



ENHANCING ENERGY EFFICIENCY FOR CLOUD COMPUTING ENVIRONMENTS

Ms. M. Shanmugapriya
Assistant Professor,
Department of CSE,
PSNA CET
Dindigul

Ms. S. Naganandhini
Assistant Professor,
Department of CSE,
PSNA CET,
Dindigul

Ms. S. T. Senbagavalli
Assistant Professor,
Department of CSE,
PSNA CET,
Dindigul

Abstract - Cloud Computing has changed the way in which people use the IT resources today. Cloud Computing is transforming information technology. As information and processes are migrating to the cloud, it is transforming not only where computing is done, but also fundamentally, how it is done. Now, they can use the services offered by Cloud Computing with reasonable costs based on a “pay-per-use” model. Cloud Computing solves many problems of conventional computing, including handling peak loads, installing software updates, and, using excess computing cycles. However, with the wide adoption of Cloud Computing, the costs for maintaining the Cloud infrastructure have become a vital issue for the providers, especially with the large input of energy costs to underpin these resources. Thus, this paper proposes a system architecture that can be used to profile the resources usage in terms of the energy consumption. From the profiled data, the application developers can improve their energy-aware decisions when generating the applications to be more energy efficient. This paper also presents an adapted existing Cloud architecture to enable energy-aware profiling based on the proposed system. The results of the conducted experiments show energy-awareness at physical host and virtual machine levels.

Keywords: *Cloud Computing, Energy Efficiency, Energy-Aware Profiling, Energy Efficiency Metrics.*

I. INTRODUCTION

The radical adoption of Cloud Computing technology has exposed a significant overhead in maintaining its infrastructure, which has become a major issue for the Cloud providers due to the associated high operational costs, such as energy consumption. It has been stated that a data centre may consume about 100 times more energy than a typical office of the same

size [1]. So, efficiently managing the power consumed by the servers would improve the overall consumption; in the sense that as the servers consume less power, the heat generated by these servers would be reduced, which would then reduce the need for cooling resources that consume large amount of energy as well and result in more energy savings. Many researchers have investigated the new ways to manage cloud infrastructure as a means of enhancing the energy efficiency. A number of techniques have been already proposed and deployed for better resource management. For example, Data Voltage and Frequency Scaling (DVFS) and Virtual Machines (VMs) allocation have been widely studied and deployed to manage the resources more efficiently [3]. Nonetheless, there is still a need to make the whole stack of Cloud Computing more energy-aware and not only focusing on the resource management aspects. There are a large number of different Cloud applications with different requirements of resources; some of them are data-intensive whereas others are compute intensive. So, depending on the taxonomy of the application, the energy consumption of the resources that underpin these different applications can vary. The properties of Cloud applications are derived from the characteristics of Cloud Computing. Fehling et al [4] stated that Cloud applications should be able to support the characteristics of Isolate state, Distribution, Elasticity, Automated management, Loose coupling (IDEAL), all of which reflect the patterns of Cloud Computing environments. Depending on the behavior of users and submitted tasks, these applications can experience different patterns of workloads, which are depicted based on the utilization of IT resources hosting the applications. These workloads can be categorized as static workload that has equal utilization of resources over time, periodic workload that has repeating peak utilization at interval time, once-in-a-lifetime workload that has a peak utilisation once over time, unpredictable workload that has a frequent and random peak utilisation over time, and continuously



changing workload that has a utilisation increases or decrease continuously over time [4]. These different types of application workloads can have different impact of energy consumption depending on usage of the resources component. Thus, this research is aimed to add value to the Cloud Computing energy efficiency by investigating energy efficiency modelling in terms of energy-aware profiling and energy efficiency metrics. Energy-aware profiling is studied in order to understand how the energy is consumed by the infrastructure components, like CPUs, when the application is in operation. Thus, the output measurements of energy-aware profiling and energy efficiency metrics will be combined to form KPIs for the running application. Also, these KPIs will be further analysed and used to facilitate the decision-making of application developers with better energy-aware programming. The main contributions of this paper include:

- A proposed system architecture for profiling and assessing the energy efficiency of Cloud infrastructure resources.
- Implementing the proposed architecture in an existing Cloud testbed to enable energy-aware profiling.
- Introducing an energy modeller to enable energy-awareness at VM level.

II. RELATED WORK

The rest of the paper is structured as follows: various streams of improving energy efficiency in Cloud Computing will be reviewed in Section 2, and Section 3 will further review some aspects in energy efficiency modelling; Section 4 will present the proposed system architecture for improving the energy efficiency assessment of Cloud Computing infrastructures; Section 5 will discuss the implementation of the proposed architecture in an existing Cloud testbed to enable energy-aware profiling at physical host and VM levels; Section 6 will present the experimental set up and design to validate the architecture of the testbed, and finally Section 7 will conclude the paper and discuss future work.

2 Energy Efficiency in Cloud Computing For the Cloud Computing stack, energy efficiency has been extensively studied in the literature and has focused on a large number of different topics, like virtualisation, requirement engineering, programming models, and resource management. In terms of virtualisation, a number of studies proposed different approaches for allowing resource utilisation, server consolidation and live migration of virtual machines

[5,6,7], which all can offer significant energy and costs savings [1]. With the advancement of software-intensive systems for self-adaptive systems to meet the growing needs for autonomic computing [9], requirements engineering for self-adaptive software systems ensuring energy aspects has received less attention [1]; as that can be justified with the challenges to encounter when dealing with uncertainties associated with the operating environment [1]. Optimising energy efficiency at different layers of Cloud stack is considered significantly important, as argued by Djemame et al [2]. They therefore have proposed a Cloud architecture that addresses energy efficiency at all layers of the Cloud stack and throughout the whole Cloud application lifecycle. In terms of programming models, there are a number of platforms used for the development and deployment of Cloud applications and services, like Hadoop [3], Windows Azure [14], Microsoft Daytona [5], Twister [6], Manjrasoft Aneka [1], and Google App Engine [1]. Yet, these platforms lack consideration for energy efficiency, whereas a work presented in [1] proposed a general-purpose programming environment to simplify and help the developers make energy-efficient decisions for constructing energy-aware applications. Most of the attention in the literature has focused on enhancing the energy efficiency of Cloud Computing through better resource management to avoid some issues like excessive power consumption and SLAs violation reliability [3]. Therefore, many developments have been introduced like, Dynamic Power Management (DPM) and DVFS (Dynamic Voltage and Frequency Scaling) techniques to control the power consumption of servers in accordance with the workload [2], virtual machine consolidation policies to optimise the hosts by migrating VMs from one host to another [3], some models for better prediction of the power consumption for the servers [1], task consolidation model for maximizing resource utilisation [2], a holistic framework called Mistral for optimizing the power consumption for the physical hosts [3], a CPU re-allocation algorithm that combines both DVFS and live migration techniques to reduce the energy consumption and increase the performance in Cloud datacentres [4]. However, there is a lack of research that tackles the issue of properly ensuring I. Alzamil et al. / Electronic Notes in Theoretical Computer Science 318 (2015) 91–108 93 energy-awareness from the design stage and not only through resource management of the Cloud Infrastructure. So, there is still a need for modelling the energy efficiency of Cloud infrastructures to gain a better understanding of energy efficiency and to



feed the decision-making at the service design stage, which will be discussed in the following section.

III. ENERGY EFFICIENCY MODELLING

It is important to model energy profiling techniques and introduce new metrics to inform the providers how energy efficient their infrastructure is to make strategic decisions, such as creating and configuring energy-aware application and forming new energy-aware pricing mechanism, accordingly.

3.1 Profiling

Having such tools that would help understand how the energy has been consumed in a system is essential in order to facilitate software developers to make energy-aware programming decisions. Schubert et al [5] state that the developers lack the tools that indicate where the energy-hungry sections are located in their code and help them better optimize their code for enhancing energy consumption more accurately instead of just relying on their own intuitions. In their work, they proposed eprof, which is a software profiler that narrates energy consumption to code locations; therefore, it would also help developers make better energy-aware decisions when they re-write their code [5].




Fig., cloud infrastructure

For example, with storing data on a disk, software developers might choose between storing the data in an uncompressed format or a compressed format, which would require more CPU resources. Compressed data has been commonly suggested as a way to reduce the amount of I/O needed to be performed and therefore reducing the energy based on the hypothesis that the CPU can process the task of compression and decompression with less energy than the task of transferring large data from and to the disk [6]. However, that would depend on the data being processed. In fact, some conducted experiments in [5] with eprof profiling tool show that the process of compressing and decompressing the

data consumes significantly more energy than the process of transferring large amount of uncompressed data because the former would use more CPU resources than the latter. So, it can be a controversial issue depending on the application domain. Thus, having such tools identifying where the energy has been consumed would help software developers to make more energy-aware decisions. Moreover, a new framework called Symbolic Execution and Energy Profiles (SEEP) has been introduced in [2] as an approach to help software developers make well informed decisions for energy optimization from early stages at the code level. To illustrate, SEEP is designed to provide the developers with energy estimations to make them more energy-aware while they are programming. 94 I. Alzamil et al. / Electronic Notes in Theoretical Computer Science 318 (2015) 91–108 3.2 Metrics Energy efficiency in Clouds can be assessed by different metrics. In terms of Cloud infrastructure, the well-known Power Usage Effectiveness (PUE) metric has been introduced by the Green Grid organization to help the providers assess and improve the energy efficiency of their data centres [28]. However, despite the fact that the PUE metric has been successful and widely used, Grosskop [29] argues that it is restricted as an indicator for energy efficiency to the infrastructure management only and not considering the optimization at the software levels to enhance the efficiency of the whole stack. Also, Bozzelli et al [3] have reviewed a number of software metrics and emphasised the importance to assess the energy efficiency not only from the hardware side but also from early stages of the software lifecycle in order to make such energy savings. Additionally, as stated by Wilke et al [31], analysing software's energy consumption is considered an important requirement for such optimizations. So, Grosskop proposed a new metric called the Consumption Near Sweet-Spot (CNS) that identifies how well the system's energy efficiency optimum and its utilisation are aligned by calculating the ratio between the average consumption and optimum consumption for a system to deliver a particular unit of work [2]. Moreover, other works have looked at other metrics for energy efficiency measurements, like utilisation percentage and SLA violation percentage. For example, in the work conducted by Beloglazov et al [3], they evaluate the efficiency and performance of their proposed algorithms by using some metrics, namely the total energy consumption, the average number of SLA violations, and the number of VM migrations. Recently, some works have started to measure the energy consumption in more detail, like measuring



energy consumption for each VM in a physical machine. Research conducted in [3] introduces a VM power model to measure the estimated power consumption of VM with using performance events counter. They argue that the results of their proposed model can get on average about 97% accuracy. Nonetheless, as mentioned earlier, there is a limited number of metrics to measure the energy efficiency of Clouds from different layers other than the infrastructure only. In terms of fine-grain measurement, there is a need to map the energy consumption for each single VM in a server, which indicates the importance to fill this gap by introducing new suitable metrics for measuring and mapping the energy consumption to each VM.

IV. ENERGY-AWARE PROFILING

Ensuring energy efficiency from different layers in Cloud Computing has become inevitable, especially with the increased energy costs. We propose in this paper to have energy-aware profiling for the Cloud infrastructure to better understand how the energy has been consumed and assess its energy efficiency in order to help the software developers from the application layer enhance their decision-making in terms of energy-awareness when optimizing their applications and services.

4.1 Proposed System Architecture

The scope of this proposed system architecture would be in the IaaS layer where the operation of services takes place. The main components of this model consist of Resource Monitoring Unit (RMU), Energy Unit (EPU), Reporting and Analysis Unit, as can be shown in Figure. 1. Fig. 1. Proposed System Architecture This proposed system architecture would have the RMU to dynamically collect the energy consumed by the hardware components and observe the number of assigned VMs. After that, EPU would have appropriate algorithms to calculate the energy consumed by each VM and hardware components, and it would then profile and populate these measurements as KPIs to a database. This data can be further analysed by the Reporting and Analysis Unit to provide the software developers energy-aware reports in order to enhance their awareness of the energy consumption when making programming decisions. For example, it might be interesting to know whether the CPU or the memory of the hardware component would consume more energy, so that the developer can create applications that would depend more on components with less

energy consumption, without compromising performance.

4.2 Illustration

An application can run on the Cloud to deliver services for the end users.



Data-intensive tasks would depend more on using disk storage for processing large amounts of data and data retrieval or update, which would require high disk I/O bandwidth to maintain performance, whereas computation-intensive tasks would depend more on using the processors to perform more computation [3].

When the service is configured in the application development tool with descriptions of the allocated software and hardware resources, and is deployed in the service deployment environment and goes through VM management, the proposed system would then start with the RMU to capture and monitor the energy consumed by the infrastructure that underpins and operates that service. The captured data (as input to the system) will be dynamically collected by the EPU for appropriate measurements and profiling in terms of energy efficiency. Next, EPU would populate the profiled data as KPIs to a database. Hence, these KPIs (as output of the system) can be further analysed and reported in a meaningful format to the application developers to enhance their energy-aware decisions when making and configuring new services. The next section will provide the implementation of this proposed system architecture into an existing Cloud testbed to enable energy-aware profiling.



V. IMPLEMENTED ENERGY-AWARE CLOUD ARCHITECTURE

In this section, we discuss how an existing Cloud architecture, Leeds Testbed, has been adapted to support energy-awareness at physical host and VM levels.

5.1 Leeds Testbed

The software architecture of the Cloud testbed is illustrated in Figure 2. Each node at the time of writing runs Centos version 6.6 for its operating system. The XEN [34] hypervisor version 4.0.1 is also deployed along side the Linux Kernel version 2.6.32.24 as the Virtual Machine Manager (VMM). Version 3.8 of OpenNebula [35] is used as the Virtual Infrastructure Manager (VIM). In addition, Virtual machines instances are configured to leverage Hardware Assisted Virtualisation (HVM) and the QEMU [6] device module for enhanced performance and interoperability.

From a hardware perspective, the Leeds testbed is comprised of a cluster of Dell commodity servers. For the purpose of this research, four of these were used with energy meters. Each server consists of a four core X3430 Intel Xeon CPU, running at the default clock speed of 2.40GHz and a total of 8GB of RAM (four modules of 2GB DDR3 at 1333Mhz). Additionally, each server utilised a single 3.5 inch Western Digital RE3 250GB SATA HDD (Model: WD2502ABYS), with 16MB of cache and a spindle speed of 7200 RPM. The machines connect via Gigabit Ethernet using a Broadcom PCI-E NIC (Model: BCM95722). This connectivity provides shared access to a NFS share running on the cluster headnode. The NFS share is backed by four 500GB HDDs running in Raid 0, providing a total of 2TB storage for VM images.

VI. EXPERIMENT SET UP AND DESIGN

In this section, we present some experiments that have been conducted on Leeds testbed. The overall aim of the experiments conducted is to validate that the testbed has been setup correctly as a Cloud environment that supports energy aware profiling both physical host and VM levels. In order to design such experiments, a software testing tool that represent real patterns of Cloud applications is needed. Cloud9, a software testing benchmark, has therefore been setup on the testbed to generate real scale-out workloads. The generated workloads by

Cloud9 reflect real Cloud applications patterns [39]. Cloud9 is capable of scheduling a task or set of tasks to run on one or multiple VMs, and these tasks can be configured to run in parallel or in stages after each other [40] to represent real pattern of elastic Cloud application. The following experiments are designed differently to show various aspects of Cloud Computing patterns. Each one has been repeated 10 times to get the average mean value of the power consumption and eliminate any anomalies of the results.

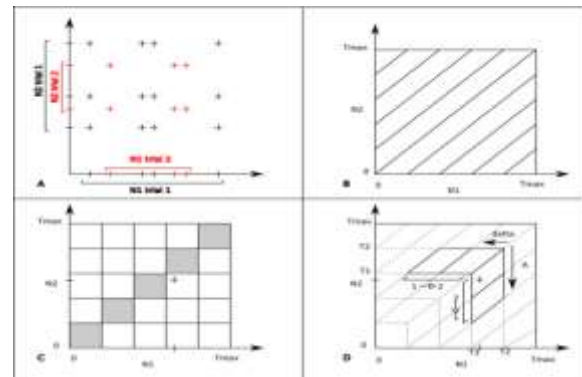


Fig 2. Experiment result for VM

VII. CONCLUSIONS AND FUTURE WORK

This paper has proposed system architecture for Cloud Computing environment. This system architecture can help software developers understand how their applications are using the infrastructure resources and consuming energy in order I. Alzamil et al. / Electronic Notes in Theoretical Computer Science 318 (2015) 91–108 105 to enhance their decision-making when creating and configuring new applications. We have adapted an existing Cloud architecture based on the proposed system to enable energy-aware profiling and presented the energy modeller that allows energy awareness at VM level. The conducted experiments showed that Leeds testbed has been setup correctly as a Cloud environment supporting energy-awareness at both physical host and VM levels. Future work will include investigation on energy efficiency modelling to identify new metrics and form KPIs to better understand to what extent a running application is energy efficient in relation to these KPIs. Implementing the Analysis and Reporting Unit of the proposed system to provide a meaningful feedback of these KPIs to the application developers to enhance their programming decisions with energy-awareness. Finally, when these KPIs are identified, further



research would investigate how to identify new energy-aware pricing mechanisms to charge the users for the offered Cloud services based on these measurements. So, the end-users are being charged more precisely based on their actual resource usage of Cloud services considering the energy consumption as well, which would contribute efficiently to the overall business model of Cloud Computing.

VIII. REFERENCES

- [1] P. Scheihing. Creating Energy-Efficient Data Centers. In *Data Center Facilities and Engineering Conference*, Washington, DC, 18th May., 2007.
- [2] A. Berl, E. Gelenbe, M. Di Girolamo, G. Giuliani, H. De Meer, M. Q. Dang, and K. Pentikousis. Energy-Efficient Cloud Computing. *The Computer Journal*, 53(7):1045–1051, August 2009.
- [3] A. Beloglazov, J. Abawajy, and R. Buyya. Energy-aware resource allocation heuristics for efficient management of data centers for Cloud computing. *Future Generation Computer Systems*, 28(5):755–768, May 2012.
- [4] C. Fehling, F. Leymann, R. Retter, W. Schupeck, and P. Arbitter. *Cloud Computing Patterns: Fundamentals to Design, Build, and Manage Cloud Applications*. Springer, 2014.
- [5] Bertion, E., Paci, F., & Ferrini, R. (2009). Privacy-Preserving Digital Identity Management for Cloud Computing. *IEEE Computer Society Data Engineering Bulletin*, pp. 1-4, March 2009.
- [6] K. Ye, D. Huang, X. Jiang, H. Chen, and S. Wu. Virtual Machine Based Energy-Efficient Data Center Architecture for Cloud Computing: A Performance Perspective. In *Green Computing and Communications (GreenCom), 2010 IEEE/ACM Int'l Conference on & Int'l Conference on Cyber, Physical and Social Computing (CPSCoM)*, pages 171–178, 2010.
- [7] Y. Dodis, S. Vadhan, and D. Wichs, “Proofs of retrievability via hardness amplification,” in *Proc. of the 6th Theory of Cryptography Conference (TCC'09)*, San Francisco, CA, USA, March 2009.
- [8] Winkler, V.J.: *Securing the Cloud: Cloud computer Security techniques and tactics*. Elsevier, 2011.
- [9] G. Ateniese, R. D. Pietro, L. V. Mancini, and G. Tsudik, “Scalable and efficient provable data possession,” in *Proc. of SecureComm'08*, 2008, pp. 1–10.
- [10] Popovic K; Hocenski Z; (2010), “Cloud computing security issues and challenge”, 5533317searchabstractMIPRO, 2010 Proceedings of the 33rd International Convention , pp 344,24-28 May 2010.
- [11] C. Erway, A. Kupcu, C. Papamanthou, and R. Tamassia, “Dynamic provable data possession,” in *Proc. of CCS'09*, 2009, pp. 213–222.
- [12] J. Hardy, L. Liu, N. Antonopoulos, W. Liu, L. Cui, and J. Li. Assessment and Evaluation of Internet-Based Virtual Computing Infrastructure. In *Object/Component/Service-Oriented Real-Time Distributed Computing (ISORC), 2012 IEEE 15th International Symposium on*, pages 39–46, 2012.