

# Monitoring and assessing energy consumption and CO<sub>2</sub> emissions in cloud-based systems

Cinzia Cappiello, Sumit Datre, MariaGrazia Fugini  
Paco Melia', Barbara Pernici, Pierluigi Plebani  
Dipartimento di Elettronica, Informazione e Bioingegneria  
Politecnico di Milano  
Milano, Italy  
name.surname@polimi.it

Michael Gienger, Axel Tenschert  
High Performance Computing Center Stuttgart  
Universität Stuttgart  
Stuttgart, Germany  
surname@hirs.de

**Abstract**—The development and maintenance of cloud sites are often characterized by energy waste and high CO<sub>2</sub> emissions. Energy efficiency and the decrease of the CO<sub>2</sub> emissions in cloud-based systems can be only obtained by adopting suitable actions and techniques (e.g., utilization of green energy sources, reduction of the number of physical and virtual machines, usage of the greener machines). In order to evaluate the suitability of these different actions, it is necessary to define a measure for greenness of the whole system. For this reason, this paper defines a set of metrics to assess the greenness of a cloud infrastructure. In order to provide a detailed view of the behaviour of the system and to facilitate the identification of the causes of the energy waste, metrics have been defined at different layers of the system (i.e., application, virtualization, infrastructure layers). The monitoring infrastructure that is necessary to retrieve all the data required for the assessment of the identified set of metrics is also described.

**Index Terms**—CO<sub>2</sub> emissions, energy efficiency, cloud computing

## I. INTRODUCTION

According to several papers (e.g., [1]), the impact of cloud data centers on the energy consumed worldwide is getting more and more relevant. As reported, these centers may consume the equivalent of 180,000 homes and waste a lot of energy during their lifecycle. The PUE (Power Usage Effectiveness) is the traditional metric adopted to measure the efficiency of a data center with respect to the energy waste. As discussed in [2], this metric is valuable to give an immediate and intuitive way for classifying data centers regarding their greenness, but PUE is not enough to clearly identify which are the most critical elements and which are the interventions that can mitigate the energy waste. On this basis, the Green Grid Consortium [3] enlarged the set of metrics considering additional aspects like productivity and efficiency of the IT system as well as the resource saturation. At the same time, measures to consider the CO<sub>2</sub> emissions have been also defined [1].

Although the proposed metrics are very useful, they are mostly focused on the infrastructure of the cloud data center. On the contrary, less attention is paid to the estimation of the greenness of the applications that are running on the cloud centers. The goal of this paper is to fill this gap proposing:

(i) a set of metrics, which extends the ones currently present in the literature, able to assess the energy consumption in a federated cloud also considering the virtualization and the application level; (ii) the analysis of the relationships among the metrics in order to figure out how the energy consumed is distributed among different components; (iii) a monitoring system that supports the calculation of the metrics that are considered relevant in a given scenario.

This work is one of the initial results of the EU *ECO<sub>2</sub>Clouds* Project<sup>1</sup>. The project aims to propose a set of methods and tools to reduce the CO<sub>2</sub> emissions of federated clouds. In that context, the BonFIRE platform<sup>2</sup> represents the technological infrastructure adopted to validate the work.

The rest of the paper is structured as follows. After discussing, in Section II, related work in this topic, Section III defines the way in which applications run on federated clouds. Section IV contributes the definition of the set of metrics and Section V illustrates the adopted infrastructure monitoring. The way the metrics influence monitoring and the vice-versa are discussed in Section VI. Finally, Section VII concludes the paper with some hints for possible future work.

## II. RELATED WORK

In the literature, energy efficiency is considered a critical factor and it is often measured through the so called Eco Metrics. Such metrics are often considered as an extension of the Key Performance Indicators (KPIs) that measure how well an organization performs an activity which is success-critical.

In general, there is still no widely accepted metric set for measuring and monitoring energy consumption in data centers and in the cloud. In order to clarify the energy consumption assessment procedures, the Standard Performance Evaluation Corporation (SPEC) released a power and performance methodology to let performance benchmark designers and implementers integrate a power component into their benchmark [4]. This document presents guidelines for designing new benchmarks to provide a more complete view of energy

<sup>1</sup><http://www.eco2clouds.eu>

<sup>2</sup><http://www.bonfire-project.eu>

consumption. Other recent research contributions ([5], [6], [7]) propose different ways of monitoring and representing energy and other parameters such as QoS. In the industrial context, the Greenhouse Gas (GHG) protocol [8] provides calculation tools and guidelines to quantify and manage emissions in commercial data centers and cloud platforms.

Focusing on metrics, a set of layered eco metrics, named Green Performance Indicators (GPIs) [9], is proposed in the GAMES project [10] to analyze energy efficiency of data centers and service-based applications. GPIs regard system components (CPU, memory, I/O channels, disk) and the life-cycle of applications (development costs, Quality of Service, energy efficiency during run-time). Upon detection of energy leakages through monitoring, green actions are set in place (e.g., substitution of services, storage operation change) to remove or reduce energy loss. [11] illustrates the results given by GAMES through monitoring of GPIs on data centers; it suggests that both energy efficient components should be selected and the usage of existing components should be improved. For this second aim, a set of usage-centric GPI metrics is proposed.

[12] proposes energy efficiency and low carbon emission IT frameworks for complex server farms, using energy saving techniques like virtualization and green metrics. The framework applies traditional green metrics like PUE, DCE (Data Center Effectiveness) and CEC (Carbon Emission Calculator).

[6] reviews methods and technologies adopted for an energy-efficient use of hardware and network infrastructures. Best practices and relevant literature are reviewed and key research challenges are identified, considering metrics for measuring hardware energy-efficiency, energy-aware scheduling and techniques such as self-aware runtime adaptation, based on the trade-offs that can result in an optimal balance between performance, QoS and energy consumption.

[13] is an extensive survey of power and energy-efficient design of data centers and cloud computing systems. A review is given on the relationship between various metrics, such as workload and power consumption metrics, architectural metrics (like IPC and MPC), the energy-per-transaction metric which depends on both CPU and disk utilizations, energy-performance metrics and user-centric metrics that encompass both performance and fairness.

In [14], metrics for measuring and improving data center efficiency are explored. Metrics at different levels, from the infrastructure components to the entire data center are reviewed. The primary data center efficiency metric, PUE, is discussed as well as its variants. The challenges of defining a metric around computing output or data center work are also presented.

[15] considers Green Cloud Environments and minimization of its energy consumption. An analysis of energy consumption is proposed based on energy consumption patterns, grounded on measurable metrics computed on runtime tasks to compare the relation between energy consumption and cloud workload, computational tasks as well as system performance. A set of energy consumption profiling metrics is proposed. Finally, as already mentioned in Section I, [1] provides a new set of

metrics that considers energy consumption and CO<sub>2</sub> emissions.

In ECO<sub>2</sub>Clouds, metrics are needed for a complete view of the energy consumed and environmental impact of cloud-based systems. Since cloud federations provide an environment to deploy cloud applications across multiple cloud platforms, the heterogeneity and lack of standards in current clouds, together with the incomplete or inconsistent support for high-level operations, make this an open issue that is only partly addressed by current technologies and systems [16]. A set of layered and multifaceted metrics has to be investigated also in accordance to pressing challenges [17]. The first results of the ECO<sub>2</sub>Clouds metrics are presented in the next sections.

### III. FEDERATED CLOUD MODEL

Before introducing the set of metrics and the monitoring infrastructure, this section clarifies the model adopted to represent a federated cloud. First of all, we define a federated cloud as a set of hardware resources (e.g., hosts, storage, network devices) that are organized in distinct sites under the supervision of various operators and administrators.

As shown in Figure 1, the characteristics of the sites and the resources available in each of them represent the *infrastructural layer*. On top of this layer, there is the *virtualization layer* that contains the hypervisor and the Virtual Machines (VMs). The role of the VMs is to host the applications (or part of applications) that are used to manage the cloud infrastructure (e.g., the monitoring system) or that are offered to the final users (e.g., business applications). The applications running on the VMs compose the *application layer*.

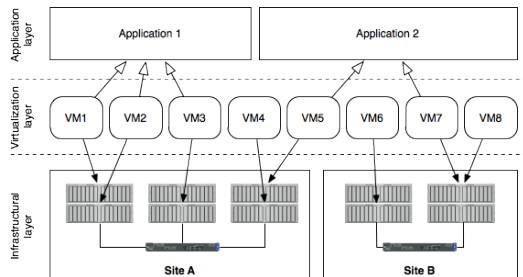


Fig. 1. Federated cloud model

Focusing on the application layer, we consider different kinds of applications. On the one hand, a siloed application runs on a single VM where all the resources required by the application are represented by the VM. On the other hand we consider federated applications modeled as business processes that can be seen as the composition of different tasks where not only the composing tasks could run on different VMs, but also a given task can be reused in different business processes.

Based on this model, hereafter we will adopt the following notation:

- $\{A_i\}$  is the set of applications deployed on the federated cloud.
- Each application is composed by  $j$  tasks deployed on  $k$  VM, i.e.,  $A_i = \langle \{T_{ij}\}, \{VM_{ik}\} \rangle$

<i>Metric</i>	<i>Definition</i>
CPU utilization	Average utilization of the processors inside a host. For each processor, this metric indicates how much of the processor capacity is in average in use by the system
IOPS	Total number of I/O operations per second (when performing a mix of read and write tests)
Availability	The probability that a request is correctly served by a specific host within a maximum expected time frame. In order to assess this metric, it is necessary to compare the number of satisfied requests with the total amount of the requests received by the analysed host. Note that, at the host level, the request concerns the deployment of a VM on a given host
Energy Consumption	The energy consumed by the analysed host in a specific time period

TABLE I  
INFRASTRUCTURE LAYER METRICS – HOST

- $P_{ijk}$  represents the power consumed by the  $j$ -th task of the  $i$ -th application running on the  $k$ -th VM.
- $P_{ik}$  represents the power consumed by the  $k$ -th VM used by the  $i$ -th application also considering the possible other tasks of other applications running on the same VM.

#### IV. ENERGY METRICS

Energy consumption and CO<sub>2</sub> emissions in a federated cloud infrastructure can be decreased by considering several actions such as the utilization of green energy sources, the reduction of the number of physical and virtual machines and the utilization of the greener machines. In order to evaluate the suitability of these different actions, it is necessary to define the measure for energy efficiency of the system. For this reason, a set of metrics for measuring the greenness of an application running in the analysed cloud infrastructure has to be defined. This set of metrics should reflect the energy efficiency of IT systems from a holistic perspective and should allow the derivation of the interrelation between different components of IT cloud infrastructure.

On the basis for the deployment model of applications in a federated cloud infrastructure described in Section III, we adopted a layered metric approach. Infrastructural, virtualization and application layers can be seen as organized hierarchical structures. In the following subsections we present the metrics to consider for each level.

##### A. Infrastructural layer metrics

The infrastructure layer includes the host layer and the site layer. The former focuses on the analysis of the behaviour of the single host inside a site. The latter considers the whole cloud site providing an overall picture of its greenness.

1) *Host layer metrics*: The metrics defined at the host level are defined in Table I. CPU utilization provides information about the load of a specific host. It is important to consider this metric to avoid low host utilization. Furthermore, the analysis of CPU utilization together with the host energy consumption provides information about the energy efficiency of a specific host. Data about the I/O operations per second are useful to understand the type of application that is running on a specific host (e.g., interactive, batch, high performance computing). Finally, the availability is an index of the host reputation.

<i>Metric</i>	<i>Definition</i>
PUE	Measure that compares the power used by the entire infrastructure with the power used for computation.
Site utilization	Average utilization of the power drawn by the IT equipment with respect to the power capacity of the site
Storage utilization	Percentage of storage used with respect to the overall storage capacity within the site.
Availability	The probability that a request is correctly served by a site within a maximum expected time frame
Green Efficiency Coefficient (GEC)	Percentage of energy consumed by the site that is produced by green energy sources. This metric is calculated as the ratio between the green energy consumed by the site and the total energy consumed by the site.
CO <sub>2</sub> emissions	Quantity of CO <sub>2</sub> emitted by the site.

TABLE II  
INFRASTRUCTURE LAYER METRICS – SITE

In fact, the greater the availability the less the probability to have unsatisfied requests. Hence, an effective application deployment strategy should prefer highly available hosts.

2) *Site layer metrics*: In a federated cloud infrastructure, metrics that characterize the site activity are needed in order to evaluate the most suitable cloud site the application should be deployed on. The metrics defined at the site level are listed and defined in Table II. They aim to evaluate the energy efficiency of the cloud site, as well as its overall resource usage and CO<sub>2</sub> emissions. For the emissions, it is also important to monitor the percentage of consumed energy that is provided by green energy sources (i.e., GEC metric).

The metrics are the basis for selecting the site on which the VM composing an application should be deployed. Indeed, this selection firstly depends on the utilization of a site and of the storage. In fact, site congestions should be avoided. Another factor to consider is the site availability. Similarly to the host availability, the site availability provides indications about the reliability of the site. Moreover, in order to achieve a greener application deployment, sites with high Green Efficiency Coefficient (GEC) should be preferred.

##### B. Virtualization layer metrics

Metrics at the virtualization layer aim to characterize the virtual machines on which the applications are running (see Table III). Information about the usage of VM resources has to be analysed in order to evaluate if the current deployment can be further improved and thus optimized. The analysis of the VM energy consumption aims to understand how the energy consumed by the host is distributed among the VMs deployed on it. Moreover, we also define new metrics that are inspired by the data center metrics proposed in the last years, especially by The Green Grid Consortium.

The idea is to redefine the classical infrastructural metrics, like PUE and Data Centre Energy Productivity (DCeP), at virtualization level to measure the impact of the application tasks in terms of energy consumption and carbon emissions. The original PUE compares the power used by the entire infrastructure divided by the power used for the IT infrastructure. This metric evaluates the power wasted to feed devices that are not directly involved in the computation (e.g., cooling,

<i>Metric</i>	<i>Definition</i>
CPU Usage	Processor utilization percentage for a running application over a run time interval. It is calculated by using the ratio between the amount of used CPU and the amount of allocated CPU.
Storage Usage	Storage utilization percentage for data-read and -write operations on the corresponding storage device, computed as the ratio between the used disk space and the allocated disk space.
I/O Usage	Percentage of process execution time in which the disk is busy with read/write activity.
Memory Usage	Ratio of the average size of the portion of memory used by the process to the total amount of memory available for the application.
Energy Consumption	The energy consumed by the analysed VM in a specific time period.
VM-PUE	Measure of how efficiently a VM uses the provided power
VM-EP (VM Energy Productivity)	Ratio between the output of the VM in a certain time interval and the energy consumed
VM-GE (VM Green Efficiency)	Information about the portion of energy consumed by the VM that is produced by green energy sources

TABLE III  
VIRTUALIZATION LAYER METRICS

lighting). VM-PUE<sub>ik</sub> compares the  $P_{ik}$  that is the total amount of power required by the VM<sub>ik</sub> with  $P_{ijk}$ , the power used to execute the application tasks:

$$\text{VM-PUE}_{ik} = \frac{P_{ik}}{\sum_j P_{ijk}}$$

In the literature, there are some contributions that describe how to obtain the power consumed by a task, i.e.,  $P_{ijk}$  (e.g., [18]). Moreover, it can also be computed considering the system processes running on the VM related to the task. Given these system processes, tools like `ptop` can be used to estimate the power they consume.

We also considered DCeP that is the ratio between the work output of the data center and the total energy for the data center. At the virtualization level, this metric can be redefined as the VM Energy Productivity (VM-EP):

$$\text{VM-EP}_{ik} = \frac{NTrans_{ik\Delta t}}{P_{ik} * \Delta t}$$

where  $NTrans_{ik\Delta t}$  is the number of completed executions of the tasks deployed on the VM<sub>ik</sub> in a time interval  $\Delta t$ .

Finally, we are also interested in the greenness of the VM in terms of how much green energy is used to run VM<sub>ik</sub>. To do this, we consider the Green Efficiency Coefficient (GEC) factor and we multiply it by the energy consumed by the virtual machine VM<sub>ik</sub>:

$$\text{VM-GE}_{ik} = \text{GEC} * (P_{ik} * \Delta t)$$

### C. Application layer metrics

The application layer can be analyzed by using the metrics contained in Table IV.

On the basis of the PUE, Energy Productivity index and Green Efficiency Coefficient defined at virtualization level, we define the same metrics at application level as described in the following. An indication of the energy efficiency related to the execution of an application  $A_i$  is:

$$\text{A-PUE}_i = \frac{\sum_k P_{ik}}{\sum_{jk} P_{ijk}}$$

<i>Metric</i>	<i>Definition</i>
Task Execution Time	The time taken to execute the specific task.
Application Execution Time	The time taken to execute the whole application.
Energy Consumption	The energy consumed from the analysed application in a specific time period. This metric is calculated by aggregating the energy consumed by the VMs through which the application is deployed.
Response Time	The average time taken to handle user requests. This metric is particularly relevant for interactive application. Note that for batch application the response time will coincide with the application execution time.
Throughput	Number of executions of an application within a specific time frame.
A-PUE	Measure of how efficiently an application uses the provided power
A-EP (Application Energy Productivity)	Ratio between the number of execution of an application in a certain time interval and the energy consumed
A-GE (Application Green Efficiency)	Provide information about the portion of energy consumed to execute a specific application that is produced by green energy sources.

TABLE IV  
APPLICATION LAYER METRICS

The Application Energy Productivity of an application  $A_i$  is defined as the ratio between the number of executions of the whole application and the total energy consumed by all the VMs on which the application tasks are deployed:

$$\text{A-EP}_i = \frac{NTrans_{i\Delta t}}{\sum_k (P_{ik} * \Delta t)}$$

The greenness of the application is calculated also by means of the Application Green Efficiency factor defined as:

$$\text{A-GE}_i = \sum_k \text{VM-GE}_{ik}$$

## V. MONITORING INFRASTRUCTURE

The assessment of the metrics defined in the previous section is enabled by the monitoring infrastructure. This section presents an overview of the monitoring infrastructure and its realization within the ECO<sub>2</sub>Clouds project.

For the assessment of the metrics described in Section IV, it is necessary to monitor the power consumption of the hosts using power distribution units (PDUs) and different aspects on the host, virtualization and application layer. Therefore, besides the layers described in Section III, it is necessary to consider PDUs that are additional external hardware power devices distributing electric power to the physical compute components and network devices. Furthermore, these PDUs allow gathering monitoring data for the power consumption of each physical host. Measuring specific parameters on the host layer provides information about the overall performance of the different hosts within the correlation of energy data and allows the assessment of additional metrics about the VMs deployed on them. Finally, analysing the trend parameters monitored at the virtualization layer provide knowledge about the VMs themselves and the application running currently.

The monitoring infrastructure is mainly focusing on these layers in order to collect suitable data about the performance

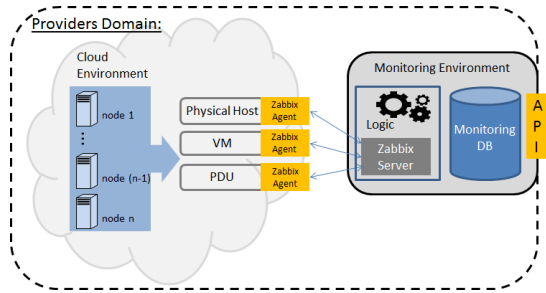


Fig. 2. Collecting Data via the Monitoring Environment

and the energy consumption. Figure 2 shows a schematic representation of the monitoring environment including its components.

In order to ensure a well working monitoring infrastructure, a dedicated monitoring server with the corresponding agents will be used to capture monitoring data. In particular, ECO<sub>2</sub>Clouds will use the on-hand Zabbix monitoring tool as each BonFIRE infrastructure provider already offers a Zabbix server with Zabbix agents to capture the data and reduce the monitoring overhead. These agents can be used to retrieve data for physical and virtual resources and the energy related data. The installation and the configuration of the agents is the most important goal to establish a scalable and flexible monitoring solution: Zabbix is adapted for the use in the ECO<sub>2</sub>Clouds project to support such objectives. For the virtual resources, a decoupled dynamic framework based on Zabbix will allow the configuration of the virtual machines in order to retrieve the proposed metrics. For that purpose, an additional Zabbix aggregator needs to be started in order to monitor the defined metrics and link to the physical ones. For the PDUs, the native SNMP protocol support of Zabbix will be used to collect the energy metric values.

The general monitoring infrastructure is presented in Figure 3. Each cloud facility provides physical nodes with attached PDUs in order to host virtual resources. With the help of collector scripts written in bash, executed through the Zabbix agents, a dedicated host used for operating the Zabbix server is handling these monitoring issues and receives the required monitoring data. This Zabbix server provides a complete framework for monitoring: it offers the collection of data and also the functionality of storing monitoring information into a database backend to analyse the data afterwards. The used MySQL backend is also maintained by the Zabbix server: clean-up scripts are keeping the size small according to the defined settings and enable a data mining process. In addition to the web based graphical user interface (GUI) to manage the infrastructure, an application programming interface (API) is provided by Zabbix. With the help of that, the ECO<sub>2</sub>Clouds Accounting Service will be able to extract data from the monitoring databases into the Accounting Service database to provide monitoring information for the Scheduler as well as for the ECO<sub>2</sub>Clouds Portal.

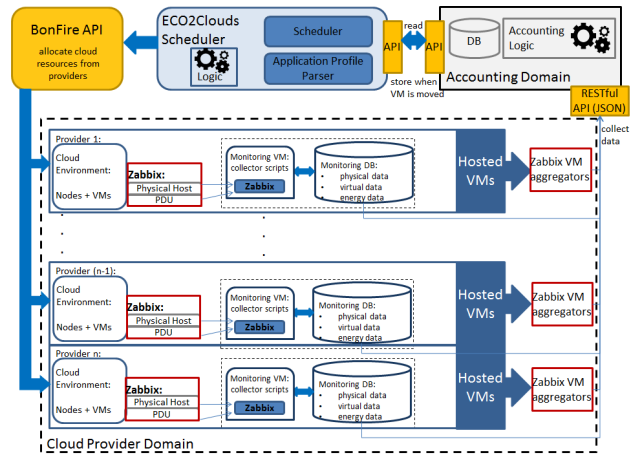


Fig. 3. The general Monitoring Infrastructure

## VI. GUIDELINES FOR MONITORING AND ASSESSING GREEN METRICS

As explained in Section V, the monitoring infrastructure allows the retrieval of most of the data required to assess the set of metrics defined in Section IV.

Note that the adoption of a hierarchical structure suggests also a hierarchical relationship between layers: the application layer can be defined as the top layer, the infrastructure layer as the bottom layer. A layer serves the layer above it and is served by the layer below it. For example, the assessment of the metrics at the application levels exploits the metrics defined for the other two layers. Figure 4 presents a simplified view of the architecture needed for the assessment of the defined metrics. Note that the bidirectional arrows that connect infrastructure and virtualization metrics to the component that contains the assessment logic mean that the metrics defined at this level are used to calculate other metrics.




Fig. 4. Metrics Assessment Infrastructure

Figure 4 also highlights that the assessment phase often requires additional data, as for example the evaluation of the CO<sub>2</sub> emissions. In fact, the carbon footprint measures the total amount of greenhouse gas (GHG) emitted during the full lifecycle of a product, service or system. It is calculated in terms of carbon dioxide equivalent (CO<sub>2</sub>e) using the relevant 100-year global warming potential [19]. The overall carbon footprint of a data center includes 1) the carbon embedded in the building hosting the center, facilities and IT equipment;

2) GHG emissions due to electricity and/or fuel consumption during the operational phase; and 3) the decommissioning of the center at the end of its operating life and the recycling or final disposal of all the materials [20]. The relative contribution of the embedded carbon depends upon the characteristics of the infrastructure, but for most data centers emissions associated with the operational phase are 2 to 3 orders of magnitude larger than the contribution of the infrastructure [20]. Emissions during operation are due to the energy consumption of the IT equipment (to process information) and of the cooling systems (to keep temperature stable). The carbon footprint of electricity consumption depends upon the energy production mix (the proportion of different power generation technologies) and the efficiency of the distribution grid. The energy mix is the most significant factor affecting GHG emissions. Fossil fuels (and in particular coal) are the most carbon-intensive energy sources. On the contrary, nuclear power and renewable sources are often considered as carbon neutral, because they emit virtually no GHGs from the operation. However, even those sources leave a footprint during the construction and decommissioning phases (see [21] for a review). In particular, a general agreement is still lacking on the emission factors of nuclear power (e.g., [22], [23], [24], [25]). Anyway, in general, strategies for reducing CO<sub>2</sub> in a federated cloud system include the increase of energy efficiency and fuel switching.

In our approach, the evaluation of the CO<sub>2</sub> emissions is based on the emission factors (gCO<sub>2</sub>e/kWh) provided by the national grids. Assuming that we know the average power consumption (AP) for a specific site, the energy (kWh) consumed in one year of operation can be calculated by multiplying AP by the number of hours in a year. CO<sub>2</sub> emissions results multiplying the energy consumed by the emission factor.

## VII. CONCLUSIONS

This paper has shown the usefulness of defining metrics for monitoring different layers of an eco-aware cloud system. The proposed set of metrics is mainly assessed on the basis on the values of the physical and virtual-related parameters gathered by using the Zabbix tool and additional PDUs. By using PDUs, accurate assumptions regarding the energy consumption of the physical compute nodes and the hosted virtual machines of distributed cloud environments related to the CO<sub>2</sub> emission become possible. Future work focuses on the refinement of the metrics for calculating the energy consumption of the VMs. Furthermore, applications will be deployed and executed in a cloud-based system in order to assess the proposed metrics and analyze relationships among metrics and the presence of significant patterns.

## ACKNOWLEDGMENT

This work has been supported by the EU-FP7-ECO<sub>2</sub>Clouds project (<http://eco2clouds.eu/>) under contract number 318048.

## REFERENCES

- [1] Greenpeace, "How Clean Is Your Cloud?" Tech. Rep., 04 2012.
- [2] Nlyte Software, "Why PUE alone is not Enough!" Tech. Rep., 2011.

- [3] D. Azevedo, J. Cooley, M. Patterson, and M. Blackburn, "Data Center Efficiency Metrics: mPUE, Partial PUE, ERE, DCcE," GreenGrid, Tech. Rep., 2011.
- [4] SPEC, "SPEC Power and Performance, Benchmark Methodology," Tech. Rep., 2011.
- [5] I. Dumitru, S. Ploix, I. Fgran, and S. Iiescu, "Energy efficiency dependency analysis for a data center," in *Advances in Intelligent Control Systems and Computer Science*, L. Dumitrache, Ed. Springer Berlin Heidelberg, 2013, vol. 187, pp. 307–321.
- [6] 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*, vol. 53, no. 7, pp. 1045–1051, 2010.
- [7] E. Curry, G. Conway, B. Donnellan, C. Sheridan, and K. Ellis, "Measuring energy efficiency practices in mature data center: A maturity model approach," in *Computer and Information Sciences III*, E. Gelenbe and R. Lent, Eds. Springer London, 2013, pp. 51–61.
- [8] M. Meinshausen, N. Meinshausen, W. Hare, S. C. Raper, K. Frieler, R. Knutti, D. J. Frame, and M. R. Allen, "Greenhouse-gas emission targets for limiting global warming to 2°C," *Nature*, vol. 458, no. 7242, pp. 1158–1162, 2009.
- [9] A. Kipp, T. Jiang, M. Fugini, and I. Salomie, "Layered green performance indicators," *Future Generation Computer Systems*, vol. 28, no. 2, pp. 478–489, 2012.
- [10] A. Kipp, T. Jiang, J. Liu, M. Fugini, M. Vitali, B. Pernici, and I. Salomie, "Applying green metrics to optimise the energy consumption footprint of it service centres," *International Journal of Space-Based and Situated Computing*, vol. 2, no. 3, pp. 158–174, 2012.
- [11] D. Chen, E. Henis, R. I. Kat, D. Sotnikov, C. Cappiello, A. M. Ferreira, B. Pernici, M. Vitali, T. Jiang, J. Liu, and A. Kipp, "Usage centric green performance indicators," *SIGMETRICS Performance Evaluation Review*, vol. 39, no. 3, pp. 92–96, 2011.
- [12] M. Uddin and A. A. Rahman, "Energy efficiency and low carbon enabler green it framework for data centers considering green metrics," *Renewable and Sustainable Energy Reviews*, vol. 16, no. 6, pp. 4078 – 4094, 2012.
- [13] A. Beloglazov, R. Buyya, Y. C. Lee, and A. Y. Zomaya, "A taxonomy and survey of energy-efficient data centers and cloud computing systems," *CoRR*, vol. abs/1007.0066, 2010.
- [14] M. K. Patterson, "Energy efficiency metrics," in *Energy Efficient Thermal Management of Data Centers*, J. Yogendra and K. Pramod, Eds. Springer, 2012.
- [15] A. Uchechukwu, K. Li, and Y. Shen, "Improving cloud computing energy efficiency," in *Cloud Computing Congress (APCloudCC), 2012 IEEE Asia Pacific*. IEEE, 2012, pp. 53–58.
- [16] B. Rochwerger, D. Breitgand, A. Epstein, D. Hadas, I. Loy, K. Nagin, J. Tordsson, C. Ragusa, M. Villari, S. Clayman *et al.*, "Reservoir-when one cloud is not enough," *Computer*, vol. 44, no. 3, pp. 44–51, 2011.
- [17] R. Moreno-Vozmediano, R. Montero, and I. Llorente, "Key challenges in cloud computing to enable the future internet of services," *IEEE Internet Computing*, vol. 99, no. PrePrints, 2012.
- [18] A. Kansal and F. Zhao, "Fine-grained energy profiling for power-aware application design," *SIGMETRICS Performance Evaluation Review*, vol. 36, no. 2, pp. 26–31, 2008.
- [19] L. A. Wright, S. Kemp, and I. Williams, "Carbon footprinting: towards a universally accepted definition," *Carbon Management*, vol. 2, no. 1, pp. 61–72, 2011.
- [20] Schneider Electric, "Allocating Data Center Energy Costs and Carbon to IT Users," Tech. Rep., 2011.
- [21] M. Lenzen, "Current state of development of electricity- generating technologies: A literature review," *Energies*, vol. 3, no. 3, pp. 462–591, 2010.
- [22] B. K. Sovacool, "Valuing the greenhouse gas emissions from nuclear power: A critical survey," *Energy Policy*, vol. 36, no. 8, pp. 2950 – 2963, 2008.
- [23] J. Beerten, E. Laes, G. Meskens, and W. Dhaeseleer, "Greenhouse gas emissions in the nuclear life cycle: A balanced appraisal," *Energy Policy*, vol. 37, no. 12, pp. 5056 – 5068, 2009.
- [24] E. S. Warner and G. A. Heath, "Life cycle greenhouse gas emissions of nuclear electricity generation," *Journal of Industrial Ecology*, vol. 16, pp. S73–S92, 2012.
- [25] Storm van Leeuwen, J.W. , "Nuclear power, energy security and CO2 emission," Tech. Rep., 2012. [Online]. Available: <http://www.stormsmith.nl>