Scientific Bulletin of the „Petru Maior" University of Tîrgu Mureş
Vol. 10 (XXVII) no. 1, 2013
ISSN-L 1841-9267 (Print), ISSN 2285-438X (Online), ISSN 2286-3184 (CD-ROM)

VIRTUALIZATION, The next step for online services

Cristian Nicolae MĂRGESCU[1], HALLER Piroska[2]

"Petru Maior"University of Tîrgu Mureş

Nicolae Iorga street, no.1, Tîrgu-Mureş, Romania

[1]cristian.nicolae.margescu@stud.upm.ro
[2]phaller@upm.ro

Abstract

Virtualization allows sharing and allocating the hardware resources to more virtual machines thus increasing their usage rate. There are multiple solutions available today such as VMware vSphere, Microsoft Hyper-V, Xen Server and Red Hat KVM each with its own advantages and disadvantages. Choosing the right virtualization solution largely depends on the used applications and their resources requirements. The comparative analysis of the available virtualization solutions shows that it is essential to establish performance criteria's and minimum and maximum resources usage thresholds over a given period of time. The coexistence of different services in different virtual machines that use different amount of resources allows a more efficient use of the available hardware resources.

Keywords: Virtualization, Cloud Computing, Distributed Systems, Hypervisor

1. Introduction

Diversity in the software world is a reality and the majority of applications are depended to a certain operating system, to a certain library. It is common in public institutions to see numerous computers running different operating systems due to the extremely cost of the application's migration to a new computer, new platform.

Server virtualization improves the effectiveness and the availability of IT resources and applications; it creates an environment where the speed and scalability with which a deployment to a service is made can be broken down to a few simple steps. It frees system administrators of repetitive administrative tasks that can lead to security gaps. Due to the drop of equipment prices and the increases of the server's performances, the market of the virtualization solutions has extended rapidly making available a series of virtualization platforms.

Choosing the correct virtualization solution and configuring the systems largely depends on the applications that are going to be used, on the amount of resources needed and on identifying the bottlenecks and the scalability of the system. This paper aims to perform a detailed analysis of the different solutions available from a performance point of view also taking into account the fault tolerance and scalability.

Until recently it was common practice to use a single server on a single machine but with each application using the minimum amount of resources needed, the system administrator can join multiple servers into a single machine running a number of virtual services. For companies using hundreds of servers the need for physical space can drop considerably.

Another important factor that led to the wide spread of virtualization is REDUNDANCY that became common practice without the need to acquire new hardware. By running the same type of application on multiple servers, the safety of the system is increased and offers protection in case of a software error.

Virtualization can offer programmers an isolated and independent environment where they can develop and test applications and even operating systems, without the need of acquiring a dedicated machine. An administrator can create a virtual server independent of the physical machine where he can create a deployment without affecting other applications.

Server hardware will become, regardless of technology at some point obsolete and migrating from one system to another can create difficulties in the continuity of services. Virtualizing the machine to modern servers ensures the continuity of the service, the transition being unnoticeable; applications will behave like nothing has changed. Migration in virtualization has become common practice today. It offers the possibility to move a virtual server environment from one virtual environment to another regardless of the location. In the past this was possible only with servers that had the exact same architecture.

2. Existing platforms

When making the decision to move public services to virtual platforms the first step is to choose the optimal Business class platform the company requires. There are four major competitors: VMware vSphere, Microsoft Hyper-V, Citrix Xen Server and Red Hat KVM each with its own advantages and disadvantages. Ten years ago the argument around virtualization was simple and orbited around VMware as it was the only service available. Nowadays its reign is put to the test by Microsoft, Citrix and Red Hat.

The last version of Hyper-V from Microsoft brings substantial improvements compared to VMware detailed in Aidan Finn's paper [1]. Hyper-V has a few extra facilities compared to VMware such as Server Attached Storage, the possibility to migrate live between websites, failover capability and the ability to access more virtual CPU's and memory. These benefits of these aspects can be seeing in the progress made by the number of licenses issued in 2012. Hyper-V grew 62% last year compared to ESX and Citrix that grew only 21% and 25% according to IDC (International Data Corporation) [2]. The yearly increase is also due to the fact that the number of x86 servers converted to virtual machines has increased dramatically. According to Gartner [3] the number of x86 virtual machines has increased from 11 million in 2009 to 55 million in 2010 and will reach 100 million in 2013.

Meanwhile Citrix, which has a market share approximated at 6% has reduced virtual desktop costs by using physical servers more efficient allowing them to be cheaper than a traditional desktop and still offer increased security. The company came out with a new product called XenServer 6.0 with better support in processor sharing between virtual machines and up to 16 virtual processors and 128 GB RAM virtual memory. Citrix continues collaboration with Microsoft to support joint management of the two technologies.

Red Hat, throw acquiring Qumranet has won a virtualization technology based on Linux which will prove to be the obvious choice for those using Open Source. KVM (kernel-based virtual machine) implemented in the operating system is a unique solution allowing the OS to act as a hypervisor.

VMware was the first skeleton of a Hypervisor x86 in 2001 currently at the 5th generation. It was conceived with a single goal: virtualization. Competitors have created supervisors based on different partitions usually under the form of a parent partition. Being built with a single goal ESXi completely eliminates the need for an operating system, thus eliminating the millions of lines of code with a 144MB footprint, reducing the range of attack at a code level and drastically reduces the number of update packages. VMware integrates different security systems in the supervisor, with innovative solutions dedicated to virtual environments.

3. Virtualization architecture

XEN [4] provides a structure based on a Linux distribution that lies above the manager layer, processing all of the system calls made by one of the virtualized operating systems. XEN Virtualization provides a method called Para-Virtualization by abstracting the physical layer below and provides hardware access to multiple virtual machines.

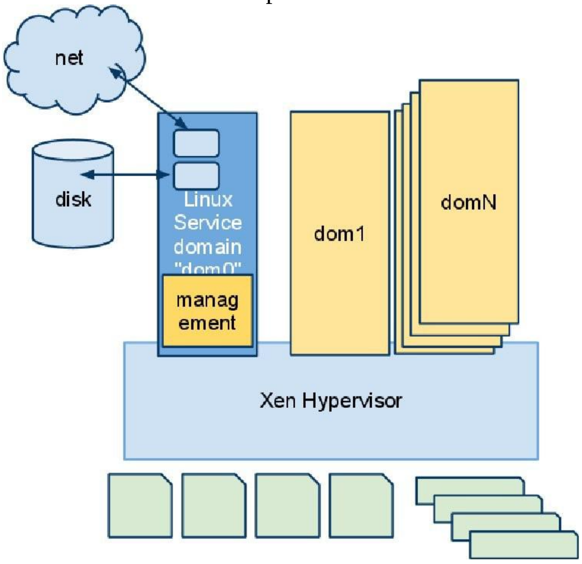

Figure1. XEN Architecture

As shown, the VM (virtual machine) relies exclusively on the operating system to communicate with hardware. System calls being made in a User Space mode that offers a high degree of security but can facilitate a quick overload of I/O requests of the CPU due to the Kernel and the process management mode.

The same idea was used to develop the Hyper-V manager [5]. Resources are monitored and distributed by the management system to each virtual machine.

Performances are directly influenced by the resources available in the operating system. It relies on the operating system installed on a parent partition. The virtualization stack runs in the parent partition having direct access to hardware and the virtualized system runs in a logical partition without access to the processor or memory and without handling interrupts. It runs in a virtual region of the memory isolated from the physical side.

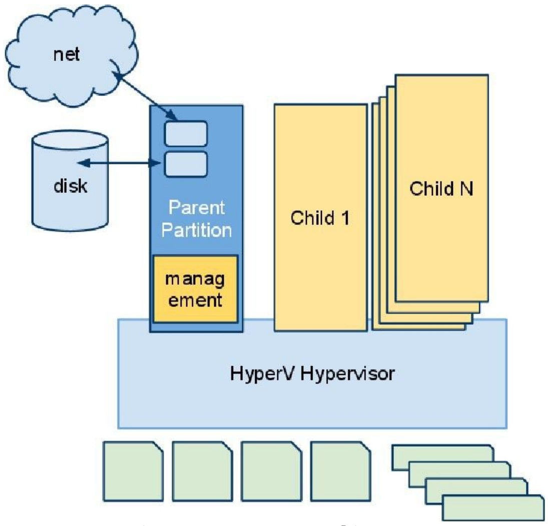

Figure 2. Hyper-V Architecture

ESXi [6] architecture proposes a different approach based on a layer within the stack called VMkernel. All the processes run throw it including management applications, agents and virtual machines. It has total control of the server hardware and handles the management of application resources.

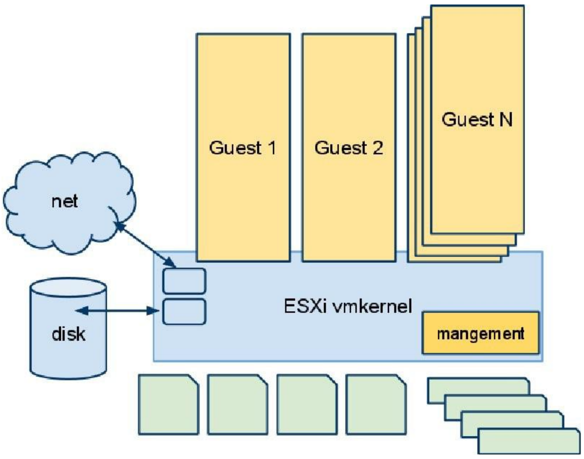

Figure 4. VMware ESXi Architecture

The main processes that run on him are:

Direct console User Interface – low level configuration manager.

Virtual machine monitor (VMM) – It assures the environment of execution of a virtual machine and *Virtual Machine Executable* (*VMX*), the last one is responsible for communicating with user interfaces.

Common information model (CIM) – The Interface that initialize the management on hardware level for remote applications via an API.

RedHat is proposing a Kernel-Based virtual machine (KVM) [7]. It is implemented as a kernel module which transforms the Linux Kernel in a hypervisor. KVM is relied on two principles:

Mainly because of the fact that it was developed after hardware virtualization, it takes full advantage of Intel-VT – X and AMD – V and it is fully optimized whiteout requiring any modifications on the OS level.

Secondly a VMM requires a memory manager, processes scheduler, I/O Stack, drivers etc.

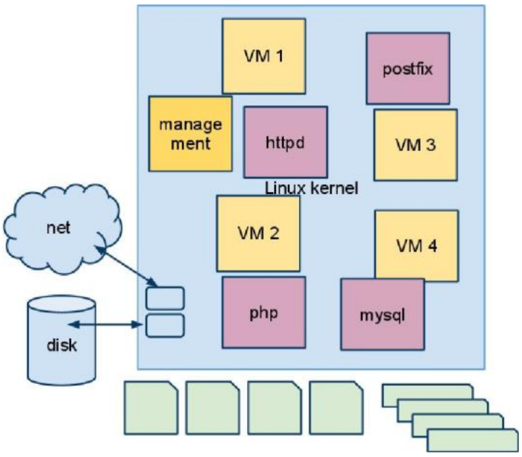

Figure 3. KVM Architecture

It is actually a specialized operating system which does not run applications but virtual machines. Inside KVM a virtual machine is implemented as a normal process and selected to run by the standard scheduler of the kernel. Device emulation is offered by a modified version of the QEMU.

4. Performance evaluations

The main aspects that raise a question and they must be of concern in the moment that we are faced with the decision on moving into virtualization are how many resources are consumed by the execution of the virtualization environment instead of the application themselves. We have executed a series of tests under the same hardware and software conditions to identify how many resources are consumed by a Hypervisor (VMM) and to identify the benefits of their usage. The tests are of two types of services: intensive computational applications and interactive file transfer. The target is to load the system to the extreme to better observe the border between a virtualized OS and a VMM. The machine is an Intel Xeon 8x2.5 GHz, 12 GB RAM, and 160 GB HDD on which it is installed ESXi 5.0. On the VMM it created a Resource Pool of 8CPU's, 8GB Ram, and 130 GB HDD. Inside the resource pool are installed two Server 2008 operating systems each with 4 Core, 4 GB Ram and 60 GB HDD. The first server will offer a transcoding service based on IIS Smooth Streaming to Silverlight clients and the second will offer interactive movie transfer application with a series of movie playlists to web clients. For both services it will be observed CPU Load, Memory Load, the Transfer Rate on to the HDD and average response time. The two types of application that I chose have a totally different behavior and resource consumption, thus allowing us determine the optimal settings of the resource quantum.

4.1 Analyzing data

The tests are performed under the same conditions as follows: on start there are ten clients connecting to each service and after each 20 seconds another client is attempting to connect until the limit of 30 clients. It is important to identify the system capabilities.

Figure 5 is indicating the total resources consumed by the Virtualized System (VS) and the VMM together for the first service.

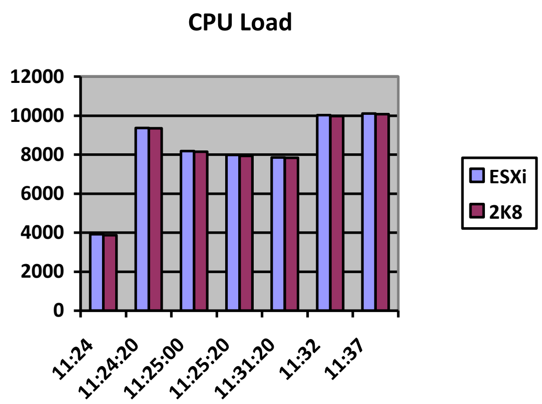

Figure 5.Transcoding CPU Load (MHz)

As you can observe in the moment the workload started the CPU is at 3912 MHz representing around 39.12 % of the maximum value. Increasing the number of transcoding requests it can be observed an increase to 9362 MHz representing around 92% load after which it stabilized on 7800MHz. In this period the VMM has a maximum consumption of 166 MHz of CPU which leads us to the conclusion that on a full virtualized environment the resource consumption of a VMM is negligible and on today's machines that quantity is insignificant and it leads us to think further on how efficient we can distribute resources inside a virtual environment.

By setting a maximum of CPU resources offered to a Virtualized System we can assure a stabile running environment without losing quality of service towards the client even if there is more than one Virtualized System running in parallel on the same platform. Inside resource pools we can set the CPU affinity which means that VS with 4 cores can execute processes on all available cores on the physical machine without disturbing other VS that requires fewer resources. The detailed analysis of the gathered data shows that the workload is uniformly distributed on the physical platform which can lead to a 20 – 25 % increase in processing speed even if the Virtualized System has only 2-Cores available.

The proposed method on resources organization is the creation of resource pools, where each VS inside the pool, can have access to a limited number of resources but each can benefit of an increase of allocated resources in the critical moments when are needed. It can be observed a smooth increase of resource usage in the moment when HDD write latency appears and the other system does not utilize more than 10% of its reserved resources. To set the maximum threshold of allocated resources will perform detailed running test of the Virtualized System with its services. It is measured the optimum value of consumed resources and it is set the maximum threshold that it can reach on maximum load, thus determining the optimum resources needed from a resource pool.

Analyzing the second type of service which will offer video streaming we can be observe the same results regarding the footprint of VMM even in the case of low CPU consumption. The measurements show that the presence of an OS inside a virtualized environment may bebenefical from the resources point of view. The main benefits can be observed on the maintenance and running costs which can be reduced by up to 60% in a time frame compared to a dedicated system.

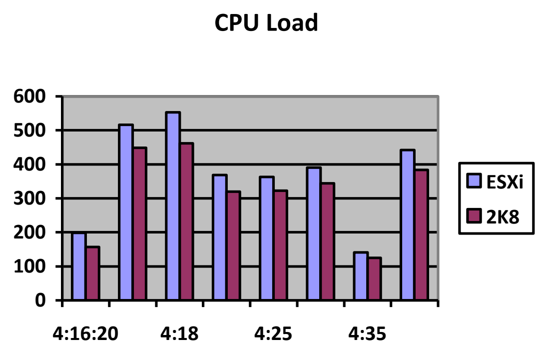

Figure 6.Web Server CPU Load (MHz)

Memory and Virtual disk management inside a virtualized system is done mainly by setting maximum thresholds, but in the moment that these limits are reached by the VS, the VMM choses an optimum method of resolving this issue depending on the administrator configuration. The most frequent methods are to create a buffer in RAM for the data that are to be written on the Hard Drive. Using this technique, even when the physical HDD has reached extreme write latency's the VS is not affected by them having the felling that the write or read of data is in normal parameters. Thus the Hard Drive's I/O channel is freed facilitating the read of data from him without disrupting the execution of the running services; afterwards the data are flushed on to the HDD. The next figure exemplifies the read/write rates on the physical HDD in a giving amount of time.

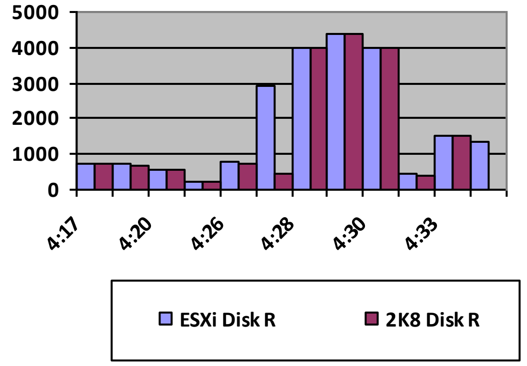

Figure 7. Disk read and write rates (kbps)

The latency's emerged on the HDD level are automatically detected by the VMM and it is starting to take the appropriate measures to avoid an eventual congestion. The next table indicates the read rate from the physical HDD expressed in kbyte/s in a giving amount of time.

Table 1. Read rate for different services

Transcoding	Webserver	Total	Percentage
1438	2430	3868	37.8/62.8
7454	3968	11422	65.3/34.7
3050	160	3210	95.1/4.9
3284	367	3651	90/10
3355	44	3399	98.8/1.2
2438	1162	3600	67.8/32.2
120	6494	6641	2.3/97.7
97	9180	9277	1.5/98.5
1927	1954	3881	49.7/50.3
1189	1954	3143	37.9/62.1
2984	2099	5083	58.8/41.2
4052	1952	6004	67.8/32.2
2652	361	3013	82.1/11.9

In the same amount of time the latency's on the virtual disk are in the region of 4 to 26 msec. per VS. We compare the latencies measured on each virtual disk with the latencies measured directly on the physical HDD to determine the exact impact of each Virtualized System to the physical machine. Comparing the latency table and memory load table we can observe that in the moment that the write latency becomes greater than 16ms the VMM allocated more RAM to that specific VS creating a memory buffer and decongestion the HDD so the memory usage increases from 4096MB allocated memory to 4194MB. An increase of only 1.2 %, having a minimal impact on the physical machine, has a major performance impact on the Virtualized System.

Table 2. Latencies for different services

Transcoding Virtual Disk	Web Server Virtual Disk	Physical HDD
8ms	0ms	10ms
4ms	0ms	5ms
18ms	0ms	19ms
6ms	0ms	8ms
12ms	0ms	16ms
11ms	0ms	15ms
10ms	0ms	12ms
10ms	0ms	13ms
15ms	0ms	16ms
14ms	0ms	16ms
11ms	0ms	14ms
26ms	0ms	29ms
12ms	0ms	16ms
9ms	0ms	12ms

We can clearly observe that the coexistence of two or more Virtualized System under the same resource pool can be done because of the fact that with a good resource planning and optimization of usage we can provide an optimal platform with a minimal cost of maintenance and deployment.

On a memory level the allocation of resources it is done to the maximum configured on each Virtualized System and in special cases the VMM can allocate a surplus to overcome a critical situation. In our case we can observe that for transcoding service have been allocated 4096 MB RAM but in the moment that congestion had appeared on the HDD I/O level VMM allocates the extra needed memory for HDD Buffering (figure 8.) .

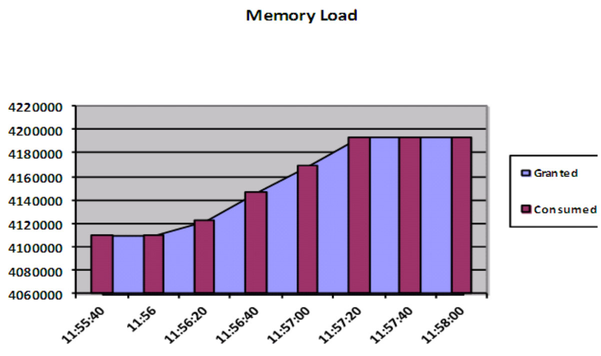

Figure 8. Memory usage (KB)

4.2 Problems detected

The main problem that I have noticed in a virtualized environment is the number of Input Output operations generated by the hosted operating systems towards the HDD, thus creating a bottleneck and after it resulting in an unresponsive system. To overcome this problem many of the IT industry leaders have researched the integration of flash drives in enterprise level systems. The study of the flash drives performance in a virtualized system is the topic for my next work in progress.

5. Conclusions

The paper outlines the advantages and disadvantages of running an OS and different type of services inside a virtualized environment. The design of a virtualized environment in which to host two or more OS's on maximum performance and with minimum cost in which to run their installed services requires a thorough analysis. It is essential to establish the minimum and maximum thresholds of consumed resources, the time interval that the systems are at maximal load and the metric with which we can establish the Resource Pool with sufficient resources to host our needs. Most important is to determine the maximal number of clients that are going to use one instance of the service, to provide the platform in which everything runs smoothly.

Today's tendencies to cut down on costs in online services together with the increase of computational power have created a pressure over the industry to reduce costs and so it was introduced the term virtualization. The statistics have shown that around 60% of computational resources are unused because of the dedicated systems, configured to support their maximum load even if this are present in around 5% of the time. More and more x86 systems are virtualized to reduce the costs of maintenance.

References

[1] *Aidan Finn ,Mastering Hyper-V Deployment*, John Wiley & Sons, 09.11.2010.

[2] *International Data Corporation* , http://www.idc.com

[3] *Gartner research institute* http://www.gartner.com/.

[4] Paul Barham, Boris Dragovic, Keir Fraser and Steven Hand , *Xen and the art of virtualization*, ACM SIGOPS Operating Systems review, Volume 37,issue 5, December 223.

[5] *Liang Liu, Hao Wang, Xue Liu, Xing Jin: GreenCloud, A new architecture for a green data center*, ACM New York 2009.

[6] *VMware White Paper: The Architecture of VMware ESXi*, Revision: 20081024 WP-030-PRD-02-02**,** 2007 VMware Inc.

[7] *RedHat White Paper, KVM architecture:* June 27, 2007, Ottawa, Ontario Canada.

[8] Springer Berlin, *Performance evaluation of HPC Benchmarks on VMware ESXi Server*. Heidelberg, Berlin 2012

[9] Open MPI: *Open Source High Performance Computing* , 2011

[10] Waldspruger, C.A., *Memory resource management in VMware ESX Server*, OSDI 2002