ISSN: 2278 – 1323
International Journal of Advanced Research in Computer Engineering & Technology (IJARCET)
Volume 2, Issue 2, February 2013

A cost effective Online Examinations model using server virtualization with virtual machine clusters.

Balaji Kandregula,Associate Professor, Deaprtment of Computer Science & Engineering, Dadi Institute of Engineering & Technology, Anakapalli.

ABSTRACT

Over the past 5 years, most of the examinations conducted by state level and country level firms are becoming online. The costs related to implementing and managing data centre resources can be a significant drain on organization‘s overall budget. Virtualization is one of the most significant technologies to impact computing in the last few years. This paper explores about virtual machines that has the ability to run multiple loads on the same physical computer at the same time. which leads to increased hardware utilization and a number of additional benefits, and on the approach can be applied by running a single workload spread across multiple physical computers commonly referred as clustering. This paper focuses on Server virtualization technology with clustering that can be implemented to increase the scalability, performance and reliability in a cost effective manner.

Keywords

Server Virtualization, Clustering, Datacenter, Virtualmachine

1. INTRODUCTION

Virtualization is one of the more significant technologies to impact computing in the last few years. With roots extending back several decades, today its resurgence in popularity has many industry analysts predicting that its use will grow expansively in companies over the next several years. Promising benefits such as consolidation of infrastructure, lower costs, greater security, ease of management, better employee productivity, and more, it’s easy to see why virtualization is poised to change the landscape of computing. Virtualization refers to the process of decoupling the hardware from the Operating system on a physical machine. It turns what used to be considered purely hardware into software. Put simply, you can think of virtualization as essentially a computer within a computer, implemented in software. This is true all the way down to the emulation of certain types of devices, such as sound cards, CPUs, memory, and physical storage. An instance of an operating system running in a virtualized environment is known as a virtual machine. Virtualization technologies allow multiple virtual machines, with heterogeneous operating systems to run side by side and in isolation on the same physical machine. By emulating a complete hardware system, from processor to network card, each virtual machine can share a common set of hardware unaware that this hardware may also be being used by another virtual machine at the same time. The operating system running in the virtual machine sees a consistent, normalized set of hardware regardless of the actual physical hardware components.

In this paper, I present the server Virtualization Architecture Later in detail I discuss about current online examination model and its drawbacks in terms of cost and maintainability. Then I discuss about server virtualization with clusters of virtual servers to show case a new cost effective online examination model.

1.1 SERVER VIRTUALIZATION

One inescapable "cost of admission" related to working with virtualization is that there will always be some amount of performance overhead. In theory, the simplest and most efficient method for implementing virtualization is to run directly on the host computer's hardware. The idea of server-level virtualization is just that: Rather than requiring a full-fledged host OS, it includes a thinner layer of virtualization management software that handles basic hardware management functions. The concept is that—because many of the features and functions of the host OS are unnecessary for virtualization—there is no need to waste resources on it. Instead, drivers and system services can be optimized primarily for the purpose of supporting virtual machines. This approach is also referred to as a Hypervisor—a type of system supervisor that sits above the hardware and between the virtual machines that are running on the system. Figure 1 provides an architectural overview of this model

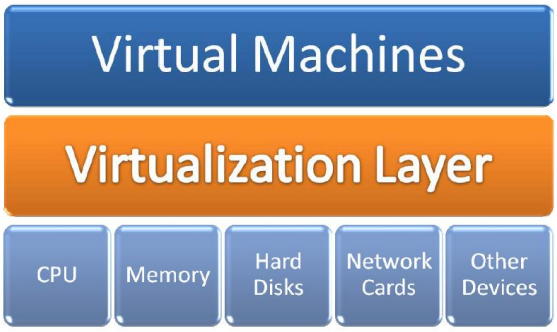

Figure- 1: Server Virtualization

2. CURRENT SOLUTIONS

The current online examination model is based on either client/server model or a mainframe based model. Mainframe model is more like a centralized environment in which both web server and database server are deployed on a single physical machine as shown in figure -2

DFigure-2: Mainframe Approach

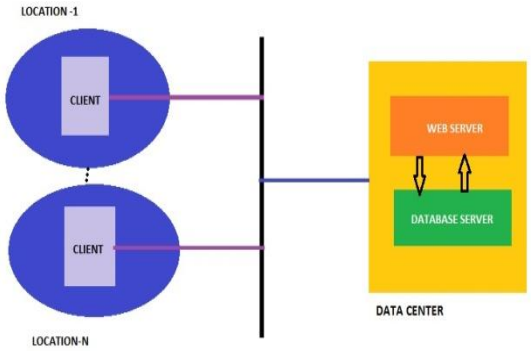

As mainframes are more legacy, integrating new technologies are ardous. To sustain the load due to more number of users we have to deploy regional number of servers which increases the cost of infrastructure.

2.1 Client –server model:

In client- server model regional level proxy servers will act as offline clients which maintain the local dumps of examination modules to be conducted in particular regions. The local dumps will be updated in the mainframe server which is located centrally on timely basis .With increased load a single mainframe server cannot furnish more number of regional proxy servers. The proxy servers either they have to be fixed permanently or moved physically on demand or it also increases maintenance and operational cost.

3. PROPOSED MODEL

Server Virtualization combines the best economic properties of data centre and client/server computing. The mainframe era was characterized by significant economies of scale due to high up-front costs of mainframes and the need to hire sophisticated personnel to manage the systems. As required computing power – measured in MIPS (million

instructions per second) – increased, cost declined rapidly at first (Figure-3), but only large central IT organizations had the resources and the aggregate demand to justify the investment. Due to the high cost, resource utilization was prioritized over end-user agility. Users' requests were put in a queue and processed only when needed resources were available.

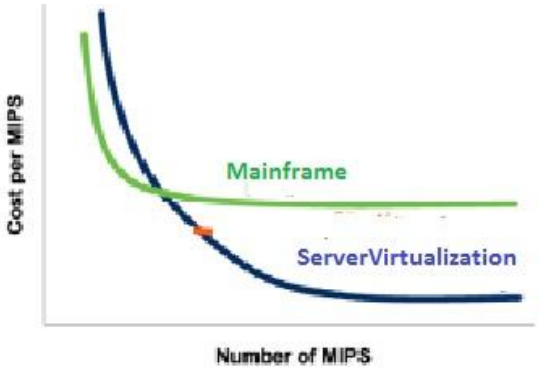

Figure-3 MIPS vs Cost per MIPS

Clustering is designed to ensure that no data is lost in the event of a software or hardware failure. Clustering has typically been offered by application vendors as an add-on to their base product, with some attendant drawbacks like extra expense, redundant solutions, and infrastructure complexity.

Part of the extra expense reflects the fact that you need extra hardware standing by, with the mirrored system on standby, ready to take over should the primary system fail. You don't need to be a genius to recognize that buying a second set of hardware makes clustering significantly more expensive. However, if you've got millions of U.S. dollars worth of transactions occurring on your system, keeping a redundant server ready may be a worthwhile investment. And techniques exist that can allow you to run other work on the stand-by server until it is needed. How does clustering work? Essentially, the coordinating virtualization software runs two VMs on separate machines. The VMs are identical in terms of the OS and application configuration, but differ, naturally, in the details of their network connections and local hardware. The virtualization supervisor constantly communicates with the clustered VMs to confirm they are working (this is usually referred to as a *heartbeat*, signifying the continued existence of the entity).

One VM is the primary server and is the system that users interact with, and the second VM serves in a backup capacity, ready to stand in should the primary server go down. The primary server constantly sends any changes to the secondary server so that its state reflects that of the primary VM at all times. If the primary VM goes down, the virtualization supervisor notes its unavailability, and transparently switches users to the backup server. New users connecting after the switch don't see anything different — they're just connecting to what looks like the same application, and are unaware it's running on a different VM. Users that have been connected to the original VM that is no longer available also are unaware of the switch, because the virtualization software has been sending their state to the secondary machine all along. They may notice a short break in responsiveness while the switch is made, but it's usually so quick that no-one notices.

Figure-4 server virtualization with clusters of virtual servers

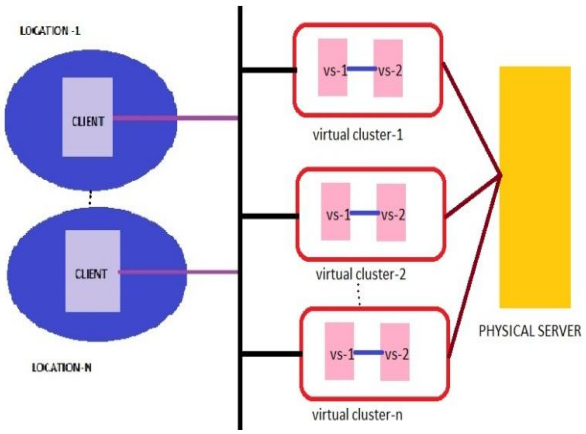

4. CONCLUSION

The proposed model provides a solution to conduct multiple examinations of different firms simultaneously by using virtual machine clusters which reduces the operational and maintenance overhead as illustrated in figure-3.

REFERENCES

[1] Murat Kantarcioglu, Alain Bensoussan, SingRu(Celine) Hoe,Impact of Security Risks on Cloud Computing Adoption, Forty-Ninth Annual Allerton Conference Allerton House, UIUC, Illinois, USA September 28 - 30, 2011

[2] M. Rosenblum, T. Garfinkel. ―Virtual Machine Monitors: Current technology and future trends‖ , Computer, 38(5), Los Alamitos, CA, IEEE Computer Society Press, pp. 39-47, 2005 .

[3] The Open Group, ―Service Oriented Infrastructure Reference Framework‖ , http://www.opengroup.org/projects/soasoi/uploads/40/19218/soi-V1-5-P1.pdf , 2009

[4] Khanvilkar, S., Khokhar, A., "Virtual Private Networks: An Overview with Performance Evaluation", IEEE Communication magazine, vol. 42 (10), pp. 146-154, October 2004

[5] M. Rosenblum and T. Garfinkel. Virtual Machine Monitors: Current Technology and Future Trends. IEEE Computer, 35(5):39-47, 2005

[6] Kyle E. Stewart, Jeffrey W. Humphries, and Todd R. Andel Developing a Virtualization Platform for Courses in Networking, Systems Administration and Cyber Security Education

[7] Anil Desai, The definitive guide to virtual platform management, RealTime Publishers