



Workload Considerations for Virtual Desktop Reference Architectures

INFORMATION GUIDE

Table of Contents

Introduction	3
Description	3
Architecture	4
RAWC Workload.....	5
Applications	5
Active Directory Groups	5
RAWC GUI	6
Configuration Tab	6
Test Variables and Considerations.....	6
Application Launch Order	6
Application Launch Delay	7
Total Number of Virtual Machines Under Test	7
Test Mode	8
Logoff Users	8
Workload Variables and Considerations.....	9
Number of emails +1.....	9
Typing Speed in Seconds	9
Number of Random Words	9
PDF Pages to Browse.....	9
Results.....	9
Conclusion.....	9
About the Author.....	10
Acknowledgements.....	10
White Papers.....	10

Introduction

As interest in virtual desktop infrastructure continues to accelerate, so do the needs for improved assessment, planning, and deployment methodologies and tools. The historical methodology for determining deployment parameters, (that is the impact of VDI on network, storage, and compute resources) typically failed to truly assess the impact of the workload that would eventually end up being generated by the desktop virtual machine once it went into production. The reason for this gap was simple: There were few, if any, workloads that accurately represented what a user actually did during their workday, that is, the applications they used, the frequency of use, and the “intensity” of use. So professional services organizations typically estimated network and storage impacts based on their experience, and then architected or modified an environment based on those estimates. The problem, of course, is that such qualitative assessments can underestimate the impact of a deployment and result in storage contention or network bottlenecks, or overestimate the impact and result in excess costs.

The purpose of this paper is to introduce a desktop workload tool that generates a realistic, adjustable workload with various applications in the desktop virtual machine. The results gathered (regarding CPU usage, memory utilization, storage, and network) can then be analyzed to identify appropriateness / readiness of a given environment to run virtual desktops. This tool, called the Desktop Reference Architecture Workload Code (RAWC), has been used in numerous studies to simulate application workloads for various user types. It can be configured to simulate light, medium, or heavy user characteristics, including the types of applications used in a typical Windows desktop environment.

RAWC can be used to evaluate server and storage performance, validate configurations, and perform scalability studies and proof of concepts.

Description

The RAWC workload runs on a Windows 7 or XP guest operating system and is executed on each desktop virtual machine on one or more ESX hosts. The RAWC workload has a set of functions that performs operations on common desktop applications including Microsoft Office, Adobe Reader, McAfee Virus Scan, Windows Media Player, Java, and 7-Zip. The applications are called randomly and perform operations that mimic those of a typical desktop user, including open, save, close, minimize and maximize windows, view an html page, insert text, insert random words and numbers, conduct a slideshow, view a video, run a virus scan, send and receive email, and compress files.

The RAWC workload uses a configuration file that is created via the RAWC GUI and writes application open/close times and any errors to log files in a shared network folder. Various test variables can be configured via the RAWC GUI, including a start delay for creating ‘boot storms,’ and density (delay between application operations) resulting in applications under test to be run together faster, number of emails created and sent, and typing speed.

The combination of applications and test variables that are configured can either increase or decrease the workload on virtual machines and VMware ESX server(s). The goal of doing this is to evaluate how many virtual machines per solution can be run.

In addition to running native applications, RAWC can also run VMware® ThinApp™ applications hosted on either a local or remote site.

RAWC supports the use of Active Directory Groups. Large-scale deployment evaluations will find this feature extremely valuable. Based on the virtual machine’s Active Directory Group membership, a workload can be configured specifically for that Group via the RAWC GUI. This allows virtual machines to run varying workloads to simulate a more realistic large-scale work environment.

RAWC was recently updated to support a number of new features including duration or time-based testing, running applications in static order for repeatability, and the ability to log the user out after a test is completed.

Architecture

The RAWC Architecture was designed with the following in mind:

- Simplicity – Minimal number of components and software packages to install.
- Ease of use – GUI used to configure workloads, create log folders and launch and cleanup configuration files.
- Scalability – Unlimited number of virtual machines under test.
- Active Directory aware – Able to determine Group membership and locate the correct configuration file for the test.
- Policy based workload – Configure realistic workloads based on Group membership for large scale enterprise testing.

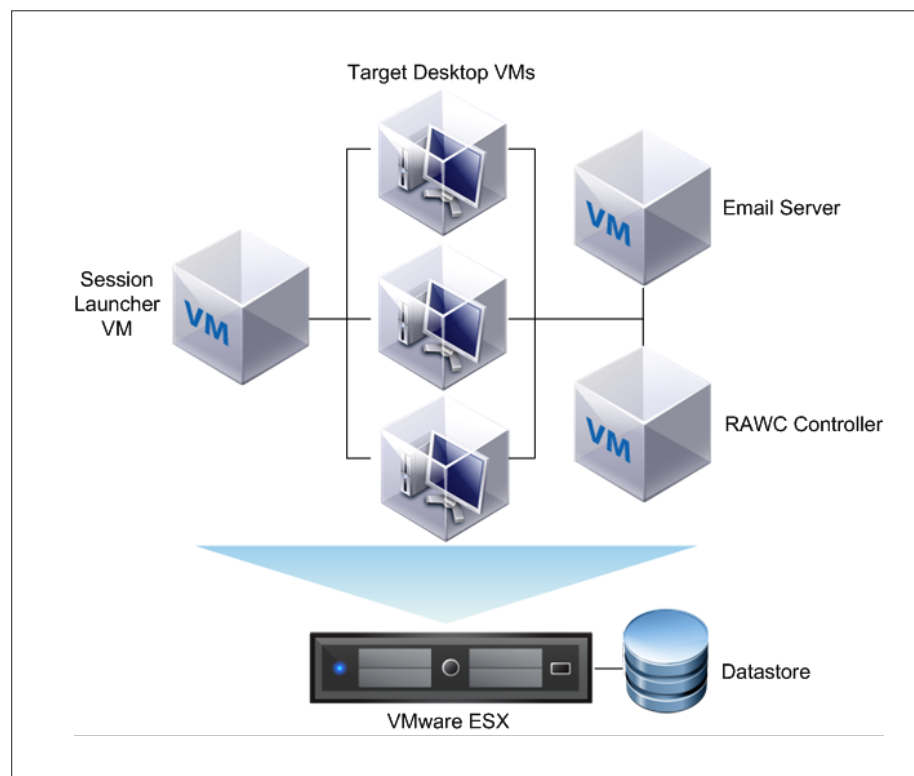


Figure 1 - RAWC Architecture

The components of the RAWC Architecture include the following:

Session Launcher Virtual Machine(s): One or more session launcher virtual machines must be set up to support the launching of desktop sessions. Each session launcher virtual machine can support up to 20 sessions.

Target Desktop Virtual Machines – Workload: The RAWC code resides on each of the desktop virtual machines.

Email Server Virtual Machine: The email server is required only if you are running Microsoft Outlook. You may use the email server that is provided or supply your own.

RAWC Controller Virtual Machine: The RAWC Controller hosts the RAWC GUI and the shared network folder for the configuration and log files. The RAWC Controller can be a physical or virtual machine.

RAWC Workload

Applications

The RAWC workload consists of a number of applications most commonly used in a Windows desktop environment. These applications are called randomly and the operations are described below.

- Microsoft Word (1 & 2) – Open, minimize and maximize the window, insert text, follow the cursor, save modifications, and close
- Microsoft Word (Random) – Open, minimize and maximize the window, write random words and numbers, save modifications, and close
- Microsoft Excel – Open, minimize and maximize the window, write random numbers in cells, insert and delete a column and row, copy and paste formulas, save modifications, and close
- Microsoft PowerPoint – Open a presentation, minimize and maximize the window, conduct a slide show presentation, and close
- Microsoft Outlook – Setup of a user account
- Microsoft Outlook – Open, send an email, minimize and maximize the window, send additional emails, execute a send/receive, and close
- Internet Explorer – Open, browse an html page, minimize and maximize the window, and close
- Windows Media Player – Open, view a video, and close
- Adobe Acrobat Reader – Open a PDF document, browse random pages, minimize and maximize the window, and close
- Java – Open, run a java application, and close
- McAfee Anti-virus – Execute a real time virus scan
- 7-Zip – Open, compress large files, and close

Active Directory Groups

The RAWC workload, if configured for Active Directory Groups, will query the Active Directory for the virtual machine's Group membership, then locate the proper configuration file and begin the test. This feature allows workloads to be defined based on the virtual machine's Group membership, thus enabling multiple workloads to run in the same test. For large-scale tests, this allows for a more realistic representation of an enterprise. For example: All groups need access to email, Internet Explorer and Adobe Reader, but Finance needs access to Excel and PowerPoint, whereas an individual in Manufacturing may not.

RAWC GUI

The RAWC GUI is used to configure RAWC. The GUI enables you to save and retrieve test variables, create log folders, and define unique workloads based on Active Directory Groups. In addition, you can use the GUI to increase or decrease the load or adjust the user behavior, such as the number of words per minute. The Configuration tab contains information about the test, logging information, email variables, and applications selected for each Group, if using Active Directory Groups.

Configuration Tab

The Configuration tab contains two sections, General and Workload.

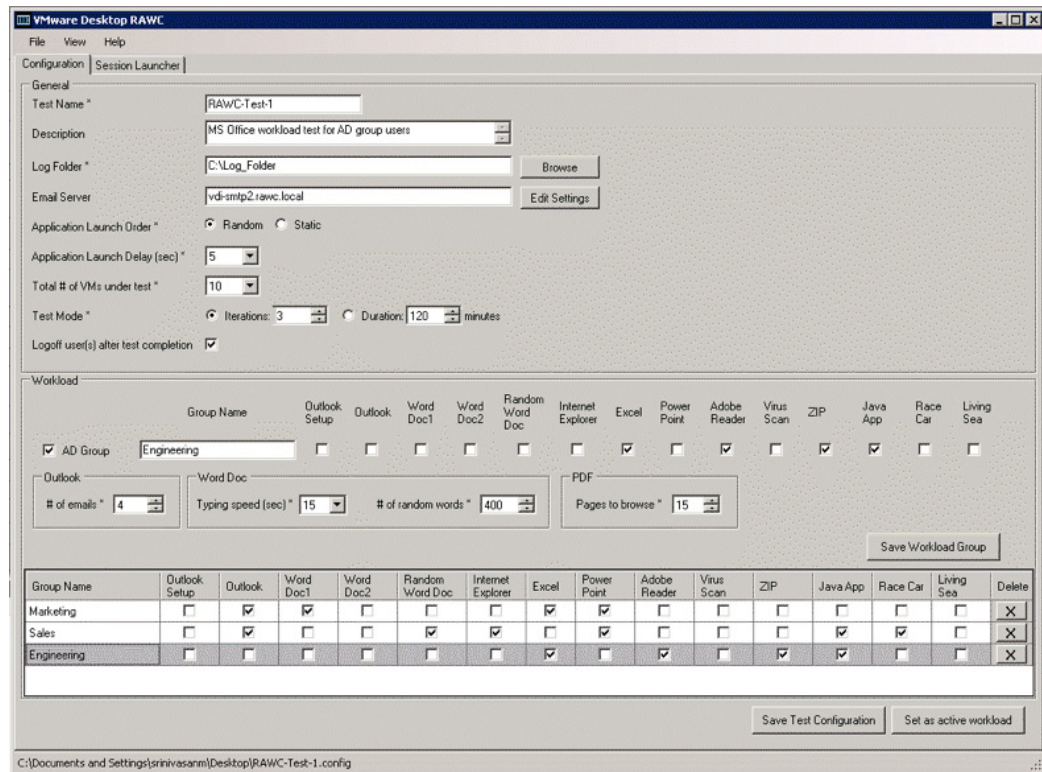


Figure 2 - Configuration Tab

Test Variables and Considerations

The General section allows you to configure the test variables.

Application Launch Order

This feature allows you to determine if applications will run in a random or static order.

Random Order

Random order is the standard mode in RAWC. By running applications in a random order, applications will be executed in a different (random) order on different virtual machines. When running applications in a random order, it is important to understand the following:

- Applications run in a random order.
- A varying Start Delay (explained in the next section) is in effect, depending on the number selected for the 'Total number of virtual machines under test.'
- The random delay between applications (Application Launch Delay) is calculated on the fly. When running applications in random mode, the user chooses the upper limit number used in the random delay calculation.

Static Order

Static order is a new feature in RAWC. By running applications in a static order, all virtual machines will run the exact same test, including the applications running in the same order and with fixed delays between the applications. This feature gives you the ability to repeat a test over and over again. When running applications in a static order, it is important to understand the following:

- The applications selected will run in the exact same order for each iteration on each virtual machine.
- There is no Start Delay (explained below) in effect. A varying Start Delay in this situation is irrelevant, as the exact same applications are running on each virtual machine, so if the test started in one minute or 12 minutes or 40 minutes, all virtual machines would start together.
- The random delay between applications (Application Launch Delay) is now a fixed delay. When running applications in static order, the upper limit number that is selected becomes the actual delay between the applications.

Application Launch Delay

The application launch delay is used to either increase or decrease the load on the server by running the applications closer or farther apart. Keep the following in mind:

- When running applications in a random order, the number selected for the application launch delay becomes the upper limit of a random delay that is calculated on the fly. For example, the random delay that is calculated could be 1 second or 119 seconds or any number in between if 120 seconds is selected.
- When running applications in a static order, the number selected for the launch delay, becomes the actual fixed delay between applications.
- Typically, the greater the number of virtual machines under test, the larger a number you will want to select, say above 120 seconds.
- By selecting a lower number, you can generate a heavier load by making the applications run closer together.

Total Number of Virtual Machines Under Test

The total number of virtual machines under test is used to determine the 'Start Delay' of the virtual machine. This feature allows you to simulate a normal login scenario or create a more intensive 'Boot Storm.' A 'Boot Storm' occurs when more virtual machines start their workloads closer together, thus creating a larger load on the server. A 'Boot Storm' scenario is likely to occur after systems have recovered from a power outage or logging in first thing in the morning.

- A normal login scenario allows virtual machines to start their workloads over a period of time. Select the number of virtual machines that is closest to the number of virtual machines you have under test.
 - 10 virtual machines under test - No start delay
 - 64 virtual machines under test - 0 to 12 minutes
 - 250 virtual machines under test - 0 to 24 minutes
 - 500-1000 virtual machines under test - 0 to 54 minutes
- A 'Boot Storm' scenario can be created by simply selecting a lesser number of virtual machines under test. For example: If you are running a 500 virtual machine test, select '250' for the 'Total number of virtual machines under test.' This means that 500 virtual machines will start within 0 to 24 minutes of each other.

Test Mode

The GUI allows you to select either the number of iterations or the duration (minutes) that a RAWC test will run.

Iterations

The GUI allows you to select the number of iterations you want the virtual machines to run. If you are interested in the open and close times of applications, you will want to run at least three iterations. Applications take longer to open the first time. By the second iteration, applications/pages have been cached and the open times are much shorter.

Duration

Duration is a new feature in RAWC. The GUI allows you to select the number of minutes you want the test to run. You no longer have to guess at how many iterations it will take to run a four hour or even eight hour test. When using duration, it is important to understand the following:

- Virtual machines do not finish at exactly the same time or exactly on the duration time.
- Time checks are performed before and after random (or static) delays between applications.
- Time checks are performed at certain points while an application is performing. For example:
 - A time check occurs before and after an email is sent, but not during the creation of an email.
 - A time check occurs before and after the slide presentation.
- When the duration time has been met, the log files are closed and the desktop is cleaned of any open applications.

If you are using Active Directory Groups and running in the duration test mode, you will notice that lighter workloads (2-3 applications) will run more iterations of the applications than heavier workloads (5+ applications). This is fairly accurate, as task workers are limited in the number of applications they run, and more often than not, use these applications more often.

Logoff Users

Logoff users is a new feature in RAWC. You can now specify via the GUI if the user logs out of the virtual machine after the test completes or remains logged in.

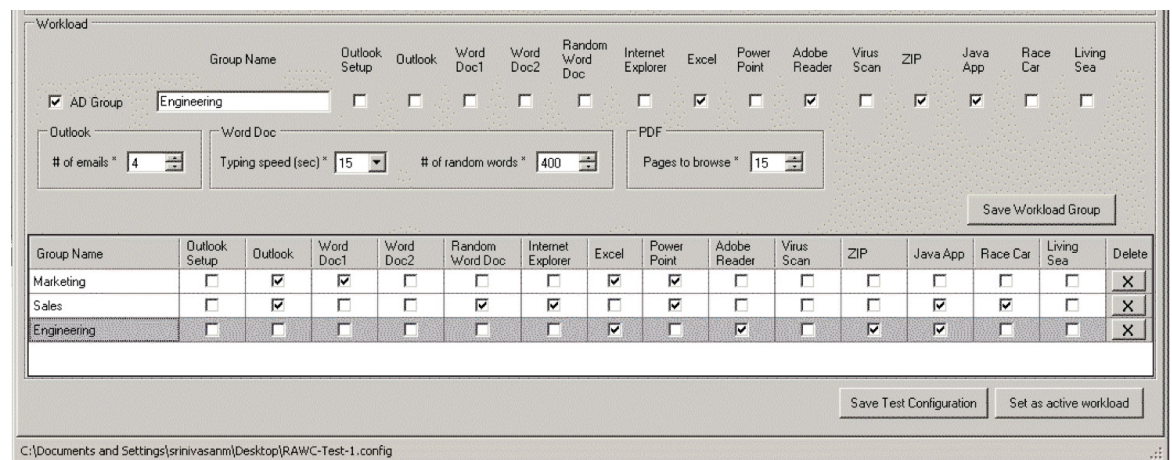


Figure 3 - Workload Section of Configuration Tab

Workload Variables and Considerations

The Workload section allows you to define a unique workload and workload variables per Active Directory Group. If your environment does not include an Active Directory, a group name of 'NA' can be used. The variables in this section are more reflective of user behavior, but can also be used to create an additional load on the server.

Number of emails +1

You can configure the number of emails that are created and sent, via the GUI. Outlook is always the first application that is opened, with an email being created and sent. Other applications are then opened randomly and operations performed. Outlook is then randomly called later in the test and the number of emails specified in the GUI are created and sent.

Typing Speed in Seconds

You can manipulate user behavior by selecting either a fast, medium, or slow typing speed. By selecting the fast typing speed, you can move through the word applications quickly.

Number of Random Words

The random word document allows you to define the number of random words and numbers that are written to the document. This particular application is able to create an additional load on the virtual machine since there is a calculation to determine what word or number is written to the document.

PDF Pages to Browse

You can define the number of random pages that are browsed in the PDF document. To increase the load, decrease the number of pages browsed so that the next application is accessed faster.

Results

The esxtop tool was used to monitor in real time and gather data about the state of the physical server running an ESX Server. The esxtop tool is invoked from the command line and provides information on the CPU utilization of each physical processor, memory utilization, and disk and network bandwidth for each disk and network device available to the ESX Server machine.

Information about interpreting esxtop statistics can be found here:

<http://communities.vmware.com/docs/DOC-9279>

Application open and close times can be gathered from the timing logs deposited by each virtual machine in the shared network folder on the Test Master virtual machine. This data can be manipulated and graphed in Excel.

Conclusion

As the desktop virtualization world evolves, so will the desktop workload tools. Know your environment (network, servers, storage and desktops), and your users (applications, light/heavy use), before determining which workload tool to use. Research the various tools and even conduct an assessment to determine if virtualization is right for your company. Make sure the tool is easy to install and use and that it simulates the applications you run most. Be sure to allow enough time for a proper assessment, testing and interpretation of the results, as this will aid in the proper planning and deployment of your virtual desktop infrastructure.

About the Author

Fred Schimscheimer is a Senior Technical Marketing Manager at VMware. In this role, he works as part of the product marketing team as an expert in storage and workloads for virtual desktop solutions. Fred is the architect and developer of the RAWC workload.

Acknowledgements

I would like to thank Mason Uyeda for his continued support and encouragement in developing the RAWC workload. I would like to give special thanks to Srinivasan Muthu for his expertise in designing and coding the RAWC GUI and packaging. I would also like to thank Scott Drummonds for introducing me to workload automation and providing the initial ideas for the workload.

White Papers

The following white papers used the RAWC workload:

Vblock Powered Solutions for VMware View 4.5 Reference Architecture

<http://www.vce.com/pdf/solutions/vce-vmware-view-reference-architecture.pdf>

VMware View 4.5 on Cisco Unified Computing System and EMC Unified Storage Design Guide

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/ucs_view_emc.html

VMware View 4.5 on Cisco Unified Computing System and NetApp Storage Design Guide

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/ucs_view_netapp.html

Simplifying enterprise desktop deployment and management using Dell EqualLogic storage and VMware View

<http://www.vmware.com/files/pdf/Dell-VMware-View-RA.pdf>

Desktop Sizing Guide: VMware View 4.0 and VMware vSphere 4.0 Update 1

http://www.mainline.com/_web/pdfs/mrktgMaterials/VMware-WP-VMsizing-050310.pdf

50,000-Seat VMware View Deployment

<http://www.vmware.com/files/pdf/VMware-View-50kSeatDeployment-WP-EN.pdf>

VMware View Reference Architecture – A Guide to Large-scale Enterprise VMware View Deployments

<http://www.vmware.com/files/pdf/resources/vmware-view-reference-architecture.pdf>

VMware ThinApp Reference Architecture – A Guide for Enterprise VMware ThinApp Deployments

<http://www.vmware.com/files/pdf/thinapp-ref-arch.pdf>

EMC CLARiiON Storage – A Guide to deploying EMC CLARiiON CX4-240 FC with VMware View

<http://www.vmware.com/files/pdf/view-deployment-emc-clariion-fc.pdf>

EMC CLARiiON Storage – A Guide to deploying EMC CLARiiON CX4-240 iSCSI with VMware View

<http://www.vmware.com/files/pdf/view-deployment-emc-clariion-iscsi.pdf>

EMC Celerra Unified Storage – A Guide to deploying EMC Celerra NS20 storage with VMware View

<http://www.vmware.com/files/pdf/resources/vmware-view-ns20-deployment-guide.pdf>

Scalability Study for Deploying VMware View on Cisco UCS and EMC Symmetrix V-Max Systems

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/App_Networking/vdiucswp.html

Using Local Disks in a VMware View Deployment

http://www.vmware.com/files/pdf/view_local_disk.pdf

The VMware Reference Architecture for Cisco UCS and EMC CLARiiON with VMware View 4
<http://vmware.com/go/vce-ra-brief>

Deploying Virtual Desktops with VMware View 4 on EMC Celerra NS-120 Validation Test Report
<http://www.vmware.com/files/pdf/VMware-View4-EMC-NS-120-VTR.pdf>

VMware View on NetApp Storage Solution Brief
<http://www.vmware.com/files/pdf/VMware-VMwareView-NetAppStorage-WP-EN.pdf>

