Global Journal of Computer Science and Technology

Energy Efficiency in Datacenters through Virtualization: A Case Study

K. Thirupathi Rao[1]P.

Sai Kiran[2] L.S.S.Reddy[3]

GJCST Computing Classification
H.3.4, E.5, H.2.8 ,D.4.3, F.2.2

***Abstract*-Data centers have turned into a big concern for enterprises. Data centers have exhibited steadyenergy demand growth, and electric utilities have received sizeable requests for electrical power for new facilities. Electricity usage costs have become an increasing fraction of the total cost of ownership for data centers. There are many challenges and opportunities for researchers to explore the issues to quantify the electricity savings and provide methods for improving the energy efficiency. This paper systematically explores the opportunities and adoptable methods for achieving the energy efficiency in datacenters. We have presented a Study on Data Center Virtualization and scaled it to realize significant savings for an enterprise.**

I. Introduction

Energy efficient data center is gathering momentum as organizations have started realizing its importance in energy conservation and sustainable development. It is applied to new technologies that can help in cutting down data center energy costs and saving energy, which is synonymous to saving money. It has a big role to play in reducing the power consumption in the data center. [7] explores key capabilities to look for when investing in data center power management and discussed the associated business value that can be derived from energy management solutions. Data center is a special facility that physically houses various equipment, such as computers, servers (e.g., web servers, application servers, database servers), switches routers, data storage devices, load balancers, wire cages or closets, vaults, racks, and related equipment. It can store, manage, process, and exchange digital data and information and provide application or management services for various data processing applications

Virtualization has the promise to maintain or increase computing power and data center performance while controlling costs and extending the value of existing data center facilities. Virtualization enables partitioning, whereby a single physical server runs multiple virtual machines, each with its own independent and secure application and operating system. The same work gets done, but there's far less idle capacity. I is possible to adapt workload Variations by shifting resources and priority allocations among virtual machines. A software failure in a virtual machine does not affect other virtual machines thus providing high reliability and availability. It is possible to achieve cost reductions by consolidation smaller servers into more powerful servers. Virtualization overview and approaches are provided in [8].

About[1-]*Associate Professor, Department of Computer Science and Engineering*
About [2]*Director, Lakireddy Bali Reddy College of Engineering*
Lakireddy Bali Reddy College of Engineering, Mylavaram, Krishna District, Andhra Pradesh, India
[1]*email2ktr@gmail.com*, [2]*psaikiran@hotmail.com*, [3]*director@lbrce.ac.in*

II. Background and Related Work

Datacenter facilities are become common and essential to the functioning of business, communications, academic, and governmental systems since every office shift from paper-based to digital information management. Data centers are found in nearly every sector of the economy: financial services, media, high-tech, universities, government institutions, and many others use and operate data centers to aid business processes, information management, and communications functions. Day to day there is an increasing demand for data processing and storage. This demand is driven by several factors like the increased use of electronic transactions in financial services, such as on-line banking and electronic trading, the growing use of internet communication and entertainment, the shift to electronic medical records for healthcare, the growth in global commerce and services, and the adoption of satellite navigation and electronic shipment tracking in transportation. There has been mounting interest in opportunities for energy efficiency government sector. The energy used by the servers and data centers is significant. It is estimated that this sector consumed about 61billion kilowatt-hours (kWh) in 2006 for a total electricity cost of about $4.5 billion. The energy use of the nation's servers and data centers in 2010 is estimated to be more than double the electricity that was consumed for this purpose in 2000. The United States (U.S.) Environmental Protection Agency (EPA) developed a report which assesses current trends in energy use and energy costs of data center sand servers in the U.S. and outlines existing and emerging opportunities for improved energy efficiency. It provides particular information on the costs of data centers and servers to the federal government and opportunities for reducing those costs through improved efficiency. It also makes recommendations for pursuing these energy-efficiency opportunities broadly across the country through the use of information and incentive-based programs [1][5]. Electricity usage costs have become an increasing fraction of the total cost of ownership for data centers. It is possible to dramatically reduce the electrical consumption of typical data centers through appropriate design of the network-critical physical infrastructure and through the design of the IT architecture. [2] explains how to quantify the electricity savings and provides examples of methods that can greatly reduce electrical power consumption Datacenter inefficiency is a widespread and growing concern. Soaring costs and ever increasing environmental footprints are impacting the corporate investment landscape, threatening profitability and inviting board and regulatory scrutiny. [3] Outlines ways for organizations to address twin challenges of rising datacenter spend and green gas emissions. Organizations are now more concerned about environment and energy saving models.[4] addresses how IT professionals can do better design of power and cooling for efficiencies in their data centers through the strategic application technologies and best practices. The response of many IT departments are positive because they realized that energy efficient datacenter is implemented in the short term often with very little expense and minimal staff involvement that can have a positive impact and help protect the environment at the same time Worldwide, the total power consumption of datacenters is expected to double between 2010 and 2015. Doubling server consumption would require additional capacity equal to more than 10 additional1,000-megawatt power plants. So there are plenty of reasons to cut datacenter power consumption. Servers and storage systems are not the only things that use energy in the datacenter. Cooling equipment often uses as much power as the systems themselves. Add to that the energy used to light the datacenter, the power distribution loss, and other factors, and you'll see that the majority of power coming into the datacenter is used for something other than IT equipment. [6] will give a closer look at how power is actually consumed within the datacenter. Each type of power usage is an opportunity to reduce total power consumption and CO2 emissions, so it's important to consider the entire range.

III. Challenges of Datacenters and IT Industry

a) Energy Challenges of Datacenters: Achieving energy efficiency in data centers is a challenging goal and there are large opportunities for savings. These savings will be easy to achieve by addressing the challenges of today's datacenters.

b) Energy costs: As electricity becomes increasingly expensive due to accelerating demand and carbon emissions regulations, data center operators have less control over their fixed data center energy costs. It is challenging to keep operators remain competitive while reducing their cost per user without cutting vital headcount.

c) Managing Operations: Maintaining control of network operations has historically required running all data center servers at capacity to ensure reliable, adequate service levels. Methods have to be explored to dynamically manage capacity to save costs without compromising availability. The mainframe is one part of the wider problem facing data centre managers, as they are bulky and demand plenty of cooling
and air-conditioning.

d) CO2 emissions: Eco-managing of datacenter is mostly about using energy efficiently, because when you cut power consumption you simultaneously cut power costs and CO2 emissions. Every system uses energy, and the production of usable energy is CO2-intensive. Multiply the individual online interactions and transactions by billions cause huge costs in terms of datacenter infrastructure, power,
and environmental impact. Carbon emissions are receiving increased scrutiny, and some sort of Capand Trade or carbon tax appears likely. These initiatives will punish inefficient users of energy economically and tarnish their corporate images. Data center operators minimize these emissions.

e) Service uptime: As operators attempt to improve service levels across multiple data centers, they are
also looking for ways to improve energy efficiency, including during system failures. Researchers
investigate tools to in achieving this goal.

IV. Challenges of IT Industry

a) Consolidation and Dynamic Provisioning of Resources

Capital expenditure increases enormously

to meet the greater demand of computing power to meet the business challenges. Application
compatibility issues resulting in significant server sprawl. Provisioning new servers is a lengthy, labor
intensive process and increases operating expenditures as well as power and cooling costs. At the same
time, servers are often underutilized. Typically, server workloads consume only 5 percent of their total
physical capacity, wasting hardware, space, and electricity. It is a challenge for IT to keep pace with the much faster rate of business growth and change.

b) Business Continuity/Disaster Recovery: Businesses to maintain continuity, disaster recovery needs to
be at the core of any IT strategy, but implementing a reliable, rapid recovery strategy can be time-consuming and expensive.

***c) High Availability*:** Since downtime cost organizations to lose money, high availability has become vital ingredient of IT strategies in a world where businesses need to operate 24 hours a day, seven days week. The disruption of IT services can be fatal to a business.

***d) Managing Centralized and Optimized Desktop*:** Varying end-user needs for applications and services require organizations to develop, deploy, manage, and support dozens of desktop images. Managing hundreds or even thousands of desktops, applications, and servers is incredibly challenging, complicated and requires vast resources.

V. Opportunities for Improving Energy Efficiency in Datacenter

There is significant potential for energy-efficiency improvements in data centers. Many technologies are either commercially available or will soon be available that could improve the energy efficiency of microprocessors, servers, storage devices, network equipment, and infrastructure systems. Still thereare plenty of unexplored, reasonable opportunities to improve energy efficiency. Selection of efficient IT equipment and reducing mechanical infrastructure increases the energy efficiency. Improvements are possible and necessary at the level of the whole facility i.e. system level and at the level of individual components. It is not possible to optimize data center components without considering the system as a whole, still it is true that efficient components are important for achieving an efficient facility for instance, efficient servers generate less waste heat which reduces the burden on the cooling system. For achieving greatest efficiency here we will provide a comprehensive approach that explores theopportunities for improvement in many areas of the IT and mechanical infrastructure systems.

***a)* Virtualization and Consolidation:** Virtualization is the process of presenting a logical grouping or subset of computing resources so that they can be accessed in ways that give benefits over the original configuration. This new virtual view of the resources is not restricted by the implementation, geographic location or the physical configuration of underlying resources.

System utilization rates of 10 to 15 percent are still commonplace in enterprise datacenters, and servers use most of their peak power whether they are 15 percent loaded or 80 percent loaded. Companies can reap instant savings by consolidating underutilized physical data servers onto virtual machines that act like physical computers, but don't require the space, management time or energy of individual servers.

***b)* State-of-the-art Technologies:** Adopting best practices and using available state of the art technologies for aggressively consolidating servers, storage will lead to achieve the maximum energyefficiency savings

***i) Specify efficient server equipment*:** Efficient server equipment can reduce the energy use hence savings are possible at total infrastructure requirements, cooling loads, electricity consumption, smallerUPS.

ii) Select systems which are more inherently efficient: If you can save a few watts with the systemsyou deploy, that will often save more watts that would have been used for cooling or lost in delivery. Advances such as chip multi-threading (CMT), slower disk drives, and automated power-downtechnology are options for driving down the power consumption of datacenter systems.

iii) Use rating systems to calculate power consumption: Rating systems helps to compute the expected power consumption of specific server configurations. These systems provide some head-to-head comparisons between products from multiple vendors; and they can also help you size equipment since over-sizing those systems can lead to very large inefficiencies.

***iv) Code Tuning for efficiency*:** Less code — executed more efficiently — means fewer CPU cycles expended in processing workloads. For large compute farms a 10 percent improvement in efficiency can mean 10 percent less machines and 10 percent less energy. Software engineers are encouraged to minimize code and maximize execution efficiency wherever possible.

***v) Refresh your technology regularly*:** Compute technology continues to follow an exponential improvement curve. Combine that with new features and just upgrading older equipment can often result in enough energy savings. Other innovative ways of making IT more environmentally friendly include ultra-thin, high-density server designs, hydrogen fuel-cells as alternative green power sources, and nanofluid-cooling systems for the IT datacenter.

***c)* Power and cooling systems:**The cost of data centers is going up as a result of the increased power capacity required. Server powerdensity will continue to increase and data centers will have to scale their infrastructure to support thisincrease. There can be huge savings in this energy cost by focusing on optimizing the power and cooling in the data center.

i) Find more efficient power distribution: Find ways to obtain more efficient high-voltage powerdistribution, and compares high-voltage AC and DC and the potential energy savings of each.

Minimize the number of power conversions from the utility feed to the equipment

ii) Cooling: Cooling is suddenly an area of major innovation exploring new options such as integrating cooling technology directly into server racks, even simply using variable-speed fans rather than singlespeed fans, and making better use of cool external air

iii) Equalize heat and cooling balance*:* Data centers waste enormous amounts of energy byovercooling the majority of their data cabinets. This is because they make macro-cooling decisionsbased on the heat generated by their hottest cabinets. By matching cooling resources (from floor vents or liquid cooling units) to the actual needs of each individual cabinet, data centers can realizesignificant savings on energy use. In addition, further savings can be achieved by balancing heat loads on an intra-cabinet basis. By grouping servers in a minimum of 2 and preferably 3 zones within a cabinet and moving towards equalized heat loads between the zones.

iv) Use Equipment Racks with Integral Coil: Transferring IT equipment waste heat to a cooling waterloop directly at the rack allows for the complete elimination of heat recirculation. The heat is capturedprior to mixing with the room air at a higher air temperature. This allows a correspondingly higher cooling water temperature to be used in thecool, allowing significant plant efficiency opportunities.

d) Manage to the metrics: As data centers add, move, and change servers, many on a daily basis, theyneed to continue to monitor and manage heat generation and cooling requirements. Solutions has to beprovided for managing energy usage of data centers over the long term to achieve maximum energyefficiency. Investigate how IT equipment energy consumption varies with computation loads anddevelop quantitative metrics. Refine metrics and measurement protocols for benchmarking server

VI. DATA CENTER OPTIMIZATION THROUGH VIRTUALIZATION

Virtualization is basic technological innovation that allows experts and skilled IT managers to deploy creative solutions to nowadays business challenges such as cost effective utilization of IT infrastructure, business agility etc. Broadly the term 'virtualization' describes the separation of resource or request for a service from the underlying physical delivery of that service. Virtualization allows youto run the several application environments on the same machine in such a way that these environmentsare completely isolated from each otherVirtualized System can be divided roughly into three components, the Hardware, the Virtual Machine Monitor (VMM), and the Virtual Machines as shown below in Fig.1.

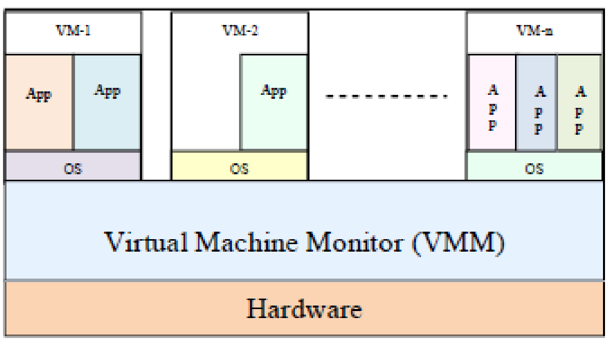

Fig.1. Abstract view of components in Virtualized System

a) The Hardware**:** provides basic computing resources CPU, memory, I/O devices etc.

b) Virtual Machine Monitor (VMM): It is a hardware abstraction layer which provides an interfacebetween hardware and Virtual Machines. It controls access to the resources shared by all virtualmachines. The VMM schedules the virtual machines, in a manner similar to how an operating systemschedules processes, and allocates processor cycles to them. The VMM layer can map and

remapvirtual machines to available hardware resources at will. Load balancing among a collection ofmachines thus becomes trivial, and there is a robust model for dealing with hardware failures or forscaling systems. When a computer fails and must go offline or when a new machine comes online, theVMM layer can simply remap virtual machines accordingly. Virtual machines are also easy to replicate, which lets administrators bring new services online as needed.The VMM virtualizes all resources and allocates them to the various virtual machines that run on top of the VMM.

c) Virtual Machines (VM): A representation of real machine using software that provides an operating

environment which can run or host a guest operating system and can support for many applicationenvironments.

VII. CASE STUDY

A typical data center can have 'n' different application environments (App) and 'm' servers. Each App may execute several classes of transactions hence require several servers for processing. Servers can bedynamically provisioned among Apps with the goal of optimizing server utilization for the data center. Consider a typical Medium Business enterprise, running 500 applications, one application per server,each server operating at about 10 to 15 percent CPU utilization. In a typical scenario, about 450 ofthese applications would be candidates for virtualization, at an average rate of five applications perserver. Server consolidation would lead to the configuration down to 90 physical servers in avirtualized environment and 50 conventional servers each running a single application, for a total of140 hardware boxes. The savings are significant, as shown in the table 1.

Savings Through Virtualization	Before Virtualization	After Virtualization
Data center size	500 servers	140 servers
Server power draws	500 @ 200W	50 @ 200W, 90 @ 270W
Total power required	100 kW	34.30 kW
Cooling	25.6 tons	12.7 tons
UPS/electrical loss	15.2 kW	8.20 kW
Electrical cost	Rs 4 per kWhr	Rs 4 per kWhr
Yearly electrical costs (IT only)	Rs 3504000	Rs 1201872
Additional costs (UPS/electrical)	Rs 532608	Rs 287328
Additional costs (cooling)	Rs 2547548.16	Rs 1263822.72
Total yearly electrical spend	Rs 6584156.16	Rs 2753022.72
Total savings per year		**Rs 3831133.44**
Power improvement		**65 percent**
Reduction in cooling		**12.9 tons**
Reduction in UPS/electrical loss		**7.0 kW**

Table1: Savings through Virtualization

The experiment is conducted in a data center with 10 servers and scaled up to realize the significant benefits of server virtualization

VIII. Conclusion and Future Scope

It has been observed in our study on Data Center, Virtualization can lead to savings for an enterprise upto 65%. Energy-efficiency strategies could be implemented in ways that do not compromise data centeravailability, performance or network security. It requires a vision of the long term standard architecture that can be applied strategically to enable, accelerate and save costs for multiple projects and business initiatives. Overall data center power consumption will be lower, but each server will draw more power which may lead to cause of costly downtime in a data center. There will be fewer servers, but each one will be more critical than ever. Applications can be dynamically moved around as needed, but thesupport infrastructure cannot do the same. Data center footprint will be smaller, but overall data centerfficiency might still be suboptimal. A virtualized datacenter that is not well managed can be lessreliable and perhaps even more expensive than its non virtualized counterpart. Further research can becarried out to address the above challenges.

IX. References

1) Report to Congress on Server and Data Center Energy Efficiency", Public Law 109-431,U.S.Environmental Protection Agency, ENERGY STAR Program, Aug 2, 2007

2) White Paper #114:"Implementing Energy Efficient Data Centers", Neil Rasmussen, AmericanPower Conversion. APC-2006.
3) Revolutionizing Data Center Energy Efficiency", James M Kaplan,William Forrest, Noah Kindler, McKinsey & company, july 2008.
4) White Paper: "Energy Efficient Data Center Solutions and Best Practices" Cisco July 2007.

5) Energy Efficiency in Data Centers: Recommendations for Government-Industry Coordination U.S EPA's National Data Center Energy Efficiency Strategy Workshop, Jul 8 to Oct 16, 2008.
6) Website: http://www.thegreengrid.org

7) IBM Energy management White paper on Realizing the value of the green data center byintegrating facilities and IT' May 2008.

8) Global knowledge, Expert reference series of White Papers: "Virtualization Overview", @ 2006vmware.