Implementing Green IT Through Virtualisation: A Case Study

Tom Townsend, Masoud Mohammadian
School of Information Systems and Accounting, University of Canberra
Canberra, Australia
tom.townsend@canberra.edu.au

***Abstract*—Demand for Information Technology industry is increasing worldwide. The expansion of digital communication has seen a large growth in the use of IT and increased capacity for data collection, storage and processing for internal and external use in business and industry. IT demands at the university sector have been increasing rapidly with the introduction of new teaching and learning methods requiring anywhere anytime teaching and learning. Data centres have emerged as an approach that enables accesses to shared computing resources. However, the increasing trends of resource usage and cost of maintaining a large number of servers and electrical consumption within data centres are of a growing concern. There are several challenges and opportunities within a data centre realm to lower the resource usage and cost of data centres to improve the cost and energy efficiency. At the beginning of 2008, the University of Canberra (UC) completed a major re-organisation involving the centralisation of all IT departments. This case study looks at how virtualisation helped rationalise that dispersed environment and how virtualisation has continued to enable dramatic affects at UC. Topics covered include; how virtualisation has affected the physical Data Centre, how the management of the IT systems has changed and what benefits and challenges were realised.**

***Index Terms*—Green IT; Virtualization.**

I. Introduction

Demand for Information Technology industry is increasing exponentially worldwide. With the increased expansion of digital communication, there has been a large growth in the use of information technology in organisations. This, in turn, has increased capacity for data collection, data storage and processing by business and industries for internal and external purposes. University sector is not an exception and university sector has a large demand for use of information technology in innovative ways for new ways in teaching and learning methods. This requires large and powerful information technology resources for anywhere anytime teaching and learning. Other examples include online financial services, mobile facilities; social websites and internet access markets. Data centres provide a clean, secure and stable environment for storage of data and servers for applications and processing of data. Data centres are maintained physically and virtually at all hours of the day and they house a large number of servers (i.e from a dozen to a few hundred servers).

In 2010 over 11 million servers were installed in the U.S. alone [1]. These servers are maintained for peak performance and are kept in physical environments that are very clean and maintained within specific temperature, humidity and air quality standards. The National Science Foundation (NSF) and NASA, through “The Green Data Centre Project” provide details of how a data centre can achieve energy efficiency while maintaining reliability standards. The term Green IT is now well defined and recognized [1, 2]. Major IT manufacturers have accepted the challenge of considering their products impact on sustainability.

The number of academics and academic conferences and publications on Green IT are increasing each year [1, 2, 3]. The University of Canberra (UC) has gone from using a hand full of virtual machines (VMs) to over two hundred and fifty VMs in the space of three years using Microsoft Hyper-V. In that time UC had reduced power consumption by almost 10% and reduced their data centre physical space requirements by six racks. This case study will look at how the environment has changed, what new capabilities were gained and the design decisions that have been made as a result of virtualisation. Some of the benefits gained through this process include;

- Increased flexibility through quick deployments of new VMs.
- Simplified physical server environment.
- Improved capacity management.
- Self-Service Portal for non-administrator staff.
- Improved uniformity through the use of templates.

This paper presents the University of Canberra (UC) case study which considers how server virtualisation helped rationalise that dispersed computing environment and how virtualisation has continued to enable dramatic affects at UC. Topics covered include; how virtualisation has affected the physical Data Centre, how the management of the IT systems has changed and what benefits and challenges were realised.

II. Related Work

Virtualization is the process of presenting a set of logical computing resources that will be accessed and shared regardless of geographic location or physical configuration [6-10]. Virtualization is under constant exposure to the media and large organisations as a solution for Green IT.

Virtualization was introduced in the early 1960s by the IBM Corporation as a technology for simultaneous timesharing of mainframe computers [5, 6]. IBM idea on virtualization was then developed to incorporate a hardware abstraction layer (i.e. Virtual Machine Monitor (VMM)) to provide interaction between the hardware and software layers [6].

With the adaptation of VMware, the concept of virtualization was transferred from mainframes to industry standard. As a result, servers could then have the capabilities of being partitioned into a large number of virtual machines. Virtualization allowed the concurrent processing of different Operating Systems and software applications in an independent fashion. Virtualizations brought a large number of advantages and disadvantages for Green IT, business and industry. The ability to have a large number of VMs on a server could reduce hardware costs and IT department overhead. However, Kappel et al, [3] argue that this potentially increases the failure risk of IT in organizations as these VMs are solely depending on a single physical server to function correctly.

There are two types of VMMs. Type I Hypervisor (OS Level Virtualization) is considered as one of the most common methods for VMs as a physical server. This architecture is also known as OS sharing. Hypervisor Virtualization is a popular method for dedicated servers with a primary purpose of running virtual servers. Type II Hypervisor Virtualization does not rely on a Host Operating System as its Virtualization Layer. It directly interacts with hardware resources. From the research and review, it could be concluded that server virtualization technologies can contribute to the Site Infrastructure and IT optimization.

III. Case Study: Green IT Through Virtualisation At University Of Canberra

Prior to 2008, the University of Canberra had a fairly typical server environment for the time; providing the vast majority of services from standalone physical servers, with many services sharing physical hardware. The University completed an IT centralisation exercise in 2009 which resulted in a server environment of two hundred and fourthly eight physical servers spreading across four separate locations; one primary data centre and three smaller satellite ‘server rooms’. The environment had a wide range of server brands, types and operating systems. With server hardware ranging from HP, Dell, IBM, SUN, Apple, Cisco and generic ‘no name’ hardware; several different types of SAN technology and two different styles of blade servers. Some of the hardware was over thirteen years old, and still running in production. The range of different operating systems was very difficult to manage, with OS’ ranging from Windows Server 2000 and 2003, Fedora, Red Hat, FreeBSD, MacOS and Ubuntu.

In addition to the physical environment, UC had a fledging virtualisation environment with twenty-nine VMs spread across four physical hosts. The virtualisation technology product used at the time was Microsoft Virtual Server 2005. Virtual Server 2005 was Microsoft’s first virtualisation platform, based on technology purchased with Connectix in 2003; and by all accounts failed to make much of an impact. This was also the case at UC, with a difficult to use web-only, management interface, marginal reliability and no 64-bit guest operating system support.

IV. Choosing Virtualization

The environment described above posed a number of issues for IT support staff, including: The management of multiple support contracts with different vendors and resellers; duplicate services and systems; inconsistent OS configurations; multiple backup and recovery environments; and finally, unpredictable interactions between applications and services sharing common hardware were also a huge issue. There were two general options available to UC, the first a physical hardware consolidation project and the second a virtualisation project.

Virtualisation was chosen as it posed several key benefits over and above a traditional hardware consolidation project. The key reasons behind the decision were around making the environment portable in the future, providing quick deployments of new servers from template VMs and the ability to standardise and consolidate much of the physical hardware while maintaining service separation. The decision to invest in Virtualisation was made and coincided with the release of Windows Server 2008, which included the first revision of Microsoft Hyper-V. After looking at the alternatives and the TCO for each of the major products, the decision was made to implement Hyper-V in conjunction with the rollout of Server 2008.

Initial steps - The initial deployment was quite modest, with the four IBM x366’s being used to host the existing 29 VMs using Virtual Server 2005 R2 being cycled out for Server 2008 RTM and Hyper-V Beta3. The initial strategy was to migrate all the VMs off one of the existing hosts, re-build with Hyper-V and then migrate enough VMs to the new platform until a second host was free and so on. Once our existing VMs were migrated and Virtual Server 2005 had been decommissioned testing began with the new Physical to Virtual (P2V) migration tools. By the time Hyper-V was in full release (June 26, 2008) UC had migrated several workloads including DHCP, Script Hubs, License Servers and Web Apps to the new virtual environment. In order to leverage one of virtualisation most prominent advantages work began on creating standard templates for each of our preferred operating systems; Windows Server and Red Hat Enterprise Linux. In addition to basic templates which included just basic customisations, UC also developed templates for the standard operating environment, to enable quick access to test machines and a LAMP (Linux, Apache, MySQL and PHP) template for quick deployment of new web apps.

V. The Changing Environment

Within the first year of deployment, UC had exchanged approximately fifty physical servers with virtual servers. With more than double the number of virtual servers than physicals three years after deployment (See “Table 1. Physical and Virtual Server Numbers”).

Table 1
Physical and Virtual Server Numbers

Time	Number of Physical Servers	Number of Virtual Servers
Quarter 1 2008	248	29
Quarter 1 2009	199	85
Quarter 1 2010	147	129
Quarter 1 2011	142	287
Quarter 1 2012	149	291

Between the beginning of 2010 and 2011, the number of physical servers didn’t reduce greatly, however the number of virtual servers did increase significantly. To the point where at the beginning of 2011 there was more than twice the number of virtual servers than physicals. During 2011 the environment stabilised, while the virtual environment continued to grow and change, support staff had also become better at managing the environment and cut down on 'VM sprawl'.

VI. Hardware Consolidation And Standards

As stated earlier, in 2008 the server environment was represented by no less than six different hardware vendors, and while in the first year the makeup of the physical environment didn't change dramatically by 2011 the vast majority of hardware had been replaced with little or no interruption to services. Table 2 shows the changes in the Physical Environment from 2008 to 2012. While the total number of servers between 2010 and the beginning of 2012 didn't fluctuate greatly, it is clear from the breakdown of hardware vendors that the vast majority of the environment has changed. There are several ways in which virtualisation has been able to assist in the standardisation of hardware, the first and easiest to identify, is through direct one-to-one replacements of physical servers with virtual servers. Secondly, virtualisation enabled UC to create test and development environments for systems that didn't have them previously, either due to the cost or practicality. With what are essentially 'disposable' servers which are quickly available and practically free, UC was able to test if an application would work on a VM or in a particular configuration with minimal or no risk to the production systems. The practical result of which was that systems were replaced quicker and easier because there was more confidence in our ability to test.

Having the bulk of the server environment virtual also allowed for decommissioning of old hosts quickly and easily through the ability to easily migrate VMs between hosts. This effectively reduced the hosts to a commoditised, and easily replaceable, RAM, Processor and Storage resource. This same flexibility allows UC to continuously update and improve the physical environment with minimal impact on the services being presented.

Table 2
Changes in The Physical Environment From 2008 to 2012

Vendor	2008	2009	2010	2011	2012
Dell	59	47	30	95	118
IBM	63	63	61	30	21
Sun	65	59	37	14	7
Apple	22	10	5	3	3
HP	9	7	3	0	0
Other	30	13	11	0	0
Total	248	199	147	142	149

In 2008, all of the circuits in the UPS distribution board were being used; therefore UC was unable to deploy new racks without either decommissioning an old rack first or upgrading the distribution board. Over the first few years, nine server racks were decommissioned, with three new ones added to support replacement equipment; the net result being six fewer racks in the data centre. The reason for implementing new racks as opposed to reusing space in some of the nine decommissioned racks is to begin to reconfigure the room to allow for more efficient air flow from a cooling perspective. Since 2010, several more racks have been decommissioned and subsequently replaced with new racks; there are currently twelve free circuits in the original distribution board. Reclaiming this space allows UC to grow again, without needing costly upgrades to the data centre facility.

Since embarking on the original virtualisation project power consumption has been dramatically reduced. 2010 was the first year that the entire server environment was in a central location and able to be monitored effectively in terms of power consumption. Prior to 2010 equipment was still distributed among several locations inherited during UCs IT centralisation process in 2007. The reduction in power is arguably far more dramatic than demonstrated in Table 3. One hundred and one physical servers were decommissioned between the beginning of 2008 and 2010; a large portion of which was virtualised. The remainder of the 101 servers were decommissioned as duplicate services.

Table 3
Reduction in Power Consumption

	L1	L2	L3	Percentages
Q1 2010	62	76	90	76.00%
Q1 2011	53	81	68	67.33%
Q1 2012	60	75	59	63.75%
Power Usage Change				-12.25%

Table 3 shows the load (as a percentage) across each line or phase through the UPS. Over a two year period, there has been a reduction of 12.25%. This drop is primarily due to the further consolidation through virtualisation and the replacement of older server hardware with newer, more powerful and more energy efficient servers. The powers savings shown in Table 3 are further diminished by a number of new large systems begin commissioned during that time period; including almost a dozen new iSCSI San's. As discussed earlier, the capacity existed for these deployments, was made possible through virtualisation.

VII. Hardware Decisions

Throughout the virtualisation process, UC has made some hardware decisions that are uncommon for virtualisation environments. The 'norm' for virtual environments involves the hosts being clustered with a SAN technology supplying the storage for VMs. This approach is usually preferred as it allows for 'virtually' no downtime when moving VMs between hosts, VMotion (VMware) or Live Migration (Hyper-V). However, it comes at a high cost in terms of initial capital expenditure and increased complexity. UCs approach to their virtualisation environments was to primarily use medium sized rack-mounted servers with a sufficient amount of internal storage when compared to average VM sizes and the other resources in the host. The average specification, for host purchased over the past four years is a 2 RU Rack Mounted Server with 32-96GB Ram, Dual Quad Core Xeon Processors, 2x600GB 15k SAS HDDs configured as a mirror for the host OS and as a backup/restore space and 4x600GB-900GB 15k SAS HDDs configured in a RAID 5 for VM storage. Each host can run, on average, between fifteen and thirty virtual servers, depending on the age/sizing of the host and the requirements of the VMs.

As with many other IT departments, one of the driving factors behind virtualisation was to save costs, not only in ongoing operational costs but also in the initial capital outlay of hardware. With that in mind, it didn't seem sensible to invest heavily in SAN technologies to support HA for all VMs or features like Live Migration without a strong business case. To that end, two key questions were asked: "*Do all VMs need HA?*", "*If 'Live Migration' was available will, or can, we use it?"*

"*Do all VMs need HA?*" After analysing the list of services that were intended for virtualisation, UC determined that none of the services in the first round of virtualisation was required HA. Four years later, less than 10% of the production virtual servers are preferred to be HA. There is a large variety of reasons why a virtual server was found to not need to be highly available. For example, domain controllers are already clustered at the application layer, as are many other systems like mail filtering and some web services. Other systems were simply not deemed important enough to the business to warrant high availability. Conversely, it was determined that in the vast majority of cases it was preferable to provide HA for services at the application layer as opposed the infrastructure layer; the downside in many cases being increased complexity for the application, but the benefits include the removal of the SAN as a potential single point of failure, more flexibility when the application was 'cluster aware' and finally there was a reduction in capital costs.

"*If 'Live Migration' was available will, or can, we use it?'*" As with many comparable sized organisations, any changes to a production environment must be approved through a Change Control process. A sample of business owners and IT managers were asked if they felt comfortable migrating systems during business hours; despite a very low risk of any problems and effectively no downtime, the answer was overwhelming "*no*". It was determined that there was no strong business case to implement our virtualisation environment all with all of the VMs stored on SANs. Instead, internal storage was implemented exclusively for the first two years. In the third year, UC implemented a small Hyper-V HA cluster to house the small percentage of VM that were deemed to require high availability and were unable to use application layer clustering tools. Subsequently, after the release of Hyper-V R2, quick migration became available, which allow for a VM hosted on local storage to be moved between hosts with only a few seconds of downtime. This new feature has been heavily used since its release in UCs test and development environments, but not in production.

There are bound to be several benefits and improved efficiencies gained in any implementation of virtualisation. The 'big' gains being; the ability to easily consolidate hardware platforms; reclaiming physical space, whether that's some rack units or entire racks; and the reduction of operational costs through reduced power consumption, cooling requirements etc. Depending on the hardware and infrastructure decisions made, you should also be savings on capital outlay. Each of these benefits was realised at UC through this virtualisation project; however, in addition to these major points several key benefits were made; each of which was either unexpected or it wasn't fully realised how much impact they would make. These are; VM templates, Quick deployment times, the self-service portal, easier capacity management and centralised, and easier management of the virtual environment.

Prior to virtualisation, UC was using tools like PXE boot servers to deploy operating systems via a network and using group policy to configure new servers. After Hyper-V was deployed, support staff continued to use the same workflows and processes they were familiar with. Creating templates stayed for several months on the 'I know I should but I don't have time right now' list. Once the first sets of templates were created, it was evident that they would have a huge impact on the workflow. Prior to templates, it would take at least a few hours to build a new system, sometimes longer, once the entire configuration had been completed and the many progress bars had been endured. Using a VM template, the time to build a new system was reduced to 20-30 minutes. Moreover, the actual staff time was reduced far more dramatically, while previously a staff member would need to visit the process periodically to perform the next step. Using templates, all of the decisions are made at the beginning of the process means that a staff member could take 2-3 minutes to start the deployment, leave to work on something else and when they return it's complete. The template includes any domain bindings, and subsequently group policies and any standard tools and drivers etc. making the process very efficient in terms of staff time and in maintaining uniformity. The above points presume that the new physical server is present, racked and cabled. The difference becomes even more dramatic when comparing an unplanned physical vs. and unplanned virtual server. Prior to virtualisation if the last-minute requests came in for a new server/service there were two possibilities, the first was to add the requested software to an existing server, effectively increasing is the complexity and management requirements of that system; the merits of maintaining service separation will be discussed later. The second option is to order a new physical server, wait for approvals, wait for delivery, rack and cable the server before getting to a point where you can compare the process with a virtual server deployment. Arguably, there could be circumstances where you need new virtual server hosts and require completing the same process. Since UC took the virtualisation path, there has not been an instance where a VM deployment had to be delayed due to lack of hardware; this is due to effective capacity management and forward planning.

In addition to the time savings and the other benefits gained through using templates, there were further efficiencies to be gained through the centralised management of the virtual environment. In UCs deployment that came in the form of Microsoft's System Centre Virtual Machine Manager (SCVMM), which allowed administrators to have a good view of the entire environment from a single interface. SCVMM also managed the templates, provides a central location for ISO images and helped to make placement decisions during VM deployments.

As UCs virtual environment has grown, SCVMM has proved to be an increasingly important tool. Providing a clear division between different host groups; for example, UC divides their hosts into three broad groups, production, test and development and disaster recovery. One of the key benefits of using this central management tool is during deployment, when the 'new VM wizard' is triggered the admin can choose which group the VM should be in and then SCVMM ranks all the available hosts within that group based on the size and configuration of the VM, for example, which has the most RAM, disk space and required network connectivity. As a side note, separating out hosts into groups also helped to simplify our backup and recovery configuration by allowing us to treat entire groups of hosts differently with the knowledge that any new VM will be backed up appropriately for the environment it's in.

As mentioned above, the effect of templates was greater than expected. Once 'the word got out' about the ease and speed of deployment, requests starting coming thick and fast for new virtual servers. Almost overnight every crazy idea, development project and potential improvement or upgrade to existing system seemed viable, at least for a trial. Almost as quickly, the requests for, "can you: "turn off my VM", "turn on my VM", "mount this CD", "check this setting", etc. had completely consumed all of the server administrator time saved through quick deployments and centralised management.

To help ease this constant stream of requests UC decided to try implementing the self-service portal (SSP). Within SCVMM administrators are able to assign resources, in the form of a host or group of hosts, to either a user or an AD group. UC was also able to assign 'quota points' to each member of the group, allowing them to 'own' a limited number of virtual machines, so not to consume all of the resources shared among a group of staff members. The self-service portal gives non-administrator staff members to deploy new VMs from templates, turn them on and off, remotely control the virtual console and mount media all from a web interface. The implementation of the SSP was quite dramatic, with almost all requests stopping altogether, once again freeing up time for the server administrators.

Virtualisation made forward planning and capacity management much simpler than previous years. Rather than interviewing every business area and planning for hardware purchases a year in advance for every system, no matter how small. In recent years UC has been able to simply maintain a percentage of headroom in each virtual environment which caters for the majority of small systems, and then only plan specific hardware environments for the larger systems that either couldn't be virtualised or had separate funding. A point to note, at the beginning of 2012, there was approximately 50% headroom in all three virtual environments; meaning that the number of VMs could be doubled before all of the virtual server hosts were filled. Take that into consideration with the 12% reduction in power consumption and 100 fewer physical servers than the beginning of 2008. It is clear that the server environment is not only dramatically different in terms of the hardware makeup and ease of planning, but it is also far more flexible.

Traditionally a single physical server would host a number of different, and over time often unrelated, applications, databases and services. Leading up to the virtualisation project, UC often struggled with this style of resource sharing. It is difficult to perform maintenance on one application when it is sharing resources with other services. It can also result in a complete inability to upgrade a key system due to the limitations of another application. For example, if a single physical server hosts two applications on a windows server 2003 R2 32bit system; and the latest update to one of those applications requires a move to a 64bit operating system and the other was only supported on 32bit, the upgrade would not be possible. This same stalemate can occur with support platform versions such as .NET or Java. Historically, the only option would be to separate the two applications either onto separate physical servers or forgo the upgrade. In the same situation using virtual machines; both 32bit and 64bit operating systems can happily co-exist side by side, and independent of each other.

With every new implementation, there are bound to be new problems and responsibilities, virtualisation is no exception. The most notable challenges faced at UC relate to the explosion of new VMs and how to prevent duplicate services sneaking into the environment. There are also some new tasks that are required periodically to maintain the new functionality gained; for example, the virtual machine templates need to be updated regularly with new patches or base configuration changes. Other additional workloads include an increase in the number of operating systems to patch and secure, while helped through keeping the templates up-to-date and having the basic configuration consistent, there is still more maintenance to be done. UCs original set of VM templates included an automatic domain binding, this resulted in a huge number of stale AD objects for development systems created and deleted within a short time span. As a result, it was determined that the templates would not automatically bind to the domain, exchanging the constant clean up the task for the odd request to bind a new VM to the domain.

Once the confidence had grown in the new Hyper-V platform, the requests for new VMs came thick and fast; with the deployment of the SSP, developers and other staff members were given access to create and manage their own virtual machines. This new found freedom to create new machines, on a whim, resulted in a tremendous number of VMs being created and not always being deleted. Several steps had to be taken to manage this 'VM Sprawl'. The most effective tool in managing the sprawl was asking each area consuming this new service to either fund their own host(s), or at least pre-commit to the number of resources they require. By separating each 'sandpit' into physical groups of hardware it provided a nice segregation which was simple to manage. It also resulted in the managers of certain areas to feel responsible for their finite resources, motivating them to ask their staff to clean up old VMs or to archive off machines while they are not being used. This in conjunction with assigning quota points to individual team members resulted in each area is responsible for managing their own 'sprawl' and proved to be quite effective. This strategy of physically separating environments worked so well in their environment, UC implemented in a broader fashion. UC created three large groups; two logically segregated groups for production and for testing/development, and a smaller physically separated environment for disaster recovery. Each host either funded or reserved for a team became a part of the larger testing/development environment as a sub-group and assigned, through SCVMM, to each area to maintain the level of ownership and self-management created.

Similar to VM Sprawl, service duplication refers to multiple VMs sneaking into the environment that provides the same function. In UCs experience this was more common for web servers. For example, UC has created a 'set' of web servers for all small sites not large enough to warrant their own system, i.e. one with IIS and .NET and another with Apache and PHP. Service duplication occurs when multiple servers are allowed into production without any discernible difference in the service they are providing. If a new service requires IIS and .NET and a server with the correct versions is already running then why create a second? Conversely, a balance needs to be struck between service duplication and service separation. Often the decision whether to create a new VM or use an existing one be decided by a single technical staff member. This is also a decision that cannot be managed through technical means. At UC, the separation of the production and testing/development environments between separate hosts has helped to formalise the move form a test/development system into production via the UC change control process. While not perfect, the process allows for a considered decision to be made, and even if it 'on the line' at least the discussion has occurred. This appears to be working well so far.

VIII. CONCLUSION

After four years of aggressive virtualisation and consolidation, the virtualisation environment is likely to remain pretty much the same in the near future. There are a small number of remaining physical servers earmarked for replacement with VMs, and each time a physical system reaches its 'end of life', the decision will be reassessed if it should remain physical or be moved to virtual servers. It is expected that the percentage of virtual machines in UCs environment to plateau around 70%, including the host servers. In summary, some of the most significant changes include some dramatic changes to the environment. The physical changes include; a reduction of physical servers from 248 to 149, a reduction in power consumption of 12.25%, after the bulk of servers was rationalised, and removing 6 racks worth of equipment and reclaiming future growth in a previously 'maxed out' facility. Some of the less tangible but equally dramatic changes include; the ability for non-administrator staff self-provision and manage their own virtual machines, and the built-in overhead to double the number of virtual machines without the provision of new hardware. All pro's and con's considered, it is clear that virtualisation at the University of Canberra has been a great success. Providing the university with the ability to change and adapt quickly while reducing power consumption, reclaiming floor space and standardising hardware; a good trade for a few extra challenges around service duplication, VM sprawl and some additional maintenance tasks.

This research has few limitations. This paper presents an analysis of VMs. But to make further conclusions the study needs to follow a more systematic method of qualitative analysis. The challenges that need consideration when performing research in the area of Green IT is the time-frames for data collection in years to come and calculate the benefits received in more details in terms of sustainability. It is intended that this research will continue by using Green IT metrics to evaluate VMs benefits in future research.

REFERENCES

[1] G. Boccaletti, M. Loffler, J. M. Oppenheim, "How IT can cut carbon emissions,"*McKinsey Quarterly*, 2008. https://www.mckinseyquarterly.com/How_IT_can_cut_carbon_emissi , 2008.

[2] R. J. Creasy, "The Origin of the VM 370 Time-Sharing System", *IBM. Journal of Research and Development,* vol. 25, pp. 483-490, 1981.

[3] J. A. Kappel, T. J. Velte,, "*Microsoft virtualization with Hyper-V*", New York: McGraw Hill Professional., 2009

[4] S. S. Jadhav, M. S. Kukreja, "Virtualization Impact for Green Technology", *International Journal of Research in Advent Technology* (E-ISSN: 2321-9637) Special Issue National Conference "ACGT 2015", 13-14 February 2015

[5] K. T. Rao, P. S. Kiran, L. S. S.Reddy, "Energy Efficiency in Datacenters through Virtualization: A Case Study. ," *Computer Science and Technology* vol. 10, pp. 2-6, 2010.

[6] J. E. Smith R. Nair, "The Architecture of Virtual Machines," *IEEE Computer,* pp. 32-38, 2005.

[7] D. S. Reddy, A. Govardhan, D. Anji Reddy, "Nodal Relational Approach to Database Virtualization", International Conference in Clod Computing and e-Government 2012.

[8] Y. Ghorpade, S. Ghorpade, T Bennur, H. S. Acharya, "Server Virtualization: A cost Effective and Green Computing Approach Towards Educational Infrastructure Management", *International Journal of Advanced Computational Engineering and Networking,* Volume- 1, Issue- 3, 2013

[9] V. R. Vedula, V. R. Mandapati , "Server Virtualization and Network Virtualization in Cloud Computing", *International Journal of Clod Computing and Service Science*, Vol 3, No 4. 2014

[10] M. Uddin, A. Shaha, A. Abubakara, I. Adelekeb, "Implementation of Server virtualization to Build Energy Efficient Data Centers", Journal of Power Technologies 94, (2) pp 1–10, 2014.