A Survey of Performance Modeling and Analysis Issues in Resource Management Across x86-based Hypervisors in Enterprise Cloud Computing Deployments

Melynda Eden, meden@wustl.edu (A paper written under the guidance of Prof. Raj Jain)

Download

Abstract:

In virtualized environments, physical resources are partitioned into virtual resources so that they can be shared among individual virtual machines. This sharing of virtualized resources can reduce waste by increasing used capacity thereby reducing the need for as many physical machines. Currently, there is an increased desire to use virtualized systems in enterprise cloud computing in order to more efficiently utilize resources thereby reducing costs. However, it is challenging to accurately model virtualized systems in order to analyze performance issues. Traditional modeling methods that were developed for physical machines do not seamlessly work in virtualized environments. This paper explores the complexities associated with correctly modeling and measuring the performance of virtualized systems, and surveys various new approaches in performance analysis for virtualization including recently developed benchmarks, new modeling methodologies, and updated simulation techniques.

Keywords: virtualized environments, performance modeling, virtualization benchmarks, simulation framework, cloud computing, hypervisors, VMware, Xen, server partitioning, Benchvm, VMmark, vConsolidate, VSCBenchmark, Performance Metric M, Linear Parameter Varying, Fuzzy models, Artificial Neural Networks, probabilistic performance model, CloudSim

Table of Contents

1. Introduction

Virtualization is gaining popularity in enterprise cloud computing environments because of realized cost savings and better management that result from resource sharing and server consolidation. In an effort to improve virtualized systems, researchers look to modeling and simulation techniques in order to pinpoint performance bottlenecks and resource contention. Many challenges unfold in the development of accurate performance models, as virtualized environments add levels of complexity beyond models suitable for single physical machines [Benevenuto06]. This paper first defines virtualization, cloud computing and hypervisors; then investigates the multifaceted problem of performance analysis issues, and finally explores recent solutions in measurement, modeling and simulation for virtualized environments.

1.1 Overview of Virtualization

Virtualization allows for the partitioning of physical resources into virtualized containers. These containers, more commonly referred to as virtual machines, are functionally nothing more than sets of files that represent virtual hardware, executing in the context of a hypervisor. This virtual hardware provides a platform on which an operating system and applications can be installed and can be configured to provide almost any service that a physical server typically provides [Ardagna10]. These services include email, database, hosting, file and print, monitoring, management, and more. The virtualized environment consists of the underlying hardware, the hypervisor, the Virtual Machine Monitor (VMM), the virtual machines (VMs), and the operating systems and applications installed on these virtual machines.
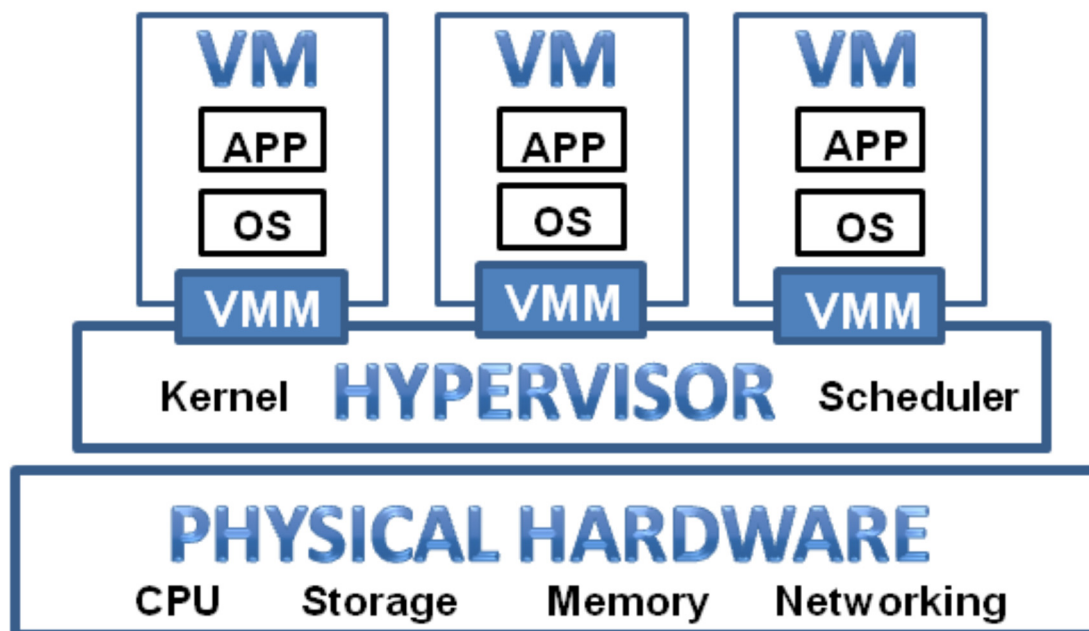
Figure 1: Virtualized Environment Representation [Hostway11]

Next, the VMM and hypervisor are discussed.

1.2 Terminology Challenges: Virtual Machine Monitor

The components of the virtualized environment include the hypervisor and the VMM. Unfortunately, terminology for what goes on inside a hypervisor is not consistently applied, most notably with the concept of the VMM. Most resources agree that the VMM is involved with the scheduling of virtual resources on their underlying physical counterparts, however in some cases the VMM is referred to as the hypervisor itself while in others this is merely a component of a hypervisor in a broader context. To avoid confusion, the VMM will not be a specific point of focus in this paper, since its meaning would change based on context. The following section explores various hypervisors.

1.3 x86 Bare-Metal Hypervisors

Hypervisors are an important component of virtualized environments. Hypervisors are programs that allow multiple operating systems, known as guests, to run in virtual machines in an isolated fashion, and thus share a single physical machine, or host. Many different hypervisors have been developed, a few of which are listed in the following table:

Table 1: Common x86 Bare-Metal Hypervisors [Lu06] [VMware10]

x86 Hypervisors	Vendor	Type	Licensing	Pricing
XenServer	Citrix	Para, Full	Open Source	Basic implementation is free
ESXi	VMware	Para, Full	Proprietary	Basic implementation is free
Hyper-V	Microsoft	Full	Proprietary	Requires OS purchase
KVM	RedHat	Full	Open Source	Basic implementation is free

Some x86 hypervisors can support both fully virtualized and para-virtualized operating systems. Fully

virtualized systems are unaware they are running in a virtualized environment, and no modifications to the core operating system are required in order for them to function. Para-virtualized systems, on the other hand, are specifically modified to work in concert with the hypervisor. Prior to recent advances in CPU virtualization technologies, para-virtualized systems enabled performance improvements because the hypervisor could grant a para-virtualized system direct access to the underlying hardware resources on the host. However, para-virtualized operating systems cannot readily be migrated across different hosts in a heterogeneous environment [Virtualization06]. Recent advances in CPU hardware acceleration have exceeded performance gains garnered by para-virtualized systems, and as such, some vendors are phasing out support for them in their product stack [GuestOS11].

Another form of para-virtualization is device-based para-virtualization, in which a fully virtualized operating system is configured with special para-virtualized device drivers. These device drivers improve performance for specific devices, such as virtual network cards or virtual storage controllers, without sacrificing mobility. This form of para-virtualization is expected to prove beneficial for the foreseeable future [GuestOS11]

This paper specifically focuses on â€œbare-metalâ€ hypervisors which are installed directly onto physical hardware. All hypervisors listed in the table above are bare-metal hypervisors. Another class of hypervisors, known as â€œhostedâ€ architectures, requires an underlying operating system such as Windows, Linux, or Mac OS in order to function. These are typically not deployed in modern enterprise datacenters due to reduced performance characteristics. The next section provides an overview of the concept of cloud computing.

1.4 Introduction to Cloud Computing

Cloud computing is a model for pooling IT resources to provide real-time, on-demand, self-provisioned services to business users on an as-needed basis. From the userâ€™s perspective, this model requires minimal management and interaction with IT staff, streamlined provisioning processes, and provides dramatic cost savings over traditional IT [Calheiros09]. On the back end, cloud computing requires the establishment of complex networking, storage, and server configurations that optimally are configured to be self-monitoring and self-healing. Clouds are self-monitoring in that they should be able to automatically balance workloads across physical assets during peak periods, while conversely decommissioning idle resources to conserve energy and reduce costs during non-peak periods. Clouds are self-healing in that the failure of any individual physical hardware or software component is remediated in the shortest possible time, and this implies rapid recovery even in the event of a complete datacenter outage.

Figure 2: The Concept of Cloud Computing [Wiki11]

In the following section, the three categories of cloud computing are explained.

1.5 Cloud Computing Categories

There are three categories of cloud computing: external, internal, and hybrid. External clouds are those provided by companies like Amazon and Google that are used to provide commodity IT services at prices that individual entities (businesses, research groups, etc.) could not match using internal resources. Internal clouds are used by companies that are not ready to embrace external clouds over concerns about data security and loss of centralized control over critical business processes. The emerging hybrid cloud model is being employed by companies that desire to benefit from the cost efficiencies of external clouds for their own commodity services while at the same time maintaining control over IT services that have critical, sensitive, or proprietary natures.

Figure 3: Cloud Computing Categories Representation [vCloud11]
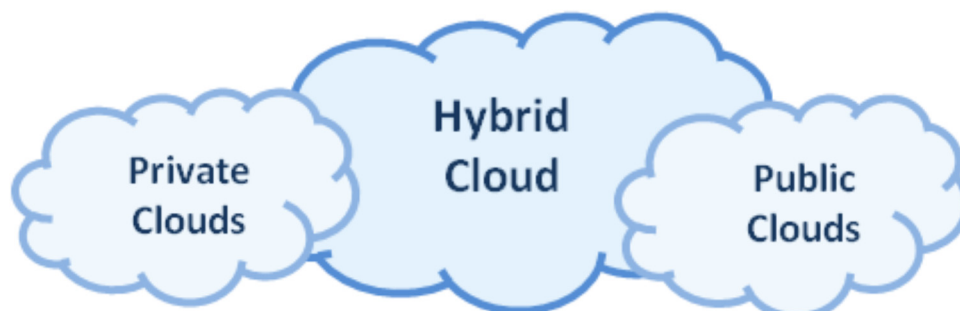
This section introduced the concepts of cloud computing and virtualization, and explored categories of cloud computing as well as the components of virtualized environments. The next section defines the complex issues that are encountered in performance analysis of these virtualized environments.

2. Issues with Performance Analysis in Virtualized Environments

By design, virtual machines share resources and thus compete for resources. It is difficult to accurately model a system with resource contention when many factors, including number of virtual machines, type of hypervisor, differences in underlying and support hardware, type of application, and overall workload characteristics can vary widely even in what appear to be otherwise homogenous environments. These components and the numerous interactions between them can affect performance in ways that are not yet fully understood [Huber10]. Specific challenges with performance analysis of virtualized environments are hypervisor differences, resource interdependence, workload variability, and disparity in modeling. The following section examines hypervisor differences.

2.1 Hypervisor Differences

Every hypervisor is different by design. A model of one system with a given hypervisor may not accurately model another system with a different hypervisor. Many research groups use the Xen hypervisor because it is open source and supports both full virtualization and para-virtualization. That said, Xen is not currently the most feature-rich hypervisor platform available on the market.

VMware's proprietary ESXi hypervisor, which serves as a cornerstone for their vSphere cloud computing platform, provides a host of capabilities not currently available with any other hypervisors. These capabilities include High Availability (the ability to recover virtual machines quickly in the event of a physical server failure), Distributed Resource Scheduling (automated load balancing across a cluster of ESXi servers), Distributed Power Management (automated decommissioning of unneeded servers during non-peak periods), Fault Tolerance (zero-downtime services even in the event of hardware failure), and Site Recovery Manager (the ability to automatically recover virtual environments in a different physical location should an entire datacenter outage occur) [Hostway11]. Thus, even though a significant portion of academic research is done using open source hypervisors such as Xen, commercially VMware has a larger share of the commercial cloud computing market than all of its competitors combined. Now that we have discussed differences in hypervisors, the following section looks at how resource interdependence can complicate performance modeling.

2.2 Resource Interdependence

Performance analysis in computing focuses on one or more of the four basic computing resources: CPU, memory, storage, and networking. When analyzing a physical computer, it is important to recognize the way these resources interact in order to determine true performance characteristics. For example, a system experiencing high CPU or storage bandwidth utilization may in fact be suffering from a lack of enough physical memory. Thus, evaluating any of the four resources independent of the other three can result in incorrect conclusions.

This situation is complicated by an order of magnitude in a virtualized environment, which is made up of many virtual machines sharing resources on a single a physical machine. These virtual machines have the same interplay of resources as a physical machine, but in addition compete for each of these resources against each other [Tickoo09]. Based on their configurations, certain virtual machines may be given preferential treatment for certain physical resources while losing out when competing for others. In addition to these issues arising from resource interdependence, the next section examines how workload variability also adds complexity to performance analysis.

2.3 Workload Variability

Workload is a term that encapsulates the overall usage of computing resources by an application. In a virtual environment this generally includes entire groups of virtual machines and their individual applications in aggregate. When evaluating workloads, it is important to remember that the same application, or set of

applications, can display different performance characteristics based on a number of variables, including the way the application was configured, the number of users of the system, how the application is specifically being used, and the nature of the underlying hypervisor and hardware platform. Small changes in any of these factors can produce results that vary widely from one research project to another. The following section describes the three dimensions of virtualization modeling.

2.4 Three Dimensions of Virtualization Modeling

Virtualization can be evaluated in one of three dimensions. In the first dimension, we evaluate the performance of a single virtual machine running on a single host to evaluate the effect of the hypervisor on performance without the possibility for resource contention. In the second dimension, we evaluate multiple virtual machines running on a single physical host to evaluate performance in an environment where contention for resources may be a factor. In the third dimension, we evaluate the performance of multiple virtual machines running across multiple physical hosts, which most closely models the desired scenarios in cloud computing. These various dimensions lead to the next section, which explores differences in modeling frameworks.

2.5 Modeling Framework Differences

Given this background of virtualization modeling, there are numerous possibilities of various virtualization frameworks. Research groups have developed models representing various virtual environments, such as enterprise, small business, hosting, and cluster [Jang07]. However, even after defining an assortment of systems it is unlikely, for example, that a given server consolidation model can accurately model the full gamut of possible objectives, configurations, and workloads for all server platforms. To further complicate matters, most of the available research focuses on performance analysis in a two-dimensional framework whereas in commercial applications, performance of virtual machines across clusters of physical hosts is the more salient issue.

Resource allocations, workloads and system behaviors can fluctuate widely, and performance can further vary based on system configurations. Despite the assorted issues with performance analysis of virtualized environments, researchers have made much headway in recent years developing tools and techniques for measuring, modeling and simulating virtualized environments. The next section examines some of the newly developed benchmarks and metrics that can be successfully used in virtualized systems.

3. Virtualization Measurement Tools

Various research groups have developed performance metrics and benchmarks specifically to measure the performance of virtualized systems. In some cases, research groups have been able to use custom implementations of existing benchmarks when analyzing virtualized systems. For example, a particular implementation of SPECweb2005 benchmark, which is the standard benchmark used to determine the performance of Web servers, can be successfully used in performance testing of virtualized applications [Ardagna10]. However, most benchmarks that were originally developed for physical machines cannot be used for reliable performance testing of virtualized environments due to the diverse and complex nature of these virtualized systems. The following section provides a brief overview of some benchmarks and metrics specially developed to quantify the performance aspects of virtualized environments.

3.1 Benchvm benchmark

Benchvm is an open source benchmark suite designed to compare virtualization systems. Benchvm was developed during qualitative analysis testing of Xen and KVM, and was specifically created to automate

testing and to facilitate repeatable testing [Deshane08]. Benchvm assesses overall performance of virtualized environments, as well as scalability and performance isolation measures. Scalability measures reflect how well a given system can add virtual machines without downgrading performance. Performance isolation measures refer to how well one guest is shielded from the resource consumption of other guests in the virtualized environment. The next section introduces another measurement tool, VMmark.

3.2 VMmark benchmark

VMmark, is a tile-based workload developed by VMware. In this case, a tile refers to a workload unit, and VMmark is comprised of many different workloads concurrently running on individual virtual machines. The collection of workloads are representative of workloads found in a typical datacenter, such as email server, file server, database, web server, Java, as well as Standby, which is an extra virtual machine with no workload [McDougall10]. While the workloads run on the virtual machines, a score is recorded to reflect the ability of the system to consolidate services and run various workloads on multiple virtual machines based on the number of tiles running and their performance [VMware11]. The following section describes another tile-based benchmark, vConsolidate.

3.3 vConsolidate benchmark

vConsolidate is another server virtualization benchmark whose workload is ideal for a server consolidation framework [Apparao08]. Like VMmark, vConsolidate is a tile-based workload and its aim is to emulate server consolidation workloads. It is comprised of SPECjbb, a Java server; Sysbench, a database server; Webbench, a web server; and a mail server, running concurrently on individual virtual machines on a server platform [Tickoo09]. These multiple workloads of vConsolidate are optimal for server consolidation models, which can then be used to model the performance of each workload in various virtualization scenarios to analyze resource contention between virtual machines. The next section looks at the benchmark suite VSCBenchmark.

3.4 VSCBenchmark

VSCBenchmark is a server virtualization benchmark designed to run dynamic workloads on multiple virtual machines. VSCBenchmark consists of three parts, the VSCB Control Center and the VSCB Server, both of which run on the Control VM, and the VSCB Workload Client, which runs on the Workload VM [Jun08]. These components generate various workloads, and the Workload VM is created and killed as configured by the user via the VSCB Control Center. Many scenarios can be configured and tested with this benchmark, providing a dynamic performance testing environment. VSCBenchmark was developed after the advent of VMmark and vConsolidate, as a means to effectively assess the dynamic performance of virtual machines and extensively evaluate resource contention. The following section covers the last measurement tool, performance metric M.

3.5 Performance Metric M

Performance metric M can be used to compare the performance of virtualized versus non-virtualized environments. Performance metric M "â€¦represents the CPU waste rate of virtualized environment, where the performance of application running on each VM is consolidated and normalizedâ€▦ [Jang07]. Performance metric M is equivalent to the wasted capacity of a system over the throughput of the applications running on the system. The waste rates of the virtualized and non-virtualized systems are normalized by comparing the two systems at the same performance level.

The benchmarks and metrics discussed previously were distinctively designed as thorough measurement tools

for virtualized environments. Benchvm can be used to compare virtualized systems and hypervisors, VMmark and vConsolidate are both tile-based workloads ideal for consolidated server performance analysis, VSCBenchmark is designed for dynamic performance testing, and performance metric M is used to compare the performance of virtualized and non-virtualized systems. The next section focuses on some of the performance models and simulation frameworks developed with the complexities of virtualization in mind.

4. Modeling and Simulation in Virtualized Environments

Many challenges ensue as researchers attempt to accurately model the complexities of virtualized environments in order to assess performance issues. Virtualization adds levels of complexity as the virtual machines interact not only with the underlying physical machine but also with each other [Iyer09]. Models that have been developed for single physical machines and for non-virtualized environments cannot account for the complexities and variabilities of virtualized environments [Watson10]. For example, regression techniques that have been used to model application performance in non-virtualized systems are not necessarily suited for predicting application performance in virtualized systems [Kundu10]. The following sections provide a concise review of several modeling and simulation techniques developed to assess performance of virtualized environments. The first modeling framework explored in this paper is the Linear Parameter Varying model.

4.1 Linear Parameter Varying (LPV)

Black box models based on Linear Parameter Varying, or LPV, framework have been used for performance modeling of virtualized web hosting systems in order to predict the performance of virtual machines in the hosting environments. The LPV models allow non-linear systems to be treated as linear systems with varying parameters. A dynamic model was developed based on the LPV framework that represents the correlation between the allocation of resources to the virtual machines and the virtualized serverâ€™s response time. These models were validated by assessing the performance of Web service applications in virtualized environments [Ardagna10]. The next section discusses an alternative model useful for virtualized systems, the fuzzy model.

4.2 Fuzzy Models

Another method for modeling non-linear systems, such as virtualized environments, is fuzzy models. Fuzzy logic is used to develop fuzzy models that characterize the relationship between application performance and resource allocation. Unlike Boolean logic, or â€œcrispâ€ logic where values are either 0 or 1, fuzzy logic allows for approximate values to exist in the continuum between 0 and 1. Fuzzy logic also includes rule based, or if-then logic. A fuzzy model is based on fuzzy rules, and represents the relationship between system input and system output. The fuzzy if-then rules describe system behavior, and these rules can be used to produce linear equations representing this behavior. Fuzzy models have been developed and used to determine resource demands based on the relationship between workloads, resource needs, and desired application performance [Xu07]. The following section introduces models based on artificial neural networks.

4.3 Artificial Neural Networks (ANN)

Artificial neural networks, or ANN, are statistical systems patterned after biological neural networks. Using artificial neurons, or nodes, these networks can be used to model non-linear systems. A specific implementation of an ANN based model has been used to predict the performance of applications in virtualized environments at a given level of allocated resources. In order to accomplish this, the models first had to undergo an iterative training process, and the training data set was then followed by a testing data set. A variety of benchmarks have been tested using this model with low modeling error. An ANN based model

â€œcan quickly adapt to changes in application and system behaviors, and it can be used at any time by data center managers to predict application performance at desired parameter configuration of the VMâ€ [Kundu10]. The next section talks about another model suited for virtualization, the probabilistic performance model.

4.4 Probabilistic Performance Model

Probabilistic performance modeling models the â€œprobability distribution of an application performance metric conditioned on one or more variables that we can measure or control, such as system resource utilization and allocation metrics. We consider the influence of hidden factors that we cannot readily measure to be implicit in the probability distribution of performanceâ€ [Watson10]. In applying the probabilistic performance model to virtualized environments, this model can be used to assess the performance of applications and evaluate the effects of resource allocation and CPU contention. This is accomplished by taking a given application and workload, and then adjusting the allocation of CPU resources. In addition to these models which have been specifically developed for virtualized environments, the next section describes a simulation framework created for cloud computing.

4.5 CloudSim

CloudSim is â€œâ€¦a new generalized and extensible simulation framework that enables seamless modeling, simulation, and experimentation of emerging Cloud computing infrastructures and management servicesâ€ [Calheiros09]. In past years, Grid simulators, which simulate Grid computing environments, have been used as simulation frameworks for high performance and data-intensive systems. Grid computing can be defined as a distributed computing system in which workloads that do not directly interact are executed across relatively low-power, heterogeneous computing systems [Business11]. While Grid computing is useful for distributing mathematical and analytical tasks, Grid computing (and thus Grid simulators) are not able to support the complex, multi-tiered interactions between virtualized systems that are associated with cloud computing environments.

The CloudSim framework consists of a datacenter node as well as a virtualization engine that can create virtual machines on this node. CloudSim is capable of simulating the advanced infrastructures of cloud computing such as Cloud brokers and Cloud exchange, which model real-time interactions between service providers and clients. Developers can utilize this simulation framework to determine performance bottlenecks relating to service delivery and provisioning policies. CloudSim is a customizable and extensible framework, and leaves much possibility for future development and increased potential [Calheiros09].

CloudSim, as well as the non-linear models mentioned in previous sections (LPV, fuzzy, ANN, and probabilistic), have specifically been developed to overcome the complex performance modeling challenges and accurately model virtualized environments. Models based on LPV framework can characterize Web service application performance, and fuzzy models can be used to estimate resource demands. ANN based models and probabilistic performance models can both be used to predict application performance. CloudSim can be used to determine performance bottlenecks before live deployment.

5. Summary

Virtualization has become more attractive in recent years as cost controlling measures are sought out in enterprise environments. In an effort to improve performance of these virtualized systems, accurate performance measures, modeling, and simulation are needed. This paper explored the multifaceted issues in performance modeling and analysis in resource management of virtualized environments. These challenges include hypervisor differences, resource interdependence, workload variability, and diversity in modeling.

This paper also provided a survey of new benchmarks uniquely created for virtualization, including benchvm, VMmark, vConsolidate, VSCBenchmark, and performance metric M. Several modeling techniques have been specifically developed for performance analysis of virtualized environments, including those based on LPV, fuzzy logic, ANN, and probability distribution. Finally, CloudSim, a new simulation framework built for virtual cloud computing infrastructures, was described. Research efforts continue, and researchers are working to further develop and expand upon measurement tools, models, and simulations especially for virtualized environments [Tickoo09] [Apparao08] [Watson10] [Kundu10] [Huber10].

References

1. [Tickoo09] Tickoo, Omesh; Iyer, Ravi; Illikkal, Ramesh; Newell, Don; "Modeling Virtual Machine Performance: Challenges and Approaches", ACM SIGMETRICS Performance Evaluation Review, Volume 37 Issue 3, December 2009, http://portal.acm.org/citation.cfm?doid=1710115.1710126, Articulates the issues related to modeling the performance of virtual machines and suggests using VPA core and cache usage estimation for performance modeling of virtual machines.

2. [Ardagna10] Ardagna, Damilo; Tanelli, Mara; Lovera, Marco; Zhang, Li; "Black-box Performance Models for Virtualized Web Service Applications", WOSP/SIPEW '10 Proceedings of the First Joint WOSP/SIPEW International Conference on Performance Engineering, ACM, 2010, http://doi.acm.org/10.1145/1712605.1712630, Suggests Linear Parameter Varying (LPV) models for performance analysis of web service applications in virtualized environments.

3. [Calheiros09] Calheiros, Rodrigo N.; Ranjan, Rajiv; De Rose, Cesar A. F.; Buyya, Rajkumar; "CloudSim: A Novel Framework for Modeling and Simulation of Cloud Computing Infrastructures and Services", CloudSim ICCP 2009, http://www.gridbus.org/reports/CloudSim-ICPP2009.pdf, Introduces and emphasizes the benefits of CloudSim, a new customizable modeling and simulation tool developed specifically for cloud computing.

4. [Watson10] Watson, Brian J.; Marwah, Manish; Gmach, Daniel; Chen, Yuan; Arlitt, Martin; Wang, Zhikui; "Probabilistic Performance Modeling of Virtualized Resource Allocation", ICAC '10 Proceeding of the 7th International Conference on Autonomic Computing, pp.99-108, ACM, 2010, http://doi.acm.org/10.1145/1809049.1809067, Proposes a model for application performance in a virtualized system based on the probability distributions of performance metrics.

5. [Kundu10] Kundu, Sajib; Rangaswami, Raju; Dutta, Kaushik; Zhao, Ming; "Application Performance Modeling in a Virtualized Environment", High Performance Computer Architecture (HPCA), 2010 IEEE 16th International Symposium on , pp.1-14 Jan. 2010, http://dx.doi.org/10.1109/HPCA.2010.5463058, Implemented a particular execution of artificial neural network (ANN) model to predict the performance of applications running on virtualized systems.

6. [Apparao08] Apparao, Padma; Iyer, Ravi; Newell, Don; "Towards Modeling & Analysis of Consolidated CMP Servers", ACM SIGARCH Computer Architecture News, Volume 36 Issue 2, May 2008, pp.38-45, http://doi.acm.org/10.1145/1399972.1399980, Presents a consolidation performance model for the performance analysis of consolidated servers, and utilizes the benchmark vConsolidate in a case study illustrating this model's potential.

7. [Jang07] Jang, Jiyong; Han, Saeyoung; Kim, Jinseok; Park, Sungyong; Bae, Seungjo; Choon Woo, Young; "A Performance Evaluation Methodology in Virtual Environments", Computer and Information Technology, International Conference on, pp. 351-358, 7th IEEE International Conference on Computer and Information Technology, 2007, http://www.computer.org/portal/web/csdl/doi/10.1109/CIT.2007.179, Defines four performance models representing different virtualized systems, introduces a new performance metric, and use one model and performance metric M to evaluate the performance of virtualized versus non-virtualized environments.

8. [Hostway11] VMware ESXi Cloud Simplified, Hostway UK, http://www.hostway.co.uk/small-business/dedicated-hosting/cloud/vmware-esxi.php, Comprehensive explanation of the features and benefits of VMware ESXi hypervisor.

9. [GuestOS11] "Guest OS Install Guide", VMware Community Page, http://blogs.vmware.com/guestosguide/2009/09/vmi-retirement.html, Web page announcing that VMware will no longer support their Virtual Machine Interface (VMI) technology.

10. [Huber10] Huber, Nikolaus; Von Quast, Marcel; Brosig, Fabian; Kounev, Samuel; "Analysis of the Performance-Influencing Factors of Virtualization Platforms", On the Move to Meaningful Internet Systems, OTM 2010, Springer-Verlag Berlin, 2010, pp. 811-828; http://dx.doi.org/10.1007/978-3-642-16949-6_10, Offers a benchmark based approach to predict the performance of a Xen virtualized environment.

11. [Xu07] Xu, Jing; Zhao, Ming; Fortes, Jose; Carpenter, Robert; Yousif, Mazin; "On the Use of Fuzzy Modeling in Virtualized Data Center Management", Proceedings of the Fourth International Conference on Autonomic Computing (ICAC), IEEE Computer Society, p. 25, June 2007, http://portal.acm.org/citation.cfm?id=1270385.1270747, Proposes a local and global system based on fuzzy logic for the management of virtualized resources.

12. [Lu06] Lu, Jie; Makhlis, Lev; Chen, Jianjiun; "Measuring and Modeling the Performance of the XEN VMM", International CMG Conference 2006, pp. 621-628, http://svn.assembla.com/svn/biSTgsRbOr3y0wab7jnrAJ/trunk/artikels/Measuring_and_Modeling_the_Performance_of_XEN_VMM.pdf, Presents Xen as a server virtualization option, discusses why traditional modeling methods will not work with a virtualized system, and suggests new modeling techniques.

13. [Iyer09] Iyer, Ravi; Illikkal, Ramesh; Tickoo, Omesh; Zhao, Li; Apparao, Padma; Newell, Don; "VM3: Measuring, Modeling and Managing VM Shared Resources", Computer Networks: The International Journal of Computer and Telecommunications Networking, Volume 53 Issue 17, December, 2009, http://dx.doi.org/10.1016/j.comnet.2009.04.015, Models virtual machine performance on a consolidated chip-multiprocessor platform (CMP), measuring the effects of server consolidation with the benchmark vConsolidate, and discusses methods to manage shared resources.

14. [Jun08] Jun, Hai; Cao, Wenzhi; Yuan, Pingpeng; Xie, Xia; "VSCBenchmark: Benchmark for Dynamic Server Performance of Virtualization Technology", IFMT '08, Proceedings of the First International Forum on Next-generation Multicore/Manycore Technologies, ACM, 2008, pp. 5:1-5:8, http://doi.acm.org/10.1145/1463768.1463775, Discusses the VSCBenchmark for analyzing server consolidation and compares this benchmark with vConsolidate and VMmark benchmarks.

15. [VMware11] "Features of VMmark, Virtualization Benchmark", VMware Product Page, http://www.vmware.com/products/vmmark/features.html, Web page listing the features and benefits of VMmark, a tile-based benchmark.

16. [Deshane08] Deshane, Todd; Shepherd, Zachary; Matthews, Jeanna N.; Ben-Yehuda, Muli; Rao, Balaji; Shah, Amit; "Quantitative Comparison of Xen and KVM", Xen Summit, Boston, MA, June 23–24, 2008, http://www.todddeshane.net/research/Xen_versus_KVM_20080623.pdf, Developed benchvm, a virtualization benchmarking suite, and used this benchmark to compare two hypervisors, Xen and KVM.

17. [McDougall10] McDougall, Richard; Anderson, Jennifer; "Virtualization Performance: Perspectives and Challenges Ahead", ACM SIGOPS Operating Systems Review, Volume 44, Issue 4, December 2010, pp.40-56, http://doi.acm.org/10.1145/1899928.1899933, Discusses the performance issues of virtualized systems with an emphasis on VMware's vSphere virtualization platform.

18. [Benevenuto06] Benevenuto, Fabrício; Fernandes, César; Santos, Matheus; Almeida, Virgílio; Almeida, Jussara; Janakiraman, G. (John); Santos, José Renato; "Performance Models for Virtualized Applications", Lecture Notes in Computer Science, Volume 4331, Frontiers of High Performance Computing and Networking – ISPA 2006 Workshops pp.427-439, 2006, http://dx.doi.org/10.1007/11942634_45, Discusses methodology for building models for performance prediction of applications migrated from a non-virtualized system to a virtualized environment.

19. [vCloud11] Cloud Computing Services with VMware Virtualization Cloud Infrastructure, http://www.vmware.com/solutions/cloud-computing/index.html, Provides an overview of vCloud, VMware's approach to cloud computing.

20. [VMware10] â€œUnderstanding Memory Resource Management in VMware ESX 4.1â€, VMware, Inc. Performance Study, 2010, http://www.vmware.com/files/pdf/techpaper /vsp_41_perf_memory_mgmt.pdf, This published performance study describes the basic memory management concepts in ESX, the configuration options available, and provides results to show the performance impact of these options.

21. [Virtualization06] â€œVirtualization Overview White Paperâ€, VMware, 2006, http://www.vmware.com/pdf/virtualization.pdf, Overview and definitions of virtualization concepts and terms including para-virtualization.

22. [Wiki11] Cloud Computing, Wikipedia, http://en.wikipedia.org/wiki/Cloud_computing, Overview of cloud computing concepts.

23. [Business11] What is Grid-Computing? Definition and Meaning, Business Dictionary, http://www.businessdictionary.com/definition/grid-computing.html, Provides a definition of grid-computing.

List of Acronyms

Acronym	Definition
VM	Virtual Machine
VMM	Virtual Machine Monitor
ANN	Artificial Neural Networks
LPV	Linear Parameter Varying
KVM	Kernel-based Virtual Machine

Last modified on April 24, 2011
This and other papers on latest advances in performance analysis are available on line at
http://www1.cse.wustl.edu/~jain/cse567-11/index.html
Back to Raj Jain's Home Page