ISSN:0975-9646

Yuvapriya Ponnusamy et al, / (IJCSIT) International Journal of Computer Science and Information Technologies, Vol. 3 (5) , 2012,5101 - 5107

Application of Green Cloud Computing for Efficient Resource Energy Management in Data Centres

[1]Yuvapriya Ponnusamy, [2]S Sasikumar

[1]Lead, Infosys Limited, Mysore. India.
[2]Associate Professor, Saveetha Engineering College, Chennai. India.

Abstract - **The perception of cloud computing has not only reshaped the field of distributed systems but also fundamentally changed how businesses utilize computing today. Cloud computing is offering utility oriented IT services to users worldwide. It enables hosting of applications from consumer, scientific and business domains based on pay-as-you-go model. However data centres hosting cloud computing applications consume huge amounts of energy, contributing to high operational costs and carbon footprints to the environment. With energy shortages and global climate change leading our concerns these days, the power consumption of data centres has become a key issue. The area of Green computing is also becoming increasingly important in a world with limited energy resources and an ever-rising demand for more computational power. Therefore, we need green cloud computing solutions that can not only save energy, but also reduce operational costs. In this paper, an architectural framework and principles that provides efficient green enhancements within a scalable Cloud computing architecture with resource provisioning and allocation algorithm for energy efficient management of cloud computing environments to improve energy efficiency of the data centre. Using power-aware scheduling techniques, variable resource management, live migration, and a minimal virtual machine design, overall system efficiency will be vastly improved in a data centre based Cloud with minimal performance overhead.**

Keywords - **Cloud Computing, Green Computing, Virtualization, Energy Efficiency, Resource Management, Virtualization, Scheduling.**

I. INTRODUCTION

The vision of computing utilities based on a service provisioning model anticipated the massive transformation of the entire computing industry whereby computing services will be readily available on demand. Similarly, users (consumers) need to pay providers only when they access the computing services. In addition, consumers no longer need to invest heavily or encounter difficulties in building and maintaining complex IT infrastructure. This model has been referred recently as Cloud computing in which users access services based on their requirements without regard to where the services are hosted. Later it denotes the infrastructure as a “Cloud” from which businesses and users can access applications as services from anywhere in the world on demand. Hence, Cloud computing can be classified as a new paradigm for the dynamic provisioning of computing services supported by state-of-the-art data centres that usually employ Virtual Machine (VM) technologies for consolidation and environment isolation purposes.

Cloud computing delivers infrastructure, platform, and software (applications) as services, which are made available to consumers as subscription-based services under the pay-as-you-go model. Clouds aim to drive the design of the next generation data centres by architecting them as networks of virtual services. So that users can access and deploy applications from anywhere in the world on demand at competitive costs depending on their requirements. Cloud computing offers significant benefits to IT companies by freeing them from the low-level task of setting up basic hardware and software infrastructures and thus enabling focus on innovation and creating business value for their services.

Green Computing is defined as the study and practice of using computing resources efficiently through a methodology that combines reducing hazardous materials, maximizing energy efficiency during the product’s lifetime, and recycling older technologies and defunct products. Green Computing enables companies to meet business demands for cost-effective, energy-efficient, flexible, secure & stable solutions while being environmentally responsible. Every data center transaction requires power. Efficiency, equipment disposal and recycling, and energy consumption, including power and cooling costs, have become priority for those who manage the datacenters that make businesses run.

Modern data centres, operating under the Cloud computing model are hosting a variety of applications ranging from those that run for a few seconds to those that run for longer periods of time on shared hardware platforms. The need to manage multiple applications in a data centre creates the challenge of on-demand resource provisioning and allocation in response to time-varying workloads. Normally, data centre resources are statically allocated to applications, based on peak load characteristics, in order to maintain isolation and provide performance guarantees. Until high performance has been the sole concern in data centre deployments, this demand has been fulfilled without paying much attention to energy consumption. Data centres are not only expensive to maintain, but also unfriendly to the environment. And now it drive more in carbon emissions than both Argentina and the Netherlands .High energy costs and huge carbon footprints are incurred due to massive amounts of electricity needed to power and cool numerous servers hosted in these data centres.

Cloud service providers need to adopt measures to ensure that their profit margin is not dramatically reduced due to high energy costs. Lowering the energy usage of data centres is a challenging and complex issue because computing applications and data are growing so quickly that increasingly larger servers and disks are needed to process them fast enough within the required time period. Green Cloud computing is envisioned to achieve not only efficient processing and utilization of computing infrastructure, but also minimize energy consumption. It is essential for ensuring that the future growth of Cloud computing is sustainable. Otherwise, Cloud computing with increasingly pervasive front-end client devices interacting with back-end data centres will cause an enormous escalation of energy usage. To address this problem, data centre resources need to be managed in an energy-efficient manner to drive Green Cloud computing. In particular, Cloud resources need to be allocated not only to satisfy QoS requirements specified by users via Service Level Agreements (SLA), but also to reduce energy usage.

II. ARCHITECTURE OF GREEN – CLOUD COMPUTING

People in IT industry are reassessing data center strategies to determine if energy efficiency should be added to the list of critical operating parameters.

Issues of concern include:

1. Reducing data center energy consumption, as well as power and cooling costs
2. Security and data access are critical and must be more easily and efficiently managed
3. Critical business processes must remain up and running in a time of power drain or surge

These issues are leading more companies to adopt a Green Computing plan for business operations, energy efficiency and IT budget management. Green Computing is becoming recognized as a prime way to optimize the IT environment for the benefit of the corporate bottom line – as well as the preservation of the planet. It is about efficiency, power consumption and the application of such issues in business decision-making. Simply stated, Green Computing benefits the environment and a company’s bottom line. It can be a win/win situation, meeting business demands for cost-effective, energy-efficient, flexible, secure and stable solutions, while demonstrating new levels of environmental responsibility.

A. Cloud:

Cloud computing is becoming one of the most explosively expanding technologies in the computing industry today. It enables users to migrate their data and computation to a remote location with minimal impact on system performance. These benefits include:

1. Scalable - Clouds are designed to deliver as much computing power as any user wants.
2. Quality of Service (QoS) - Unlike standard data centres and advanced computing resources, a well-designed Cloud can project a much higher QoS than typically possible.
3. Specialized Environment - Within a Cloud, the user can utilize custom tools and services to meet their needs.
4. Cost Effective - Users finds only the hardware required for each project.
5. Simplified Interface - Whether using a specific application, a set of tools or Web services, Clouds provide access to a potentially vast amount of computing resources in an easy and user-centric way.

B. Cloud Infrastructure:

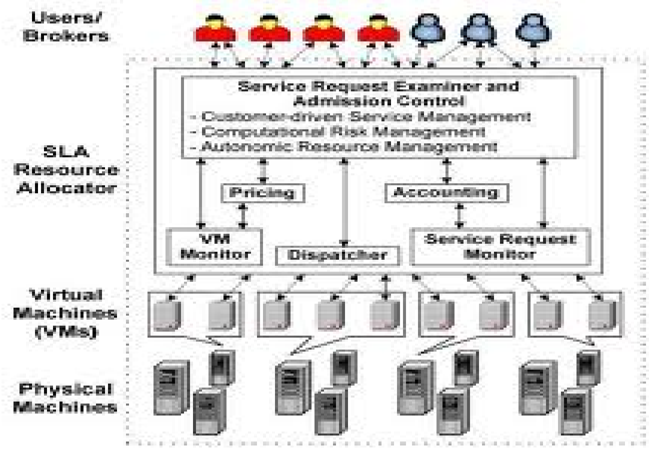

Fig 1. The high level system Architecture.

In Cloud computing infrastructure, there are four main entities involved:

1. Consumers/Brokers: Cloud consumers or their brokers submit service requests from anywhere in the world to the Cloud. It is important to notice that there can be a difference between Cloud consumers and users of deployed services.
2. Green Resource Allocator: Acts as the interface between the Cloud infrastructure and consumers. It requires the interaction of the following components to support energy-efficient resource management:

   Green Negotiator: Negotiates with the consumers/brokers to finalize the SLA with specified prices and penalties between the Cloud provider and consumer depending on the consumer’s QoS requirements and energy saving schemes.

   Service Analyser: Interprets and analyses the service requirements of a submitted request before deciding whether to accept or reject it.

   Consumer Profiler: Gathers specific characteristics of consumers so that important consumers can be granted special privileges and prioritized over other consumers.

   Pricing: Decides how service requests are charged to manage the supply and demand of computing resources and facilitate in prioritizing service allocations effectively.

   Energy Monitor: Observes and determines which physical machines to power on/off.

   Service Scheduler: Assigns requests to VMs and determines resource entitlements for allocated VMs. It also decides when VMs are to be added or removed to meet demand.

   VM Manager: Keeps track of the availability of VMs and their resource entitlements. It is also in charge of migrating VMs across physical machines.

Accounting: Maintains the actual usage of resources by requests to compute usage costs. Historical usage information can also be used to improve service allocation decisions.

3. VMs: Multiple VMs can be dynamically started and stopped on a single physical machine to meet accepted requests, hence providing maximum flexibility to configure various partitions of resources on the same physical machine to different specific requirements of service requests. Multiple VMs can also concurrently run applications based on different operating system environments on a single physical machine.
4. Physical Machines: The underlying physical computing servers provide hardware infrastructure for creating virtualized resources to meet service demands.

C. Green Cloud Architecture:

As discussed above, cloud computing platform as the next generation IT infrastructure enables enterprises to consolidate computing resources, reduce management complexity and speed the response to business dynamics. Improving the resource utilization and reduce power consumption are key challenges to the success of operating a cloud computing environment. To address such challenges, we design the Green - Cloud architecture and the corresponding Green Cloud exploratory system. The exploratory system monitors a variety of system factors and performance measures including application workload, resource utilization and power consumption, hence the system is able to dynamically adapt workload and resource utilization through VM live migration. Therefore, the Green-Cloud architecture reduces unnecessary power consumption in a cloud computing environment.

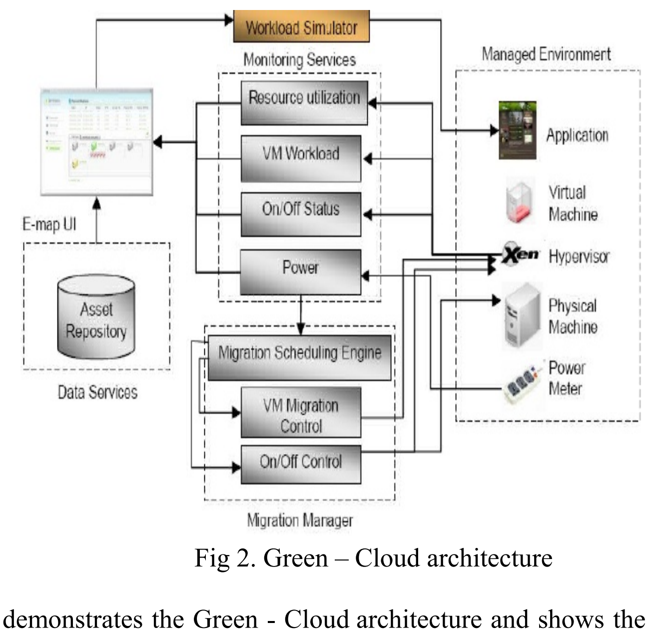

Fig 2. Green – Cloud architecture

demonstrates the Green - Cloud architecture and shows the functions of components and their relations in the architecture.

Monitoring Service, monitors and collects comprehensive factors such as application workload, resource utilization and power consumption, etc. The Monitoring Service is built on top of IBM Tivoli framework and Xen, where the IBM Tivoli framework is a CORBA-based system management platform managing a large number of remote locations and devices; Xen is a virtual machine monitor (VMM). The Monitoring Service serves as the global information provider and provides on-demand reports by performing the aggregation and pruning the historical raw monitoring data to support to intelligent actions taken by Migration Manager.

Migration Manager, triggers live migration and makes decision on the placement of virtual machines on physical servers based on knowledge or information provided by the Monitoring Service . The migration scheduling engine searches the optimal placement by a heuristic algorithm, and sends instructions to execute the VM migration and turn on or off a server.

Managed Environment includes virtual machines, physical machines, resources, devices, remote commands on VMs, and applications with adaptive workload, etc.

E-Map is a web-based service with Flash front-end. It provides a user interface (UI) to show the real-time view of present and past system on/off status, resource consumption, workload status, temperature and energy consumption in the system at multiple scales, from high-level overview down to individual IT devices and other equipment E-map is connected to the Workload Simulator , which predicts the consequences after a given actions adopted by the Migration Monitor through simulation in real environment.

Workload Simulator accepts user instructions to adapt workload e.g. CPU utilization, on servers, and enables the control of Migration Manager under various workloads. Then,

E-Map collects the corresponding real-time measurements, and demonstrates the performance of the system to users. Therefore, users and system designers will verify the effectiveness of ascertain algorithm or adjust parameters of the algorithm to achieve better performance.

Asset Repository is a database to store the static server information, such as IP address, type, CPU configuration, memory setting, and topology of the servers.

III. ENERGY USAGE IN DATA CENTERS

With the growth of cloud computing, large scale data centers have become common in the computing industry, and there has been a significant increase in energy consumption at these data centers, which thus becomes a key issue to address. As most of the time a data center remains underutilized, a significant amount of energy can be conserved by migrating virtual machines (VM) running on underutilized machines to other machines and hibernating such underutilized machines. As the data center industry grows increasingly obsessed with energy efficiency, cloud computing presents a compelling opportunity to reduce data center power bills, according to a leading expert on IT power issues. Energy use is a central issue for data centers. Power draw for data centers ranges from a few kW for a rack of servers in a closet to several tens of MW for large facilities. Some facilities have power densities more than 100 times that of a typical office building.

A. Energy Efficiency:
The most commonly used metric to determine the energy efficiency of a data center is power usage effectiveness (PUE). This simple ratio is the total power entering the data center divided by the power used by the IT equipment.

$$PUE = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

Power used by support equipment, often referred to as overhead load, mainly consists of cooling systems, power delivery, and other facility infrastructure like lighting.
Four areas where cloud computing have power efficiency advantages:

1. Diversity: Spreading computing loads across many users and time zones can improve hardware utilization.
2. Economies of Scale: Computation is cheaper in a large shop than small shop, as fixed costs can be spread over more servers and users.
3. Flexibility: The management of virtual servers in cloud apps is easier and cheaper than managing physical servers. It also has reliability advantage that can create savings in the data center. If you can void outages using software to route around problems, you don't need to buy two power supplies for each server.
4. Enabling Structural Change: The shift to a cloud model enables broader efficiencies in a business that can save money over time.

The concepts inspired by the notion of utility computing have recently combined with the requirements and standards and is defined as, "A large-scale distributed computing paradigm that is driven by economies of scale, in which a pool of abstracted, virtualized, dynamically-scalable, managed computing power, storage, platforms, and services are delivered. On demand to external customers over the Internet As new distributed computing technologies like Clouds become increasingly popular, the dependence on power also increases. Therefore, it is Imperative to enhance the efficiency and potential sustainability of large data centers. One of the fundamental aspects of virtualization technologies employed in Cloud environments is resource consolidation and management.
Using hypervisors within a cluster environment allows for a number of standalone physical machines to be consolidated to a virtualized environment, thereby requiring less physical resources than ever before. While this improves the situation, it often is inadequate. Large Cloud deployments require thousands of physical machines and megawatts of power. Therefore, there is a need to create an efficient Cloud computing system that utilizes the strengths of the Cloud while minimizing its energy footprint. The framework provided in this paper represents many promising ways to reduce power consumption; true sustainable development also depends on finding a renewable and reliable energy source for the data center itself. When combined, many of today's limits in the size of data centers will begin to deteriorate.

IV. VIRTUALIZATION

There are a number of underlying technologies, services, and infrastructure-level configurations that make Cloud computing possible. One of the most important technologies is the use of virtualization. Virtualization is a way to abstract the hardware and system resources from a operating system. This is typically performed within a Cloud environment across a large set of servers using a Hypervisor or Virtual Machine Monitor (VMM) which lies in between the hardware and the Operating System (OS). From here,
One or more virtualized OSs can be started concurrently. It is the Cloud's job to exploit this capability to its maximum potential while still maintaining a given QoS. All of the clouds leverage the power of virtualization to create an enhanced data center.

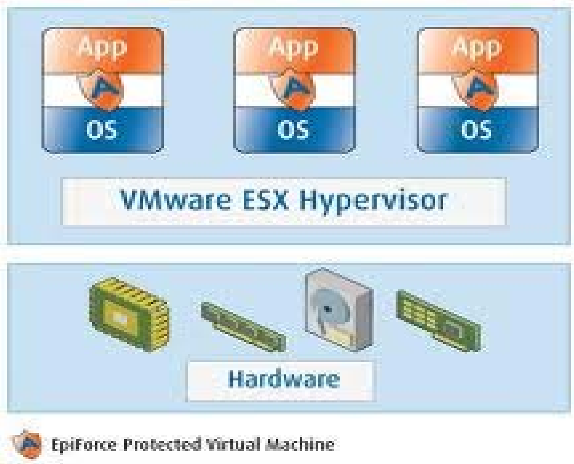

Fig 3. Virtual Machine Abstraction

The past few years has seen an increase in research on developing efficient large computational resources. Supercomputer performance has doubled more than 3000 times in the past 15 to 20 years, the performance per watt has increased 300 fold and performance per square foot has only doubled 65 times in the same period of time. This extended period of time in computing history has created the need for more efficient management and consolidation of data centers.

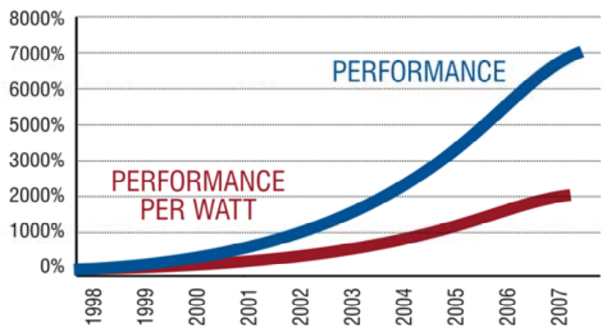

Fig 4. Performance increases much faster than performance per watt of energy consumed.

One technique being explored is the use of Dynamic Voltage and Frequency Scaling (DVFS) within Clusters and Supercomputers. By using DVFS one can lower the operating frequency and voltage, which results in decreased power consumption of a given computing resource considerably. A power-aware Cluster supports multiple power and performance modes on processors with frequencies that can be turned up or down. This allows for the creation of an efficient scheduling system that minimizes power consumption of a system while attempting to maximize performance. The scheduler performs the energy-performance trade-off within a cluster. Combining various power efficiency techniques for data centers with the advanced feature set of Clouds could yield drastic results, however currently no such system exists.

V. NEED OF GREEN CLOUD FRAMEWORK

There is a pressing need for an efficient yet scalable Cloud computing system. This is driven by the ever-increasing demand for greater computational power countered by the continual rise in use expenditures, both economic and environmental. Both business and institutions will be required to meet these needs in a rapidly changing environment. We present a novel Green computing framework that is applied to the Cloud in order to meet the goal of reducing power consumption. This framework is meant to define efficient computing resource management and Green computing technologies can be adapted and applied to Cloud systems.

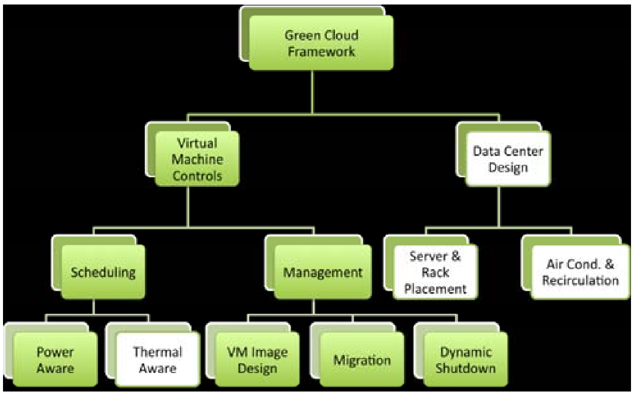

Fig 5. Green Cloud Framework.

Green Cloud framework maximizes performance per watt within a Cloud. This outlines the major areas are VM scheduling, VM image management, and advanced data center design. Within the framework, it expand functioning of virtual machines in a cloud environment to bring more efficient scheduling system for VMs. The Scheduling section addresses the placement of VMs within the Cloud infrastructure while minimizing the operating costs of the Cloud itself. This achieved by optimizing either power of the server equipment itself or the overall temperature within the data center. The design of the virtual machine images can also lead to a drastic power savings. Using more efficient Air Conditioning units, employing exterior ”free” cooling, using completely separated hot and cold isles, or simply picking more efficient power supplies for the servers can lead to incremental but substantial improvements. It provides a sustainable development platform which has the largest potential impact factor to drastically reduce power requirements within a Cloud data center.

VI. VIRTUAL MACHINE SCHEDULING & MANAGEMENT

In many service oriented scientific Cloud architectures, new VMs are created to perform some work. The idea is similar to sand boxing work within a specialized environment.

A. Power-aware VM Scheduling:

Currently, there are two competing types of Green scheduling systems for Supercomputers; power-aware and thermal aware scheduling. In thermal-aware scheduling, jobs are scheduled in a manner that minimizes the overall data center temperature. The goal is not always to conserve the energy used to the servers, but instead to reduce the energy needed to operate the data center cooling systems. In power-aware scheduling, jobs are scheduled to nodes in such a way to minimize the server’s total power. The largest operating cost incurred in a Cloud data center is in operating the servers. VM scheduling algorithm that minimizes power consumption within the data center.

```
Algorithm - Power based scheduling of VMs
FOR i = 1 TO i _ jpoolj DO
        pei = num cores in pooli
END FOR
WHILE (true)
        FOR i = 1 TO i _ jqueuej DO
                vm = queuei
        FOR j = 1 TO j _ jpoolj DO
                IF pej _ 1 THEN
                IF check capacity vm on pej THEN
                schedule vm on pej
        pej □ 1
                END IF
                END IF
        END FOR
        END FOR
wait for interval t
END WHILE
```

B. VM Management

Another key aspect of a Green Cloud framework is virtual machine image management. By using virtualization technologies within the Cloud, a number of new techniques become possible. Idle physical machines in a Cloud can be dynamically shutdown and restarted to conserve energy during low load situations. This concept of shutting down unused machines will have no effect on power consumption during peak load as all machines will be running. However in practice Clouds almost never run at full capacity as this could result in a degradation of the QoS. Therefore by design, fast dynamic shutdown and start-up of physical machines could have a drastic impact on power consumption, depending on the load of the Cloud at any given point in time.

The use of live migration features within Cloud systems is a recent concept. Live migration is presently used for proactive fault tolerance by seamlessly moving VMs away from failing hardware to stable hardware without the user noticing a change in a virtualized environment. Live migration can be applied to Green computing in order to migrate away machines. VMs can be shifted from low load to medium load servers when needed. Low load servers are subsequently shutdown when all VMs have migrated away, Thus conserving the energy required to run the low load idle servers.

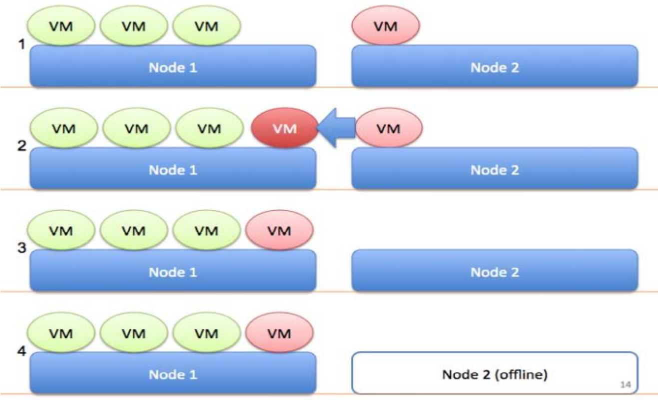

Fig 6. Virtual Machine management dynamic shutdown technique

This process of dynamically allocating and de- allocating physical machines is complementary to our scheduling system. When load increases, we use Wake on LAN (WOL) to start them back up. This control can be easily monitored and implemented as a daemon running on the Cloud head node or scheduler. This effectively displays the goal of the Green Cloud Framework: while any single power saving technique can be beneficial, the calculated combination of multiple techniques from a systems-level perspective can yield significant power savings when compared to their individual implementations.

C. Service Oriented Virtual Machine Image

While scheduling and management of virtual machines within a private Cloud environment is important, one must realize what is actually being scheduled. These VM instances contain much more than they need to in order to support a wide variety of hardware software and varying user tasks. While this is ideal for a desktop based environment, it leads to wasted time and energy in a server based solution. A hypervisor provides the same virtualized hardware to each VM and each VM is typically designed for Specific task. In essence, we want the OS within the VM to act only as a light wrapper which supports a few specific but refined tasks or services, and not an entire desktop/application suite. so most modern operating systems including Linux are able to detect various hardware and load modules on the fly upon start-up. This is not an issue with a virtual machine environment since the hardware is standardized and known in advance. It is common for the boot to spend 15 seconds running mod probe to load only a single module. The modules in the system and many of the time consuming probing functions can be reduced upon boot up within a VM environment. Boot time can be further improved by creating a new order which maximizes both the CPU utilization and I/O throughput. The use of boot chart can profile where boot up system inefficiencies occur and to allow for optimization of the boot sequence

VII. CONCLUSION AND FUTURE WORK

As the prevalence of Cloud computing continues to rise, the need for power saving mechanisms within the Cloud also increases. This paper presents a Green Cloud framework for improving system efficiency in a data center. To demonstrate the potential of framework, presented new energy efficient scheduling. Though in this paper, we have found new ways to save vast amounts of energy while minimally impacting performance. Not only do the components discussed in this paper complement each other, they leave space for future work. Future opportunities could explore a scheduling system that is both power-aware and thermal-aware to maximise energy savings both from physical servers and the cooling systems used. Such a scheduler would also drive the need for better data center designs, both in server placements within racks and closed-loop cooling systems integrated into each rack. While a number of the Cloud techniques are discussed in this paper, there is a growing need for improvements in Cloud infrastructure, both in the academic and commercial sectors.

VIII. REFERENCES

[1] B. Alexander, “Web 2.0: A new wave of Innovation for Teaching and Learning?” Learning, vol. 41, no. 2, pp. 32–44, 2006.

[2] R. Buyya, C. Yeo, and S. Venugopal, “Market-oriented cloud computing: Vision, hype, and reality for delivering it services as computing utilities,”in Proceedings of the 10th IEEE International Conference on High Performance Computing and Communications (HPCC-08, IEEE CS Press, Los Alamitos, CA, USA), 2008.

[3] I. Foster, Y. Zhao, I. Raicu, and S. Lu, “Cloud Computing and Grid Computing 360-Degree Compared,” in Grid Computing Environments Workshop, 2008. GCE’08, 2008, pp. 1–10.

[4] J. Koomey, “Estimating total power consumption by servers in the US and the world,” Final report. February, vol. 15, 2007.

[5] “Report to congress on server and data center energy efficiency - public law 109-431,” Webpage, July 2007. [Online]. Available: www.energystar.gov/ia/partners/prod_development/ downloads/EPA Datacenter Report Congress Final1.pdf

[6] L. Wang, G. von Laszewski, A. Younge, X. He, M. Kunze, and J. Tao, “Cloud computing: a perspective study,” New Generation Computing, vol. WangLYAH, to appear in 2010.

[7] G. von Laszewski, F. Wang, A. Younge, X. He, Z. Guo, and M. Pierce, “Cyberaide JavaScript: A JavaScript Commodity Grid Kit,” in GCE08 at SC’08. Austin, TX: IEEE, Nov. 16 2008. [Online]. Available: http://cyberaide.googlecode.com/svn/trunk/papers/ 08-javascript/vonLaszewski-08-javascript.pdf

[8] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. L. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield, “Xen and the art of virtualization,” in Proceedings of the 19th ACM Symposium on Operating Systems Principles, New York, U. S. A., Oct. 2003, pp. 164–177.

[9] VMware, “Understanding Full Virtualization, Paravirtualization, and Hardware Assis,” VMware, Tech. Rep., 2007. [Online]. Available: http://www.vmware.com/files/pdf/VMware paravirtualization.pdf

[12] Amazon, “Elastic Compute Cloud.” [Online]. Available: http://aws. amazon.com/ec2/ D. Nurmi, R. Wolski, C. Grzegorczyk, G. Obertelli, S. Soman, L. Youseff, and D. Zagorodnov, “The Eucalyptus Open-source Cloud-computing System,” Proceedings of Cloud Computing and Its Applications, 2008.

[10] J. Fontan, T. Vazquez, L. Gonzalez, R. S. Montero, and I. M. Llorente, “OpenNEbula: The Open Source Virtual Machine Manager for Cluster Computing,” in Open Source Grid and Cluster Software Conference, San Francisco, CA, USA, May 2008.

[11] S. Adabala, V. Chadha, P. Chawla, R. Figueiredo, J. Fortes, I. Krsul, A. Matsunaga, M. Tsugawa, J. Zhang, M. Zhao, L. Zhu, and X. Zhu, “From virtualized resources to virtual computing Grids: the In-VIGO system,” Future Generation Comp. Syst., vol. 21, no. 6, pp. 896–909, 2005.

[12] J. Chase, D. Irwin, L. Grit, J. Moore, and S. Sprenkle, “Dynamic virtual clusters in a grid site manager,” in 12th IEEE International

Symposium on High Performance Distributed Computing, 2003. Proceedings, 2003, pp. 90–100.

[13] K. Keahey, I. Foster, T. Freeman, X. Zhang, and D. Galron, "Virtual Workspaces in the Grid," Lecture Notes in Computer Science, vol. 3648, pp. 421–431, 2005. [Online]. Available: http: //workspace.globus.org/papers/VW EuroPar05.pdf

[14] I. Foster and C. Kesselman, "Globus: A Metacomputing Infrastructure Toolkit," International Journal of Supercomputer Applications, vol. 11, no. 2, pp. 115–128, 1997, ftp://ftp.globus.org/pub/globus/papers/globus.pdf. W. chun Feng and K. W. Cameron, "The Green500 List: Encouraging Sustainable Supercomputing," IEEE Computer, vol. 40, no. 12, pp. 50–55, 2007.

[15] M. Wagner, "The efficiency challenge: Balancing total cost of ownership with real estate demands," Cherokee International, Tech. Rep., 2008.

[16] K. Barker, K. Davis, A. Hoisie, D. Kerbyson, M. Lang, S. Pakin, and J. Sancho, "Entering the petaflop era: the architecture and performance of Roadrunner," in Proceedings of the 2008 ACM/IEEE conference on Supercomputing. IEEE Press Piscataway, NJ, USA, 2008.

[17] J. Dongarra, H. Meuer, and E. Strohmaier, "Top 500 supercomputers," website, November 2008.

[18] S. Sharma, C.-H. Hsu, and W. chun Feng, "Making a case for a green500 list." in IEEE International Parallel and Distributed Processing Symposium (IPDPS 2006)/ Workshop on High Performance - Power Aware Computing, 2006.

[19] J. Laros III, K. Pedretti, S. Kelly, J. Vandyke, K. Ferreira, C. Vaughan, and M. Swan, "Topics on Measuring Real Power Usage on High Performance Computing Platforms," in IEEE Cluster, 2009.

[20] C.-H. Hsu and W. chun Feng, "A Feasibility Analysis of Power Awareness in Commodity-Based High-Performance Clusters," in CLUSTER, 2005, pp. 1–10