S. Jayanthi *et al*, International Journal of Computer Science and Mobile Applications,
Vol.3 Issue. 2, February- 2015, pg. 11-21 **ISSN: 2321-8363**

Green Cloud Computing - Resource Utilization with Respect to SLA and Power Consumption

S. Jayanthi[1], Srinivasa Babu[2]
[1]M.E Scholar, Department of Computer Science and Engineering, Adhiyamaan College of Engineering, Hosur, India
[2]Professor, Department of Computer Science and Engineering, Adhiyamaan College of Engineering, Hosur, India

Abstract

Many organizations are working towards reducing the carbon footprint of their data centers; i.e. reducing their power consumption. Server virtualization is used to decrease power consumption by consolidating multiple servers onto a few physical machines. Virtualization provides increased flexibility by providing a means to dynamically move virtual machines from one physical machine to another. Using resource utilization as a proxy for power, we build models of power consumption for individual server types, and use this information along with business value and SLA information, to efficiently allocate virtual machines to physical machines.

***Keywords**:* Green computing, Cloud computing, data center, virtualization, SLA.

1. Introduction

Cloud computing is a style of computing in which dynamically scalable and often virtualized resources are provided as a service over the Internet. A number of advantages of cloud computing for the enterprise have been recognized. Among these are improved TCO, reduced infrastructure costs, improved business agility, converting fixed costs to variable costs, and of course, with respect to power, acquiring IT services which can be scaled as needed, even on a temporary basis to deal with peaks in IT requirements, and which do not impose power capacity constraints [18]. Although enterprises have been hesitant thus far to adopt cloud computing broadly as a source of IT services, there are cases where cloud computing might be advantageous, without raising the typical concerns about security, reliability, meeting SLAs, and other issues raised by cloud computing which have not yet been fully resolved. [18] suggests that there are four cases where the cloud should be considered by CIOs, namely, for: (1) new initiatives where budgets are very constricted; (2) business processes which have widely varying or unpredictable load patterns; (3) services provided by non-core systems which are commoditized; and (4) systems where infrastructure management and operations costs are high. Over the years, as data centers consume more and more power, their carbon footprint has increased. Figure 1 shows a graph from McKinsey"s report[28] that places data centers" carbon emissions close to the emissions of the airline industry and steel plants. Data centers need to be more energy efficient and environment friendly.

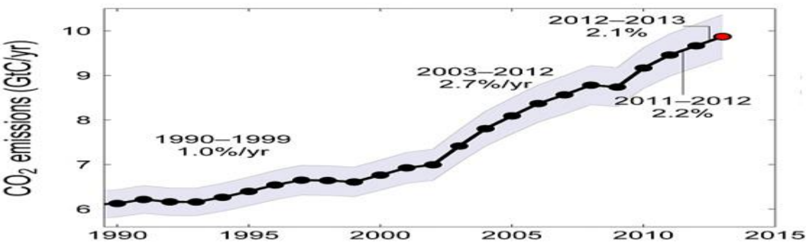

Figure 1: Carbon dioxide emissions of data centers. Source

Based on the above, clearly there is a need for the enterprise to reduce the growth in data centre power demand. This requires a more comprehensive and effective approach to power management than the piecemeal approach of employing various best practices, with little or no attention to coordination. Such a comprehensive management methodology would benefit the enterprise by: (1) increasing the life of the data center; (2) increasing profitability for the enterprise by maximizing ROI for the data center; (3) reducing the power consumed by idle servers; (4) increasing revenue by allowing the enterprise to better leverage the capacity of its data center resources; (5) preventing "free riding" within the enterprise by allowing development of an improved chargeback model, which more accurately charges business units for the capacity that they use; and (6) promoting a greener data center, which may have public relations and perceived corporate environmental responsibility benefits. How this approach offers all these benefits is described in "Chapter 5: Research Contributions". In the following section however, the factors in data centers which have led to the current crisis in capacity management and power management are characterized.

In the past, data centers (DCs) were designed and managed with availability, reliability, and contribution to business value as the principle goals. Capacity limitations were not generally significant, at least with respect to power; i.e., power did not constitute a constraint on the operation of the DC in practice. If power was considered at all in the design or management of the data center, it was only with respect to achieving a data center lifetime which was considered reasonable, typically in the range of 15 - 20 years [1].

One of these developments being that the steady increase in the use of information technology (IT) in today‟s enterprise has led to a rapidly increase in the demand for power in data centers. In contrast to the expectation of a 15 - 20 year life cycle of the past, Figure 2 dramatically illustrates the change. Even data centers which have significantly expanded capacity within the last five years are still facing the prospect of insufficient capacity in the next four years.

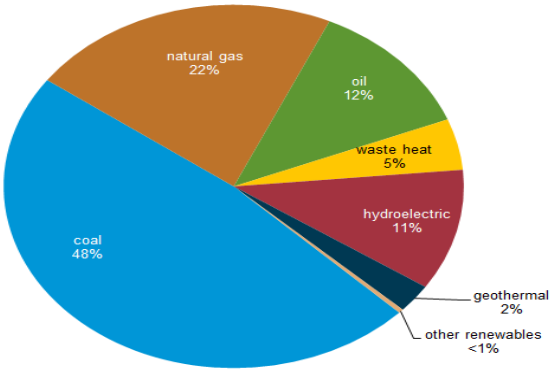

Figure 2: Percentage of Data Centers that anticipate the need for additional capacity

Now, these data centers that have seen a huge increase in capacity have another problem. It is difficult to trace how the power is being used. Once a request enters the data center, we usually do not have information about which application is servicing the request. In addition we neither have information about which particular server is being used, nor do we know the effect of servicing that request on power consumption. In other words we do not have traceability with [1] notes that processing requirements in data centers, and the attendant need for power, grew more than 15 times between 1990 and 2005.

Further, successful enterprises involve constant change due to innovation or attempts to increase efficiency. Indeed, it is often this continuous change that has helped the organization to survive. Constant changes to the enterprise mean constant changes to the different systems that form the enterprise. These changes are reflected in terms of infrastructure changes due to additions to the existing architecture by extending it in lieu of redesigning it, and other implementation changes.

2. Problem Statement

As the enterprise grows, along with becoming more complex, these systems also grow in size and require the addition of new physical infrastructure. This causes the datacenter to grow in size. Typically, data centers are built with the possibility of expansion in mind, but it is not only additional space, but also electrical capacity, that is needed. The ability to expand is always limited and thus at some point the enterprise must build a new data center once the capacity of its existing data centers is reached.

The complex nature of today"s enterprise is reflected in the data center. As the enterprise grows, it is constantly adapting to changes in the environment. The acquisition of new companies and merging of IT systems is one example of change the enterprise faces. Another reason for change can be the adoption of new technologies. These changes over time, lead to a heterogeneous data center, which grew in an ad hoc manner, or sometimes in a more systematic manner. Either way, it leads to a very complex system where different systems interoperate with each other in a very knotty configuration to produce the desired service.

This complex nature of the enterprise leads to the inability to measure and manage power in a comprehensive way. The result is that the business can neither track nor optimize ROI (Return on Investments) or TCO (Total Cost of Ownership) for investments in data center infrastructure (hardware, software, facilities, human resources, etc.); but the business has no metrics to determine power costs associated with a particular infrastructure element. This gap in traceability also means that the business cannot make good decisions with respect to infrastructure investment or management.

The enterprise is unable to understand the dynamics of the trade-offs between performance and power consumption, and therefore cannot make sound choices about power expenditures versus value added for the enterprise. On the capacity management front, this amounts to an inability to make informed choices regarding DC lifetime, because the enterprise cannot quantify the tradeoff between DC lifetime reduction and the benefit to the enterprise of making a particular change in DC management or operation which adds some amount of value for the enterprise. This moves the DC closer to "power out," i.e., a state where the power infrastructure capacity of the data center has been reached.

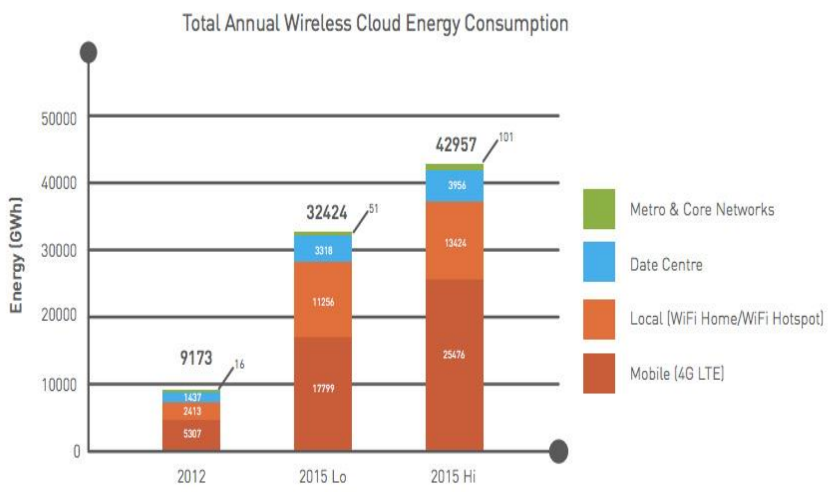

Figure 3: Power Consumption of Data Centers

2.1 Data Center Management Problem

Another problem for data center managers is the difficulty in understanding the dynamics of the trade – offs between performance and power consumption. Having such detailed information can enable the enterprise to look into the Service Level Agreements (SLAs) set for the various applications, and would force them to justify the requirement of a particular SLA. Studying the business value would help the enterprise see the value generated by a particular SLA, and how much the business value would be impacted if there was a reduction on the SLA. Reducing the SLA a little bit could result in huge power savings – this is because higher SLAs could mean more availability, more redundancy, etc. which directly implies more power consumption.

There is therefore a need to model the relationship between service level agreements (SLAs), operating level agreements (OLAs), and the power consumed. Such a model would allow the enterprise to understand the relationship between IT performance and power consumed, and therefore to make more informed business decisions about how power is used.

3. Problem Analysis

This section looks at the problem in more detail and describes the problem and the questions the solution is expected to answer.

3.1 Data Center Scenario

Figure 4 attempts to describe the problem at the data center. The data center consists of a huge collection of servers, or physical machines, which have to be provided with power and cooling. The lowest block represents the data center. The four blocks on top of the data center represent four physical machines.

Now, if we consider the data center environment to be consolidated and virtualized, then every physical machine will contain one or more virtual machines. If there is no virtualization, then we can treat the machine as if it is housing a single virtual machine, as can be seen in the figure – the backward right side physical machine does not have a hypervisor installed on it – it can thus run only one operating system instance, but it can still run multiple applications.

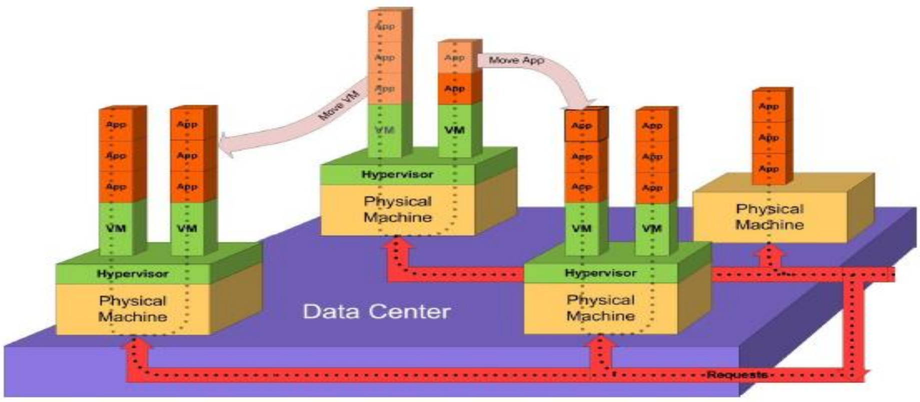

Figure 4: The data center scenario

Every virtual machine runs at least one application, and these applications service transactions, which are incoming requests. The requests are shown as black dots in the figure.

The incoming requests, which are represented as dots in Figure 4, have business value associated with them. We want to be able to trace this business value down to the level of resource allocation and power consumption and make more informed decisions when allocating virtual machines to physical ones, or deciding which applications go together on a single physical server.

3.2 Lack of Power Modeling

Power modeling, or the ability to predict power consumption by hardware, and power requirements of applications, is necessary to evaluate the current operation of the data center, and also to predict the effects of changes which are being considered. Any approach which cannot predict power with reasonable accuracy must be based to a large extent on guesswork.

There are two general approaches to the measurement of power. The first approach is direct measurement of the power consumed using an instrument such as a wattmeter. While the most direct, this approach has several limitations: (1) The additional hardware cost; (2) The time delay in collecting and aggregating all the data collected from the watt meters; (3) The fact that very few DCs have been built with such direct measurement capability, and therefore it would have to be retrofit in virtually all data centers.

3.3 Inability to trace the use of power

Another difficult challenge is tracing the use of power - the ability to identify how power is used to do IT work, or execute transactions, in the data center. Since a given transaction is typically executed by various physical, and perhaps virtual, servers, in a distributed environment, accurate modeling and sufficiently granular measurements are required to enable this type of tracing.

Once the subsystem resource use of a given service component running on a particular piece of hardware can be modeled, tracing the power consumed in processing service requests - using the power model described above - also becomes possible.

3.4 Inability to curtail rapid growth in power demand

As new hardware and applications are added to the data center, the management is unable to determine precisely how power capacity will be affected. Although various best practices have been followed in most data centers, these are to some extent "one size fits all" solutions which may not work well, or at least, may not be optimum, in a given data center. Business Problem, best practices have not succeeded in significantly reducing the rapidly increasing demand for power. The inability to measure, model, and trace power prevents the enterprise from determining if a given best practice will result in a net benefit or detriment with regard to the demand for power in its data center.

3.5 Inability to quantify the tradeoffs when using new green technologies

The data center management teams have to deal with the adoption of new green technologies. New technologies provide great benefits, but also pose new management problems. Unless there is a means to quantify the benefit gained by using a particular technique, there is no way for the management team to quantify the power savings. As an illustration, consider virtualization, which is a technique which enables a single piece of hardware to run multiple virtual machines. These virtual machines (VMs) can be dynamically moved from one physical server to another, to provide dynamic load balancing.

3.6 Inability to Match Applications to Hardware

The aim behind application to hardware matching in the data center is to optimally match applications to hardware with respect to the amount of resources consumed by the application. This is because the types of hardware running in the data center can be heterogeneous. For example, some servers could be very efficient in performing CPU intensive computations, but consume a lot of power for I/O intensive applications. Thus, a lot of power can be saved by using the most appropriate hardware for a particular application. In this sense, this type of matching is similar to what Nathuji et al. call *platform heterogeneity* [12].

3.7 Need for a combined approach of IT and cooling

The two key factors listed below, make the use of static measures alone insufficient for managing power in the data center. Accordingly, the thesis proposes that a Sense and Respond (SaR) architecture is necessary for optimizing power use.

First, power consumption in modern data centers varies widely, even over short periods of time. Two major consumers of power in data centers are IT and cooling [20]. Without even considering the power used by the cooling infrastructure, the dynamic variation due to IT power fluctuations is significant. [21] reports a range of variation between 45% and 106% for typical enterprise class servers.

Second, although current servers use relatively less power at idle than previous generations of hardware, every server consumes some power at idle, while at the same time it does no useful IT work. As can be seen from the range of power variation cited above, even modern hardware typically consumes at least 50% of its peak power consumption even at idle, as can be seen in Figure 19. For this reason, significantly reducing server idle time, in order to increase server utilization as much as possible, has become critical to minimizing power use while still maintaining required performance levels.

3.8 Benefits of Traceability

To understand what traceability means, and what its actual implications are, let"s look at an example. Consider an online site that sells cell phones. Customers visit the site and order a phone, and the business has the phone shipped to their home. The system managing these transactions has many subsystems, for example, an inventory management system, a system connected to the suppliers of the phones, the courier company, customer support, etc. For simplicity, we"ll look at only two systems in our example. One system is the one that supports the website that the customer uses to place the order, and the other is a website the suppliers use to indicate that they have sent a shipment to the company. Figure 5: Need for traceability illustrates these two systems at an abstract level.

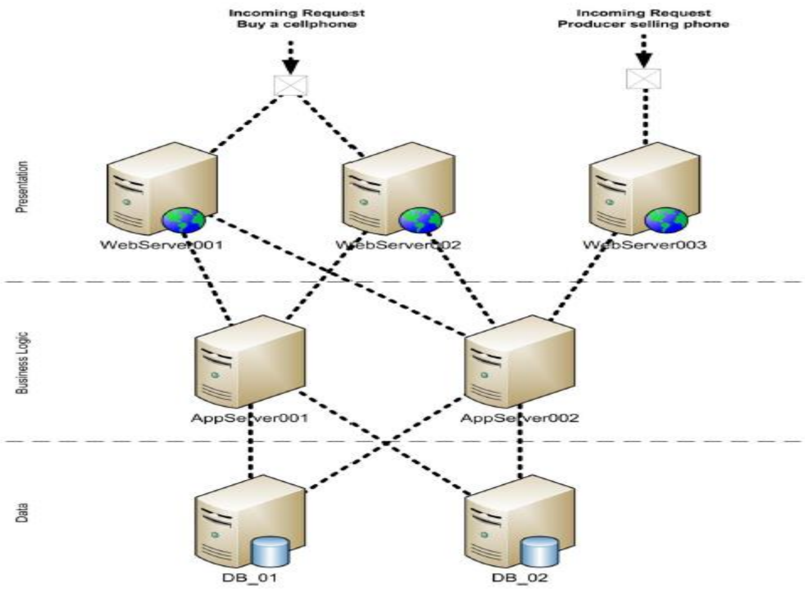

Figure 5: Need for traceability – two intermingled systems

From the figure we can see that the customer requests are load-balanced and split into two web servers, WebServer001 and WebServer002. The Supplier requests are serviced by WebServer003. The business logic layer consists of two servers, AppServer001 and AppServer002. The requests from WebServer001 are serviced by AppServer001 and AppServer002. However AppServer002 services requests from all the three WebServers. Finally, the data stored in database servers DB_01 and DB_02 are used by both the AppServers.

Traceability gives the organization many benefits. The unit of execution that can be traced is called a transaction. Tracing the workload and treating it as transactions allows us to do the following:

3.9 Power consumption of a transaction

The resource utilization data collected can be used as a proxy to determine the amount of power consumed. Quantifying the amount of power consumed by a transaction allows us to calculate the cost associated with servicing the requests. If the fixed and variable costs of the resources are taken into consideration, we could have a means of quantifying the energy spent, and justifying if it is worth the business value it will generate.

3.10 Carbon footprint of a transaction

Once we quantify the power consumption of a transaction, we can convert this into carbon equivalent values, and hence we can calculate the carbon footprint of transactions. This information could be used to provide new services – like inform the customer of the carbon footprint of the transaction, etc, and can be also used to make the data center more competitive, by having them inform the user of how power-efficient .

3.11 Green Practices

Existing research in green computing has led to several recommendations for making the data center green. These green practices have made their way to being best practices for any organization. This section lists these green practices and describes each of them.

3.12 Server Consolidation

Most of the servers running in a data center are running close to idle. The servers are not very efficient when run at idle, and thus consume a lot of power. However, at higher utilization levels, the servers are more power efficient. Thus, if we consolidate many applications onto a fewer number of servers, the servers can be run at higher utilizations and would be more power efficient.

3.13 Dynamic Voltage and Frequency Scaling

Dynamic voltage and frequency scaling (DVFS) is another method for reducing power, and can reduce both unnecessary server idling, which consumes power with no benefit, and also over performance of server systems, where the server executes the application at a higher performance level than is required by the relevant SLA, and therefore also uses power unnecessarily. This method is highly dependent on hardware characteristics.

Bircher and L. John [13] investigate dynamic frequency scaling, specifically in multi-core processors, but also illustrates the very significant power savings that can be achieved with little or no loss of performance. [14] shows a useful approach to control voltage and frequency scaling dynamically using a feedback control loop, and accompanying significant reductions in power consumption.

3.14 Smarter Cooling Solutions

Two papers make differing, but persuasive, arguments that data center power consumption cannot be optimally controlled without managing both IT power and cooling power in a coordinated fashion. Nathuji et al. argue that managing IT power and cooling power independently may lead to less than optimal results [16].

Niles makes the somewhat different argument in [17] that, not only must IT power and cooling power be considered and managed in a coordinated fashion, but beyond this, that cooling infrastructure in the data center must provide for row-based cooling, i.e., more granular control of cooling. This is necessary, Niles argues, because virtualization, along with high-density servers, makes dealing with the presence of hot spots in the data center a constant challenge.

3.15 Continuous Monitoring and Redeployment (Dynamic Migration)

S.Niles, in an APC white paper, discusses benefits of virtualization, including the fact that it enables dynamic migration and consolidation of workloads based on IT resource demands [7].

When a physical server is being utilized at a low level, its virtual servers can be migrated toanother physical server. This provides the opportunity for physical servers to be dynamically powered up or shut down in response to changing loads, which reduces total data center power consumption [8]. Such an ability to migrate and consolidate workloads supports a dynamic architecture for data center power management, because it enables dynamic allocation of data center resources based on the current demand, which is subject to significant fluctuation.

Recent research has focused on models and mechanisms for managing virtual machines (VMs) in the data center, in order to manage dynamic migration, and to take maximum advantage of the power savings that virtualization and consolidation offer [2, 9, 10, 11].

3.16 Heterogeneity Awareness

Nathuji et al. discuss yet another advantage of virtualization [12]; namely, it provides a looser coupling between the IT load and the underlying physical platform. This loose coupling can be leveraged by seeking optimal matches between workload characteristics and the hardware on which it is run. In typical data centers, the heterogeneous nature of the physical platforms present in the data center presents the opportunity to save significant power by matching load characteristics and hardware. Nathuji et al. report an average reduction in power use of 20% for one such approach [12]. Our approach to power modeling also supports such matching, by allowing characterization of the load in terms of resource requirements, and by facilitating the identification of hardware, through examination of its power consumption in terms of resource utilization, which is optimal for the load.

3.17 Strategic Replacement of Hardware

Strategic hardware replacement attempts to replace older, less efficient hardware with newer systems which have been designed with features which support power conservation. Once power usage can be accurately modeled, we can determine if replacement of certain hardware would result in a large enough power reduction, and if it will offset the total cost of the hardware and the person hours required to replace it. Such a strategy can only be undertaken, though, once power usage can be reliably modeled at the level of server system resource utilization. Several major hardware manufacturers are offering server systems that are significantly more energy efficient than in the past [15]; again, though, the more information DC managers have about the resource utilization of their applications, the better the choices they can make with respect to the optimum characteristics of hardware they choose to replace less efficient systems.

4. Solution Approach

The approach taken is to have a model of the datacenter, where we can trace the use of power down to the level of individual transactions performed on the data center. Doing this helps us to quantify the value generated by the power spent, and hence provide the business stakeholders with this information – specifically the power consumed by applications. The can use this data to make decisions concerning the data center.

This approach uses resource utilization as a proxy for power. The incoming requests need some resources for their execution. The incoming requests can be converted into their corresponding resource utilization, and then a model for power can be used to convert the resource utilization into power. This section gives an overall view of how this process is carried out.

S. Jayanthi *et al*, International Journal of Computer Science and Mobile Applications,
Vol.3 Issue. 2, February- 2015, pg. 11-21 ISSN: 2321-8363

As requests for service are made, these incoming requests are monitored and a trace is maintained of all the applications that are used to execute the request. A request usually interacts with more than one system. Figure 6 describes how the execution of every incoming request can be split up into executions of transactions serviced by various applications.

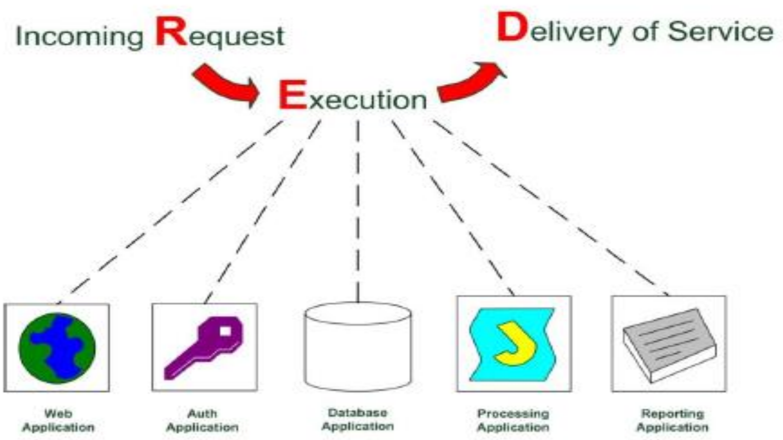

Figure 6: Tracking the application usage of each incoming request

We can therefore see that every incoming request can be split up into multiple transactions which act as requests for service for various applications. Now, each application ideally runs on at least one server, and in turn requires resources to execute the transaction

This section describes how we can use the profiling of applications and machines to understand the patterns of application usage and power consumption. We can use this traceability to optimize data center operations.

4.1 Machines running at higher optimizations - close to the SLA

By studying the SLA of applications, we can determine the minimum resources needed to maintain this SLA. Usually, in today‟s enterprise the applications are given more resources than required to maintain the SLA. When a single application resides on a machine that is designed for peak load, we are in a situation where most of the time, we are actually doing better than the SLA. This may or may not add any business value.

With consolidation and virtualization, we have multiple applications on the same physical machine. We are now in a position to run applications closer to the SLA. One of the ways to allocate applications/VMs to physical machines is to allocate high business value apps first, and then proceed to add the lower business value apps, while trying to minimize the total number of servers used while maintaining the SLA of each application. In other words, we are using the business value to allocate resources to the transactions that are executed by the apps.

4.2 Business value Profiling for SLA

Another area for improvement is aligning the SLA to better reflect business value. Many times, an enterprise assigns a random SLA to the application. The SLA must be dependent on the business value and the organization must be able to quantify and justify the SLA requirements of an application. We call this process "Business Value Profiling for SLA". It involves studying the effects of resource availability on the SLA and trying to meet the SLA with the least resources possible.

5. Conclusion

Green computing is constantly becoming more relevant, and many organizations are working towards reducing the carbon footprint of their data centers. This reduction in carbon footprint is achieved by reducing the data center‟s power consumption, which in turn results in savings for the organization. Many new techniques have been used to achieve this reduction in power. One of them is virtualization.

Virtualization helps us consolidate multiple servers onto a few physical machines, which increases their utilization, and decreases their power consumption. However, virtualization also provides us new possibilities – it allows us to increase flexibility in the data center by allowing us to dynamically move virtual machines from one physical machine to another. However, such a complex system needs a good management system to be effective. We need a system that takes power consumption and cooling power into account, in addition to server load, while allocating virtual machines to servers. There is thus a need to measure power consumption, and an even greater need to be able to predict the power consumption effects of changes in the data center.

We saw how resource utilization can be used as a proxy for power – how it could be used to measure current power, and also how it could be used to predict future power consumption. Putting this resource utilization information together with cooling information, business value and SLA information and of course the traditional server workload information – results in a mine of information that could be used to make intelligent and efficient allocations in the data center. Future implementations can take this information and implement the comprehensive data center management architecture specified in this thesis.

References

[1] MacKinnon, C. The changing data center life cycle: one day here & the next day gone. Processor, December 22, 2006.
[2] S. Kumar, V. Talwar, P. Ranganathan R. Nathuji, and K. Schwan. M-channels and m- brokers: new abstractions for co-ordinated management in virtualized systems. In Proceedings of the Workshop on Managed Many-Core Systems (MMCS), June 2008.
[3] R. Nathuji, Mechanisms for coordinated power management with application to cooperative distributed systems. PhD thesis, School of Electrical and Computer Engineering, Georgia Institute of Technology, 2008. [4] D. Narayanan. Power by proxy: a new aprpoach to measurement. August 2002. research.cs.ncl.ac.uk/cabernet/www.laas.research.ec.org/cabernet/workshops/radicals/2003/papers/paper.pdf_5.pdf.
[5] C. Isci and M. Martonosi. Runtime power monitoring in high-end processors: methodology and empirical data. In Proceedings of the 36th International Symposium on Microarchitecture (MICRO-36‟03), December 2003.
[7] S. Niles. Virtualization: optimized power and cooling to maximize benefits. American Power Conversion White Paper, 2008.
[8] J. Chase, D. Anderson, P. Thakar, A. Vahdat, and R. Doyle. Managing energy and server resources in hosting centers. In Proceedings of the 18th Symposium on Operating Systems Principles (SOSP), 2001.
[9] R. Nathuji and K. Schwan. VirtualPower: coordinated power management in virtualized enterprise systems. In Proceedings of Twenty-First Symposium on Operating Systems Principles (SIGOPS), October, 2007.
[10] S. Kumar, V. Talwar, P. Ranganathan, and K. Schwan. vManage: coordinated management in virtualized systems. www.hpl.hp.com/personal/Partha_Ranganathan/papers/2008/2008_hotac_vmanage.pdf
[11] J. Stoess, C. Lang, and F. Bellosa. Energy management for hypervisor-based virtual machines. In Proceedings of the USENIX Annual Technical Conference, June 2007.
[12] R. Nathuji, C. Isci, and E. Gorbatov. Exploiting platform heterogeneity for power efficient data centers. In Proceedings of the IEEE International Conference on Autonomic Computing (ICAC), June 2007.
[13] W. Bircher and L. John. Analysis of dynamic power management on multi-core Processors. In Proceedings of the 22nd Annual International Conference on Supercomputing (ICS), June 2008.
[14] D. Suleiman, M. Ibrahim, and I. Hamarash. Dynamic voltage frequency scaling (DVFS) for microprocessors power and energy reduction. In 4th International Conference on Electrical and Electronics Engineering, December 2005.
[15] P. Kurp. Green computing: are you ready for a personal energy meter? Communications of the ACM. October 2008.
[16] R. Nathuji, A. Somani, K. Schwan and Y. Joshi. CoolIT: coordinating facility and IT management for efficient datacenters. In Proceedings of HotPower ‟08. December, 2008. Retrieved March 3, 2008 at http://www.usenix.org/events/hotpower08/.

[17] S. Niles. Virtualization: optimized power and cooling to maximize benefits. American Power Conversion White Paper, 2008.
[18] J. Thethi. Realizing the value proposition of cloud computing: CIO‟s enterprise IT strategy for cloud. Infosys White Paper, 2009.
[19] J. Ramanathan and R. Ramnath. Co-Engineering Applications and Adaptive Business Technologies in Practice Enterprise Service Ontologies, Models, and Frameworks. Igi Global , 2009.
[20] Report to congress on server and data center energy efficiency. Technical report, U.S. Environmental Protection Agency, Energy Star Program, August 2007.
[21] J. Spitaels. Dynamic power variations in data centers and network rooms. American Power Conversion White Paper, 2005.
[22] Carnegie Mellon website, http://www.sei.cmu.edu/architecture/cbam.html [23] L. A. Barroso and U. Hölzle. The Case for Energy-Proportional Computing – Google.
[24] N. Rasmussen. An Improved Architecture for High-Efficiency, High-Density data centers. The Green Grid White Paper #126.
[25] Jack Pouchet, "Creating Energy Efficient Data Centers," Energy & Power Management, January 2007; [26] R Development Core Team (2009). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL http://www.R-project.org.
[27] Ramanathan, J. (2005 May). Fractal Architecture for the Adaptive Complex Enterprise. Communications of the ACM, 48(5), 51-57
[28] James M. Kaplan, William Forrest and Noah Kindler, Revolutionizing Data Center Efficiency, McKinsey & Company 2008
[29] Rong Ge, Xizhou Feng, Kirk W. Cameron, Performance-constrained Distributed DVS Scheduling for Scientific Applications on Power-aware Clusters University of South Carolina, Columbia, SC
[30] Kazman, R. Klein, M. Barbacci, M. Longstaff, T. Lipson, H. Carriere, J. The architecture tradeoff analysis method, Software Eng. Inst, Carnegie Mellon Univ, Pittsburgh, PA;
[31] Hardwiring Green into Infrastructure and Data Center Investments - Data Center Operations Council.

Author Biography

S.Jayanthi – B.E Degree in Computer Science and Engineering, Sri Andal Alagar college of Engineering , Chennai, Tamil nadu in 2012,where she is pursuing the M.E Degree . She is currently visiting student at Adhiyamaan College of Engineering, Krishnagiri, Tamil nadu. She has a membership in Computer Society of India. Her current interest in cloud computing.