

Energy Efficiency and Power Measurements: an Industrial Survey in Data Centers

Giuseppe Procaccianti
Computer Science Department
Vrije Universiteit Amsterdam
Amsterdam, the Netherlands
Email: g.procaccianti@vu.nl

Aristeidis Routsis
Computer Science Department
Vrije Universiteit Amsterdam
Amsterdam, the Netherlands
Email: arisroutsis@gmail.com

Abstract—Context. In the last few years, the energy consumption of data centers increased substantially. Currently, energy costs are the main expenditure in data center ownership. This creates the need for more precise quantitative metrics and measurements for energy management in the data center industry. However, despite the effort of multiple public and private agencies to address this problem, there is still little knowledge among practitioners on such topic.

Aim. The goal of our work is to investigate the state-of-the-Practice of power measurements and energy efficiency techniques in the data center industry.

Method. We interviewed 7 practitioners and stakeholders from 4 companies of different size. We used semi-structured interviews of 1 hour as our research methodology.

Results. Our results show that all the companies we interviewed measure power consumption, but only up to rack-level granularity. Several energy efficiency practices have been adopted, with positive results. However, the only metric considered to evaluate the impact of such practices is the Power Usage Effectiveness (PUE), although the representativity of such metric has been widely questioned in scientific literature and by our interviewees as well. We observed a lack of awareness with respect to other metrics, which could also be used to share such information across different stakeholders (e.g. suppliers, infrastructure managers, customers).

Conclusions. Our contribution motivates further efforts from the data center industry to go beyond the PUE and adopt more advanced metrics to evaluate the energy efficiency of their data centers. This is a required step to ensure that real improvements are made towards sustainability goals.

Index Terms—Power Measurements, Energy Efficiency, Datacenters, Industrial Survey

I. INTRODUCTION

In the last years, the energy impact of data centers became a hot topic. Worldwide, the total number of server has increased approximately 150 times in less than a decade [23]. During this period, the average computational power per server has increased by a factor of 10 and smaller form factors have packed that power into much smaller spaces [23]. Higher density servers have higher heat generation, and this results in higher power consumption. According to the U.S. Environmental Protection Agency [12], in the last decade, the power consumption of data centers is steadily increasing; more specifically, the amount of energy consumed from data centers has doubled between 2000-2006. More recent figures calculate IT power consumption

reaching almost 5% of worldwide electricity consumption (920 TWh) [16]. Data centers are responsible for about one-third of this amount and reducing their energy impact is a priority not only for environmental reasons, but also for their economic sustainability: operational electricity costs account for 20% of a data center Total Cost of Ownership (TCO) [17].

Nowadays, many governments and companies globally have started adopting practices and initiatives to increase IT energy efficiency. An example is the European Code Of Conduct (EU COC) [2], a voluntary initiative which aims to bring interested stakeholders, manufacturers, vendors, consultants and utilities together, in order to adopt and abide by a set of agreed commitments. The COC also aims to inform data center operators/owners and also to encourage them to decrease their data center's energy consumption. In order for this to be attained, data centers managers have to understand, manage, analyze, and measure the power usage and consumption in their facilities.

As regards the state-of-the-Art, a significant amount of techniques have been described in scientific literature [25] [13] [20] and several metrics have been proposed. In this research, we will instead investigate the *state-of-the-Practice* of data center power measurements and energy efficiency. More specifically, this research concentrates on two main aspects: first, we will elicit what energy efficiency techniques and measurements are being used by data center managers, and secondly we will assess the awareness of the stakeholders in the data center industry with respect to the techniques and metrics already available in literature.

The structure of this paper is as follows: in the following Section II we present the problem and the research questions we address in this research. Section III provides a catalog of proposed metrics for energy efficiency, and a brief overview of surveys related to our study. In Section IV, we describe our research methodology. In Section V the results of our semi-structured interviews are presented. In Section VI we discuss the results and in Section VII we identify the limitations of our research and how we tried to overcome them. Finally, Section VIII concludes the paper, answering the research questions and also identifying topics for future research.

II. PROBLEM STATEMENT AND RESEARCH QUESTIONS

It is common knowledge that "You cannot manage what you cannot measure", i.e. good decision making requires timely and insightful information. Information about the availability, usage and quality of service of IT resources is essential in identifying areas of improvement in order to boost efficiency and reduce costs. However, proper techniques and metrics must be used. In the next Section, we present a number of energy efficiency techniques and metrics that have been proposed by several studies in the recent years. We pose the following problem: to what extent are those techniques and metrics *actually* being used by the data center industry?

In this sense, evidence seems counter-intuitive: the importance of measuring the energy efficiency of data centers cannot be overstated, but apparently many companies do not see the value in doing so. A survey conducted on more than 1,500 data center owners and operators, presented at the Datacenter Dynamics conferences in 2008 [7] indicate that fewer than one in three (461 out of 1,523 respondents) were using a data center efficiency metric, whereas the remaining (1,062, 69.7 percent) reported they did not. When the same group was asked whether they plan to employ efficiency metrics in the future, 642 (42.2 percent) indicated they did, whereas 492 (32.3 percent) indicated they did not. The remaining 389 (25.5 percent) did not reply or else indicated that they did not know [7].

We then set forth to investigate the current practices of data centers regarding the energy efficiency techniques and metrics, and the awareness of stakeholders towards this problem. We formulated these goals as two separate research questions (RQs) and sub-questions, presented below:

- 1) *What is the state-of-the-practice regarding energy efficiency and power measurements in the data center industry?*
 - Which power measurements are taken and how?
 - Are those measurements used in appropriate indicators or metrics regarding the overall data center energy efficiency?
- 2) *What is the level of awareness of the stakeholders in the data center industry towards power measurements?*
 - Who are the stakeholders and what is the concern of each stakeholder?
 - Are the power measurements considered useful from the stakeholders?

III. LITERATURE REVIEW

In this Section, we provide a literature review in energy efficiency metrics and measurements as a background. In addition, we also provide some references to other related surveys on energy efficiency practices and awareness in industry.

A. Metrics and Measurements

Figure 1 shows the points where energy use can be measured in the data center to assist in understanding the efficiency, losses and impacts of changes and improvements. As noted in [12] measurements can be categorized in three groups:

- 1) Simple measurements: these are the most basic levels of measurement, required to deliver a PUE/DCIE metric of the ratio of utility power to IT power for the facility. These measurements are simple to calculate and can be extremely informative for the operator. It is important for these measurements to record the external and data floor set temperatures along with the electrical measurements.
- 2) Detailed measurements: these mainly concentrate on the points where power loss is observed during the power delivery, as well as on the power redirected to non-IT loads. These measurements provide an insightful view on how to improve the facility and minimize power loss.
- 3) IT measurements: these are related to the power draw of specific IT devices. These are useful in situations where a physical device is dedicated to a specific task, but are typically regarded as less informative in a virtualized or grid infrastructure, or where resources such as network or storage are shared.




Fig. 1. Power measurements on data centers. Adapted from [12].

Many agencies worldwide have started to develop and apply metrics for measuring energy efficiency on data centers. The most influential metrics proposed by The Green Grid [1], McKinsey¹, and Uptime Institute² are presented in the following sections. An attempt to harmonize and consolidate this set of metrics is the European Code of Conduct for data centers (EU COC)³.

PUE: Power Usage Effectiveness (PUE) is the most widely used metric nowadays. PUE divides a facility's total

¹<http://www.mckinsey.com/>

²<https://uptimeinstitute.com/>

³<http://iet.jrc.ec.europa.eu/energyefficiency/ict-codes-conduct/data-centres-energy-efficiency>

power draw by the amount of power used solely by the data center's IT equipment:

$$PUE = \frac{\text{Total Facility Energy}}{\text{IT Equipment Energy}} \quad (1)$$

IT Equipment Energy considers the energy consumed by servers, networking devices, storage units and peripheral items. *Total Facility Energy* includes IT equipment and all data center-related primary electrical systems (power distribution units and electrical switching gear), standby power systems (Uninterruptible Power Sources [UPS]), cooling components, and other infrastructure elements. The closer the value is to 1 (the minimum possible score, all the power provided is used by IT equipment), the more efficient the data center is considered to be: a lower PUE indicates that a greater portion of the power going to a data center is used for its essential tasks and is not going wasted. For example, a PUE of 1.5 means a data center needs half as much additional power as is solely needed for the IT equipment to operate, whereas a PUE of 3.0 means a data center needs twice as much additional power for non-IT elements as it does for IT hardware.

However, PUE has major limitations: namely, it can greatly vary depending on *when* measurements are taken. For example, a high IT load increases power consumption of IT systems, however typically cooling systems do not scale accordingly. Hence, a data center might exhibit a lower PUE when under high load [11]. Also, if measured during winter, the PUE value will be lower than summer, when cooling is working all day. Hence, if not complemented with sufficient contextual information, PUE is not a reliable metric for energy efficiency [11]. In spite of that, it is a *de facto* standard metric for the data center industry.

DCIE: the Data Center Infrastructure Efficiency (DCIE) is the inverted PUE, where IT equipment power is divided by total facility power:

$$DCIE = \frac{\text{IT Equipment Energy}}{\text{Total Facility Energy}} \quad (2)$$

DCIE is expressed as a percentage: For example if a data center's PUE is 1.5, DCIE shows that 2/3 (1/1.5=66 percent) of the data center power is consumed by the IT equipment. DCIE is sometimes preferred to the PUE because increasing values indicate a higher energy efficiency (i.e. higher is better).

DCP: Data Center Productivity not only measures the consumption of a data center-related resource, but also quantifies the useful work that a data center produces based on the amount of energy it consumes. In short, DCP tries to define what a data center accomplishes in relation to what it consumes. The formula for DCP is:

$$DCP = \frac{\text{Useful Work produced by data center}}{\text{Resource Consumed producing the work}} \quad (3)$$

However, while DCP seems a simple and straightforward metric, it is hard to know what each organization will consider as *useful work* for a data center, so this makes

it a subjective metric, unreliable when comparing different data centers.

CADE: McKinsey & Co. and the Uptime Institute introduced the Corporate Average Data Center Efficiency (CADE) metric to measure data center performance in a way that encompasses both IT and facilities' technologies. CADE is composed by four elements:

- *Facility energy efficiency:* how much of the power drawn from the electric grid by the data center is being used by IT equipment;
- *Facility asset utilization:* how much of the data center's maximum electrical capacity is in use;
- *IT asset utilization:* the average CPU utilization of servers in the data center;
- *IT energy efficiency:* this measurement has not been precisely formulated yet, but is intended to describe how efficiently servers, networking equipment and storage units use the power they are drawing to perform their functions

Combining the first two factors determines the efficiency of the facility; combining the second two determines the efficiency of the IT assets. Each factor is expressed as a percentage and then multiplied by the other:

- FE = Facility Energy Efficiency x Facility Utilization
- AE = IT Energy Efficiency x IT Utilization
- CADE = FE x AE

SI-POM: The Site Infrastructure Power Overhead Multiplier (SI-POM) determines the amount of overhead a data center consumes to power its critical IT equipment. It is defined as follows:

$$\text{SI-POM} = \frac{P_{DC}}{P_{IT}} \quad (4)$$

where P_{DC} expresses the data center consumption at utility meter and P_{IT} expresses the total hardware power consumption at the plug for all IT equipment. SI-POM, proposed by the Uptime Institute, is basically equivalent to the PUE, as it includes all the conversion losses in transformers, UPS, PDU and critical power distribution losses, as well as cooling systems, lights and other minor building loads. The main difference is that SI-POM explicitly mentions "overhead" rather than "efficiency", hence lower values are more intuitively linked to higher efficiency.

H-POM: The Hardware Power Overhead Multiplier (H-POM) determines the amount of power wasted in power supply conversion losses or diverted to internal fans, rather than in useful computing power. For a single device, H-POM is defined as follows:

$$\text{H-POM} = \frac{P_{AC}}{P_{DC}} \quad (5)$$

where P_{AC} expresses the Alternating Current hardware load at the plug and P_{DC} expresses the Direct Current hardware load before the power supply.

DH-UR: Deployed Hardware Utilization Ratio (DH-UR) determines which fraction of IT equipment is not running any application or handling important data. This metric is designed because most IT equipment is always switched on -unless specifically intended not to - regardless whether a given component is doing something important or not. As a result, such 'dormant' equipment can waste significant amount of power during their lifetime. DH-UR can be defined for both servers, as follows:

$$\text{DH-UR (servers)} = \frac{S_{live}}{S_{total}} \quad (6)$$

where S_{live} indicates the number of server running live applications and S_{total} indicates total number of deployed servers, or for storage elements as follows:

$$\text{DH-UR (storage)} = \frac{Data_{accessed}}{Data_{total}} \quad (7)$$

where $Data_{accessed}$ indicates the number of terabytes of storage holding frequently accessed data and $Data_{total}$ indicates the total terabytes of deployed storage.

DH-UE: Deployed Hardware Utilization Efficiency (DH-UE) is a dimensionless ratio that expresses the level of underutilization of servers and storage units. This metric is designed because, in non-virtualized environments, servers typically run a single application, using only 10-30% of their computing load - a situation commonly known as "server sprawl" [19]. Since servers running at low computation loads often draw nearly as much power as those running at high loads, a large number of such partly loaded servers can quickly consume valuable UPS and HVAC capacity and raise the electricity bill. When the DH-UE is significantly below 1 (the theoretical maximum), implementing virtualization is advised in order to increase the utilization rate. DH-UE for servers is defined as follows:

$$\text{DH-UE (servers)} = \frac{S_{min}}{S_{total}} \quad (8)$$

where S_{min} indicates the minimum number of server needed to handle peak computing load and S_{total} indicates total number of deployed servers.

B. Related Surveys

In 2012, the Uptime Institute conducted a survey regarding the importance and progress of energy efficiency on public facing companies that provide cloud services. The results of the survey showed that the progress on energy efficiency has been slowed in the last years and, is not considered a high priority. The survey was conducted on more than 1,000 global data center operators, however, only half of the North American respondents said they considered efficiency to be very important. Most operators answered that a 1.65 PUE (the average for the survey) is good enough [3].

In 2014, another survey from Uptime Institute within its Network members, indicated that nearly three fourths (72%) of respondents measure PUE, while 82% of the

executives report that metric to their corporate management. Moreover, the same survey tracked that PUE was improved from 2.5 in 2007 to 1.7 in 2014 [6].

Another survey in 2013, which was based on the US Data Center Infrastructure Management (DCIM) adoption, showed that only a 15 to 40 percent of data centers in the US are using DCIM. According to Uptime institute's survey, the biggest obstacle for adopting a DCIM solution is the cost, with 60% of the respondents asserting that its cost is too high. Two other issues for not adopting a DCIM solution were that, firstly, it is rarely planned for and many data center operators do not have a large budget for software implementation and secondly, because it is not clear who should pay for the IT or the facilities department [4].

One other important aspect is the ongoing improvement of energy efficiency of data centers. A survey conducted from Digital Realty Trust in January 2013 indicated that only 20 percent of the 300 North American data center companies with revenues of at least 1 billion and more than 5,000 employees have a PUE below 2.0 (average is 2.9). Any improvement below such threshold is hard to achieve and inelastic, requiring a higher expenditure in the budget which is hard to be justified from the management [5].

Virtualization is the most used technique for energy efficiency. A research from Gartner and VMware in 2012 estimates overall adoption rates to be 50-75%. However, despite the wide adaptation of virtualization, the server utilization rate from 2006 to 2012 has remained unchanged between 12 and 18 percent. This occurs due to a combination of factors, such as growth in hardware performance and capacity continuing to outpace the consolidation of workloads. Moreover, data center operators tend to take a cautious approach to consolidation ratios, mindful of unpredictable workloads and concerned about potential Service Level Agreement (SLA) violations as performance could be affected by resource contention [26].

IV. METHODOLOGY

In order to answer our research questions, we adopted qualitative empirical approach, by means of *semi-structured interviews* [24]. This method was chosen for multiple reasons. First, interviews have the potential to overcome the poor response rates of a questionnaire survey [8]. Secondly, in order to investigate the awareness of the interviewer on a specific subject, as noted by Bailey [9], it is crucial "*to ensure that the respondent is unable to receive assistance from others while formulating a response*".

In addition, a semi-structured interview allows a structured set of questions without compromising the flexibility of the conversation or interview's input freedom. It allows the interviewers to ask questions in order to receive the right information in case the answer of the interviewee is not satisfactory and, additionally, to react and follow up on emerging ideas and unfolding discussions.

A more practical reason for choosing a semi-structured interview was due to the busy schedule of the interviewees:

the chance of interviewing them a second time was highly unlikely, so according to Bernard [10] semi-structured interviews are best used in this occasion. The participation of interviewees was voluntary, i.e. no financial or any other kind of incentive was offered to them in order to participate.

A. Participants

In order to select our participants, we had to make sure that they would have the relevant experience and knowledge to answer the questions of the interview. The main requirement to select our participants was that they had to be involved in the power management process of their data centers. The age and gender of the participants were not considered relevant factors for the scope of this research.

Initially, an email was sent to the prospective participants including a consent form, which contained a small description of the study and the request to participate. In many cases, the email was forwarded from the prospective participants to specific colleagues that possessed certain knowledge. In total, 7 people have been interviewed from 4 companies. In order to minimize the effort required from the participants, one interview was conducted per company and the participants from the same company gave their input to the same interview.

1) *Company A*: a small-medium sized company with approximately 250-300 employees which provides IT solutions to their clients. The interview was conducted with the facilities' manager who is responsible for the facilities and the infrastructure parts of the data center as well as for the connectivity, cooling, security and related issues. In addition, the engineer responsible for the software developments, forensic analysis and applications, participated in the interview.

2) *Company B*: a large sized company that develops and provides a wide variety of hardware components as well as software and related services to consumers, small and medium sized businesses and large enterprises. It serves customers in the government, health and education sectors. Three individuals participated in the interview conducted, the Program/Policy Manager, the Global Environmental Manager and the Data Center Manager of operations who is also responsible for the operations in two data centers.

3) *Company C*: a small size company based in the Netherlands, employing approximately 10 employees, which provides hosting solutions. Company C does not own a data center but rents a space in another data center. The interview was conducted with the CTO of the company who is responsible for the technical, hardware and software implementations.

4) *Company D*: a medium sized company based in the Netherlands. It employs approximately 5.000 employees and provides software solutions, IT outsourcing and possesses its own data centers at various locations in the Netherlands. The participants of the interview were the facility manager of the data center and the manager

involved in the division of power in the equipment inside the facility.

B. Interviews

The duration of the interviews is 60 minutes in average, with some variations depending on the input, knowledge and enthusiasm of the interviewee. The interviews were conducted online (Skype, Lync, Hangouts) so as to provide the interviewer and the interviewees flexibility on scheduling, cost efficiency and reducing travel time, as many of the interviewees were in different cities. The disadvantage of this method is that the amount of information gathered from online or telephone interviews is less in comparison to face to face interviews. However, this was considered an acceptable trade-off considering the other benefits.

The questions were divided in two parts so that each part could deeply focus on each research question separately. The first part of the questions concerned the awareness of stakeholders regarding energy power measurements and it also focused on finding out who those stakeholders were and which were their concerns. The second part of the questions were mostly technical and focused in identifying the state of practice regarding energy efficiency techniques and metrics that are in use. A pilot interview was conducted in order to test the interview questions. After the pilot, some questions were redefined and rewritten to improve their clarity.

C. Data Analysis

For the data analysis, we adopted Data Display, an analysis procedure based on [22]. This procedure consists of three key activities: data reduction, data display and drawing conclusions.

Data Reduction: the first step of the data display analysis, where the mass of qualitative data is reduced and organized. The first step of this process was to create codes [22] from the interview transcripts which aided in creating summaries of the interview data and to get an overview of the main topics during the interviews.

Data Display: in this step, the findings are organized in matrices, in order to visualize differences and match concepts among different groups. In particular, matrices were created to match different energy efficiency practices with the state-of-the-Art and determine the level of awareness of the stakeholders.

Interpreting results and drawing conclusions: in the last step of data display process, after processing and reviewing the data of the first and second step, the results are interpreted and conclusions are presented regarding the research questions of this paper.

V. RESULTS

A. State-of-practice of energy efficiency and power measurements

From the interview summaries, we extracted the following data: a) where is power actually measured (measurement points); b) in what metrics are these measurements

Measurement Points	Incoming Electrical Power
	PDU
	UPS
	Rack Server
	Transformer
	Switch Boards
Metrics	Distribution Boards
	PUE
Green IT Practices	Consolidation
	Virtualization
	Participating to organizations (e.g. EU COC)
	Efficient Cooling
	Energy Management Systems
	Monitoring Server Utilization
	ISO Certifications
	Hardware Upgrades
	Demand Management
	Workload Estimation
	Dynamic Power Scaling
	Recycling Equipment
	Increase Density
	Renewable Energy Sources
	Switch off unused devices
	Energy-efficient hardware
Green IT budget	
Dedicated Facilities Manager	

TABLE I
PRACTICES, MEASUREMENTS AND METRICS EXTRACTED FROM THE INTERVIEWS

used; c) which green IT and energy efficiency practices are in use by the organization. This data is summarized in Table I.

In Table II we show which practice is adopted by each of the companies we interviewed. A checkmark (✓) indicates that the practice is fully implemented in the day-to-day activity of the company. A bullet (●) indicates that the practice was used but is not being used anymore, because no significant change was observed or it was hard to manage.

In Section III we introduced a classification for energy

Practices	Companies			
	A	B	C	D
Consolidation	✓	✓	✓	✓
Virtualization	✓	✓	✓	✓
Efficient Cooling	✓	✓	✓	✓
Energy Management Systems	✓	✓	✓	✓
Monitoring Server Utilization	✓	✓	✓	✓
Demand Management	✓	✓	✓	✓
Workload Estimation	✓			
Dynamic Power Scaling	●			
Renewable Energy Sources		✓	✓	
Switch off unused devices	●	●		✓
Energy-efficient hardware	✓	✓	✓	✓
Green IT budget	●	✓		
Dedicated Facilities Manager	✓	✓		✓
Participating to organizations	✓	✓		
ISO Certifications		✓		
Hardware Upgrades		✓		✓
Recycling Equipment		✓		
Increase Density	✓			✓

TABLE II
PRACTICE ADOPTION AMONG COMPANIES

measurements: simple, detailed and IT [12]. We matched this classification with the actual measurements taken by the companies we interviewed. The matching is provided in Table III. In the Table, a bullet (●) indicates measurements not currently in use, but that interviewees plan to adopt in the very near future.

B. Stakeholders and Level of awareness

Through the interviews, it was crucial to understand and extract who the stakeholders are and what are their concerns and interests. These are presented in Table IV:

- *Environmental Managers*: they have to ensure that the company is aligned with the environmental policies and current regulations and legislation. They are also responsible for acquiring the required certification for their products. From the interviews conducted, it can be argued that environmental managers can be found mostly in big companies, for which the regulations and procurement policies are more complex. The concerns of the environmental managers are to always try to meet the government regulations and policies and to be informed about any possible alteration that may occur in these regulations.
- *Data center Managers*: responsible for data center operations. In addition, they have to ensure continuous availability of resources and the data center energy efficiency. They receive all the data, reports and the measurements in the data centers and make the required decisions for the data center to be energy efficient. The concerns of the data center managers are to monitor the energy in order to prevent low performance and downtime on the machines and to maintain the energy cost to a limit. Data centers managers existed in all companies that were involved in the interviews and that own data center facilities.
- *Infrastructure managers*: from the interviews it was noted that depending on the size of the company, whether it is small or big, there are two options for the infrastructure; when the company is big, infrastructure managers are typically outsourced consultants, while when the company is small they are internal employees. Their role is to ensure the quality, functionality and efficiency of all the installed components in the data center. They are also responsible for the installation of the monitoring solution and the capacity and growth planning of the data center.
- *Account Teams*: during the interviews, some of the interviewees referred to them as account teams and other as administration teams. Their role is to provide support to customers. Their interests are to ensure customer satisfaction and to provide them with the suitable IT solution based on their needs. They are also responsible for the monthly bills that are sent to customers. Similarly to the infrastructure managers, the administration teams/accounts teams were found in companies that had their own data center.

Company	Simple		Detailed				IT			
	Incoming	Electrical Power	Cooling	Transformers	UPS	PDU	Rack	Server	Device	VM
Company A		✓	✓	✓	✓	✓	✓	✓	•	•
Company B		✓	✓	✓	✓	✓	✓	✓		
Company C		✓	✓	✓	✓	✓	✓	✓	•	•
Company D		✓	✓	✓	✓	✓	✓	•		

TABLE III
ENERGY MEASUREMENT LEVEL AMONG COMPANIES

Stakeholders	Role	Interests	Concerns
Environmental Managers	Align with all environmental policies, legislation and current regulations	Meeting environmental regulations and policies	Failing to meet regulations and policies, complexity on procurement policies, new legislation
Data Center Managers	Manage data center operations and ensure the availability and energy efficient usage of resources	Managing resources as energy efficiently as possible, lower costs, keep continuous availability, monitor energy consumption	Low performance, downtimes, excessive energy costs
Infrastructure Managers	Ensure the functionality, quality and efficiency of all the installed components in a data center	Improve the infrastructure, develop new energy management and monitoring solutions	Continuous availability, power capacity, growth
Account Managers	Provide customer support	Customer satisfaction, provide IT solutions to customers based on their needs, Quality of Service	Miscalculation of customer needs, miscalculation of bills, failing to provide the right IT solutions to customers
Suppliers	Provide resources (e.g. power, hardware)	Keep providing resources and renewing contracts with data center companies	Failure to provide the requested resources, termination of contract
Governmental Organization	Regulate data center operation regarding power and carbon emissions	Ensure that regulations are followed	Violation of regulations
Non-Governmental Organization	Support the industry and provide standards towards energy efficiency	Setting standards and metrics for energy efficiency	Fail to establish widespread knowledge and usage about their standard and metrics
Customers	Profit for the company	Continuous availability, value for money, satisfaction	Loss of data, loss of performance, extra costs

TABLE IV
MAIN STAKEHOLDERS AND CONCERNS

- *Suppliers*: suppliers have the responsibility to provide continuous resources to data centers. These could be suppliers for energy or suppliers for hardware which are subject to a contract with the data center's company. Their scope is to continue providing resources to companies according to company's needs.
- *Governmental Organizations*: Specific governmental bodies that control and apply regulations to the data centers' operations. They define thresholds in power and emission and they are also responsible for collecting taxes.
- *Non-Governmental organizations*: entities such as the EU COC and RenewIT which are trying to improve the conditions in energy efficiency of the data centers. They are trying to accomplish that by attracting members to their organizations and by establishing several standard operations and metrics which in turn will enhance the evaluation of the data centers with regards to energy efficiency.
- *Customers*: customers are regarded as the most important stakeholders. Their ultimate goal is that the services they receive from the data centers meet their needs and that their business is continuously available without any downtimes and with the maximum performance. In addition, any possible miscalculation of the energy consumed by the customers from the account

and administration teams, will affect them from an economic point of view, either positively or negatively.

VI. DISCUSSION

In this section, the outcomes of the results will be further analyzed and discussed in relation to our RQs.

A. What is the state-of-the-practice regarding energy efficiency and power measurements in the data center industry?

According to Table II, the most widely used energy efficiency practices are Consolidation, Virtualization, Cooling, Demand Management and Energy-efficient Hardware. Among the companies there are several similarities and differences:

- *Virtualization*: all the interviewees mentioned that their environments are virtualized, which allows them to decrease physical servers utilization.
- *Consolidation*: some of the companies, such as company A, are more supportive of the value of consolidation. For instance, company A managed to achieve the utilization of 60-70%. In order to accomplish this, they took several measurements and calculated the optimal memory vs CPU balance, allowing them to install up to 80 VMs into one physical server.

- Cooling: managing cooling was mentioned as the most crucial means of energy efficiency of the data centers. Companies B and D mainly use outside air which, after filtering, enters the data center. This gives them the ability to use free cooling for several days during the year, when the outside temperature is very low. Company A follows a different path; it stores the cold air that is being produced underground, so as to use it even when the outside temperature is high. Thanks to this practices, they claim to have decreased cooling energy by 86%.
- Demand management: all the companies are able to identify the low and high peak demands to be sure that they have the necessary energy when the high peak demand occurs. On the other hand, most of the companies do not recognize any significant differences between the low and high peak demands during the day. The interviewees mentioned several reasons for this. Company C stated that during the night, when demand decreases, backups start working, which in turn have a great impact on the consumption of resources. Company A mentioned that they are able to reach a “flat” demand curve thanks to an extensive use of consolidation.
- Energy-efficient hardware: according to the interviewees, all companies purchase energy-efficient hardware. In addition, measurements allow them to determine which hardware is more energy-efficient, and tune their future hardware procurement strategy accordingly.

In addition, Company A mentioned that they used Dynamic Power Scaling on some of the hardware components but this did not bring the desirable results, hence their decision not to adopt such practice in the future.

Which power measurements are taken and how: the measurements used are shown in Table III. Here it is worth noting that all companies take measurements up to Rack level. After the Rack level, taking measurements becomes more difficult. This occurs because when using more fine-grained measurements, the amount of data increases abruptly and it is then difficult to manage all the information gathered. Companies A, B and C perform measurements at server level, but none of them measures at the level of virtual machines.

Usage of appropriate indicators or metrics regarding the overall data center energy efficiency: all companies mentioned that they use such measurements in order to extract the PUE value and assess the overall data center energy efficiency. However, there was no specific common technique from all the companies in the way they calculate energy efficiency. Moreover, the information was limited because many data centers receive the PUE value from external companies or they are not responsible for the PUE’s calculation. Other reporting techniques include graphic solutions (e.g. dashboards) and reports with the raw data.

Apart from PUE, no other energy efficiency metrics such as those we identified in the literature review were mentioned. All the companies have a PUE of 1.2-1.5 which is considered to be a great score. On the other hand, most of our interviewees expressed doubt about the reliability of the PUE. This leads to questioning whether PUE is indeed an important metric or its role and goal has reversed to a commercial/advertising one. For this reason, the use of additional metrics is considered to be very important. From the interviews, it can be argued that many measurements are used but we could agree with one of the companies that stated that the measurements are good but they should be used in a way that is understandable for all and it should facilitate decision making. So, from the moment that the measurements exist, the proper metrics should be properly introduced in order to enhance the evaluation purposes. Taking a significant amount of measurements without using them is a waste of time and resources.

B. What is the level of awareness of the stakeholders in the IT industry towards power measurements?

Identified stakeholders and concerns: in Table IV we present the stakeholders identified from our interviewees. A difference can be observed between big and small companies. The big companies tend to outsource many tasks and this results in them having less control over the infrastructure and operations of their data center. Because of this, some of the interviewees seemed to be unaware of several tasks and urged us to contact other individuals in order to find the required information.

The EU COC and the other non-government organizations dealing with the energy efficiency have a crucial part to play. The companies with their own data centers were members of the EU COC or other similar organizations. Some of the companies recognized the role of EU COC as important, while others expressed their displeasure because in these organizations, they mentioned that there is a lot of discussion but no actions taken.

If this is true, then a question arises as to what extent these organizations reach their goal or if they stand as an obstacle to that goal: some endorsers might join such activities and take place in the discussions purely for ‘greenwashing’ reasons and as a mean to avoid governmental regulations.

Usefulness of power measurements according to the stakeholders: all of our interviewees considered taking measurements to be useful, mostly for two reasons:

- Strategic decision-making: some examples of strategic decisions that the companies stated are whether further consolidation is needed and whether hardware needs to be upgraded.
- Quick response: some companies are using power measurements in real time, to be able to make quick decisions when something goes wrong. For instance, company D has monitors inside the facilities which allows them to see the power measurements at rack

level. If something wrong is observed, they can quickly proceed to different actions.

An important issue is the measurements in resources that are shared among different stakeholders, i.e. colocation. Progress has been made in other sectors such as cooling, for which great research has been conducted, leading to generation of positive outcomes, as also seen in the interviews. However, there should be more actions towards the next level which is the measurements in shared IT components, especially when dealing with virtualization. The barrier towards that direction is the lack of suitable software solutions: as VMs are software entities that can be freely moved across a data center, determining their specific energy impact, according to our interviewees, is a complex problem. This is an issue that demands special attention in the coming years as this will allow significant energy savings.

VII. THREATS TO VALIDITY

In order to mitigate threats to validity in our research, we adhered to the following criteria [18]:

- 1) During the interview, the interviewer tried to be open minded and to avoid showing any possible emotions or reactions and most importantly, to not express any personal opinion so as to not mislead the interviewees.
- 2) The interviewer avoided to ask the questions in a way that would drive the interviewee towards specific answers.
- 3) The meetings were scheduled whenever it was convenient for the interviewees, in order to avoid any feelings of pressure and for the interviewee to be relaxed.

Furthermore, the validity of this research is threatened from biased opinions of the interviewees [14]. The interviewees were one or two per interview and for this reason, it does not mean that their opinion represents that of the whole company. However, as in this research we are investigating the awareness of the stakeholders towards a specific topic, we consider their perceptions and personal bias as part of the results. Additionally, according to [14], the transcripts were sent back to the interviewees in order to avoid any possible misinterpretation of their statements by the interviewer which can in turn cause biased data.

One of the four companies that participated in the interviews does not own a data center which is why only information about the racks it rents could be gathered.

An additional limitation in the results was the low number of the companies participating. Big effort was given to approach various companies but without positive outcome, due to lack of time and availability. Specifically, from the seven companies approached, only four of them accepted the invitation for an interview. Finding these participants and successfully conducting the arranged interviews were

both time consuming processes. In addition, the busy schedule of the interviewees was another obstacle, as it prevented us to perform follow-up interviews.

VIII. CONCLUSION AND FUTURE WORK

The aim of this research is to investigate the current state-of-practice regarding energy efficiency practices and power measurements in data centers. The first research question is stated as: “*What is the state of practice regarding energy efficiency and power measurements in the data center industry?*” For the period examined, the data centers were using several energy efficiency practices, allowing them to achieve PUE values between 1.2-1.5. To be more specific, the most widely used techniques were virtualization, consolidation, managing low and high peak demands, using eco-hardware and most importantly cooling, where many of the companies invested most resources. Regarding the power measurements, all of the companies take measurements until Rack level and use graphic and monitoring solutions. After the rack level, taking measurements becomes difficult and measuring power at VM or application level is still a challenge. Moreover, the only efficiency metric being used in order to see the overall data center energy efficiency was PUE. In addition, there were many doubts regarding PUE’s effectiveness, something which was clearly demonstrated from some of the interviewees and also from the literature review.

The second research question is stated as: “*What is the level of awareness of the stakeholders in the data center industry regarding power measurements?*” From the interviews conducted we identified several stakeholders: environmental managers, data center managers, infrastructure managers/outside consultants, government organizations, non-governmental organizations, suppliers, administration teams/account teams and customers. In addition, it was observed that the small companies have complete control over their facilities, without having any outside consultants, which in turn gives them a better overview of the possible improvements they could make. The IT specialists consider the measurements useful because they enable better decision-making processes in strategic as well as immediate decisions. However, a lack of awareness was observed with respect to advanced measurements and metrics at IT level.

Several topics emerged from this work and deserve further research efforts. The first topic concerns the role of non-profit organizations such as the EU-COC. More precisely, the importance and relevance of these organizations is still subject to debate in industry, and more transparency in their processes would be helpful. However, their role is crucial in collaboration with regulators, to establish standards and metrics that will have to be enforced.

In order to ease the adoption of Green IT best practices, we developed a number of tools [15], [21] that provide insights and estimations on their economic and environmental impacts. We also have an online library of such practices⁴

⁴<http://greenpractice.few.vu.nl/>, last visited on July 13, 2016

that can be used as a knowledge base for practitioners who want to have a concrete overview of actionable strategies for reducing their carbon footprint – and energy bill.

Another important topic is power measurements in software components. This topic will be more and more relevant in the future, since the improvement of hardware and infrastructure energy efficiency has probably reached a tipping point, but energy costs are still a major concern.

ACKNOWLEDGEMENT

This work has been supported by the RAAK-MKB innovation fund from the Stichting Innovatieve Alliantie (SIA), the Dutch National Taskforce for Applied Research, as part of the Greening the Cloud project, and by the Netherlands Enterprise Agency (RVO) as part of the GreenServe project.

REFERENCES

- [1] Data center power efficiency metrics: PUE and DCiE. Technical report, The Green Grid, 2007.
- [2] Code of conduct on data centres energy efficiency (version 2): Endorser guidelines and registration forms. Technical report, European Commission, 2009.
- [3] Data center industry survey. Technical report, Uptime Institute, 2012.
- [4] Data center industry survey. Technical report, Uptime Institute, 2013.
- [5] North american campus survey results. Technical report, Digital Realty Trust, 2013.
- [6] Data center industry survey. Technical report, Uptime Institute, 2014.
- [7] D. Alger. *Grow a greener data center*. Pearson Education, 2009.
- [8] E. K. Austin. *Guidelines for the Developing of Continuing Education merings for Nurses*. Appleton-Century-Crofts., 1981.
- [9] K. Bailey. *Methods of social research*. Simon and Schuster, 2008.
- [10] H. R. Bernard. *Research methods in anthropology: Qualitative and quantitative approaches*. Rowman Altamira, 2011.
- [11] G. A. Brady, N. Kapur, J. L. Summers, and H. M. Thompson. A case study and critical assessment in calculating power usage effectiveness for a data centre. *Energy Convers. Manage.*, 76:155–161, Dec. 2013.
- [12] R. Brown et al. Report to congress on server and data center energy efficiency: Public law 109-431. *Lawrence Berkeley National Laboratory*, 2008.
- [13] R. Brown et al. Report to congress on server and data center energy efficiency: Public law 109-431. *Lawrence Berkeley National Laboratory*, 2008.
- [14] T. Daim, J. Justice, M. Krampits, M. Letts, G. Subramanian, and M. Thirumalai. Data center metrics: an energy efficiency model for information technology managers. *Management of Environmental Quality: An International Journal*, 20(6):712–731, 2009.
- [15] H. Fernández, G. Procaccianti, and P. Lago. Economic aspects of green ICT. In *Green in Software Engineering*, pages 107–127. Springer International Publishing, 2015.
- [16] E. Gelenbe and Y. Caseau. The impact of information technology on energy consumption and carbon emissions. *Ubiquity*, 2015(June):1:1–1:15, 2015.
- [17] C. Gough, I. Steiner, and W. Saunders. Data center management. In *Energy Efficient Servers: Blueprints for Data Center Optimization*, pages 307–318. Apress, Berkeley, CA, 2015.
- [18] S. A. Hutchinson, M. E. Wilson, and H. S. Wilson. Benefits of participating in research interviews. *Image: The Journal of Nursing Scholarship*, 26(2):161–166, 1994.
- [19] G. Khanna, K. Beaty, G. Kar, and A. Kochut. Application performance management in virtualized server environments. In *Network Operations and Management Symposium, 2006. NOMS 2006. 10th IEEE/IFIP*, pages 373–381. ieeexplore.ieee.org, Apr. 2006.
- [20] J. Liu, F. Zhao, X. Liu, and W. He. Challenges towards elastic power management in internet data centers. In *Distributed Computing Systems Workshops, 2009. ICDCS Workshops '09. 29th IEEE International Conference on*, pages 65–72. IEEE, 2009.
- [21] K. Lundfall, P. Grosso, P. Lago, and G. Procaccianti. The green practitioner: A Decision-Making tool for green ICT. In *Proceedings of 29th International Conference on Informatics for Environmental Protection (EnviroInfo 2015) and Third International Conference on ICT for Sustainability (ICT4S 2015)*. Atlantis Press, 2015.
- [22] M. B. Miles and A. M. Huberman. *Qualitative data analysis: An expanded sourcebook*. Sage, 1994.
- [23] L. Minas and B. Ellison. *Energy efficiency for information technology: How to reduce power consumption in servers and data centers*. Intel Press, 2009.
- [24] I. Seidman. *Interviewing as Qualitative Research: A Guide for Researchers in Education and the Social Sciences*. Teachers College Press, 2013.
- [25] J. Shuja, K. Bilal, S. A. Madani, and S. U. Khan. Data center energy efficient resource scheduling. *Cluster Computing*, 17(4):1265–1277, 2014.
- [26] J. Whitney and P. Delforge. Scaling up energy efficiency across the data center industry: Evaluating key drivers and barriers. Technical report, Natural Resources Defense Council (NRDC) - Anthesis, 2014.