PERFORMANCE OF HYPERVISORS IN THE CLOUD BASED ON INFORMATIION USING SIGAR FRAMEWORK

Gopathoti Sujanarao[1],Kamuluri Murali2
1: Asst Professor Department of CSE,Princeton College of Engineering and Technology,TS,India
2: Asst Professor , Department of CSE,DBS Institute of Technology, Kavali, AP, India

Abstract

Hypervisors using virtualization technology enable multiple operating systems to run on one physical server. Cloud computing model is less expensive because it streamlines the delivery of services by providing a stage for optimizing complex IT resources in a scalable manner with the help of virtualization technology and hypervisors. Selecting a suitable hypervisor for their organization's private cloud is herculean task for modern CIOs. Hypervisor vendors do claim that they have negated virtualization overhead completely compare to native system, but still there exists minute virtualization overhead because virtual machines have to communicate with middle-layer hypervisor to access the underlying physical hardware and also there is an impact of other virtual machines running on the same hypervisor. Hypervisors are developed using different virtualization techniques like full virtualization, para-virtualization and hybrid model virtualization. This paper evaluates the performance of three hypervisors ESXi, XenServer and KVM using SIGAR framework for system information and Passmark for system workloads in the private cloud environment. Private cloud has been designed using open source cloud computing software CloudStack. Hypervisors are deployed as hosts in the CloudStack. This paper recommends best suited hypervisors for respective workloads in the private cloud based on the performance of system information and system workloads.

1. Introduction

Cloud computing as a model enables on demand access to servers, networks, applications and provides the option to pay as you use manner [1]. The major benefits of cloud computing are flexible and scalable infrastructures, reduced implementation and maintenance costs, IT department transformation and increased availability of high performance applications.

Cloud computing model encourages availability and is composed of four deployment models. In which, Private Clouds are deployed behind the firewall of a company and the cloud infrastructure is operated solely for an organization. Private cloud deployment model creates proprietary computing architecture behind a firewall with full control over infrastructure. This paper uses private cloud model for experiment.

Virtualization is a technology that combines or divides computing resources to present many operating environments using methodologies like hardware and software partitioning, machine simulation, emulation, timesharing and etc., [2].

2. Hypervisor Models

All three hypervisors that are used in the experiment are briefly described along with their virtualization technique.

2.1. Paravirtualized Hypervisor

XenServer- Citrix XenServer is server virtualization platform built on the Xen Hypervisor. Xen [8] uses para-virtualization technique. Para-virtualization modifies the guest operating system. XenServer is a virtual infrastructure solution that gives the flexibility of management console, and the tools needed to move applications, desktops, and servers from a physical to a virtual environment [9]. XenServer hypervisor claims that it completely negates virtualization overhead gives near native application performance.

2.2. Full virtualized Hypervisor

ESXi Server **-** VMware ESXi is a Hypervisor designed for server virtualization environments capable of live migration of Virtual Machine (VM) using VM motion. VMware ESXi supports full virtualization [10]. The hypervisor implements shadow versions of system structures such as page tables and maintain consistency with the virtual Tables by trapping every instruction that attempts to update these structures. Therefore, an extra level of mapping is in the page Table. The virtual pages are mapped to physical pages throughout the guest operating system‛s page Table [11]. TheHypervisor then translates the physical page to the machine page, which ultimately is the right page in physical memory. This helps the ESXi server to manage the overall memory and enhance the overall system performance [12].

3. Experiment Design – Private Cloud: CloudStack with Hypervisors

The experimental design contains private cloud infrastructure created using CloudStack. CloudStack is an Infrastructure as a service (IaaS) cloud based software which builds private cloud environments. CloudStack supports multiple hypervisors. CloudStack has the ability to build cloud environments with different hypervisors with web interface for users and administrators. CloudStack is open source software written in java that is designed to deploy and manage large networks of virtual machines as a highly available, scalable cloud computing platform. CloudStack provides a web interface, command line to manage the cloud environment.

Two machines are needed to implement private cloud using CloudStack. One machine is Management Server, runs on a dedicated server or a VM. It controls allocation of virtual machines to hosts and assigns storage and IP addresses to the virtual machine instances. The Management Server runs in a Tomcat container and requires a MySQL database for persistence. In the experiment, Management Server (a Virtual Machine with hardware configuration of 4GB RAM and 100GB hard disk) is installed on Ubuntu (12.04 64-bit) operating system. Second machine is the host machine where hypervisors are installed on a bare metal with hardware configuration of AMD FX 8150 – 8 Core 3.6 GHz processor, 32 GB RAM, 1 TB hard disk and 2 NICs for the test environment. Front end will be any base machine to launch CloudStack UI using web interface (with any browser software IE, Google Chrome) to provision the cloud infrastructure by creating zone, pod, cluster and host in the sequential order.

In our test environment XenServer 6.0, ESXi 4.1 and KVM (Ubuntu 12.04) hypervisors are deployed as hosts and virtual machine (VM1- Windows 2008 R2) is installed on all three hypervisors in the private cloud as depicted in Figure 1.

The virtual machine (VM1) Windows 2008 R2 is installed on each hypervisors and system information performance is gathered using SIGAR and system workloads performance evaluated in detail using Passmark. After the Windows VM is installed on all three hypervisors, CPU, Memory, Disk I/O and Network performances are measured using SIGAR Framework. SIGAR (System Information Gatherer and Reporter) is a platform independent tool for accessing system level information in Java and other programming languages. In the experiment, Java program has written to gather system information using SIGAR API by deploying sigar-amd64-winnt.dll for Windows.

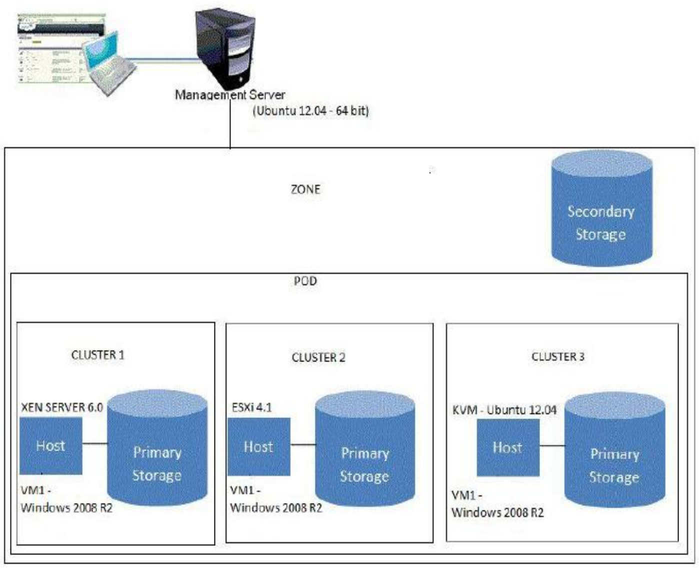

Figure 1. Experiment Design – Private Cloud (CloudStack with Multiple

4. Results

This section provides the detailed results of all the performance tests which are executed on three hypervisors using SIGAR API and Passmark. Detailed results are outcome of each of the benchmarks run. All of the results have been normalized to native performance measures. Native performance is normalized at 1.0 and all other various benchmark results are shown relative to that number. Hence benchmark results of 90% of the native performance would be shown as 0.9 on the scale in the graph. Higher numbers indicate better performance of the particular virtualization platform, unless indicated otherwise. Near-native performance also indicates that more virtual machines can be deployed on a single physical server, resulting in higher consolidation ratios.

4.1. Sigar

Available CPU details are captured through java program using SIGAR API on the virtual machine operating system for each hypervisor. CPU availability close to native indicates the better performance for a hypervisor. The below Table 1 indicates available CPU performances for native and respective hypervisors.

Table 1. Available CPU of Hypervisors

	Available CPU (%)
Native	97
ESXi 4.1	97
Xen 6.0	96.5
KVM(Ubuntu 12.04)	93.5
Native	1
ESXi 4.1	1
Xen 6.0	0.99
KVM(Ubuntu 12.04)	0.96

ESXi with windows virtual machine (VM) as guest operating system shows available CPU equal to native without any virtualization overhead. XenServer with windows VM shows 1% available CPU overhead compare to native. KVM with windows VM shows 4% available CPU overhead compare to native.

5. Conclusion

The intent of this paper is to evaluate the performance of three hypervisors, VMWareESXi Server, XenServer and KVM for system information using SIGAR and for system workloads using Passmark in the private cloud environment. Virtual machine (VM) windows 2008 R2 is deployed on each hypervisor in the private cloud. CloudStack is used to create a private cloud. Once entire experiment setup is ready, system information is gathered using SIGAR API to compare the performance of three hypervisors. Among three hypervisors, for system information, VMWare'sESXi shows better performance in available CPU, available memory, disk I/O and network performance compare to other two hypervisors. KVM needs to improve in all four system resources performance point of view. For system workloads Passmark is used to evaluate three hypervisors performance. Among three hypervisors, for system workloads, VMWare'sESXi shows better performance in CPU mark, and network performance compare to other two hypervisors. XenServer shows better performance in memory mark, and disk I/O performance compare to other two hypervisors. KVM needs to improve in all four system resources performance from system workloads point of view. System workload environment and hardware configuration is same for all three hypervisors hence the hypervisors which are lacking in the respective system resources performance need an improvement. Hypervisor vendors may concentrate on the loopholes exhibited in respective performance tests and may improve their products to mitigate virtualization overhead which are captured at system information level. This system information indicates the gap which should be filled for improvement of hypervisors.

References

[1] P. Mell and T. Grance, "The NIST Definition of Cloud Computing", National Institute of Standards and Technology, Information Technology Laboratory, Version 15, (2009) October 7.

[2] S. Nanda and T. Chiueh, "A Survey on Virtualization Technologies", Technical report, Department of Computer Science, SUNY at Stony Brook, New York, (2005), pp. 11794-4400.

[3] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg and I. Brandic, "Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility", in: Future Generation Computer Systems, Elsevier B. V., (2009).

[4] "VMware Understanding Full Virtualization, Paravirtualization and Hardware Assist. VMware", white paper, (2007) November 10.

[5] "SIGAR", [Online] https://support.hyperic.com/display/SIGAR/Home.

[6] "Passmark", [Online] http://www.passmark.com/products/pt.htm.

[7] "CloudStack", [Online] http://cloudstack.apache.org.

[8] "Xen—How does Xen work", Xen Organization, (2009).

[9] "Fujitsu Technology Solutions", Data Sheet Citrix Xen Server.

[10] "Hostway UK VMware ESXi Cloud Simplified", Comprehensive explanation of the features and benefits of VMware ESXi Hypervisor.

[11] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt and A. Warfield, "Xen and the art of virtualization", Proceedings of the Nineteenth ACM Symposium on Operating systems Principles, ACM Press, New York, (2003), pp. 164–177.

[12] "VMware, —The Architecture of VMware ESXi", white paper, (2007).

[13] J. Che, Q. He, Q. Gao and D. Huang, "Performance Measuring and Comparing of Virtual Machine Monitors", College of Computer Science, Zhejiang University, Hangzhou 310027, China, IEEE/IFIP International Conference on Embedded and Ubiquitous Computing, (2008).

[14] "VMware (2007) A Performance Comparison of Hypervisors VMware", White paper, (2007) February 1.

[15] "XenSource (2007) A Performance Comparison of Commercial Hypervisors. XenEnterprise vs. ESX Benchmark Results", XenSource, (2007).

[16] "Standard Performance Evaluation Corporation (SPEC)", [Online] http://www.spec.org/cpu2000/.

[17] "FUJITSU", Benchmark Overview-vServCon, white paper, (2010) March.

[18] P. Apparao, S. Makineni and D. Newell, "Virtualization (2006) Characterization of network processing overheads in Xen", Technology in Distributed Computing, VTDC, (2006).

[19] C. Jianhua, H. Qinming, G. Qinghua and H. Dawei, "Performance Measuring and Comparing of Virtual Machine Monitors", Embedded and Ubiquitous Computing, EUC '08, (2008).

[20] A. Menon, et. al., "Diagnosing Performance Overheads in the Xen Virtual Machine Environment", Conference on Virtual Execution Environments (VEE'05), (2005).

[21] Z. Shan and H. Qinfen, "Network I/O Path Analysis in the Kernel-based Virtual Machine Environment through Tracing", Information Science and Engineering (ICISE), (200p).

[22] K. T. Moller, "Virtual Machine Benchmarking", Diploma Thesis, Karlsruhe Institute of Technology, (2007).