


Virtual Reality Check



**Hyper-V 2008R2, vSphere 4, XenServer 5.5
on Intel 'Nehalem' Xeon 5500**

"Project VRC: Phase II"

Author(s) :	Ruben Spruijt, Jeroen van de Kamp
Version:	1.0
Date:	February 2010

©2010 PQR and Login Consultants, all rights reserved.

All rights reserved. Specifications are subject to change without notice. PQR and Login Consultants, the PQR and Login Consultants logo and its tagline Eenvoud in ICT are trademarks or registered trademarks of PQR and Login Consultants in the Netherlands and/or other countries. All other brands or products mentioned in this document are trademarks or registered trademarks of their respective holders and should be treated as such.

DOCUMENT OVERVIEW

HISTORY

Version	Date	Author(s)	Remarks
0.95	January 2010	Ruben Spruijt & Jeroen van de Kamp	
0.96	January 2010	Ruben Spruijt & Jeroen van de Kamp	Update Triple Memory Channel Tests
0.97	January 2010	Ruben Spruijt & Jeroen van de Kamp	Review VMware
1.0	February 2010	Ruben Spruijt & Jeroen van de Kamp	

REVIEWERS

Version	Date	Reviewers	Company
0.95	January 2010	Herco van Brug	PQR

CONTENT

1.	Executive Summary	3
2.	Introduction Project VRC.....	5
2.1	Project VRC objectives	5
2.2	Intended audience.....	5
2.3	Better together.....	6
2.4	Vendor Involvement	6
2.5	Contact.....	6
3.	About the authors	7
3.1	About Login Consultants	7
3.2	About PQR.....	7
3.3	Team Members	8
4.	The Login VSI Benchmark.....	9
4.1	VSI overview.....	9
4.2	VSI 2.1 Medium workload	10
4.3	Randomization	11
4.4	VSImax 2.1.....	12
5.	The VRC platform.....	14
5.1	Hardware configuration	14
5.2	Launcher Configuration.....	15
5.3	External Clock.....	15
5.4	Test Constants.....	15
6.	Project VRC Methodology & Best Practices.....	16
6.1	vCPU's	16
6.2	vCPU overcommit & dedicated hardware	16
6.3	Transparent Page Sharing	16
6.4	Interpreting Project VRC results	17
6.5	VRC, VSI and Clocks Reviewed.....	17
6.6	VSImax 2.1.....	18
6.7	Test Descriptions	18
7.	Miscellaneous Phase 2 Results	19
7.1	Office 2007 Service Pack 1 & 2	19
7.2	IE7 vs IE8	19
7.3	HP Services.....	20
7.4	XenApp CPU Fairshare	20
7.5	Xenserver Xenapp Optimization on Nehalem.....	21
7.6	vSphere 4.0 vMMU impact.....	21
7.7	Triple Memory Channels	22
8.	NEXT GEN Hypervisor and Hardware: Impressive.....	23
9.	Virtualizing Windows 2003x86 TS Workloads.....	25
9.1	vSphere HaltingIdleMsecPenalty	25
9.2	vSphere 2vCPU vs 4vCPU vs HT with VSI 2.1 & VSI 1.1.....	25
9.3	Virtualizing TS 2003x86 Workloads Not Utilizing Hyper-threading	26
9.4	Virtualizing TS 2003x86 Workloads Utilizing Hyper-threading.....	27
10.	Virtualizing TS 2008x64 Workloads.....	28
10.1	4 VM's with 4vCPU's vs Bare-Metal	28
10.2	"Big Iron" Tests vs Bare-Metal.....	29
11.	Overview of All Tests	30

1. EXECUTIVE SUMMARY

This Project VRC phase 2 whitepaper focuses completely on analyzing Terminal Services (TS) workloads running on the latest generation hardware and hypervisors. While in phase 1 of project VRC VMware ESX3.5, Citrix XenServer 5.0 and Microsoft Hyper-V 1.0 were tested on a HP DL385G5 with AMD “Barcelona” processors, this whitepaper reviews vSphere 4.0, XenServer 5.5 and Hyper-V 2.0 running on HPDL380G6 with Intel “Nehalem” processor.

Project VRC made a deliberate choice not to directly compare Hypervisors in phase 1 (and published results in separate whitepapers), the community and the vendors, logically, did. Therefore, we decided that Project VRC will release this single whitepaper for the phase 2 results, including the three hypervisors. This is also possible with Login VSI 2.1, the product independent VDI & TS benchmark, which allows the external timing and calibration of the workload, ensuring that the published results are not influenced by the potential clock drifts within the virtual machines.

The most important conclusions of the phase 2 tests are:

- Virtualizing 32-bit Terminal Server and Citrix XenApp workloads remains highly recommended on all hypervisors (when using the latest generation AMD or Intel processors): running x86 TS workload on a hypervisor in comparison to bare-metal configuration is far more scalable.
- Virtualizing 64-bit Terminal Server and Citrix XenApp workloads requires a little more thought. Windows Server 2008x64 Terminal services has been designed from the ground up to scale well on the latest hardware generation. In contrast to x86, x64 Terminal Server is not automatically more scalable when virtualized. For example, the maximum amount of vCPU’s possible today per VM is 8, as a result it is necessary to setup 2 VM’s to utilize all 16 logical processors.
- While it is possible to setup many relatively small x64 TS VM’s, this contradicts with the concept of running x64 TS, which has been build and designed from the ground up to host many users on a single Windows instance. It still makes a lot of sense to virtualize x64 TS: but not specifically for performance reasons. The management and consolidation benefits virtualization brings are also valid for x64 Terminal server workloads.
- The Intel Nehalem processor architecture brings impressive performance increases for Terminal server workloads on all hypervisors. The move from HP’s AMD (G5) to Intel (G6), with similar configurations, almost doubled the amount of users on every platform. Both Hyper-Threading and EPT-D introduced by Intel Nehalem have an impressive positive impact on Terminal Server workloads. This is also logical, because Terminal Server workloads are typified by the rather extreme amount of processes, threads and context switches within the VM. Having more logical processors available and the performance optimization addressing memory page tables is specifically beneficial to terminal server workloads.
- It is also safe to conclude that these spectacular performance increases are not caused by specific developments on a hypervisor level. Intel’s innovations in the Nehalem architecture can be almost solely accredited for the performance improvements seen with TS workloads. It is up to the hypervisors to support/enable the hardware innovations.

- When comparing hypervisors, performance is close to equal throughout when no Hyper-threading is used by the VM's. In all test the hypervisors perform with a 5% range with the Terminal Server workloads, with a slight edge for vSphere 4.0. Utilizing Hyper-Threading, on all platforms a performance increase is seen, but vSphere 4.0 trailing slightly by 15% in comparison to XenServer 5.5 and Hyper-V 2.0. These differences are only visible under full load, under normal operating conditions all hypervisors perform practically equal. When VMware releases specific updates to improve the performance of Terminal Server workloads, Project VRC will review and update the published results.
- Strikingly, XenServer 5.5 and Hyper-V 2.0 perform almost identical in all tests. The only differentiator between these two is that XenServer 5.5 (and vSphere 4.0) support 8vCPU's, where Hyper-V 2.0 has a maximum of 4vCPU per VM.

2. INTRODUCTION PROJECT VRC

Welcome to “Project: Virtual Reality Check (VRC)”!

If you are looking for an independent advise and a ‘Reality Check’ in relation to Virtualizing Terminal Server and Desktop workloads, if you are curious about the impact of different hypervisors and the performance differences with various hardware and if you are searching for best practices for your virtual Desktops ... Project VRC whitepapers are a must read!

PQR and Login Consultants started this unbiased and independent R&D project early 2009. The goal of Project VRC is to analyze the developments in the Application- and Desktop Virtualization market and to objectively present the results. All together over 400 tests have been carried out (Q1-2010).

In the haze of extreme innovation rate in the virtualization market and corresponding marketing promises this information is appreciated.

Therefore we published our methods and conclusions in various whitepapers which can be downloaded on www.virtualrealitycheck.net.

2.1 PROJECT VRC OBJECTIVES

The overall goal of Project VRC is to investigate, validate and give answers to the following questions and much more:

- What is the true impact of innovations on a hardware and hypervisor level?
- Which performance optimization on the host and guest virtualization level can be configured, and what is the impact of these settings on user density?
- With the introduction of the latest hypervisor technologies, can we now recommend running large scale TS/CTX workloads on a virtualization platform?
- How does a VDI infrastructure scale in comparison (virtualized) Terminal Server?
- How do the two usage scenarios compare, that is Microsoft Terminal Server [TS] only, versus TS plus XenApp?
- How does various Microsoft Windows Client OS's scale as a virtual desktop?
- How do x86 and x64 TS platforms compare in scalability on bare metal and virtualized environments?
- What is the best way to partition (memory and vCPU) the Virtual Machines the hypervisor host, to achieve the highest possible user density?
- What is the impact of the latest and greatest hardware on (virtualized) terminal servers and desktops?

Project VRC is not finished, and probably never will be. We look forward to evaluate new innovations in the hypervisor arena, hardware level, Application Virtualization and Remoting Protocols.

Project VRC publishes their findings on www.virtualrealitycheck.net.

2.2 INTENDED AUDIENCE

This document is intended for IT Managers, Architects, (Performance) Analysts, System Administrators and IT-Pro's in general who are responsible for and/or interested in designing, implementing and maintaining virtualized Terminal Server and Virtual Desktop Infrastructures.

2.3 BETTER TOGETHER

“..The two largest and most focused competitors in the Dutch Virtualization and Application Delivery market space are working together on project: Virtual Reality Check...” PQR and Login Consultants started this joined-venture to share insights with the virtualization community with Project: Virtual Reality Check. There are several reasons for PQR and Login consultants to execute this project together:

- The Project leaders, Ruben Spruijt and Jeroen van de Kamp know each other for a long time from the virtualization community and share the same passion for these technologies.
- Project VRC is a huge undertaking, PQR and Login consultants individually do not have the resources, or time, to execute this project on their own. Thus is it logical to cooperate, share the workload and deliver the results together;
- Both organizations share the same technical vision, which is critically important in complicated projects like these.

2.4 VENDOR INVOLVEMENT

All major vendors whose products are covered by Project: Virtual Reality Check, such as VMware, Microsoft and Citrix have been approached in advance to create awareness of Project VRC and discuss the results.

2.5 CONTACT

All information about Virtual Reality Check can be found at www.virtualrealitycheck.net Contact details are:

PQR

Tel: +31 (0)30 6629729

E-mail: info@pqr.nl

www.pqr.com

Login Consultants

Tel: +31 (0)20 3420280

E-mail: info@loginconsultants.nl

www.loginconsultants.com

We try to provide accurate, clear, complete and usable information. We appreciate your feedback. If you have any comments, corrections, or suggestions for improvements of this document, we want to hear from you! Please send e-mail to Jeroen van de Kamp (j.kamp@loginconsultants.nl) or Ruben Spruijt (rsp@pqr.nl). Include the product name and version number, and the title of the document in your message.

**THIS DOCUMENT IS PROVIDED "AS IS"
WITHOUT WARRANTY OF ANY KIND
FOR REFERENCE PURPOSES ONLY**

COPYRIGHT PQR & LOGIN CONSULTANTS

IT IS NOT ALLOWED TO (PARTIALLY) PUBLISH OR DISTRIBUTE CONTENT WITHOUT APPROVAL

3. ABOUT THE AUTHORS

3.1 ABOUT LOGIN CONSULTANTS

Login Consultants (Login) is an international IT-service provider specialized in virtualization, migration, desktop-deployment and application-delivery. In this field, we are recognized as the experts in the technologies of Microsoft, Citrix, VMware, RES, AppSense en Symantec/Altiris.

Technical innovation, like virtualization, of end-user infrastructures can bring significant benefits in the areas of costs, flexibility, safety, stability, and license management. As a result of the rapid technological advances and the technical complexity of implementations, only the very best specialists in this field are able to fully realize the business advantages of these innovations.

Login has an experienced team of more than 100 consultants in The Netherlands, Belgium and Germany. Our consultants have been accredited by Microsoft (4 MVP's), Citrix (2 CTP's) and VMware (1 vExpert). Our specialists are well-regarded as speakers at (inter)national technical seminars. They are active in technical blogs and involved as experts in several IT-publications.

Login's passion for technology and innovation is also materialized in our successful suite of point solution virtualization software tools in use by thousands of organizations around the world.

The service portfolio of Login Consultants is grouped around our three core service areas: consulting, projects and support. With our services we support our customers with the deployment of traditional and hosted desktops, with application virtualization and server virtualization.

In our role of independent advisor we support our customers with their selection of the architecture and software products that will best address their specific needs. Next to this role our specialists help with the design, implementation, migration and management of innovative desktop- and application infrastructures. We are in this business for more than 7 years.

Login Consultants has over 100 customers in The Netherlands, Germany and Belgium. Especially large and midsized organizations benefit of our expertise. We have a very large customer base in healthcare, government, financial services, transport, industry and references in all other verticals.

3.2 ABOUT PQR

"It is easy to complicate simple matters" Very few people have the ability to simplify something that is complicated. Consider the rubber band created by the British inventor Stephen Perry in 1845, for example. Complex and yet straightforward at the same time. PQR stands for the same straightforwardness. But in a different field, namely ICT infrastructures, with the focus on:

- Server & Storage Solutions;
- Application and Desktop delivery;
- Virtualization.

"Simplicity in ICT", experience how PQR can make ICT manageable and predictable via solutions that are linked to one another, geared to the future, flexible, inventive and solid at the same time. Work together with a company that likes the result-oriented approach and with personnel who ensure that a solution simply works. ICT has never been that straightforward!

PQR delivers advanced infrastructures with a focus on Server & Storage and Application & Desktop Delivery solutions and the associated migration, consolidation and virtualization paths including network and security. PQR is a Cisco Partner, a Citrix Platinum Solution Advisor, a CommVault Value Added Reseller, and HP Enterprise Specialist Partner, an HP ProCurve Elite Partner, a threefold Microsoft Gold Partner, a NetApp Gold Reseller, an RES Platinum Partner, a VMware Premier Partner and a Websense Platinum Partner. PQR's customers are active in all sectors of society and a significant part of the sales is realized with non-profit organizations, the health care sector, education and local and national government.

PQR is headquartered in De Meern and counts meanwhile over 100 employees. In fiscal year 2007/2008 the company posted sales of € 80.3 million and a net after tax profit of € 5.8 million. PQR's clients are active in all sectors of society. A significant part of our sales is achieved by non-profit organizations, the health care industry, education and local and federal government. www.pqr.nl

3.3 TEAM MEMBERS

Ruben Spruijt, Technology Officer PQR

Ruben Spruijt, born in 1975, studied Computer science and started his career as a Systems Engineer at A-Tree Automatisering. He has been working as a Solutions Architect at PQR since 2002.

Focusing on Server & Storage, Virtualization and Application Delivery solutions, PQR implements and migrates advanced ICT-infrastructures and has achieved the highest certifications of its most important partners: HP Preferred Partner Gold, Microsoft Gold Certified Partner, Citrix Platinum Solution Advisor, VMware Premier and Consultancy Partner.

In his job, Ruben is primary focused on Application and Desktop Delivery, hardware and software Virtualization. He is a Citrix Certified Integration Architect (CCIA), Citrix Certified Enterprise Administrator (CCEA) as well as Microsoft Certified Systems Engineer (MCSE+S). Ruben has been awarded with the Microsoft Most Value Professional (MVP), Citrix Technology Professional (CTP), VMware vExpert and RES Software Value Professional (RSVP) title.

At various local and international conferences Ruben presents his vision and profound knowledge of 'Application- and Desktop Delivery' and Virtualization solutions. He is initiator of PQR's conceptual modes of 'Application and Desktop Delivery solutions' and 'Data and System Availability solutions' and originator of www.virtuall.eu, the solutions showcase of PQR. He has written several articles that have been published by professional magazines and informative websites. To contact Ruben directly send an email to rsp@pqr.nl.

Jeroen van de Kamp, CTO Login Consultants

As Chief Technology Officer, Jeroen van de Kamp is responsible for defining and executing the technical strategy for Login Consultants. From the start, Jeroen has played a critical role in the technical growth and accreditation Login has accumulated over the years. He has developed several core solutions which allow Login Consultants to easily differentiate in the infrastructure consulting market. The most important ones are Infrastructure 2.0; this is the unconventional strategy for IT services to establish the agile IT infrastructure foundation to support the constant changing business demands and Solution4 which is the best practice and automation methodology for enterprise Citrix environments in high density data centers. Jeroen is also responsible for several well-known publications like the Flex Profile Kit, TCT templates & "The black hole effect". Because of his contribution to the technical community Van de Kamp is recognized as a thought-leader in the application delivery industry and has become a residential speaker for seminars like BriForum, Citrix Solution Summit and many others. He is one of the 25 members worldwide who participate in the exclusive "Citrix Technology Professional" program.

Jeroen is still engaged with strategic key accounts for Login Consultants, defining and realizing an all encompassing strategy for the application, desktop and server delivery infrastructure. Previous to his position as CTO at Log*in Consultants Jeroen held positions as Infrastructure Architect at Login Consultants; IT Consultant at QFace ICT and IT specialist at ASG de Veer. To contact Jeroen send an email to j.kamp@loginconsultants.nl.

4. THE LOGIN VSI BENCHMARK

For Project VRC, the free Login Virtual Session Indexer (Login VSI 2.1) methodology was used. Login VSI is a benchmarking methodology which calculates index numbers based on the amount of simultaneous sessions that can be run on a single physical machine, running either bare metal or virtualized operating systems. To keep the results representative it is imperative that identical tests are run on different types of systems. Therefore Login VSI does not allow any customization of the load scripts.

Login Virtual Session Indexer is freeware and can be downloaded from: www.loginconsultants.com.

The philosophy behind VSI is different to conventional benchmarks. In general, most system benchmarks are steady state benchmarks. These benchmarks execute one or multiple processes, and the measured execution time is the outcome of the test. Simply put: the faster the execution time or the bigger the throughput, the faster system is according to the benchmark.

Login VSI is different. VSI is specifically not a steady state test. VSI, a solution for benchmarking TS or VDI workloads, loads the system with simulated user workloads. VSI focuses on how much users can run on the system before it saturates. VSI is similar to investigating the maximum amount of seats on the bus, airplane or lifeboat by trial and error. This is the “Virtual Session Index (VSI)”. With Virtual Desktop Infrastructure (VDI) and Terminal Services (TS) workloads this is very valid and useful information. This index simplifies comparisons and makes it possible to understand the true impact of configuration changes on hypervisor host or guest level.

It is also important to understand why specific VSI design choices have been made. The most important, and also debated, design choice is to execute the workload directly on the target system. The scripts simulating the workloads are performed as a compiled AutoIT script on every target system, and initiated at logon within the simulated user’s desktop session context.

From a benchmarking point of view this is not ideal: the AutoIT script has an unnatural overhead, and both sleeps and timings are executed on the platform itself which can reduce accuracy.

An alternative to the VSI method is to generate user actions client side through the remoting protocol. Remoting protocols like Microsoft RDP and Citrix ICA support this. However, such solutions are complicated to build and maintain. Far more important: these methods are never product and vendor independent. With every change on a protocol level or with every new remoting solution, VSI needs to be revised/expanded to support these changes. Even with a huge development budget, it is practically impossible to support every single remoting protocol. Some protocols simply do not have a method to script user actions client side.


Another solution is to execute and time the scripts within each session remotely through an external server, using a native remote scripting protocol. Although technically not impossible (e.g. Microsoft PowerShell 2.0 allows the remote execution of scripts), it is far more complicated to develop. Additionally, the overhead would be substantial since the window information and all timed events has to be sent and translated over the network.

For Login VSI the choice has been made to execute the scripts completely server side with AutoIT. This is the only practical and platform independent solution for a benchmark like Login VSI. Also the relative overhead and footprint AutoIT is small enough (1-5% range) for Login VSI’s purposes.

4.1 VSI OVERVIEW

VSI consists of 4 components:

- AD Domain controller for user accounts and standard policies
- A file share for central configuration and logging
- Launcher workstations (Master and Slaves) to initiate the sessions
- Target platform (VDI or SBC) where the user load script are installed and performed



4.2 VSI 2.1 MEDIUM WORKLOAD


The medium workload is the only workload available in VSI Express and also available in VSI PRO. The medium workload is the default workload in VSI.

- This workload emulated a medium knowledge working using Office, IE and PDF.
- Once a session has been started the medium workload will repeat every 12 minutes when the load is normal. When the system is getting close to maximum capacity (VSI_{max}), a single loop may take around 13 minutes.
- During each loop the responsetime is measured every 2 minutes.
- The medium workload opens up to 5 apps simultaneously.
- The typerate is 160 ms for each character.
- The medium workload in VSI 2.0 is approximately 35% more resource intensive than VSI 1.0.
- Approximately 2 minutes of idle time is included to simulate real-world users.

Each loop will open and use:

- Outlook 2007, browse 10 messages.
- Internet Explorer, one instance is left open (BBC.co.uk), one instance is browsed to Wired.com, Lonelyplanet.com and heavy flash app gettheglass.com.
- Word 2007, one instance to measure response time, one instance to review and edit document.
- Bullzip PDF Printer & Acrobat Reader, the word document is printed and reviewed to PDF.
- Excel 2007, a very large randomized sheet is opened.
- PowerPoint 2007, a presentation is reviewed and edited.
- 7-zip: using the command line version the output of the session is zipped.

This is an overview of the medium workload:



4.3 RANDOMIZATION

Since Beta 0.5 of VSI, randomization is introduced within the user load simulation. This is an important feature as optimizers on a memory or network level operate on the principle of de-duplicating and compressing repeated patterns. If every single session is doing exactly the same, de-duplication of memory and compression on a presentation protocol level becomes unrealistic. For a more realistic load testing and randomization is crucial. This prevents overly optimistic benchmark results caused by unrealistic optimization of the workload.

Building randomization into a benchmark needs special consideration. From a system resource load and execution timing perspective, randomization would harm the repeatability of the results. This will not make sense from a performance benchmarking perspective as this one of the first requirements of VSI.

It is therefore chosen only to randomize the data set for each session. As a result, the workload, including all applications, how and when they are executed is exactly the same for each session. Only the documents, presentations, mailboxes and excel sheets are randomized. All random paragraphs and pictures used in each document, presentation or e-mail are generated to have the same size and structure.

VSI has a pool of 100 randomly generated documents, presentations and pst files which differ no more than 5% in size.

Overview of randomization in VSI:

Item	Randomized Pool Amount	Refreshed at each loop start	File Size Range (KB)
Word Document	100	Yes	+/- 265 KB
PowerPoint	100	Yes	+/- 1195 KB
Outlook inbox (PST file)	100	Yes	1257 KB
Excel Sheet	1	Yes, each cell uses the "Random" function. Every time the sheet is opened completely different values are generated.	1325 KB
Internet Explorer	No Randomization	n/a	n/a

4.4 VSImax 2.1

With VSI 2.0, a new index *VSImax* was introduced which replaced the OPI of VSI 1.0.

VSI 2.0 updated *VSImax* to 2.1 which improved accuracy (e.g. *VSImax* is reached when CPU is around 100%)

To measure *VSImax* the response time of 5 operations are measured (instead of only two):

- Maximizing Microsoft Word** (also used for OPI)
 This operation will measure the responsiveness of the Operating System itself. Maximizing Word is not an operation on an application level, but for a large part a system operation.
- Starting the "File Open" dialogue** (also used for OPI)
 This operation is handled for small part by Word and a large part by the operating system. The file open dialogue uses generic subsystems and interface components of the OS. The OS provides the contents of this dialogue. When the disk I/O is extensive or even saturated, this will impact the file open dialogue considerably.
- Starting the "Search and Replace" dialogue** (new in *VSImax*)
 This operation is handled within the application completely; the presentation of the dialogue is almost instant. Serious bottlenecks on application level will impact the speed of this dialogue.
- Starting the "Print" dialogue** (new in *VSImax*)
 This operation is handled for small part by Word and a large part by the OS subsystems, as the print dialogue is provided by the OS. This dialogue loads the print-subsystem and the drivers of the selected printer. As a result, this dialogue is also dependant on disk performance.
- Starting "Notepad"** (new in *VSImax 2.1*)
 This operation is handled by the OS (loading and initiating notepad.exe) and by the Notepad.exe itself through execution. This operation seems instant from an end-user's point of view.

For the measurements of *VSI_{max}* no sleep, “winwaitdelay” or “sendkeydelay” actions are included in the measured events. For this, temporarily “winwaitdelay” (default 250ms waits between window operations) or “sendkeydelay” (160ms waits between each send keystroke) has been configured to zero when the response time is measured.

The measured operations with VSI do hit considerably different subsystems such as CPU (user and kernel), Memory, Disk, the OS in general, the application itself, print, GDI, etc. These operations are specifically short by nature. When such operations are consistently long: the system is saturated because of excessive queuing on any kind of resource. As a result, the average response times will then escalate. This effect is clearly visible to end-users. When such operations consistently consume multiple seconds the user will regard the system as slow and unresponsive.

The new response time measurement in VSI 1.1 and higher versions (*VSI_{max}*) excludes the timing of sleeps/idles/autoIT overhead completely. *VSI_{max}* measures: Maximizing MS Word, starting the “File Open”, “Find and Replace”, “Print” dialogues and starting “Notepad”, and nothing else. *VSI_{max}* is reached when 6 consecutive averages response times are higher than 1000ms, and 6 consecutive maximum response times are higher than 2000ms. There is no relationship with the 2000ms of the Optimal Performance Index used in VSI/VRC 1.0.

In comparison to OPI used in VSI 1.0 *VSI_{max}* is much more accurate. While OPI had a deviation up to 10%, the *VSI_{max}* deviation proves to be to around 5%.

OPI and *VSI_{max}* must be interpreted differently. OPI stands for the amount of users you could have on the system comfortably. As a result, OPI represents the maximum amount of VSI users which do not interfere with each other from a performance point of view. When executing the same workload, *VSI_{max}* typically represents a higher number of users. *VSI_{max}* represents the maximum amount of users the system can handle before serious performance degradation. This is one of the key differences between VSI 1.0 and VSI 1.1 and above.

5. THE VRC PLATFORM

Login Consultants and PQR have built the benchmark platform for Project VRC at PQR in a datacenter in Amsterdam, The Netherlands. Login VSI 2.1 was used to create transparent, reproducible and stable performance tests on Terminal Server and virtualized desktop workloads. To effectively demonstrate the scalability of the Hypervisor platforms the benchmark environment has been built-up with the latest hardware- and software technologies.

5.1 HARDWARE CONFIGURATION

The bare-metal, Citrix XenServer 5.5, Microsoft Windows Server 2008R, Hyper-V 2.0 and VMware vSphere 4.0 update 1 (build 208167) platforms are tested on the following server hardware.

Note: all tests in this document were performed on 64GB of memory. The memory controller will default back to single channel in this configuration, because the memory channels banks are not equally configured. Review section 7.7 “Triple Memory Channel” for more information.

Component	Details
Server Brand/Model	HPDL380G6
BIOS version	P62 07/24/2009
CPU	2 x Intel Quad core x5550@2.67GHz
CPU cache	1Mb L2, 8Mb L3
Memory	64GB; 1333MHZ
Disk	8 x 146Gb, 820.2Gb, dual port 10.000RPM Serial SCSI
RAID level	RAID-5 with online spare (75% Read / 25% Write)
RAID controller	HP Smart Array P400i, with 512Mb and Battery Backed Write Cache
RAID controller	Firmware 5.20
Integrated Lights Out (iLO) v2	Firmware v1.79
Network Interface	NetExtreme II, Gb



5.2 LAUNCHER CONFIGURATION

All the VSI launchers are installed and configured within Virtual Machines running VMware.

All the VSI launchers have been installed on Windows Server 2008 x86 Enterprise Edition SP1 with 2vCPU and 3GB memory. The Microsoft Remote Desktop Client is included in the OS, no special configuration settings are applied. The Citrix XenApp plug-in for hosted apps (ICA Client) version 11.0.0.5357 has been installed.

The screen resolution for the RDP/ICA connection to the target machines was set to:

- 1024x786 Resolution
- 16 Bit Color Depth
- Speed Screen accelerators disabled
- Client Drives are disabled
- Client Printing is disabled
- Clear Type is not configured

5.3 EXTERNAL CLOCK

All tests are completely externally clocked and calibrated unless specifically noted. For this a dedicated bare-metal SQL 2008 x64 server was used without any other services or applications. Before each test the database was reset and the server rebooted to guarantee untainted results.

Clock Server configuration (not virtualized):

- Intel Xeon 3.0 GHZ dual core
- 2GB memory
- GB Network
- Windows 2008 x64 SP2
- SQL Server 2008 x64 (10.0.2531)
- No other services or applications

5.4 TEST CONSTANTS

All tests were performed with the following settings, unless stated otherwise:

- vSphere: TPS is disabled by setting Mem.ShareScanGhz from 4 to 0.
- vSphere: vMMU enforced for Intel EPT-D.
- XenServer: XenApp optimization = ON.
- XenServer: /USEPMTIMER switch in boot.ini to prevent drift issues.
- Launchers: RDP client on Windows 2008x64 SP1.
- All tests are externally clocked and calibrated on SQL.
- All launchers, SQL server, Host and VM's are fully rebooted before each test.
- Windows 2003 x86 SP2 Enterprise edition was used for all 32-bit Terminal Server tests.
- Windows 2008 x64 SP2 was used for all 32-bit Terminal Server tests.

6. PROJECT VRC METHODOLOGY & BEST PRACTICES

Although the best-practices have been discussed in the first Project VRC whitepapers and presented on events worldwide after the release of the phase 1 whitepapers, it is still important to understand what Project VRC's methodology is and how the result should be interpreted.

6.1 vCPU's

A Terminal Server is similar to a highway: it is shared by its users. And just like highways require two lanes to be practical and not continuously congested, Terminal Server require at least two vCPU's. If a highway consists only of one single lane, every slowdown or accident will immediately affect all users. Having two lanes on a highway results not only in the doubling of the total capacity, the impact of slowdowns or accidents is strongly reduced, its users now have the possibility to overtake on the other lane.

This is not fundamentally different with Terminal Server workloads. Configuring only a single vCPU for each VM could be possibly more efficient (or faster). However, in the real world of Terminal Server workloads this is highly undesirable. Such workloads vary greatly in resource utilization and are typified by hundreds and even thousands of threads within a single VM. Consequently, with a dual vCPU setup is Terminal Server users have much better protection as a single vCPU would show congestion issues much faster. As a consequence, a minimum of 2 vCPU's per server is configured for all tests within Project VRC, specifically not choosing a single vCPU per VM, even when this has proven to be more efficient.

6.2 vCPU OVERCOMMIT & DEDICATED HARDWARE

Another important best practice, which is valid for all tested hypervisors, is not to overcommit the total amount of vCPU's in relationship to the available logical processors on the system. For example, on a system with eight logical processors, no more than eight vCPU's should be assigned in total to all VM's running on this host.

This is only important when the primary goals to maximize user density. Various tests in phase 1 of Project VRC have proven that overcommitting vCPU's negatively affects performance. This is not completely surprising since multiple VM's now have to share individual logical processors, which will create additional overhead.

As a result, it is also recommended to use dedicated server hardware for Terminal Server and Xenapp workloads, so it is easier to control VM configuration and assign vCPU's in relationship to the available processors.

6.3 TRANSPARENT PAGE SHARING

vSphere's ability to overcommit VM memory and memory de-duplication through transparent page sharing (TPS) is highly useful for the consolidation of many VM's on a single server. Nevertheless, one of the older Terminal Server best practices floating around the internet communities was to disable TPS. Project VRC phase 1 showed that disabling actually improved performance by 5%. This is understandable, since TPS is possible through a background process which is scanning memory, and this consumes a modest amount of CPU. However, the performance impact of TPS is only visible with full CPU loads. TPS has no performance impact under normal conditions.

When it is the primary objective to maximize the amount of users with TS workloads and there is enough physical memory available, it is recommended to disable TPS. As a result, Project VRC performs tests by default with TPS disabled, unless stated otherwise.

However, this VRC recommendation should not be understood as an overall recommendation to disable TPS. For instance, when maximizing the amount of VM's is the main goal (this is quite common, e.g. VDI and rather typical server consolidation efforts), TPS can be very helpful.

It is important to note that VMware does not recommend disabling TPS, their publications have shown TPS does not impact performance.

6.4 INTERPRETING PROJECT VRC RESULTS

Project VRC uses the product independent Login Consultants VSI 2.1 benchmark to review, compare and analyze desktop workloads on TS and VDI solutions. The primary purpose of VSImax is to allow sensible and easy to understand comparisons between different configurations.

The data found within Project VRC is therefore only representative for the VDI & TS workloads. Project VRC results cannot and should never be translated into any other workloads like SQL, IIS, Linux, Unix, Domain Controllers, Network, etc...

Also, the "VSImax" results (the maximum amount of VSI users), should never be directly interpreted as real-world results. The VSI workload has been made as realistically possible, but, it always remains a synthetic benchmark with a specific desktop workload.

Real world TS and VDI performance is completely dependent on the specific application set and how and when these applications are used. To include specific applications or customize the VSI 2.1 workload, VSI PRO must be used.

It is important to stress that no benchmark or load testing tools can 100% predict the real-world capacity of an environment, even when the specific applications are used. Real-world users are simply not 100% predictable in how and when they use applications.

6.5 VRC, VSI AND CLOCKS REVIEWED

Project VRC's whitepaper "VRC, VSI and Clocks Reviewed" is a review and reflection on previous Project VRC publications, the benchmark: "Login Virtual Session Indexer (VSI)" and Windows clock behavior within virtual machines. This discussion is fueled by the fact that results from the individual Project VRC whitepapers are set side-by-side to compare hypervisors. Project VRC has been in discussion with both vendors and community, and performed additional research in this context. Before Project VRC can publish new results in this whitepaper, it is important to address any questions, review the impact of this discussion and improve VSI where possible.

The major conclusions in the "VRC, VSI and Clocks Reviewed" whitepaper are:

- Clocks drifts within VM's are possible on every hypervisor. This has two implications:
 - The clock is responsible for the response time measurements: the reported results can be less accurate.
 - The clock is used with every timed event such as sleep/wait/delay/pause/etc.: this can have a impact on weight of the workload because the activities are stretched over time.
- The use of the /USEPMTIMER in the boot.ini is required for Windows 2003 guests running on XenServer. This will fix the consistent 10% clock drift/accuracy issues seen on XenServer.
- The measured drift/accuracy issues are much lower (a consistent 10-20ms per timed event) with VMware vSphere 4.0 and Microsoft Hyper-V 2.0.
- Login VSI 2.0 introduces a new indexing method called VSImax. The new method is more accurate because any form of scripted or build-in delay/sleep is completely excluded from the measured response time.
- Because of the mentioned improvements in the scripting method used in Login VSI 2.0, the reliability and robustness has been greatly enhanced: the VSI workload is now stable under extreme conditions.

- VSI is vendor and product independent. For this reason AutoIT remains the preferred and most practical method to generate the VSI workload.
- With VSI 2.0 it is now possible to measure the response time on an external Microsoft SQL server, and calibrate sleeps within the AutoIT script during runtime.
- External clocking and calibration does affect results by a relatively small margin (up to 6% was measured).
- The switch from Optimal Performance Index to VSImax has potentially a much bigger influence on the results. As a result, it is safe to assume that the ratio between VMware ESX 3.5 and Citrix XenServer 5.0 reported in VRC1.0 would be comparable to the VSI 2.0 results we found with vSphere 4.0 and XenServer 5.5 if VSImax was used.

It is worthwhile to use calibration when hypervisors are compared. However, when evaluating different configurations on a single hypervisor, the results are always relative, and calibration would not change conclusions. Consequently, these findings do not change the recommendations and best-practices discovered and published in VRC 1.0 and the 1.1 update.

From now on, Project VRC will only publish the specific VSImax outcome when the test is calibrated using an external clock. Any “uncalibrated” results will only be used to highlight the differences (in %) within the context of a single hypervisor.

6.6 VSIMAX 2.1

VSImax, the maximum session/user density designation, was first introduced when VSI 2.0 was released in the late summer of 2009. During the analysis of phase 2 it became clear VSImax could be tuned further to ensure it is a true representative of the maximum amount of user sessions.

For this reason all VSImax result in this whitepaper are based on VSImax 2.1, which is also released with Login VSI 2.1 Express and Pro. The improved VSImax 2.0 will have on average have an 18% lower number in comparison to the original VSImax (2.0).

6.7 TEST DESCRIPTIONS

In this whitepaper most tests descriptions displayed in the graphs have the following convention:

VMware vSphere 4.0_4_2_4_8_2003_x86_sp2

- VMware
- vSphere
- 4.0
- 4
- 2
- 4
- 8
- 2003
- x86
- sp2

- Special Settings/Configuration
- 32 bit or 64 bit Windows TS VM
- TS OS used in VM
- Swap file size in VM
- GB Memory configured in VM
- Amount of vCPU per VM
- Total Amount of VM
- HV Platform Version
- Hypervisor Platform

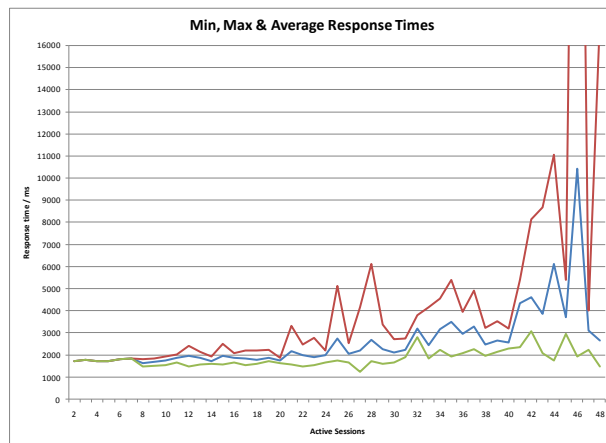
7. MISCELLANEOUS PHASE 2 RESULTS

This whitepaper focuses on virtualizing TS workloads. Nevertheless, additional, less critical findings were possible about:


1. Office 2007 Service Pack 1 & 2
2. IE7 vs IE8
3. HP Services
4. XenApp CPU Fairshare
5. XenServer XenApp Optimization
6. vSphere 2.0 vMMU impact

7.1 OFFICE 2007 SERVICE PACK 1 & 2

One of the key findings of the first Project VRC whitepapers is that Office 2007 Service Pack 1 has, because of a bug in Outlook's preview pane, a considerable impact on performance. We compared response time using a single VM:



Office 2007 Service Pack 1



Office 2007 Service Pack 2

It is clear from the response time tests above that Service Pack 2 fixed the high CPU utilization bug in Outlook's preview pane.

7.2 IE7 vs IE8

With every new version of Microsoft's Internet Explorer, the most commonly used browser in enterprises today, there is a potential impact on performance. In this test the impact on performance of migrating from IE7 to IE 8 is compared.


Tests were run on 8x 2003x86 TS VM's with 2 vCPU and 4GB mem, hosted on Hyper-V 2008 R2. It is important to note that with these tests the only bottleneck is CPU capacity, these tests do not provide insight on Internet Explorer's memory consumption since there is plenty memory available in the 8 VM's.



Which such a marginal difference between IE7 and IE8 average scores, it is safe to conclude that from a CPU point of view IE8 does not impact performance negatively.

7.3 HP SERVICES

In preparation of the Windows 2008 Bare-Metal tests the HP services as part of the HP Support Pack were installed as part of the default configuration. When reviewing the Bare-Metal results, the impact of these HP services was unknown. On the VM's running on all hypervisors, these services are not installed. As a result, the question was raised, how do these HP services impact the amount of users you can run on a system?

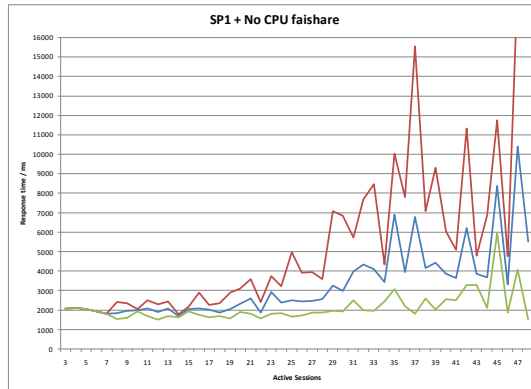


Tests were run with and without the default HP services started. The difference is minimal (services on: 149, services off: 148.5), therefore, it is safe to conclude that the default and unconfigured HP services have no specific impact on performance.

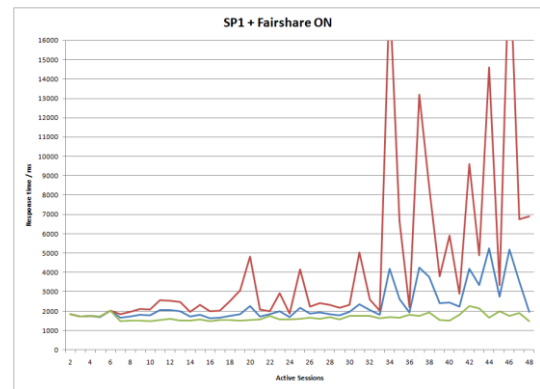
7.4 XENAPP CPU FAIRSHARE

In the original Project VRC whitepapers the results with CPU Fairshare enabled were surprisingly a little lower. The original explanation was that Citrix XenApp's CPU Fairshare cannot optimize processes which are behaving correctly. Although a logical conclusion, it is not satisfying the value of CPU Fairshare was not confirmed. In practice CPU Fairshare (and other CPU throttling and scheduling solutions) do have a proven positive impact on performance, especially with rogue applications which consume too much of the CPU time.

In hindsight, it was realized that Project VRC should use Office 2007 with service pack 1, to demonstrate the value of CPU management, as this causes a high CPU utilization bug when using Outlook's preview pane. A test to compare response times was performed on a single Windows 2003x86 TS VM with 2vCPU's and 4GB memory:



Office 2007 SP1 without CPU Fairshare



Office 2007 SP1 with CPU Fairshare

These graphs clearly show what Citrix Xenapp CPUFairshare does improve response times with the buggy Office 2007 SP1 installed. The average response times start rising after 21 sessions without CPU Fairshare. When enabled, the average response time is not rising until 33 sessions.

7.5 XENSERVR XENAPP OPTIMIZATION ON NEHALEM

The Citrix Xenserver feature to optimize performance for XenApp workloads (which also works for Terminal Server workloads) has proven effective on processors from the previous generation without AMD RVI or Intel EPT-D. Since Xenservers Xenapp optimization is basically a software variant of RVI/EPT-D, it would be interesting to see how Terminal Server workloads would perform on Intels Nehalem with and without XenApp optimization enabled.

Xenserver tests (8x 2003x86 TS VM's with 2 vCPU and 4GB mem) were performed with and without XenApp optimization enabled:



These tests show that enabling XenApp optimization does not improve or degrade performance. This is because XenApp optimization is not used when EPT-D is available.

7.6 VSPHERE 4.0 VMMU IMPACT

After the release of the first VRC whitepapers and the resulting discussions with the community and VMware's performance team, it became clear that the vMMU setting in ESX 3.5 on "Automatic" would not enable hardware vMMU (utilizing AMD's RVI) for Windows 2003 guests by default. Additional tests were executed, and when the hardware based vMMU was enforced for each VM, a considerable improvement (20-25%) was noticeable. As a result, the new results were published in the updated VMware whitepaper (v 1.1).

Although it is already clear that hardware based vMMU is very important for TS/XenApp workloads, it is interesting to see how hardware based vMMU (utilizing Intel's EPT-D) results compare to the software vMMU on Intel Nehalem running vSphere 4.0.

In comparison to the phase 1 test it is clear that Intel's EPT-D has an even bigger positive influence on Terminal Server workloads. On vSphere a 62% improvement over software vMMU is a rather substantial improvement.



These figures show the importance of AMD RVI and Intel EPT-D for Terminal Server workloads. It is therefore important to emphasize that hardware based vMMU is only supported on the latest generation of processors from AMD and Intel.

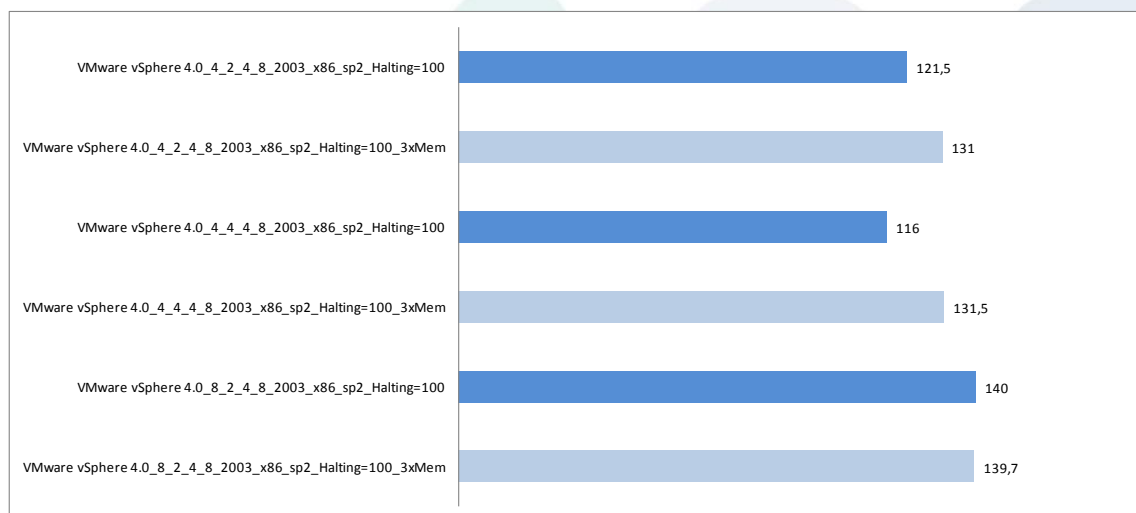
As a result, Terminal Server workloads performance on the still relatively young processors from the previous generation will be substantially lower. This detail is overlooked in practice quite often, as even recent server hardware do not always support hardware vMMU, in practice this is not always expected.

Consequently, when expectations for Terminal Server workload performance are not correct because hardware vMMU performance is anticipated, Terminal Server workload performance with software vMMU can be quite disappointing.

7.7 TRIPLE MEMORY CHANNELS

When we offered this whitepaper for review to our peers in the community and vendors, we were hinted by some that the 64GB memory configuration was not typical for Nehalem set-ups. For the best performance a triple memory channel setup is required to maximize memory throughput. As a result, the memory modules should be configured per three banks.

The 64GB configuration Project VRC used for all tests in this whitepaper is not a triple memory channel configuration. The memory controller on the server will then default to single channel, which reduces overall memory throughput. To review the impact of triple memory channel configuration opposed to the single memory channel used in all tests, additional tests were performed on vSphere 4.0.



In the illustration above, the dark blue bars represents the typical 64GB single channel configuration used for all tests in this whitepaper. The light blue bars represents the 48GB triple channel configuration.

As expected, configuring triple channel will increase performance approximately 10%. Interestingly, the difference is minimal in the 8VM + 2vCPU tests, the difference averages about 10% with both 4VM + 2vCPU and 4VM + 4vCPU tests.

8. NEXT GEN HYPERVISOR AND HARDWARE: IMPRESSIVE

This should not come as a big surprise, but the move to the latest hypervisor and the most recent Intel processor line does impact overall capacity of a single server. Already many results are published by reviewers around the world. But, the latest Intel Nehalem processors include one specific optimization that greatly benefits Terminal Server workloads: EPT-D.

“Extended page tables (EPT). When this feature is active, the ordinary IA-32 page tables (referenced by control register CR3) translate from linear addresses to guest-physical addresses. A separate set of page tables (the EPT tables) translate from guest-physical addresses to the host-physical addresses that are used to access memory. As a result, guest software can be allowed to modify its own IA-32 page tables and directly handle page faults. This allows a VMM to avoid the VM exits associated with page-table virtualization, which are a major source of virtualization overhead without EPT.”

In the first series of Project VRC whitepapers, tests were performed on the HPDL385G5 hardware with AMD (Barcelona) processors along with 8 logical processors and RVI. In the second phase tests were performed on HPDL380G6 hardware and the latest Intel (Nehalem) with 16 logical processors and EPT-D.

The following tests were not calibrated with an external VSI clock. For this reason no total numbers are published for these comparisons, only the relative performance improvement is shown in the graphs. For each test on the G6 Hyper-threading was enabled and 8 vm's with 2vCPU were configured to get the most out of the new hardware.

8.1.1 G5 and ESX 3.5 vs G6 and vSphere 4.0


When ESX 3.5 and vSphere 4.0 were tested, vMMU was enforced on both AMD and Intel.



The difference in total capacity is more than 90%, when the old and new generation of hardware and hypervisors are compared.


8.1.2 G5 and XenServer 5.0 vs G6 and Xenserver 5.5

XenApp optimization was enforced on both G5 AMD and Intel G6 tests.



The difference in total capacity is more than 95%, when the old and new generation of hardware and hypervisors are compared.

8.1.3 G5 and Hyper-V 1.0 vs G6 and Hyper-V 2.0



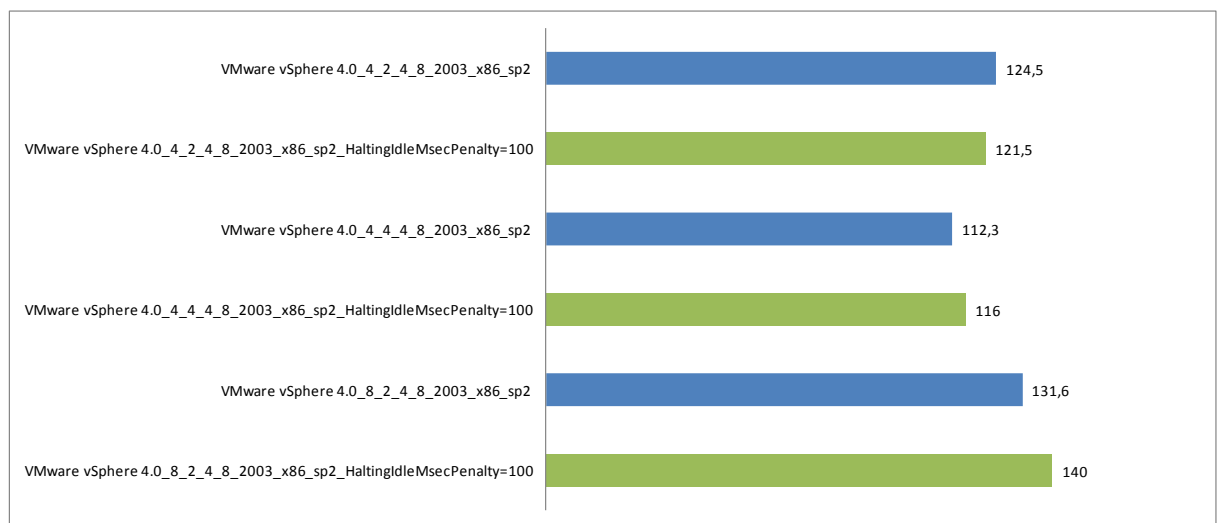
Looking at the Hyper-V results, the difference of 154% can be considered spectacular. But it is in many ways also logical, since Hyper-V 2.0 introduces support for hardware extended (or nested) page table virtualization by supporting Intel EPT-D and AMD RVI. Similar to vSphere, support for EPT-D and RVI does bring considerable performance value with Terminal Server workloads. When Hyper-V 1.0 would be supporting RVI, it is safe to conclude that the move to G6 increases total capacity around 100%, very similar to vSphere and Xenserver (Unpublished tests of the Hyper-V 2.0 beta on the G5 saw a 40% improvement over Hyper-V 1.0).

9. VIRTUALIZING WINDOWS 2003X86 TS WORKLOADS

The first whitepapers of project VRC already have proven that virtualizing TS workloads is highly recommended. Running Windows 2003 standard 32-bit Terminal Server workloads without the use of a hypervisor does not make sense anymore. 32-Bit Windows is simply not able to get the most out of current hardware because of its memory limitations. By virtualizing Windows 2003 Standard Edition 32-bit TS workloads, the amount of users running on a single server can easily be doubled.

9.1 vSPHERE HALTINGIDLEMSecPENALTY

When reviewing the vSphere results VMware suggested the configuration of HaltingIdleMsecPenalty setting to improve multi core and processor performance. With ESX 3.5 this was always set to 20, vSphere 4.0 now has a default value of 50.

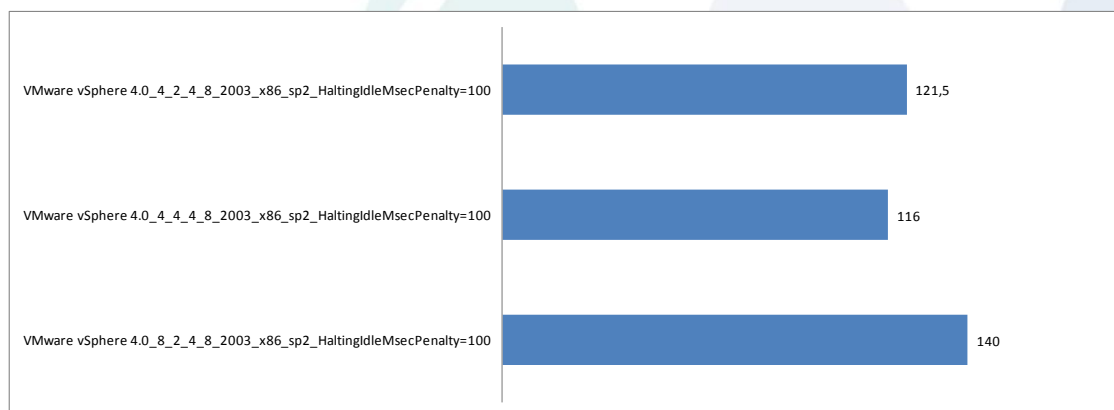


The first test with 4VM's and 2vCPU's not utilize Hyper-threading. Configuring the setting HaltingIdleMsecPenalty to 100 does not provide specific performance improvements here.

Although the differences are small, configuring HaltingIdleMsecPenalty to 100 on vSphere 4.0 does provide small performance improvements when Hyper-threading is utilized with all available logical processors. It would be fair to state that only the 8VM test show a tangible 5% improvement.

9.2 vSPHERE 2vCPU vs 4vCPU vs HT WITH VSI 2.1 & VSI 1.1

When reviewing VSI 2.1 tests 1 on the G6 platform with HT enabled slightly surprising conclusion can be made:



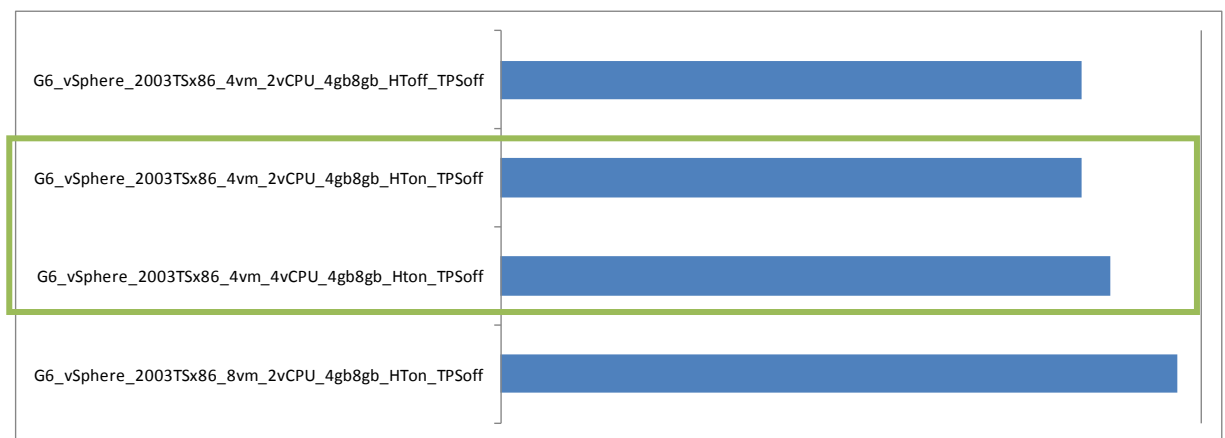
In the previous graph three different VMware vSphere 4.0 tests are shown with HaltingIdleMsecPenalty set to 100. All test types have exactly the same setup. In the first test 4 VM's are configured with 2vCPU, in the second 4 VM's with 4vCPU's are setup to utilize hyper-threading (2 logical processors per core), and in the third 8 VM's are set to use 2vCPU each, also utilizing the hyper-threading.

From these test, it can be concluded that Terminal Server workloads do not scale well when a 4vCPU per VM setup is combined with the use of Hyper-threading on vSphere 4.0. This combination does provide performance benefits over a 2vCPU for each of the 4 VM's.

However, the 8 VM with 2vCPU's setup do give a considerable performance boost over 4 VM's with 2vCPU.

It important to note that the VSI 2.1 workload is considerably more heavy on the CPU in comparison to VSI 1.x. This is caused by the inclusion of quite heavy flash app in VSI 2.x.


vSphere tests were also performed with VSI 1.1, which is much more CPU friendly and is not spiking CPU utilization like 2.1.



These tests show that switching from 2vCPU to 4vCPU, and thus utilizing hyper-threading with 4vCPU's per VM, does prove to be slightly beneficial and be dependent on the intensity of the workload.

9.3 VIRTUALIZING TS 2003x86 WORKLOADS NOT UTILIZING HYPER-THREADING


In the test below vSphere 4.0, Xenserver 5.5 and Hyper-V 2.0 are compared on the G6 platform, specifically not utilizing the Hyper-threading processors:



When not using hyper-threading, the optimal configuration is 4 VM's with 2vCPU's for all hypervisors. From the tests above it is clear that VMware has a small 5% margin over the other two hypervisors. Hyper-V and XenServer are very similar in performance.


9.4 VIRTUALIZING TS 2003x86 WORKLOADS UTILIZING HYPER-THREADING

In the test below vSphere 4.0, Xenserver 5.5 and Hyper-V 2.0 are compared on the G6 platform, specifically utilizing Hyper-threading logical processors:



When using the Hyper-threading processors, the optimal configuration is 8 VM's with 2vCPU's for all hypervisors. From the tests above it is clear that Hyper-V and XenServer have a 15% margin over vSphere utilizing Hyper-threading. Hyper-V and XenServer are again very similar in performance. It seems that VMware vSphere 4.0 does scale up with Hyper-threading processors, but Microsoft's and Citrix's hypervisors seem to make a little better use of the Hyper-threading processors.

In the graph below the average response times are displayed of the 8VM + 2vCPU 2003x86 TS test on vSphere (3x test averages), XenServer (3x test averages) and Hyper-V (3x test averages).



By comparing how the response time develop the following trends are clear:

- Up to 50 sessions all three platforms are identical in performance.
- From 50 to 100 sessions a light difference in avg. response times occur between vSphere and Hyper-v versus Xenserver. This difference is negligible in the real world.
- From 100 sessions and up Hyper-v and XenServer equalize in avg. response times, and vSphere response times increases more quickly with ever additional session.

Up to 100 sessions all hypervisors perform practically equal, only after 100 sessions the difference is more pronounced.

10. VIRTUALIZING TS 2008X64 WORKLOADS


Choosing 64-bit for Windows Terminal Servers is commonsense. The primary driver for choosing the x64 Windows architecture is the far better upscaling potential it has in comparison to Windows x86. Not being limited by the x86 memory restrictions takes away the key bottleneck for multiuser workloads. The only reason why x86 is still dominant today in the Terminal Server world is application and driver compatibility. This is expected to change in the upcoming years.

While virtualizing x86 TS workload does make a lot of sense, virtualizing x64 TS is a little more complicated and sometimes a paradox. The reality is that, in sharp contrast to x86 workloads, 64-bit Terminal Services has been build to scale up. However, if the primary bottleneck memory is taken away, the processor becomes the predominant bottleneck again. This is where in today's reality a new dilemma occurs in the hypervisors from VMware, Citrix and Microsoft.

While it is sensible to have multiple smaller machines with x86 Terminal Services, does this still make sense with x64? VMware's vSphere 4.0 and Citrix XenServer 5.5 do support a maximum of 8 vCPU's, Microsoft Hyper-V 2.0 only support 4 vCPU's. Now this seems like a lot, but with current hardware 16 logical processors is normal. To utilize all 16 logical processors a minimum of 2 VM's is required with vSphere and XenServer, and even 4 Vm's with Hyper-V.

10.1 4 VM's WITH 4vCPU'S VS BARE-METAL

First a 4 VM's with 4vCPU's configuration is compared. This is the only configuration which allows direct comparisons on all hypervisors since Hyper-V only supports a maximum of 4 vCPU's per VM.




From the tests in the graph above it is possible to conclude that Bare-Metal outperforms hypervisors with x64 Terminal Server workloads. This is to be expected as the x64 architecture has been created to scale up on large hardware setups.

Again Hyper-V and XenServer achieve strikingly similar results. VMware's vSphere 4.0 outcome is now approximately 21% lower than the XenServer and Citrix.

10.2 “BIG IRON” TESTS VS BARE-METAL

In the big iron test the VM’s are equipped with 8 vCPU’s. Aim of such a configuration is to try to minimize the amount of VM’s required to run x64 TS workloads. This is a reasonable preference, since every extra VM would require the purchase of an additional OS license and additional system tools in comparison to the single instance required for the x64 Terminal Server running natively on the server.




When only one VM is configured with 8 vCPU’s, and consequently not utilizing Hyper-Threading, XenServer and vSphere score strikingly similar. Both hypervisor score a *VSI_{max}* of around 100.

When two VM’s with 8 vCPU are configured to utilize Hyper-threading, XenServer shows approximately a 35% improvement over using only a single VM. It looks like vSphere is having difficulties utilizing Hyper-Threading, similar described seen in chapter “Virtualizing TS 2003x86 Workloads Utilizing Hyper-threading”.

11. OVERVIEW OF ALL TESTS

This is an overview of all average VSImax 2.1 results tests performed in phase 2 of Project VRC. Please note that most tests have been performed multiple times (mostly 3-4 times, some critical test were even executed 8 times) to improve the accuracy.





Login Consultant B.V.
De Entree 11-13
1101 BH Amsterdam
The Netherlands

Tel: +31 (0)20 3420280
Fax: +31 (0)20 6975721
E-mail: info@loginconsultants.nl
www.loginconsultants.com



Eenvoud in ICT

PQR B.V.
Rijnzathe 7
3454 PV De Meern
The Netherlands

Tel: +31 (0)30 6629729
Fax: +31 (0)30 6665905
E-mail: info@pqr.nl
www.PQR.com
www.VirtualLl.nl