

Allocation of Virtual Machine Resources in Server Virtualization Environment

Thiri Thitsar Khaing[#], Hla Yinmin^{*}, Nang Noom Yein Sang[#]

^{#*}Faculty of System and Technologies, Computer University (Pin Lon)

¹thirithitsarkhiang@gmail.com, ²hlayinmin@gmail.com

Abstract— These days, the IT world is lit by lighting of the use of virtualization technology in various aspects. We observe that, we can not truly benefit from virtualization if we neglect the process of virtual machine resource allocation. Motivated by this issue, we propose a model of resource management scheme for virtual machines in server virtualization framework. This paper presents a study on hypervisor-based virtualization technology in order to get flexible management of physical resources of a host. In this work, we created three different settings of resource pools for three level of priority of virtual machines in accordance with different level of CPU and memory resource requirement. The results can provide several guidelines for server virtualization administrators.

Keywords— cloud computing, virtualization, hypervisor, resource allocation, resource pools

I. INTRODUCTION

The growing awareness of the advantages provided by virtualization technology is brought about by economic factors of scarce resources, government regulation, and more competition [1]

Virtualization is a combination of software and hardware engineering that creates virtual machines (VMs) which is an abstraction of the computer hardware that allows a single machine to act as if it were many machines. Without VMs, a single operating system owns all hardware resources. But with VMs, multiple operating systems, each running its own VM, share hardware resources. Thus virtualization enables multiple operating systems to run on the same physical platform as shown in Fig 1. It allows you to run multiple operating systems on a single computer including Windows, Linux, and more.

Virtualization is a framework or methodology of logically dividing computer resources such as

hardware, software, time-sharing, and others into separate VMs executing instructions independent from the operation system. The number of VMs that you can run on a single computer depends on the hardware specification [2].




Figure 1. Traditional vs Virtualization architecture

In order to better understand a virtualized computer environment, it is beneficial to compare the basic computer organization of a typical physical computer to that of a computer running a virtualization platform and virtualized environments. The common arrangement of a virtualized computing system is one in which the virtualization platform is installed directly onto the computer's hardware or installed on bare-metal such as VMware ESX. Fig. 2 shows this arrangement.

A key use of virtualization technology is server virtualization, which uses a software layer called a hypervisor to emulate the underlying hardware. This often includes the CPU's memory, I/O and network traffic. By using server virtualization we can utilize the resources up to the maximum use. So these unutilized resources are used by server virtualization by creating multiple VM on the physical machines as shown in above figure. In this way the cost of the hardware as well as software can be reduced up to maximum by demanding server virtualization[3].

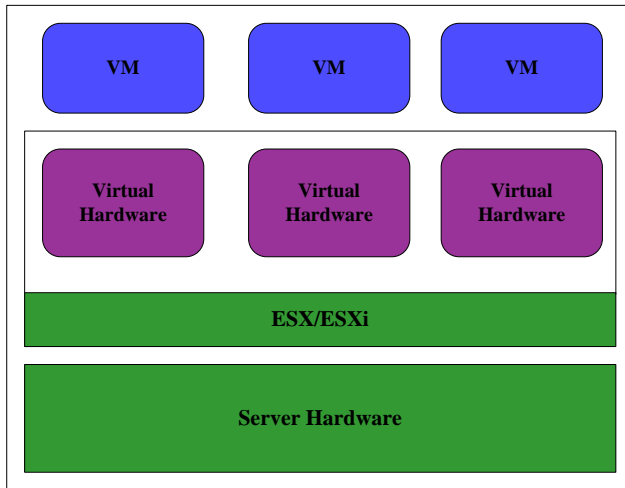


Figure 2. VMware ESX/ESXi Platform

In the context of cloud computing, we understand resource management as the process of allocating computing, storage, networking and (indirectly) energy resources to a set of applications, in a manner that seeks to jointly meet the performance objectives of the applications, the infrastructure (i.e., data center) providers and the users of the cloud resources. [4]. Cloud computing has taken the degree of efficiency and agility realized from virtualization.

Virtualization helps efficient use of hardware resources. Different VMs can run different operating systems and multiple applications on the same physical computer[5]. Hence VMs are allocated to the user based on their job in order to reduce the number of physical servers in the cloud environment. But most VM resources are not efficiently allocated based on the characteristics of the job to meet out Service Level Agreements (SLA) [6].

In this paper, we are going to configure a specific resource management model of ESXi host allocating shares, reservation, and limit.

The rest of the paper is organized as follows. In the next section, Section II, related works regarding resource allocation of server virtualization in cloud computing will be discussed. The resource management concept is discussed in Section III. The setup of the

experimental environment are elaborated upon, followed by some evaluation in detail as in Section IV. Finally, Section V concludes this work.

II. RELATED WORK

Recent studies has primarily focused on allocating resources from resource providers to resource consumers. Aaqib and Amit Chaturvedi [7] have introduced the types of resource pools like physical server pools, virtual server pools, storage pools, network pools. It has been analysed and presented some common resources pool as a study approach.

As per Gaston et al [8], the introduction of virtual server replication with experimental results has been discussed. As per empirical research of Saraswath et al [6], a dynamic VM allocation model based on the characteristics of the job, which can dynamically reconfigure virtual resources has been described to increase the resource utilization. They proposed a new method for execution of high priority job.

This white paper composed by John A. David [9] examines several scenarios based on actual customer implementations where resource pools were expected to be useful. Some scenarios describe examples of poorly configured pools. Those examples include a description of the undesired results and recommendations for correcting the situation. There is also explanation of how Shares, Limits, and Reservations settings on a resource pool are applied.

C.haul [10] invoked the PowerCLI SpongeBob graphic to balance resource pool shares. He introduced the very basic script which connects to the vCenter server and cluster specified and looks at the resource pools within. It then reported on the number of VMs contained within and offers to adjust the shares value based on an input you provide.

III. RESOURCE MANAGEMENT

Resources are the foundational elements that get virtualized by the different virtualization technologies. Resources include the following:

- Hardware such as the CPU, memory, or devices
- The network
- The operating system
- The desktop
- A general software layer

Resource management limits access to shared resources, but it also monitors resource consumption and collects accounting information.

The management of resources is important, because many consumers, such as VMs, zones, containers, or virtual desktops, request resources. Consolidating different workloads on one system often entails combining workloads that have different service level agreements and different needs for throughput, response time, and availability.

But resources are always limited, and they are shared among many virtualization environments on one IT system; therefore, it is important to restrict access to specific shared resources, isolate resources from being used by certain workloads, or at least limit shared resource consumption of workloads. By doing that, we can guarantee a service level for each virtualized environment or influence its performance.

Without resource management, all workloads would be handled equally, based on their resource requests. The result could be that one VM consumes so much memory during runtime that other VMs on the same system get blocked and important memory requests can no longer be served, due to no available memory.

Another example of the importance of resource management is the ability to determine how many resources (for example, CPU) should be shown to or seen by a virtualized environment. This could be important for license or software-behavior reasons.

So the goals of resource management related to zones, containers, or VMs are as follows:

- Prevent entities from consuming unlimited resources

- Be able to change a priority, based on external events
- Balance resource guarantees against the goal of maximizing system utilization

A. *Resource Allocation Shares*

Shares specify the relative priority or importance of a VM. If a VM has twice as many shares of a resource as another VM, it is entitled to consume twice as much of that resource when these two VMs are competing for resources.

Shares are typically specified as High, Normal, or Low and these values specify share value with a 4:2:1 ratio, respectively. It can be selected as Custom to assign a specific number of shares to each VM.

Specifying shares makes sense only with regard to sibling VMs or resource pools, that is, VMs or resource pools with the same parent in the resource pool hierarchical siblings share resources according to their relative share values, bounded by the reservation and limit. When assigning shares to a VM, the priority has to be specified for that VM relative to other powered-on VMs.

B. *Resource Allocation Reservation*

A reservation specifies the guaranteed minimum allocation for a VM.

vCenter Server or ESX/ESXi allows you to power on a VM only if there are enough unreserved resources to satisfy the reservation of the VM. The server guarantees that amount even when the physical server is heavily loaded. The reservation is expressed in concrete units (megahertz or megabytes).

For example, assume you have 2GHz available and specify a reservation of 1GHz for VM1 and 1GHz for VM2. Now each virtual machine is guaranteed to get 1GHz if it needs it. However, if VM1 is using only 500MHz, VM2 can use 1.5GHz.

Reservation defaults to 0. You can specify a reservation if you need to guarantee that the

minimum required amounts of CPU or memory are always available for the VM.

C. Resource Allocation Limit

Limit specifies an upper bound for CPU, memory, or storage I/O resources that can be allocated to a VM.

A server can allocate more than the reservation to a VM, but never allocates more than the limit, even if there are unused resources on the system. The limit is expressed in concrete units (megahertz, megabytes, or I/O operations per second).

CPU, memory, and storage I/O resource limits default to unlimited. When the memory limit is unlimited, the amount of memory configured for the VM when it was created becomes its effective limit in most cases.

In most cases, it is not necessary to specify a limit. There are benefits and drawbacks:

- Benefits — Assigning a limit is useful if you start with a small number of virtual machines and want to manage user expectations. Performance deteriorates as you add more VMs. You can simulate having fewer resources available by specifying a limit.
- Drawbacks — You might waste idle resources if you specify a limit. The system does not allow VMs to use more resources than the limit, even when the system is underutilized and idle resources are available. Specify the limit only if you have good reasons for doing so.

D. Resource Pools

A resource pool is a logical abstraction for flexible management of resources. Resource pools can be grouped into hierarchies and used to hierarchically partition available CPU and memory resources.

Each standalone host and each DRS cluster has an (invisible) root resource pool that groups the resources of that host or cluster. The root resource pool does not appear because the

resources of the host (or cluster) and the root resource pool are always the same.

Users can create child resource pools of the root resource pool or of any user-created child resource pool. Each child resource pool owns some of the parent's resources and can, in turn, have a hierarchy of child resource pools to represent successively smaller units of computational capability.

A resource pool can contain child resource pools, VMs, or both. You can create a hierarchy of shared resources. The resource pools at a higher level are called parent resource pools. Resource pools and VMs that are at the same level are called siblings. The cluster itself represents the root resource pool. If you do not create child resource pools, only the root resource pools exist [11].

Resource pools can be a very powerful tool that helps an organization achieve very aggressive CPU and Memory utilization. Properly configured resource pools ensure that cluster resources are intelligently allocated during times of contention. Critical VMs can be given preferential access to compute in the form of shares and/or reservations, while less critical VMs, with fewer shares, will receive fewer CPU cycles and less access to memory[12].

IV. EXPERIMENTAL SETUP

A. Test Virtualization Environment

Basic components and configuration of virtualization lab has been described in Table 1. We use a single physical server of VMware ESXi host in a private rack in a data center. VMware ESXi provides a virtualization layer that abstracts the processor, memory storage, and networking resources of the physical host into multiple VM. Hence, three VMs of running Windows 2008 Server, Windows 7 and OpenSuSE have been set up on that ESXi host. A single point of control in this data center is provided by vCenterServer which can provide ease of watch over resource utilization of these virtual machines.

The configuration scenario is simple as below. Assume we have three teams of network, software, and testing. The network team uses Windows Server 2008 VM, the software teams uses Windows 7 whereas the testing team uses OpenSuSE VM. Because the network team needs larger amounts of CPU and memory, the software team needs a medium amount resource, but the testing team just runs no special resource requirement, the administrator creates one resource pool for each group.

TABLE 1
ENVIRONMENT CONFIGURATION

Component	Hardware Configuration	Software Configuration
Server	Intel Core i7(4 th Gen)4770/3.4GHz 4GHz, 1TB	VMware ESXi 5.0
Virtual Machine	Intel Core i7(4 th Gen) 4770/3.4GHz 4GHz, 1TB	Windows 7 OpenSuSE 13-2 CentOS
vCenter Server	Intel Core i7(4 th Gen) 4770/3.4GHz 4GHz, 1TB	VMware vCenter Server Windows 2012

The administrator sets CPU Shares to High for the network team pool and to Normal for the software team pool and to Low Shares to low priority testing team pool. So that the network team users can run in its resource. The second resource pool with fewer CPU and memory resources is sufficient for the lighter load of the software staff. Whenever the network team is not fully using its allocation, the software team can use the available resources. If the allowable resource allocation of the software team is not fully utilized, the testing team can get the available resources. This scenario is shown in Fig 3. The numbers show the effective allocations to the resource pools.

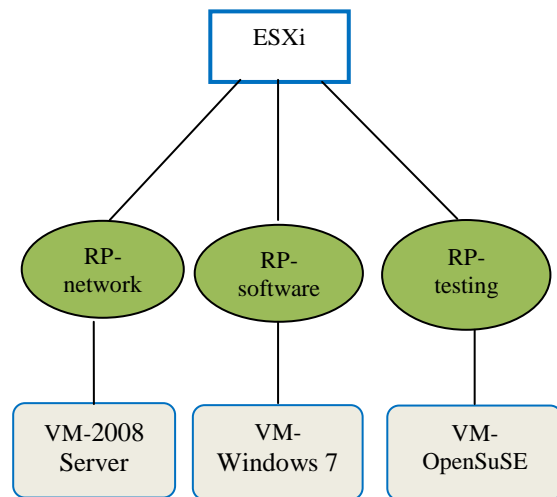


Figure 3. Test Network Scenario

B. Result and Discussion

After setting each VM to their respective resource pools with High, Normal, or Low, all the VMs are powered on. Then we test with CPU busy script to confirm whether the resource allocation is in harmonize. The following graphs show the CPU consumption and memory usage of each VMs on each resource pool.




Figure 4. CPU Consumption of Windows Server 2008



Figure 5. CPU Consumption of Windows 7

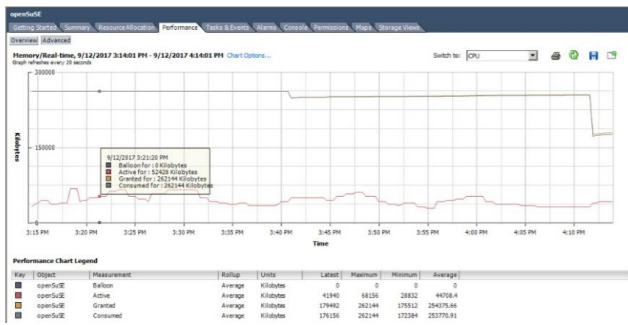


Figure 6. CPU Consumption of OpenSUSE Linux



Figure 7. Memory Usage of Windows Server 2008



Figure 8. Memory Usage of Windows 7



Figure 9. Memory Usage of OpenSUSE Linux

As the physical host we used is a multicore processor, all four cores support harmonically

and the allocation of resource is not a problem, so there is no significant impact on CPU resource usage. To get support of a single CPU processing, we configure the multicore support to ONE in BIOS setup setting. Then the obvious result. Then the obvious result can be seen after restarting the ESXi host and all VMs.

V. CONCLUSIONS

In this work, we have implemented the evaluation environment and conducted resource allocation tests on three different OSs virtual machine on a single ESXi host. The result shows that it can successfully manage virtual machine resources in a server virtualization framework. However, our work can still be improved in several aspects. In the future, we are planning to learn on aggregation of resources which may be configuring host resources to part of a cluster resources. That may be a way of creating cluster-wide resource pools and enforcing cluster-level resource allocation policies. We believe that more interesting result can be discovered.

REFERENCES

- [1] Thomas Burger "The Advantages of Using Virtualization Technology in the Enterprise", Intel Developer Zone, MENU documentation, 2012.
- [2] O.Kulkarni, S.Nagul, D.Gawali, P.Kumar, "Virtualization Technology: Leading Edge," IJCA, Issue 2, Volume 2, ISSN:2250-1797. 2012.
- [3] Y.Ghorpade, T.Bennur "Server Virtualization: A Cost Effective and Green Computing Approach Towards Educational Infrastructure Management," Volume 1, Issue 3, May 2013.
- [4] B. Jennings, R. Stadler, "Resource Management in Clouds: Survey and Research Challenges," Journal of Network and System Management manuscript, 2013.
- [5] L. Langid, N. Sharma. K. Rathi, "A Survey on Secure the Cloud Environment using Hypervisor-based virtualization Technology," IJICSE, Volume 3, Issue 3, 2014.
- [6] A.T.Saraswathi, Y.R.A.Kalaashri, S.Padmavathi "Dynamic Resource Allocation Scheme in Cloud Computing," Procedia Computer Science, volume 47, 2015.
- [7] A. Rashid, A. Chaturvedi "A Study on Resource Pooling, Allocation and Virtualization Tools used for Cloud Computing," Volume 168-No 2, June, 2017.
- [8] G. Keller, H. Lutfiyya "Dynamic Resource Management in Virtualized Environments through Virtual Server Allocation" Volume 3, No 3-4, 2010.

- [9] John A. Davis, "Recommended Use of Resource Pool in VMware vSphere DRS Clusters", Expert Reference Series of White Papers.
- [10] C. Hhal, "Understanding Resource Pool in VMware vSphere," General Tech, February 1, 2012.
- [11] https://pubs.vmware.com/vsphere-4-esx-vcenter/index.jsp?topic=/com.vmware.vsphere.resourcemanagement.doc_41/getting_started_with_resource_management/c_shares.html.
- [12] T. Ferris, "When Bad Resource Pools Happen to Good People," Green Pages Technology Solution, March 1, 2012.