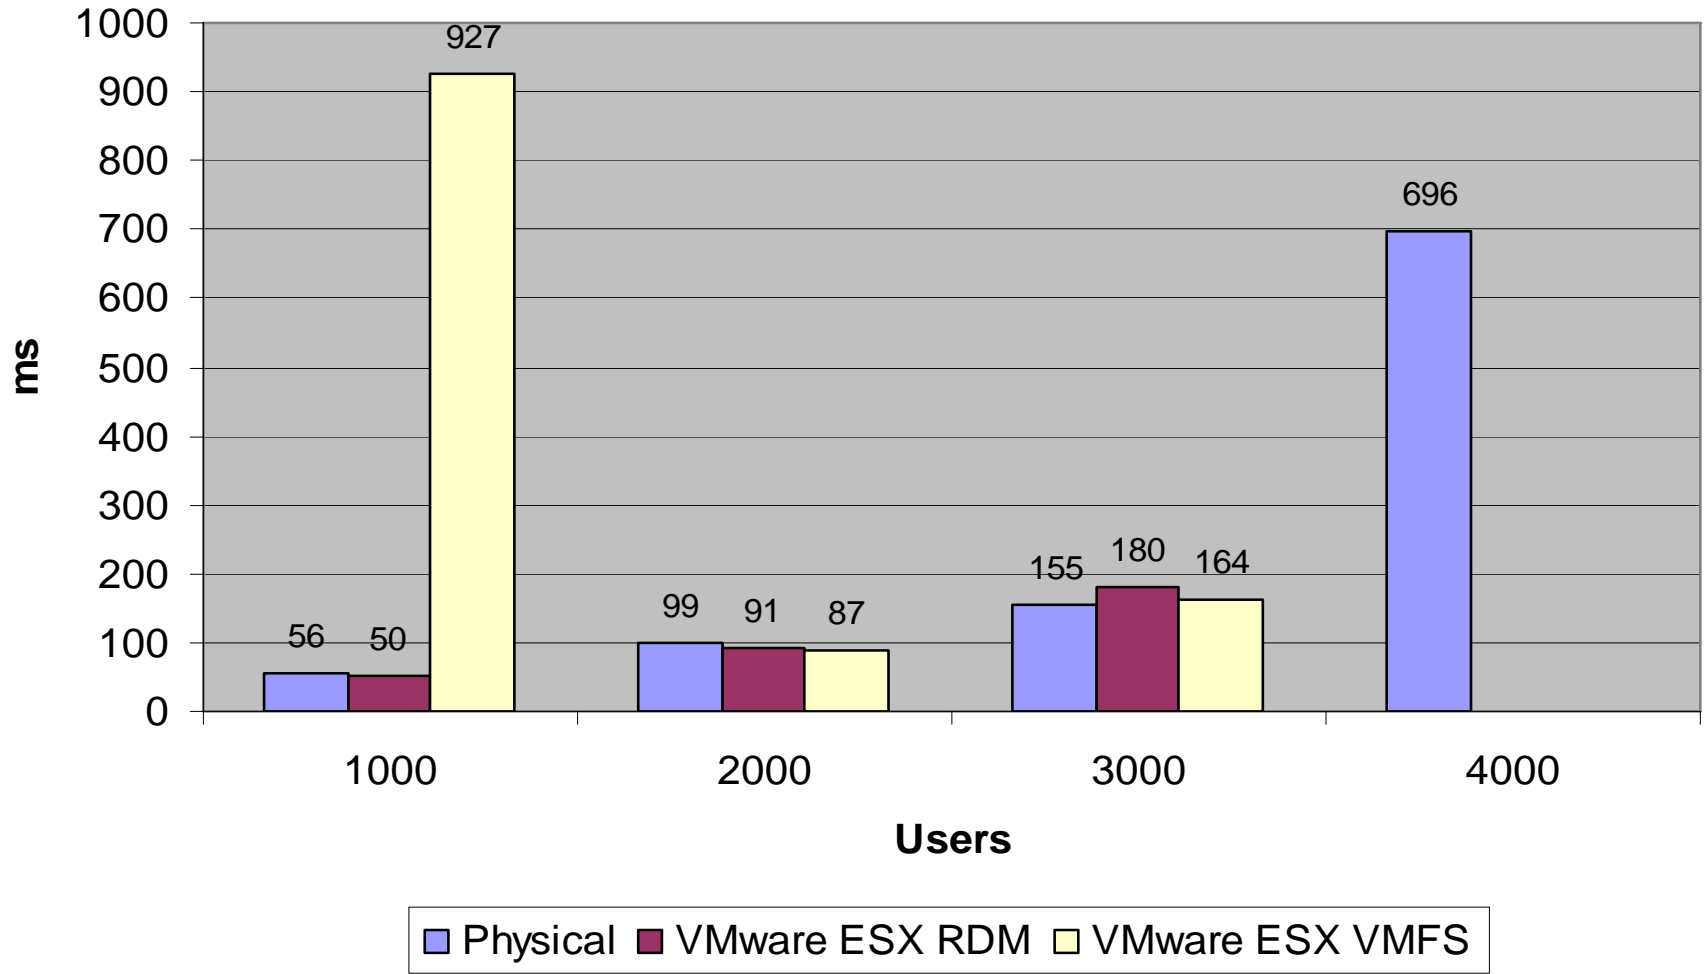
IBM System X 366

Average CPU Utilization



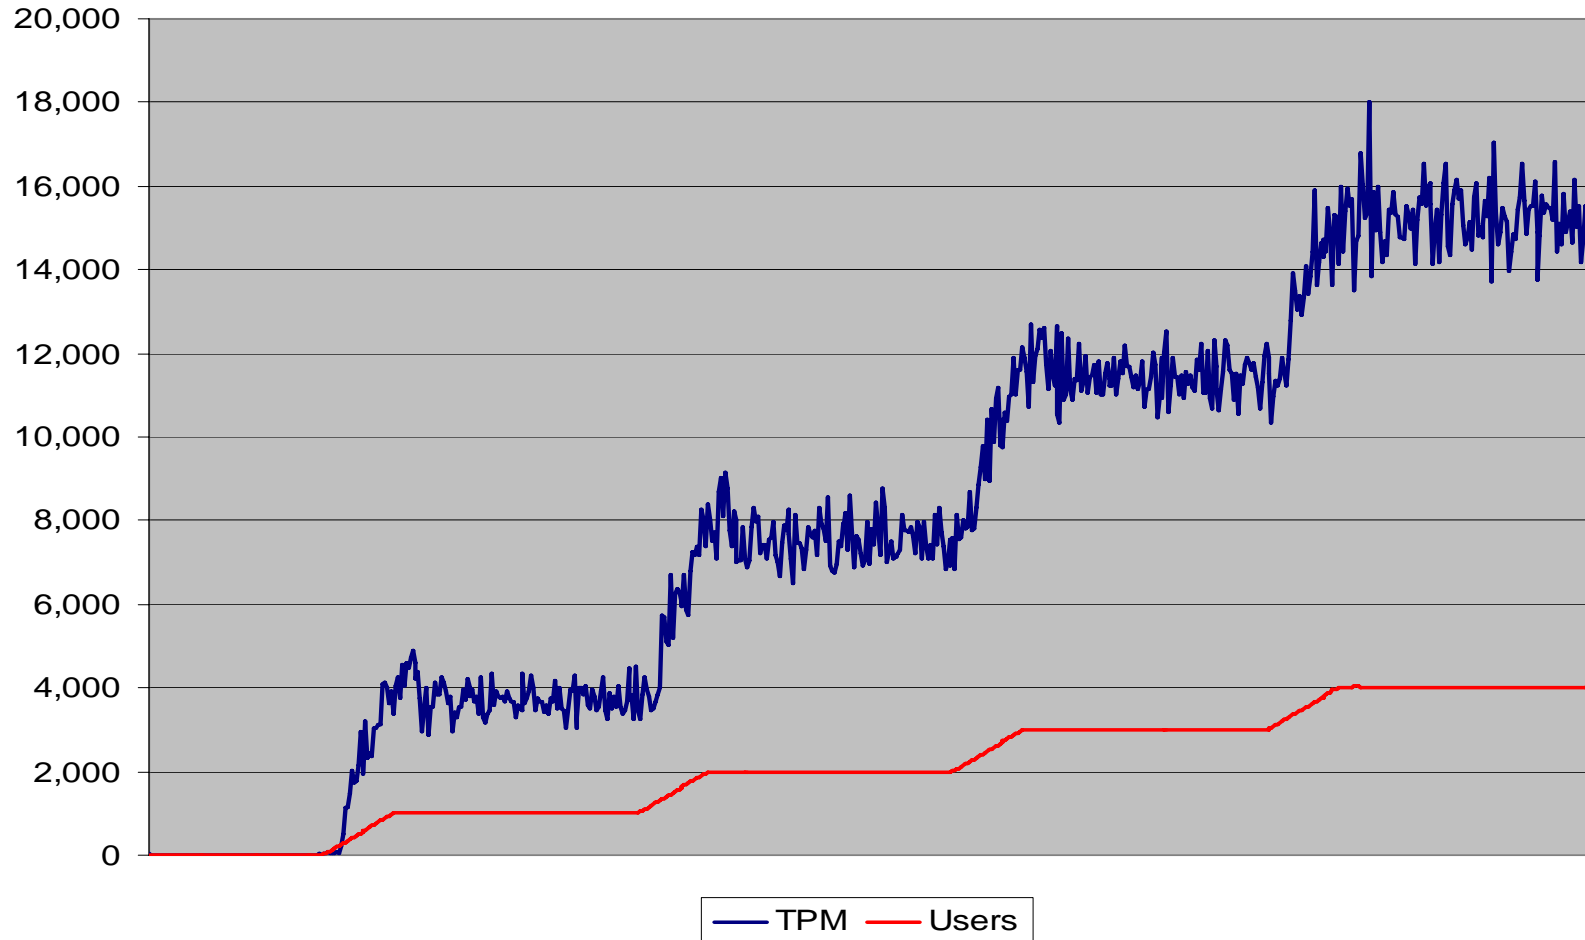
IBM System X 366

Average Response Time



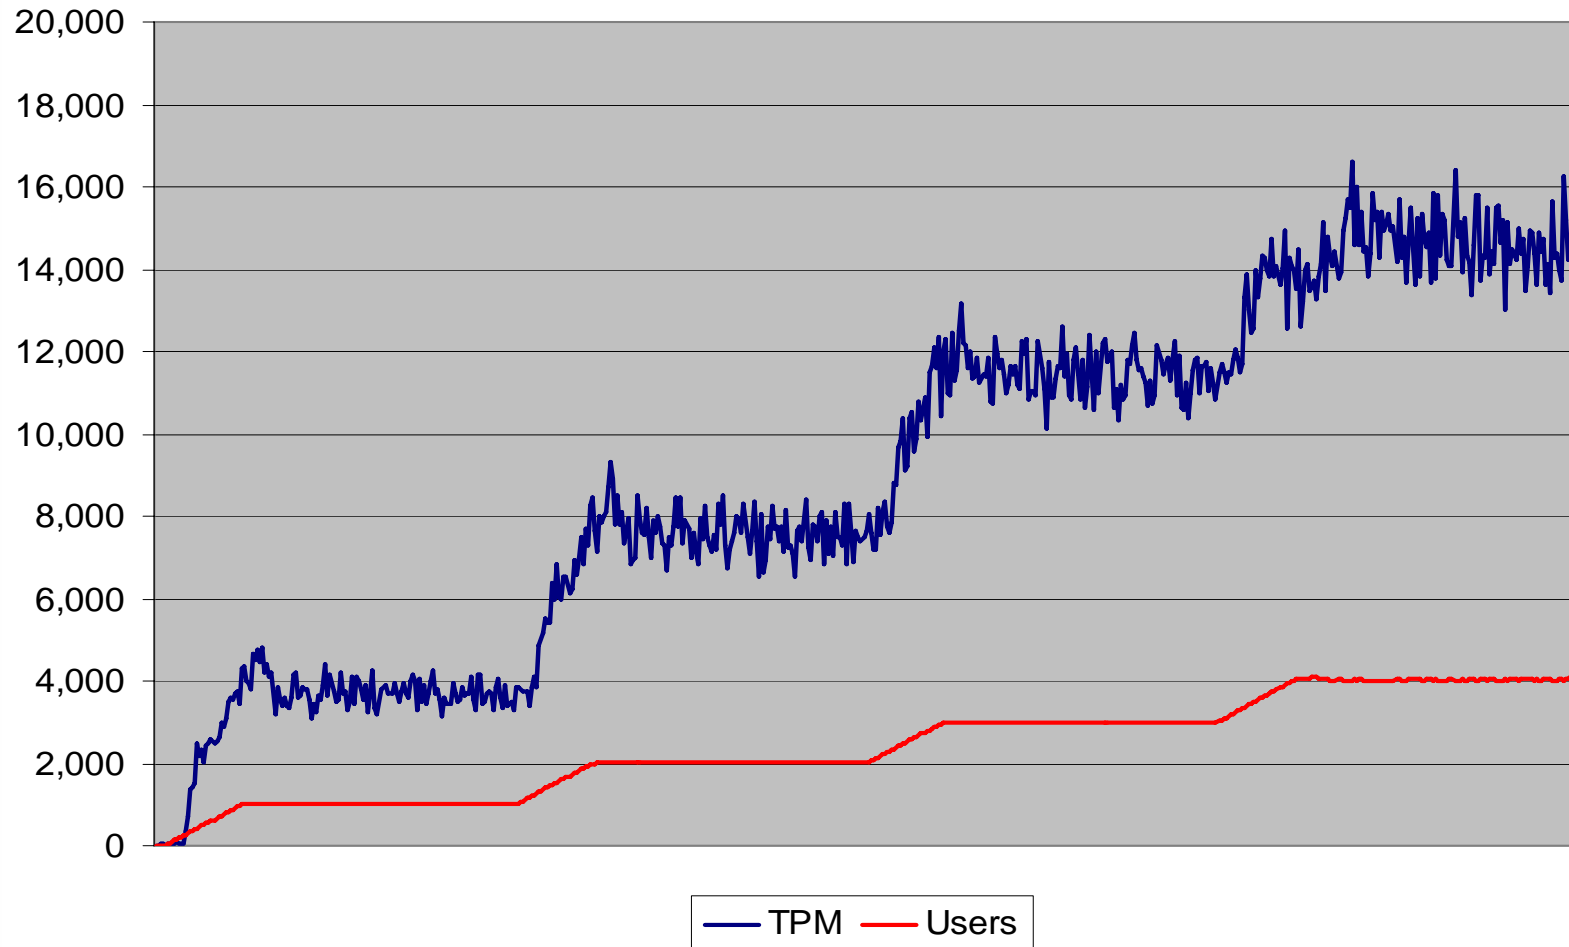
IBM System X 366

Transactions/Minute – Physical Test



IBM System X 366

Transactions/Minute – VMware – RDM



IBM System X 366 Results Comments

- **CPU utilization is 2.5x when using RDM**
 - > Mostly caused by virtualization of disk I/O
- **Response time is 16% higher at 3,000 users**

Intel Clovertown whitebox details

OS: Microsoft Windows 2003 R2 Enterprise Edition SP2 64 bit

○ Physical test

- > 4 CPUs
- > HT disabled
- > 7.25 GB RAM
- > Boot disk & OS C: (local)
- > Program Directory E: (SAN RAID-5)
- > Data Directory F: (SAN RAID-5)
- > Transaction logs G: (SAN RAID-1)
- > Page file on E: configured for 8192-20000 MB
- > /numproc=4 /burnmemory=8960 in boot.ini

○ VMware test

- > 4 vCPU
- > HT disabled
- > 7.25 and 14.5 GB RAM
- > Boot disk & OS C: (SAN RAID-5)
- > Program Directory C: (SAN RAID-5)
- > Data Directory F: (SAN RAID-5)
- > Transaction logs G: (SAN RAID-1)
- > Page file on C: configured for 8192-20000 MB

- > Single VM running on VMware ESX 3.0.1

Intel Clovertown results



1,000-6,000 users

Clovertown CPU utilization	1K users	2K users	3K users	4K users	5K users	6K users
Physical	3.50%	7.50%	11.40%	17%	23.40%	
VMware ESX RDM	13.70%	25.90%	32.50%	34.20%	45.30%	
VMware ESX RDM 2x RAM	8.30%	16.20%	31.10%	47.20%	61.30%	
VMware ESX VMFS	11.10%	22.10%	31.20%	39.90%	52.80%	
VMware ESX VMFS 2x RAM	9.10%	18.60%	29.70%	46.00%	64.60%	

Clovertown response time (ms)	1K users	2K users	3K users	4K users	5K users	6K users
Physical - 7.25 GB RAM	31	73	56	96	177	
VMware ESX RDM - 7.25 GB RAM	93	78	127	546		
VMware ESX RDM - 14.5 GB RAM	31	197	103	384		
VMware ESX VMFS - 7.25 GB RAM	45	68	70	113	491	
VMware ESX VMFS - 14.5 GB RAM	22	87	861	65	124	

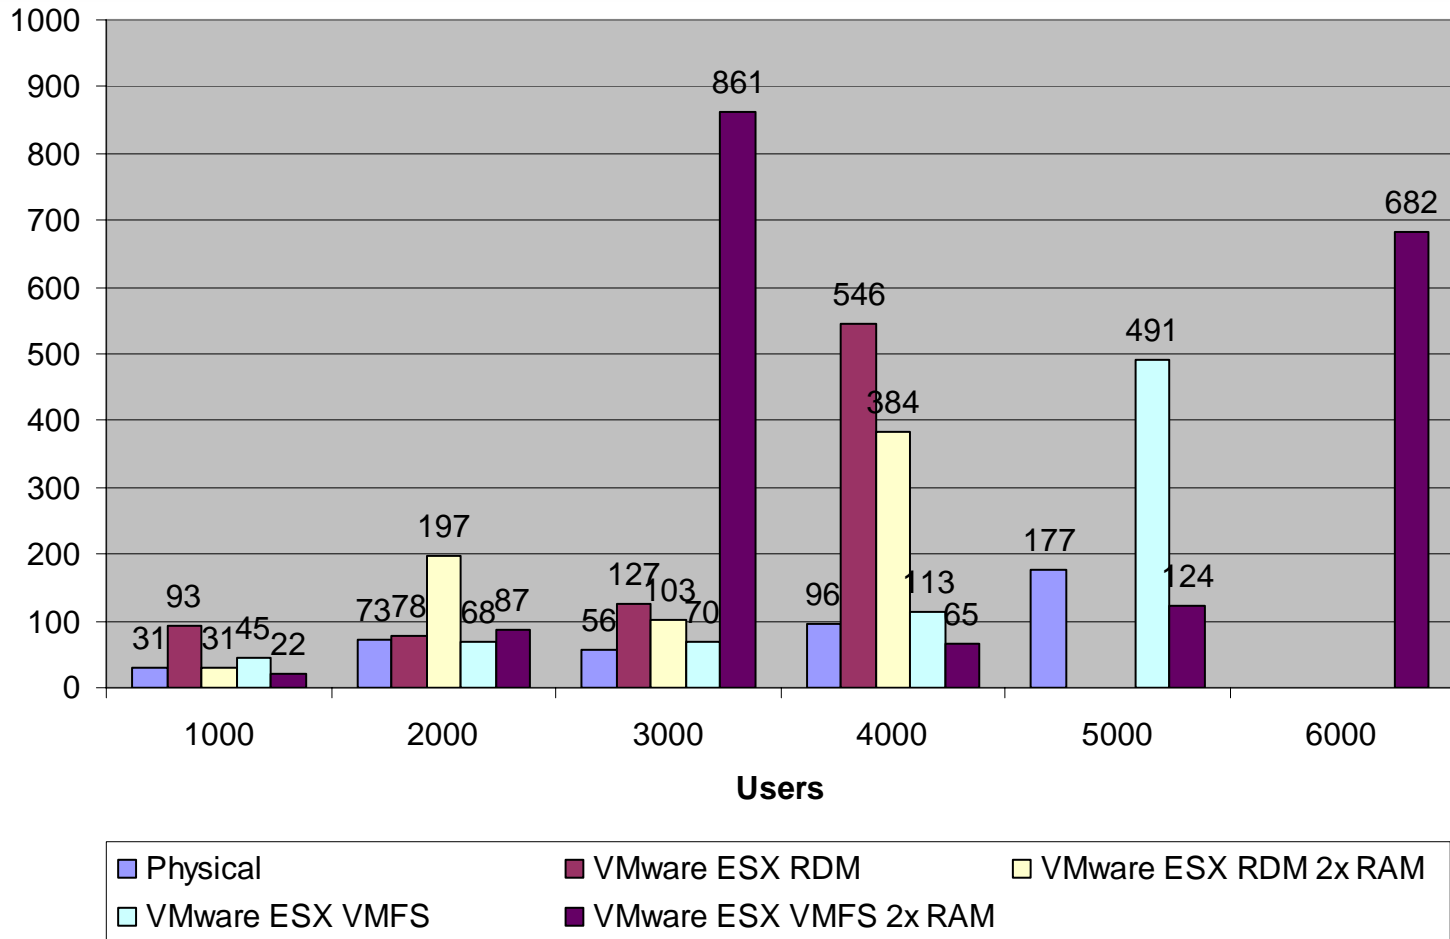
Intel Clovertown

Average CPU Utilization




Intel Clovertown

Average Response Time




Intel Clovertown

Transactions/Minute – VMware – RDM



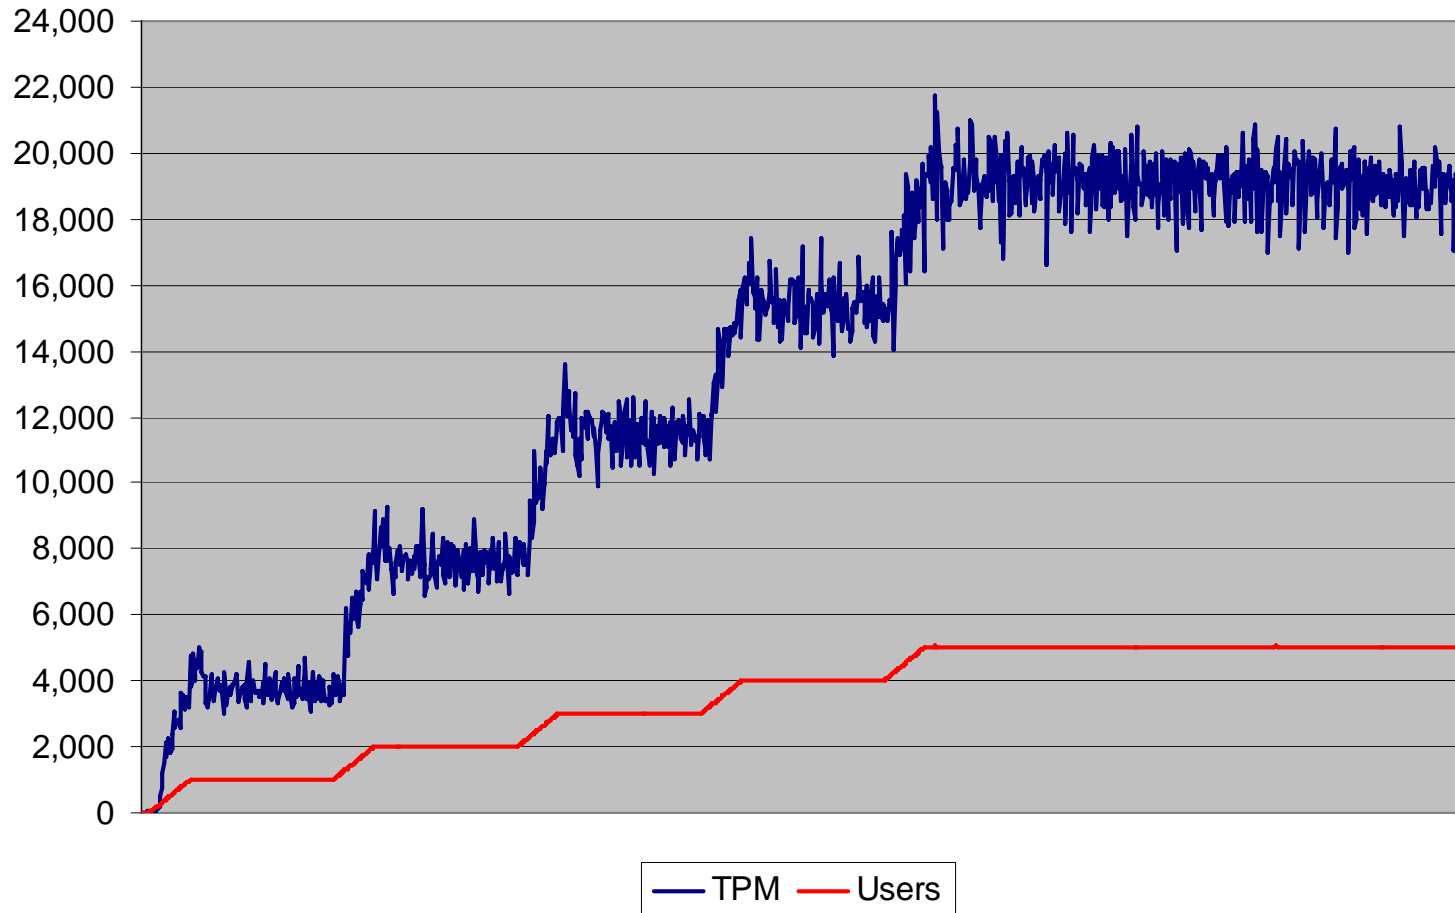
Intel Clovertown

Transactions/Minute – VMware – RDM – 2x RAM



Intel Clovertown

Transactions/Minute – VMware – VMFS



Test Results Comments – High-End System

○ **VMFS provided better performance than RDM**

- This could be different on other environments and for different workloads than mail NRPC
- VMware engineers are investigating these results, and should have more conclusive statement in an upcoming whitepaper

○ **RDM used less CPU than VMFS**

- May be good in some situations or for different workloads
- However, for NRPC mail it gives worse response time than VMFS

○ **On W64 doubling the RAM for a VM increased performance (faster response time)**

- 64 bit OSES can take advantage of additional RAM and use it for file caching
- I/O is very expensive in a virtualized environment
- Avoiding I/O by increasing RAM for the VM is often the easiest and cheapest way to improve performance
- Combined with VMFS it allows up to 6K users with good response time (no comparison with native physical system was performed)

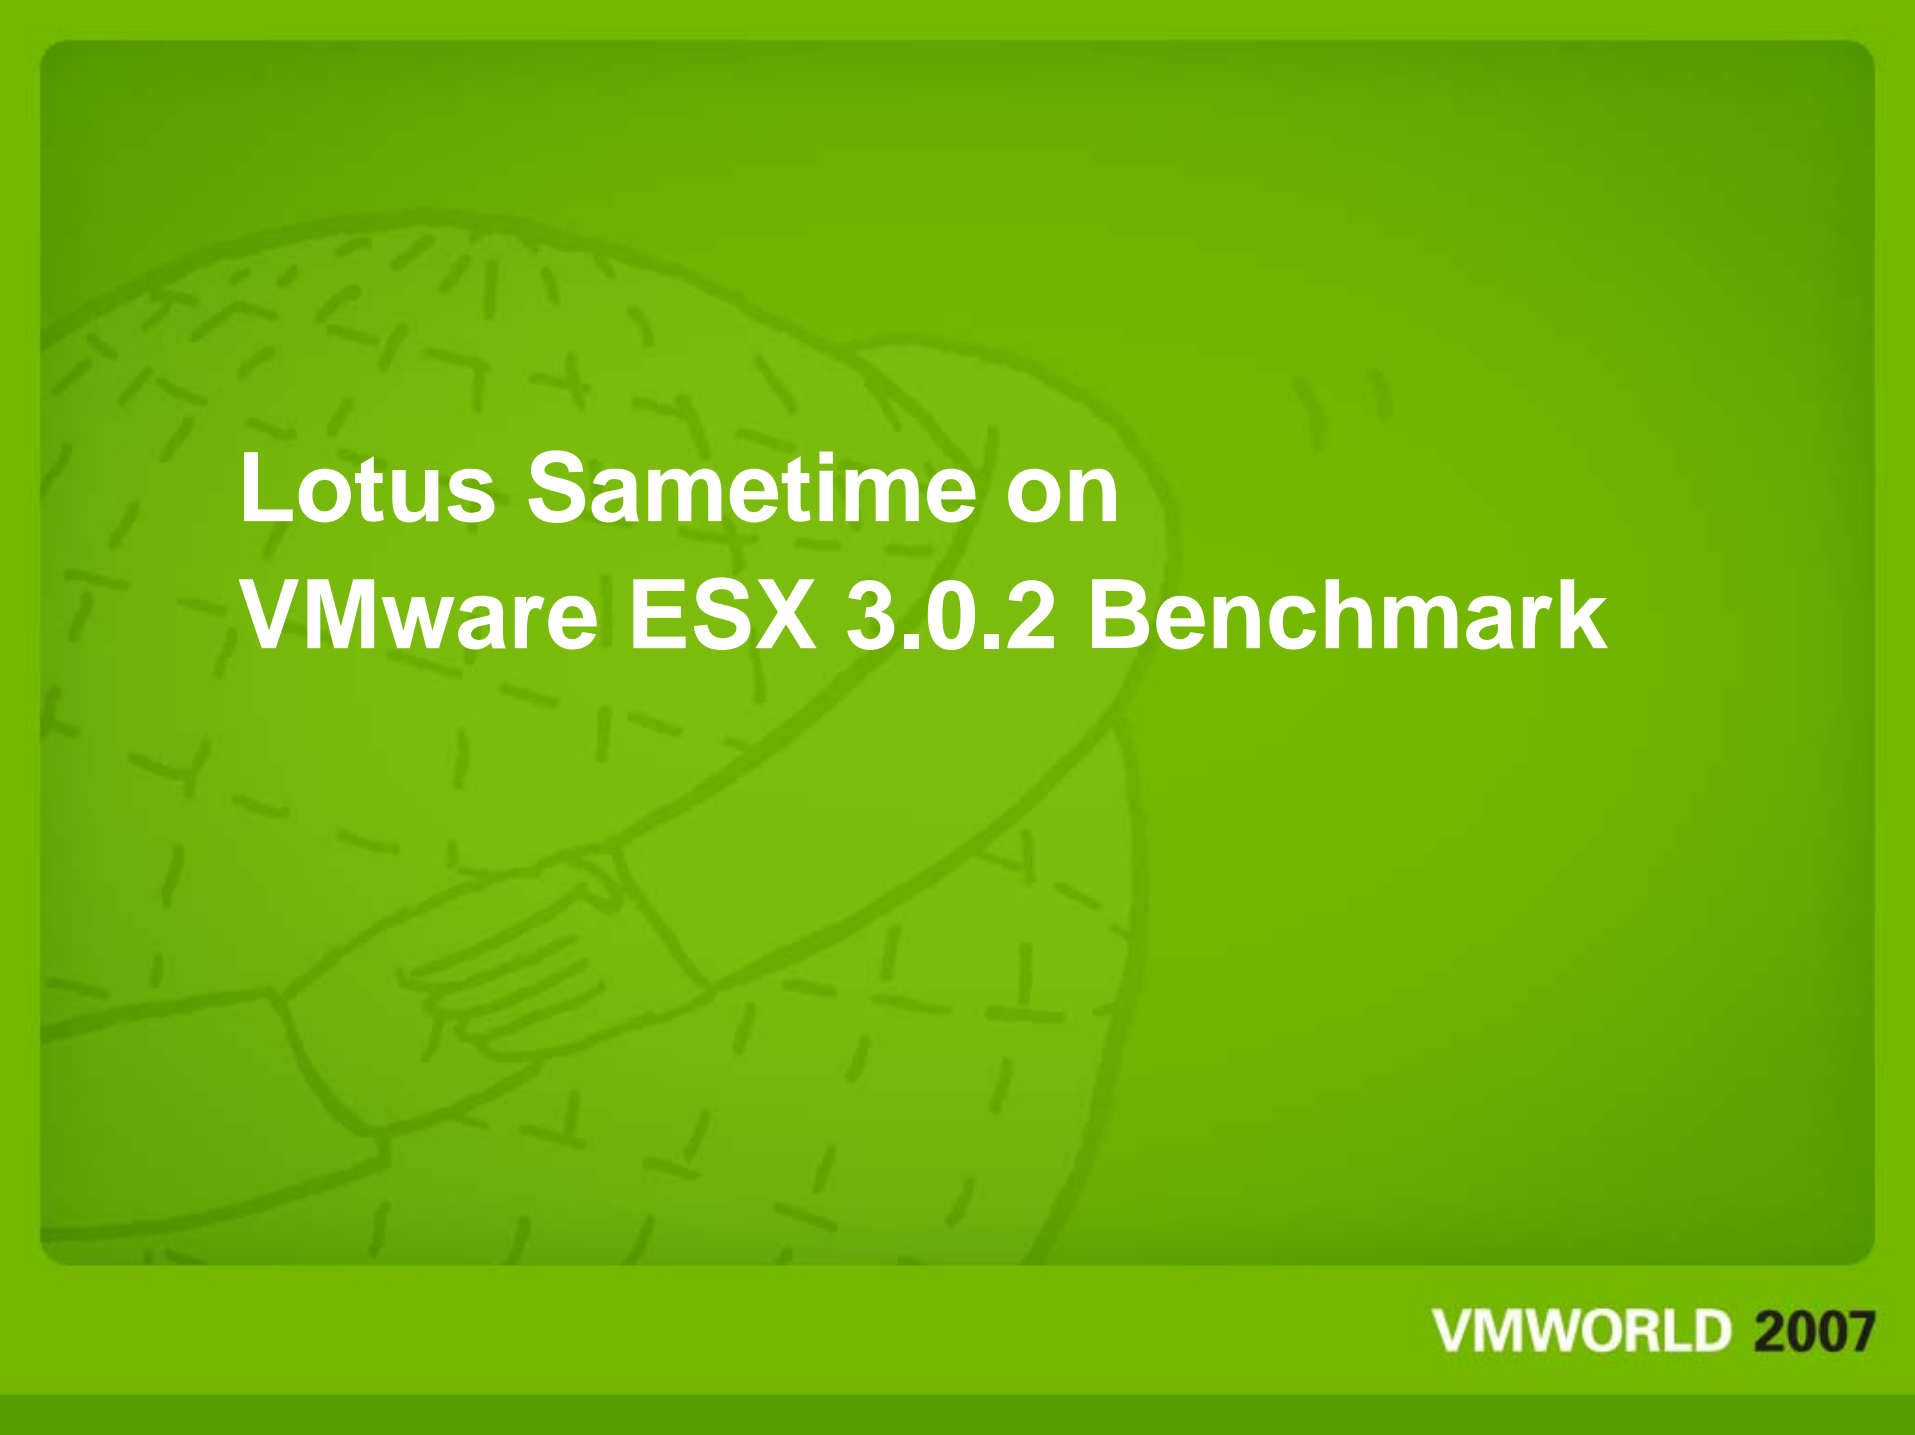
Old vs. New Hardware

A Migration Path to VMware ESX

Response time (ms)	1K users	2K users	3K users	4K users	5K users	6K users
x366 Microsoft Windows 2003 EE 32 bit Physical	56	99	155	696		
Clovertown Microsoft Windows 2003 R2 EE 64 bit VMware ESX RDM 7.25 GB RAM	93	78	127	546		
Clovertown Microsoft Windows 2003 R2 EE 64 bit VMware ESX VMFS 7.25 GB RAM	45	68	70	113	491	
Clovertown Microsoft Windows 2003 R2 EE 64 bit VMware ESX VMFS 14.5 GB RAM	22	87	861 (*)	65	124	682

- Identical workload and SAN storage configuration
- Numbers in **green** represent better results on new virtualized environment vs old physical environment

(*) we suspect this result value to be incorrect, we didn't have time to repeat this run on time for this presentation

A faint, light-colored illustration of a hand holding a globe is visible in the background of the slide. The hand is positioned at the bottom left, with fingers wrapped around the base of the globe. The globe shows latitude and longitude lines. The entire background is a solid green color.

Lotus Sametime on VMware ESX 3.0.2 Benchmark

IBM Lotus Sametime Benchmark goal

- **Verify if a complex configuration of IBM Lotus Sametime 7.5.1 software can scale on VMware ESX Server 3.0.2 using a single powerful server**
- **This is a unique solution not possible in the physical world**
 - > 3-5 physical servers would be required to scale beyond 35K concurrent users
- **Target: 100,000 concurrent chat users on a single ESX server**
 - > In production you may want to host this on multiple ESX Servers for fault tolerance and disaster recovery
- **Chat (instant messaging) only**
 - > No instant meetings (collaboration) nor scheduled meetings (Web conference)

Hardware

○ IBM System X 3650

- › 24 GB RAM
- › 2 QuadCore Intel Xeon CPUs 2.33 GHz
- › QLogic QLA2432

○ SAN DS4800

- › 8 GB cache
- › 15K RPM
- › 36 GB disks

System X 3650 Configuration for IBM Lotus Sametime Benchmark on ESX 3.0.2

○ 2 Lotus Sametime Community Server

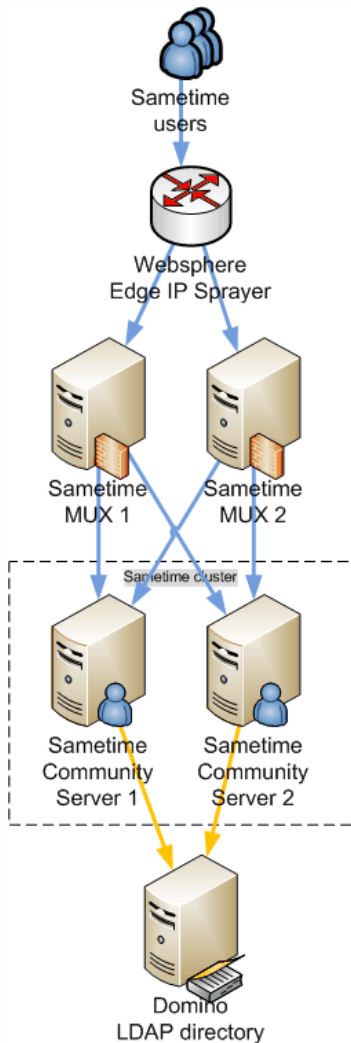
- > Microsoft Windows 2003 Standard Edition SP 1 32 bit
- > Lotus Domino 7.0.2 FP2
- > Lotus Sametime 7.5.1
- > 2 vCPU
- > 4 GB RAM

○ 2 Lotus Sametime Multiplexors

- > Microsoft Windows 2003 Standard Edition SP 1 32 bit
- > Lotus Domino 7.0.2 FP2
- > Lotus Sametime 7.5.1
- > 1 vCPU
- > 2 GB RAM

HT was disabled in the BIOS during this test

IBM Lotus Sametime Benchmark



○ On VMware ESX:

- > 2 x IBM Lotus Sametime Community servers in a cluster
 - 4 GB RAM, 2 vCPUs, Microsoft Windows 2003 Standard Ed.
- > 2 x IBM Lotus Sametime Multiplexers
 - 2 GB RAM, 1 vCPUs, Microsoft Windows 2003 Standard Ed.
- > Storage: All virtual disks were on the same VMFS LUN
- > Directory: clustered Lotus Domino LDAP 7.0.1

○ Benchmark scaled only for chat using Lotus Sametime Connect client

○ No scheduled or instant meetings included in this test


IBM Lotus Sametime Benchmark

Average Operations per Minute During the Benchmark

- **419 users logged in**
- **900 users change status, resulting in 90,000 status change notifications**
- **3,160 IM channels created**
- **4,714 small IMs sent**
- **3,160 large IMs (40 KB; i.e., screen capture) sent**
- **900 User Info requests (business card)**
- **400 reads and 267 updates made from/to storage (buddylist 30 KB)**
- **16 n-way (3 or more parties) chats are started**

Transaction Peak in a 5-Minute Interval


Community Server 1



■ ConcurrentLogins ■ ImMsgs ■ NewImCnls ■ NumberOfStatusChanges

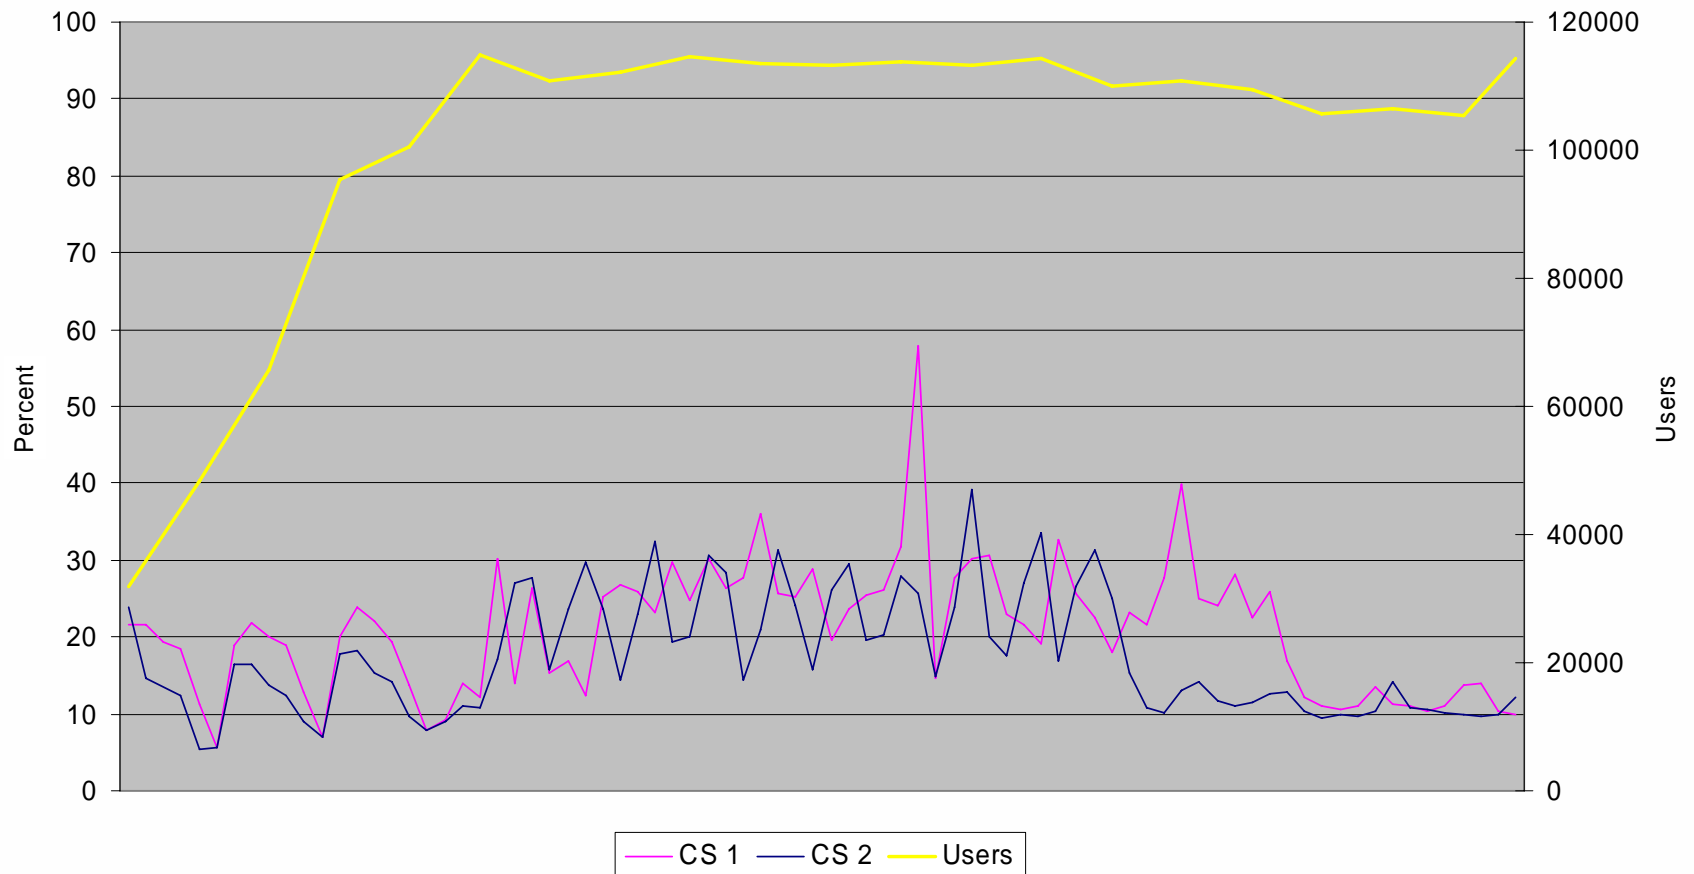
Transaction Peak in a 5-Minute Interval

Community Server 1




■ ConcurrentNWChats ■ FailedLoginOperations ■ NumberOfBuddyListChannelCreations ■ SuccessfulLoginOperations

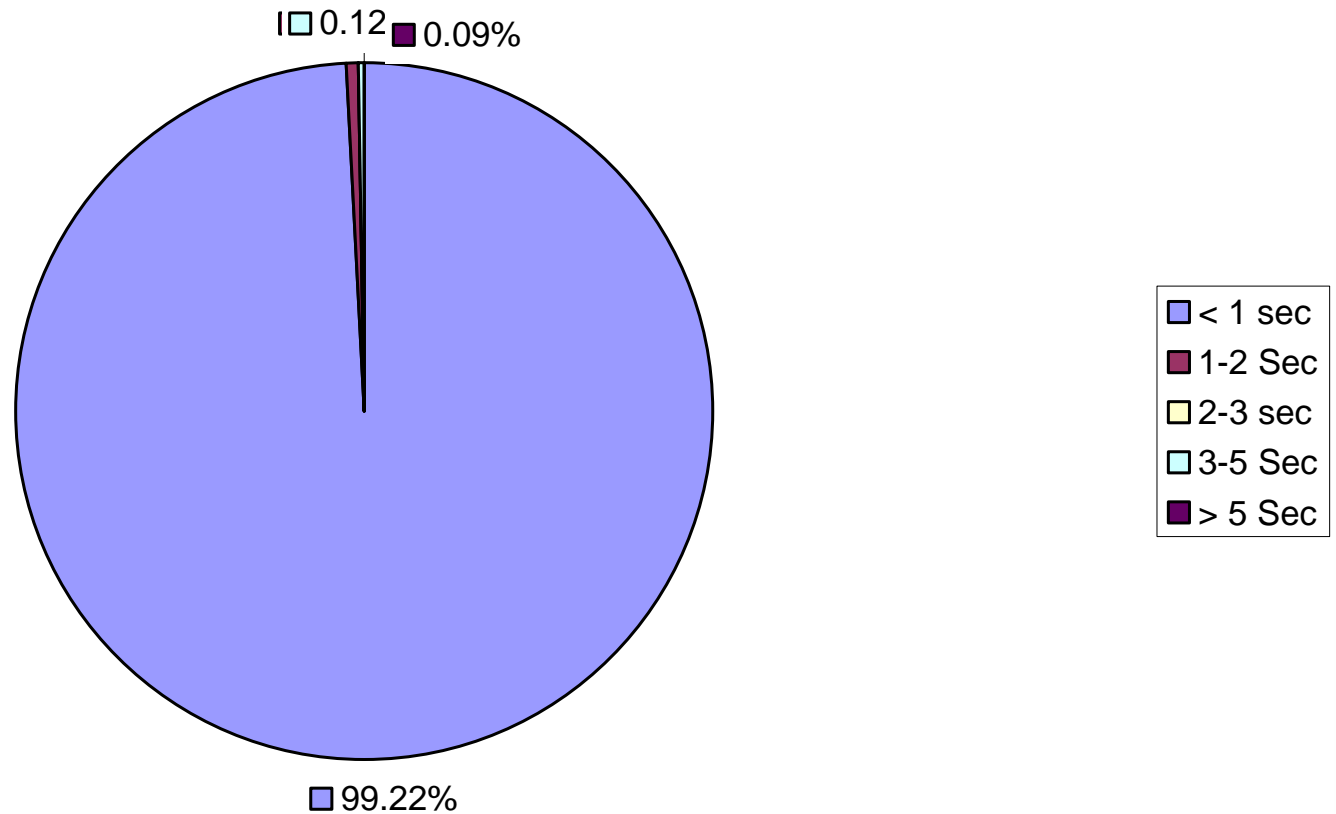
Peak of CPU Utilization for the Lotus Sametime Community Server



Peak CPU Utilization for Lotus Sametime Multiplexers



Response Time for all Transactions



Comments on IBM Lotus Sametime benchmark

- **In this benchmark we used a clustered remote LDAP directory**
 - > No pictures were present in this directory. If present it would increase load on Sametime server as more data is being transferred.
- **When using local Lotus Domino Directory keep in mind that:**
 - > CPU utilization could be higher, as authentication is performed locally
 - > Disk speed and configuration is important for best performance (as learned from Lotus Domino benchmark)
- **Notes client is heavier in nature than the Lotus Sametime client**
 - > <http://www.ibm.com/support/docview.wss?rs=203&uid=swg21222797>
- **If multiple clients are used (Lotus Sametime, Notes, Web enabled applications using stlinks) consider each client connection as a separate user for sizing purposes**
 - > For additional information on sizing refer to IBM Redbook (currently in draft) <http://www.redbooks.ibm.com/abstracts/sg247410.html>

A faint, light green line drawing of a globe is centered in the background. A hand is shown pointing towards the globe from the bottom left. The globe has a grid of latitude and longitude lines, and the hand is rendered with simple lines for fingers and a wrist.

Conclusions

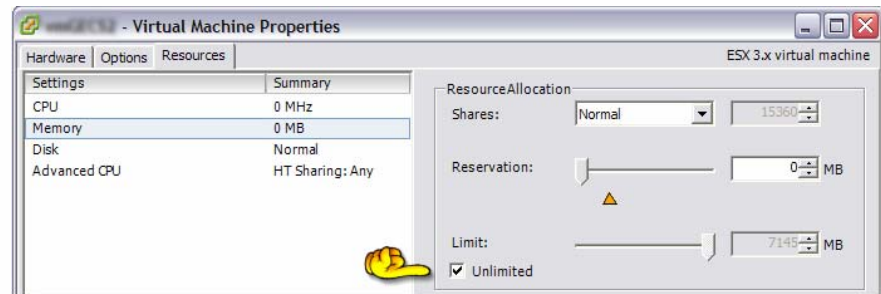

Lotus Domino and Lotus Sametime on VMware ESX in Production Environments

○ Assign the correct resources to your VMs

- VMware ESX does not reinvent the laws of physics: CPUs and GHz, available RAM, network speed, disk subsystem speed and latency do not disappear
- Don't be fooled by free memory of a VM running Domino
 - Lotus Domino allocates more RAM for caching based on total RAM available
 - Example: 1 GB VM → 40% free, same VM with 2 GB → 45% free
- Additional RAM can be used by the OS as disk cache: very important to have plenty of it – nothing is wasted

○ Don't limit the RAM the VM can effectively use

- Lotus Domino calculates its cache size based on OS available memory, it isn't aware of artificial limitations imposed by VI 3
- In Memory >> Resources "Unlimited" really means "up to the amount allocated to the VM", therefore is a correct setting for our VMs



Lesson Learned from Lotus Domino Benchmark and Support Incidents

- **Use dedicated LUNs with VMFS for Lotus Domino data with a single VMDK virtual disk on it**
 - Single virtual disk typically provides better performance, helps to avoid potential VMFS wide locks
 - If you need more storage, use more LUNs, each with a single virtual disk
- **Response time is your success rate**
 - Make sure users don't have to wait longer to open their mail or they will complain to the help desk
 - Migrating from older to newer hardware on VMware is good as it may give you same or better performance. Some customers reported up to 30% decrease in response time
 - Virtualizing on the same hardware is not always a good idea: if load is high response time will increase and users won't be happy; plan correctly
- **Pay attention to the disks and network layout**
 - Having multiple Lotus Domino VMs sharing the same local disk isn't good. SAN is typically a better option
 - Even with SAN make sure you have enough bandwidth for concurrent activity from all VMs; verify your SAN can handle all hosts (hosts = number of active VMs, not just ESX servers).
 - Use esxtop to verify SAN queue and latency (press d for disks, f for fields and select A, H, I and J) – maximize window to see all values
 - Check the queue length for both esx (esxtop) and Lotus Domino platform statistics for problems
 - Consider giving a dedicated NIC to busy and critical VMs

Lesson Learned from Lotus Sametime Benchmark and Support Incidents

- **LotusSametime instant messaging servers typically scale well and aren't heavy on network utilization**
 - > Meetings are more network and CPU intensive
- **Using separate MUXes can dramatically increase scalability of the whole environment**
 - > MUXes typically don't need the same amount of memory and vCPU as the Lotus Sametime Community Servers
- **To cluster Lotus Sametime you must use an IP sprayer or equivalent solution**
 - > We used IBM WebSphere® Edge Network Dispatcher
 - > DNS roundrobin can be a less expensive alternative, even if isn't typically as robust

Lotus Domino and Lotus Sametime Software on VMware ESX in Production Environments

- Giving more resources to the VM will often prevent difficult to troubleshoot issues
- Overcommit resources can be a dangerous game
- Make sure to read the ESX tuning guide for best practices
 - > http://www.vmware.com/pdf/vi_performance_tuning.pdf
- Having additional vCPUs available can help to recover from network outages
 - > Think what happens if your WAN link goes down for 30 seconds: users get disconnected and will attempt to reconnect all at once

Considerations

- ◊ In the benchmark we simulate the user activity
- ◊ While the workload is designed to mimic the user activities, it is standard practice to divide the results by a factor of 2

“Your mileage may vary.” May cause addiction. If symptoms last more than 4 hours seek professional counseling. Do not use in direct sunlight, expose to extreme temperature or to rain. It’s amazing how people keep on reading this. If you can read this you don’t need glasses.

We'd Like to Thank

Garrett Drysdale, Yashmeet Khopkar, Joseph Ellis, Jantz Tran, and Sreekanth Keesara of Intel for their invaluable advice, support, and for the hardware provided for this benchmark

also:

Kristen Edwards, Hemant Gaidhani, Robert Campbell, and Murli Kotteyil of VMware for their help and support

and finally:

Emmet Clifford, Brendan Furey, Joseph Grace, James Powers, Frank Sperlinga, Carol Stout, and William Hankard of IBM for always supporting and helping us

Without them all we wouldn't be here today.

© IBM Corporation 2007. All Rights Reserved.

The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this publication or any other materials. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

IBM, the IBM logo, Lotus, Lotus Domino, Sametime, WebSphere, xSeries and System X are trademarks of International Business Machines Corporation in the United States, other countries, or both.

Microsoft and Microsoft Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.



Q&A

**We appreciate your comments –
please fill in your evaluations**

**Session WV21:
IBM Lotus Domino and IBM Lotus
Sametime on VMware Infrastructure 3:
From Planning to Successful Deployment**

**Ivan Dell'Era and Purvi Trivedi
IBM**



VMWORLD 2007

EMBRACING YOUR VIRTUAL WORLD

BREAKOUT SESSION