Introducing a performance measurement tool for managing the environmental performance of cloud computing

K. (Koen) van Schijndel
Delft University of Technology
Faculty of Technology, Policy and Management
Jaffalaan 5
2628 BX Delft, The Netherlands
+31(0)614771300
koenvanschijndel@gmail.com

30th of April, 2015

Abstract
Despite the cloud's potential for its customers to realize significant energy savings and cost reductions compared to their on-premises infrastructure, the exploding demand for cloud computing services drives enormous energy consumption and CO_2-emissions in cloud computing data centers. This intensification of cloud data centers challenges the potential of cloud computing and specifically puts the overall environmental performance of cloud computing services under pressure. To be able to manage and control the environmental performance of the cloud, the deliverable of this research is a design for a performance management tool, which is based on the Deming-wheel for quality management and a set of design requirements derived from theory and practice. This way, the performance measurement tool provides a structured and comprehensive approach for dealing with the environmental performance of cloud computing. Future research into using the measurement tool for calculating the environmental performance of other parts of the cloud computing life-cycle is needed.

Keywords: cloud computing, energy efficiency, data center, performance management

1. Introduction

Cloud computing provides major benefits to businesses. It eliminates the problem of both over- and under provisioning of IT resources since it promises companies 'infinite' capacity and instantaneous scalability (Armbrust et al., 2010). The elimination of excessive capacity, results in significant reductions of energy consumption and costs for the cloud service provider. Moreover, the customer only cloud computing has the ability to *"convert capital expenses to operational expenses"* (Armbrust et al., 2010, p. 53), because it does not require large up-front investments and only charges the user when it uses the services (Grossman, 2009).

At the supply side however, *"the growing demand of cloud infrastructure has drastically increased the energy consumption of data centers, which has become a critical issue"* (Garg & Buyya, 2011, p. 1). Cloud services are provisioned in large complex data centers, that function as the backbone of cloud computing (Elgelany & Nada, 2013). As these data centers comprise of a large amount of servers, a data center can use as much energy as a small city (EPA, 2009). Data centers currently account for up to 3% of the global amount of energy produced and the generation of 200 million metric tons of CO_2 (Lavallée, 2014).

Given the idea that the cloud computing market will continue to grow, the energy consumption of cloud data centers will further increase leading to a negative effect on the environmental performance of cloud computing. Currently, there are no tools available that support cloud service providers in managing environmental performance, which results in a lack of control. Therefore, the

Design Oriented Approach of Verschuren and Hartog (2005) was used to develop a prototype for a performance measurement tool. This approach was completely followed and described in the graduation work of van Schijndel, which is titled 'Designing a performance measurement tool for the environmental performance of cloud computing'. In this paper, the focus is on the presentation of the design that has been developed for the measurement tool.

The structure of this paper is as follows. First, a problem description is presented followed by the scope to which the design presented in this paper applies. Next, a brief analysis of the current measurement effort at cloud service providers is presented, followed by an overview of the structural specifications. These structural specifications form the basis of the design for the measurement tool which is presented next. This paper concludes with conclusions, discussion and some recommendations.

2. Problem description

Uddin and Rahman (2012, p. 4078) state that *"the increasing demand for storage, networking and computation has driven intensification of large complex data centers that run many of today's Internet, financial, commercial and business applications"*. The pressure on cloud data centers is expected to become even bigger, as the cloud computing market continues to grow rapidly. A survey performed by 2nd Watch indicates that companies will spend at least 15% more on public cloud infrastructure in 2015, based on the observed trend of companies moving their IT from privately owned data centers to public cloud (2nd Watch, 2014). The Global Cloud Index of Cisco (2014) predicts that in 2018, 78 percent of all workloads will be processed by cloud data centers. Moreover, a report of Market Research Media (2014) forecasts that the cloud computing market will grow at an compound annual growth rate (CAGR) of 30% to comprise of 270 billion dollars in the year 2020. While the cloud computing market keeps growing, managing and controlling the environmental performance of cloud computing is an ongoing challenge, because energy consumption and CO_2-emissions will further increase. The lack of a performance management tool to support the cloud service provider in this challenge, makes it even more difficult. As the old saying goes: *"you cannot manage, what you cannot measure"* (Uddin & Rahman, 2012, p. 4083). The lack of control over the data center's CO_2-emissions is reflected through uncertainty about the effectiveness and return-on-investment of interventions that have been done to improve the energy efficiency of the data center (Alger, 2009). This makes it difficult to set benchmarks for future improvements (Jenkin, Webster, & McShane, 2011). The inability to manage CO_2-emissions is not the only negative consequence. Energy consumption accounts for the highest operational costs of a data center (Alger, 2009). To prevent the profit margins of cloud service providers from being significantly reduced, interventions to increase energy efficiency should be implemented (Buyya, Beloglazov, & Abawajy, 2010), which again invokes the problems with measuring the effectiveness of the interventions.

So, to give cloud service providers control over CO_2-emissions and energy costs caused by their cloud services, they must be able to measure and evaluate interventions that are done for improvement. Therefore, the objective of this paper is to present the design and the added value of a performance measurement tool that supports the cloud service provider in managing environmental performance in a structured way.

3. Scoping

Chou and Chou (2012) define four complementary paths in the environmental impacts of IT: green design, green manufacturing, green use, and green disposal. As the energy consumption of the cloud has been identified as a major concern, this research shall focus on the path 'green use', which reflects the larger part of the cloud's energy consumption. More specifically, For cloud computing, the focus is on the IaaS-layer, as it *"consumes a huge part of total energy in a cloud computing system"* (Jing, Ali, She, & Zhong, 2013, p.445). Environmental performance can be measured and expressed in multiple ways. In this paper, environmental performance is referred to as environmental impact, which can be defined as: *"the degree to which an organization's business processes, activities and* *operations positively or negatively affect the natural environment"* (Jenkin et al., 2011, p. 19). Evidently, the indicator for the environmental impact of cloud computing energy consumption is the carbon footprint. Bringing this all together yields that the scope of this research is the environmental impact caused by the energy consumption of IT-resources needed for provisioning cloud services. This scoping implies the following elements not to be included in the scope of this research:

- The energy consumed by cooling facilities and other auxiliary equipment.
- Energy consumption of network resources outside the data center.
- The *materials* used to build the data center, IT infrastructure and manufacture IT hardware.
- *Waste* that is produced by the data center (such as cooling water).

4. Analysis of current measurement effort

To analyze the current measurement effort for measuring the environmental performance of cloud computing a case study including six cloud service providers was done. The cloud service providers included in the case study, are of different size (in terms of total data center floor surface) and geographical orientation (national versus international, based on data center locations). Using semi-structured interviews and analysis of open source information, the current measurement efforts for measuring environmental performance were analyzed, to have an understanding of current best practices. Having performed this analyses, the most important findings are the following.

None of the six cloud service providers has a process in place to measure the environmental performance of their services, but they all do measure energy consumption. Energy consumption is also an indicator of environmental performance, but a rather indirect one. Nonetheless, these energy consumption measurements are needed to determine the environmental impact in terms of CO_2-emissions. All cloud service providers included in the case study mention to use sensors at rack or Power Distribution Unit level for measuring energy consumption. Since cloud computing resources are actually *virtual* resources, measuring energy consumption only at physical locations is not enough. But, one cloud service provider uses an integrated cloud platform which is able to measure energy consumption and allocate this energy consumption to the virtualization level. Another cloud service provider mentioned to switch to such a platform soon. It is therefore assumed that cloud service providers will be ready in the future for performing the needed energy consumption measurements.

5. Structural specifications

In the graduation research of van Schijndel, requirements were derived by means of a case study including six cloud service providers and a literature study on relevant theoretical concepts for measuring environmental performance of cloud computing. Table 1 presents structural specifications that are based on the requirements which are specified or translated into characteristics, aspects or elements that the artefact should comprise of (Verschuren & Hartog, 2005). For detailed information about the establishment of the requirements and structural specifications, the graduation thesis of van Schijndel should be consulted. For the purpose of this article, the use of the structural specifications of Table 1 for developing the design of the measurement tool is briefly discussed.

First of all, the requirement 'could follow the 'Plan-Do-Check-Act cycle of Deming' is dominant in determining the structure of the measurement tool. Although using the Deming-circle is optional, it seems a very relevant option, since it is commonly accepted and used for quality management and problem solving in organizations. Moreover, according to Basu (2004), this cycle represents 'continuous improvement', which was one of the requirements (16). More specifically, It is proposed to use a variant of the Deming-Wheel: the Observe-Plan-Do-Check-Act-cycle, because observation of the current performance is also necessary to determine whether action is needed or not. Next step is to determine the process steps to be taken within each phase. To do so, the functional requirements can be used as they specify what the measurement tool should *do*. The functional requirements indicate that the measurement tool should include the following steps (1) measure energy

Table 1: structural specifications

Type	Requirement	
Functional [R_f]	The measurement tool should follow logical performance measurement process steps:	
	▪ Should measure the energy consumption of a cloud computing service.	1
	▪ Should translate the energy consumption into useful impacts	2
	○ Should measure the environmental performance of cloud computing as the carbon footprint caused by energy consumption.	3
	○ Should translate energy consumption into economic performance.	4
	○ Should measure the environmental performance of cloud computing quantitatively.	5
	- Should use a quantitative metric for calculating environmental and economic impact.	6
	▪ Should interpret the meaning of the resulting impacts.	7
Non-functional [R_{nf}]	▪ Should have the right level of granularity. ○ Must measure energy consumption on the virtualization level (constraint).	8
	▪ Should collect data on regular time intervals.	10
	▪ Should be independent of the time of measurements.	11
	▪ Could use a benchmark or baseline value.	12
	▪ Should use standardized benchmark values.	13
	▪ Could follow the Plan-Do-Check-Act-Cycle of Deming (1982).	15
	▪ Should be aimed at continuous improvement.	16
	▪ Should look for greener solutions and continue the greening process.	17
User [R_u]	▪ Should provide comparable results.	18
	○ Should have objective and measurable standards.	19
	○ Should enable environmental friendly purchasing.	20
	- Should facilitate vendor selection based on environmental performance.	21
	○ Should contribute to the marketing of cloud services for cloud service providers.	22
	▪ Should be understandable.	23
	○ Should be self-explanatory.	24
	○ Should not require a lot of knowledge.	25
	○ Should guide the user in using the instrument.	26
	▪ Should be transparent.	27
	○ Should enable tracing back outputs to their original inputs.	28
	○ Should facilitate the transparent reporting of outputs.	29
	▪ Should appeal to customers.	30

consumption, (2) translate energy consumption into environmental and economic impact and (3) interpret these impacts. These functional requirements belong to the 'Do' and 'Check' phase. The other steps in the phases

Table 2: phases of the cycle and steps

Phase	Steps	#
Observe	Observe current performance	1
Plan	Set targets	2
	Determine how to reach targets	3
Do	Execute plan	4
	Measure energy consumption	5
Check	Translate energy consumption into impacts	6
	Interpretation of impacts	7
Act	Go back to 1	n/a
	Go back to 3	n/a

of the cycle are determined on the basis of the purpose of these phases, which yields the overview in Table 2.

The non-functional and user requirements, have been covered in the design of the measurement tool as a certain design construct or as a characteristic of a design construct depending on the nature of the requirement. For example, requirement 12 implies the use the *design construct* 'standardized benchmark value' in the model, while requirement 11 specifies a *characteristic* of the measurement tool. Table 1 also contains requirements that do not directly apply to a certain element of the measurement tool, but rather apply to the measurement tool as a whole. These include the requirements on understandability and transparency (req. 23-29) and

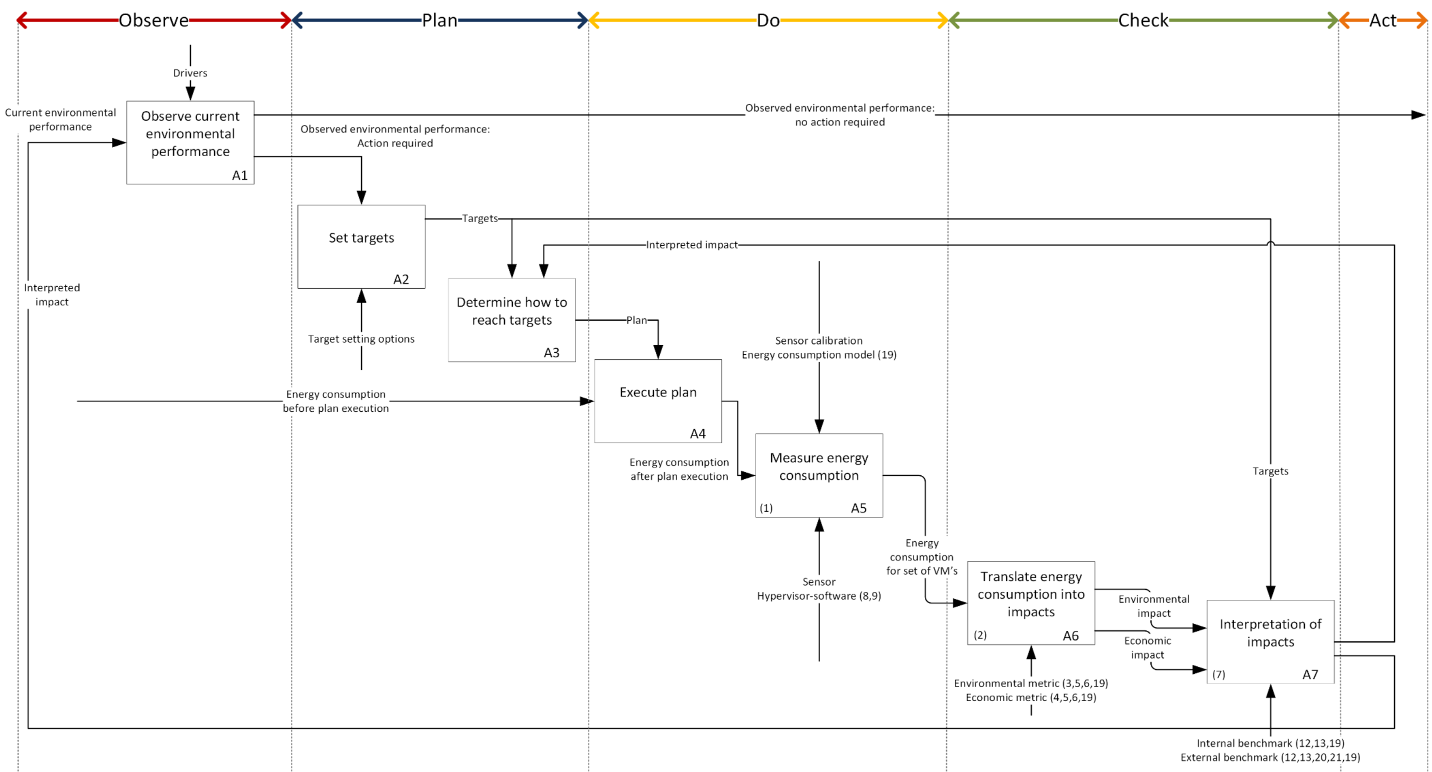

Figure 1: visualized structure of the measurement tool using IDEF0-modeling

requirement 30 which implies that the measurement tool should appeal to the cloud customer. As explained, requirements 15 and 16 that imply the use of the Deming-wheel and a focus on continuous improvement are embedded in the overall structure of measurement tool. Using the structural specifications as described for creating the design of the measurement tool, the result can be modeled using the IDEF0-modeling technique. Doing so, yields the structure of the measurement tool as presented in Figure 1. The phases Observe, Plan, Do, Check and Act can be clearly recognized, just as the seven steps presented in Table 2. Following the IDEF0-semantics, each step presented as a 'box' in Figure 1 (e.g. A1), can contain input (incoming arrow), output (outgoing arrow), controls (top arrow) and mechanisms (bottom arrow). Each arrow is described using a label that optionally also includes references to requirement numbers. For example, the bottom arrow in step A6, contains the labels 'environmental metric' and 'economic metric', followed by the numbers of the requirements that imply the use or characteristics of these metrics. The structure presented in Figure 1 is still generic; it is for example unclear *how* energy consumption should be measured or *what* metrics exactly should be included. Therefore, this structure was translated into an actual design of the measurement tool. This is done through identifying those elements of Figure 1 that need to be replace by 'real-life' design constructs. The details of creating this design can be found in the graduation thesis of van Schijndel. For now, the final design is presented and described in the next paragraph.

6. Final design

The final design of the measurement tool is presented in Figure 2. One may notice that the final design consist of only six steps, while the structure presented in Figure 1 contains seven. This is the case because step A3 and step A4 have been merged into one step (step 3 in Figure 2).

6.1 Step 1: evaluate current carbon footprint

This step uses the output of the previous iteration (T_7) as input. The user should evaluate the current carbon footprint with the help of the drivers: costs, greening IT, customer, government and organization. If these drivers

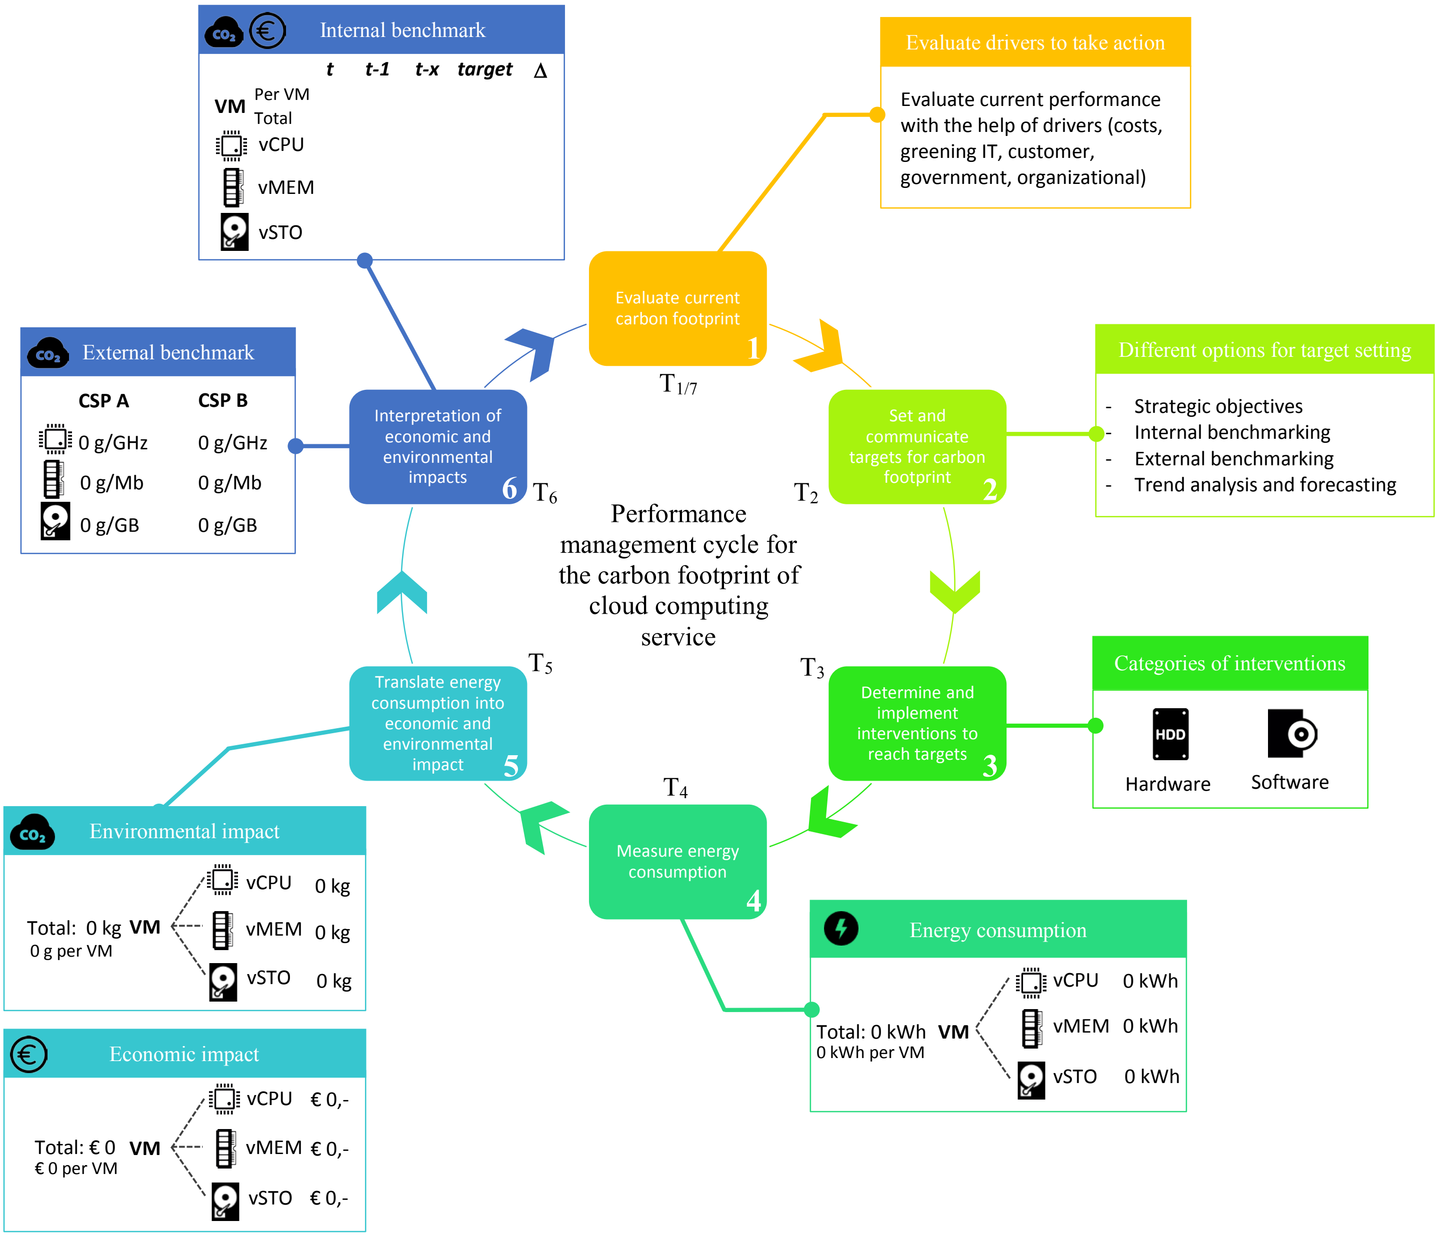

Figure 2: final design of the performance measurement tool

motivate the user to improve the carbon footprint, he or she should proceed to step 2. If not, the carbon footprint should be evaluated again at a later moment in time.

6.2 Step 2: different options for target setting

The user should set targets for carbon footprint. To do so, strategic objectives, internal benchmarking, external benchmarking and trend analysis and forecasting can be used.

6.3 Step 3: determine and implement interventions

In this step, the user should develop and implement a plan of interventions to increase the energy efficiency of the IT resources. The measurement tool presents two categories in which interventions can be chosen: hardware and software.

6.4 Step 4: measure energy consumption

The user should retrieve the needed energy consumption data from the hypervisor-software. To do so, the user needs to provide two inputs: (1) the time interval over which energy consumption data needs to be gathered (e.g. February or 2015) and (2) the set of virtual machines of which the energy consumption data is needed. This could be a set that is used for the provisioning of a specific service, a set that is used for provisioning services for a specific customer or all the virtual machines used by the cloud service provider. The outputted energy consumption data provided by the hypervisor-software should be presented as displayed in the connected rectangle to step 4 in Figure 2 which means that the total amount of kWh consumed by the set of virtual machines, the average amount of consumed kWh per virtual machine and the total amount of consumed kWh for CPU, memory and storage should be presented.

6.5 Step 5: translate energy consumption into economic and environmental impact

In this step, the raw energy consumption data of step 4 is translated to environmental and economic impact, using several equations. The following steps should be taken:

Calculate equation 1 to determine the carbon footprint of a set of virtual machines (CF_{set}) in which VM is Virtual Machine, *e* is energy consumption and α is grams of CO_2 per kWh.

$$(1)\ CF_{set} = \sum_{i=1}^{x} VM_{ei} * \alpha$$

Calculate equations 2, 3 and 4 to determine the total carbon footprint (CF) allocated to the CPU, memory and storage resources. In which vCPU/MEM/STOR present the virtual CPU, memory and storage respectively, *e* is energy consumption and α is grams of CO_2 per kWh.

$$(2)\ CF_{CPU} = \sum_{i=1}^{x} vCPU_{ei} * \alpha$$

$$(3)\ CF_{mem} = \sum_{i=1}^{x} vMEM_{ei} * \alpha$$

$$(4)\ CF_{stor} = \sum_{i=1}^{x} vSTOR_{ei} * \alpha$$

Calculate equation 6 to determine the energy costs (E_{cost}) incurred by the energy consumption of the set of virtual machines in which VM is Virtual Machine, *e* is energy consumption and β is euro per kWh.

$$(6)\ E_{cost} = \sum_{i=1}^{x} VM_{ei} * \beta$$

Calculate equations 7, 8 and 9 to determine the energy costs (E_{cost}) incurred by the energy consumption, allocated to CPU, memory and storage resources.

$$(7)\ E_{costCPU} = \sum_{i=1}^{x} vCPU_{ei} * \beta$$

$$(8)\ E_{costMEM} = \sum_{i=1}^{x} vMEM_{ei} * \beta$$

$$(9)\ E_{costSTOR} = \sum_{i=1}^{x} vSTOR_{ei} * \beta$$

The calculated carbon footprints and energy costs should be visually represented as in the two rectangles connected to step 5 in Figure 2 for environmental and economic impact respectively.

6.6 Step 6: interpretation of economic and environmental performance

To be able to interpret the environmental and economic impacts of the previous step, two benchmarks are provided:

- Internal benchmark: calculate equations 13 (for t ≤ t – 1), 14 (for t ≥ t – 1), 15 (for t ≤ t – 1) and 16 (for t ≥ t – 1), to benchmark the carbon footprint of the virtual machines and the virtual machine resources of the current time period with data from other relevant time intervals and/or the target that has been set in step 2. To correct for changing demand, the benchmark also contains the amount of CO_2-emissions and costs as an average per virtual machine.

$$(13)\ 1 - \frac{CF_{set}t}{CF_{set}(t-1)} \qquad (14)\ \frac{CF_{set}t}{CF_{set}(t-1)} - 1$$

with:
$CF_{set} = carbon\ footprint\ of\ a\ set\ of\ VM's$
$t = time\ interval$

$$(15)\ 1 - \frac{CF_{vRES}t}{CF_{vRES}(t-1)} \qquad (16)\ \frac{CF_{vRES}t}{CF_{vRES}(t-1)} - 1$$

with:
CF_{vRES}
$= carbon\ footprint\ of\ CPU, memory\ or\ storage$
$t = time\ interval$

- External benchmark: calculate equations 17 and 18 to benchmark the carbon footprint per unit of measurement (i.e. grams of CO_2/GHz) per type of virtual resource (CPU, memory and storage). To be able to calculate the equations 17 and 18 (example for CPU, similar equations apply to memory and storage), equations 10, 11 and 12 need to be calculated first to calculate the CO_2/Ghz, Mb and GB for each individual cloud service provider included in the benchmark.

$$(10)\quad CF_{Ghz} = \frac{\sum_{i=1}^{x} vCPU_{ci}}{\sum_{i=1}^{x} vCPU_{gi}}$$

$CF = carbon\ footprint$
$vCPU = Virtual\ CPU$
$g = consumed\ Gigahertz$
$c = carbon\ footprint$

$$(11)\quad CF_{Mb} = \frac{\sum_{i=1}^{x} vMEM_{ci}}{\sum_{i=1}^{x} vMEM_{mbi}}$$

$CF = carbon\ footprint$
$vMEM = Virtual\ Memory$
$mb = consumed\ megabyte$
$c = carbon\ footprint$

$$(12)\quad CF_{GB} = \frac{\sum_{i=1}^{x} vSTOR_{ci}}{\sum_{i=1}^{x} vSTOR_{gbi}}$$

$CF = carbon\ footprint$
$vSTOR = Virtual\ Storage$
$gb = consumed\ gb$
$c = carbon\ footprint$

$$(17)\quad 1 - \frac{CF_{Ghz}at}{CF_{Ghz}bt}$$

$with:$
$For\ a \leq b$
$CF_{Ghz} = carbon\ footprint\ per\ GHz$
$a = cloud\ service\ provider\ A$
$b = cloud\ service\ provider\ B$
$t = time\ interval$

$$(18)\quad \frac{CF_{Ghz}at}{CF_{Ghz}bt} - 1$$

$with:$
$For\ a \geq b$
$CF_{Ghz} = carbon\ footprint\ per\ GHz$
$a = cloud\ service\ provider\ A$
$b = cloud\ service\ provider\ B$
$t = time\ interval$

To use this external benchmark, the amount of Gigahertz, Megabytes and Gigabytes consumed by the CPU's, memory and storage units of the virtual machines respectively for the same time interval as used for retrieving the energy consumption data must be retrieved from the hypervisor-software.

7. Conclusion

The goal of this paper was to present a design for a performance measurement tool and its added value. The added value of the design can be viewed both from a scientific and practical point of view. Previous research is mainly aimed at the development of detailed and specific solutions for improving data center efficiency. In contrast, this measurement tool provides a comprehensive, structural approach for dealing with data center efficiency to improve the environmental performance of cloud computing. Moreover, the added value of the measurement tool is in the metrics and equations the measurement tool combines for calculating and interpreting the environmental and economic impacts of cloud computing together with setting goals and creating plans to improve these impacts.

From a more practical point of view, this measurement tool adds value to cloud service providers, because it enables them to develop a well-considered strategy for improving the environmental performance of their services. The possibility to benchmarks against other providers can be used for improvement of the quality of services based on environmental performance. For the cloud customer, this external benchmark can be used in the process of vendor selection to judge cloud service providers based on the environmental performance of their services, which is currently not possible.

5. Discussion

There are some limitations that need to be discussed. First of all, the measurement tool is based on a narrow scope, which means that it applies to a particular part of the cloud computing life-cycle, while neglecting other impacting parts of this life-cycle. Secondly, the reliability and validity of this measurement tool has not been properly tested yet due to a lack of data, despite this being a basic condition for the development of measurement methodologies.

9. Recommendations and future research

Following the discussion points, it is recommended to perform research on possible ways to gather data for testing the reliability and validity of the measurement tool. Next to that, the environmental impact of other elements of the cloud computing life-cycle should be investigated, to work towards a more comprehensive model for managing the environmental performance of cloud computing.

References

2nd Watch. (2014). Public Cloud Workloads Gain Steam According to Latest 2nd Watch Survey.

Retrieved 3-3, 2015, from http://www.marketwired.com/press-release/public-cloud-workloads-gain-steam-according-to-latest-2nd-watch-survey-1976058.htm

Alger, D. (2009). *Grow a greener data center*: Pearson Education.

Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., . . . Stoica, I. (2010). A view of cloud computing. *Communications of the ACM, 53*(4), 50-58.

Basu, R. (2004). *Implementing quality: A practical guide to tools and techniques: Enabling the power of operational excellence*: Cengage Learning EMEA.

Cisco. (2014). Cisco Global Cloud Index: Forecast and Methodology, 2013-2018. Retrieved March 3rd, 2015, from http://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/Cloud_Index_White_Paper.pdf

Deming, W. E. (1982). *Quality, productivity, and competitive position*: Massachusetts Institute of Technology Center for Advanced En.

Elgelany, A., & Nada, N. (2013). Energy Efficiency for Data Center and Cloud Computing: A Literature Review. *Energy, 3*(1).

EPA. (2009). EPA ENERGY STAR program requirements for computer systems - draft 4. Washington, DC: Environmental Protection Agency.

Garg, S. K., & Buyya, R. (2011). Green cloud computing and environmental sustainability. *Harnessing Green IT: Principles and Practices (2020)*, 315-340.

Grossman, R. L. (2009). The case for cloud computing. *IT professional, 11*(2), 23-27.

Jenkin, T. A., Webster, J., & McShane, L. (2011). An agenda for ‘Green’information technology and systems research. *Information and Organization, 21*(1), 17-40.

Jing, S.-Y., Ali, S., She, K., & Zhong, Y. (2013). State-of-the-art research study for green cloud computing. *The Journal of Supercomputing, 65*(1), 445-468.

Lavallée, B. (2014). Undertaking the Challenge to Reduce the Data Center Carbon Footprint. Retrieved 25th of March, 2015, from http://www.datacenterknowledge.com/archives/2014/12/17/undertaking-challenge-reduce-data-center-carbon-footprint/

Market Research Media. (2014). Global Cloud Computing Market Forecast. Retrieved March 3rd, 2015, from http://www.marketresearchmedia.com/?p=839

Uddin, M., & Rahman, A. A. (2012). Energy efficiency and low carbon enabler green IT framework for data centers considering green metrics. *Renewable and Sustainable Energy Reviews, 16*(6), 4078-4094.

Verschuren, P., & Hartog, R. (2005). Evaluation in design-oriented research. *Quality and Quantity, 39*(6), 733-762.