J. Basic. Appl. Sci. Res., 2(5)4910-4920, 2012
© 2012, TextRoad Publication

ISSN 2090-4304
Journal of Basic and Applied Scientific Research
www.textroad.com

VIRTUALIZATION: A COST EFFECTIVE APPROACH TO INFRASTRUCTURE MANAGEMENT IN DATA CENTER NETWORKS.

(A Case for Elastic Sky X Server)

[1]Okafor Kennedy, [2] Udeze Chidiebele. C,. C, [3]Prof. H. C. Inyiama, [4]Dr. C.C Okezie

[1&2]Electronics Development Institute (FMST-NASENI), Awka. Nigeria.
[3&4]Department of Electronics and Computer Engineering, Nnamdi Azikiwe University Awka.

ABSTRACT

Business continuity is a priority objective to organizations running data center networks. Hence, operations must maintain a near zero downtime running 24x7 in this type of network. Best practices suggest using infrastructure redundancy to establish multiple data centers each with replicated operating systems, applications and data services. Does this redundancy address cost-effective utilization of IT infrastructures; responsiveness in supporting new business initiatives; and flexibility in adapting to organizational changes. This paper adopted a candidate scheme known as infrastructure data center virtualization (IDCV) in an elastic Sky X (ESX) server architecture to fix the challenges of infrastructure duplication in data center networks thereby saving cost at large. Using virtualization for different service agents, it is possible to achieve a redundancy system for all the services running on a data center. This new approach of using a software application called a hypervisor to divide one physical server into multiple isolated virtual environments called virtual machines (VMs) provides high availability and cost reduction for IT infrastructure by exploiting the features of the virtualization layer between physical hosts.

Key words: Business Continuity, Infrastructure Redundancy, IDCV, Virtualization, Virtual Machines.

INTRODUCTION

Contemporarily, Modern data centers are comprised of tens of thousands of servers, and perform the processing for many enterprise as well as Internet business applications, such as those used by financial institutions, Google, twitter, facebook etc. However, the cost of deployment as well as maintaining the IT infrastructure in data center networks is very immense and hence calls for a better approach to cost reduction as well as service availability. Today several operating systems (OS) are used for high performance computing (HPC): Linux on clusters, CNK on Blue Gene/L [1], and catamount on Cray [2, 3]. Virtualization offers a comprehensive remedy to the cost of OS setup and management in an enterprise data center network.

Internet and business applications are increasingly being moved to large data centers that hold massive server and storage clusters. Current data centers can contain tens of thousands of servers, and plans are already being made for data centers holding over a million servers [4]. Virtualization affects datacenter networks in two important ways. The first is the demand for considerable bandwidth. Consolidated server platforms require higher bit-rate connections to support multiple processes. Storage further adds to the demand for bandwidth. Multimedia content and accumulated application data inflate storage overhead, and more data directly drives the need for high-speed access. The result is an accentuated need for a high-bandwidth datacenter network.

The massive amounts of computational power required to drive these systems results in many challenging and interesting distributed systems, as well as IT cost resource management problems. In this paper, server virtualization is proposed as a remedy to infrastructure demands and also an attempt is made to evaluate virtualization vis-a-vis data center domain, with a particular emphasis on how new virtualization technologies can be used to simplify deployment, improve resource efficiency, and reduce the cost of reliability.

***Corresponding Author:** Udeze Chidiebele. C, Electronics Development Institute (FMST-NASENI), Awka. Nigeria.
E –mail address: udezechidi@yahoo.com

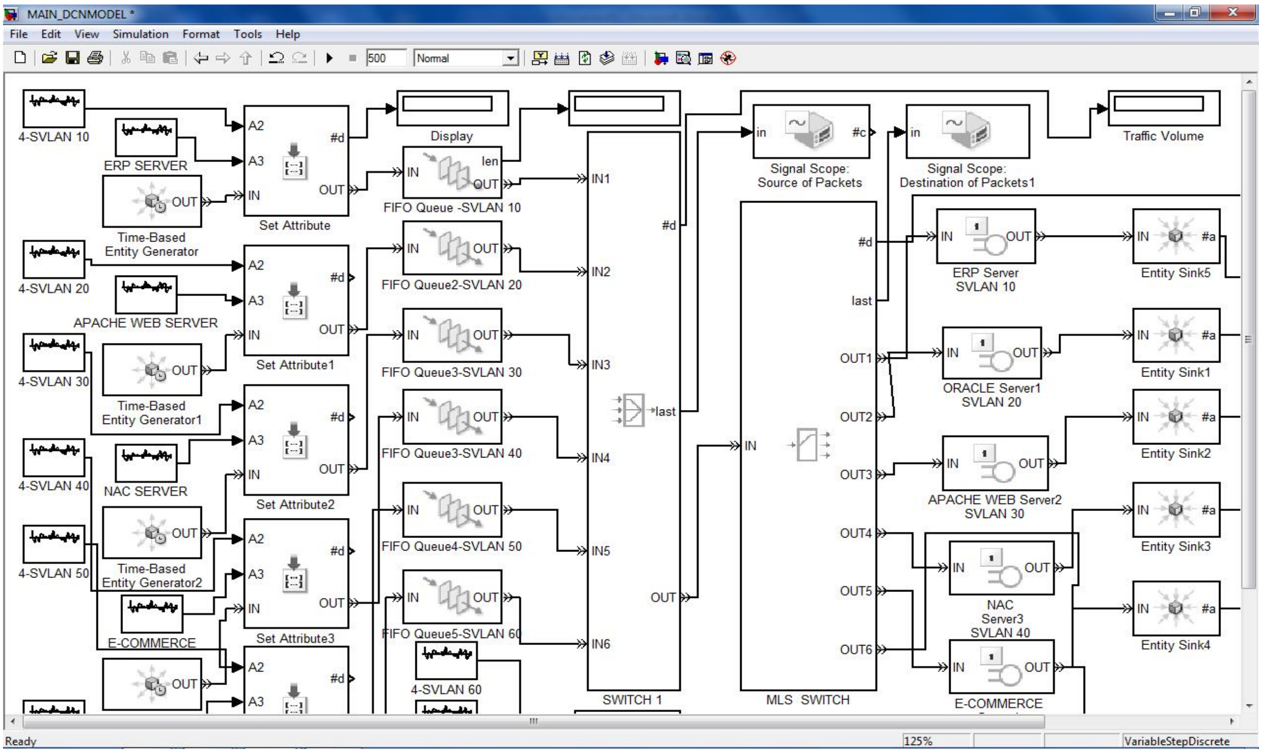

Fig. 1: An Unvirtualized Datacenter Process Model with MATLAB Simevent.
(**Source**: Okafor kennedy C.,A synthesis VLAN approach to congestion management, 2010).

Fig. 1&2 depicts a star topology model runing enterprise applications on a separate web server (Apache),enterprise server (SAP ERP), network admission control server (NAC), E-commerce server in a datacenter model.This paper observes that the cost of infrastructure management for the on-premises applications and services is very high (capital expenditure). This is a major concern to IT administrators and bussiness managers.

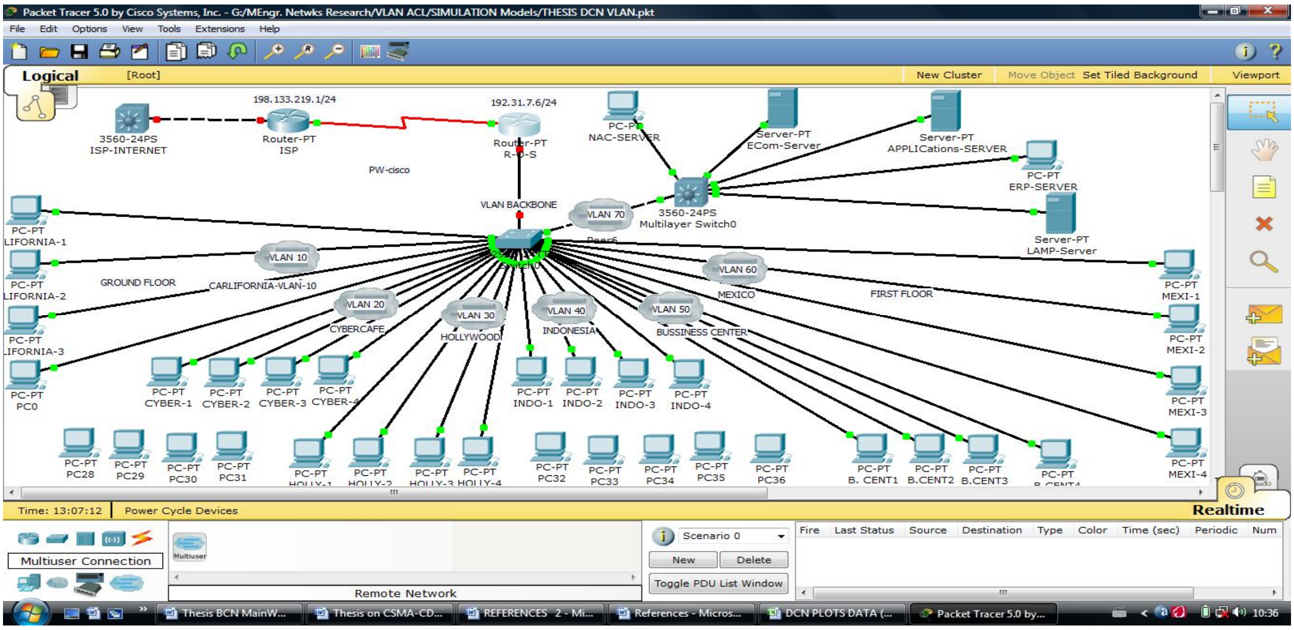

Fig. 2: Network Model of Unvirtualized DCN.
(**Source**: Okafor kennedy C.,A synthesis VLAN approach to congestion management, 2010).

2. **RELATEDWORK**

According to [5], the term virtualization broadly describes the separation of a resource or request for a service from the underlying physical delivery of that service. With virtual memory, for example, computer software gains access to more memory than is physically installed, via the background swapping of data to disk storage. Similarly, virtualization techniques can be applied to other IT infrastructure layers - including networks, storage, laptop or server hardware, operating systems and applications. This blend of virtualization technologies (virtual infrastructure) provides a layer of abstraction between computing, storage and networking hardware, and the applications running on it, (see Fig. 3,5). The deployment of virtual infrastructure is non-disruptive, since the user experiences are largely unchanged. However, virtual infrastructure gives administrators the advantage of managing pooled resources across the enterprise, allowing IT managers to be more responsive to dynamic organizational needs and to better leverage infrastructure investments [5]. Virtualization has been applied to operating systems both commercially and in research for nearly thirty years. IBM VM/370 [6, 7] first made use of virtualization to allow binary support for legacy code. VMware [8] and Connectix [9] both virtualize commodity PC hardware, allowing multiple operating systems to run on a single host. All of these examples implement a full virtualization of the underlying hardware, rather than paravirtualizing and presenting a modified interface to the guest OS. A small number of changes had to be made to the hosted operating systems to enable virtualized execution in the server architecture. At present, we are aware of two other systems which take the paravirtualization approach: IBM presently supports a paravirtualized version of Linux for their zSeries mainframes, allowing large numbers of Linux instances to run simultaneously. Denali [10], is a contemporary isolation kernel which attempts to provide a system capable of hosting vast numbers of virtualized OS instances.

In addition to Denali, we are aware of two other efforts to use low-level virtualization to build an infrastructure for distributed systems. The vMatrix [11] project is based on VMware and aims to build a platform for moving code between different machines.

As vMatrix is developed above VMware, they are more concerned with higher-level issues of distribution that those of virtualization itself. In addition, IBM provides a "Managed Hosting" service, in which virtual Linux instances may be rented on IBM mainframes. The PlanetLab [12] project has constructed a distributed infrastructure which is intended to serve as a testbed for the research and development of geographically distributed network services. The platform is targeted at researchers and attempts to divide individual physical hosts into *slivers*, providing simultaneous low-level access to users. The current deployment uses VServers [13] and SILK [14] to manage sharing within the operating system. We share the same motivation with authors in existing literature that with the operating system extensibility, IT resource management will be very efficient, hence the proposed DCN virtualization as a cost effective approach to infrastructure management.

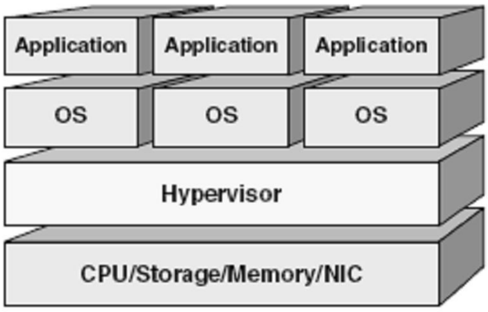

Fig 3 : Virtualized Server with Hypervisor Layer [6]

This paper focuses on server-IT resource virtualization on a DCN architecture as well as its cost budget effect to organizations.

2.1 Contextualization - Virtual Machine

A virtual machine is an integrated isolated software bundle, container that can run its own operating systems and applications as if it were a physical computer. A virtual machine behaves exactly like a physical computer and contains its own virtual (i.e. Software based) CPU, RAM hard disk and network interface card (NIC). VMware Server allows quick provisioning of virtual machines by supporting several guest operating systems, including

Windows Server 2008, Windows Vista and various Linux distributions. Figures 4 and 5 depicts DCN infrastructures in enterprise organizations with virtualization in figure 5.

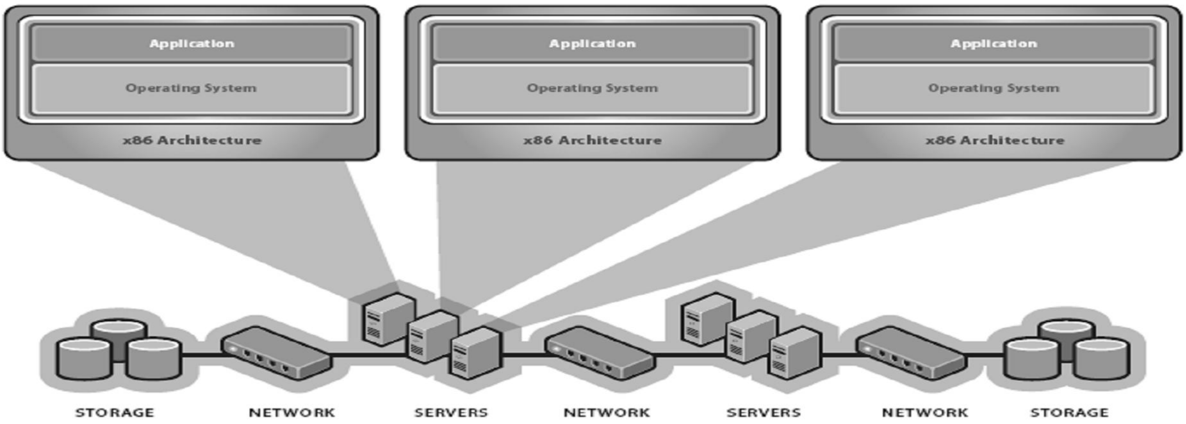

Fig. 4: A Model of DCN Traditional Infrastructure

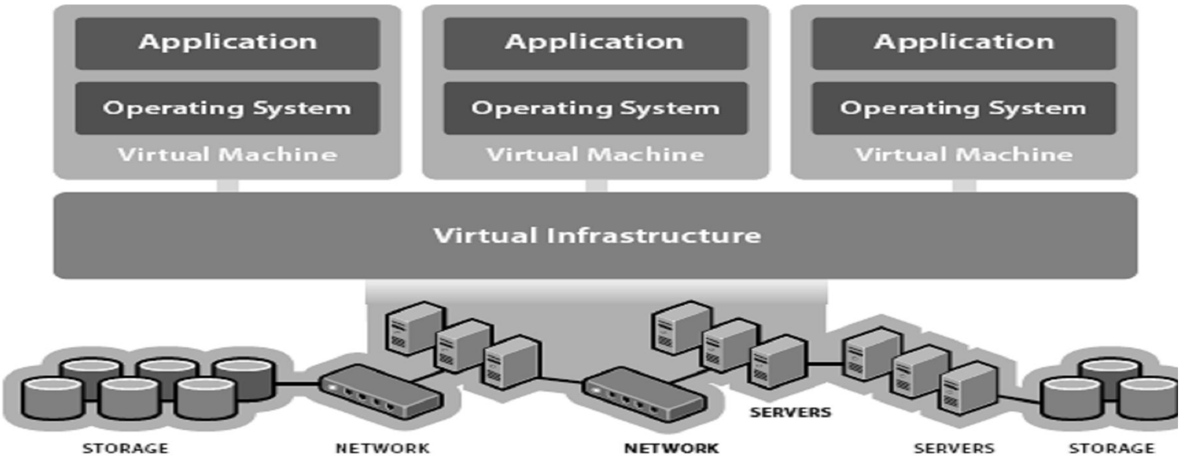

Fig. 5: A Model of VMware Virtual Infrastructure

3. Virtualization Topologies and Cabling

The physical network must adapt to the requirements and advantages created by virtualization, specifically higher utilization and higher bandwidth. To do this, forward-thinking network professionals in [15] implemented the End-of-Row ("EoR") or Top-of-Rack ("ToR") topologies in their datacenter networks. This work focuses on EOR virtualization for the DCN design owing to its benefits over TOR Topology.

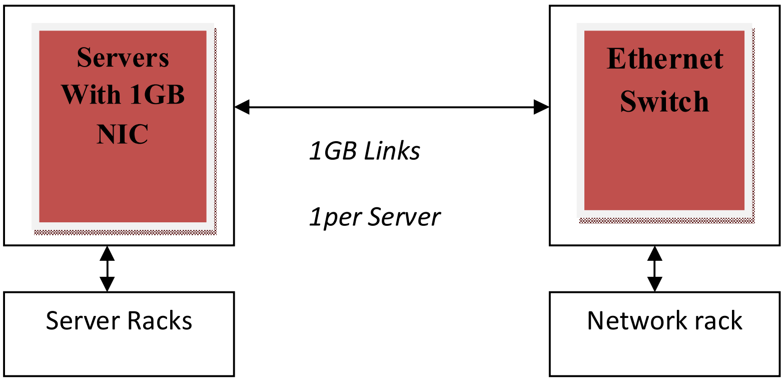

Fig 6: Data center Basic Topology

As shown in Figure 6, datacenter pre-virtualized have each resource (server, storage device, etc.) individually linked to an Ethernet switch. This topology uses structured cable connections that is difficult to modify. Since virtualization facilitates change, a network architecture that inhibits it is inherently problematic. The topology with 1 Gigabit links is incongruent with consolidated servers that need fewer and faster connections in a DCN.

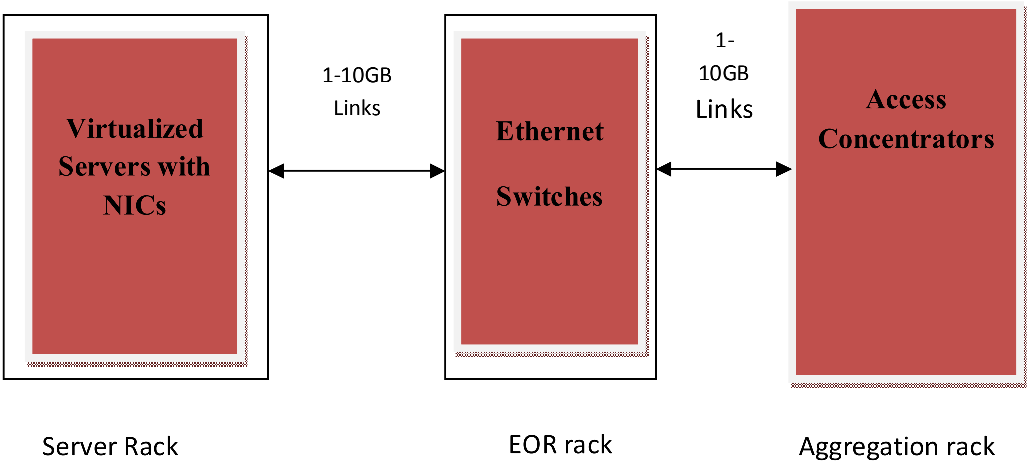

Fig 7: DCN Virtual End-of-Row Topology

The End-of-Row network topology, as shown in Figure 6, this addresses the limitations of a conventional Datacenter network topology shown in figure 1 by dedicating an Ethernet switch to each row of equipment racks [15]. The virtualized assets in each rack, in each row, are linked to a switch in the EoR rack. The switch also provides a trunk connection to a datacenter concentrator.

The EoR topology divides the switch fabric and physical connections from one tier into two, making the network more adaptable. EoR limits the length of the cables in the lower tier to the length of a row of racks. Shorter cables are generally easier to install and easier to change. EoR topology confines the impact of asset reconfiguration to a row of racks, instead of across an entire datacenter. EoR usually reuse some elements of the existing physical network, although major changes and upgrades are likely. This is not the case with ToR topology. Topology and economics are the prime determinants of connectivity in our data center network, but for a reliable cable infrastructure Table 1 shows the cable infrastructure for the DCN.

Table 1: EoR and ToR Cables

Class	Use	Bit rate	Max. Distance	EoR	ToR
Intra-rack	Device to TOR switch	1-10Gbps	5 meters	Nil	yes
Rack-to-Rack	Device to EoR switch	1-10Gbps	50meters	yes	Nil
Cross-Datacenter	ToR or EoR switch to Concentrators	10Gbps	300meters	Yes	yes

Some device-to-switch links can work at 1 Gigabit today, but adapting to higher bit rates is the best option for future investments. Table 2 shows other ways to support 10-Gigabit traffic for the DCN

Table 2: Shows the most common 10-Gigabit solutions

Media	Cable	Type	Max. distance	EoR	ToR
Copper	Twinax	SFP+Copper	10m	Yes	Nil
Copper	Twinax	10GBASE-CX4	15m	Yes	Nil
Copper	Cat 6 UTP	10GBASE-T	50m	Yes	Nil
Copper	Cat6A UTP	10GBASE-T	100m	Yes	Yes
Copper	Cat 7 UTP	10GBASE-T	100m	Nil	yes
Fiber	850/Multimode	10GBASE-SR	80-300m	Nil	yes
Fiber	1300/Multimode	10GBASE-LRM	220m	Nil	Yes
Fiber	1310/Single mode	10GBASE-LR	10km	Nil	Yes

4. ESX Server Deployment Architecture for Virtualized DCN

Figure 8 shows the sequence of deployment of an ESX server for a multi-agent data center network as implemented in this work. In this context, this paper first determine the workload profile and since it proposes to run a Web server application (Apache/Asp), database application (Oracle), ERP Application (SAP Ag), Content Management System (CMS) application and E-commerce applications, hence five (5) virtual machines (VM) were created on the ESX server platform. The memory management capabilities of VMware ESX Server provides a unique and sophisticated way to maximize the usage of physical memory within a single box (hardware). Basically, for many workloads, memory is the limiting factor, and effective memory management enables more virtual machines to share a single server, increasing return on investment (ROI) for consolidation [6]. Advances in virtualization, CPU, and memory technology make the addition of memory one of the most effective investments for maximizing the utilization of an ESX Server host.

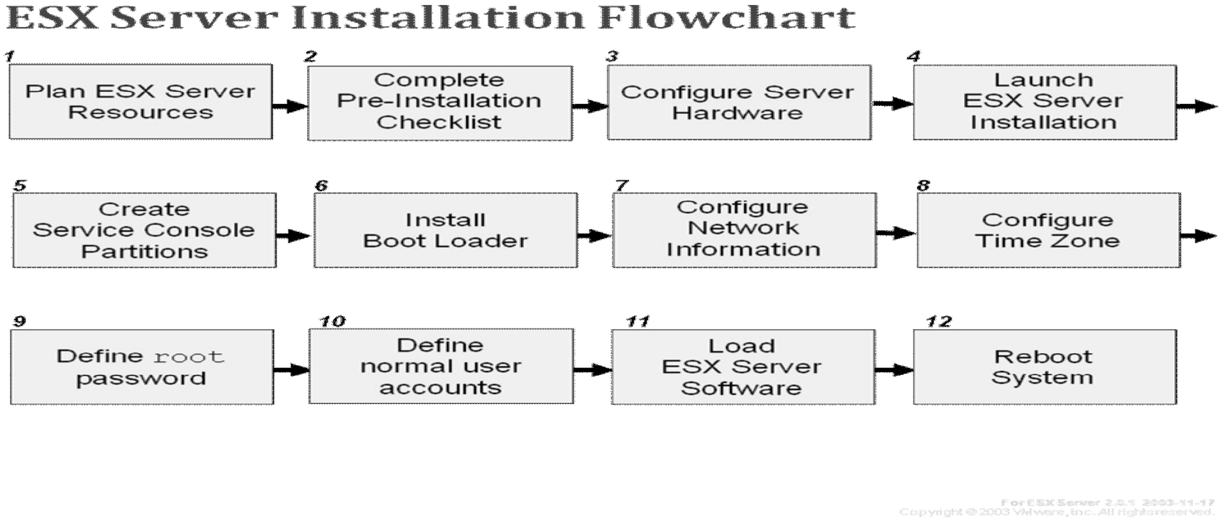

Fig. 8: An ESX Server Deployment Flowchart

After the deployment, five virtual machines were created for the proposed applications, the observed performance stimulated a comparison between a virtualized platform and non virtualized platforms as shown in Table 3

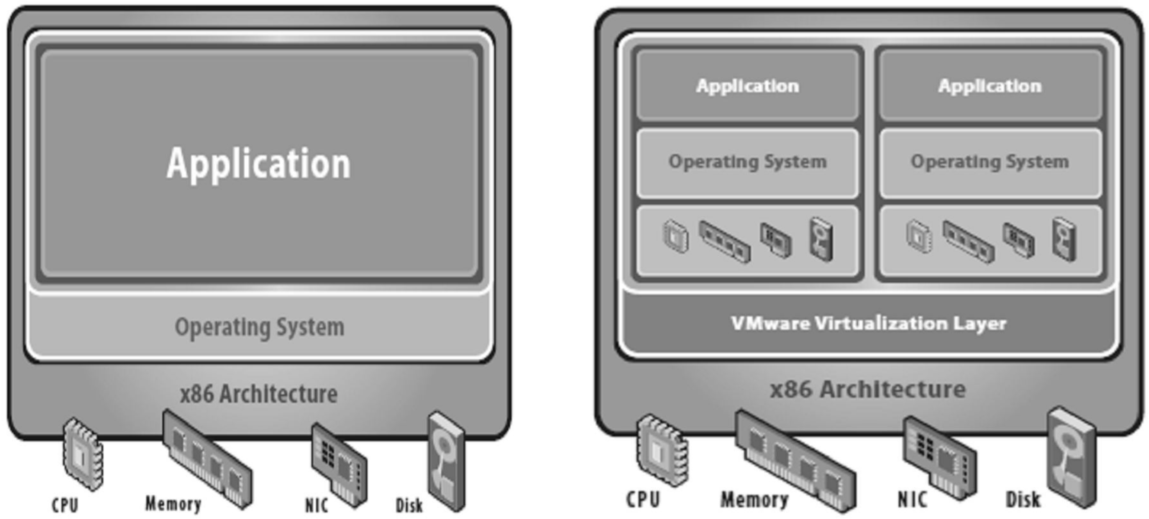

Fig. 9a) Before Virtualization: Fig. 9b) After Virtualization:

Figure 9: The Basic virtualization model.

Considering figures. 9a) and 9b, Table 1 was derived for deployment evaluation in enterprise domain.

Table 3: Comparison on Virtualization States

	Before Virtualization:	**After Virtualization**
1.	Single OS image per machine	Hardware-independence of operating system and applications.
2.	Software and hardware tightly coupled	Virtual machines can be provisioned to any system.
3.	Running multiple applications on same machine often creates conflict	Can manage OS and application as a single unit by encapsulating them into virtual machines
4.	Underutilized resources	High and Effective resource Utilization
5.	Inflexible and costly infrastructure	Flexible less expensive overall infrastructure

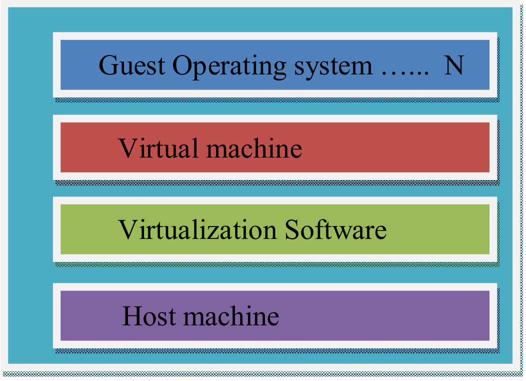

Fig 10: Basic Virtualization Model

Fig. 10 shows a basic virtualization model which was explained in figs 3, 5. In full virtualization, the virtual machine simulates enough hardware to allow an unmodified guest operating system to be run in isolation. In this case, there is an abstraction between guest operating system and the underlying hardware (completely decoupled). Full virtualization offers the best isolation and security for virtual machines, and allows simple procedures for migration and portability as the same guest operating system instance can run virtualized or on native hardware. In this work, the use of virtual switches was implemented in ESX server as a flexible way of assigning network connectivity to virtual machines for our applications. While virtual switches were created with a single physical NIC. The failure of that NIC would cause the virtual machines using it to lose their network connectivity. To prevent this downtime, ESX server allows a bonding up to 8 gigabit Ethernet adapters (up to 10 10/100 adapters) together to present to virtual machines. Basically, a virtual switch emulates a 32-port switch for the guests that are configured to utilize it. Each time a virtual machine references a virtual switch in its configuration it utilizes one port. Virtual switches also load-balance virtual machines across all physical NICs used to create the switch. If one network switch port or VM NIC were to fail, the remaining VM NICs in the bond that makes up the virtual switch would pick up the workload.

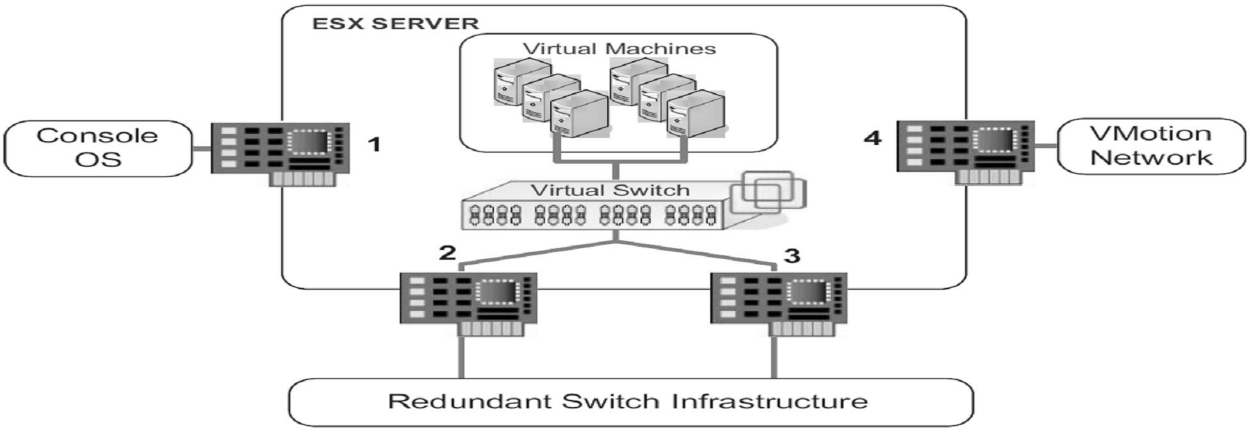

Fig. 10a) A Virtual Network (VMnet) Configuration

In fig. 10b, with VMnets in ESX server, a virtual switch methodology implements what this paper refers to as virtual networks, which provides the capability to create a private network that is visible only to other hosts configured on the same VMnet on the same physical host. In this regard, VMnets are simply virtual switches that have no outbound adapters assigned to them. Using this feature, implementation of a multi-network environment on a single host is feasible, hence we refer to it as VM multiagent support for DCN. In the fig. 10b , traffic coming in through the external virtual switch would have no way to directly communicate to the Virtual Machine 3 and would only be able to interact with Virtual Machine 1 or 2, thereby enhancing switching security in DCNs.

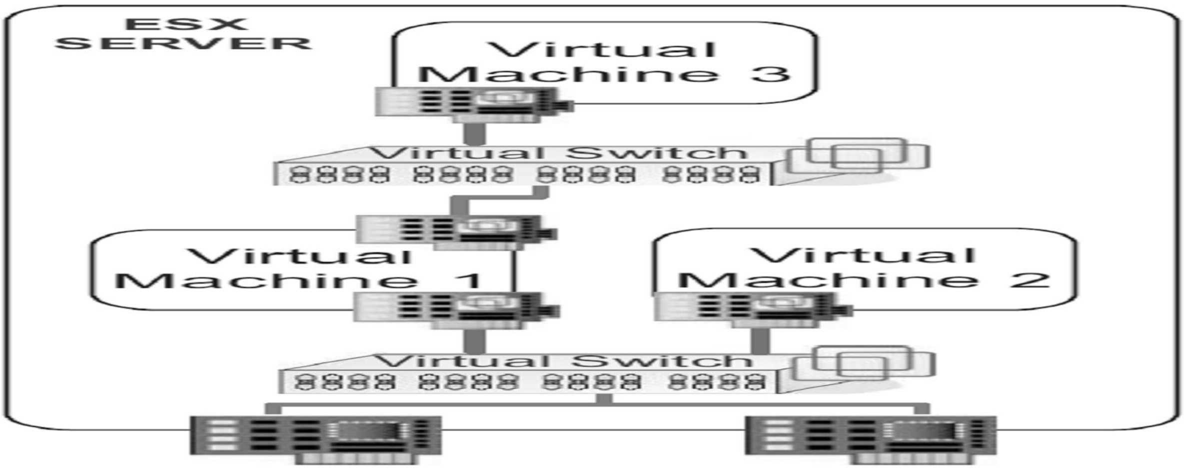

Fig. 10b: VM Switching in Virtual Network (VMnet) Setup

5. DISCUSSION

For large scale enterprises, using virtualization improves general system performance and efficiency while simultaneously minimizing the costs associated with the softwares including the number of licenses required. The results of this testing have direct application in the area of database virtualization, a field where little progress has been made in the past. Also, the very large 1TB memory capacity of the ESX server series offers great potential for use in memory-resident database platforms. Owing to the resource provisioning and demands on a data center network, this paper presents virtualization as a cost effective approach to resource management. With virtualization, a guest system itself can run unmodified on top of the hypervisor engine. Because using unmodified drivers in the guest system uses up some system resources, it was observed that VMkernel controls and manages most of the physical resources on the hardware, including:

- Memory.
- Physical processors
- Storage controllers.
- Networking
- Keyboard, video, and mouse.

Basically, the VMkernel includes schedulers for CPU, memory, and disk access, and has full-fledged storage and network stacks. The virtual machine monitor is the component actually responsible for virtualizing the CPUs as shown in figures above. When a virtual machine is powered on, control transfers to the virtual machine monitor, which begins executing instructions from the virtual machine. The transfer of control to the virtual machine monitor involves setting the system state so that the virtual machine monitor runs directly on the hardware. The VMkernel manages all machine memory except for the memory that is allocated to the service console. Individual virtual machine process threads are handled directly by the VMkernel based on our envisaged workload profile, thereby eliminating the need to allocate service console memory for each virtual machine. It was necessary to allocate more memory to the service console, since service agents run there for monitoring or backup. Based on the data center setup, the VMkernel dedicates part of its managed machine memory for its own use, while the rest is available for use by virtual machines that run the applications. As shown figure 11, a VMware Server partitions a physical server into multiple virtual machines.

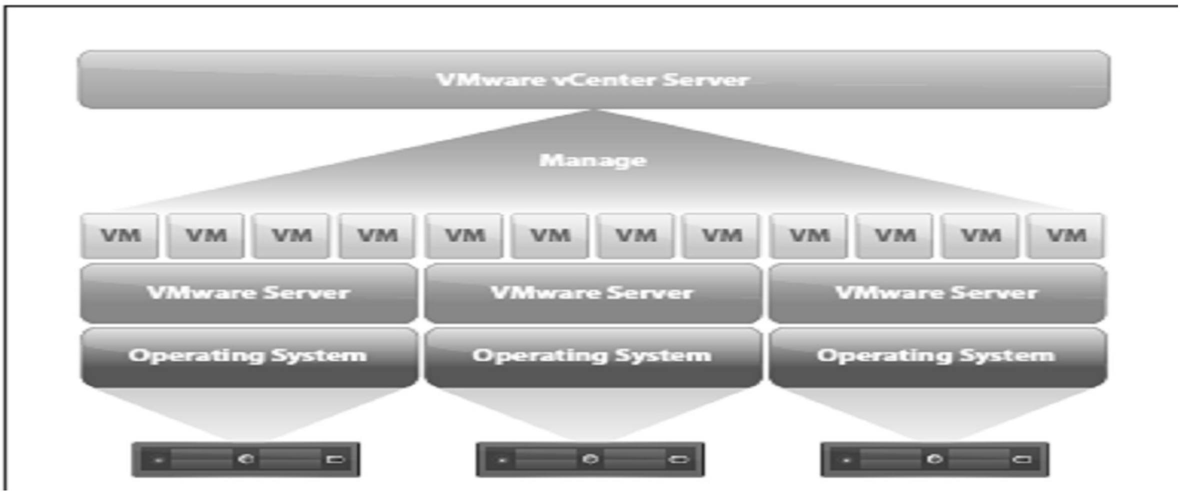

Fig.11: Multi-agent service provisioning with VM Server-ESX server.

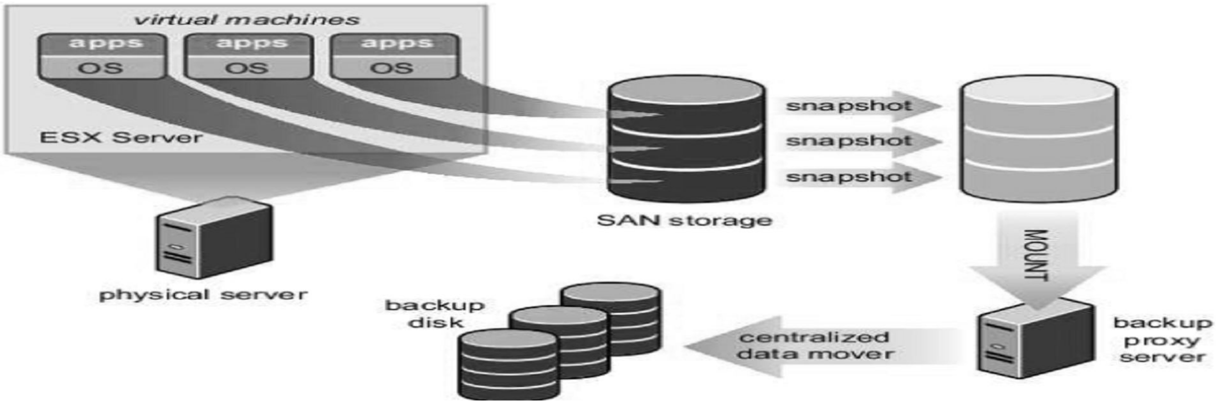

Fig.12: Centralized DCN Backup infrastructure

Fig. 12 shows a complex storage and backup system, able to provide a centralized storage for the virtual machines, and a full backup of the virtual machines snapshots in a DCN.

Experience with production workloads in DCN has shown that in modern DCNs, x86-based hardware memory is often the most used resource for OS management. Table 4: shows the average utilization rates of the four main computing resources for a typical DCN system.

Table 4: Average Utilization Rates.

Resource	Utilization
CPU	6%
MEMORY	40%
NETWORK I/O	<5%
DISK I/O	<5%

Utilization highly depends on the particular workload being run. In addition, although the average utilization might be very small, utilization rates experience peaks that occur as a result of either normal usage cycles during the business day or unexpected demands. In all cases, memory usage remains large for the host machine.

The utilization of a workload directly translates into the utilization of a virtual machine running that workload. Since the utilization of the four resources on one virtual machine determines how much of those resources is available for other virtual machines on the same host, a doubling of the amount of physical memory on an ESX Server host enables twice as many virtual machines to be put on that host.

6. Conclusion and Future work.

This work presents the importance of virtualization in data center networks with emphasis on ESX server and its underlying architecture for EoR DCN topology . It has portrayed the implication of sticking to the traditional data center design methodology which imposes high cost budget on organizations. A good foundation for the possibility of migrating to cloud computing technology in Nigeria has also been laid. Besides, the combination of ESX Server with the newer generation servers based on two-way, multicore CPUs delivers compelling platforms for enterprise virtualization and consolidation efforts. Multicore CPUs, high-bandwidth, higher memory capacities, and other architectural improvements increase the return on investment for virtualization efforts and reduce life-cycle costs for IT managers by enabling greater use of existing servers. Future work will investigate the effect of dynamically adding and removing hosts to the VMkernel via a dynamic VM management mechanism in the ESX server machine.

REFERENCES

[1] T. B. Team. An overview of the blue gene/l supercomputer. In ACM Supercomputing Conference, 2002.

[2] S. M. Kelly and R. Brightwell. Software architecture of the light weight kernel, catamount. In 2005 Cray Users' Group Annual Technical Conference, Albuquerque, New Mexico, May 2005.

[3] S. M. Kelly, R. Brightwell, and J. VanDyke. Catamount software architecture with dual core extensions. *In 2006 Cray Users' Group Annual Technical Conference, Lugano, Switzerland, May 2006.*

[4] Katz, Randy. IEEE spectrum: Tech titans building boom. http://www.spectrum.ieee.org/green-tech/buildings/tech-titans-building-boom.

[5] VMware- http://www.vmware.com.

[6] P. H. Gum. System/370 extended architecture: facilities for virtual machines. *IBM Journal of Research and Development*, 27(6):530–544, Nov. 1983.

[7] L. Seawright and R. MacKinnon. VM/370 – a study of multiplicity and usefulness. *IBM Systems Journal*, pages 4–17, 1979.

[8] S. Devine, E. Bugnion, and M. Rosenblum. Virtualization system including a virtual machine monitor for a computer with a segmented architecture. *US Patent*, 6397242, Oct. 1998.

[9] Connectix. Product Overview: Connectix Virtual Server, 2003. http://www.connectix.com/products/vs.html

[10] A. Whitaker, M. Shaw, and S. D. Gribble. Scale and performance in the Denali isolation kernel. In *Proceedings of the 5th Symposium on Operating Systems Design and Implementation (OSDI 2002)*, ACM Operating Systems Review, Winter 2002 Special Issue, pages 195–210, Boston, MA, USA, Dec. 2002.

[11] A. Awadallah and M. Rosenblum. The vMatrix: A network of virtual machine monitors for dynamic content distribution. In *Proceedings of the 7th International Workshop on Web Content Caching and Distribution (WCW 2002)*, Aug. 2002.

[12] L. Peterson, D. Culler, T. Anderson, and T. Roscoe. A blueprint for introducing disruptive technology into the internet. In *Proceedings of the 1st Workshop on Hot Topics in Networks (HotNets-I)*, Princeton, NJ, USA, Oct. 2002.

[13] J. Gelinas. Virtual Private Servers and Security Contexts, 2003. http://www.solucorp.qc.ca/miscprj/s_context.hc.

[14] A. Bavier, T. Voigt, M. Wawrzoniak, L. Peterson, and P. Gunningberg. SILK: Scout paths in the Linux kernel. Technical Report 2002-009, Uppsala University, Department of Information Technology, Feb. 2002.

[15] White paper: Connectivity in the virtualized Datacenter: How to Ensure Next generation Services,2010.

[16] http://www.vmware.com/vmi.

[17] http://www.vmware.com/stand

[18] R. J. Adair, R. U. Bayles, L. W. Comeau, R. J. Creasy, "A Virtual Machine System for the 360/40", IBM Corporation, Cambridge Scientific Center Report No. 320-2007, 1966

[19] R. P. Goldberg, "Survey of Virtual Machines Research", *IEEE Computer*, 1974, pp. 34-45

[20] B. Bacci, "Virtualization - TOSSLab Tuscany open source software laboratory", Pisa 2009

[21] Parallels - http://www.parallels.com

[22] VirtualBox - http://www.virtualbox.org

[23] AMD-V - http://www.amd.com/virtualization

[24] VT-x - http://http://www.intel.com/technology/itj/2006/v10i3/1-hardware/ 6-vtx-vt-i-solutions.htm

[25] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, A. Warfield, "Live migration of virtual machines", Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation (NSDI'05), Boston 2005, Vol. 2

[26] C. P. Sapuntzakis, R. Chandra, B. Pfaff, J. Chow, M. S. Lam, M. Rosenblum, "Optimizing the migration of virtual computers", Proceedings of the 5th symposium on Operating systems design and implementation (OSDI'02), Boston 2002

[27] A. Awadallah and M. Rosenblum. The vMatrix: A network of virtual machine monitors for dynamic content distribution. In Proceedings of the 7th International Workshop on Web Content Caching and Distribution (WCW 2002), Aug. 2002.

[28] A. Bakre and B. R. Badrinath. I-TCP: indirect TCP for mobile hosts. *In Proceedings of the 15th International Conference on Distributed Computing Systems (ICDCS 1995), pages 136–143, June 1995.*

[29] G. Banga, P. Druschel, and J. C. Mogul. Resource containers: A new facility for resource management in server systems. *In Proceedings of the 3rd Symposium on Operating Systems Design and Implementation (OSDI 1999), pages 45–58, Feb. 1999.*