


# Virtual Reality Check

## Platform Performance Index

---



### Whitepaper -

### VRC, VSI and Clocks Reviewed

Author(s) : Ruben Spruijt, Jeroen van de Kamp

Version: 1.0.3

Date: September 2009

© 2009 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
1.0	September 2009	Ruben Spruijt & Jeroen van de Kamp	
1.0.3	September 2009	Ruben Spruijt & Jeroen van de Kamp	Processed comments vendors, added executive summary

### REVIEWERS

Version	Date	Reviewers	Company
1.0	September 2009	Herco van Brug	PQR
1.0	September 2009	Dr. Bernhard Tritsch	Immidio

# CONTENT

- 1. Executive Summary ..... 3
- 2. Introduction Project VRC..... 4
  - 2.1 Project VRC objectives ..... 4
  - 2.2 Intended audience ..... 4
  - 2.3 Better together ..... 4
  - 2.4 Vendor Involvement..... 5
  - 2.5 Contact ..... 5
- 3. About the authors ..... 6
  - 3.1 About Login Consultants ..... 6
  - 3.2 About PQR..... 6
  - 3.3 Team Members ..... 7
- 4. VRC, VSI and CLOCKS..... 8
  - 4.1 Introduction ..... 8
  - 4.2 The Login VSI Methodology ..... 8
  - 4.3 VSI 1.0 Optimal Performance Index ..... 9
  - 4.4 “Stuck” Sessions ..... 10
  - 4.5 Response time measurement method and clock drift ..... 11
  - 4.6 Using SQL as reference clock..... 11
  - 4.7 Login VSI 2.0 ..... 13
  - 4.8 Reviewing Citrix XenServer and VMware ESX/vSphere ..... 14

## 1. EXECUTIVE SUMMARY

This whitepaper 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, it is important to address any questions, review the impact of this discussion and improve VSI where possible.

The major conclusions in this document 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 *VSI<sub>max</sub>*. 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 *VSI<sub>max</sub>* 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 *VSI<sub>max</sub>* 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.

Project VRC will from now on only publish the specific *VSI<sub>max</sub>* outcome when using an external clock. Any “uncalibrated” results will only be used to highlight the differences (in %) within the context of a single hypervisor.

## 2. INTRODUCTION PROJECT VRC

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

More and more people ask advice about the hardware virtualization solutions, particularly in a Terminal Server and Virtual Desktop Infrastructure context. PQR and Login Consultants have decided to compare the various hardware virtualization platforms in relationship to the end-user performance experience using the freely available benchmarking methodology of Login Consultants, Login VSI. Jointly they will deliver the outcome of the investigations and benchmarks to a broader audience, as part of the joint venture “Project: Virtual Reality Check (VRC)”.

Running Terminal Server workloads on virtual hardware is generally not recommended, but recent developments give grounds to a re-evaluation of current best practices. By benchmarking these and the Virtual Desktop workloads on physical servers and various virtualization solutions, Project VRC will give you valuable and most importantly unbiased experience and insights.

### 2.1 PROJECT VRC OBJECTIVES

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

- How does various Microsoft Windows Client OS's scale as a virtual desktop?
- How does a VDI infrastructure scale in comparison (virtualized) Terminal Server?
- 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 do the two usage scenarios compare, that is Microsoft Terminal Server [TS] only, versus TS plus XenApp?
- 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?

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

Project VRC publishes their findings on [www.virtualrealitycheck.net](http://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](http://www.virtualrealitycheck.net) Contact details are:

PQR

Tel: +31 (0)30 6629729

E-mail: [info@pqr.nl](mailto:info@pqr.nl)

[www.pqr.com](http://www.pqr.com)

Login Consultants

Tel: +31 (0)20 3420280

E-mail: [info@loginconsultants.nl](mailto:info@loginconsultants.nl)

[www.loginconsultants.com](http://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](mailto:j.kamp@loginconsultants.nl)) or Ruben Spruijt ([rsp@pqr.nl](mailto: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***

## 3. ABOUT THE AUTHORS

### 3.1 ABOUT LOGIN CONSULTANTS

Founded in 2002, Login Consultants has grown to become one of the leading IT infrastructure consulting firms in access, virtualization and deployment solutions. Currently Login employs up to 100 skilled and motivated team members in the US, Germany, Belgium and the Netherlands.

Login Consultants started in the Netherlands as a consulting firm with 100% focus on virtualization solutions. Although the Dutch IT market lost its pace at the time, Login Consultants was able to grow rapidly. This was in large part because of many organizations launching strategic projects to lower IT costs through centralization and consolidation. This proved more challenging in practice than originally expected. Login Consultants contributed its vast virtualization expertise, skills and knowledge to help customers achieve true success.

Over the years Login Consultants has become a well-respected partner to many IT organizations and vendors. This has driven growth and enabled expanding operations to Belgium (2004), Germany (2005) and the United States (2006), allowing Login to extend its services and success to international oriented organizations.

In the application-centric infrastructure, there are many strategic technologies that help IT departments to introduce agility in their internal service delivery. Login Consultants has embraced these technologies for access, virtualization and deployment, thus fulfilling the promise of a dynamic infrastructure.

Many technology vendors and strategic partners subscribe to the success Login Consultants achieves for its customers. Citrix, Microsoft & VMware all have accredited Login as a consulting partner.

With no specific focus on hardware or software selling, Login is an independent solution provider for many IT organizations who trust Login for objective advice and quality consulting. The ability to help define and realize a comprehensive strategy for application delivery, makes Login unique and of great value to its customers.

### 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.com](http://www.pqr.com)

### 3.3 TEAM MEMBERS

#### **Ruben Spruijt**, Solutions Architect 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](http://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](mailto: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](mailto:j.kamp@loginconsultants.nl).

## 4. VRC, VSI AND CLOCKS

### 4.1 INTRODUCTION

Although we have already many new results we would like to share with the community, the Project VRC team decided to first release this intermediate whitepaper. The reason is simple: in the background there have been extensive discussions and research with both vendors and community about VRC 1.0's accuracy and the methodology used: the Login Consultants Virtual Session Indexer (VSI). This discussion centers mostly about a phenomenon called "clock-drift" within virtualized guests and the response time accuracy.

For Project VRC to move forward and publish new results, it is important to fully understand the impact of this conversation and improve the methods used where possible. In this document the following topics are addressed:

1. The Login VSI Methodology;
2. VSI 1.0 Optimal Performance Index;
3. "Stuck" Sessions;
4. Response time measurement method and clock drift;
5. Using Microsoft SQL Server 2005 as a reference clock;
6. Login VSI 2.0;
7. Reviewing Citrix XenServer and VMware vSphere 4.0;

### 4.2 THE LOGIN VSI METHODOLOGY

Before we continue, it is important to understand what the methodology is behind Login VSI. 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. The impact of this overhead and accuracy are explained later in this document.

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.

### 4.3 VSI 1.0 OPTIMAL PERFORMANCE INDEX

The first Project VRC whitepapers, released in January 2009, were based on Login Virtual Session Index (VSI) 1.0. To compare configurations Project VRC used the "Optimal Performance Index" (OPI) provided by VSI 1.0.

The response time of the "OPI" included the build-in AutoIT overhead. For example, a response time of 1750 ms contained both the build-in delays for window manipulation (about 1500ms) and the actual response time (about 250ms) of the "Maximize Word" and "Open File Dialogue" commands.

The inclusion of the build-in AutoIT sleeps, "winwaitdelay" or "sendkeydelay" in the reported response time does not add value to the results and makes them less accurate. Also, only measuring two operations is a somewhat limited representation of the response time.

#### 4.3.1 VSImax

With VSI 1.1 (not publicly released), a new index VSImax is introduced which replaces the OPI. 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 "Calculator"** (new in VSImax)

This operation is handled by the OS (loading and initiating calc.exe) and by the Calc.exe itself through execution. This operation seems instant from an end-user's point of view.

For the measurements of *VSI<sub>max</sub>* 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<sub>max</sub>*) excludes the timing of sleeps/idles/autoIT overhead completely. *VSI<sub>max</sub>* measures: Maximizing MS Word, starting the "File Open", "Find and Replace", "Print" dialogues and starting "Calculator", and nothing else. *VSI<sub>max</sub>* is reached when 3 consecutive averages 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<sub>max</sub>* is much more accurate. While OPI had a deviation up to 10%, the *VSI<sub>max</sub>* deviation proves to be to around 5%.

OPI and *VSI<sub>max</sub>* 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<sub>max</sub>* typically represents a higher number of users. *VSI<sub>max</sub>* 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.

*VSI<sub>max</sub>* proves to be "Event Horizon" with every test. Beyond this point the system is clearly saturated and end-user performance is degraded badly. By experiment we shifted *VSI<sub>max</sub>* calculation 200ms down and 500 ms up (in the range 1800 and 2500ms) on every result we have: this does not change the *VSI<sub>max</sub>* results significantly. Over 30 different benchmark configurations were compared. Interestingly there are two trends very consistent with *VSI<sub>max</sub>*. The minimum response time clearly starts rising, and the maximum response time is almost doubled after *VSI<sub>max</sub>*.

*VSI<sub>max</sub>* is much more accurate and the deviation with retests is much lower in comparison to OPI. To achieve such consistency with an unconventional benchmark like VSI is crucial for the validity of follow-up VRC results.

#### 4.4 "STUCK" SESSIONS

With the first version of VSI approximately 5% of sessions hang during execution. In the 1.0 whitepaper this was called a "stuck session". This was generally caused by AutoIT missing or being out of sync of Windows events, this effect seemed magnified under high system load. Although these stuck sessions were calculated in the reported OPI in VRC 1.0, stuck sessions are undesirable for the quality of the test.

Further investigation in 2009-Q2 learned the following best practices that help to almost completely eliminate stuck sessions, even beyond *VSI<sub>max</sub>*:

- Under certain circumstances, the "Minimize Window" command within resulted in Z-ordering issues (order of appearance) with active windows. As a result, a window can be minimized and

active at the same time. Normally this does not happen, active Windows are always visible (not minimized). AutoIT could not handle this, and kept on waiting for the Window indefinitely. This is only happens when the minimized window was the last active application before another application is closed. The fix for this is to launch a new window, which will reset the z-order.

- A major reduction in stuck session was possible if the VSI always checks first if a specific window is not already active, before the activation command is given. Activating (focusing) windows proved to be consistently more reliable this way.
- And lastly, another large reliability improvement was gained by consistently waiting 2 seconds before “typing” in to active window. 100% of the scripted activities within VSI are performed through keyboard shortcuts. These are performed with scripted keystrokes. Under load, the system Graphical Device Interface (GDI) still can be rendering sub-windows when the main window it is already considered active by AutoIT (an application often consists of multiple windows). When the “typing” is performed too early, chances are high AutoIT runs out of sync of what is truly happening on a window level. The 2 second wait after activating windows was responsible for 80% reduction in stuck sessions.

Since these improvements, over 60 tests on different platforms have been performed. It is now safe to conclude that Login VSI 1.1 and 2.0 are not affected by seemingly random stuck sessions anymore.

## 4.5 RESPONSE TIME MEASUREMENT METHOD AND CLOCK DRIFT

The clock is important in two specific ways within the VSI Benchmark:

- The clock is responsible for the response time measurements;
- The clock is used with every timed event such as sleep/wait/delay/pause/etc;

Login VSI 1.0 uses the clock from the guest OS to determine the response time and to execute every sleep event. As a result, the response time is measured on the target (guest) platform itself. Normally, in non virtualized environments, this poses no great problems, as the clock of the OS is precise enough.

However, there are a couple of concerns which are magnified in virtualized environments. First: this method is similar to having the individual athletes measuring their laps with their own stopwatch, and not by an independent referee. Secondly, when virtualizing systems, the clock has a tendency to drift, and so the results are less accurate. Typically when the Windows clock drifts on a virtualized OS, time slows down. This will impact the VSI load on the system.

With clock drifts, the workload execution can be stretched over time, lengthening the total length of the loop for each user. This affects the overall load on the system in two ways: the load is possibly lower because the script is executed slower or the workload is higher because resource intensive applications (for instance a flash enabled webpage, or apps with animations / video’s) are running longer.

When reviewing Citrix XenServer results running VSI 1.0 on Intel Xeon 5500 ‘Nehalem’ platform, VSI returned even negative numbers. These findings did fuel the clock drift discussion considerably. The data of the VSI 1.0/XenServer/Nehalem made clear a external clock was required to validate the results.

Consequently, a rudimentary method to reference check timings with an external clock was build into VSI 1.1. Additional tests were performed with an external SQL server running as reference clock.

## 4.6 USING SQL AS REFERENCE CLOCK

Microsoft SQL Server 2005 was chosen because it is a highly scalable data solution and SQL timestamps are known be 3ms precise. This SQL server was running on HP Proliant enterprise class, bare metal,

server hardware without other specific roles. Network was 1Gbps full-duplex, latency was less than 1ms. Consequently there were absolutely no bottlenecks on the external “clock” infrastructure.

VMware vSphere 4.0, Citrix XenServer 5.5 and Microsoft Hyper-V 2.0 (RC1) were evaluated running Windows 2003 Terminal Services. The AutoIT “timediff” results were compared with the time difference measured on the SQL server. Both sleep events (in the range of 50ms to 2000ms) and VSI’s response time: maximizing MS Word, starting the “File Open”, “Find and Replace”, “Print” dialogues and starting “Calculator”.

- On Hyper-V AutoIT’s “timediff” function reported about 10ms average quicker per timed event than the reference external SQL clock. Example: AutoIT reports 200ms, in reality this is 210ms. This difference between what AutoIT reported and the reference clock was very consistent with both low and high loads on the system, also beyond *VSI<sub>max</sub>*.
- On vSphere 4.0 the AutoIT “timediff” function reported about 20ms average quicker per timed. Example: AutoIT reports 100ms, in reality this is 120ms. This difference between what AutoIT reported and the reference clock was very consistent with both low and high loads on the system, also beyond *VSI<sub>max</sub>*.
- On XenServer the AutoIT “timediff” function reported about 10% average quicker per timed event. Example: 100ms reported by AutoIT is in reality 110ms, and 500ms reported is in reality 550ms. When XenServer becomes saturated, these differences increase considerably. Extreme drifts were visible beyond *VSI<sub>max</sub>* when the system is heavily saturated.

There is more to be considered. By experiment, execute the AutoIT the script below multiple times on a bare metal machine with speed step or CPU throttling disabled:

```
MsgBox(0,"Start","This test will now start a single sleep of 100ms")
$ClockTimerInit=timerinit()
Sleep(100)
$TimeStamp=round(TimerDiff($ClockTimerInit), 0)
MsgBox(0, "Result", "This one sleep measured "&$TimeStamp)
MsgBox(0, "New test", "We will now test 5 sleeps in a row")
$ClockTimerInit=timerinit()
Sleep(100)
Sleep(100)
Sleep(100)
Sleep(100)
Sleep(100)
$TimeStamp2=round(TimerDiff($ClockTimerInit), 0)
MsgBox(0, "Result", "These sleeps totalled "&$TimeStamp2&" averaging "&Round($TimeStamp2/5,0))
```

Interestingly, the sleep function in AutoIT is not always accurate. There seems to be a consistent 5-10 ms deviation per timed event when multiple consecutive sleeps are executed. Consequently the differences measured with the SQL reference clock (described three paragraphs ago) are 5-10ms lower when you consider this native AutoIT deviation.

Looking at the numbers measured on Citrix XenServer 5.5 running on Intel Xeon 5500 ‘Nehalem’ using an external clock, two specific considerations were needed. The accuracy of AutoIT “timediff” function we used to measure response time and the possible impact of these drifts on sleep execution (slower) and ultimately the VSI workload.

A fix was needed for the XenServer drift. Because it is technically possible to review timings after the test, the impact would be clarified.

The XenServer team was asked: “Why is this drift happening and how can we fix it?”. At first the XenServer team was surprised because any drifting issues should have been fixed specifically in the XenServer 5.5 release. One of their suggestions was to configure /USEPMTIMER in the boot.ini.

An article on MS technet explains the /USEPMTIMER switch:

<http://blogs.technet.com/perfguru/archive/2008/02/18/explanation-for-the-usepmtimer-switch-in-the-boot-ini.aspx>

According to this article the PM timer should be enabled by default on Windows Server 2003 SP2. However, sometimes this is not correctly detected. A manual configuration of this switch is then required in the boot.ini. It is clear also XenServer is affected.

*“Unless you are running SP2 for W2K3 (where we use the PM timer for MP APIC ACPI HALs by default now), you must force the use of the PM timer on AMD machines via the /USEPMTIMER switch. The decision to use the PM timer versus the TSC is actually a check done on boot to query the BIOS, and determine whether or not it will properly support the PM timer functions - this check is not 100% accurate on AMD chipsets, and that is why forcing /USEPMTIMER is required.”*

After setting the /USEPMTIMER switch the timings of AutoIT are much more accurate.

Normally the loop time for Login VSI 1.1 is around 1040 seconds:

- Before /USEPMTIMER AutoIT reported on Xenserver a VSI 1.1 looptime of 964;
- After /USEPMTIMER we got AutoIT reported on Xenserver a VSI 1.1 looptime of 1041.

This also explains the negative VSI 1.0 numbers seen on the VSI1.0/XenServer/Nehalem combo described earlier. The Optimal Performance Index (OPI) response time was calculated by measuring the total execution time of the Word operations. The VSI 1.0 timer script included many sleeps to make this method reliable: approximately 14 seconds of additional sleeps were included. For the calculation of the OPI response time, these sleeps were deducted. When AutoIT reported 15500ms execution time, 14000ms was deducted to retrieve the response time within VSI 1.0. However, on the XenServer/Nehalem combo this timer was clearly broken by not using the /USEPMTIMER switch. As a result, AutoIT reported 10% less: about 14000ms. Then VSI specifically deducts the 14000ms worth of sleeps. 14000ms (reported) -14000ms (Sleep) would result in a response time of “0”.

It is now possible to conclude that there is a deviation on all platforms and even within AutoIT. On the VSI 1.0/ XenServer/Intel Nehalem without the /USEPMTIMER switch the OPI results are clearly skewed. The switch to *VSI<sub>max</sub>* improves the response time accuracy considerably in comparison to OPI.

However, this data does not explain what the true impact is on the *VSI<sub>max</sub>*. To fully understand this discussion we need to perform tests with externally clocked response time and sleep execution.

## 4.7 LOGIN VSI 2.0

July 2009 the development of Login VSI 2.0 was started.

Login VSI 1.1 (not publicly released) had the same workload as VSI 1.0, the only major difference in Login VSI 1.1 was the introduction of *VSI<sub>max</sub>* which completely replaced OPI.

The big change with Login VSI 2.0 is the completely redesigned workload(s). The PRO version of VSI also includes a light/heavy and multimedia workload. The new medium workload in VSI 2.0 is considerably more CPU intensive in comparison the VSI 1.x workloads. Typically, with the VSI 2.0 medium workload will result in about 35% lower *VSI<sub>max</sub>* compared to *VSI<sub>max</sub>* with VSI 1.1. This choice has been made because evaluation of VSI 1.x workloads proved to be too subtle on the system. The biggest change in

the VSI 2.0 workload is the inclusion of a very rich flash based website: [www.gettheglass.com](http://www.gettheglass.com). This greatly varies the impact of the new 2.0 workload in comparison the relatively steady 1.x workload.

The clock drift and accuracy discussion did require reviewing Login VSI's clocking methods. Preferably VSI should be able to address this discussion without the requirements to reference check the results afterwards. Although this is a method to validate results, it would also be cumbersome and would leave room for discussion. Such a reviewing method will not make clear what the impact of deviation is.

Within VSI 2.0 it is now also possible to measure the response time for *VSI<sub>max</sub>* using an external SQL server as clock. Although this has less impact on the *VSI<sub>max</sub>*, it is relevant to measure the response time externally when hypervisors are compared. It is relatively easy to implement SQL as an external source for response time measurements.

However, it is too complicated to outsource every single timed event (sleeps, winwaitdelay, sendkeydelay) to an external time source like SQL for each session: there over 1000 events per loop (sendkeydelays, winwaitdelays and sleeps). Also, we cannot use the remoting protocol to perform the workload without losing practicality and platform independency.

egg-shaped lad stood in stunted silence, slowly growing gay, who took himself too serious night school rhythmically. The girl had a secret admirer sometimes, but not always. The wack cle gang lounged around. The rotating smart aleck flipped burgers in the summer. The dorky ie tennis instructor picked raisins out of a bowl of cereal and then studied it.



With Login VSI 2.0 it is therefore now possible to calibrate clock drift every 2 minutes for every individual session. This is done by continuously (every 2 minutes) sampling AutoIT's sendkeydelays, winwaitdelays and sleeps with the external SQL clock, and then (re-)calculate the drift factor. This drift factor is then applied to subsequent sendkeydelays, winwaitdelays and sleeps. Since drifting is not completely random, this method will correct any trended deviations. Through this method, it is possible to safeguard the VSI workload execution rate and the impact it should have on the system.

## 4.8 REVIEWING CITRIX XENSERVER AND VMWARE ESX/VSPHERE

The ultimate question is of course, what is the impact of all these considerations?

For this reason both Citrix XenServer 5.5 and VMware vSphere 4.0 are tested with VSI 2.0. (Hyper-V results will be added in the near future).

All tests were performed with:

- HP Proliant DL380G6 with 2x Intel Xeon 5500's, 8 x 146GB HD RAID 5, 64GB memory; BBWC module;
- 4VM's running 2 vCPU's and 4GB of memory assigned;
- Microsoft Windows Server 2003 TS x86 + 8GB page file and Microsoft Office 2007 SP2 installed.
- XenServer: XenApp optimization = ON
- vSphere: vMMU = Enforced for Intel EPT

Below is the extraction of logging with calibration with the use of /USEPMTIMER for XenServer. These log entries show the drift factor (multiplier) of measured by continuously sampling the average of four 1000ms sleeps:

```

47: 30-08-2009;00:48:28;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.00075;Finished
85: 30-08-2009;00:50:49;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.0015;Finished
119: 30-08-2009;00:53:24;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 0.999;Finished
142: 30-08-2009;00:56:03;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 0.99575;Finished
174: 30-08-2009;00:58:31;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.00075;Finished
221: 30-08-2009;00:59:34;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.0025;Finished
242: 30-08-2009;01:00:48;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.0015;Finished
279: 30-08-2009;01:02:59;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 0.999;Finished
313: 30-08-2009;01:05:30;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1;Finished
336: 30-08-2009;01:08:05;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 0.99825;Finished
368: 30-08-2009;01:10:32;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 0.999;Finished
415: 30-08-2009;01:11:35;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 0.999;Finished
436: 30-08-2009;01:12:50;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.00075;Finished
473: 30-08-2009;01:15:01;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.00075;Finished
507: 30-08-2009;01:17:33;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.0025;Finished
530: 30-08-2009;01:20:11;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 0.9965;Finished
562: 30-08-2009;01:22:39;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1;Finished
609: 30-08-2009;01:23:43;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 0.9975;Finished
630: 30-08-2009;01:24:58;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1;Finished
667: 30-08-2009;01:27:07;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.0015;Finished
701: 30-08-2009;01:29:40;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.0015;Finished
724: 30-08-2009;01:32:20;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.0015;Finished
756: 30-08-2009;01:34:47;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 0.9975;Finished
803: 30-08-2009;01:35:53;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.0165;Finished
824: 30-08-2009;01:37:08;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.00075;Finished
861: 30-08-2009;01:39:23;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 0.9965;Finished
895: 30-08-2009;01:42:04;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.0025;Finished
918: 30-08-2009;01:44:52;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.024;Finished
950: 30-08-2009;01:47:20;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1;Finished
997: 30-08-2009;01:48:28;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.004;Finished
1018: 30-08-2009;01:49:47;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.005;Finished
1055: 30-08-2009;01:52:07;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1;Finished
1089: 30-08-2009;01:55:07;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.005;Finished
1112: 30-08-2009;01:58:03;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.0025;Finished


From here VSImax is reached: 145 sessions.

1144: 30-08-2009;02:00:35;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.00075;Finished
1191: 30-08-2009;02:01:45;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 0.999;Finished
1212: 30-08-2009;02:03:14;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 0.99825;Finished
1249: 30-08-2009;02:05:42;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.0025;Finished
1283: 30-08-2009;02:09:06;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.0125;Finished
1306: 30-08-2009;02:12:42;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1;Finished
1338: 30-08-2009;02:15:35;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.00825;Finished
1385: 30-08-2009;02:17:25;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.0075;Finished
1406: 30-08-2009;02:19:31;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1;Finished
1443: 30-08-2009;02:23:30;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 0.99825;Finished
1477: 30-08-2009;02:27:15;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 1.0015;Finished
1500: 30-08-2009;02:30:25;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 0.999;Finished
1532: 30-08-2009;02:33:00;Login_VSI1;TS60;VSI Timer Calibration;New Drift Factor is 0.99825;Finished

```

From the logs above it is clear that the /USEPMTIMER fixes clock drift issues. The average drift is consistently lower than 1%.

These are the *VSI<sub>max</sub>* results measured with VSI 2.0 PRO on XenServer 5.5 and vSphere 4.0.



From these results the following conclusions are possible:

- For XenServer 5.5 + VSI 2.0 measured
  - Without calibration and not using the /USEPMTIMER fix results in *VSI<sub>max</sub>*: 150
  - With calibration and /USEPMTIMER XenServer results in *VSI<sub>max</sub>*: 142
- For vSphere 4.0 + VSI 2.0 measured
  - Without calibration resulted in *VSI<sub>max</sub>*: 150
  - With calibration resulted in *VSI<sub>max</sub>*: 147
- Calibrating and externally clocking the VSI workload does have a bigger influence on *VSI<sub>max</sub>* with XenServer 5.5 (-5.5%) than it has on vSphere 4.0 (-2%).

The differences between calibrated (using an external clock) and non calibrated tests (using the target's clock) are relatively small. However, the switch from OPI to *VSI<sub>max</sub>* has potentially a much bigger influence on the results. As a result, it is safe to assume that the ratio between ESX 3.5 and 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 *VSI<sub>max</sub>* was used. Still, all recommendations and best-practices discovered in VRC 1.0 remain unchanged.



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](mailto:info@loginconsultants.nl)



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](mailto:info@pqr.nl)

[www.PQR.com](http://www.PQR.com)