



UvA-DARE (Digital Academic Repository)

Greening the Cloud

van den Hoed, R.; Hoekstra, E.; Procaccianti, G.; Lago, P.; Grosso, P.; Taal, A.; Grosskop, K.; van Bergen, E.; Center for Applied Research Technology

Publication date

2016

Document Version

Final published version

[Link to publication](#)

Citation for published version (APA):

van den Hoed, R., Hoekstra, E., Procaccianti, G., Lago, P., Grosso, P., Taal, A., Grosskop, K., van Bergen, E., & Center for Applied Research Technology (2016). *Greening the Cloud*. (Technology Knowledge Center of the Amsterdam University of Applied Sciences series; No. 09). Amsterdam University of Applied Sciences, Faculty of Technology.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.



**Amsterdam University
of Applied Sciences**

CENTER FOR APPLIED RESEARCH TECHNOLOGY

GREENING THE CLOUD



Robert van den Hoed
Eric Hoekstra
Giuseppe Procaccianti
Patricia Lago
Paola Grosso
Arie Taal
Kay Grosskop
Esther van Bergen



GREENING THE CLOUD

Earlier publications from the Technology Knowledge Center of the Amsterdam University of Applied Sciences series:



01 Vertical farming
Technologie en innovatierichtingen voor de toekomst



02 Duurzaam bewaren
Simulatiemodel en technologieën voor energiesbesparing



03 Extreme neerslag
Anticiperen op extreme neerslag in de stad



04 Beter beheer met BIM
Van informatiemodel naar informatiemanagement



05 Het stedenbouwkundig bureau van de toekomst
Spin in het web



06 (Terug)schakelen naar ketendenken
Innovaties realiseren bij logistiek MKB in mainports



07 Maintaining your competitive edge
Planesense process improvement in aviation maintenance



08 Biobased plastic
Unieke eigenschappen en inspirerende toepassingsmogelijkheden

PUBLICATION SERIES Amsterdam University of Applied Sciences (AUAS) FACULTY OF TECHNOLOGY

In this publication series, the faculty of Technology of the Amsterdam University of Applied Sciences bundles the results of applied research. The publications are focused on professionals and provide access to knowledge and expertise obtained via applied research of the AUAS in the metropole region Amsterdam. This publication series gives the reader tools to realise improvement and innovation in the engineering profession.

Faculty of Technology

The faculty of Technology of the Amsterdam University of Applied Sciences is the largest engineering college institution of the Netherlands. The faculty consists of eight engineering training courses with varied learning routes and specializations. The range of courses is diverse, from Engineering to Logistics, from Built Environment to Forensic Research and from Maritime Officer to Aviation.

Research plays a central role in the faculty of Technology. This research is rooted in professional practice and contributes to the continuous improvement of the quality of education and innovations in practice. The practice-oriented research of the HvA has three functions:

- Developing knowledge
- Innovation of the professional practice and the professional society
- Innovation of educational programmes

The faculty of Technology has 3 research programmes, all closely connected to a part of the educational programmes:

1. Aviation
2. Forensic Research
3. Urban Technology

Editorial

The publication series are published by the faculty or Technology of the AUAS. The editorial board consists of professors of this faculty. Every publication has its own authors, consisting of AUAS employees, sometimes supplemented with representatives of companies and other knowledge institutions.

COLOPHON

Publication:

Amsterdam University of Applied Sciences

Faculty of Technology

Urban Technology

www.hva.nl/urbantechnology

Editorial board:

Robert van den Hoed

Eric Hoekstra

Giuseppe Procaccianti

Patricia Lago

Paola Grosso

Arie Taal

Kay Grosskop

Esther van Bergen

Translation:

Splend!

Lay out/infographics:

Nynke Kuipers

Print:

SDA Print+Media, Amsterdam

© Faculty of Technology, Amsterdam University of Applied Sciences, July 2016,



CONTENTS

1 Introduction Greening the Cloud	11
1.1 Why Greening the Cloud?.....	11
1.2 Knowledge gap.....	13
1.2 Goal and approach.....	14
1.4 Partners in this project.....	15
2 Greening of the ICT sector	17
2.1 Electricity consumption of the ICT sector	17
2.2 Data centers in Amsterdam.....	20
2.3 Energy use of data centers	20
2.4 Energy loss chain	24
2.5 Conclusions.....	26
3 Software Energy Footprint Lab	31
3.1 System requirements	31
3.2 Details of the first measurement system.....	32
3.3 Measuring modern servers.....	32
3.4 Recent measuring components	33
3.5 Development challenges	34
3.6 Conclusions and further research	41
4 GreenCloud Case Studies	43
4.1 Introduction to cases.....	43
4.2 Case studies regarding virtualisation	44
4.3 Power estimation models.....	49
4.4 Software engineering practices.....	55
4.5 Conclusions.....	63
5 Standing on the shoulders of PUE	65
5.1 Introduction	65
5.2 PUE strengths and limitations.....	65
5.3 GreenCloud Model	67
5.4 Industry benchmarking: making sense of numbers	70
5.5 Conclusions.....	71
6 Knowledge valorisation	73
6.1 Development of a cost module	73
6.2 Translation into education.....	82
6.3 Amsterdam Sensor Lab.....	83
7 Conclusions and Next Steps	87
7.1 Main project findings	87
7.2 General observations.....	88
7.3 Collaboration Agenda for Green Software and Clouds	90
Acknowledgements	92
Literature list	94



1 INTRODUCTION GREENING THE CLOUD

1.1 Why Greening the Cloud?

The cloud has become an essential part of our daily lives. We use it to store our documents (Dropbox), to stream our music and films (Spotify and Netflix) and without giving it any thought, we use it to work on documents in the cloud (Google Docs). The cloud forms a massive storage and processing system for an increasing amount of our digital activities. Not only consumers, but companies too are using the cloud more and more for the storage of data and to host software applications in. Behind the virtual cloud is a world of physical data centers, connected globally, packed with servers that enable us to store documents, use processing power and send data regardless of where we are or which device we use (as visualised in Figure 1.1 - Cloud infrastructure enables flexibility in our work and private lives). These data centers consume a large amounts of energy, comparable with the amount consumed by the aviation industry.

Although data centers work on saving energy, in particular by means of efficient cooling techniques and energy efficient servers, the growing demand for data center capacity makes it hard to realise absolute energy savings. Whereas on the one hand servers are rapidly becoming more energy efficient, this effect is cancelled out by software that consumes increasing amounts of energy. It is expected that as a result the energy consumption of (new) data centers will keep rising in the future. Solutions with energy efficient cooling techniques and renewable energy are insufficient. Most energy savings of hardware can be achieved by looking at the source of the energy load: the software that runs on the hardware. After all, it is the software that controls the processors, the memory, the network ports and all other components of the server.



Figure 1.1 Cloud infrastructure enables flexibility in our work and private lives.

This report is the result of a research effort which was focused on supporting companies that are active in the cloud and desire to make smarter, energy efficient choices for configuring their cloud. There is much to gain from energy efficient software, with smart choices about how cloud environments are designed and which applications and features are used. However, the industry requires more knowledge about the measures that can be taken to realise energy savings. In addition, there is little scientific knowledge about the guidelines and tools that can be used in the field to make clouds greener. An important ambition of this project was to develop practical knowledge by means of carrying out concrete experiments, in order to make clouds more energy efficient. In doing so, the researchers hope to contribute to a more energy efficient ICT infrastructure in the future, while doing more research into green software and energy efficient clouds.

During its course of two years, this project has contributed the following results for use in practice, science and education:

- 10 case studies with concrete results about energy effects of design decisions in software or cloud environments.
- The GreenCloud Model, which provides a broader vision on the factors that need to be considered to assess a cloud environment for its energy consumption.
- A multi-server platform, able to measure the energy consumption of cloud applications and cloud environments in a controlled environment. It serves as testing bed for companies as well as a lab facility for students to create more awareness for the energy consumption of software.
- A cost model that translates energy savings into cost effects for organisations.
- Increased awareness in industry about the possibilities to achieve (significant) energy savings by making smart design choices.

1.2 Knowledge gap

The Metropolitan Region Amsterdam (MRA) is one of the largest Internet hubs of Europe and hosts more than 35 big data centers and a large number of hosting companies. A lot of international ICT companies are represented here, partly thanks to great digital connections via the Amsterdam Internet Exchange (AMS-IX) and the reliable energy supply. In the Amsterdam region, over 19,000 companies are active in the ICT sector. Given its large energy footprint, one of the challenges for the industry is to showcase green practices. Apart from low hanging fruit such as efficient cooling technologies that data centers can apply, the industry has limited knowledge how to green cloud environments. This research focuses on energy efficiency measures that small and medium-sized enterprises (SME-companies) active in the "cloud" can apply, the so-called cloud service providers.

Cloud service providers

Cloud service providers are companies that offer connectivity, storage and processing capacity, applications and services via Internet to both companies and individuals. But there are also hundreds of SME-companies in the Netherlands that offer services, processing capacity and software applications to customers via Internet. Two important parties who provide cloud related services are so-called (i) hosting providers and (ii) cloud application developers. Hosting providers offer computing capacity, hosted in one or more data centers, to all kinds of customers. For instance, to companies that want to build a website or web shop. Cloud application developers are software developers who develop applications that run in the cloud and that are used by individuals or companies via the cloud.

Hosting providers are confronted with concerns about the rising energy consumption in the ICT chain. The energy cost can rise up to a third of the overall operating cost, which offers significant opportunities for cost savings. Sessions with hosting providers have also shown that they see opportunities to distinguish themselves with a green product, and there is a small but growing market with customers (mainly municipalities and NGOs) that prefer to host their ICT services at a 'green' hosting company. Hosting providers currently do not have sufficient knowledge about the possibilities to configure their cloud environment more energy

efficient. Although in literature, there are (limited) guidelines for energy efficient design strategies, these are usually not validated. In addition, programmers do not have an effective and accurate measuring platform to test design guidelines or software versions in a controlled environment. Typical questions of these companies include: How can I configure my cloud environment as energy efficient as possible? What variables are most likely to reduce the energy consumption of the cloud and applications? Hosting providers would truly reap the benefits from guidelines to configure their cloud more energy efficient.

Cloud application developers develop software that runs in cloud environments. For instance, software to apply for permits online. There are hundreds of software developers in the Netherlands, however they all have limited knowledge about the consequences of certain software architectures and solutions for energy consumption. Performance and speed of the software are usually the predominant concern when software is designed. The impact this has on the energy consumption of the software is usually not given any consideration. Most developers do not have any knowledge about the relationship between performance and energy consumption. Software developers would benefit from establishing what the energy footprint of their software is as an additional quality aspect of software, which could support them in distinguishing themselves from competing software firms.

Software Energy Footprint Lab (SEFLab)

To make objective energy measurement of software possible, the Amsterdam University of Applied Sciences (HvA), in collaboration with the Software Improvement Group (SIG), founded the Software Energy Footprint Lab (SEFLab) in 2012. In this physical measuring lab, various servers are equipped with measuring sensors, so that very accurate energy footprints of software can be calculated in a controlled environment. In several years, the SEFLab has gained knowledge about the deployment and development of the right sensors to be able to measure all energy consuming components in servers. For the project that is described in this report, the challenge was in modelling a cloud environment, for which a network of clients and servers had to be interconnected (a so-called client-server setup) including an effective data acquisition system and a high level of accuracy.



Figure 1.2 Students at work in SEFLab.

1.3 Goal and approach

The project was focused on developing practical knowledge about the energy consumption of cloud environments and cloud applications. Within this project, this was translated to carrying out 10 case studies. In these studies, concrete questions from SME companies (both hosting companies and cloud application developers) were translated into experiments in which different configurations or options can be compared. A necessary requirement thereto was the availability of a professional measuring platform to measure the energy consumption of cloud environments and applications in a controlled environment. To this end, the SEFLab was deployed and made suitable to measure energy consumption of cloud environments and applications.

Specific challenges for the project

From the start, the context of the project provided its particular challenges that have shaped the activities and detailed goals of the project:

Priority: Although there has been growing interest, greening of software or cloud environments still remained a relatively low priority for a large share of the industry (compared to, for instance performance, cost reduction, uptime). Saving energy largely translates into efficient cooling rather than looking at software. This project has researched whether green software could become a growing factor in industry.

Awareness: Many software developers, systems integrators, data centers or end users had limited insight in the actual energy consumption of their software, cloud or ICT operations. This project aimed to increase this level of awareness.

Complexity: Measuring energy performance of clouds is much more complex given that clouds use distributed networks (rather than one particular server), with difficulties in allocating particular processes to resources and therefore to energy usage. This project aimed to harness part of this complexity through analysing key indicators of energy usage.

Energy costs: Low energy costs limits the potential cost savings that can be achieved through applying greening software or cloud strategies. This project focused on hosting providers to establish whether their relative high energy bill could provide an incentive to invest in greener design strategies.

Metrics for Green Clouds: Metrics for the efficiency of datacenters have been insufficient in measuring and monitoring the energy performance of cloud environments. This project aimed to develop more accurate metrics that entail a broader part of the complete stack.

Split incentives: Split incentives amongst key players in the ICT value chain has limited the uptake of energy and cost saving measures. Software developers were

rarely requested or incentivized to adhere to energy performance criteria. End clients commissioning software development had limited awareness about energy costs of their ICT operations and did not steer on energy performance. Hosting providers had limited influence on software choices of their clients, but also usually deployed a business model that stimulates selling more computational resources, with limited incentives to stimulate green design strategies. In this project we researched opportunities to bring incentives between stakeholders more in line.

Research methodology

The research plan consisted of a 'mixed method' approach, in which literature/benchmark studies (establishing the state of art), case studies and experiments (preparing and implementing tests to establish energy footprints), interviews for the development of a cost module and GreenCloud Model were utilised. Per case / SME, the predefined software variables that could easily be manipulated and measured for all cases were studied. These variables were translated into experiments, and tested within the measuring platform in the SEFLab.

1.4 Partners in this project

Within this project, there has been a close collaboration with companies, knowledge institutions and trade associations. Coordinator was the Amsterdam University of Applied Sciences (HvA), in the form of a collaboration between the Urban Technology research programme of the Technology Faculty and the CREATE-IT research programme of the Digital Media and Creative Industry Faculty.

Academic input came from the research group Software and Services of the Vrije Universiteit Amsterdam (VU) and the research group System and Network Engineering of the University of Amsterdam (UvA). VU contributed its knowledge and expertise in the field of green design strategies for software and made a significant contribution to the design and implementation of a number of experiments within the Green Lab¹. UvA contributed its knowledge and expertise about the configuration of hardware platforms, network configurations and virtualisation techniques and carried out a number of case studies within the SEFLab. Both partners contributed greatly to the development of the GreenCloud Model.

In total, 10 SME-companies have been actively involved (See Acknowledgements) in this project, namely Software Improvement Group, Greenhost, REM Automatisering, CobraSystems, Schuberg Philis, Almende, CloudProvider, Diesveld Query Technology, VKA, Transfer Solutions and GreenInvents. They provided the knowledge questions that led to the implementation of practice-relevant case studies.

The trade associations ISPCconnect (hosting providers), Nederland ICT (ICT sector) and Green IT Amsterdam Region (a.o. data centers) have deployed their network to gather knowledge questions, to organise events and they have shared results of the project to their supporters.

1. www.s2group.cs.vu.nl/green-lab/3.



2 GREENING OF THE ICT SECTOR

The ICT sector is a major consumer of electricity. An estimated 2% of the global power consumption is at the expense of the ICT sector, comparable to the power consumption of the aviation industry. The ICT sector is making significant investments in energy savings resulting in a clear decoupling of data growth and energy consumption. But with the expected growth of Internet of Things and large-scale (big) data generation and storage, the question is how the power consumption of the sector will develop in the (near) future. Where greening of ICT is currently mainly being implemented by using innovative cooling and energy efficient hardware, it must also be investigated what opportunities energy efficient cloud environments and energy efficient software code can provide in this regard.

2.1 Electricity consumption of the ICT sector

Energy use by ICT in the Netherlands

CE Delft recently made an analysis of the overall energy use of ICT in the Netherlands (CE Delft 2016) and translated the trends towards 2030. In total, the electricity consumption of ICT in the Netherlands was 9.4 TWh in 2013, which represents 8% of the Dutch electricity consumption. 75% of that (6.9 TWh) is at the expense of households and corporate ICT devices; such as computers, laptops, digital television and WiFi networks. 25% (2.5 TWh) is consumed within the ICT sector itself: 1.4 TWh by data centers, 1.0 TWh by telecom companies and 0.1 TWh by the ICT service providers. In the Netherlands, this also matches 2% of the national power consumption.

WORLDWIDE DATA CENTER FACILITIES - POWER NEEDS IN GW




Figure 2.1 Worldwide ICT electricity consumption of data centers per region. Source: Korhan Cengiz. Towards Green Communications. *Advances in Wireless Communications and Networks*. Vol. 1, No. 1, 2015, pp. 1-5. doi: 10.11648/j.awcn.20150101.11

Compared to the last broad study into energy consumption by ICT (Tebodin, 'ICT stroomt door' in 2007) the overall energy consumption of ICT in the Netherlands has remained almost stable since 2007 (a 3% reduction). This is striking, because in 2007, a significant increase of over 50% was expected. However, the sector has managed to keep the energy consumption stable, partly by implementing efforts set out in the MJA3 energy covenant with the Dutch national government, whereas the data traffic grew tens of percentage points per year. In the covenant, the industry committed to an energy efficiency improvement of 30% between 2005 and 2020. This goal has already been achieved.

For the subsector data centers, representing over half of the energy consumption within the ICT sector, there is an annual growth of the absolute power consump-

tion of about 3%, mainly due to the strong growth of data storage and processing. According to CE Delft, more connected devices, Internet of Things and more data consumption will result in an increase of the energy consumption. On the other hand, the trends of consolidation in the cloud and efficiency improvements of ICT equipment will persist and have a strong saving effect. A typical development is the shift from the energy consumption from users (so called business users with their own server space or data center) to the provider side (data centers that take over or host data center facilities).

Within data centers, investments in (energy efficient) free cooling, use of residual heat, combined with higher temperatures inside data centers contribute (in addition to more efficient hardware) to the limited growth of the energy consumption, compared to the rapid growth of services provided by the sector. In the

ELECTRICITY CONSUMPTION ICT IN THE NETHERLANDS, 2013




Figure 2.2: Electricity consumption ICT in the Netherlands.

report of CE Delft it is expected that the current trend of 4% increase in energy consumption per year for data centers will continue, resulting in an expected energy consumption of 1,670 GWh in 2020. It must be noted that 70% of the energy is currently being procured sustainably by the sector, leading to an overall CO₂ footprint of 0.34 Mt CO₂ (less than 1% of the overall CO₂ emission in the Netherlands).

2.2 Data centers in Amsterdam

The Amsterdam region is in the top 3 data center hubs in Europe after London and Frankfurt. The Amsterdam Internet Exchange (AMS-IX) is one of the largest in the world. A large portion of the European Internet traffic runs through the Metropolitan Region Amsterdam (MRA). Data centers are a crucial part of our ICT infrastructure. Amsterdam hosts over 35 big data centers, the larger ones being part of the MJA3 covenant of the ICT sector. With a consumption of 460 million kWh in 2013 (equivalent of more than 125,000 households) they are jointly responsible for 11% of the overall power consumption of all industries in the Amsterdam region. With the MJA3, the data centers in Amsterdam have agreed to realise energy savings of 68 million kWh in the period 2013–2016, which is a saving of 15%². To this end, the focus is put mainly on energy efficient cooling and energy efficient servers (for more information see Dutch Datacenter Report on Green IT)³.

2 <https://www.amsterdam.nl/gemeente/volg-beleid/agenda-duurzaamheid/datacenters-besparen/>

3. <http://www.dutchdatacenters.nl/dutch-datacenter-report-on-green-it-2015.html>

4. Blog by Prof. Joost Visser (SIG) about green software <http://ibestuur.nl/weblog/groene-software>

5. Note that not all data centers have influence on reducing energy consumption of hardware or software. Data centers that focus on co-location provide facilities for companies to place hardware and software of these companies without involvement and influence their settings. This project focuses on data centers which its own server hardware and have some span of control regarding the software and cloud environment.

6. <http://kngs.wikidot.com>

2.3 Energy use of data centers

One of the most frequently used indicators for energy performance of a data center is the Power Usage Efficiency (PUE). The PUE is determined by dividing the amount of energy entering a data center by the energy used to run the computer infrastructure within the data center. PUE is thus as a ratio, with overall efficiency improving as the quotient decreases towards 1. Until several years ago, a lot of data centers had a PUE between 1.5 and 2. A PUE of 2 means that the data center uses as much energy for the primary process of ICT equipment as it does for facilities such as cooling and back-up power. Over the past years, the bar has been raised for data centers. Currently a PUE of 1.5 is fairly common for average data centers. Modern data centers may reach a PUE of 1.2, meaning that 20% of the overall data center energy is used for facilities ('overhead').

Although the PUE is frequently used in the industry, it is also quite limited, because it mainly encourages energy savings in facilities and support services. Because of that, there is less focus on energy savings in the primary (IT) process, such as the purchase of energy efficient equipment, energy efficient software or the smart configuration of cloud environments. In particular, when it concerns the greening of data centers, a lot can be gained from switching to energy efficient software^{4,5}.

After all, when servers are used more efficiently, they generate less heat, significantly reducing the need for cooling. Case studies demonstrate that efficiently written software can realise two-digit percentage reduction in energy in terms of consumption⁶



Jeroen van der Tang
Manager Sustainability



What does your company do?

Nederland ICT is the trade association with over 550 ICT companies as members in the Netherlands. We are a primary contact for the government when it comes to ICT, energy and sustainability. Concrete measures have been taken in this area with the creation of the MJA3-ICT energy efficiency covenant and the SER Energy Agreement. As Manager Sustainability, I am both involved in Greening of IT and Greening by IT – the role of ICT in energy savings in other sectors and in the energy transition. Additionally, I am manager of ICT Milieu, the collective for the collection of ICT e-waste.

Why did you participate in the Greening the Cloud project?

Within Nederland ICT, green software has been on the agenda for quite a while. For instance in the form of the MJA3 knowledge network green software #KNGS, in the ICT 2030 Roadmap and as topic at meetings of the Software VOC and the MJA3-ICT business days. This is why we have supported the Greening the Cloud project from the very start and are involved as dissemination partner. Several members have also directly participated in the studies of the project and we will share the results within the relevant networks within Nederland ICT and the wider communities.

What are the most valuable learning experiences from the project for you?

The design of the project with the inclusion of three Amsterdam knowledge institutions, a dozen of enthusiastic SME ICT companies and the three dissemination and network partners has worked very well during the run of the project. The concrete and appealing results from the various studies had a motivating effect on companies, students and their supervisors. This has led to a lot of spin-off publications, presentations and integration into the curriculum.

What would you recommend for a follow-up study regarding greening of software and clouds?

Formulating the specific follow-up research questions is in great hands with the researchers of the VU, UvA and HvA. Together they form a unique knowledge base in Amsterdam around the topic of green software. What we as Nederland ICT consider important – in addition to knowledge development and knowledge sharing among the new generation of developers – is the connection and alignment with practical application. That is why we have taken the initiative to – as a follow-up to Greening the Cloud – expand from SME companies to bigger companies in the MJA3-ICT with the new project GreenServe.

“As a follow-up of the Greening the Cloud we have taken the initiative to expand from SME companies to bigger companies in the MJA3-ICT with the new project GreenServe”

PARTNER PERSPECTIVE

Energy Loss Chain in Data Centers




1. Energy Loss Chain Data Center

2. Uninterrupted Power Supply (-8%)

7. Fans (-8%)

8. Power Supply Unit (-14%)




3. Cooling (-23%)

4. Lighting (-1%)

9. Other Server Components (-16%)

10. Inefficient software (-10-70%)



5. Power Distribution Unit (-3%)

6. Supporting ICT Equipment (-10%)

11. Improvement potential (upto factor 10)

12. Overview + SEFLab

Figure 2.3 Energy Loss Chain by SEFLab <https://www.youtube.com/watch?v=Ae31N1LAOKk>

2.4 Energy loss chain

Figure 2.3 shows the energy loss chain of a data center, indicating where energy losses occur and how much percent of the inbound power is used for the primary process. Assuming a data center PUE of 1.5 (standard for data centers), major losses take place in cooling, lighting and backup power of the data center (33%). In the so-called server rack, where several servers are housed, there is loss in the power supply (~20–30%), fans in the servers (~8–10%) and other network equipment (5–15%). For an average data center, of each inbound 100 Watts of power, just 15–20 Watts is used for the primary process, a loss of 80–85%⁷

Efficiency of server equipment doubles every 18 months (Kooamey's Law)

The energy loss chain illustrates where the energy leaks of a data center are, indicating what the smartest investments are for realising savings. Within the server environment, the most logical saving opportunities lie in (i) energy efficient servers, (ii) virtualisation / consolidation, (iii) avoiding idle running ('comatose') servers and (iv) energy-efficient software.

Illustration: Costs of underutilisation

The Dutch publication 'Automatiseringsgids'¹⁰ calculated that an average server uses about 4,000 kWh for computing tasks and 1,750 kWh when standby. That makes the *standby* energy consumption of servers almost half of the maximum consumption. Including cooling (at a PUE of 1.5) the standby energy consumption is increased by 50% to 2625 kWh/year. If these idle computers were to be turned off, this would result in annual savings of more than € 280 per server (at an energy price of 11 eurocents per kWh). Since data centers easily have hundreds of servers, a significant number of which are standby, the saving potential is (tens of) thousands of euros. For this reason, several companies are active in the industry that offer data centers to switch servers on and off in a controlled manner via specialized software.

Energy efficient servers

Purchasing energy efficient servers is a great opportunity to save energy for users of data centers. Because of the ever advancing development of processors, the efficiency of server equipment doubles every 18 months (Kooamey's Law). As a result of this, it is recommended to replace servers after 4 to 5 years with servers of the latest generation (Harryvan, 2015).

An important focus point in servers is power management settings. A study of Certios shows that a large part of the consulted data centers, disable power management settings as standard practice (Harryvan 2015). Because of this, opportunities are missed to disable servers that are used relatively little and which are standby most of the time. Further investigation is required into the motivation thereto.

Virtualisation

In addition, there are significant opportunities with the advent of virtualisation, where physical machines are subdivided into virtual machines, so that the physical resources of the server (computing power, storage, network capacity) are shared by various customers. Various studies show that the capacity of servers in the industry is significantly underused, varying from 15% to 33% utilization of servers (Harryvan 2015, Ansett 2015⁸, IDC from AG article⁹). Mainly in smaller data

Case: The StateCloud (RijksCloud)

A good example of how virtualisation and consolidation can lead to major energy savings, is the 'Rijkscloud'. With the RijksCloud, the national government bundles its data centers of various departments and organisations, due to which the number is reduced from 64 to 4. By merging the 64 data centers, the government expects to realise savings of 107 GWh per year – based on the first completion in 2020 – from 235 GWh now to 128 GWh in the future, a 45.5 percent reduction. This is realised by means of innovative cooling, reuse of heat and advanced virtualisation. The number of hectares of servers is reduced from 1.9 to 0.3.¹²

centers with relatively little virtualisation, the utilisation of servers is low; bigger data centers with a high level of virtualisation often have a higher utilisation rate for servers, but often not higher than 50%. That is striking, because from an energetic point of view, servers operate at an optimal extent at a utilisation rate of 70–90% of the maximum capacity (Harryvan 2015).

For some industries, having idle systems may make more sense, given that putting ICT systems in sleep-states comes at a penalty due to latency times. However, many business processes do not have strict reaction times, which could legitimize putting systems in sleep-state rather than keep running idle.¹¹ A lot of customers that have physical servers under their own management or customers that use physical servers via co-location, prefer uninterrupted continuous service and accept the extra cost for overcapacity as collateral damage. However, a study by Logius and SIG show however that 'standby' servers can be real money guzzlers. The reason is that in addition to energy costs, operational costs such as licenses, management and hardware amortization must be considered as well. According to initial calculations, the savings per kWh are not just 0.11 eurocents (regular energy costs) but

can be as high as 24 euro per kWh (more than a factor 200 higher!)¹³. This cost model has not yet been validated in other cases, but does offer perspective to convince data centers and customers to focus more on power management and virtualisation.

There are big opportunities for the reduction of the number of 'idle' servers, configuring power management settings and the virtualisation of servers, so that physical servers can be used more efficiently.

Energy efficient software

Where there is a growing focus on virtualisation and power settings to realise energy savings, energy efficient software remains a relatively neglected topic. There is a lot to gain here. Because of the energy loss as outlined in the energy loss chain and a lower utilisation rate of the physical servers, only 10 to 15 Watt of the 100 Watt is effectively used for computing power. Vice versa, this means that more efficient software or a more efficient cloud configuration translates into energy savings with factor 7 to 10 at data center level. More energy efficient software has major leverage on energy savings.

7. Energy Loss Chain by SEFLab <https://www.youtube.com/watch?v=Ae31N1LAOKk>

8. <http://knngs.wdfiles.com/local--files/start/scan00556920151113101359.pdf>

9/10/11. http://knngs.wdfiles.com/local--files/knngs%3Aag-9-2015/WebAG_P12_AG-09-19-2015.pdf

12. https://www.sig.eu/index.php/download_file/506/948/

13. <https://www.deingenieur.nl/artikel/rijkscloud-halveert-energieverbruik>

Another reason to investigate savings potential in software is that – despite the fact that hardware is becoming more efficient, software is growing bigger and more complex, due to which it consumes more and more energy. This is illustrated by Wirth's law (Wirth, 2015)

"Software is getting slower more rapidly than hardware becomes faster."

Scientific literature frequently refers to the importance of energy efficient software and applications (such as Ranganathan 2010 and Berl et al., 2010). Knowledge development concerning green software is still limited however. There are green design guidelines for software, but these mainly focus on mobile or embedded applications, given the urgency to increase the battery life for these devices (Flinn & Satyanarayanan, 1999; Thiagarajan et al., 2012, Simunic et al., 2000; Peymandoust et al., 2002). The VU Amsterdam and the University of Amsterdam, also partners in this project, has done a lot of research into green design strategies and best practices for software and cloud infrastructure (such as Lago et al., 2014; Procaccianti 2015) and energy performance in high performance computing clusters (GreenClouds).¹⁴

Still, knowledge development in respect of guidelines for cloud applications is in its infancy. The set of guide-

lines available for software developments is far from complete, these guidelines are at various levels (varying from programming languages to operating systems) and a lot of guidelines from literature have not yet or to a very limited extent been tested in practice. The greening the cloud project offers the opportunity of testing a number of the design guidelines and validating them in practice with real world cases and companies.

2.5 Conclusions

A lot of progress is being made in greening data centers, for instance with more efficient cooling and efficient hardware. Regardless, the energy consumption of data centers is expected to keep rising, due to the growth of workloads that servers will have the process. A step that is currently being taken in the industry is energy savings by virtualisation, whereas, in the short term, the deployment of power management options can further reduce energy consumption.

The biggest leverage can be achieved by programming energy efficient software at the source. In addition to concrete guidelines for designers, it is first important to map the savings potential with clear results of simple design adjustments and to make the ICT chain more aware of the possibilities offered by green software. The Greening the Cloud project provides possibilities to contribute to this specifically, and has therefore consciously chosen use cases that study energy effects of virtualisation and green software practices. These are presented in chapter 4.

14. http://wiki.cs.vu.nl/greenclouds/index.php/Main_Page



Simon Besteman
director of ISPCONnect and
portfolio holder Education of
the umbrella organisation DINL



What does your company do?

ISPCONnect is the trade association of hosting companies and cloud providers. DINL is the umbrella organisation of the entire digital infrastructure.

Why did you participate in the Greening the Cloud project?

Our supporters are major cloud users. We are the companies that place their servers in the data centers. From the perspective of the sector, we have two important challenges. As major consumers of energy, we want to work more sustainably, we acknowledge that we must take our responsibility to make sure – through collaboration and innovation – that the carbon footprint of the internet is reduced. In addition, power is one of the main cost items for our members. We are actively looking for ways to deal with the server loads we have to run in smarter and more efficient ways.

What are the most valuable learning experiences from the project for you?

Firstly, the fact that this project came to fruition, is a great success. The fact that businesses and the educational sector have worked together to explore this reasonably unknown area and to look for room for improvement together is an achievement in itself. Moreover, demonstrable results have been provided. A solid foundation has been laid to build on with the identification of the potential for improvement and to convert this into actual practical application.

What would you recommend for a follow-up study regarding greening of software and clouds?

Eventually, you will win the sector over by putting concrete things on the table. Follow-up studies must focus even more on the practical, turnkey solutions companies can work with. This means that insight has to be provided in the relationship between the investment and the benefits, and what can be done concretely to start working on it. I think that's the direction we need to head at, to take these promising results to the next level.

PARTNER PERSPECTIVE

"A solid foundation has been laid to identify potential improvements in green software and to convert this into actual practical application"

10 Best practices for a green IT System

With support of the Knowledge Network Green Software, the Software Improvement – within the project Cluster Green Software (UvA, VU, HvA) – presented a list of 10 best practices for green IT systems. The list was formed after a practical study, in which the energy footprint of generic IT infrastructure of various ICT users has been reviewed as well as what the generic possibilities are for energy savings. The best practices are intended for procurers, suppliers and policy makers.

1. Apply virtualisation	Where possible, use a virtualized environment: Virtual servers can be 10 times as energy efficient. When defining the architecture, consider the fact that all components can be virtualized.
2. Use energy settings	Make use of the energy efficiency settings of the hardware and the virtualisation layer: Footprint definitions have shown that the use of energy efficient hardware settings (for instance in the use of the CPU C states) is not trivial.
3. Enable measuring	Provide a measuring infrastructure for the determination of energy KPIs during the rollout: Enabling energy measurements in hindsight is not cost efficient and is difficult to achieve. Having the energy measurement available is crucial for steering on energy efficiency.
4. Investigate	Dare to investigate what approach leads to more energy efficient solutions: Via the measuring infrastructure, trends and consequences of changes can be made visible. In many current systems, one is afraid to make changes and energy inefficient situations remain in effect.
5. Upgrade hardware	Replace old hardware with new hardware in a timely manner: Old hardware (three years and over) is less efficient than the latest equipment. Moreover, the capacity and computing power of hardware (and software) is still increasing year over year.

6. Limit oversizing	Projects and systems often use a horizon of a couple of years to determine the load, with significant oversizing and inefficiencies in the beginning as a result. Choose a gradual capacity buildup.
7. Reconsider availability requirements	Users that fear risks tend to define the desired availability very broad, to be on the safe side for their application. In hindsight, these requirements sometimes appear to be much broader than was strictly necessary.
8. Optimise test environment	Only activate the test and failover environment on demand: Design the test and failover environment in such a way that they are only activated when they are actually being used. It often happens that test environments are unnecessarily constantly activated, which consumes a lot of energy.
9. Optimise for performance	Improving the hardware resource consumption (CPU/memory etc) per amount of work (transactions) often leads to a reduction of the energy consumption. There are a lot of tools and experience available for performance analysis: use them.
10. Understand your workload	Know your workload and take it into account in the sizing of the system: A lot of systems show a pattern in the workload: for instant a very constant load, or a peak once per day, week, month or year. Gain insight into the expected workload and configure your system as such that it can respond accordingly.

- Measure the instantaneous power consumption with sufficient sample rate, to enable signal reconstruction and to adhere to the research goals.

An important challenge was to identify where the power lines lie within the server, which component is connected to which power line, determine the most effective way to interrupt the power line, establish voltage levels of these power lines (given that they differ for different components) and what the whole measurement chain should look like.

Illustrative for the different voltage levels, the component voltage levels of the first servers that were instrumented ranged from 1.2V to 12V. The voltages between the power supply unit and other components such as motherboard, hard disks and so forth were 3.3V, 5V and 12V.

The measurement chain entails all the sensors used, the signal processing components (filtering, level of amplification, reduction of noise sensitivity, analog/digital conversion), the data acquisition system (DAQ), read out computer and complementary software for reading data from the sensors and visualising them in a correct way. The main challenge was to do so within the limits set for errors (2%). Given the typical characteristics of a server (i.e. several low voltage levels, high currents, unknown layout of the power grid) and the measurement requirements within the SEFLab, it was concluded that off the shelf sensors would not suffice; we needed to develop our own dedicated sensor systems.

3.2 Details of the first measurement system

The first two measurement setups were a Dell PowerEdge SC1425 and a SUN Fire X4100. In both systems we interrupted the ATX power cables (an ATX cable is a standardized cable for power connectors) between the computer's power supply unit (PSU), the motherboard, CPUs, and hard drives to measure its energy consumption. Between both wire ends a sensor was placed. The first sensor module was quite rudimentary, it consisted of passive components only, being a shunt resistor and a filter circuit. The shunt resistor enabled us to measure the current that flows through it and the filter circuit filtered out undesired high frequency noise from the signals that we wanted to measure. The voltages

of the power distribution lines were measured non-intrusively by the setup. Also, a dedicated filter circuitry was applied.

The measured voltages were digitized by a National Instruments Data Acquisition system (DAQ), read into the computer by Labview, and processed and visualised by Matlab. Within Matlab the components' current was calculated by dividing the measured shunt voltage with the shunt resistance. Subsequently, we determined the component's power consumption by multiplying the calculated current with the measured component voltage. An illustrative study was a benchmark between different browsers (presented in: Ferreira et al, 2013).

3.3 Measuring modern servers

Modern servers do not allow interrupting the ATX-cable as was done in the first SEFLab system. In modern servers the connection between the PSU and the motherboard contains no wires anymore as this induces too high voltage drops due to the high currents between both components. As a solution, thick solid power connectors are used to connect the PSU and the motherboard.

Therefore, it was necessary to invent a new measurement methodology to measure the power consumption. For this purpose, we invented a new flexible and modularized approach using specialized *sensor modules*. First, the components of interest for power measurements were determined, in our case mainly the two CPUs, four memory banks and multiple hard drives. Depending on the research goal and power grid layout, the fans, the chipset, video card, network adapter, or a PCI-card are also measured.

According to our new methodology, connectors are installed into the current path of the selected components. The connectors house specialized sensor modules that enable measuring the power consumption through these paths. The advantages of installing a connector rather than soldering a sensor module into the current path include the following:

- Sensor modules are disconnected from the server itself, and can therefore be further developed and updated after the measurement setup is finalized.

- Sensor modules can be recalibrated when needed, as they can be easily disconnected from the server. Malfunctioning sensors can be replaced easily.

Within modern servers it is not easy to locate the component's power lines on the motherboard. However, most components have dedicated DC-DC converters or other voltage regulators which reduce the main PSU voltage to the component's operation voltage. As these circuits are in the current path it was decided to install the connector into the current path by temporarily desoldering a circuit element, (e.g. a coil, a fuse, a transistor or a ferrit) from the motherboard. Subsequently the connector is installed into the current path and finally the circuit element is placed back. For hard drives another methodology was followed, given that they usually gain their power through wires. To implement the connectors, the wire is cut, and the connector is placed in-between both ends.



Figure 3.1 Wires soldered on motherboard to measure energy consumption of memory.

Drawbacks of the above presented method are that the connector and the sensor module add a small resistance (10 to 50 mΩ) to the current path (although largely within the allowed limits), that in-depth knowledge of the motherboard layout is necessary (making instrumenting a new server somewhat more time-intensive) and that in a small number of cases the chip's energy consumption cannot be measured due to specific board layout choices. The latter limits the amount of servers that can be instrumented.

3.4 Recent measuring components

In the past couple of years, a lot of work has been put into developing new sensors with more specific features for effectively measuring energy flows. In addition, two new types of sensors have been developed that can be placed in the connectors installed in the server.

The I2 Sensor

The first sensor, named I2, connects seamlessly to the measurement chain as described in section 3.1. Compared to passive sensors, this active sensor has a number of advantages:

- **Kelvin Bridge:** the measurement circuit has its own connection wires to the shunt resistor. As a result the measurement remains accurate at high component current-levels.
- **Amplification:** the measured shunt-voltage is amplified locally (resulting in lower noise sensitivity).
- **Voltage stabilization:** the sensor voltage is stabilized locally, so that the local amplification suffers less noise caused by long power cables.
- **Temperature measurement:** the sensor temperature is measured, so that measurements can be corrected for fluctuations in the sensor temperature. Temperature fluctuations occur because of self-heating of the sensor due to the high component currents through it, affecting the sensor accuracy by several percent.

AD2 Sensor and the UD2 module

The second sensor we developed is the AD2 sensor, an advanced version of the analog 12 sensors. In the AD2 sensors, an analog-digital convertor is integrated into the module, as a result of which the exit signal of these sensors is a digital signal instead of an analog signal. For these digital sensors, a proprietary DAQ module was developed (the UD2 module) to which four sensors can be connected simultaneously. As can be seen in Figure 3.2, the AD2 sensors communicate with the UD2 modules via a 3.3V Serial Peripheral Interface (SPI) bus. The UD2 module bundles the measurement data and sends the data to the measurement computer. The voltage and the temperature are sent from the AD2 sensor to the UD2 module in analog form and

converted to digital values with the ADC of the UD2 microcontroller. The sample rate of the UD2 module is configured at 1 kHz (1,000 times per second). The bandwidth of the AD2 sensor is configured at a cut-off frequency of 179 Hz, to limit aliasing (signal distortion caused by sampling a time continuous signal).

3.5 Development challenges

At the start of the Greening the Cloud project, a relatively professional measuring setup had been built within the SEFLab. However, there was room for improvement regarding the digital sensors in which the need arose to measure both clients and servers simultaneously, so that the rapidly growing industry of cloud

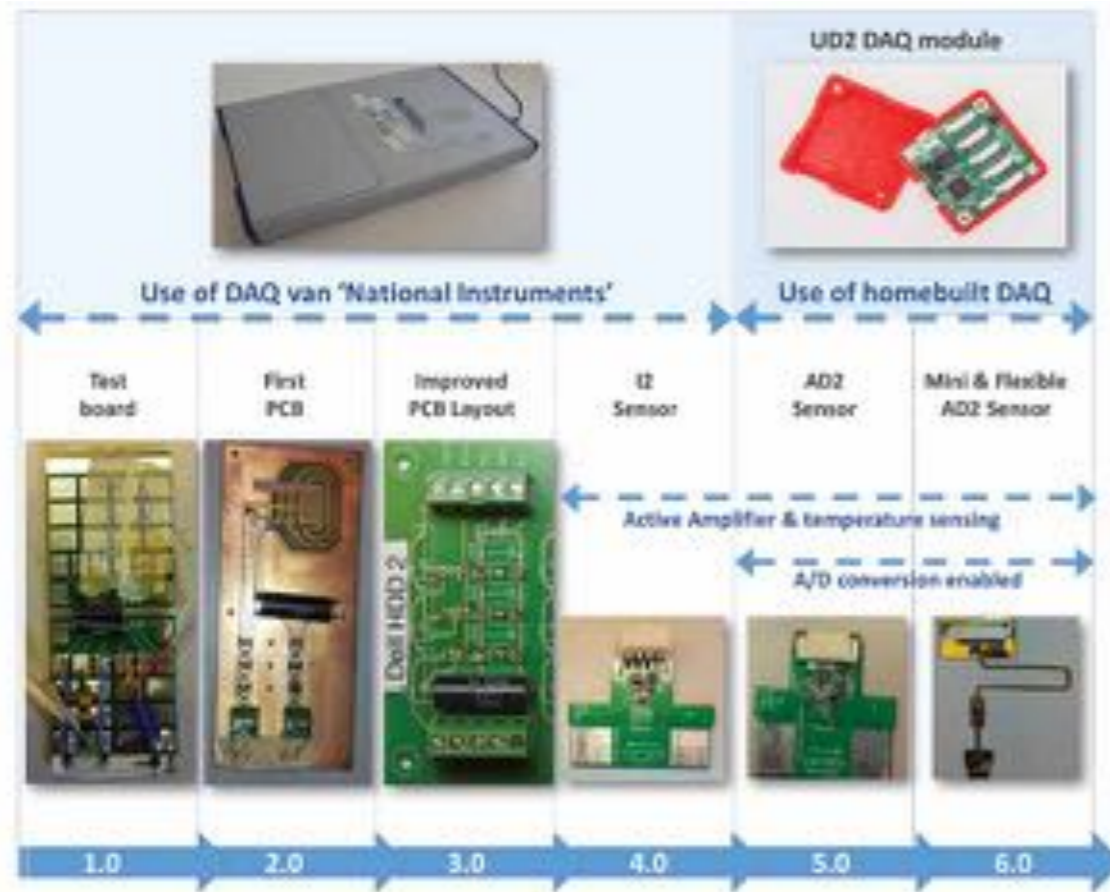


Figure 3.2 Sensor board development steps.

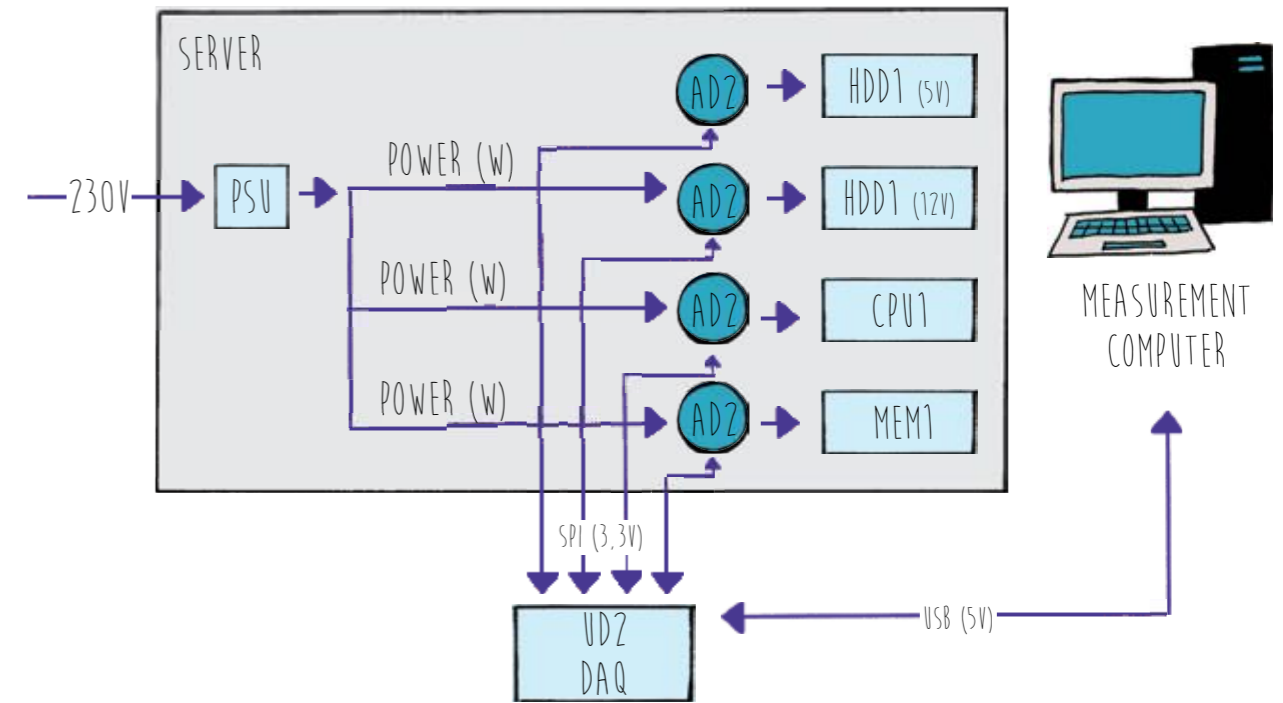


Figure 3.3 Block diagram of measurement system with digital sensors.

environments could be investigated. In addition, there was much to gain from automatic calibration to increase accuracy, making available a larger number of servers on which the measurements could be performed and there was the desire to allow for remote login on the SEFLab ("remote testing"). The latter would offer (intern)national companies and institutes the opportunity to have tests carried out remotely and review the results of the tests online. The Greening the Cloud project made it possible to further professionalize the SEFLab.

Increasing measurement accuracy

To increase the accuracy of the measurements, the communication robustness was researched. The AD2 sensors, combined with the UD2 data acquisition modules, provided valid measuring data with one sensor connected. However, with two or more sensors, it caused unexpected peaks in the measuring results.

The reason for this was poor shielding, which was resolved by improving the shielding of the digital communication cables. Further improvement is possible by amplification of the digital signals (with line drivers).

Client-server setup

Key within this project was the synchronous measurement of the energy consumption of both clients and servers. To this end, software was written that can read several UD2 DAQs simultaneously, allowing for the synchronous measurement of both client and server (UD2 reader software).

With regard to the structure of the software, multi-threading was chosen to prevent issues in the receipt of the data from the UD2 modules due to the displaying of the results. Per module, the software contains one thread that reads the data and one thread that processes and writes the data. The U2D reader software

Client –Server systems

A client-server system is a combination of a computer system (client) and another system (server) connected through a network. The client initiates the communication with the server, for instance for the purpose of requesting data, transferring data or carrying out an action on the server. Client-server systems represent cloud environments in which laptops and desktops (clients) are set up for input and control, and in which servers in data centers or in the cloud focus on carrying out tasks and delivering processing power. A client-server measuring setup is indispensable when calculating the energy consumption of clouds.

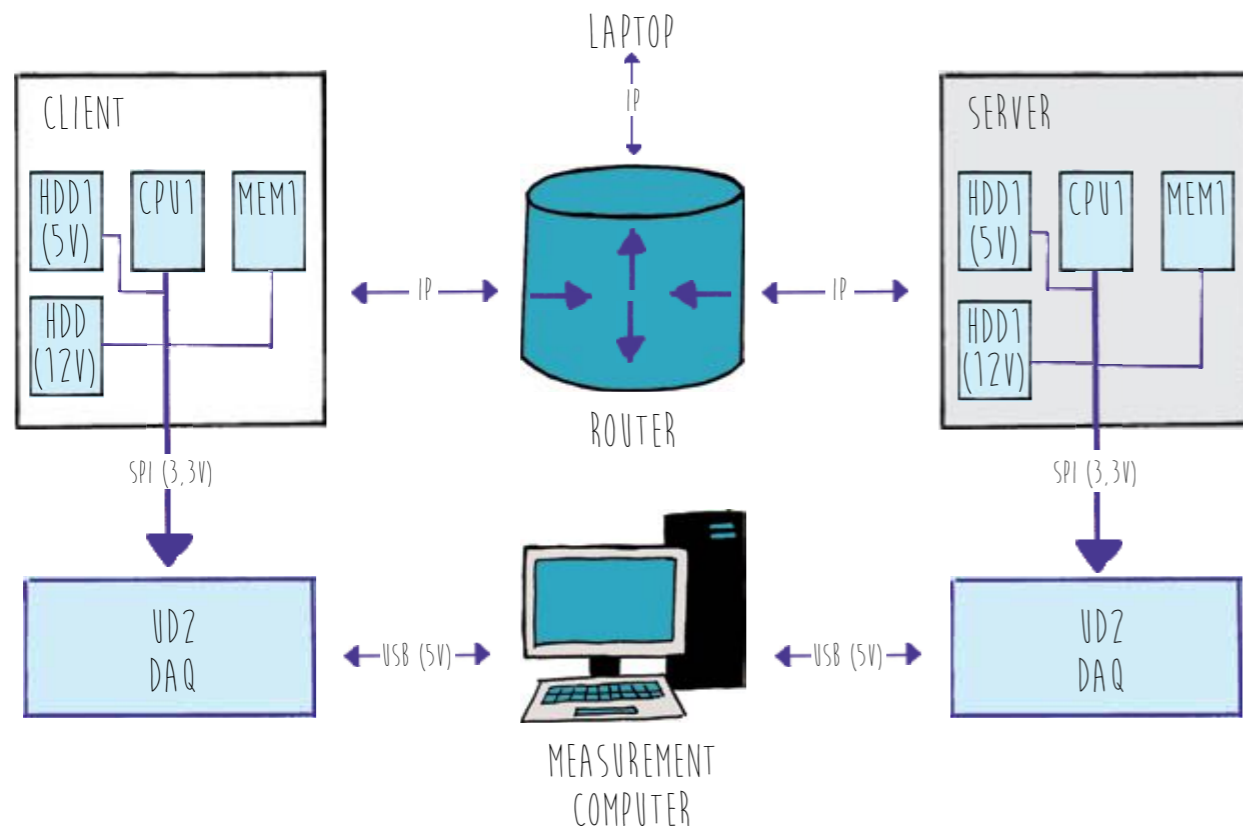


Figure 3.4: Design client server measurement system.

corrects the measured power and voltage based on calibration tables (see the paragraph 'Calibration' below) and writes the files in a common format (.csv). Figure 3.4 shows a client-server measurement setup. This setup is used to send a file via FTP, see Figure 3.5. At the start (phase 1) the server and client are booted, after which (at phase 2) the virtual machine (Debian 6) is booted. At phase 3, the file is placed on the server for

the FTP transfer, which is why the client does not show increased CPU activity. Subsequently, SSH was used to login to the client, after which an FTP connection was made with the server. Phase 4 shows that a file was sent from server to client four times. The results in the graphs are the measured data after filtering and averaging.




Figure 3.5 Client server functional test case: FTP file sent.



Emile Nijssen
 Founder and Lead enigeer



What does your company do?

At GreenInvents, we have a passion for making daily life as energy efficient as reasonably possible. Our aim is to provide designers of electronic products the tools to accurately measure and profile the power consumption of their products without the big service contracts or high instrument prices that traditionally plague the test equipment sector. Besides just energy efficiency in operation, GreenInvents also provides life cycle analyses and electromechanical design services to improve energy consumption during production and decommissioning. GreenInvents consists of the founder and lead engineer (Emile Nijssen) along with a small team dedicated to put the products out there in the world (See also www.greeninvents.com)

Why did you participate in the Greening the Cloud project?

As a young and ambitious company we consider it a challenge to assist in designing greener server systems and cooperating with companies and organisations that have the same vision and goals as we have - i.e. use available energy as efficiently as possible in electronic systems. To this end, we have contributed in the design and deployment of the instrumentation of a few server systems.

What are the most valuable learning experiences from the project for you?

Aside from the results from our direct involvement, the partnerships within the 'Green IT' community as well as results from the various work packages of the project have given us invaluable insight into those areas that are ripe for further 'greenification', and the instrumentation necessary to get there. This has given our company the opportunity to focus on a few specific products targeting those niches.

What would you recommend for a follow-up study regarding greening of software and clouds?

Maybe not surprisingly, but as a Dutch saying goes: "Meten is weten" - Measuring is knowledge, loosely translated to 'knowing is half the battle'. In our opinion, existing measurement interfaces like Nagios don't go far enough in effectively profiling energy consumption. Specifically, temporal resolution is key; we shouldn't go forward relying on measurement systems with 1-second accuracy.

"The partnerships within the Green IT community as well as results from the workpackages of the project have given us invaluable insight"

Calibration

The UD2 software works with calibration tables. The sensors readout is inaccurate with a few percent due to scattering in components used (resistors, voltage regulators). Furthermore, the sensors are sensitive to temperature differences, in such a way that it significantly impacts the accuracy of the measurements. Temperature difference in the server arise as a result of heavy loads, and the sensor heats up due to high component currents. This can result in tens of degrees of sensor temperature differences, with significant effect on the measuring values (+/- 2% accuracy). For these reasons, a calibration table has been drawn up with which the measuring values of the sensors can be

corrected for component scattering and temperature fluctuations. To establish the effects of the temperature on the measuring values of the sensors, a calibration oven was developed.

Remote access

The desire to be able to remotely login to the measuring systems is significant. The reason is that it would allow researchers to upload their software under test to the lab, carry out the measurements remotely and download the measurement data. Following the project, parties can login to the measuring systems online. The first functional tests have been carried out, after which remote login will become available for external parties.

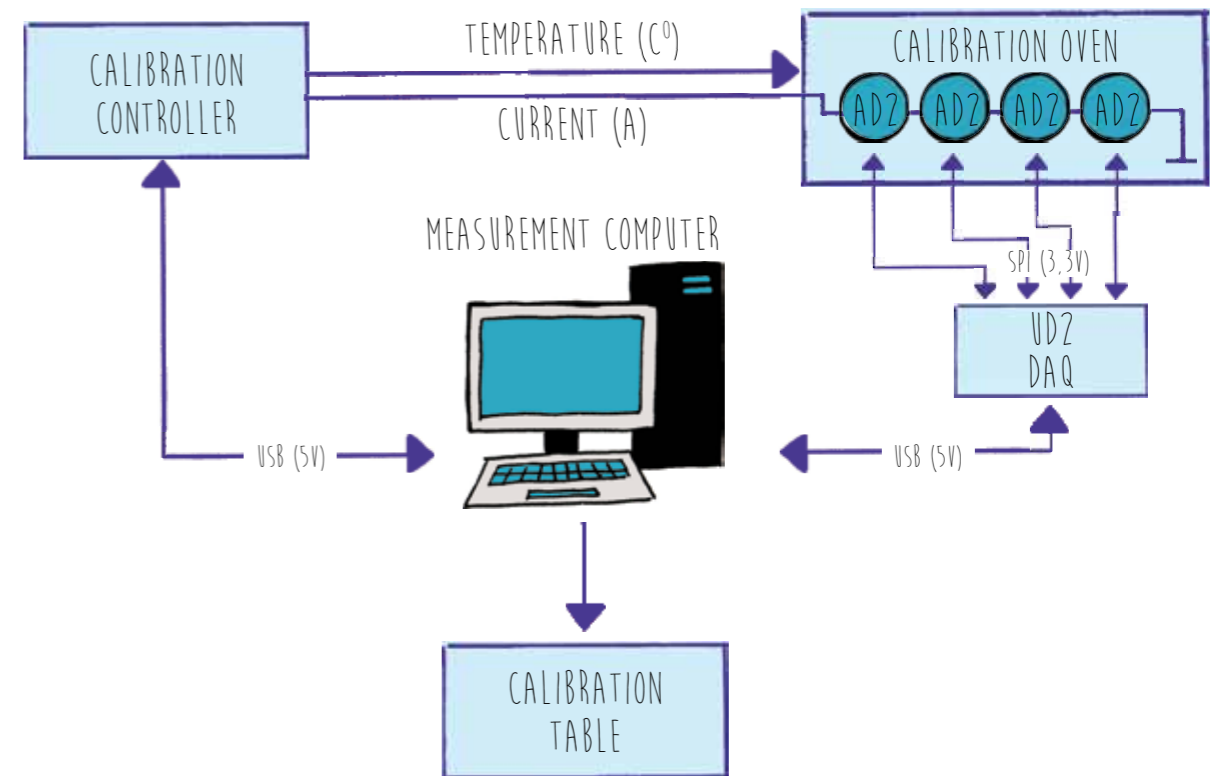


Figure 3.6 Schematic flow to generate calibration tables.

Current server park

During the Greening the Cloud project, a number of servers was converted into a measuring platform. As a result, the server park of the SEFLab has grown to a total of nine servers at the end of the project, five of which are operational measuring platforms. In addition, there is one measuring platform in the Green Lab.

Server	Status	System Information
Dell PowerEdge SC1425	Obsolete	System is out of date
SUN Fire X4100	Obsolete	System is out of date
HP ProLiant DL380 G5	Not available	Instrumentation failed due to motherboard lay-out
HP ProLiant DL360 G7 – no 1	Available	Is situated at Green Lab
HP ProLiant DL360 G7 – no 2	Available	Is situated at SEF Lab
Dell PowerEdge M620	Available	Blade server; Is situated at SEFLab
Dell PowerEdge R620	Used as remote login server and file server.	Instrumentation failed due to motherboard lay-out
IBM X3550 M4 – no 1	Available	Is situated at SEFLab
IBM X3550 M4 – no 2	Available	Is situated at SEFLab

Table 3.1 Current server park of the SEFLab.

3.6 Conclusions and further research

The Greening the Cloud research project has allowed for the professionalization and further development of the SEFLab. A client-server setup with high quality sensors has been developed (including the necessary software) which allows for the measuring of cloud applications. The server park has been expanded to five operational measuring platforms equipped with sensor systems developed in-house. Furthermore remote access was realised and a calibration set up has been developed to allow calibration of temperature fluctuations.

Further work should focus on continued validation of the client-server setup for various cloud applications as well as for the remote access setup, through extensive testing with external parties. Furthermore improving the measuring accuracy of power sensors in servers remains relevant by increasing the bandwidth of the measuring system, and for instance by developing more knowledge about representative wave forms of server components (e.g. CPU, Memory, and hard drive) to predict measurement accuracy of new sensor systems upfront.



Figure 3.7 SEFLab measuring platform: the HP ProLiant DL360 G7.



4 GREENCLOUD CASE STUDIES

4.1 Introduction to cases

Knowledge is power. Many professionals in the field have ideas, based on years of experience, about which software applications require a lot of resources (i.e. energy) or what options could be more efficient. Scientists have also formulated a large number of hypotheses about the energy impact of cloud environments, operating systems, applications and options within those applications. These hypotheses are often not validated or only validated to a limited extent.

Because of this, the value of these hypotheses is unclear. The SEFLab and Green Lab offer the opportunity to conduct energy measurements of all kinds of software and cloud related variables. These tests are at the heart of the Greening the Cloud project. In total, 10 case studies have been carried out (see table 4.1). Six of these cases are discussed in more detail in this chapter. The cases can be subdivided into the topics Virtualisation (4.2), Power estimation models (4.3) and Software engineering practices (4.4).

4.2 Case studies regarding virtualisation

Virtualisation has a crucial role in the delivery of clouds services. Hypervisors, or virtual machine managers, are essential pieces of software used by cloud providers. A hypervisor runs on a physical machine and allows multiple instances of virtual machines (VM) or guest

operating systems to run simultaneously and independently from each other on the physical host. The effect of virtualisation techniques on energy and power consumption is partially uncharted territory, especially when it comes to comparing different hypervisors and their energy efficiency. Any effort to create a greener cloud should take the efficiency of hypervisors into consideration. Because of this, we decided to measure the energy usage of two open source hypervisors.

Topic	Partner	Assignment	Knowledge partner
Virtualisation – Hypervisors / containers (4.2)	Schuberg Philis	Research into the effects of different hypervisors on energy usage	UvA
Virtualisation – hypervisors / containers (4.2)	Cloud Provider	Comparing energy usage between containers versus virtual machines	UvA
Power estimation models (4.3)	VMware	Validation of the energy modelling function within VMware	VU
Power estimation models (4.3)	SURFsara	Validating energy modelling by RAPL and BMC in blade servers	HvA
Software engineering (4.4)	Diesveld Query Technology	Benchmarking different query techniques	VU
Software engineering (4.4)	REM Automatisering	Energy usage and optimisation options for an open source cloud application "OpenWave"	VU
Webservers	Cobra	Energy consumption benchmark of a WordPress website and his hardcoded version.	VU
Webservers	Transfer solutions	The energy difference between ASP (Windows), Java (Windows), Java (Linux) and PHP (Linux).	HvA
Encrypted data transfer	Greenhost	An overview of several encryption algorithms compared in their energy consumption.	HvA
E-mail server systems	VKA	Energy consumption of different options in Microsoft Exchange Server	HvA

Table 4.1: Cases in Greening the Cloud.

Focus on virtualisation techniques

To determine how virtualisation affects power-efficiency UvA students¹⁵ involved in the project initially performed a number of scientific literature studies. This revealed that the interest in virtualisation techniques of the scientific community up to 2014 does not differ from the interest since 2014 (Yu Ri Tan, 2015). About 44% of studies published concerns hardware (including the infrastructure of a data center), and 56% concerns software. Of the research in software the majority of studies addresses the software applied by cloud providers to distribute the tasks issued by users on their servers. We also learned that most surveys in the scientific literature are, not surprisingly, focused on open source hypervisors. Two well-studied open source hypervisors are Xen¹⁶ and KVM¹⁷. Comparison of the performance of both these hypervisors constitutes the main body of research, whereas the power consumption difference is less studied.

The majority of research on the power consumption of Xen and KVM (Kernel-based Virtual Machine) focuses on just one of these two hypervisors (R.A. Saez, 2015).

Most studies make use of a linear model to fit the data of the power consumption of a VM. Such a model can be used by the cloud provider as an estimation model for the power consumption of their servers as a function of the work done by the VMs deployed. However, linear models turn out to be unrealistic. Newer models try to use non-linear models to fit the data of the power consumption of a VM as a function of the workload. In these newer studies three components of each physical machine are monitored in terms of energy usage: the processor (CPU), hard disk drive (HDD), and network interface controller (NIC), since these components of the physical machine have the highest power consumption. By running specific benchmarks, or stress tests, a relation can be found with the total power consumption of the physical host. For benchmarks stressing the CPU, this relation is expressed in the number of floating point operations per second (FLOPS) per Watt, and for stressing the hard disk drive it is the number of reads/writes per Watt. Most of these studies found that Xen performs better in these respects compared to KVM, except when the number of VMs per physical host exceeds 10, in which case KVM outperforms Xen.




Figure 4.1 Virtualisation hypervisor vs containers.

15. Based on: See studies referenced under literature list by Y. R. Tan, R. A. Saez, C. van der Poll, M. Warnaar, J. van Kessel, 16. <http://www.xenproject.org/> 17. http://www.linux-kvm.org/page/Main_Page

Cases UvA: Investigation virtualisation platform aspects

Power consumption of Xen versus KVM

We decided to move our effort to test these two hypervisors against each other ourselves in a real setting. The first industrial use case we set up was at Schuberg Philis, where we studied the difference in energy consumption of Xen and KVM (C. van der Poll, 2015). During this study we found that KVM consumes about 60% less power than Xen when the physical host is idle, i.e. when zero VMs are deployed. This is due to the fact that KVM is a simple extension to the Linux kernel and is able to enter sleep states, whereas Xen is not able to do this. For the CPU-stress test the power consumption of KVM is less than the power consumption of Xen, whereas for the HDD-stress test the power consumption of both hypervisors does not differ significantly. Approximately equal consumption of power consumption is also observed for an IO-stress test, but a memory-stress test shows a difference in power consumption of about 10% in favour of Xen. Alternatively, running the different stress tests in a combined stress test shows Xen to consume about 12% less power than KVM. This study shows that for mixed applications, Xen is the best choice with respect to power consumption, but for memory intensive applications KVM is a more suited choice.

The outcome of this study differs from some published studies mentioned above. There are several reasons for this, besides the role of the hardware used, and the fact that in the studies published different versions of the hypervisors are used. In the study performed at Schuberg Philis, the power consumption was measured at the power socket of the server used.

How to determine the power consumption of a virtual machine?

In addition to the question which benchmarks are used, we realised that another important question is how to determine the power consumption of a virtual machine. Our survey on the different methods of measuring the power consumption of virtualized hardware as a function of the workload, shows that the way the power consumption of a VM is determined, is far from trivial (M. Warnaar, 2016). Two methods are usually applied, performance monitoring counter-based (PMC) and OS-performance counters. In PMC-based models, logged (virtualized) hardware event counters are used to form a model of the power consumption, whereas OS-performance counters provide numbers of CPU, disk, and memory utilisation. Both methods may result in a model for the power consumption that deviates from the power consumption of the components involved.

Figure 4.2 shows the different parts of a server that contribute to its total power consumption. From this figure it follows that relating the performance of a CPU-benchmark to the total power consumption of the server will be less precise than relating the CPU performance measure to the power consumption of the CPU itself. The number of MFLOPS per Watt may vary according to different room temperatures, because the cooling fans of the server will make more revolutions per minute as the temperature rises.

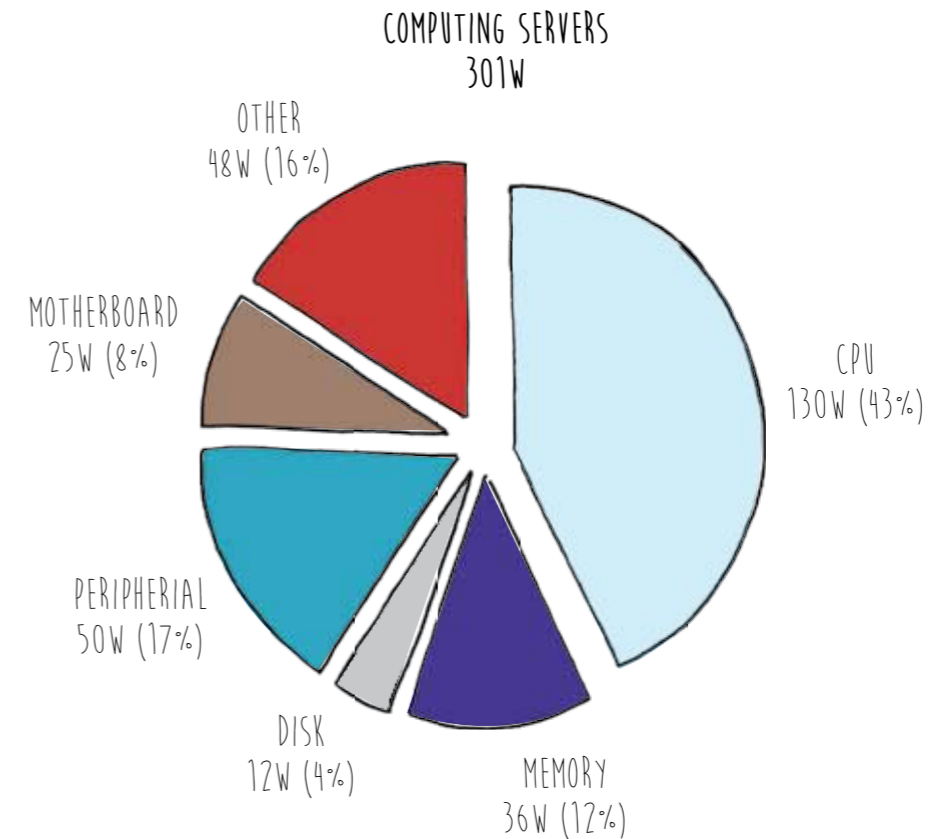


Figure 4.2 Power usage of a server per hour represented per hardware component. ©Kliazovich et al.

Real power consumption of Xen versus Docker

Besides virtualisers like Xen and KVM, containers have become popular by cloud providers. A container solution like Docker wraps up software within a complete file system that contains anything that can be installed on a server. During the course of the project we decided this was an interesting and novel virtualisation technique that needed to be investigated in relation to the model building process.

We studied the power consumption per component by running different benchmarks in a VM deployed by Xen as well in a Docker container (J. van Kessel, 2016). Our results show that in idle mode Xen, Docker and the native Ubuntu OS roughly exhibit the same average power consumption. The average CPU performance (MFLOPS) is higher for Docker than for Xen, Xen showing a performance of 14% less than Docker. However, the CPU power efficiency (MFLOPS/Watt) is about 11% better in favour of Xen. Concerning the average memory power efficiency (in operations per second per Watt) both Xen and Docker perform equally. The average HDD performance for write operations (in kb/s) is about equal for both Xen and Docker, also the HDD power efficiency for write operations (in kb/s per Watt) does not differ significantly. However, the average HDD performance for read operations is significantly better for Docker than for Xen, and the HDD power efficiency for read operations is also significantly better for Docker than for Xen.

These results show that for a single VM deployed by Xen versus a single Docker container there is no clear choice to make, based on power efficiency, one has a higher CPU performance for Docker, but Docker has a worse CPU power efficiency than Xen. Additional experiments are performed to study the performance and power efficiency when the number of VMs in Xen and the number of Docker containers increases.




Figure 4.3 Average CPU Performances with Xen and Docker.

4.3 Power estimation models

The facilities developed by the SEFLab and used in the Green Lab are hardware based; they rely on sensors either placed in the power lines of the server itself (enabling only measurement of components) or at the socket (allowing to measure the complete server). Also, software-based power estimation models are available to provide estimations on energy consumption by servers for applications as well (see for an overview Nouredine et al., 2013 or Jagroep et al., 2015). Literature shows that estimating software energy consumption on servers generally provides rough

indications of energy consumption or in case of varying loads shows relatively high error margins (Schubert et al, 2012). It also shows that these estimation models in a lot of cases only function on one type of operating system or particular hardware (e.g. intel energy checker only on intel processors) (Steigerwald & Agrawal, 2012). In this project two power estimation models were tested using the measurement systems available in the SEFLab and Green Lab: the energy estimation embedded in the VMware hypervisor (vSphere), and estimation models for high performance computers (RAPL & BMC energy).

Case: VMware Estimation of VM energy consumption

Measuring energy related data is more difficult in a virtualized environment due to extra software layer between hardware and the actual operating system. A first question therefore regards virtualisation overhead: how much energy is consumed by virtualisation itself (e.g. idle VMs)? Some data is provided by virtualisation vendors such as VMware, as part of their hypervisor interface vSphere. However, it is unclear whether such data is reliable for energy analysis, especially under different loads. This case-study aims at providing an estimation of virtualisation overhead, and contextually validate the reliability of hypervisor energy estimations.

Experiment

For the following experiment we have used a sample of common benchmark applications that offers a good coverage of the typical workflow encountered on the server side. The sample contains the following workloads: mail, database, webserver. The sample was chosen because they represent typical server tasks, such as serving a web page, executing database queries and providing email services. The main sources of data are two hardware sensors (outlet power meter WattsUp, and a fine grained monitoring system) and power data from the virtual machine hypervisor, vSphere. A WattsUp tool was connected to the server and the second one was connected to client machine. To get a sufficient amount of samples, two experiments were performed for 15 minutes each on every scenario of the experiment design. This experiment has provided data on (i) what the level of the virtualisation overhead is, and (ii) whether there is a difference between the energy values provided by the hypervisor and a physical outlet meter under different loads.

Results

We can see that the virtualisation overhead is present in each of the service types, by looking at the median value of the service as compared to the powered off state. While this is less dramatic for services such as database, power consumption, caused by the virtualisation overhead of simply starting the virtual machines, is larger for the mail service than for the database or webserver. However, the main difference between mail and the other two services is that it has Windows Server 2008 as an operating system on its VM, while all the other VMs of the rest of the services use SLES Linux. However, this is no proof, nor is it the aim of the experiment to prove differences in energy consumption generated by various operating systems, but we believe this could make a good starting point for future research.

As regards the vSphere reporting, we note that there is a systematic 10 watts overhead with respect to the values reported by the power outlet. However, apart from this error which can be easily corrected, the values are consistent across the different loads and they follow the distribution of those measured by the power meters (see figure 4.4). Correcting for the 10W overhead vSphere is within 1-2% accuracy compared to the hardware measurements. Note that vSphere is accurate regarding estimating energy usage by VM's, but is not designed for or suited to measure energy performance of particular applications.

Take-aways

Hence we can summarize our findings as follows:

1. There is a clear virtualisation overhead of leaving VMs powered on while idle, varying from 2 to 6%. Hence, we advise to shut down idle VMs whenever possible.
2. The power model of the hypervisor is precise enough to be used for rough power estimations and analysing dynamic energy behaviour under different loads. However, a validation is needed to remove potential systematic errors.
3. The estimation tool of VMware has a high level of accuracy; this tool could be used more widely by its users to estimate energy usage of VMs and possibly evaluate measures to reduce energy consumption.




Figure 4.4 Power measurements between VMware vSphere (black) and actual measurement (blue).

Case SURFsara: Power estimation in blade servers

Edit of an article of: Cedric Nugteren of SURFsara¹⁸

SURFsara wanted to explore to what extent developers of High Performance Computing software (HPCs) can help in reducing the energy consumption of HPCs. In this, the first question was to what extent existing power estimators RAPL and BMC give a good estimate of the energy consumption of supercomputers. To this end, the SEFLab installed sensors in one of the supercomputers of SURFsara.

Supercomputers are large systems that allow us to solve complex computing problems, such as weather and liquid simulations. Computers are becoming more efficient every year, but the power consumption of super computers remains a big problem. About 28% of the overall budget of a super computer (including hardware and staff) is spent on the electricity bill.

HPC programmers are currently optimizing their code in view of speed. This isn't just because speed is the most important measure, it is also the measure in which progress can be made easiest. Currently, there are two methods for measuring the power consumption: RAPL and BMC. RAPL stands for 'Running Average Power Limit'. RAPL is a software model used to estimate the power consumption based on various performance indicators in the CPU. Benefits of RAPL are the high sampling frequency



(1KHz) and the subdivision in components (such as memory, CPU). However, the model only takes the core components into account. The second method to measure the power is the Baseboard Management Controller, or BMC. This concerns an actual (so not software-based) measurement of the entire node (covering the disadvantage of RAPL). However, BMC has a much lower sample frequency (about 4Hz) and doesn't distinguish individual components.

18. <https://blog.surf.nl/en/superefficientcomputing/>

To allow for detailed power measurements, the SEFLab installed 12 sensors (for more information, please see Chapter 3) in SURFsara. The measuring system has a sampling rate of 30kHz and the sensors are placed in such a way that they can measure individual components. The Blade server was loaded with parallel running script during the test.

Results

The results of the experiments were as expected: RAPL is detailed but does not measure the complete system; and BMC is complete but not detailed. RAPL can keep up with the high sample rate of the measuring sensors. The BMC however cannot track the rapid frequency changes due to the low sample rate.

Figure 4.5 gives a summary, wherein the energy measured with RAPL is compared with SEFLab its energy measurements. In short: RAPL approximates the physically measured values for different loads relatively well, with deviations in energy consumption in respect of the SEFLab's setup of -3% to +17%, depending on component (CPU or MEM), and used load. That makes RAPL useful for rough estimates, which give programmers a first indication of the energy consumption of their application. Further research is necessary for a broader validation of the results, also under various use circumstances.

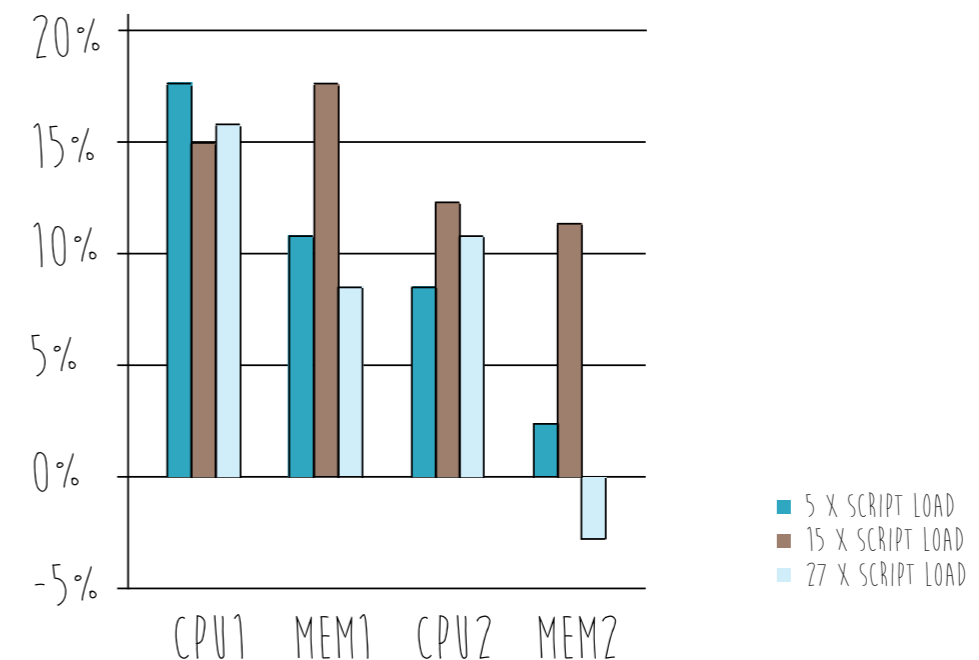


Figure 4.5 : Differences RAPL and SEFLab at different loads.



Jan-Willem Lammers
Lead Solutions Architect



What does your company do?

VMware's mission is to enable our customers' digital transformation by providing the software to run, manage, secure and connect any application across multiple clouds and multiple devices. My role is the Lead Solutions Architect for our Central Government Customers. I am also member of the CTO Ambassador programme to bridge our R&D organisation to our customers, academia and my peers.

Why did you participate in the Greening the Cloud project?

The Greening the Cloud project gives me the opportunity to test scenarios and methodologies that might lead to new or improved products. Also the ability to share my knowledge and discuss with a group of very motivated people makes this a wonderful experience.

What are the most valuable learning experiences from the project for you?

The interactions with the students and staff were very insightful. Most interactions led to follow up ideas or new insights. If we want to change anything significant we need to make messaging more relevant to decision makers so they can mandate and steer towards our common goals.

What would you recommend for a follow-up study regarding greening of software and clouds?

Look at how new application architectures using microservices and containers impact energy usage and efficiency. This is a strong movement and if we ride this wave timely we can make a big impact by putting best practices for green software for these technologies at the beginning of the adoption curve.

"The Greening the Cloud project gives me the opportunity to test scenarios and methodologies that might lead to new or improved products"

4.4 Software engineering practices

As introduced in the previous chapter, many sophisticated software power models and tools exist to estimate and predict the energy consumption of software applications through different parameters. However, this information has not been translated yet into practical guidelines for practitioners and developers to create energy-efficient software applications. A step in this direction has been made by Intel Corp (Larsson et al, 2011). that provided a number of guidelines and best practices for creating energy-efficient software. However, these guidelines targeted very low-level programming techniques (e.g. polling vs. event-based I/O). In addition, little to no validation has been performed on those practices, and their effectiveness in terms of energy consumption has not been quantified.

In this project, we conducted a number of experiments that resulted in empirical evidence on the energy impact of common software engineering practices, such as object-relational mapping (ORM), virtualisation, and tier separation. Such evidence, obviously, needs to be

contextualized in order to provide a consistent body of knowledge. Specifically, the hardware execution environment plays a very important role. For example, performing experimentation on mobile devices, servers, or embedded systems results in a variety of different insights, which can often be contradictory. This is why, aside from simply assessing the impact of best practices, we also try to explain the mechanisms behind such impact by measuring and analysing context-specific data.

We selected a small number of practices and applications as subjects – a decision that clearly poses an issue to generalisation of our results. The goal was, however, to increase the internal validity of our experimentation by focusing on the precision of measurements and rigorous data analysis. We then synthesized our results in clear, practical guidelines to create awareness amongst software developers and -architects about the energy impact of their decisions. Two cases on green software engineering practices will be presented: comparing query approaches (Diesveld Query Technology case) and comparing multitier applications in virtualized environments (REM Automatisering case).

Case: Energy efficiency of ORM frameworks (Diesveld Query Technologies)

Introduction

A large part of the facilities data centers offer is attributed to data storage, which makes it crucial to find a way to make these more energy-efficient. Structured Query Language (SQL) is one of the most used programming languages to manage data held in a relational database management system (RDBMS). Although SQL can yield very good performances, it is not structured to be easily integrated into object oriented programming (OOP) languages, often requiring the programmer to extensively use string concatenation for the creation of queries. Therefore, object-relational mapping (ORM) techniques were created in order to convert database elements that can be found in a SQL framework to objects that can be used in OOP. Even if ORMs introduce several benefits to manage data of a RDBMS, compared to a plain SQL approach, these techniques have several disadvantages. A problem that characterizes many ORMs is the creation of unnecessary objects that will significantly decrease the overall performance of the queries. The research described in this case focused on evaluating the relationship between various query frameworks used to access a database, their execution times for standard queries and their power consumption during a number of treatments.

Experiment

In order to be able to evaluate the energy consumption of RDBMS frameworks, apart from the underlying SQL, two different frameworks were selected: Propel¹⁹ and TinyQueries²⁰. Propel is a widely used ORM framework for PHP. TinyQueries is an ORM framework that is intended to solve the most common problems mentioned above, by keeping the performance and expressiveness of SQL and using a simple interface for queries, making the code more easily understandable by using an object-oriented approach. The investigated population is the relational database of the online travel guide Favoroute²¹. The chosen samples include the following tables: *GoogleTypes* (small size), *Agenda* (medium size), *Resources* (large size). The described experiment can be split into two different parts, according to its factors and relative treatments. A total of three different factors have been selected for the experiment, for a total of 9 distinct treatments.

Results

As a general consideration, effects of table sizes and query types on the outcome can be ruled out. A clear effect on performance of the three frameworks is identified, in line with the research expectations: as shown in Figure 4.6, SQL and TinyQueries have the best performance. Propel is significantly slower than the other two. This result was expected since TinyQueries and Propel are both based on the SQL programming language. Furthermore, Propel operates on objects instead of relying only on database tables. It is interesting to notice that the TinyQueries framework (created to keep both a high performance level and OOP simplicity) excelled in its purpose and it does not significantly differ, in terms of execution time, from plain SQL. On the contrary, the Propel framework was significantly slower than the other two.

19. Propel. <http://propelorm.org> Accessed: 10-04-2015.

20. Wouter Diesveld. TinyQueries. <http://www.tinyqueries.com> Accessed: 10-04-2015.

21. Favoroute, <https://www.favoroute.com> Accessed: 10-04-2015.

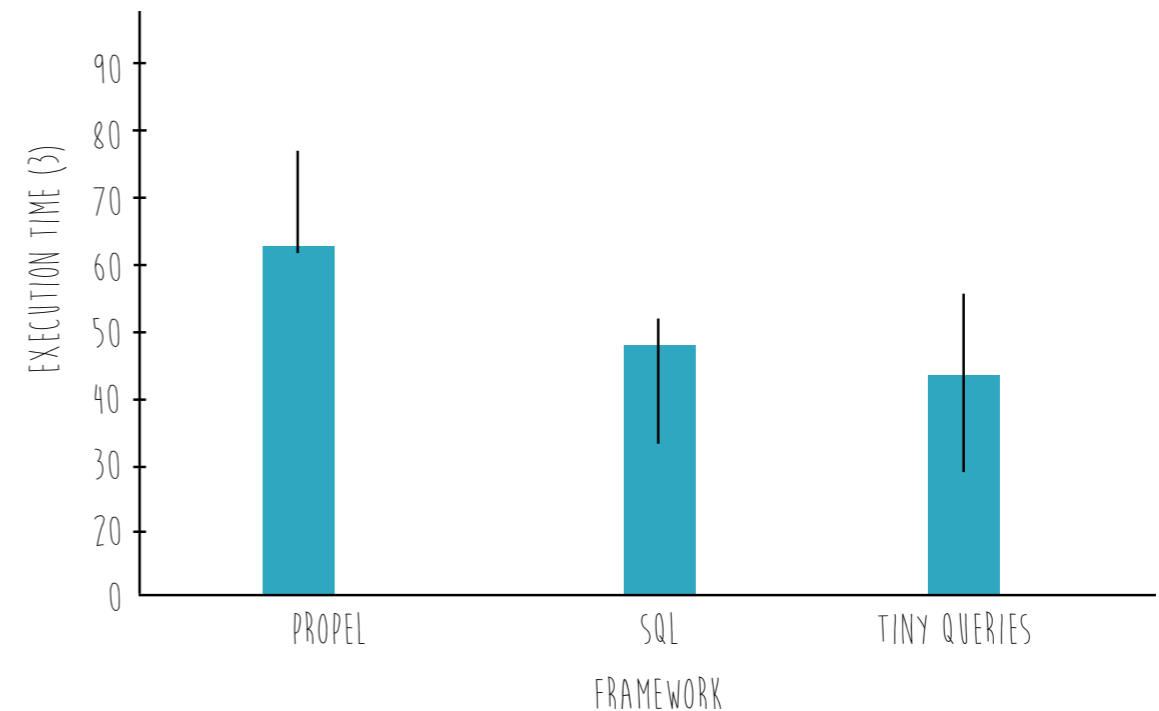


Figure 4.6 Representation of execution time per ORM framework .

In terms of power consumption, plain SQL seems to consume less on average, while Propel seems to consume more than TinyQueries. However, it is worth noticing that no statistically significant effect of the frameworks with respect to average power consumption was identified. Thus, execution time is the key factor that determines energy efficiency in this case: SQL clearly consumes less energy, as it spends less time to perform the same number of queries, while Propel is the least energy efficient. Hence, this is a clear situation where performance and energy efficiency are positively correlated. In Figure 4.7 a bar plot of energy consumption of the different frameworks is shown.

→ Case: Energy efficiency of ORM frameworks (Diesveld Query Technologies)

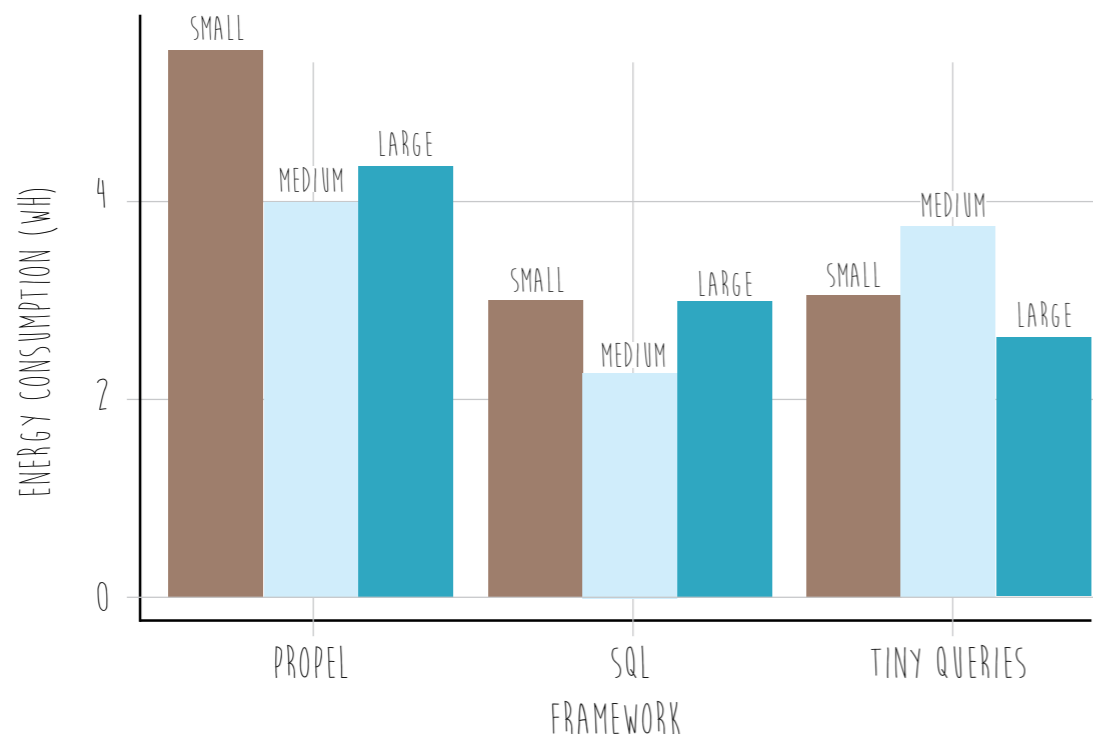


Figure 4.7 Energy consumption per ORM framework.

Conclusions

Generally, there has to be a trade-off between energy efficiency, the ease of use and implementation of an ORM framework. In the same context, trade-offs might also have to be made in terms of offered functionality of these frameworks. For almost all hypotheses, the execution time, as well as the power consumption, was generally higher for the Propel framework, while SQL was characterized by the lowest values.

Energy efficient database frameworks, even if an overhead might be necessary in order to deploy them, can have a big impact on power consumption. It is important to consider the fact that the implementation of queries is something that has to be done only once, while their energy efficiency will pay off every time a query is executed, which might be up to thousands of times per second.

Take-aways:

To achieve the best energy efficiency and performance of an application using a SQL database, implement the queries in plain SQL or adopt a “no-ORM” approach, such as TinyQueries.



Wouter Diesveld
founder and owner



What does your company do?

Query Technology specializes in optimizing database queries with the goal of making software faster by identifying bottlenecks and offering solutions to adjust the software. We focus on companies that develop software in-house. Significant speed improvements can be made here that could never be realised by purchasing new hardware. The goal of Query Technology is to optimise and utilise database queries as much as possible. Almost all databases have an enormous computing potential, it is a huge waste that so little use is made of this. In case of performance issues, the solution is often still sought in purchasing new hardware or migrating to Data Warehouses or Analytical Databases, whereas a lot of issues can be resolved by optimizing software.

Why did you participate in the Greening the Cloud project?

From my company, I issue advice about how to rewrite software in such a way that it works faster. I was curious whether – and if so, to what extent – these adjustments affect the power consumption of a computer. The Greening the Cloud project offers the perfect environment for that: the knowledge of universities and colleges can be applied to concrete situations that occur at companies. For me, measuring the exact energy demand when you adjust several variable factors in available query methods, was the primary focus of the study for me. If it turns out that this helps to save energy – and how much – that could be interesting for my customers.

What are the most valuable learning experiences from the project for you?

The case study we carried out within the project, is in line with situations I run into at customer environments. The results show that there indeed is a significant difference in energy consumption between various methods for communicating with databases, but it also shows that there is no ‘one-shoe-fits-all’ solution that applies all the time. Eventually, I could issue customers advice about how to save energy by adjusting software.

What would you recommend for a follow-up study regarding greening of software and clouds?

So far, the study was focused on relatively simple queries. I expect much bigger differences in energy consumption as the queries become more complex, since they often show that there also is a big difference in speed (execution time). A follow-up study could be very interesting and we are already exploring it with various parties.

PARTNER PERSPECTIVE

“Eventually, I could provide customers advice about how to save energy by adjusting software”

Case REM Automatisering: Different architectural setups

This experiment was aimed at finding ways to improve energy efficiency on a software level. This was done by looking at the open source OpenWave²² software. The OpenWave platform was developed in 2002 on a multitier architecture to manage license surveillance and enforcement regarding a large variety of permits, mostly governmental, in the Netherlands. The experiments are done by measuring power over time for different database configurations that the application readily supports. By looking at energy patterns with respect to this application, knowledge about general solutions might arise. Since the experiments only perform configuration changes on a database level, these changes might prove to be applicable in other software applications with multitier architecture as well.

A separation of the database tier from the application and presentation tier and their distribution over different virtual machines will be tested, as well as database memory configuration whereby updates to the database will be committed periodically. In addition to the different database configurations, the performance of the application is measured for these different configurations. This will indicate if the possible energy efficiency adjustments have any impact on the performance of the system.

Experiment:

Architectural Setup (AS):

AS-1: No change in architectural setup: The default architectural setup of the OpenWave application is used for the experiment, which means no separation of tiers will take place, i.e. all tiers will be hosted on a single VM.

AS-2: Separated architectural tiers into two virtual machines: The architectural setup of the OpenWave application is split into two parts with a focus on the database tier. The first part consists of a VM that runs the database tier only. The second part contains a VM running the API server (application tier) and the presentation server (presentation tier). The separation of the tiers is carried out by installing and setting up two distinct VMs and installing OpenWave on both of them while enabling and disabling the respective tiers on each VM in the configuration menu.

Database Configuration (DC):

DC-1: No change in database configurations: The database configuration as used by the OpenWave application is not changed in any way.

DC-2: Changed memory to database write interval: The database configuration regarding writing information to the database from the memory is changed.

The experiment is visually represented in Table 1. As can be seen from the table, the X denotes the number of trials run for each treatment of each factor. In order to balance the design an equal number of trials for each combination of factors' treatment will be conducted.

Results

The energy consumption per treatment is visualised in the barplot in Figure 1. All the values are between 8 and 9 Watthours. No significant difference is found for any of the factors. Regarding execution time a significant difference is found for both factors. In particular, the architectural setup seems to have a strong effect, with separated tiers performing better. Lastly, regarding power consumption no significant effect was found.

22. Openwave platform. <http://www.open-wave.nl/>. Accessed 09-09-2015.

		Factor 1 (AS): Architectural Setup	
		Treatment 1 (AS-1): No change in architectural setup	Treatment 2 (AS-2): Separated architectural tiers into two virtual machines
Factor 2 (DC): Database Configuration	Treatment 1 (DC-1): No change in database configuration	OpenWAVE 1...X	OpenWAVE 2X...3X
	Treatment 2 (DC-2): Changed memory to database write interval	OpenWave X...2X	OpenWave 3X...4X

Table 4.2 Experiment design of REM Automatisering case study.




Figure 4.8 Bar plot of the energy consumption per treatment.

Conclusions

Based on the results it can be concluded that there is no impact of architectural setup on energy consumption. There is a slight impact of performance when enabling database synchronisation. CPU usage will be higher when using separate virtual machines compared to a single virtual machine, while the execution time is lower when using separate virtual machines. Hence, the distributed configuration in separate virtual machines appears to be more energy-efficient than the consolidated one. This might also be due to clever optimisations in CPU usage performed by the hypervisor, as we later confirmed when interviewing a virtualisation expert.

Take-aways:

- Separate and distribute tiers over multiple VMs when possible, as the impact on energy consumption is negligible, whereas performance might significantly improve.
- Activate database synchronisation across multiple tiers, as this has no impact on energy consumption.



Albert Leenders
Technical Director



What does your company do?

Transfer Solutions is an ICT service provider, specialized in Oracle technology, Java and OutSystems. With these technologies, we realise ICT solutions that enable the business of our customers. We are located in Leerdam, but have customers throughout the country in a varying sectors such as government, industry, logistics and financial services. A great reference that has given us national fame is the realization of the ICT component for the new national large-scale topography map (<http://www.svb-bgt.nl/>).

Why did you participate in the Greening the Cloud project?

It's a great initiative to be working with our CSR goals in an entirely different way. In addition, we see that it may give you a competitive edge when offering services to potential customers, especially when you can substantiate it.

What are the most valuable learning experiences from the project for you?

At the moment of writing, we are still awaiting the end results of the study in which we are involved and have certain expectations about the energy consumption in various software technical environments. A practical learning experience for us has been that in taking part in these kind of initiatives it can sometimes prove to be difficult to free up time.

What would you recommend for a follow-up study regarding greening of software and clouds?

Keep the lines of communication open with the market parties involved and perhaps be facilitating in gaining additional incentives, such as a WBSO (Research and Development Promotion Act) subsidy.

"We see that it may give a competitive edge when offering services to potential customers, especially when you can substantiate it"

4.5 Conclusions

In this chapter we presented six case studies to analyse energy performance on the level of (i) virtualisation, (ii) power estimation models and (iii) software engineering practices. Regarding virtualisation, it was found that there were some differences in energy consumption between different hypervisors (under different conditions) and slight differences in containers versus virtual machines. Regarding the evaluated power estimation models, its relative high accuracy suggests that software developers and hosting providers should deploy these models more often in order to be able to evaluate energy performance and possible energy saving measures.

Lastly, several software engineering practices were evaluated, with two interesting outcomes: first, in line with our expectations, we showed that energy efficiency comes at a trade-off with other software qualities, such as flexibility and maintainability (e.g. in database queries); second, we found cloud applications can display counter-intuitive behaviour, showing for example that multi-tier applications distributed on separate virtual machines appear to be more energy-efficient and performant than the consolidated ones. These experiments provide first hints into energy efficiency trade-offs, which need to be further elaborated and validated in further experiments.

The case studies are described in reports and papers that are available at the Urban Technology programme.



5 STANDING ON THE SHOULDERS OF PUE

5.1 Introduction

In recent years the data center industry has developed several metrics as indicators for energy efficiency of its operations. Given the trend towards virtualisation and utilization of cloud environments for addressing our storage, computation and networking needs, it becomes increasingly important to develop indicators for assessing energy efficiency of cloud environments and cloud applications. However, the most widely adopted metric Power Utilization Efficiency (PUE) is insufficient. By focusing on loss of power between the power grid and the racks, it is agnostic of inefficiencies in the other elements in the data center: server hardware, virtualisation and application software. On the other hand, the PUE has some appealing properties, such as reducing the energy efficiency of a complex layer to a single number. In this chapter we present a new energy efficiency model that addresses all layers in

the stack by introducing a PUE-like metric for each layer. They can be used separately or combined, depending on one's sphere of influence over the stack. Thus, by empowering each stakeholder to optimise locally, the model drives overall energy efficiency.

5.2 PUE strengths and limitations

Good indicators are needed that show where inefficiencies reside, how big they are and if optimisation efforts are effective. Despite the growing research interest from academia and practice, such indicators have not been formulated in a satisfying way yet. Arguably the most successful metric to date is the widely adopted Power Utilisation Efficiency (PUE). It expresses the energy overhead (or loss) of the data center infrastructure by comparing the energy consumed from the power grid and the energy delivered to the server racks.

The PUE metric has appealing properties: it is intuitive, conceptually simple and reduces the inefficiency of various aspects of the complex data center infrastructure to a single number. Moreover it does not make assumptions about technological aspects of the infrastructure and it is designed to indicate improvements without enforcing a specific technological solution for optimisation.

PUE has a number of properties for an indicator the industry needs:

- **Partially applicable.** Allow assessment of the individual layers and provide an instrument for steering without necessarily having full-stack control.
- **Simple.** Consolidate to a single numeric indicator per layer that intuitively grasps the energy efficiency challenge of that layer.
- **Practical.** Align with current activities and efforts of practitioners. Be based on available data as much as possible.

- **Comparable.** The indicator should not only allow for relative comparison of the same installation over time but also facilitate relative benchmark comparisons.
- **Implementation agnostic.** The indicator should make as little assumptions about the technological implementation as possible, focusing on the externally observable behaviour of a layer in the stack.
- **Independent of expert opinion.** Make the assessment repeatable by avoiding dependency on expert knowledge.

A major limitation of the PUE metric, however, is that it provides only a very partial indication of the efficiency of a data center. In fact, it only considers the technology infrastructure. The energy delivered to the racks is consumed by a much more complex computational stack (see Figure 5.1) designed to eventually generate useful functionality for the end-user. There is substantial inefficiency in all of the elements in the stack,

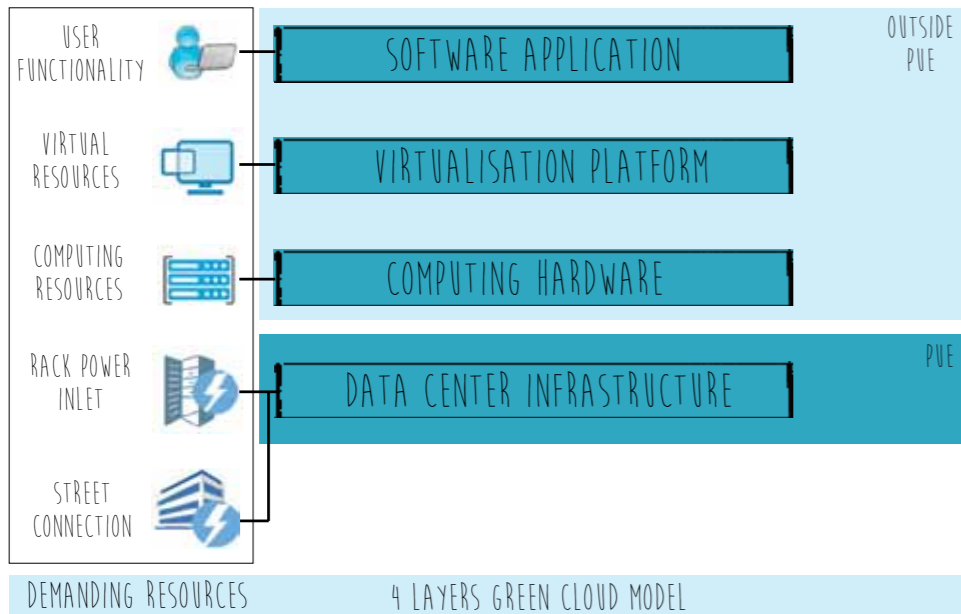


Figure 5.1 Infrastructure layers stack: What is in scope vs what is out of scope of PUE.

which include: computing hardware (server and networking hardware components), virtualisation platform (providing virtualisation and monitoring infrastructure), and software applications (including middleware and application software). All of these elements are out of the equation of PUE.

Metrics beyond PUE: state of art

Others have observed the unsuitability of the PUE to provide insight in the energy loss of the whole stack. Attempts have been made to develop a broadly usable metric to detect inefficiencies, compare progress in optimisations and compare different sites and applications. Industry attempts to arrive at suitable metrics include The GreenGrid that investigated possible Data Center energy Productivity (DCeP) (TGG, 2014) metrics for relating data center energy usage to useful work (Kooimey, 2011). However, they did not find a suitable indicator that would allow to measure productivity in a way independent of the application and that could be aggregated to the whole data center. Other metrics like FVER, CADE en DPPE (L. Newcome et al; Uptime Institute & McKinsey: 2008, Japan National Body; 2012) have been proposed which build upon the PUE in the data center by including the IT equipment. Main limitation here is that the metrics fail to account for actual delivered work to the end-user. (Visser et al., 2013) have proposed an evaluation framework to label the energy footprint of applications using three simple indicators: (i) the total energy used per year, (ii) the average energy consumed per unit-of-work, and (iii) the relation between optimal and average energy consumption per unit-of-work. For the purpose of this study we suggest to use only the last of these indicators to compare functionally different applications.

Academic research also contributed a number of metrics as extension to PUE. For example, Cappiello et al. proposed a suite of metrics as a result of the Eco2Clouds project (Cappiello et al., 2013). These metrics address three different layers of the computing stack: Infrastructure, Virtualisation and Application. Infrastructure metrics are further classified into site (i.e. data center) and host (i.e. single server). Each layer includes about four to five metrics. This division resembles the one we propose. However, the large number of metrics introduces a high degree of complexity, which makes the adoption of the proposed

metric suite difficult in practice. This creates the need for a dedicated data analytics platform that combines the metrics into meaningful indicators for decision-making. As a matter of fact, the proposed metrics are aimed at runtime monitoring, while our model aims at providing offline indicators.

5.3 GreenCloud Model

Our model for data center energy efficiency offers a set of indicators which build upon the popularity of the largely successful PUE. The model favours simplicity and pragmatism. Such properties are captured in the support for four layers and a few essential indicators provided by each of the layers. To be useful, energy efficiency indicators for data centers need to accommodate the separation of concern that is typically found in these data centers.

A main characteristic of the proposed model, is its coverage of the full stack of the data center, organised into four layers as indicated in figure 5.1:

1. data center infrastructure
2. computing hardware
3. virtualisation infrastructure
4. application software

Each of these layers has a distinct objective for optimisation, and management of these layers is typically organisationally separated. For example, hardware procurement often does not know or does not care about the concerns of application software developers. In turn, application developers will not know how virtual machines are mapped to hardware – which is the main point of hardware abstraction in cloud computing.

We propose to divide the data center stack into separate layers and assess each layer in isolation with a single indicator. This allows assessing how each layer contributes to the overall energy efficiency of the data center without having to assess the full-stack efficiency directly. This also enables the stakeholders of each layer to act independently and makes it evident where in the stack additional optimisation would be most beneficial. As we do not try to measure the

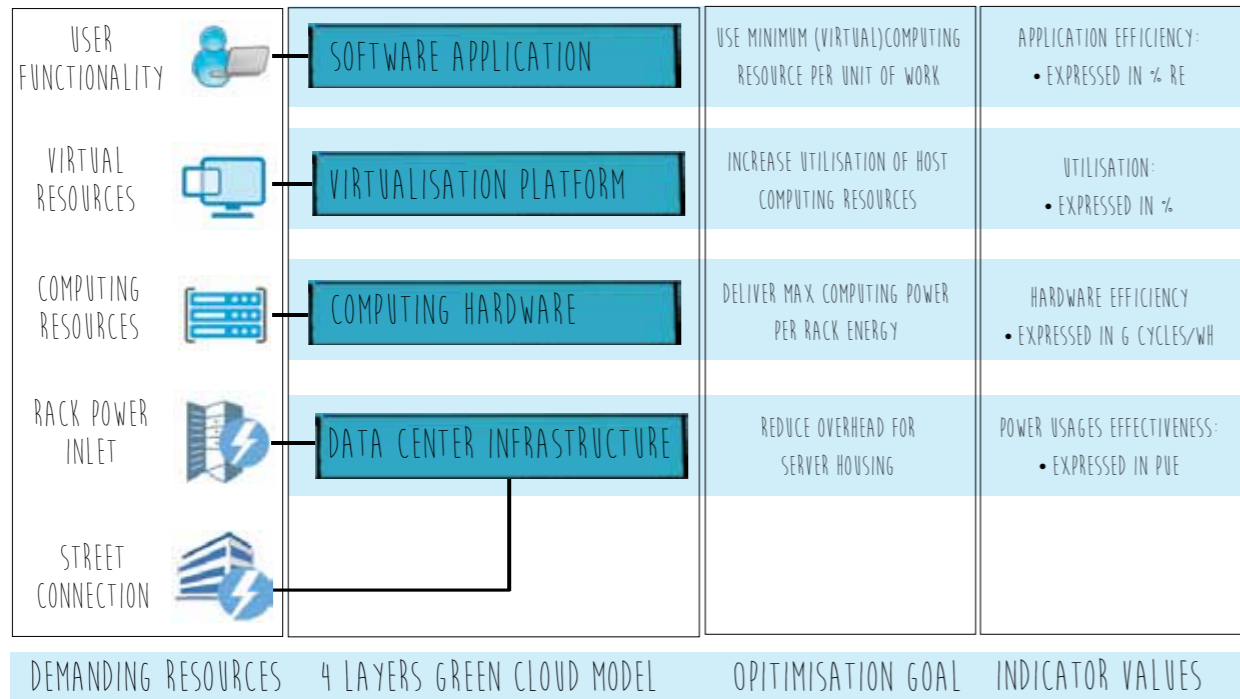


Figure 5.2 The four layers of the efficiency mode, optimisation goals and proposed indicators.

full stack as a whole, there is no need for detailed attribution of energy consumed from the grid to the particular processes in the upper layers of the stack. Optimisations in the upper layers of the stack can be made and assessed independently.

Figure 5.2 shows the four layers and how they are connected: the data center infrastructure, as characterized by the PUE metric, covers everything between the street level power grid connection and the electricity cable to the server racks. The computing hardware layer encompasses all networking, storage and computing hardware including rack-side UPSs. It consumes electricity and delivers physical computing resources to

the Virtualisation layer. The latter transforms physical into virtual resources. Finally, the software application layer consists of the whole software stack from the operating system virtual machine up to what is needed to the application logic that serves the user requests from the data network.

The four layers of the GreenCloud Model focus on the energy need of a specific application, not challenging the need or value of the application itself. Neither includes the model the quality of the energy used, what might be considered as a last effort to green the cloud; purchasing the (remaining) energy from sustainable sources.



The Data center infrastructure layer

This layer is covered by the PUE metric and is already extensively discussed. Its main objective is to reduce overhead of non-computing infrastructure. It is measured as the difference between energy consumed from the power grid and electric energy delivered at the racks. The formula is:

$$PUE = \text{Facility energy} / \text{IT energy}$$

Where a lower value is better and 1 is the optimum.



The Computing hardware layer

The computing hardware layer encompasses all computing hardware components like storage, network devices or compute servers and includes also rack-side UPS. The efficiency objective in this layer is to deliver physical computing resources (CPU cycles, memory, IOps, ...) for a minimal energy budget and the indicator would be the amount of computing resources that are used effectively versus the total energy consumed:

$$\text{Hardware efficiency} = E / \text{CPU}_{pu}$$

$$\text{CPU}_{pu} = \text{number of used physical CPU cycles}$$

E = consumed electric energy (in Joule or Wh) at the rack level.

The value is the energy used per CPU cycle. Lower values are better. Hardware efficiency is typically determined by either vendor fact sheets, benchmark workload results or actual production data. We propose to use the latter. In contrast to similar metrics, our indicator for hardware efficiency is dependent on the actual utilisation of the hardware. It measures actual used computing resources instead of potential (maximum) resources a component could deliver, similar as with the PUE. The hardware (or infrastructure in case of PUE) should be sized according to the actual needs and not to prospected levels of ambitious business plans. The indicator in this way stimulates the alignment to factual utilisation levels: if utilisation is highly fluctuating, it is important to strive for hardware that can scale its energy consumption proportionally (L. Barroso et al., 2011) with the workload (a property called energy proportionality. While workload charac-

teristics are not always known at the moment of hardware procurement, the indicator effectively points out inefficiencies in a sound way.

We choose to use CPU cycles as the proxy for overall resource utilisation, rather than other resources (e.g. memory, disk space), mainly because (i) CPU is still the dominant energy consumer in today's hardware and (ii) in a virtualized environment all other devices are mediated by the CPU.



The virtualisation layer

As mentioned above, we assume all resources are virtualized and we treat the management of that virtualisation as a separate layer. Virtualisation allows to tailor virtual machines to the need of the application and to consolidate workload across physical machines. Its main objective is to improve utilisation of the underlying hardware resources and to do so without introducing much overhead by the virtualisation mechanism itself.

The indicator we propose for this layer is again based upon utilisation, but adjusted to also cover the overhead introduced by the virtualisation platform. The virtualisation layer consumes physical computing resources and it produces virtual computing resources. Hence the formula goes as follows:

$$\text{Virtualisation layer efficiency} = \text{CPU}_{vu} / \text{CPU}_{pa}$$

$$\text{CPU}_{pa} = \text{number of available physical CPU cycles}$$

$$\text{CPU}_{vu} = \text{number of used virtual CPU cycles}$$

A high value is better and 1 (100%) is the theoretical optimum. The reason why the used and the available CPU cycles of the hardware to calculate hardware utilisation are not compared, is that it would not allow measuring the loss of the virtualisation layer itself. The objective is to only measure the used physical CPU cycles that are actually consumed by the virtual machines. Virtual resource consumption (especially CPU) is usually readily available in the virtualisation management software. CPU is used as a proxy for other computing resources following the same reasoning as above.

The application software layer

Application software has a large influence on energy consumption of data centers. However, software does not consume electricity directly and the contribution of software to energy consumption of the system is not directly observable. In order to identify the effect of software design on energy efficiency, we need to attribute the energy consumed by the hardware to the application or even an isolated piece of code that is under development. This attribution is often difficult. Especially in the presence of virtualisation, where different applications share the same hardware, it is very hard to establish a precise relationship between computing resources used by an application and the energy drawn at the rack level.

In our model we propose to focus on reducing the demand for computing resources. One challenge to measure application efficiency, is that it needs a quantification of the amount of work it executes. For this purpose, we introduce the notion of 'unit of work' or 'task'. Depending on the type of software application, a unit of work can be an email, a web request, a financial transaction, a rendered picture or anything else that represents measurable outputs of the application. Ideally, there should only be one type of unit of work per application that is representative for all the workload. If not applicable, a weighting factor can be introduced to normalize the different occurrences. In order to determine energy efficiency, the consumed resources are subsequently divided by the number of delivered units-of-work. However, this allows only to track improvements of a specific application over time. Since applications differ, they will have different units-of-work, which makes efficiency difficult to compare.

We propose to use a metric that abstracts away the differences: the relative efficiency (RE). Despite its name, the relative efficiency is a metric that allows comparing different applications. The underlying idea is that any application is typically more efficient on certain load levels than others. I.e., it has an "optimal" load level where it uses the minimal amount of resources to deliver a unit-of-work. Any other load levels is less efficient. The average over a longer period of time represents the baseline per unit-of-work. The ratio between the optimum and the average resource

consumption per unit-of-work serves as an indicator for the overall application efficiency. It is an indicator that allows for a comparison between applications:

$$\text{Relative efficiency} = \text{RTmin} / \text{RTavg}$$

RTmin = minimal amount of computing resources per unit-of-work

RTavg = average amount of computing resources per unit-of-work over a significant period.

A high value is better, and 1 (=100%) is the optimum. We again consider used (virtual) CPU cycles as proxy for all the other consumed resources. This gives us an efficiency indicator for the application that is dimensionless and does not include the application-specific unit-of-work. This in turn makes the application efficiency comparable to other applications. Still, the metric is somewhat partial in the sense that it allows for comparing applications on this aspect, but there will be application level optimisations that are not covered with this indicator. We would like to stimulate dedicated research in this area.

5.4 Industry benchmarking: making sense of numbers

While the aforementioned indicators give us a way to assess the efficiency of individual layers with a single, reasonably easy to compute metric, they do not express "how good" a specific value actually is. To this end, it would make sense to compare the values for specific types of installations with other, similar installations in the industry. While this does not express the potential for further improvements compared to an absolute optimum, it does enable stakeholders to challenge their performance on the various layers of the model.

The indicators in our model have been designed explicitly to make this comparison possible. We calculate the benchmark score for every indicator separately by comparing the actual value with the available data points in the benchmark collection of comparable installations. Thresholds are chosen in such a way that the population is divided into parts and a star-rating can be given to express whether a specific layer




Figure 5.3 An outcome example for all layers with the single indicator value and the benchmarked score.

implementation scores better or worse compared to similar cases. Figure 3 shows what the values for the model could look like if all layers of a data center have been assessed.

5.5 Conclusion

The presented GreenCloud Model provides a framework for assessing the energy performance of the complete stack of data centers. It does so with a limited amount of coherent metrics, which are relatively simple; is largely derivable from available sources in IT operations and enables the comparison of IT-applications at four different levels. As such it provides both

a monitoring instrument as well as indications for potential improvements to increase energy efficiency of the stack or of individual layers in the stack.

The definition of indicators given here are still work in progress. More precise and reliable indicators will be defined once the model has been applied and validated in more in-depth, real world case studies. Even though the details may change, we are convinced that having such a simplified visibility of the energy efficiency of the elements of the data center stack is possible and would greatly improve transparency and ultimately optimisation efforts.

6 KNOWLEDGE VALORISATION

Within research programs of universities of applied sciences the focus is on carrying out applied research and valorisation of the developed knowledge. Valorisation is the process in which developed knowledge is converted into (commercially) viable products, processes or services. Within this project, valorisation took place in several ways: the development of a cost module, the development of a measuring platform for education and the founding of the Amsterdam Sensor Lab.

6.1 Development of a cost module

Within this project, several experiments were carried out that have provided insight into energy consumption and energy saving potential by making adjustments to variables in cloud software and cloud environments. But what do possible energy savings of a specific application mean for an organisation and how does this translate into costs? To gain more insight into the cost

consequences of more energy efficient software choices, a cost module has been developed. This cost module contains the main variables that effect the way energy savings of an application are translated into a cost indication for a specific organisation. Given the complexity of this matter, the results are to be considered indicative and are mainly to be seen as a means to increase the awareness among ICT users regarding the energy cost involved in ICT.

The development of the cost module was done in close collaboration with Verdonk, Klooster & Associates (VKA). VKA consults customers in the field of ICT about optimisation, risk management and cost reduction. Developing a service focused on energy efficient software fits in this profile. Customers that have an (above) average ICT service package have been chosen as a target audience for the cost modules. The assumption that these customers have little

knowledge about the energy costs associated with ICT, was confirmed during sessions with VKA and some of its customers. There is little insight into the contribution of individual software applications to energy cost and how more efficient software choices may translate into cost savings. A cost module can contribute to increasing awareness about the energy consumption of ICT by providing insight into the relationship between energy efficiency and costs.

Realistic options

The aim of the cost module was to map realistic options for these customers. There are literally hundreds of options within frequently used cloud applications that can be adjusted and that affect energy consumption. To make a selection of a number of more realistic options, the following criteria were used:

Potential for savings: This criterion indicates the level of savings and impact that are feasible for a specific option. For instance, an application that is rarely used and / or requires little computing power, will have less potential for energy savings. Logically, the search is mainly for large savings to realise added value for customers.

Level of interdependence: The level of interdependence is an indication of how interdependent a savings option is with the entire stack. A high level of interdependence makes it harder to establish the potential energy savings, as it is difficult to assign measured energy savings to a particular software part (OS, application, driver). The choice of operating system for instance, is strongly interwoven with all kinds of applications in the stack. As a result of that, it is difficult to assign the savings potential at this level as well as to provide nuanced advice about which operating system performs substantially better with regard to energy consumption. The encryption techniques investigated at Greenhost (see description later in this chapter) are much less interdependent and therefore offer more leads for calculating realistic energy savings.

Number of options: The number of options that must be calculated forms another selection criterion. The more options (or versions) exist of an application or alternatives, the more time it will require to connect a nuanced savings potential to it. There are dozens of encryption technologies that are significantly more

time consuming to compare than the three DNS resolvers that were calculated earlier at the SEFLab for SURFnet. (A.Hankel et al, 2016)

Impact within an organisation: A fourth selection criterion is related to the resistance that may exist within an organisation against introducing a particular savings option. For instance, switching to another (more energy efficient) client management system or browser, will have major implications for an organisation if employees work frequently with these applications.

These selection criteria were used to map a number of applications so as to identify attractive options to study. Figure 6.1 provides an overview of these options in which three variables are included: number of options (x-axis), interdependence (y-axis) and savings potential (size of dot). A number of savings options have such a level of interdependence (e.g. choice of hypervisors, shift towards containers), that it is complex to come to uniform outcomes of energy savings potential. Other savings can be achieved with new query techniques, but this requires extensive testing of a large amount of available alternatives. The 'sweet spot' lies in the bottom-left hand corner: low interdependence and a limited amount of options to test; preferably with significant savings potential.

Input-variables for the cost module

The purpose of the cost module is to realise a generic model that allows for calculating the savings potential at various organisations. This requires the inclusion of generic criteria that can translate the energy savings into a cost indication. The energy savings observed during tests within the SEFLab and Green Lab must be multiplied by a number of factors:

- **Workload factor:** Relationship between the tested workload and the daily workload of the organisation under study. Here a translation was made into the energy consumption of the tested workload (e.g. 10 synchronisation steps or 100 encryption instances) to a realistic measure for how often this occurs (on average) in an organisation on a daily basis.
- **Use intensity:** Per employee, the intensity with which an application is used will generally vary. A subdivision by light, medium and heavy users was made.

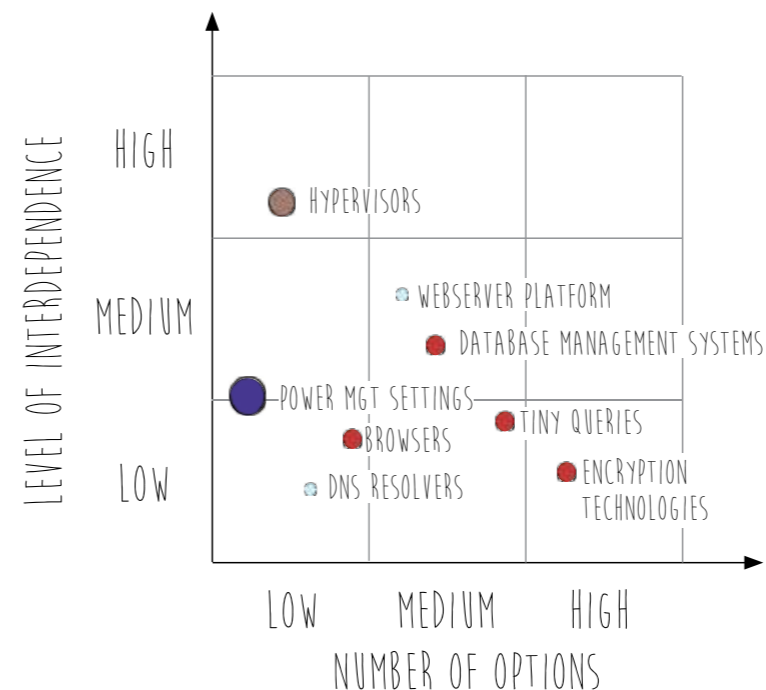


Figure 6.1 Number of options versus level of interdependence.

- **Number of employees:** The number of employees of the organisation, subdivided into light/medium/heavy user categories was included.
- **Energy aspects:** This covers both the energy price and which energy mix is used (for determination of the CO₂ footprint).

The output of the module includes an indication of (i) energy costs (both current and saved) and (ii) CO₂ emissions. This was based on an energy cost price for hosting companies (10 eurocents per kWh) and the Dutch energy mix, in which – in the module – various options (such as sustainable energy) can be selected.

Next steps

The cost module offers opportunities to translate energy savings for an application into cost indications for a company. This makes it possible to further create awareness among ICT companies about the costs of their choices. The module is rudimentary and the assumptions must be

validated further in practice. In addition, the limitations of the measuring setup apply: for each measurement it should be considered that measurements may differ on other hardware platforms and in combination with other operating systems (than the one applied). More tests are necessary to be able to formulate generic statements, both concerning energy consumption and cost implications. In addition to energy costs, complementary cost benefits could be included – in line with the Logius case²³ – such as reduced use of hardware, maintenance costs, licenses and housing-related costs.

Savings cases

Within this project, two savings options have been considered: (i) power management options within exchange, such as the effects of disabling synchronisation of agenda, mail and address book outside of office hours²⁴ and (ii) energy consumption of various encryption technologies. The latter case, provided by Greenhost, is discussed in 6.2.

23. kngs.wdfiles.com/local--files/start/DigiD.pdf.

24. The results of this study are available at the HvA and the website of Greening the Cloud.



André Hendriks
Co-owner VKA
(Verdonck, Klooster
& Associates)



What does your company do?

At VKA we are committed to make ICT works for people. We are a strategic ICT consultancy and specialize in complex issues. In the end, it is all about people, so we don't just employ technical specialists at VKA, but we also have employees who are able to 'read' an organisation. In collaboration with our customers, they realise successful projects that make sure that ICT does what it was built for: making life easier with smarter, more efficient and faster solutions.

Why did you participate in the Greening the Cloud project?

Both the use and importance of ICT (from the cloud) is growing substantially while the importance of energy (and emission) savings is also increasingly recognized. These two elements were decisive for me. I've noticed that in the ICT industry cost considerations are hardly ever made properly. We approach our customers from a more business economics perspective of ICT issues. The cost of energy must also be included more often in these analyses.

What are the most valuable learning experience of the project for you?

Interestingly enough two contradictions. On the one hand the confirmation that energy consumption of ICT really does not play a part in the choices made in the ICT domain, whereas on the other hand, attention to energy consumption is actually of great significance.

What would you recommend for subsequent studies regarding the greening of software and clouds?

Further explore which aspects can realise the most concrete impact in greening the IT of organisations. Implementing these with organisations with large-scale IT environments will make the impact resonate much more.

"Both the use and importance of ICT (from the cloud) is growing substantially while the importance of energy (and emission) savings is also increasingly recognized"

PARTNER PERSPECTIVE

Case: Greenhost and encryption technology

Introduction

In a cloud network, mobile phones, tablets, laptops and other user devices are connected – via the internet – to applications and central storage media, often hosted in a data center. This connection can be secured by means of various encryption protocols. In this process, data (text, images, films) is encrypted before it is sent to the website visitor each time a web server sends a request for this data. (see figure 6.2).

Greenhost is a hosting company that has studied the security of various encryption protocols a great deal. In addition to the level of security, Greenhost is also interested in the energy consumption involved in the various protocols. After all, encrypting the communication between client and server requires additional computing power, which results in increased energy consumption. The question is how much energy is consumed per encryption protocol, so that, in addition to security, the energy consumption can be considered in the choice for an encryption protocol.




Figure 6.2 Simple schematic of SSL encryption process.



Experiment

Within the SEFLab, the energy consumption of several common, so-called SSL/TLS²⁵ encryption algorithms has been measured. Encryption technologies consist of a combination of Key Exchange, authentication and the encryption method. In the end, ten of these combinations were studied and their energy consumption was benchmarked. The experiment covered the measurement of the energy consumption in encrypting and sending a 260 MB document, which was compared to the sending of the same, non-encrypted document. Meanwhile the energy consumption of main components of the server were measured (only CPU and memory, while energy consumption of the client and the network were not considered). To achieve significant results, the tests were carried out at least ten times and the average energy consumption was taken as an indicator.

Open SSL Name	Energy consumption (Joule)	Duration (seconds)
No encryption	97.9	4.13
RSA RC4-128 MD5	147.6	5.75
RSA RC4-128 SHA	150.2	5.81
RSA 3DES-EDE-CBC-168 SHA	502.6	15.84
RSA AES-128-CBC SHA	147.8	5.67
DHE AES-128-CBC SHA	150.9	5.92
RSA AES-256-CBC SHA	161.9	6.14
DHE AES-256-CBC SHA	163.3	6.21
ECDHE AES-128-CBC SHA	148.4	5.73
ECDHE AES-128-CBC SHA	163.2	6.20
ECDHE AES-128-GCM SHA256	245.5	9.40

Table 6.1 Overview of encryption technologies studied.

25. SSL stands for Secure Socket Layer; TLS for Transport Socket Layer; they form the most common security protocols. With the help of these software keys, the communication between several servers and / or clients can be encrypted with various encryption algorithms.

Results

Table 6.1 provides an overview of the ten tested encryption algorithms compared to the non-encrypted transfer of the same amount of data, where besides energy consumption (second column in Joule) and duration (in seconds) are displayed. It shows how most encryption techniques consume 1.5 to 5 times more energy than non-encrypted data transfers (97.8 Joule versus 145.5 to 502.6 Joule). It is also notable that various encryption algorithms consume significantly more time compared to non-encrypted data transfer, varying with a factor of 1.4 to 4 times (4.13 seconds versus 5.67 to 15.84 seconds). It was found that the process of encryption is the main cause for these differences; while key exchange and authentication prove to lead to relatively small amounts of additional energy consumption. In addition to looking at energy consumption, the relationship between speed and safety of the various algorithms was studied. For hosting providers, choosing the right encryption technology is a trade-off between security, power, consumption and speed.

Open SSL Name	Energy per year (kWh)	Costs per year
No encryption	1273	€ 127
RSA RC4-128 MD5	1919	€ 192
RSA RC4-128 SHA	1953	€ 195
RSA 3DES-EDE-CBC-168 SHA	6535	€ 654
RSA AES-128-CBC SHA	1922	€ 192
DHE AES-128-CBC SHA	1963	€ 196
RSA AES-256-CBC SHA	2105	€ 210
DHE AES-256-CBC SHA	2123	€ 212
ECDHE AES-128-CBC SHA	1930	€ 193
ECDHE AES-128-CBC SHA	2122	€ 212
ECDHE AES-128-GCM SHA256	3192	€ 319

Table 6.2 Annual energy consumption and costs based on 33.3 million daily page views at nu.nl.



From energy to cost: case of NU.nl

Encrypting a file does not require a great deal of energy and choosing more energy efficient encryption may only save several kilowatthour (kWh). Nevertheless, encryption is constantly being applied, millions of times a day, which could accumulate to more significant levels of savings. Based on the results of the study of encryption for Greenhost, the energy data of various encryption technologies was applied to a frequently used website (NU.nl) to translate into energy costs for a hosting provider. We assumed the following:

- For the average size of a message to be encrypted by a webserver, a value of 1 MB was used.
- The number of visitors combined with the number of page views is an indication for the load of the webserver(s): in the case of NU.nl it concerns 2.5 million unique visitors a day that generate 33 million page views per day.
- The size of the file and the energy consumption of the encryption algorithm are assumed to be correlate linearly – in which the figures of Greenhost (based on the test with the previously described file of 260 MB) is used as the basis (no data is available about the encryption of smaller files).

Based on these assumptions, table 6.2 provides an overview of the cost implications for the hosting provider hosting NU.nl, should it use various encryption technologies. Assuming an energy price of 10 eurocents per kWh, the differences are small: an efficient encryption technology can lead to several hundreds of euros of savings on a yearly basis – probably too little to use this as a steering argument to shift to other encryption algorithms. Possibly adding client and network costs (currently not included) would offer some additional options for cost reduction as well as the energy consumption of other components in the server than were tested in the current study (beside CPU and memory the fans and motherboard were not tested), but it is not expected that this will lead to significant cost implications. Nevertheless, the methodology of translating energy savings to costs seems plausible and should be tested and validated in other cases, preferably with higher potential for savings.

Conclusions

The case of encryption technology offers interesting insights into the possibilities of measuring the energy consumption, the relationship between security and speed and the translation into costs. Although encryption technologies seem to have little energy saving potential for one individual website, it is recommended to carry out a broader cost study (given the broad use of encryption), to include not only network and client costs, but also to consider reduced hardware, licensing and maintenance requirements. In order to strengthen the module several assumptions should be tested, such as the linearity of energy consumption and file size. Lastly, more cases should be studied in order to come to more appealing cases for the industry.





Mart van Santen
Owner Greenhost

What does your company do?
Greenhost is a webhosting company. Since 2007 we focus mainly on making our services sustainable. Together with software solutions. Take for instance our webhosting-platform where we focus on the energy efficiency of the system. Together with software solutions such as the use of virtualisation – we try to make our services as efficiently as possible and achieve the most sustainable hosting solutions. We are also working on providing more insight for our customers in their energy consumption on the web hosting platform. This increases awareness and motivates them to be more efficient in deploying their software. In addition, we commit to a free internet and freedom of speech by supporting other initiatives. Either with Free Press Unlimited or other similar projects.

Why did you participate?
I think it is important to know what is happening in the field of green software. You must keep stimulating yourself and the Greening the Cloud project seemed perfectly suited for that. Not only did it allow us to actively contribute to developments, but it also offered us the opportunity to engage with industry and academics, to share experiences and knowledge, and help ensure Sustainable IT remains significant and take it to the next level.

What are the most valuable learning experiences?
To me, an interesting result was the confirmation that certain, unsafe, old settings, appeared to be the most energy consuming settings. The most useful result is that it is more efficient to switch on an AES coprocessor of a server. This is a piece of hardware that accelerates the encryption and, as expected, it proved the process consumes significantly less energy this way. Although we were already on the right track, our aim is now to make some adjustments in our infrastructure based on these results!

What would you recommend for a follow-up study?
Greening the Cloud has motivated us to see where we can improve further, mainly by studying the other case studies (such as the Query Technology case about database query optimisation). That case study is an example of what we can address as Greenhost on the platform. In addition, we will continue with our efforts on providing more insight into the energy consumption of a website for customers. Initiatives such as Greening the Cloud make options for improvement more concrete and we hope to make even more progress with this.

"Greening the Cloud has motivated us to see where we can improve further by studying the other case studies"

PARTNER PERSPECTIVE

6.2 Translation into education

One of the goals of this project was to make a translation of the project results into education and, at the same time, engage students during the study. Both have been successfully applied within the project.

One of the main results of this project has been the completion of a practical setup, allowing software engineering and technical information students to carry

out energy measurements. Within the Faculty Digital Media and Creative Industry (FDMCI) this setup has been used for the class Project Agile Development (PAD) over the past two years. In 2015, the assignment for the 60 students (subdivided into 10 teams) was building an energy efficient video streaming app. The team that won the price was not only able to realise a strikingly low energy consumption, both absolute and in comparison to the other teams (up to 50% better), but was also able to make the streaming

software very compact. Other teams required at least twice as much storage space. The context illustrates that smart programming pays off and the significant improvements are possible²⁶. At the VU the measuring setup is used in the curriculum as well, more specifically by the Green Lab (see box).

forcement of the research qualities of students and supervising teachers. At the VU the case studies resulted in new knowledge about green design strategies that will, in time, be used in the Green Lab course.

6.3 Amsterdam Sensor Lab

In addition to the measuring setup, dozens of students of the HvA, VU and UvA have been involved in the context of the Greening the Cloud project. They have carried out literature studies, done experiments with the partners involved and contributed to the development of the cost model and the validation of the Green-Cloud Model. As such, the project contributes to rein-

A spin-off of this project has been the creation of the Amsterdam Sensor Lab (ASL). Building on the expertise acquired within the Software Energy Footprint lab (SEFLab), application specific sensors and sensor systems that are not available of the shelf will be developed within the ASL. The knowledge that was gained during the Greening the Cloud project will

Green Lab

The Green Lab is a master course in the Software Engineering and Green IT programme of the Computer Science master of the VU. The course aims at teaching students the fundamentals of Empirical Software Engineering in the context of energy efficiency. As such students work in teams to perform experiments on software energy consumption in a controlled environment. They carry out all the phases of empirical experimentation, from experiment design to operation, data analysis and reporting.

The Green Lab is also a research laboratory, currently located in the Medialab of the Network Institute. The lab hosts a number of servers, instrumented with state-of-the-art energy sensors (partly provided by the SEFLab), to run our experiments on software energy efficiency. Students contribute to the research by carrying out experiments, such as the cases of REM Automatisering, VMware, Diesveld Query Technology and Cobra Systems within the Greening the Cloud project. Yearly tens of students are enrolled in this programme; for which the Green Lab is always looking for case studies from industry to study.



Photo: The Green Lab team consists of Patricia Lake, Giuseppe Procaccianti, Remco Havermans, Nina Wolfram



26. <http://www.isptoday.nl/achtergrond/datacenter-energie-efficiency-the-next-steps/>

be deployed and the developed measuring platform will be the first project in the “showroom” of the ASL.

During this project, knowledge was gained about measuring sensors, embedded software, data acquisition systems, connection with a measuring PC and the measuring software on a PC. This knowledge is crucial for developing all kinds of measuring sensors.

The heart of ASL is technology, with the emphasis on electrical engineering and computer science. The SEFLab measuring platform is an example in which electrical engineers and computer scientists, with the support of students who use the lab as a learning environment, work on usable measurement systems.

The sensor solutions of ASL are developed in longer term projects. The primary focus is on supporting the researchers of the Urban Technology programme of the HvA in generating relevant data that may contribute to the development of new solutions for a sustainable, liveable and connected city. In addition, partnerships are sought with partners of the Urban Technology program, to support them in developing customized measuring systems. In many cases this will concern measuring processes that take place in the city.

Whether it concerns material flows, transport flows or energy flows; we find or develop the right measuring system. In particular, with the rapid development of sensor and communication options, there is a lot of demand for dedicated sensor solutions, where more standard sensor solutions lack the features that are required.

Specific attention in the ASL is paid to the traceability and repeatability of calibration, verification, maintenance, modification and commissioning. As such, the ASL provides measuring data, adheres to privacy regulations and tries to make the data of the sensors safely available to the research community.

Smart tile

Currently, within the ASL “smart tiles” are being developed: a series of tiles with extensive sensory and network capacities and power supplies that can be interconnected. In this phase, the focus lies on the pressure sensor tile, a standard Amsterdam street tile with 4 to 900 pressure sensors. The aim of the tile is to map pedestrian and cyclist flows. Other smart tiles that are part of the programme are the solar movement power tile (power for other tiles) and a talking tile (containing several interface protocols to IoT networks).



Patrick de Ridder
Project manager
and consultant



What does your company do?

Cobra Systems connects the physical with the virtual world. We do this, for instance, by investigating for our customers, how many of the online visitors have actually visited their shop and how much time they spent there or by determining the effect of the indoor climate on the well-being of the visitors. Because of the large amount of data we analyse this results in a lot of information and thus value for our customers. In practice, we use web and wireless solutions, and we do not shy away from using the soldering iron and the 3D printer either. When we need something specific that can't be supplied, we make it ourselves.

Why did you participate in the Greening the Cloud project?

Everything we do must lead to the most efficient solution. Even if it doesn't seem to have any commercial benefits (yet), or when it means that we need to engage in discussions with our stakeholders. We have seen that this can lead to interesting, and sometimes unexpected results. For us, this turned out to be the case with the 'green' theme. By using the green label more explicitly, we have discovered that it is not only a selling point but that it also helps to meet like-minded people and exchange knowledge more easily. If it also helps to make the world a better place in the process, that is a great bonus of course.

What are the most valuable learning experiences from the project for you?

A truly sustainable solution sometimes radically differs from what people are used to and often falls outside of the regular familiar market. In other words: there is no customer base for it yet. Support by means of investment in time and money and marketing are essential at this stage.

What would you recommend for a follow-up study regarding greening of software and clouds?

For a follow-up study, we think it would be interesting for the results and recommendations to be applied outside of a laboratory environment as well. This would allow for large-scale measurements of the effect on an operational environment. We also do this for our own customers in other areas. We would love to contribute this knowledge in a follow-up of this project.

"We would love to contribute [our] knowledge in a follow-up of this project"

PARTNER PERSPECTIVE

7 CONCLUSIONS AND NEXT STEPS

In this concluding chapter an overview of the most prominent findings is presented (7.1), as well as some more general observations (7.2) leading to a collaboration agenda with 10 next steps (7.3).

7.1 Main project findings

I Leveraging the potential of green software

In recent years the industry has made significant steps in energy efficiency. This has partially been achieved through efficient cooling, hardware renewal and a shift towards virtualisation. The overall energy consumption in data centers is, however, likely to keep growing significantly in the coming years. The energy loss chain of data centers shows in which way efficient cooling can be adopted as an end-of-pipe solution, whereas wasted energy through inefficient software is at the root of energy usage in data centers. As such, developing green software and green clouds have an

energy-saving potential of more than a factor 7-10 over the complete chain and provide a major leverage for significant greening of the ICT industry.

II Facilitating accurate energy measurements

Many of the partners and industry experts that were consulted during this project emphasized the importance of having measurement platforms such as the SEFLab and Green Lab to carry out experiments and benchmarks. The measuring platforms constitute a necessary condition to determine the energy efficiency of specific software solutions and evaluate green design strategies. The project has enabled further professionalization, development of a client-server set up, remote access functionality, and development of a calibration setup. As such, the lab facilities have proven their potential in providing clear experimental results and are ready to be used for large scale measurement.

III Developing green design strategies

Although green software and green clouds have been discussed amongst industry professionals as well as in academia, the amount of practical experiments and practical results is still limited. This project provides the results of ten case studies, leading to new insights about energy efficiency. New knowledge has been generated concerning the energy consumption of data hypervisors, containers compared to virtual machines, encryption technologies, query technologies and web platforms. The cases have resulted in a number of practical implications and design strategies, which require further generalisation in the future. More importantly, they provide a clear way forward in generating knowledge of green software and clouds, but also provide best practices for creating awareness in the industry.

IV Acknowledging the value of available estimation models

Software based estimation tools are available to support programmers in evaluating energy performance of for instance virtual machines (VMware vSphere) and high performance computers (Running Average Power Limit or RAPL). The tests carried out during this project have provided evidence that these models are quite sophisticated, provide rough estimations and trend lines of energy usage, and as such can support programmers to make smarter/greener choices in software development. The SEFLab and Green Lab however remain important complementary facilities to carry out more detailed analyses, develop detailed knowledge on software energy usage and green coding strategies as well as support in validating the value of energy estimation models.

V Positive effects of virtualisation techniques

Clouds are relying more and more on virtualisation technologies, either using virtual machines or software containers. We see a clear tendency to adopt the lightweight approach of containerisation, e.g. Docker solutions. The energy performance of these virtualisation techniques differ depending on the type of application deployed in the cloud environment. The studies carried out in this project, either in company settings or in the SEFLab, have identified the benefits of properly choosing the most appropriate technology as function of the computation and memory footprint.

VI Scoring clouds on energy performance

Given the current limitations of energy metrics used in industry, this project has developed a model (Green-Cloud Model) that enables the industry to evaluate and monitor the energy performance of their cloud environments. The model adds to the current PUE metric (data center level) by including a number of metrics for the hardware, virtualisation and application layer. It is designed to include a number of metrics that are readily available in most data centers and can be used to assess individual layers on energy performance. It is practical, provides steering guidelines and does not require additional data collection.

VII Translating energy savings to costs

Whereas the results from lab measurements are known in kWh, end users or hosting providers are more focused on cost savings. During this project, the first strides were made in developing a cost model that provides hints as to how energy savings through using more efficient software translate to cost savings for an organisation. In future, the cost model should include other potential cost benefits due to more efficient software, to validate the hypothesis that 1 kWh energy savings can add up to 24 euro cost savings due to reduction of hardware, licensing, maintenance and housing (see chapter 2 of this report).

7.2 General observations

In addition to the practical results of this project, more general observations can be made.

A Limited priority for greening

Based on discussions with relevant stakeholders in the industry, we conclude that the greening of software and clouds are still a too low priority in the industry. 'Being green' is used as a distinguishing factor by a number of hosting providers that were consulted, and was mentioned to have some added value to particular market segments such as governments and NGO's. Nevertheless, the majority of hosting providers mention that their customers prioritize performance, uptime and (low) costs. There are some early signs that green hosting is becoming more important, partly due to the high share of energy costs for hosting providers. For greening of clouds and software to become more mainstream, it is important to create awareness amongst end users about the energy (and cost) saving potential of greening clouds and software.

Greenit
amsterdam region

Jaak Vlasveld
Director



What does your company do?

Green IT Amsterdam is a non-profit research consortium of over 40 organisations. We make the energy transition possible with ICT. To this end, we bring parties together through networking and knowledge sharing activities, we scout new technologies, initiate pilots and create opportunities to make the Metropole Region Amsterdam more sustainable. As director of the foundation I am responsible for the performance of our project agency and the health of our foundation. In addition, I'm also responsible for identifying key innovation themes such as green software and cloud applications.

Why did you participate in the Greening the Cloud project?

The topic of green software and cloud applications was put on the agenda by several of our participants. It was previously addressed within the knowledge network and the project Cluster Green Software in which we have also participated. The interest shown by participants proves the potential of the theme and the unique strength we have in our region to realise that potential. Therefore, we would like to see that the project is strongly anchored in our region and in our other activities in the region and far beyond.

What are the most valuable learning experiences from the project for you?

The collaboration between SME and researchers is interesting to witness up close. The resilience with which the various perspectives and time lines are being approached is contagious. In addition, it actually leads to concrete and valuable cases, where practice and theory find each other. Also, outcomes show that the development of knowledge concerning green software and cloud allows for a lot of follow-up actions. The resulting cases increase the awareness among both the broader public and within the ICT industry, allowing for diverse spin-off directions.

What would you recommend for a follow-up study regarding greening of software and clouds?

We notice that there is a great need for clear cases which identify opportunities for broad application to increase impact. That focus is important and the design, in which collaboration is possible within a broad consortium, is interesting. A follow-up initiative is the GreenServe project, in which we identify and develop easily accessible cases concerning power management in collaboration with public organisations. When the SME and procurers of public organisations can find each more easily on matters of sustainable ICT, the next step in 'greening' might just be around the corner. In Europe at least, there is great demand for this, as shown in, for instance, the ICTFootprint.eu project (www.ictfootprint.eu).

"The interest shown by participants proves the potential of the theme and the unique strength we have in our region to realise that potential"

PARTNER PERSPECTIVE

B Limited awareness of energy consumption

Our research confirmed that the majority of software developers, systems integrators, data centers and/or end users have limited to no insight in the actual energy consumption of their software or cloud. Several hosting companies have some indications of their energy use at rack level and even some indications on virtual machines, but the validity of these indications remains unclear. Limited awareness of energy consumption is likely to hamper investments in innovative ways to reduce energy consumption and achieve cost reductions. It will be necessary to conduct additional experiments and to draw up and present best practices in order to stimulate and encourage the industry to look deeper into the opportunities offered by more energy efficient software and clouds.

C Split incentives in industry

This project confirmed that there are split incentives amongst key players in the ICT value chain to achieve energy and cost savings. Software developers are rarely requested (and therefore not incentivized) to adhere to energy performance criteria. End users have limited awareness about energy costs of their ICT operations, and thus have limited incentives to steer on energy performance. Hosting providers have limited influence on software choices of their clients, but also usually deploy a business model that stimulates selling more computational resources, with limited incentives to stimulate green design strategies.

A joint effort is required to line up insight and incentives for all involved parties to facilitate greener deployment of clouds and software. Creating awareness, facilitating discussion with stakeholders and presenting best practices in forums such as Knowledge Network Green Software and via industry organizations such as Nederland ICT, ISPCoconnect and Green IT Amsterdam Region can play an important facilitating role.

D Limited research

A last observation relates to the limited attention of this research topic in both the international research agendas of Horizon 2020 and in the Netherlands of the Topsector (subsidy) program. Although there is increasing awareness about the importance and potential impact of greener software (e.g. in data center, KNGS, Routekaart), there is limited traction to stimulate further

research in a more programmatic way. Apart from the GreenServe project, continuing the execution of energy testing of software, there is limited funding for this type of research. This is particularly striking given the current strong position the joint universities in the Amsterdam region (VU, UvA, HvA) have in this field and the need to maintain this leading role. We strongly believe that a more structured research programme on relevant topics and sufficient support from both industry as well as subsidy programs will provide the opportunity to strengthen the (international) knowledge position in this field, support the industry in tackling energy performance at the root of the ICT chain and work towards a new industry standard when it comes to green IT.

7.3 Collaboration Agenda for Green Software and Clouds

In order to identify next steps, a collaboration agenda between industry and academia is suggested to synchronize incentives and interests and define a number of viable actions to further strengthen the knowledge of and implementation of green design strategies for software and clouds.

- 1. Carry out experiments:** Set up an extensive experimentation effort on a range of, amongst others, representative software applications, design strategies, and cloud technologies. The set of experiments should be defined in close collaboration with the industry and should take into account the latest relevant scientific research.
- 2. Further develop measurement platforms:** Build labs and platforms that enable large scale testing where software applications, design and cloud environments can be evaluated in a controlled manner. Further work should focus on development of representative hardware platforms, increasing accuracy of sensors, test client-server systems and further extend remote access functionality.
- 3. Assess the validity of power estimation models:** Define conditions under which these models are useful and identify ways to make them helpful for programmers during software development processes.

4. Test the GreenCloud Model in practice:

Establish the validity of the GreenCloud Model and identify the value for stakeholders in monitoring the energy efficiency of their cloud environments. Also present the model in international venues with the objective to create traction for its use by industry.

- 5. Develop advanced cost models:** Translate experimental results (measured in kWh) into financial metrics. This requires developing more thorough data center cost models that include the translation of energy savings to tangible benefits, such as reduced hardware requirements, maintenance and licensing.

- 6. Line up the ICT chain programme to un-split incentives:** Gather leaders in the industry across the ICT chain in order to set up a joint effort to line up incentives to apply greener software, including defining agreements on sharing costs and benefits.

- 7. Increase participation of the industry:** Create support in (i) defining relevant research questions (demand articulation), (ii) providing hardware/platforms for testing, (iii) licensing and lowering barriers to test applications

- 8. Connect to education and students:** Build up a curriculum and training modules for students and trainers to translate research findings into practical guidelines; including the development of educational test platforms and lab facilities.

- 9. Create an awareness campaign:** Increase awareness by showcasing the energy and cost saving potential of green software and green clouds; in close collaboration with industry and industry organisations.

- 10. Set up a funding programme for research on green software and green clouds:** Enable further research that develops fundamental and applied knowledge of the topic of efficient software and clouds, in close collaboration with industry.



Photo Green IT Leaders Live 2015, an example of knowledge sharing, awareness and collaborative initiatives. Project presented early draft of GreenCloud Model here with room for feedback and discussion.

Acknowledgements

For this research we are greatly indebted to a large number of universities, researchers, partners, industry organisations and students for contributing to this applied research project.

A special thanks goes out to Giuseppe Procaccianti and Patricia Lago (of the VU) for their collaboration and contributions and Paola Grosso and Arie Taal (of the UvA) for providing the scientific angle, carrying out a large number of case study experiments and co-developing the GreenCloud Model. Also special thanks to Kay Grosskop, Michiel Cuijpers and Joost Visser of the Software Improvement Group for the lead development role of the GreenCloud Model and general support in the development of the SEFLab.

Many thanks to the partners of the project Almende, CloudProvider, Cobra Systems, Diesveld Query Technology, Greenhost, GreenInvents, Leaseweb, REM Automatisering, Schuberg Philis, SURFnet, TransferSolutions, Verdonck, Klooster & Associates and VMware.

Also the contribution of the dissemination and network partners Nederland ICT, Green IT Amsterdam Region and ISPConnect for their feedback and support in the dissemination of our results to the industry is greatly appreciated.

The following people have co-developed the lab, performed measurements or have contributed in other ways:

Researchers on this project: Eric Hoekstra, Bo Merkus, Miguel Ferreira (SIG), Emile Nijssen (Efficient Electronics), Gerben Hoogendorp, David Zeegers, Edwin Boer en Rob Schaake.

University of Applied Science (HvA) Students: Marco van Veen, Vincent Tseng, Bram Visser, Ferdinand van der Gullik, Martin Droog, Remy Bien, Allard Barends, Justin de Jong, Bram Verhoef, Patrick Kruijzer, Thomas van Ast, Horatio Wilson, Jorn van Gijn, Koen ten Brincke, Tom Flipse, Wesley Delmeer, Jesper van Tol en Corné Oudshoorn. Additionally, approximately 100 students have performed tests using SEFLab's measuring system as part of project PAD initiated by HVA DMCI domain.

Vrije Universiteit Amsterdam (VU) students: René Aparicio Saez, Anastasija Efremovska, Remco Havermans, Nina Wolfram, Gerard van Hoesel, Mihai Dodan, Mylene van der Koogh, Veaceslav Munteanu, Roberto Verdecchia, Adriatik Bedjeti, Jiafei Cai, Giulia Colacicco, Rio Essed, Bojan Filipović, Hassan Jalil, Harry Obaseki, Alex Risis.

University of Amsterdam (UVA) students: Casper van der Poll, Yu Ri Tan, Jeroen van Kessel, Sam Ansmink, Fahimeh Alizadeh.

Lastly we would like to thank SIA for making this research possible.



Literature list

CE Delft M.R. (Maarten), Afman T. (Thijs) Scholten "Trends ICT en Energie 2013–2030"; 2016

http://www.ce.nl/?go=home.downloadPub&id=1736&file=CE_Delft_3F48_Trends_ICT_en_Energie_2013_2030_DEF.pdf

Tebodin, "ICT stroomt door"; 2007

<https://www.rijksoverheid.nl/documenten/rapporten/2008/01/29/ict-stroomt-door>

Dirk Harryvan, "Zervers 'Zuinige Servers"; 2014

https://www.amsterdam.nl/publish/pages/610528/rapport_zervers_final.pdf

Niklaus Wirth, "A Plea for Lean Software", Computer 28; 1995

<http://dl.acm.org/citation.cfm?id=620152>

Parthasarathy Ranganathan "Recipe for Efficiency: Principles of Power-Aware Computing", Communications of the ACM; 2010

<http://cacm.acm.org/magazines/2010/4/81500-recipe-for-efficiency-principles-of-power-aware-computing/fulltext>

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; 2010

<http://dl.acm.org/citation.cfm?id=1852078>

Jason Flinn, M. Satyanarayanan, "Energy-aware adaptation for mobile applications"; 1999

<http://www.cs.virginia.edu/~zaher/classes/CS656/flinn.pdf>

N. Thiagarajan, G. Aggarwal, A. Nicoara, D. Boneh, J.P. Singh, "Who killed my battery?: analysing mobile browser energy consumption", In Proceedings of the 21st International Conference on World Wide Web (WWW '12), Lyon, France, 41–50; 2012

T. Simunic, L. Benini, G. De Micheli, "Cycle-Accurate Simulation of Energy Consumption in Embedded Systems", Proceedings of 36th Design Automation Conference; 1999

<http://dl.acm.org/citation.cfm?id=310090>

A. Peymandoust, T. Simunic and G.D. Micheli, "Low Power Embedded Software Optimisation using Symbolic Algebra", In IEEE Proceeding of the 2002 Design, Automation and Test in Europe Conference and Exhibition; 2002

https://www.researchgate.net/publication/232646608_Low_Power_Embedded_Software_Optimisation_Using_Symbolic_Algebra

Patricia Lago, Giuseppe Procaccianti, Grace A. Lewis, "Green Architectural Tactics for the Cloud", In 2014 IEEE/IFIP Conference on Software Architecture (WICSA); 2014

<http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=6827097>

Giuseppe Procaccianti, "Energy Efficient Software"; 2015

https://www.researchgate.net/publication/275348368_Energy-Efficient_Software

Yu Ri Tan, "Accents in literature with respect to green clouds"; 2015

<https://ivi.fnwi.uva.nl/sne/wp/wp-content/uploads/2016/02/scriptie-Yu-Ri-Tan.pdf>

R. A. Saez, "Energy efficiency of Virtualizers", VU/UVA master literature study; 2015

<https://ivi.fnwi.uva.nl/sne/wp/wp-content/uploads/2016/02/scriptie-Rene-Aparicio-Saez.pdf>

C. van der Poll, "Resource usage in hypervisors", UVA bachelor thesis; 2015

<https://ivi.fnwi.uva.nl/sne/wp/wp-content/uploads/2016/02/scriptie-Casper-van-der-Poll.pdf>

M. Warnaar, "Cloud Energy Consumption Measurement and Reduction: an Overview of Methods", VU/UVA master literature study; 2016

<https://ivi.fnwi.uva.nl/sne/wp/wp-content/uploads/2016/04/scriptie-Martin-Warnaar.pdf>

J. van Kessel, "Power Efficiency of Hypervisor-based Virtualisation versus Container-based Virtualisation" SNE UVA master thesis; 2016

<https://ivi.fnwi.uva.nl/sne/wp/wp-content/uploads/2016/04/thesis-Jeroen-van-Kessel.pdf>

Dzmitry Kliazovich, Pascal Bouvry, and Samee Ullah Khan, "Greencloud: a packet-level simulator of energy-aware cloud computing data centers", The Journal of Supercomputing, pp. 1263–1283; 2012

<http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=5683561>

A. Nouredinne, R. Rouvoy, L. Seinturier, "A Review of Energy Measurement Approaches"; 2013

<https://hal.inria.fr/hal-00912996v2/document>

E. Jagroep, J.M.E.M. van der Werf, S. Jansen, M. Ferreira, J. Visser, "Profiling Energy Profilers", Proceedings of the 30th Annual ACM Symposium on Applied Computing. New York, NY, USA: ACM pp. 2198–2203; 2015

<http://dl.acm.org/citation.cfm?id=2695825&dl=ACM&coll=DL&CFID=802224943&CFTOKEN=71918913>
Schubert, L. and Jeffery, K. (2012). Advances in clouds – research in future cloud computing; 2012
www.dphu.org/uploads/attachements/books/books_3885_0.pdf

W. Lewis, D. Dobbs, G. Veerapillai, "Software testing and continuous quality improvement"; 2008

<https://www.crcpress.com/Software-Testing-and-Continuous-Quality-Improvement-Third-Edition/Lewis/p/book/9781420080735>

B. Steigerwald, A. Agrawal, "Developing Green Software". Intel White Paper. Intel Corporation; 2011

https://software.intel.com/sites/default/files/developing_green_software.pdf

P. Larsson, "Energy-efficient software guidelines", Intel Software Solutions Group, Tech Rep. software.intel.com; 2011

<http://software.intel.com/sites/default/files/article/57237/energy-efficient-software-guidelines-v3-4-10-11-1.pdf>

Albert Hankel, Eric Hoekstra, Robert van den Hoed and Roland van Rijswijk-Deij. "Measuring Software Efficiency – Presenting a Methodology and Case Study on DNS Resolvers" In Proceedings of the 18th Mediterranean Electrotechnical Conference – MELECON; 2016

J. Koomey, "Growth in data center electricity use 2005 to 2010", Oakland, CA: Analytics Press; 2011

http://www.analyticspress.com/data_centers.html

De Marco: "You can't control what you can't measure", in Controlling Software Projects, Management Measurement & Estimation, p. 3; 1982

The Green Grid (TGG), "Harmonizing Global Metrics for Data Center Energy Efficiency", Appendix G – Data Center Energy Productivity (DCeP); 2014

<http://www.thegreengrid.org/~media/Regulatory/HarmonizingGlobalMetricsforDataCenterEnergyEfficiency.pdf?lang=en>

L. Newcombe et.al, "Data centre Fixed to Variable Energy Ratio metric DC-FVER", British Computer Society; 2012

http://www.bcs.org/upload/pdf/dc_fver_metric_v1.0.pdf

Uptime Institute & McKinsey, "CADE: Corporate average Data Center Efficiency"; 2008

<http://searchdatacenter.techtarget.com/definition/CADE-Corporate-Average-Data-center-Efficiency>

Japan National Body/Green IT Promotion Council, "DPPE: Holistic Framework for Data Center Energy Efficiency"; 2012

http://home.jeita.or.jp/greenit-pc/topics/release/pdf/dppe_e_20120824.pdf

Joost Visser, Jeroen Arnoldus, Joris Gresnigt, Kay GrossKop, "Energy efficiency Indicators for e-Services", in proceedings of the 2nd International Workshop on Green and Sustainable Software; 2013

C. Capiello, S. Datre, M.G. Fugini, M. Gienger, P. Melià, B. Pernici, A. Tenschert, "Monitoring and Assessing Energy Consumption and CO₂ Emissions in Cloud-based Systems", special session at IEEE International Conference on Systems, Man and Cybernetics (SMC) in Manchester; 2013

The Green Grid (TGG – Victor Avelar, Dan Azevedo, Alan French (Editors)), "PUE: A comprehensive examination of the metric"; 2012

http://www.thegreengrid.org/~media/WhitePapers/WP49-PUE%20A%20Comprehensive%20Examination%20of%20the%20Metric_v6.pdf?lang=en

Luiz André Barroso and Urs Hölzle, "The case for energy-proportional computing", Computer, 40(12):33–37; 2007.

<http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=4404806>

J. Koomey, S. Berard, M. Sanchez, H. Wong, "Implications of Historical Trends in the Electrical Efficiency of Computing", IEEE Annals of the History of Computing; 2010

<http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=5440129>

