International Journal of Advances in Computer Science & Its Applications
Volume 6 : Issue 1 [ISSN 2250-3765]
Publication Date : 18 April, 2016

Saving Energy by Choosing Fast the Best Server in Cloud Data Centers

Valma Prifti[1], Avjol Sakaj[2], Alban Afmeti[3], Igli Tafa[4]

***Abstract* - Virtualization technology is being used more recently in order to improve the balance between the performance of a server and its energy consumption. The use of a virtual machine also has many benefits regarding distribution of the resources, and energy aware consolidation scalability, and powerful solutions for data backup and recovery. Server consolidation is one of the ways to reduce the energy consumed in data centers between hardware and software features. During consolidation we should consider some parameters of the servers such as performance, the relation between the CPU and I/O equipments etc. VM allocation and placement are the main issues in consolidation. The work that we have done is that we have rearranging the allocation in a way that saves energy. Finally the simulation and live demonstration results are conspired. The simulation results show the efficiency of this method.**

***Keywords* – Virtualization, Server Consolidation, energy consumption, Data center**

I. Introduction

Data centers nowadays use enormous amount of energy.Jonathan G. Cooney, a consulting professor at Stanford University has estimated global data center energy used for a number of years. He found in the US data centers energy use that the consumption increased by 200% from 2000-2005 and by 36% from 2006-2010, and that data centers currently use 2% of the electricity in the United States. Studies have found that data centers should be made far more efficient in order to save energy which is a very important resource for human life. In fact, an efficient data center can use up to 80% less power than an inefficient data center, and the upshot of all this is that data centers represent an enormous opportunity to save energy. These data center energy efficiency opportunities we're talking about are not just limited to the large data centers with analyst server racks. In fact the average commercial office building spends close to 23% of annual energy costs on server rooms and closets, and for a highly efficient office building, that number increases almost half of the buildings energy footprint. So, even though the opportunity is slightly different, it may be more focused on IT equipment in the server room and server closets, there still is a large opportunity for everyone when it comes to data center energy efficiency.

High availability has always been a major problem for a data center. So far high availability is achieved by redundancy host to host, a method very expensive in terms of hardware and human costs. A new road to this problem can be offered by virtualization. With the virtualization, it is possible to reach a redundancy system for all services that are executed in a data center. This new way of high availability allows for separate virtual machines that are executed on servers, using the virtualization layer characteristics: start, stop and move virtual machines between physical hosts. Virtualization is a simulation of a hardware platform, operating system, a storage device or network resources simulation. In other words it's the creation of a virtual version of them. Practically, virtualization provides the ability to run applications, operating systems, or systems of services in a unique independent and logical environment from physical computer system. The focus of virtualization in operational logical environments instead of physical environments makes the applications, services and instances of an operating system transferable across different physical computer systems.

When we think about server consolidation on a mass scale, data center IT administrators have it large problem under hands. Consolidating multiple operating systems on each physical server allows for decreased cost for power consumption and cooling since you will be running fewer physical servers overall. It allows too, increased uptime and better redundancy. Another benefit of server consolidation is that it allows automated provisioning of virtual machines, seamless resource sharing (Consolidated I/O) and easier and more flexible Backup/Recovery.

Companies are seeing that they need for high-performance servers to meet the high requirements of new applications. At the same time server consolidation and switching from stand alone servers to rack-concentrates amounted servers is focusing systems in small spaces. To run data centers as efficiently as possible, nowadays companies are trying to balance server performance with energy consumption.

If companies continue to install more powerful servers in their rooms without making any changes, there may be more problems arise, for example the power consumption per system will increase. Secondly, heating and associated cooling become more complex problems. Processors with higher performance consume more electricity and the electricity mostly is converted to heat. By increasing this heat, equipment temperature rises and this leads to increased equipment failure, so the power management in data centers is crucial as it is not only limited to the energy cost but also the initial investment of cooling systems.

II. Related Works

The [3] presents some experiments using the Trade6 application. These experiments show that when a server cluster is being managed by the controller, it conserves 26% of the power required by the system without the management of the controller. The [7] proposes an algorithm for finding the minimum of the energy allocation of workloads to servers. In the first step the algorithm determines the optimal point. Next, the request that has arriving is allocated to a server. If the request cannot be allocated, they turn on a new server for re-allocation of all the requests using the same heuristic. [8] Presents a study on the overlap of consolidating several applications on a single server with the Xen hypervisor. The [5] makes some experiments designed for an illustration of the system's behavior in various aspects in simple scenarios. They demonstrate competitive allocation of server resources in a cluster which is inadequate for unexpected load swells. In the [10] the consolidations of virtual machines are done using live migration and they proposed architecture of energy management system for virtualized data centers.

III. Background

Nowadays a very common phenomenon is the use of more virtual machines on a single physical machine. This process is called as the process of virtualization. The reason for which becomes such action is to increase the use of resources of a server.

Given that nowadays servers are very large, in most cases they do not have enough work to justify the use of the energy from them. Thus, in order to decrease the energy consumption, some servers need to be switched and others to take their jobs, using the technique of virtualization. But what if the servers are over-utilized or under-utilized?

As for over-utilized servers the solution is that these servers work together with their virtual machines to pass into a new server which provides us with the necessary resources but without increasing more energy consumption. As for under-utilized servers, good is that their virtual machines are transferred to a server which is more efficient and provides lower consumption energetic.

A simple model of energy consumption of a server can be given to us by the following formula:

$$P_{total} = P_{idle} + U_{CPU}(P_{max} - P_{idle})$$

- P_{total} = the total power of the server
- P_{idle} = server idle power
- U_{CPU} = is the CPU power
- P_{max} = the maximum power of the server.

By performing simple calculations we can find the energy consumed during migration and finally come to formulas shown and described below:

$$Data_{mig} = \frac{VM_{datasize}}{Data_{rate}}$$

- $Data_{mig}$ = the entire migration duration
- $VM_{datasize}$ = the size of data in VM
- $Data_{rate}$ = data rate for migration

$$Cost_{mig} = \frac{P_{mig} * Data_{mig}}{T}$$

- $Cost_{mig}$ = the cost of migration for VM
- P_{mig} = power used for migration
- T = duration for the migration (in seconds).

An important question arises when a server is over-utilized or under-utilized: which of virtual machines to migrate and where to migrate it?

To answer such a question we create an algorithm which will find us server that is able to make this work but also consume much less energy.

This algorithm checks all servers which are available and makes a selection on them and selecting only some servers that will meet certain conditions defined by the need to work as performance, resources, power, speed, availability, ... (where each server keeps in his memory some data on these parameters and which are calculated from his previous work or predetermined at the time of fabrication). In addition, this algorithm has another variable, which is tolerance, to make it possible to select a certain number of servers that meet the requirements of the above, but have a tolerance.

Once the selection is made the algorithm makes them a descending order based on the above mentioned conditions. Already selected servers were growing power to switch to full functionality and select the server that is suitable for the transfer of virtual machines and their works.

The algorithm:

Input: allhostList, allVMList
Output: bestVmAlloc in allHostList

1. For all allHostList: Select hosts that meet certain conditions and make their ranking in descending order => selectHost(conditions,…. , tolerance)
2. For all the selected hosts (preceding cycle), increases the power and pull them in working condition => host.AddPower()
3. Final Step: Choose the server with less energy consumption and give him the virtual machines => HostWithMinimalPower.add (vm).

IV. Experimental Environment

As, for the part of the simulation of the above mentioned system, within a datacenter have been created 15 virtual machines and 10 different hosts. Host computer has parameters:

1- 4 GB RAM, and 500GB memory core i5
2- 4 GB RAM, and 2 * 500 GB memory core i5
3- 2 GB RAM, and 500 GB memories core i3.

While in terms of virtual machines they have 256 RAM and 20 GB of memory. Threshold value was set between 0.6 and 0.8.

The first core i5 computer is working and in a moment of time needed to move to another computer more suitable. To do this migration is used Ubuntu 13.04. For hypervisor we have used KVM/QEMU. A virtual machine manager was installed called "VMM". The hosts were connected using LAN cable and one router, and necessaries applications were installed. Also, we create the NFS (network file system) storage in server and mounted on the host machines in it. We included the utilization threshold into the open source software codes of VMM in the source system.

International Journal of Advances in Computer Science & Its Applications
Volume 6 : Issue 1 [ISSN 2250-3765]
Publication Date : 18 April, 2016

V. Experimental Phase

Some applications were executed into source virtual machine and where executed repeatedly. A shell script was created and executed in the source host machine for observing the consumption of the CPU. VMM was run in all the computers and the migration starts when the host is over-utilization. The migration in the VMM is pre-copy as there is iterative copying and migration of CPU status of the virtual machine. After the entire virtual machine is transferred to the destination, it continues to run only there.

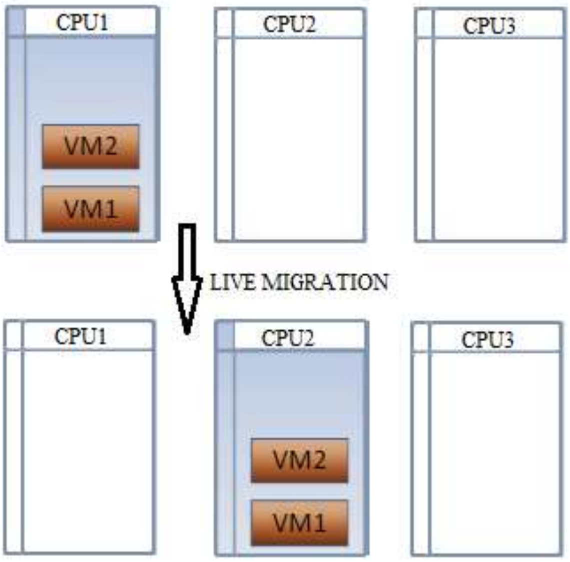

Figure 1: Live Migration

From our algorithm were chosen only hosts who satisfies the conditions (the second one because the third did not have the necessary parameters) and therefore the new host will meet the following requirements and the pre-host CPU is shouted down to keep the quality of servers and to reduce the power consummation.

VI. Interpretation of Results

After creating and configuring the necessary environment, by using our heuristics algorithm we start the simulation of the single threshold policy. In this time, during the simulation, the metric generated below is the energy consumed by the servers of single threshold policy.

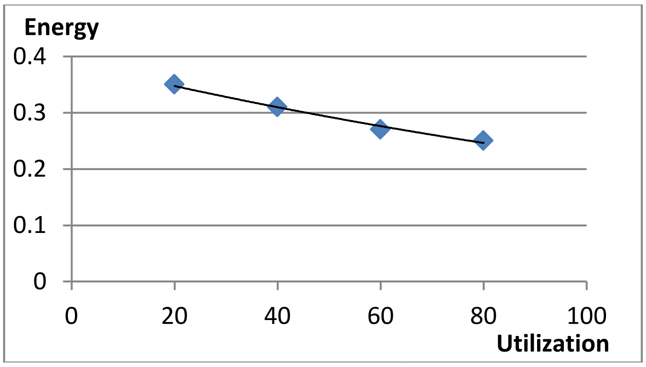

Figure 2: Energy consumed by the servers of single threshold policy

By using cloudsim we implemented the single threshold policy, and we saw what happened to the energy consumed by the server when the utilization threshold is increased. By this experiment, during migration, were derived some data and with them we create the chart below which is the live demonstration of energy consumption by the system.

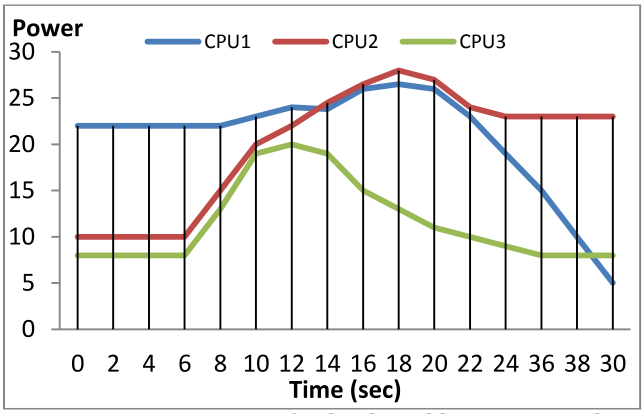

Figure 3: Energy consumption by the physical host computers during the process

This graph shows the energy consumption by the physical host computers during the initial time of the process until the end. It is noted that during the early migration of the power consumed by virtual machines is to a maximum, and then comes being reduced to switch to a state stable.
Also during migration, the use of resources always increases, causing increase the power consumption but at a very low level.

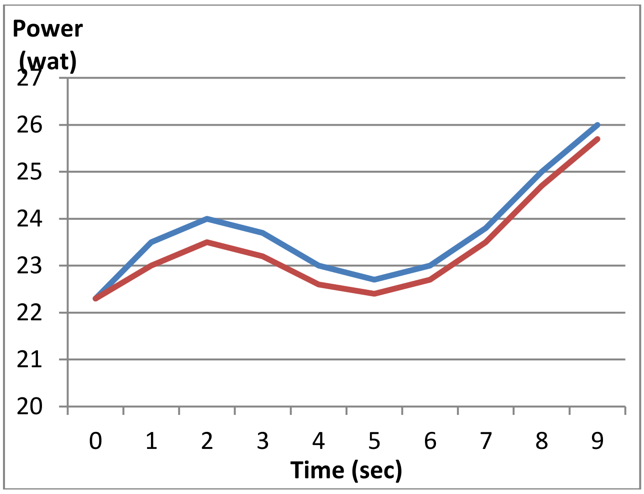

Figure 4: Energy consumption during migration in CPU1

- Blue = the worst case when all VM and hosts are close to each other in performance.
- Red = the best case when machines and hosts differ from each other in performance and destination choice for the hostname is performed quickly and without more power consumption.

Once that is done the migration of VM is created graph of energy in time. Small curve to the left is created in the beginning of the migration and finding the most appropriate host. As part of the curve to the right is the time when the host CPU is found and is now being carried migration of data.

VII. Conclusion

In order to reduce the power consumption of the system in general we use live migration to consolidate the servers and to reduce the number of physical machines that are power-on in that moment. The quality of service is another very important factor that should not be compromised and by shooting down some unused physical machine the QoS was increased.

The above described algorithm enables us to reduce energy consumption by 4% to 6%, by hanging it by the number of hosts as well as their similarities in performance.

This algorithm works better when the hosts differ from each other on in performance and power consumed by considerably reduces our energy consumption because solved soon suitable host for our work.

VIII. Future Work

By looking algorithm described above is understood that a question arises: how and how much should be the tolerance that is included in the algorithm so that it works very well but do not be left out of the selection list the hosts which may be appropriate and that consume less energy?

Our future work will be focus on this aspect, see and reestablish this algorithm tolerance in order to be more efficient as he could

References

[1] Gergo Lovasz, Florian Miedermeier, and Hermann de Meer (September 2013): "Performance Tradeoffs of Energy-Aware Virtual Machine Consolidation", Cluster Computing, Publisher: Kluwer Academic.

[2] William Voorsluys, James Broberg, Rajkumar Buyya, and Srikumar Venugopal (Nov 2009): "Cost of Virtual Machine Live Migration in Clouds: A Performance Evaluation", International Conference on Cloud Computing, Publisher: Springer-Verlag.

[3] Dara Kusic, Jeffrey O. Kephart, James E. Hanson, Nagarajan Kandasamy, and Guofei Jiang (June 2008): "Power and Performance Management of Virtualized Computing Environments via Lookahead Control", International Conference on Autonomic Computing, Publisher: IEEE Computer Society.

[4] Young Choon Lee and Albert Y. Zomaya (August 2010): "Energy Efficient Resource Allocation in Large Scale Distributed Systems", Ninth International Symposium on Distributed Computing and Applications to Business, Engineering and Science, Publisher: IEEE Computer Society.

[5] Jeffrey S. Chase, Darrell C. Anderson, Prachi N. Thakar, Amin M. Vahdat (December 2001): "Managing Energy and Server Resources in Hosting Centers", Eighteenth ACM symposium on Operating systems principles, Publisher: ACM.

[6] Richa Sinha, Nidhi Purohit, Hiteishi Diwanji (October 2014): "Energy Efficient Dynamic Integration of Thresholds for Migration at Cloud Data Centers", International Journal of Cloud Applications and Computing, Publisher: IGI Global.

[7] Shekhar Srikantaiah, Aman Kansal, Feng Zhao (December 2008): "Energy Aware Consolidation for Cloud Computing", 2008 conference on Power aware computing and systems, Publisher: USENIX Association.

[8] Padma Apparao, Ravi Iyer, Xiaomin Zhang, Don Newell, and Tom Adelmeyer (March 2008): "Characterization & analysis of a server consolidation benchmark",The fourth ACM SIGPLAN/SIGOPS international conference on Virtual execution environments, Publisher: ACM.

[9] Lionel Eyraud-Dubois, Hubert Larcheveque (May 2013): "Optimizing Resource Allocation while handling SLA Violations in Cloud Computing Platforms", IEEE 27th International Symposium on Parallel and Distributed Processing, Publisher: IEEE Computer Society.

[10] Ripal Nathuji, Karsten Schwan (October 2007): "Virtualpower: Coordinated power management in virtualized enterprise systems", Twenty-first ACM SIGOPS symposium on Operating systems principles, Publisher: ACM.

[11] Qiang Huang, Fengqian Gao, Rui Wang, Zhengwei Qi (2011): "Power Consumption of Virtual Machine Live Migration in Clouds", Third International Conference on Communications and Mobile Computing.