

# Virtual Systems Workload Characterization

## An Overview

Mohamed A. El-Refaey

Arab Academy for Science, Technology and Maritime Transport  
 College of Computing & Information Technology  
 Cairo, Egypt  
 e-mail: melrefaey@acm.org

Dr. Mohamed Abu Rizkaa

Arab Academy for Science, Technology and Maritime Transport  
 College of Computing & Information Technology  
 Cairo, Egypt  
 e-mail: m\_rizkaa@aast.edu

**Abstract**—Virtual systems and virtualization technology are taking the momentum nowadays in data centers and IT infrastructure models.

Performance analysis of such systems is very invaluable for enterprises but yet is not a deterministic process. Single-workload benchmark is useful in quantifying the virtualization overhead within a single VM, but not useful in whole virtualized environment with multiple isolated VM and varying workload on each and can't capture the system behavior.

We need a common workload model and methodology for virtualized systems so that benchmark results can be compared across different platforms.

In this paper we will present an overview of the key requirements and characteristics of virtual systems performance metrics and workload characterization which can be considered one step further in implementing virtual systems benchmark and performance model that describe the effect of the applications, host operating system and the hypervisor layer on the performance metrics of virtual workloads. An overview of Intel® vCon model and VMware VMmark will be introduced as examples for the consolidated servers' workload evaluation.

**Keywords**-Virtualization;benchmark;vConsolidate;VMmark;SPEC; Performance

### I. INTRODUCTION

The Virtualization technology is taking the momentum these days in data centers and IT infrastructure models and patterns.

Virtualization is very similar conceptually to emulation. With emulation, a system pretends to be another system. With virtualization, a system pretends to be two or more of the same system.

As shown in figure 1, the virtualization layer will partition the physical resource of the underlying physical server into multiple virtual machines with different loads. the fascinating thing about this virtualization layer is that it schedules and allocates the physical resource and makes each virtual machine think that it totally owns all the underlying hardware physical resource(processor, disks, rams etc.).

Most modern operating systems contain a simplified system of virtualization. Each running process is able to act

as if it is the only thing running. The CPUs and memory are virtualized. If a process tries to consume all of the CPU, a modern operating system will preempt it and allow others their fair share. Similarly, a running process typically has its own virtual address space that the operating system maps to physical memory to give the process the illusion that it is the only user of RAM [2].

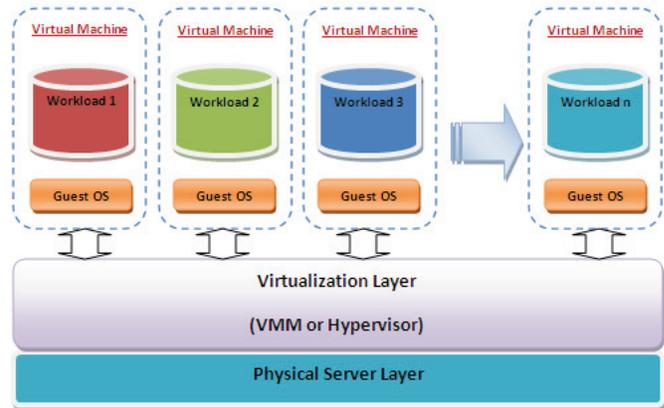


Figure 1. Virtual environment illustrating multiple VMs with corresponding workloads

Lots of institutions are now finding their assets of servers are all doing single tasks, or small clusters of related tasks. Virtualization allows a number of virtual servers to be consolidated into a single physical machine, without losing the security gained by having completely isolated environments. Several Web hosting providers are using virtualization intensively, because it let them offer each client his own virtual machine without requiring a physical machine taking up rack space in the data center.

### II. WORKLOAD CHARACTERIZATION

The dictionary defines workload as "the amount of work assigned to, or done by, a worker or unit of workers in a given time period" (The American Heritage Dictionary, 2nd Edition).

Within the confines of a network, workload is the amount of work assigned to, or done by, a client, workgroup, server,

or internetwork in a given time period. Therefore, workload characterization is the science that observes, identifies and explains the phenomena of work in a manner that simplifies your understanding of how the network is being used.

With the use of graphs and descriptive metrics, you can begin to collect useful historical information concerning your networks. This historical information, describing the volume, intensity, and patterns of workload created by your clientele, is the only accurate foundation for performance evaluations of any kind.

In order to test multiple alternatives under identical conditions, the workload should be repeatable. Since a real-user environment is generally not repeatable, it is necessary to study the real-user environments, observe the key characteristics, and develop a workload model that can be used repeatedly. This process is called workload characterization. Once a workload model is available, the effect of changes in the workload and system can be studied in a controlled manner by simply changing the parameters of the model [1].

The measured workload data consists of services requested or the resource demands of a number of users on the system. Here the term user denotes the entity that makes the service requests at the SUT (System Under Test) interface. The user may or may not be a human being. For example, if the SUT is a processor, the users may be various programs or batch jobs. Similarly, the users of a local-area network are the stations on the network. In workload characterization literature, the term workload component or workload unit is used instead of the user.

Workload characterization consists of characterizing a typical user or workload component. Other examples of workload components are as follows:

- Applications: If one wants to characterize the behaviour of various applications, such as mail, text editing, or program development, then each application may be considered a workload component and the average behaviour of each application may be characterized.
- Sites: If one desires to characterize the workload at each of several locations of an organization, the sites may be used as workload components.
- User Sessions: Complete user sessions from login to logout may be monitored, and applications run during the session may be combined.

The key requirement for the selection of the workload component is that it be at the SUT interface. Another consideration is that each component should represent as homogeneous a group as possible. For example, if users at a site are very different, combining them into a site workload may not be meaningful. The purpose of the study and the domain of the control of the decision makers also affect the choice of components. For example, a mail system designer is more interested in determining a typical mail session than a typical user session combining many different applications.

The measured quantities, service requests, or resource demands, which are used to model or characterize the workload, are called workload parameters or workload features. Examples of workload parameters are transaction types, instructions, packet sizes, source destinations of a packet, and page reference pattern.

In choosing the parameters to characterize the workload, it is preferable to use those parameters that depend on the workload rather than on the system. For example, the elapsed time (response time) for a transaction is not appropriate as a workload parameter, since it depends highly on the system on which the transaction is executed. This is one reason why the number of service requests rather than the amount of resource demanded is preferable as a workload parameter. For example, it is better to characterize a network mail session by the size of the message or the number of recipients rather than by the CPU time and the number of network messages, which will vary from one system to the next.

There are several characteristics of service requests (or resource demands) that are of interest. For example, arrival time, type of request or the resource demanded, duration of the request, and quantity of the resource demanded by each request may be represented in the workload model. Particularly those characteristics that have a significant impact on the performance should be included in the workload parameters, and those that have little impact should be excluded. For example, if the packet size has no impact on packet forwarding time at a router, it may be omitted from the list of workload parameters, and only the number of packets and arrival times of packets may be used instead.

The following techniques have been used in the past for workload characterization:

1. Averaging
2. Specifying dispersion
3. Single-parameter histograms
4. Multi-parameter histograms
5. Principal-component analysis
6. Markov models
7. Clustering

More details about the above techniques can be found in [1].

vCon model described later in the paper is based on the principal-component analysis technique (commonly used to classify workload components by the weighted sum of their parameter values. Using  $a_j$  as weight for the  $j$ th parameter  $x_j$ , the weighted sum  $y$  is:

$$y = \sum_{j=1}^n a_j x_j$$

This sum can then be used to classify the components into a number of classes such as low demand or medium demand. Although this technique is commonly used in performance analysis software, in most cases, the person running the software is asked to choose the weight. Without any concrete guidelines, the person may assign weights such

that workload components with very different characteristics may be grouped together, and the mean characteristics of the group may not correspond to any member [1].

Workload performance characterization is a very beneficial tool for systems architects, VMM developers and system administrators.

For system architects, it is very useful to project and indicate a future platform performance and how the applications can scale for future systems and platforms and it shows the virtualization impacts and overhead.

For VMM developers and designers, it provides a helpful feedback about how resources are scheduled and how they can optimize the resource scheduling and implement new scheduling techniques and algorithms.

For system administrators, this will help more in a fair share of resources provided to users and optimize the performance of their workloads. This will make an efficient use of the data center resources and also will help in rush hours, in which some virtual machines suffer from a peak load and needs to be migrated to another machine to accommodate that load.

### III. VIRTUAL SYSTEMS BENCHMARKS

In the context of virtual performance evaluation and workload characterization, there should be an overview about the benchmarks already exist in the field, along with the benchmark characteristics of the virtual machine.

Here is a minimum list of the characteristic of virtual machine benchmark required [5]:

- Able to capture the virtualized environment key performance characteristics.
- Able to define an easily understandable metric that scales with underlying system capacity.
- With this scalable metrics, the same benchmark can be used for different servers' sizes.
- The benchmark specs are platform neutral.
- The measurements can be generated using a controlled policy based on a combination of increased individual workload scores and running an increasing number of workloads.

Some details about the Intel® vConsolidate (vCon) and VMware VMmark benchmarks will be introduced below.

### IV. INTEL® vCONSOLIDATE

The vConsolidate (vCon) benchmark is one of the proposed benchmarks for virtualization consolidation developed by Intel and it can be considered a VMM agnostic.

The vCon benchmark consists of a compute intensive workload/application, a database workload, a Web server workload, and a mail server workload. Each of these workloads runs in its own VM.

In Figure 2, a high-level illustration of the system running vConsolidate is shown. At the base, there is a physical platform layer, then a virtualization layer, and some multiple numbers of VMs, each running a designated workload. vCon defined an aggregation strategy that helps consolidate workloads, define the performance metric(s) and tell how measurements will be taken.

vCon model is designed as some multiple of workloads (mentioned above) running and designate the weights of each as Weight[i]. All the weights are fixed numbers pre-defined by the workload. Each workload can be run on the system before and after any virtualization layer is included, and this can help comparing the virtualized performance of each workload with a baseline measured without virtualization on a pre-defined standard machine. This serves as a useful tool for calibrating subsequent results. And then replicate these workloads as needed based upon a set of usage model requirements. Then, calculate the ratio of virtualized/baseline as the relative performance of each workload. The performance of the virtualized environment would be

$$\sum_{i=1}^N Weight[i] * WorkloadPerf[i]$$

Where WorkloadPerf[i] is the relative performance of the i'th workload

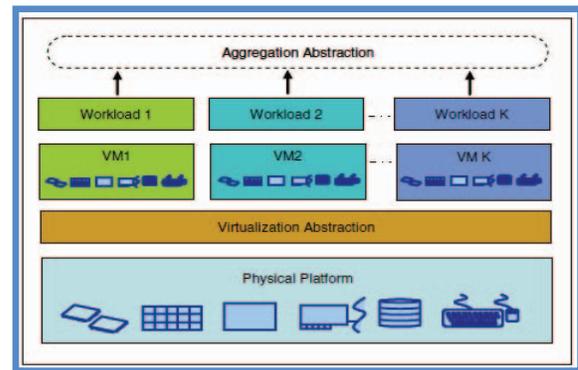


Figure 2. vConsolidate concept

To emulate a real world environment an idle VM is added to the mix since datacenters are not fully utilized all the time. The compute intensive VM runs SPECjbb. Typically SPECjbb is a CPU intensive workload that consumes as much CPU as it possibly can. However, in this environment, namely vCon, SPECjbb has been modified to consume roughly 75% of the CPU or so, by inserting in random sleeps every few milliseconds. This is to represent workloads that are more realistic. The database VM runs Sysbench; an OLTP workload running transactions against a mysql database. The Webserver VM runs Webbench which uses Apache Webserver.

The Mail VM is a Microsoft Exchange workload that runs transactions on Outlook with 500 users logged in

simultaneously. A configuration as described above with 5VMs running the different workloads comprises a Consolidated Stack Unit or known as CSU. The diagram in Figure 3 represents a 1CSU configuration [3].

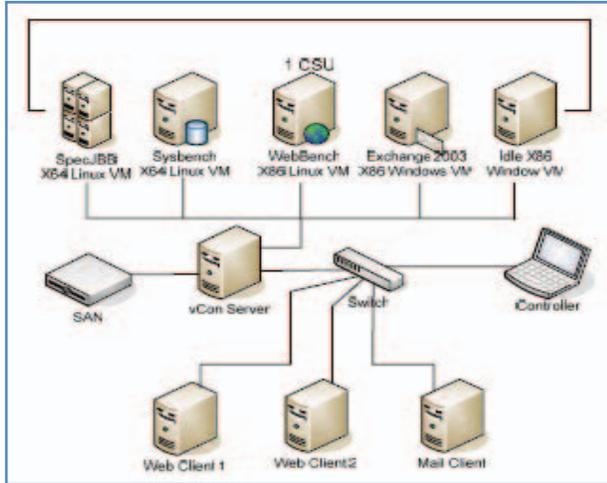


Figure 3. vConsolidate 1 CSU configuration

A. vConsolidate example [5]

- The consolidation workload consists of one Web server, one e-mail server, and one data base server.
- To mimic the real-world scenario, one idle VM is also running on the physical system. In that VM, no real workload is running. We do not take any score from this idle VM.
- The weight factors are 35% for the one Web server, 20% for the e-mail server, and 45% for the database server. This would correspond to a weight vector of (0.35, 0.20, 0.45).
- After testing each workload individually in a non-virtualized pre-defined standard machine with a specific configuration, we can get the baseline of each workload.
- At this stage, we define how the workloads were mapped into VMs, how the VMs were mapped to the underlying physical platform, and how resources were allocated amongst each. This is not defined by the workload. The user chooses the best VM configuration settings to do the measurements.
- Each of the workload components has a well-defined performance metric and a known, unvirtualized baseline result. The observed result for the virtualization of each of the component workloads is normalized to this known baseline. The resulting benchmark metric is the combination of all of the normalized workload component results. As an example, the performance results ratio vector

WorkloadPerf[i] could look like (1.8, 1.5, 2.3)<sup>2</sup>. The rollup result for the performance of the virtualized system would be

$$\sum_{i=1}^N Weight[i] * WorkloadPerf[i]$$

or  
 $(0.35, 0.20, 0.45) * (1.8, 1.5, 2.3) = 1.965$

This calculation is illustrated in Figure 4. The result becomes most useful when comparing different configurations. For example, the same consolidated set of workloads can be compared across different platforms or against two different virtualization monitors or it can be used to compare two different sets of configuration settings.

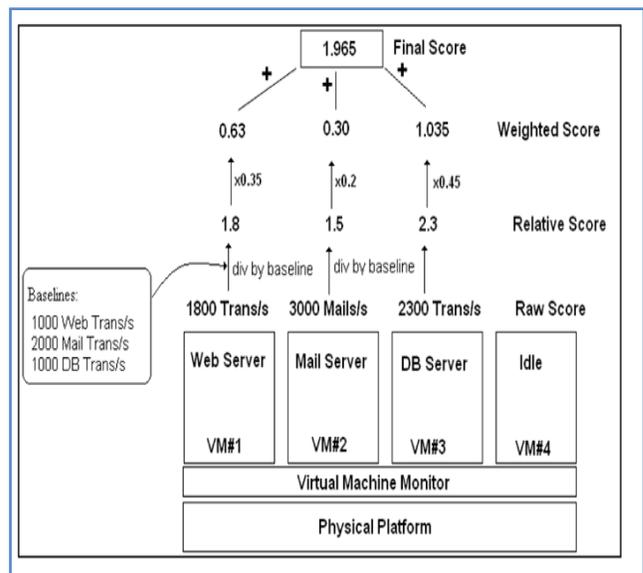


Figure 4. vConsolidate example virtualization benchmark results calculation

The vConsolidate benchmark was presented as an example implementation, highlighting the compromises required in workload selection, component definition, and metric aggregation [5].

V. VMWARE VMARK BENCHMARK

The ultimate goal of VMmark benchmark is to create a meaningful measurement of virtualization performance across a wide range of hardware platforms. Server consolidation typically collects several diverse workloads onto a single physical server. This approach ensures that all system resources such as CPU, network, and disk are more efficiently utilized. In fact, virtual environments tend to function more smoothly when demands are balanced across physical resources [4].

The unit of work for a benchmark of virtualized system can be defined as a collection of virtual machines executing a set of different workloads. The VMmark benchmark refers to

this unit of work as a tile (relatively heavyweight objects). The total number of tiles that a physical system and virtualization layer can accommodate gives a coarse-grain measure of that system's consolidation capacity.

Workload within a tile is constrained to execute at less than full utilization of its virtual machine. However, the performance of each workload can vary to a degree with the speed and capabilities of the underlying system. For instance, disk-centric workloads might respond to the addition of a fast disk array with a more favorable score. These variations can capture system improvements that do not warrant the addition of another tile. However, the workload throttling will force the use of additional tiles for large jumps in system performance. When the number of tiles is increased, workloads in existing tiles might measure lower performance. However, if the system has not been overcommitted, the aggregate score, including the new tile, should increase.

The result is a flexible benchmark metric that provides a relative measure of the number of workloads that can be supported by a particular system as well as the overall performance level within the virtual machines [4].

A tiled consolidation benchmark should be based upon a set of relevant data center workloads. Examples of workloads are mail server, Java server, standby server, web server, database server and file server

And these loads are illustrated in the figure 5 shown below:

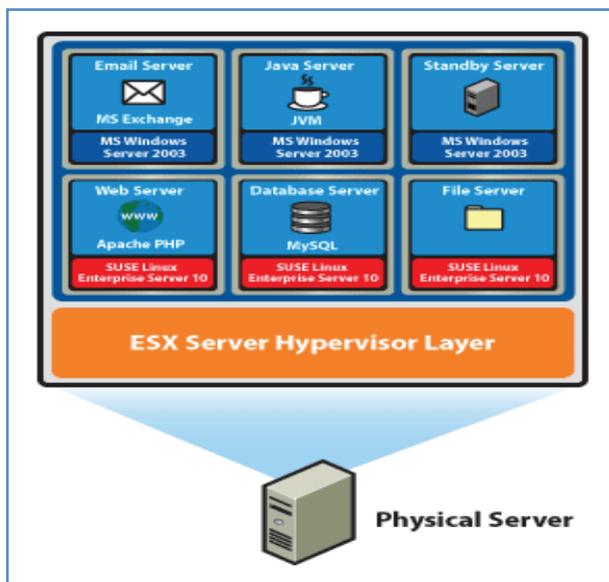


Figure 5. diversified workloads used in VMmark benchmark test (ESX Server Hypervisor layer is used)

The workloads comprising each tile are run simultaneously in separate virtual machines at load levels that are typical of virtualized environments. The performance of each workload is measured and combined with the other workloads to form a measuring score for each individual tile. Multiple tiles can run simultaneously to increase the overall

score. This approach allows smaller increases in system performance to be reflected by increased scores in a single tile and larger gains in system capacity to be captured to adding additional tiles.

#### A. VMmark Scoring

VMmark score is evaluated as follows, by the end of VMmark test cycle (run for minimum three hours with workload metrics reported every 60 seconds, each workload reports its performance metric. These metrics collected are shown in Table 1.

TABLE 1. INDIVIDUAL VMARK WORKLOAD METRICS.

Workload	Metric
Mail server	Actions/minute
Java server	New orders/second
Standby server	None
Web server	Accesses/second
Database server	Commits/second
File server	MB/second

Once all workloads have reached steady state during a benchmark test cycle, a two-hour measurement interval is taken. This steady-state interval is then divided into three 40-minute sections. For each of the 40-minute sections, the results for the tile are computed. The median score of the three sections is selected as the raw score for the tile. For multi-tile runs, the median of the sums of the per-tile scores would be used as the raw score [4]. Workload metrics for each tile are calculated and aggregated into a score for that tile. This aggregation is performed by first normalizing the different performance metrics such as MB/s and database commits/s with respect to a reference system. Then, a geometric mean of the normalized scores is computed as the final score for the tile. The resulting scores are then summed to create the final metric.

More details about the experimental setup and results for the workloads aforementioned above can be found in [4].

A final note worth to be mentioned in this context is that VMmark can only be run under licensing and publication regulations and only with a VMware solution (ESX server as an example).

So a standard benchmark for virtual systems is needed to fill this gap. For this, the SPEC (Standard Performance Evaluation Corporation) committee is developing a standard vSPEC benchmark for virtualization [6].

The SPEC Virtualization Committee is developing a new industry standard benchmark for evaluating virtualization performance for data center servers.

The committee plans to make use of the comprehensive system-level application-based benchmarks that SPEC offers. The workloads represented by these benchmarks are also representative of commonly virtualized server applications. SPEC's expertise in system level benchmarks that support a wide range of hardware architectures and operating systems will greatly benefit the committee's efforts to develop a standard methodology to evaluate the performance of servers using virtualization for server consolidation[7].

## VI. CONCLUSIONS AND FUTURE WORK

In this paper we presented an overview about the key characteristics of virtual environment benchmarks along with an overview about the workload characterization and its techniques. And introduced some of the existing benchmarks for virtual environment Intel® vConsolidate and VMware VMark.

vConsolidate is a hypervisor agnostic benchmark and VMmark can only be run under licensing and publication regulations and only with a VMware solution.

Based on this overview and the requirements of designing a virtual system benchmark it will be useful to start implementing a preliminary benchmark and performance analysis model for the virtualized environment.

## REFERENCES

- [1] R. K. Jain. Art Of Computer Systems Performance Analysis - Techniques For Experimental Design Measurements Simulation And Modeling, John Wiley & Sons (1991).
- [2] David Chisnall. The Definitive Guide to the Xen Hypervisor, Prentice Hall (2008).
- [3] Padma Apparao, Ravi Iyer, Xiaomin Zhang, Don Newell, Tom Adelmeyer. Characterization & Analysis of a Server Consolidation Benchmark.
- [4] VMmark: A Scalable Benchmark for Virtualized Systems. Technical Report VMware-TR-2006-002 September 25, 2006 [http://www.vmware.com/pdf/vmmark\\_intro.pdf](http://www.vmware.com/pdf/vmmark_intro.pdf)
- [5] Jeffrey P. Casazza, Michael Greenfield, Kan Shi - Redefining server performance characterization for virtualization benchmarking
- [6] Standard Performance Evaluation Corporation - <http://www.spec.org/>
- [7] [SPEC Virtualization http://www.spec.org/specvirtualization/](http://www.spec.org/specvirtualization/)