2011 Third International Conference on Communications and Mobile Computing

Power Consumption of Virtual Machine Live Migration in Clouds

Qiang Huang, Fengqian Gao, Rui Wang, Zhengwei Qi
School of Software, Shanghai Jiaotong University, Shanghai, China 200240
Email: {johnny1209, myzone906, wangzi584, qizhwei}@sjtu.edu.cn

***Abstract*—Virtualization Technology has been employed increasingly widely in modern data centers in order to improve its energy efficiency. In particular, the capability of virtual machine (VM) migration brings multiple benefits for such as resources (CPU, memory, et al.) distribution, energy aware consolidation. However, the migration of virtual machines itself brings extra power consumption. For this reason, a better understanding of its effect on system power consumption is highly desirable. In this paper, we present a power consumption evaluation on the effects of live migration of VMs. Results show that the power overhead of migration is much less in the scenario of employing the strategy of consolidation than the regular deployment without using consolidation. Our results are based on the typical physical server, the power of which is linear model of CPU utilization percentage.**

***Index Terms*—data center; power consumption; live migration;**

I. Introduction

In recent years, more and more data centers start to employ server virtualization strategies for resource sharing to reduce hardware and operating costs. Virtualization technologies (such as Xen[1], VMware[2], and Microsoft Virtual Servers[3]) can consolidate applications previously running on multiple physical servers onto a single physical server, via this way, the energy consumption of data center can be effectively reduced. Consequently, virtualized infrastructures are considered as a key solution to the power management of data center.

Perhaps the biggest advantage of employing virtualization technology is the ability to flexibly remap physical resources to virtual servers. virtual machine manager such as Xen can distribute the amount of physical resource (CPU, memory, disk, et al.) to the virtual servers above it. And using VMs migration technology enables the consolidation of servers spread across many locations. The hotpot in the datacenter can be handled by simply migrating the virtual server to a less loaded physical server. If QoS performance can be maintained in the consolidation, a system can be configured with a fewer number of servers and less power consumption[4]. In addition, the system can be configured so that it can handle the maximum load expected. In this paper, we focus on the evaluation on power cost of "live" or "hot" migration, which allows migrating an OS as it continues to run, as opposed to "pure stop-and-copy" or "cold" migration, which involves halting the VM, copying all its memory pages to the destination host and then restarting the new VM. The main advantage of live migration is the possibility to migrate an virtual machine with near-zero downtime, an important feature when live services are been served[5].

A. Our Contribution

In this paper, first we give a practical experimental approach to evaluate the power consumption of VM migration. And then we quantify the power cost of VM migration both for the original physical server that starts the migration and the destination physical server that accepts the transfer. Our results show that the power influence of migration to the original server decreases when the CPU usage of the migrated VM increases, but to the destination server, the influence is stable. Additionally, the time cost of migration is not impacted by the CPU usage of VM. Our study would aid researchers and practitioners currently evaluating the application of VM migration for consolidation strategy in clouds.

The rest of this paper is organized as follows: Section 2 reviews related work; And after introducing the background of the power management and live migration in data center in Section 3, we describe our objectives, experimental setup, and results in Section 4, Conclusions and future work are reviewed in Section 5.

II. Related Work

The advent of innovative technologies, such as paravirtualization[1], hardware-assisted virtualization[6] and live migration[5], have contributed to an increasing adoption of virtualization on server system. At the same time, the impact of virtualization in a variety of scenarios has been the focus of considerable attention. Live migration as a main strategy for applications consolidation, its impact has been researched in a few of studies.

Zhao & Figueiredo [7] and William Voorsluys et al.[4] specifically deal with VM migration. The former analyzes performance degradation when migrating CPU and memory intensive workloads as well as migrating multiple VMs at the time; however such study employs a pure stop-and-copy migration approach rather live migration. The latter evaluates the performance cost of virtual machine live migration in clouds, and shows that in most case, migration overhead is acceptable but cannot be disregarded, especially in systems where service availability and responsiveness are governed by strict Service Level Agreements (SLAs), However their study did not consider power consumption. Shekhar Srikantaiah et

978-0-7695-4357-4/11 $26.00 © 2011 IEEE
DOI 10.1109/CMC.2011.62

al.[8] studied the the energy performance trade-offs for consolidation of applications, but without quantifying the impact of VM live migration. The study presented by Takayuki Imada et al.[9] investigates power and QoS(Quality of Service) performance characteristics of virtual servers with virtual machine technology. They found that the live migration scheme can be applied with slight QoS performance degradation and slight increased power consumption; however such a study is based on a single benchmark workload, without given a quantified study on the virtual machine power impact and they only study the original server, without considering the migration as a system.

III. Background

In this section, we first present the modern power management strategies in data center, and then we review the technology of Live Migration of virtual machines.

A. Power Management in Data Center

power consumption of Data center is undergoing alarming growth. The EPA[10] estimates that, By 2011, U.S. data centers will cost 100 billion kWh at a cost of $7.4 billion per year, so the power management in data centers has become a critical issue in most countries. Many efforts have been made to improve the energy efficiency of data center, such as network power management, chip-Multiprocessing (CMP) energy efficiency, power capping, storage power management solutions etc.[11]. Generally, the modern approach to solve the problem is employing the virtualized technology, which enables multiple OS environments to coexist on the same physical computer, in strong isolation with each other. virtual machine technology also offers the possibility of consolidation of applications in cloud computing environments[12], [13], which presents a significant opportunity for energy optimization. Consolidation is a well-known technique to dynamically reduce the number of nodes used within a running cluster by liberating nodes that are not needed by the current phase of the computation[14]. Fig. 1 demonstrates the power consumption of data center reducing in the scenario of the employment of VM migration technology. However, the consolidation itself brings some negative impact, such as the failure to fulfill the Service Level Agreement (SLA), extra power consumption within the procedure of migration. So understanding the impact of consolidating applications is necessary to design an effective consolidation strategy.

B. Live Migration

Virtual machine migration[5], [15], which is used to transfer a VM across physical servers, has served as a main approach to achieve better energy efficiency of data centers. This is because in doing so, server consolidation via VM migrations allows more computers to be turned off. Generally, the migration of VMs can be classified into two categories: regular migration and live migration. The first moves a VM from one host to another by pausing the originally used server, copying its memory contents, and then resuming it on the destination. The

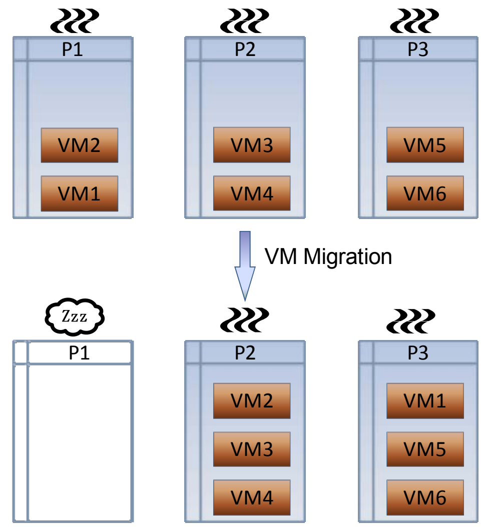

Fig. 1. Resource Consolidation in Cluster. It consists of 6 VMs (VM1-VM6) running on three physical servers (P1-P3). with the VM migration technology. VM1 and VM2 is consolidated on P2 and P3 respectively. And P1 is turned off, so the power consumption of the cluster is reduced.

second performs the same logical functionality but without the need to pause the server domain for the transition. Compare to regular migration, the live migration shows a great potential of using VM and VM migration technology to efficiently manage workload consolidation, and therefore improve the total data center power efficiency.

In this paper, we mainly focus on VM live migration, in which a VM is transferred from a physical server to another while continuously running, without any noticeable effects form the point of view of end users.

IV. Experiments and Results

In the following section, we present experimental design and results. Our main goal is to achieve a better understanding of power influence of live migration, according to the CPU utilization percentage. We consider two aspects that mainly dedicate to power cost of server: processor frequency and CPU utilization percentage, and for this reason, we have designed two preliminary experiments. The first one is to verify that the server power cost can be represented by CPU usage, specifically, is directly proportional to CPU usage. The second is to get power consumption of server in each processor frequency, which also verifies that in a fixed frequency , the power consumption can be represented by CPU utilization percentage. And then we select a fixed frequency to do the ongoing experiment. At last we have evaluated the power consumption caused by live migration.

A. Experimental Setup

In this experiment, we used three physical servers: one server is for VM hosting and operates one or two VMs, each of which handled by the workload that control the utilization

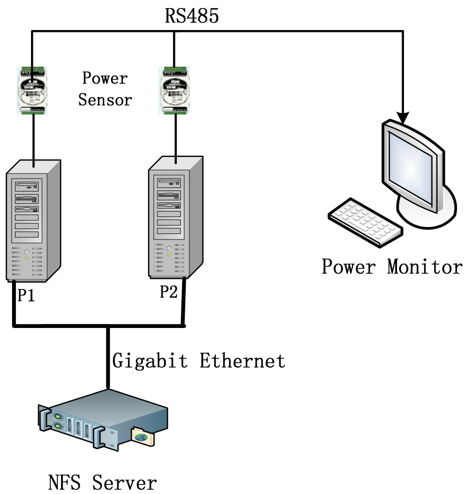

Fig. 2. Experimental Environment Deployment

percentage. The other one which is used to accepted the VM transfered by the originator. All nodes shares an NFS (Network File System) mounted storage device. Each node is equipped with Intel(R) Core(TM)2 DuoCPU E8400 and 3 Gigabytes memory. The servers are connected through a Gigabit Ethernet switch. An iPDU (Intelligent Power Distribution Unit) power meter is adopted to monitor the real-time power consumption of physical machines. Power-related parameters monitored by the power meter for a machine include Current, Voltage, Power and Kilowatt hour. To inspect energy consumption details of IT equipment and facilities in the system, the experimental environment is deployed as the following topology design as shown in Fig. 2. The parameters are collected every 2 seconds.

1) Preliminary experiment: Generally, power consumption can be expressed as percentage of the peak power across the data center. In the model of Gong Chen et.al[16], for a fixed operating frequency, the power consumption of the physical server is approximately linear model of the server utilization. Literature and our experiment indicate that power consumption is mainly determined by CPU usage. For this reason, we design a computational workload which runs in the virtual server to control its CPU utilization percentage. According to our experimental data, we verify that the power consumption increase almost linearly with CPU utilization, which is presented in Fig.3.

We formalize the linear model, using formula (1). According to our results, the power consumption of server can be represented as follows:

$$P = 0.2782 * Util + 51.2765 \tag{1}$$

where P is the power consumption of our physical server, and $Util$ is the CPU utilization.

CPU frequency of a system can be configured as several models, including powersave, userspace, ondemand etc. We evaluate the power consumption on each processor frequency, as shown in Fig. 4.

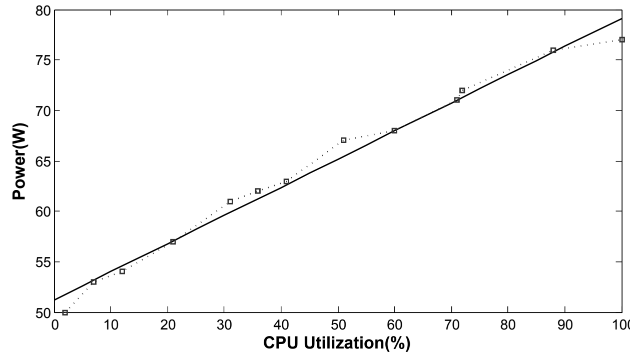

Fig. 3. Power Consumption v.s. Utilization. The actual power consumption on each CPU utilization is nearly equal to the model we set up.

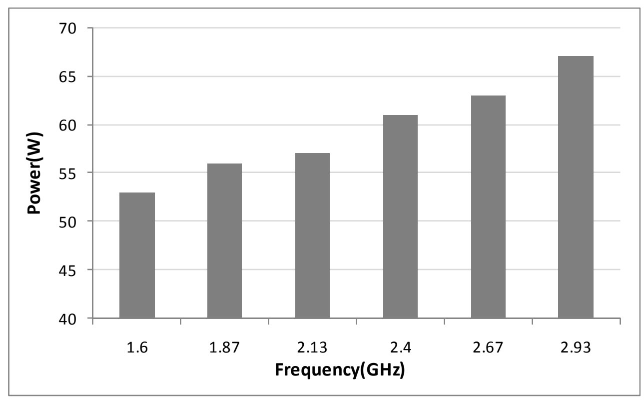

Fig. 4. Power consumption on each CPU frequency when system is fully utilized. The CPU has 2 processors, and each processor include 6 available frequency, from 1.6GHz to 2.93GHZ. The power consumption can be varied from 53Watt to 67Watt. So power consumption can be reduced by configuring CPU frequency.

2) Power Consumption of Live migration: The overall objective of our experiments is to quantify the power consumption when a virtual server is transferred from the original physical server to the destination. Specifically, this power cost of migration is comprised of two parts: the first part is the power used by the original physical server, which starts the migration; the second part is the power used by the destination server. All the cost is caused by the increase of resources, including computational resource (CPU etc.); storage resources (memory, disk,etc.)and I/O resource (Network).

To explore the power consumption of live migration according to the CPU utilization, We designed a computational workload which can singely change the CPU utilization of virtual server without influence other parts of the computer. Fig. 5 and Fig. 6 show the power consumption on both original server and destination server, where the horizontal axis indicates the average CPU utilization reported by the virtual machine manager, and the vertical axis indicates the average power consumption measured at the server power plug.

We observe two important facts. First, the power influence

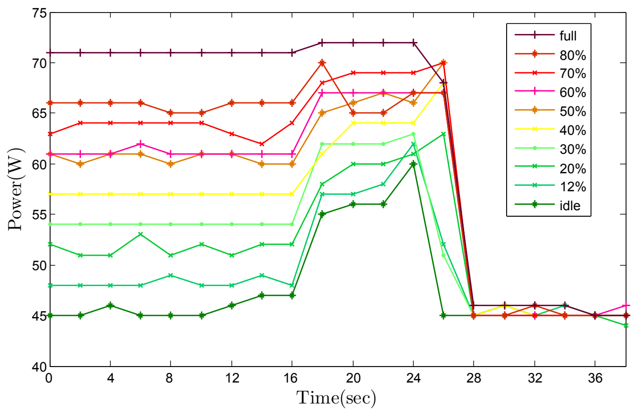

Fig. 5. Power Consumption on Original Server. In each CPU utilization percentage, The procedure of migration in the original physical cost nearly 7 seconds. The power consumption of migration decreases from 10Watt to 1Watt when the CPU utilization rises from idle to ful.

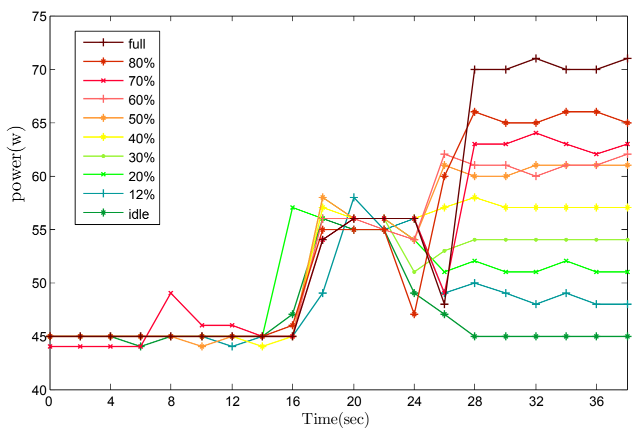

Fig. 6. Power Consumption on destination server. In each CPU utilization percentage, time cost and power consumption is almost the same. The former is nearly 7 seconds, the latter is 10Watt.

of migration on the original server goes down with the increase of CPU usage of the migrated VM, but for the destination server, the influence is stable, which is around 10-Watt power cost. Additionally, the time cost of migration is not impacted by the CPU usage of VM.

V. Conclusions and Future Work

VM migration is key to realize VM-based resource reservation and power reduction. And understanding its impact is important to make power-efficient deployment in data centers. This paper quantifies the cost of live migration for both source and destination physical servers, according to the CPU utilization percentage. Based on our results, several interesting findings are revealed: as for the original server, the power impact of live migration falls as the CPU utilization increases. However, the destination server is not influenced by the CPU usage of virtual machine transferred to it. Additionally, The time cost for both source and destination server is not affected by the CPU usage of virtual server transferred to it.

The ongoing investigation is focus on generalizing this paper's results and evaluating the migration cost. In the future work, we will try to model the power consumption in the procedure of live migration based on the results. And a benchmark workload also will be applied to verified our study.

VI. Acknowledgement

This work is supported by The National Natural Science Foundation of China (Grant No.60773093,60873209,60970107,60970108), The Key Program for Basic Research of Shanghai (Grant No.08JC1411800) , The Ministry of Education and Intel joint research foundation (Grant No.MOE-INTEL-08-11),The Science and Technology Commission of Shanghai Municipality09510701600IBM SUR Funding and IBM Research-China JP Funding.

References

[1] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. L. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield, "Xen and the art of virtualization," in *SOSP*, 2003, pp. 164–177.

[2] "Vmware, vmware inc. http://www.vmware.com."

[3] "Microsoft virtual server, microsoft coroporation. http://www.microsoft.com/windowsserversystem/virtualserver/."

[4] W. Voorsluys, J. Broberg, S. Venugopal, and R. Buyya, "Cost of virtual machine live migration in clouds: A performance evaluation," in *CloudCom*, 2009, pp. 254–265.

[5] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield, "Live migration of virtual machines," in *NSDI*, 2005.

[6] R. Uhlig, G. Neiger, D. Rodgers, A. L. Santoni, F. C. M. Martins, A. V. Anderson, S. M. Bennett, A. Kagi, F. H. Leung, and L. Smith, "Intel virtualization technology," *Computer*, vol. 38, no. 5, pp. 48–56, 2005.

[7] M. Zhao and R. J. Figueiredo, "Experimental study of virtual machine migration in support of reservation of cluster resources," in *VTDC '07: Proceedings of the 2nd international workshop on Virtualization technology in distributed computing*. New York, NY, USA: ACM, 2007, pp. 1–8.

[8] S. Srikantaiah, A. Kansal, and F. Zhao, "Energy aware consolidation for cloud computing," in *HotPower'08: Proceedings of the 2008 conference on Power aware computing and systems*. Berkeley, CA, USA: USENIX Association, 2008, pp. 10–10.

[9] T. Imada, M. Sato, and H. Kimura, "Power and qos performance characteristics of virtualized servers," in *GRID*, 2009, pp. 232–240.

[10] D. Meisner, B. T. Gold, and T. F. Wenisch, "Powernap: eliminating server idle power," in *ASPLOS*, 2009, pp. 205–216.

[11] L. Liu, H. Wang, X. Liu, X. Jin, W. B. He, Q. B. Wang, and Y. Chen, "Greencloud: a new architecture for green data center," in *ICAC-INDST '09: Proceedings of the 6th international conference industry session on Autonomic computing and communications industry session*. New York, NY, USA: ACM, 2009, pp. 29–38.

[12] R. Nathuji and K. Schwan, "Virtualpower: coordinated power management in virtualized enterprise systems," in *SOSP '07: Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles*. New York, NY, USA: ACM, 2007, pp. 265–278.

[13] F. Hermenier, X. Lorca, J.-M. Menaud, G. Muller, and J. Lawall, "Entropy: a consolidation manager for clusters," in *VEE '09: Proceedings of the 2009 ACM SIGPLAN/SIGOPS international conference on Virtual execution environments*. New York, NY, USA: ACM, 2009, pp. 41–50.

[14] F. Hermenier, N. Loriant, and J.-M. Menaud, "Power management in grid computing with xen," in *ISPA Workshops*, 2006, pp. 407–416.

[15] T. Wood, P. J. Shenoy, A. Venkataramani, and M. S. Yousif, "Sandpiper: Black-box and gray-box resource management for virtual machines," *Computer Networks*, vol. 53, no. 17, pp. 2923–2938, 2009.

[16] G. Chen, W. He, J. Liu, S. Nath, L. Rigas, L. Xiao, and F. Zhao, "Energy-aware server provisioning and load dispatching for connection-intensive internet services," in *NSDI*, 2008, pp. 337–350.