

INTERNATIONAL JOURNAL OF ENGINEERING SCIENCES & MANAGEMENT

AN INVESTIGATION OF VIRTUAL MACHINE ENERGY EFFICIENCY AND CONSOLIDATION TECHNIQUES IN IAAS CLOUD

Jitendra Soni, Ambar Dixit

Department of Information Technology Institute of Engineering and Technology, DAVV Indore, India
jsoni@ietdavv.edu.in, ambardixit@iiti.ac.in

ABSTRACT

Cloud computing is a network based, on demand access model to a pool of configurable computing resources. These resources include servers, storage, networks, applications and services. One of the major services provided through cloud computing is called IaaS (Infrastructure as a Service), where the access of infrastructure components such as storage, CPU, memory, and other devices is provided through internet. Cloud based services and web applications accessing them are rapidly growing as they promise high performance and uninterrupted services. This growth has posed a great demand of virtualized cloud data center establishments around the world. Typically, the power requirements of these data centers are huge, which leads to high operational cost as well as environmental issues. Accordingly, cloud providers started considering energy efficiency and consumption as a deciding factor for virtual machine (VM) placements. In its lifetime, a VM puts variable resource demand on physical host machine, and thus the initial placement of VM is not always a deciding parameter. In the present study, a thorough investigation has been carried out in virtual machine placement techniques in Infrastructure as a Service (IaaS) based cloud offering.

INTRODUCTION

According to NIST [1], "Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g. networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction." Cloud grants its users a cost effective model of computing with tailor made services on a subscription based pay-per-use model. Resources are scalable according to variable load, promising uninterrupted service to its end users. Cloud service reference model is broadly classified into following types:

- Infrastructure as a Service (IaaS)
- Platform as a Service (PaaS)
- Software as a Service (SaaS)

The first layer and the foundation of this stack is IaaS layer, Infrastructure as a service provides basic computing infrastructure e.g. CPU, RAM, Network, Storage. In this infrastructure as a service, cloud provider provides access to virtualized hardware via network to its client. Virtualization plays a key role in creating virtualized hardware by creating an abstraction layer of Hypervisor. A Hypervisor is a collection of computer software, firmware or hardware that creates and runs a virtual machine [2]. The hardware and operating system where hypervisor installed is known as host machine whereas the virtual resources created by the hypervisor is known as Guest machine.

These guest machines are created on the request of users using web browsers or tools provided by the cloud providers or set of commands via the internet. Users can use these virtual machine instances as per their requirements, and they can install, configure the software in the same fashion as a desktop or server at their own premises.

The layer above Infrastructure base is the Platform as a Service layer, It provides a scalable abstraction to deploy and develop applications. It has support for different runtime environments, libraries that are the part of the core middleware functionality provided by PaaS. It hides the details of underlying infrastructure services and provide a seamless and elastic environment to a user using web based management console or using APIs and libraries. The topmost layer in cloud service stack is Software as a Service (SaaS) which is built using the services and environment provided by PaaS layer.

A. Role of Data Centers

This entire service stack is hosted on virtualized data centers established by cloud providers. A typical data center consists of thousands of computing nodes, switches, routers. In order to increase the hardware life and reduce the downtime and failure of nodes, the temperature inside the data center should be kept low by using Air conditioner (A.C.), and other cooling methods like water based cooling. Hence efficient server nodes and cooling techniques play a significant role in electricity consumption. Subsequently electricity cost is a major contributor in operational expenditures. In earlier data center configurations, performance was the main concern of service providers. These servers used to run at low utilization level and

promising high Quality of Service (QoS). Usually, the power consumption is as high as 50% of the maximum, even at 10 % of CPU utilization [3]. Server Virtualization gave the solution to the problem of raising the CPU utilization. A single physical server can run multiple virtual machines, and consolidation of virtual machines on lesser physical nodes can lead to a situation having some empty servers, which can be put into sleep or hibernation to save a considerable amount of power. Due to strong environmental laws and carbon credit policies, researchers tossed a new term Green Cloud Computing, Which is defined as Green Cloud computing is envisioned to achieve not only the efficient processing and utilization of computing infrastructure, but also to minimize energy consumption. [4].

| Category | Power Consumption 2008 (GW) | Growth Rate (PJV.) | 2020 Prediction (GW) |
|-------------------|-----------------------------|--------------------|----------------------|
| Data Centers | 29 | 12% | 113 |
| PCs | 30 | 7.5% | 71 |
| Network Equipment | 25 | 12% | 97 |
| TVs | 44 | 5% | 79 |
| Other | 40 | 5% | 72 |
| Total | 168 | | 433 |
| World Electricity | 2350 | 2.0% | 2970 |
| ICT Fraction | 7.15% | | 14.57% |

Fig. 2: Power Consumption Prediction

CLOUD COMPUTING

Cloud Computing is the technology which delivers computing resources whether it is computing power, storage, network, memory, applications, software platforms over the internet; it include the services, data centers and its components [5]. Fox et al. [6] from Berkeley defined cloud Computing vision as: "Cloud computing, the long-held dream of computing as a utility, has the potential to transform a large part of the IT Industry, making software even more attractive as a service." Cloud turned the computing resources as a utility and now the user can subscribe to cloud services on the basis of "Pay as you Go" model of utilities like electricity and water.

A. Virtualization

Virtualization is the basis of Cloud Computing. It contains a set of solutions allowing the abstraction of the necessary




Fig. 3: Concept of Hypervisor.

components for computing, such as Processor, Memory, Storage, Network and other hardware. The Term Virtualization get its existence 40 years ago [7] [8], Some of the general virtual machine systems at that time were IBM VM/370 [9] but the functionalities were limited by technologies like very subtle amount of memory and low computing power. Present scenario of computing overcame the limitations of resources, therefore, virtualization become an important unit of cloud computing. Virtualization provides highly customizable and controllable IT infrastructure for users. Virtualization can be classified in two categories:

- 1) System/Hardware virtualization.
- 2) Process Virtualization.

The Most popular virtualization method is System/Hardware virtualization. In Hardware virtualization the guest system run over an abstract execution environment in terms of hardware. In this virtualization the Guest is the operating system running over virtual hardware, the host is the physical computer, the virtual machine manager is the hypervisor.

1) *Hypervisors*: The Hypervisor is a software or assembly of software and hardware that constructs the abstraction of the underlying physical server or hardware. It simulates or recreates the hardware environment. Installation of a guest operating systems can be possible over this abstraction layer.

B. Virtualization and Relation with Cloud Computing

As per discussed earlier Virtualization plays an important role in cloud computing because of its high level of customization, isolated and its secure execution, these features are essential for delivering on demand computing resources through the internet. Configurable resources can be delivered to clients using Virtualization technologies e.g. Hardware Virtualization enabled IaaS market segment, whereas programming language level Virtualization opens the door for PaaS services. Both offerings provide sandboxed, customizable environment for business enterprises having a large scale computing infrastructure with capabilities to sustain and deliver uninterrupted services despite heavy workloads.

C. VM Migration

Sandboxed virtual machines are isolated process running in a server. A Virtual machine can easily be saved in the storage, and the saved snapshots can be again copied, deployed on any other host, with same or different architecture and hardware from the previous host with same VMM layer below the hosts [10]. Migration of VM is very useful in the situation such as Disaster recovery, server replication, Hardware upgrade. This simple pure stop-and-copy [11] [12] migration process is slow and only possible after suspending the execution of VM to create a snapshot of the VM as compared to dynamic migration of VM. Dynamic or live migration is an improved technique to migrate VMs where virtualized servers migrate in running state between physical hosts without losing client and application connectivity. A VM with a load of running live service is critical and dynamic transfer occurs in a manner that fulfil the requirements of minimizing both downtime and total migration time. In the proposed work in the next chapter we are using Dynamic Live migration of VM from one host to another.

VM Migration consists of mainly three phases Clarke et al. [10]: Push phase: The source VM continues running while particular pages are pushed across the network to a new destination. To ensure consistency, pages modified during this process must be re-sent. Stop-and-copy phase: The source VM stopped, pages are copied across to the destination VM, and then the new VM is started. Pull phase: The new VM executes and, if it accesses a page that has not yet copied, this page is faulted in ("pulled") across the network from the source VM.

D. VM Consolidation

VM consolidation is a Server consolidation process to optimally moving the sandboxed machine running in one host to another host so that servers with underutilized resources can either fully release after migration or more VMs can be placed tightly inside this host. It increases resource utilization up to maximum permissible limit, resulting in best utilization of the resources in every server. Jung et al. [13] investigated the problem of dynamic VM consolidation using live migration with challenging SLA requirements, where VM is having a workload of a multi-tier web application. They proposed VM placement using bin packing and gradient search. Migration controller sense the need of VM migration and the reconfiguration based on Utility function that calculates the SLA requirements. Kumar et al. [14] have proposed a factor stability for solving dynamic VM consolidation problem. The stability factor is the probability that a VM allocation will continue to be operative for some interval in the future. A time varying probability density function is used for the prediction of the future resource load. Berral et al. [15] have used machine learning techniques for dynamic VM consolidation with VM running workloads with deadline that are fixed in SLAs. The process optimize the energy consumption and SLA using machine learning. This optimization method is intended for explicit environments e.g. High Performance computing(HPC) because HPC are deadline constraint applications.

ENERGY CONSUMPTION ANALYSIS IN DATA CENTERS

Minas and Ellison in their book [16] by Intel labs categorize several key components which consume power and presented that CPU, Memory and power supply losses are the major components of power consumption. Techniques like DVFS helped improving CPU power saving by enabling active low power mode. Due to these techniques modern servers and desktop CPUs consumes not more than 30% of their peak power in low activity mode, therefore rest of 70% power range lies in dynamic power consumption [17]. This dynamic range is comparatively low for DRAM Memory (modern RAM technology). It is less than 50% for DRAM, It is 25% for Hard

Disk, and for network devices it is 15% [18].

A. Operational Cost of a Typical Large Cloud Service Provider

Greenberg et al. [19] have shown a cost breakdown of a Data centers capital expenditure. Authors probed that 15% of the total cost used into power consumption. They estimated that the total annual power cost for a Mega data center (consist of 50,000 server with 180 W power rating) is \$ 9.3 million at the rate of \$ 0.07 per unit cost for the electricity. This shares the major part of operational expenditure of Data centers.

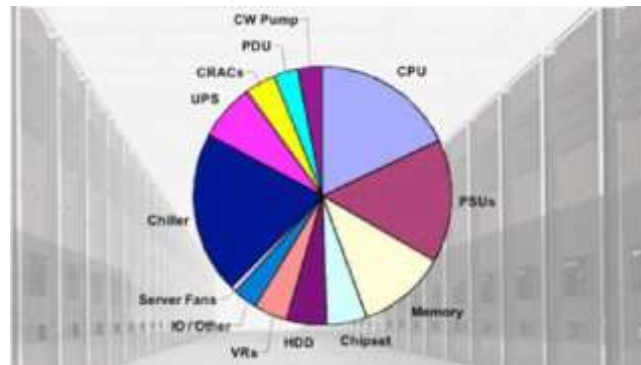


Fig. 4: Energy Consumption of a typical Data center Components.

ENERGY EFFICIENCY IN DATA CENTERS

Pinheiro et al. [20] presented the very first work that focuses on the power management and energy savings in clusters of servers and personal computers. Their methodology is applicable where there is replication of nodes occurs for load balancing; their work was to limit the workload on minimum possible servers in the cluster in a way that maximum idle nodes can be available to put into sleep mode of hibernation to save a significant amount of energy. They implemented the algorithm on the two general cluster configuration: A locality aware network server and a load balancing distributed operating system for clustered cycle servers. The load balancing is done at the application level not at the system level. The algorithm runs on the master node without any redundant node involved so it creates a single point of failure. Authors mentioned that the algorithm removes the nodes one by one, the cluster reconfiguration option is also a time consuming process therefore due to lack of parallel load balancing and configuration method it is not scalable beyond a fixed number of servers in a cluster.

Chase et al. [21] proposed an algorithm for resource management whose main focus is energy efficiency. The primary aims of this work is to assign server resources for co-hosted services in a way that automatically adjust to continuous changing workload. To increase the energy efficiency of server dusters by dynamically changing the active server set, and react to power supply interruptions or thermal events by degrading service in accordance with negotiated Service Level Agreements (SLAs).

Elnozahy et al. [22] Applied two techniques to save power 1. Switching off/on the idle/useable hosts 2. Dynamic Voltage and Frequency Scaling (DV FS), they concentrated on their evaluation environment for a single web application with a predefined response time as SLA. The Policy calculates the total CPU cycles and number of servers needed for the application execution, and it assigns the scaled frequency to CPUs of every server.

One of the significant work by Nathuji And Schwan [23] proposed Virtual Power. It is one of the very first work of power management and energy savings in Virtualized data centers. They extended the hardware level power management and frequency scaling as Soft scaling on the guest VM level.

Shrikantaiah et al. [24] modeled the consolidation of applications in virtualized data centers as bin packing problem. They examined the relationship between consolidation performances, consumption of energy and resource utilization. Authors showed that there exists an optimal point in the performance and energy consumption tradeoff due to workload consolidation.

Verma et al. [25] gave a solution for energy efficient allocation of applications in virtualized heterogeneous data centres. They formulated energy efficient assignment of the application in heterogeneous host as continuous optimization, in a fixed time interval the vm allocation is reviewed for energy and performance tradeoffs and optimize this tradeoff with appropriate migrations.

A. Existing Commercial Solutions

VMware vCloud suit is used to create a full-fledged private cloud infrastructure. Tools provided by VMware e.g. VMware DRS (distribute resource Scheduler) [26] , and VMware DPM (distributed power management) [27] can

monitor and manage the energy and power consumption of the virtualized infrastructure created by VMware vCloud suit.

Power IQ [28] is developed by Raritan a company member in green grid initiative power IQ gives customized solution to data center managers and administrators. Some powerful features of Power IQ are one click rack monitoring, power capacity forecast chart and PUE (Power usage effectiveness) monitoring.

CONCLUSION AND FUTURE WORKS

This literature provides a detailed study of data center energy consumption elements, IaaS cloud and its basic building blocks to analyze the energy consumption elements in cloud infrastructure. This work investigated almost all well known and industry accepted energy saving techniques in data centres. In this survey VM allocation methods and their efficient placement to optimize various parameters has been discussed, mainly in the context of energy savings. Some commercial VM placement algorithms are also discussed. This survey can be extended by including PaaS (Platform as a Service) based optimizations methods of reducing the energy consumption and increasing the SLA.

REFERENCES

1. P. Mell and T. Grance, "The nist definition of cloud computing (draft)," *NIST special publication*, vol. 800, no. 145, p. 7, 2011.
2. Wikipedia, "Hypervisor — wikipedia, the free encyclopedia," 2014, [Online; accessed 27-March-2014]. [Online]. Available: <http://en.wikipedia.org/w/index.php?title=Hypervisor&oldid=599741863>
3. A. Beloglazov, J. Abawajy, and R. Buyya, "Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing," *Future Generation Computer Systems*, vol. 28, no. 5, pp. 755-768, 2012.
4. D. Kusic, J. O. Kephart, J. E. Hanson, N. Kandasamy, and G. Jiang, "Power and performance management of virtualized computing environments via lookahead control," *Cluster computing*, vol. 12, no. 1, pp. 1-15, 2009.
5. M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica *et al.*, "A view of cloud computing," *Communications of the ACM*, vol. 53, no. 4, pp. 50-58, 2010.
6. A. Fox, R. Griffith, A. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, and I. Stoica, "Above the clouds: A Berkeley view of cloud computing," *Dept. Electrical Eng. and Comput. Sciences, University of California, Berkeley, Rep. UCB/EECS*, vol. 28, p. 13, 2009.
7. G. J. Popek and C. S. Kline, "The pdp-11 virtual machine architecture: A case study," in *Proceedings of the Fifth ACM Symposium on Operating Systems Principles*, ser. SOS '75. New York, NY, USA: ACM, 1975, pp. 97-105. [Online]. Available: <http://doi.acm.org/10.1145/800213.806527>
8. R. P. Goldberg, "Survey of virtual machine research," *Computer*, vol. 7, no. 6, pp. 34-45, 1974.
9. R. J. Creasy, "The origin of the vm/370 time-sharing system," *IBM Journal of Research and Development*, vol. 25, no. 5, pp. 483-490, 1981.
10. C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield, "Live migration of virtual machines," in *Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation-Volume 2*. USENIX Association, 2005, pp. 273-286.
11. C. P. Sapuntzakis, R. Chandra, B. Pfaff, J. Chow, M. S. Lam, and M. Rosenblum, "Optimizing the migration of virtual computers," *SIGOPS Oper. Syst. Rev.*, vol. 36, no. SI, pp. 377-390, Dec. 2002. [Online]. Available: <http://doi.acm.org/10.1145/844128.844163>
12. M. Kozuch and M. Satyanarayanan, "Internet suspend/resume," in *Mobile Computing Systems and Applications, 2002. Proceedings Fourth IEEE Workshop on*. IEEE, 2002, pp. 40-46.
13. G. Jung, K. Joshi, M. Hiltunen, R. Schlichting, and C. Pu, "A cost-sensitive adaptation engine for server consolidation of multitier applications," in *Middleware 2009*, ser. Lecture Notes in Computer Science, J. Bacon and B. Cooper, Eds. Springer Berlin Heidelberg, 2009, vol. 5896, pp. 163-183. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-10445-9_9
14. S. Kumar, V. Talwar, V. Kumar, P. Ranganathan, and K. Schwan, "vmanage: loosely coupled platform and virtualization management in data centers," in *Proceedings of the 6th international conference on Autonomic computing*. ACM, 2009, pp. 127-136.

15. J. L. Berral, I. Goiri, R. Nou, F. Julia, J. Guitart, R. Gavalda, and J. Torres, "Towards energy-aware scheduling in data centers using machine learning," in *Proceedings of the 1st International Conference on energy-Efficient Computing and Networking*. ACM, 2010, pp. 215—224.
16. L. Minas and B. Ellison, *Energy efficiency for information technology: How to reduce power consumption in servers and data centers*. ntel Press USA, 2009.
17. L. A. Barroso and U. Holzle, "The case for energy-proportional computing," *Computer*, vol. 40, no. 12, pp. 33—37, 2007.
18. X. Fan, W.-D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," *ACM SIGARCH Computer Architecture News*, vol. 35, no. 2, pp. 13—23, 2007.
19. A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel, "The cost of a cloud: research problems in data center networks," *ACM SIGCOMM Computer Communication Review*, vol. 39, no. 1, pp. 68—73, 2008.
20. E. Pinheiro, R. Bianchini, E. V. Carrera, and T. Heath, "Load balancing and unbalancing for power and performance in cluster-based systems," in *Workshop on compilers and operating systems for low power*, vol. 180. Barcelona, Spain, 2001, pp. 182—195.
21. J. S. Chase, D. C. Anderson, P. N. Thakar, A. M. Vahdat, and R. P. Doyle, "Managing energy and server resources in hosting centers," in *ACM SIGOPS Operating Systems Review*, vol. 35, no. 5. ACM, 2001, pp. 103—116.
22. E. Elnozahy, M. Kistler, and R. Rajamony, "Energy-efficient server clusters," in *Power-Aware Computer Systems*, ser. Lecture Notes in Computer Science, B. Falsafi and T. Vijaykumar, Eds. Springer Berlin Heidelberg, 2003, vol. 2325, pp. 179—197. [Online]. Available: <http://dx.doi.org/10.1007/3-540-36612-L12>
23. R. Nathuji and K. Schwan, "Virtualpower: coordinated power management in virtualized enterprise systems," *ACM SIGOPS Operating Systems Review*, vol. 41, no. 6, pp. 265—278, 2007.
24. S. Srikantaiah, A. Kansal, and F. Zhao, "Energy aware consolidation for cloud computing," in *Proceedings of the 2008 conference on Power aware computing and systems*, vol. 10. USENIX Association, 2008.
25. A. Verma, P. Ahuja, and A. Neogi, "pmapper: power and migration cost aware application placement in virtualized systems," in *Middleware 2008*. Springer, 2008, pp. 243—264.
26. VMware, "Cloud computing," <http://www.vmware.com/cloud-computing/hybrid-cloud.html>, March 2013. [Online]. Available: <http://www.vmware.com/cloud-computing/hybrid-cloud.html>
27. V. Infrastructure, "Resource management with vmware drs," *VMware Whitepaper*, 2006.
28. Raritan. Power iq - data center energy management software. Raritan. [Online]. Available: <http://www.raritan.com/products/dcim-software/power-iq>