

BACHELOR INFORMATICA

UvA  UNIVERSITEIT VAN AMSTERDAM

Accents in literature with respect to green clouds

Greening the cloud project

Yu Ri Tan

June 17, 2015

Supervisor(s): Dr. Paola Grosso (UvA) & Dr. Arie Taal (UvA)

Signed: Dr. Paola Grosso (UvA) & Dr. Arie Taal (UvA)

Abstract

In the last few decades the use of cloud services has increased dramatically. Therefore the energy efficiency of data centers becomes more and more important. The EU has composed a document with best practices regarding energy efficiency. In this thesis we discuss the usefulness of this document and try to tighten or improve it. First an overview is given to see where the accents lie in scientific literature with respect to data center efficiency. Then, a systematic literature review is used to see if the findings in the scientific literature correspond with the best practices proposed. The focus of the SLR lies on the management of a data center. Other areas than management must also be examined in the future to able to provide better best practices.

Contents

1	Introduction	5
2	Related Work	7
2.1	Introduction to Cloud services and energy efficiency	7
2.2	EU best Practices	9
2.3	Criticism on EU best practices	10
3	Literature overview	13
3.1	Research question	13
3.2	Finding literature	13
3.3	Classification	14
3.3.1	Overview	15
3.3.2	A more specific approach	16
4	Literature review	19
4.1	Research question	19
4.2	Methodology	19
4.2.1	Inclusion and Exclusion criteria	20
4.2.2	Search process	20
5	Research findings	21
5.1	Overview of the findings	21
5.1.1	Handling workload	22
5.1.2	Data center Hardware	22
5.1.3	Prediction and Self Learning	24
5.1.4	Metrics	24
5.1.5	Goals	25
5.2	Comparison with EU best practices	25
5.2.1	Relevant Information from the EU best practices	25
5.2.2	Improving EU best practices	25
6	Conclusion & Discussion	29
A	Primary studies SLR	33

Introduction

In the last few decades the speed of the Internet as well as the number of (high speed) networks has increased dramatically. Cloud services became therefore much more attractive. Today many of us can't live without these cloud services such as Google Drive, Dropbox or iCloud. Cloud storage is of course not the only possible type of service provided by the cloud. There are three main cloud services: SAAS (Software As A Service) services like Facebook or Dropbox, PAAS (Platform As A Service) services such as the Google App Engine and IAAS (Infrastructure As A Service) services such as Amazon EC2. These services give the user the freedom to use the resources in their own way, with their software and implementation of choice. Because of the increasing amount of services provided by the cloud, the term XAAS is regularly used. XAAS stands for Everything, Anything or "X" As A Service. This abbreviation refers to the endless amount of services delivered by the cloud, which is the core purpose of cloud computing.

The increasing amount of services resulted in a growth of the number of data centers. A data center is a facility that consists of many computers that are connected by routers and switches [1]. With the appropriate virtualization software it is possible to use these servers together as one or multiple cloud services. Even multiple data centers located in different places around the world can work together due to this software.

These data centers use a lot of energy. This represented around 1.3% of the global electricity use in 2010 [2]. According to data centerDynamics 2012 Global Census, the power requirements grew between 2011 and 2012 with 63% [3]. This becomes therefore an increasingly important topic. The power consumption is not only important for the electricity costs but also for the environmental friendliness. The terms environmental friendliness or eco-friendly will be used as synonyms. The environmental friendliness does not only depend on the energy efficiency but also on the source of the energy that is used. This approach really makes sense for the future.

The number of cloud services and the penetration of those services grow [4]. That is why it becomes more and more important to know how eco-friendly those services are. But how can we define or compare the environmental friendliness or energy efficiency of different cloud services? The European Union composed a document with best practices regarding energy efficiency of data centers. Is this document still up to date with the scientific literature?

In this thesis, three research questions are covered:

RQ 1: *Which aspects and to what extent are covered in research literature regarding greenness or eco-friendliness of cloud services?*

RQ 2: *What do we know so far regarding data center management?*

RQ 3: *In what way and how could these findings help improving the EU's best practices?*

First, a brief overview of energy consuming components of cloud services are given. After that, the best practices document composed by the European Union is reviewed. Next is a literature overview. Here the first research question will be answered by presenting classification tree, to give an overview of the accents in scientific literature with respect to green clouds. Then one subject is chosen to study in detail. This subject is "Management" as shown in Research question 2. To do this, a Systematic Literature Review is used. The results of this review will be discussed afterwards. Can we use these results to improve or tighten the EU's best practices?

This research contributes to the "Greening the Cloud" project [5]. The "Greening the Cloud" project was founded in 2012 to give more insight in the energy efficiency of cloud services. This project is emerged from a cooperation between Software Improvement Group (SIG) and the Hogeschool van Amsterdam (HvA).

Related Work

This chapter gives a brief overview about what we know of the energy efficiency of cloud services. After that, the best practices composed by the European Union are discussed and reviewed.

2.1 Introduction to Cloud services and energy efficiency

In a data center many servers and switches are installed. These components are vertically organized in racks. All these racks together are a data center.

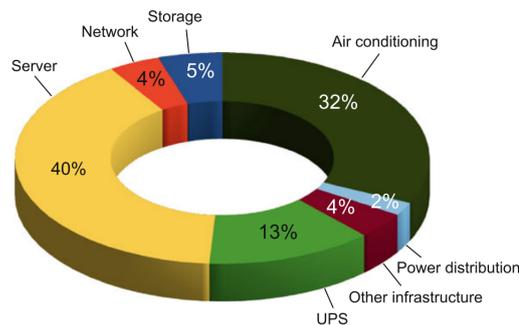


Figure 2.1: Power breakdown of a typical data center [6].

In figure 2.1 a power breakdown is shown of a typical data center. One of the key elements that contribute most to the energy consumption is cooling. The complete process of managing the temperature and humidity takes 32% of the total energy consumption [6]. Cooling a data center can be done in many ways. The traditional way is to manage a data center's temperature is to condition the air of the room, which is called CRAC (Computer Room Air Conditioner) [4]. The air conditioners are positioned on the edge of the room. Here it will lead the hot air in the upper part of the room and recirculates the conditioned cold air back to the server racks through pipes underneath the floor. When the cold air has arrived at the racks, the air will be drawn from the front of the rack to the back. The hot air, which leaves the rear end of the racks, climbs towards the ceiling of the room where it will be redirected to the air conditioner. Although this type of cooling is not that expensive, it is quite inefficient due to the mixture of hot and cold air which creates an unevenly divided temperature across the data center. But it is still one of the most used ways of cooling a data center though.

Another important element is the used energy source. The underlying goal of the greening the cloud project is to reduce the usage of fossil energy sources and to reduce the CO₂ and

other greenhouse gas emissions. All this to reduce the climate change and make a better future. According to the IEA (International Energy Agency), the amount of CO₂/kWh is a measure to compare those sources [7]. Table 2.1 gives an average overview of Implied carbon emission factors from electricity generation (grams CO₂ per kWh).

Energy source	Gram CO₂ per kWh
Anthracite	965
Coking coal	785
Other bituminous coal	860
Sub-bituminous coal	925
Lignite	1005
Coke oven coke	800
Gas works gas	420
Coke oven gas	415
Blast furnace gas	2200
Other recovered gases	2030
Natural gas	400
Crude oil	635
Natural gas liquids	540
Refinery gas	410
Liquefied petroleum gases	530
Kersene	645
Gas/diesel oil	715
Fuel oil	670
Petroleum coke	970

Table 2.1: Carbon emission per kWh produced divided by energy source [7].

As shown in Table 2.1, the CO₂ emission differs a lot depending on the energy source. Blast furnaces gas has an emission of 2200 gram CO₂/kWh, which is 5.37 times more than the emission of refinery gas (410 gram CO₂/kWh). Thus, by using a more eco-friendly energy source, the CO₂ emission can be reduced roughly up to 5 times. This has a big influence on the eco-friendliness of a data center.

The temperature and humidity also play a big part in the energy consumption of a data center. The humidity is a percentage of the amount of water vapor in the air. This percentage has influence on the amount of hardware failures and the cooling costs [8].

As mentioned earlier, a data center consists of many computers, which are connected by routers and switches. Because of the increasing complexity of data centers, the efficiency of the communication inside and between data centers also becomes more important. Although only 10-20% [9] of the total data center energy consumption is used by networking, there is much energy to save. Network equipment in idle state uses 90% of the amount of energy when it is highly used [9]. A network efficiency improvement could therefore have a significant effect on the total energy consumption of the network infrastructure.

The servers in the data center consist out of similar components as a personal computer. There is a Central Processing Unit (CPU), Random Access Memory (RAM) and Hard Drive Disks (HDD), which are all mounted on a motherboard.

In figure 2.2 the energy usage is shown of each element of the server. The CPU uses by far the most energy. This of course depends on the usage of the CPU and it's energy efficiency. Very low powered CPUs like an Intel Atom use less energy than an Intel Xeon, but the performance of the Intel Atom is a lot lower than the Intel Xeon [11].

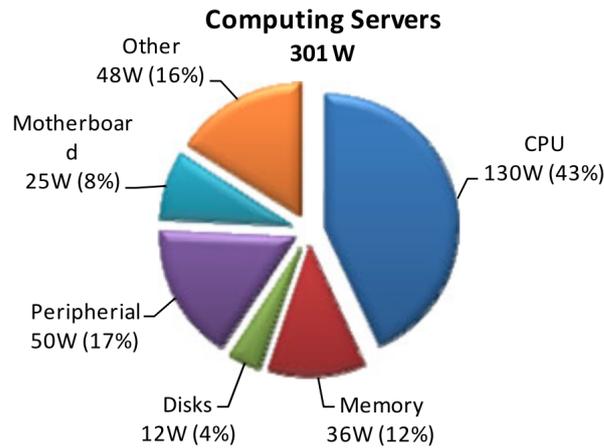


Figure 2.2: Energy usage of a server per hour represented per hardware component [10].

This information gives a brief introduction on some important aspects regarding energy efficiency of cloud computing. Since energy usage becomes a more important factor, what has the European Union to say about data centers and their energy efficiency?

2.2 EU best Practices

In 2014, the European Union composed a document [12] with best practices regarding energy efficiency. With this document, the EU gives an overview of what they think data centers must meet in terms of energy efficiency. This document is not just a list with best practices, it takes several aspects into account. The EU distinguishes applicability, several types of applicants, advice, value of practices.

First, they divided the best practices in a few categories. Practices marked green concern all existing IT, Mechanical and Electrical equipment within the data center. A practice is marked orange if it is expected during any new software install or upgrade. Yellow marked practices are practices expected for new or replacement IT equipment. Practices are marked blue if it is expected for any data center built or undergoing a significant refit of the M&E equipment from 2010 onwards. And finally, practices are marked blank (white) if it concerns practices, which are optional for participants.

Besides the color marking, the practices are ranked with a score from 1 to 5 (1 is the minimum and 5 the maximum value). This score indicates the level of benefit you would have from applying that practice.

Furthermore, the applicability is taken into account. Not all cloud services are able to honor all practices in their data center due to several constraints. For example, if you are not the owner of the entire data center. Therefore different types of applicants are defined. First, there is the operator. This applicant operates the physical part of the entire data center as well as the IT services delivered. Next, there are the colocation provider and the colocation customer. The colocation provider provides, as the name suggests, space, power and cooling capacity. The colocation customer on the other hand owns and manages IT equipment located in a data center where they purchased space, cooling capacity and power from the colocation provider. Last, there are 2 types of managed service providers: the managed service provider and the managed service provider in colocation space. The managed service provider (MSP) owns and manages just as the operator the entire data center including IT equipment, but then for the purpose of delivering IT services to customers (a third party). This is what we call outsourcing. The MSP in colocation space is equal to the regular MSP but purchases space, cooling, power and IT equipment to provide their services.

It is obvious that a different type of applicant results in a different type of responsibility. Therefore it is hard to compare cloud services mutually or to measure the degree of which a cloud service provider meets the suggested best practices. In figure 2.3 is shown who should be responsible for each sector according to the EU.

	Operator	Colo provider	Colo customer	MSP in Colo	MSP
Physical building	Implement	Implement	Endorse	Endorse	Implement
Mechanical & electrical plant	Implement	Implement	Endorse	Endorse	Implement
Data floor and air flow	Implement	Implement & Endorse	Implement & Endorse	Implement	Implement
Racks and rack air flow	Implement	Implement & Endorse	Implement & Endorse	Implement	Implement
IT equipment	Implement	Endorse	Implement	Implement	Implement
Operating System & Virtualisation	Implement	Endorse	Implement	Implement	Implement
Software	Implement	Endorse	Implement	Implement & Endorse	Implement & Endorse
Business practices	Implement	Endorse	Implement	Endorse	Endorse

Figure 2.3: Deviation of areas of responsibility of a data center according to the EU [12].

The best practices are divided into several chapters, which is shown in Table 2.2. That illuminates that the best practices are mainly focused on hardware and the organization/management of a data center.

The goal of the European Union is clear. The document should provide an overview of what elements in a data center have the biggest influence on the energy consumption. Furthermore, the document provides insight in how and when you can accomplish better energy efficiency. The idea sounds perfect, but is it?

2.3 Criticism on EU best practices

First, if you want to be able to determine if a cloud service or data center is energy efficient or not, you have to compare cloud services with each other. This is not mentioned at all in the document. What if a certain cloud service has by definition a high energy usage due to its functionality, but uses energy efficient cooling, CPU, etc. Is this cloud service a green cloud service or not? Take a really powerful and fast car for example, that has all the energy efficient solutions implemented. Although it produces a lot of power for a relative small amount of petrol, it still consumes more energy in general than a slower (energy efficient) car. Which car is greener then? The one that gets the most speed or power produced by one liter of petrol? Or the one that consumes less petrol in general bringing you from A to B? A way to compare or rank cloud

Subject	Number of best practices	Average value (1-5)
3 - Data center utilisation, Management and planning		
3.1 involvement organisational groups	1	5.00
3.2 general policies	4	3.25
3.3 resilience level and provisioning	5	3.20
4 - IT equipment and services		
4.1 selection and deployment of new IT equipment	15	4.00
4.2 deployment of new IT strategies	7	4.57
4.3 management of existing IT equipment and services	8	4.38
4.4 data management	6	3.67
5 - Cooling		
5.1 air flow management and design	13	3.23
5.2 cooling management	7	2.57
5.3 temperature and humidity settings	6	3.83
5.4 cooling plant	11	3.91
5.5 computer room air conditioners	6	3.50
5.6 reuse of data center waste heat	3	2.00
6 - Data center power equipment		
6.1 selection and deployment of new power equipment	5	2.60
6.2 management of existing power equipment	2	2.50
7 - Other data center equipment		
7.1 general practices	3	1.67
8 - Data center building		
8.1 building physical layout	5	2.40
8.2 building geographic location	5	2.40
8.3 water sources	3	1.67
9 - Monitoring		
9.1 energy use and environmental measurement	9	3.22
9.2 energy use and environmental collection and logging	4	3.75
9.3 energy use and environmental reporting	4	4.00
9.4 IT reporting	3	3.00
10 - Practices to become minimum expected		
4.1.3 expected operating temperature and humidity range	1	-
11 - Items under consideration		
11.1 air flow / delta T	1	3.00
11.2 utilisation targets	1	3.00
11.3 further development of software efficiency definitions	1	3.00

Table 2.2: Overview of subject distribution of the EU Best Practices document limited at subsection level. Per subsection, the number of suggested best practices is shown and their average value given by the EU. This value indicates the benefit you would have from implementing that best practice

services is not given by the EU.

Second, some best practices are not clearly explained. Hereby a few examples to clarify the point just made. Chapter "Other Data center Equipment", says: "Low energy lighting systems should be used in the data center". What is a low energy lighting system? What percentage of the total energy consumption should be given to lighting systems? In chapter "Deployment of New IT Services" is mentioned that you should "select and develop efficient software", "Make the energy use performance of the software a primary selection factor." or "Make the energy use performance of the software a major success factor of the project". These best practices have a maximum importance rating (5). Therefore, this should be an important aspect. Do they tell you when software is energy efficient? They just mention it so that it is taken into account. Any guidelines would be pleasant.

When they do give you hard numbers, the numbers are not specific either. In the section Equipment and Services for example, the restricted inlet temperature must be between 15°C and 32°C and the humidity between 20% and 80%. With a humidity allowance from 20% till 80% almost anything is accepted.

Another thing just mentioned in section 2.1, the importance of the energy source is not covered at all. Off course, the energy source is not directly related to the energy consumption of the data center itself. However, the goal of setting up these guidelines is to accomplish that data centers use less energy and become more eco-friendly. The energy source has a big influence on the CO₂ emission and is consequently an important factor.

Despite these (little) shortcomings, the document is a good starting point from where to focus on. But, is it possible to be more specific? At a certain point, more data centers will meet these standards. With this model you are not able to compare those data centers for example. So there is some room for improvement.

Literature overview

3.1 Research question

One of the main goals of this thesis is to give insight which aspects are covered in the research literature with respect to the greenness or eco-friendliness of cloud services. Therefore, the first research question is:

RQ 1: *Which aspects and to what extent are those aspects covered in research literature regarding greenness or eco-friendliness of cloud services?*

To know which aspects are most covered in the relevant scientific literature without reading each one individually, a classification tree is used.

3.2 Finding literature

Google Scholar is used to find literature. This data source allows you to compose queries with several properties and dependencies. For example you can use parenthesis, AND / OR operators, -intitle, "-" (Results must NOT include any words) etc. These advanced search methods are used in the queries.

Mendeley¹ is used to manage the Google Scholar results. This is an online reference (and PDF) manager. Thanks to Mendeley's Web Importer it is possible to save all Google Scholar results on the current page. Unfortunately, Google Scholar does not allow you to export all results at once. Not even with a self made script or third party software². To prevent a huge amount of mouseclicks to collect all results, an automatic mouse clicker or macro is used. This mouse clicker is called JitBit Mouse Recorder³. With these two applications it was possible to export all results into Mendeley automatically and to manage and organize the results.

There are some words considered as keywords in this literature review. Keywords such as: "energy", "power", "efficient", "cloud" and "data center" are words every a query must contain. This resulted in the following query:

(intitle:energy OR intitle:power) cloud efficiency ("data center" OR "data center").

This query resulted in 1260⁴ hits. This amount was too high to handle properly at once. Besides that, Google scholar can only show up to 1000 results even though it tells you there are 1260⁵. As

¹<https://www.mendeley.com/>

²<https://scholar.google.com/intl/en/scholar/help.html#export>

³<https://www.jitbit.com/mac-mouse-recorder/>

⁴Last visited june 2015

⁵<https://scholar.google.com/intl/en/scholar/help.html#export>

one of the goals of this thesis is to provide an overview of the accents in scientific literature with respect to green clouds, several aspects should be taken into account. If a query is too specific certain aspects could be underexposed. What can be emphasized is the part of decreasing the amount of energy used. Therefore, the keyword "reduction" is added. Since we are searching for energy efficient subjects, the keyword "metric" is also added. The energy efficiency depends on the metrics used. This resulted in this slightly more specific query:

(intitle:energy OR intitle:power) cloud efficiency reduction metric ("data center" OR "data center")

This query returned 858⁶ results. This amount of results falls within the range of results Google Scholar can offer, but is still too big to handle manually. To reduce the amount of hits with the same query, it is possible to define a time range in which an article was published. When excluding patents and all results published before 2014, the number of results reduced to 475⁷. This amount is more feasible to use.

3.3 Classification

With the method described above, it is possible to save the Google Scholar results in a BibTex file. This, unfortunately, is without content. What you can find in this BibTex file is the author, title and year of publication for example. If all titles are put into a word frequency tool, such as *WriteWords*⁸, it gives insight in which words are most used.

Since Mendeley is not capable of classifying by content, Google Scholar is also used to sort the 475 (primary) results. To be able to classify articles, several classes are necessary. Two classes or categories were made easily, being *Hardware* and *Software*. Soon, it became clear that some extra (sub)categories would be desirable because these were too general. With the potential keywords from the frequency counter in mind, 4 categories were made: Server level hardware, Data center infrastructure, Software implementation and Management. Server level hardware concerns articles about CPU, storage and other server components. But also about the allocation of those resources. Data center infrastructure is also about hardware, but of the data center as a whole. Here the communication (networking) and climate control are important aspects. The third category concerns software implementation. In other words, articles about how the software works and why you should use certain software instead of other software. Articles about virtualization algorithms and scheduling belong to this category. The last is about what you can do with software. How do you manage a data center? Subjects such as monitoring or modelling are covered here.

Per category, four keywords are chosen. For the first category, the following four keywords are used: server, component, resource and allocation. Those four keywords are added to the primary query in the following way: (server AND component AND resource AND allocation). The results from this query are a subset of the 475 primary results which contain the just mentioned four keywords. The overall results are pictured in figure 3.1 below.

The error rate of this query is $(39/475 =) 8.2\%$. The error rate is based on articles which are not related to the subject. "*A vertically discretised canopy description for ORCHIDEE (SVN r2290) and the modifications to the energy, water and carbon fluxes*" or "*A Hard X-Ray Power-Law Spectral Cutoff in Centaurus X-4*" for example, are results which Google Scholar provided. Such results are unavoidable using Google Scholar and are considered to be a real error.

⁶Last visited june 2015

⁷Last visited june 2015

⁸<http://www.writewords.org.uk/>

3.3.1 Overview

Figure 3.1 is a visual reproduction of the results of the four queries. This figure shows the distribution of the scientific literature of this subject.

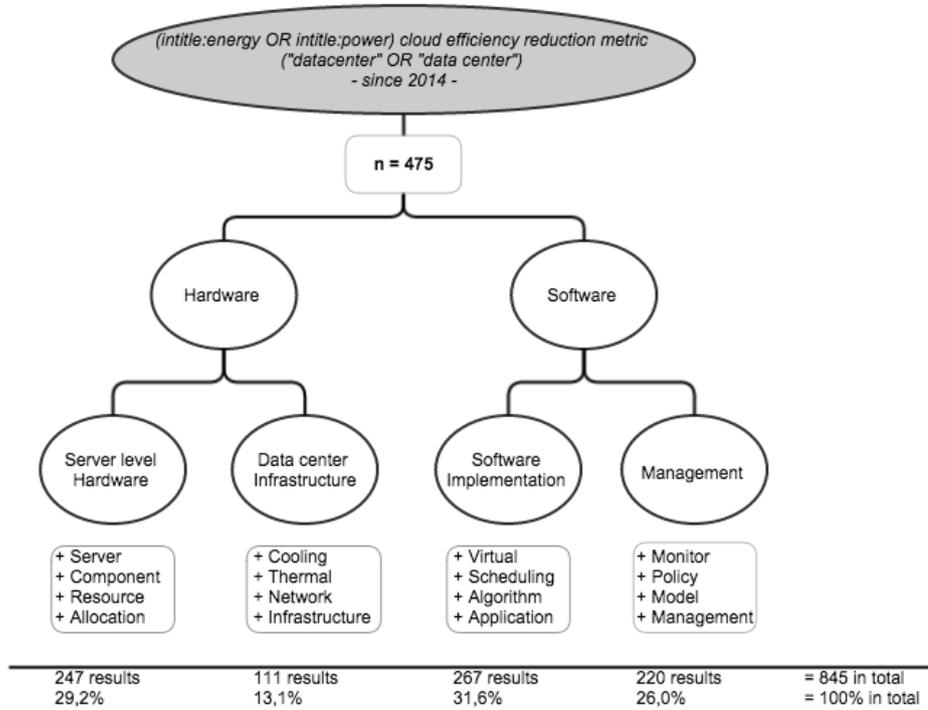


Figure 3.1: Classification tree since 2014

As shown in figure 3.1, the amount of articles of the four subsets combined is greater than the original amount of primary results. This is because an article can contain keywords from two or more categories. With this method an article could appear hypothetically in all four categories when containing all 16 keywords. Then this one article is counted four times.

Due to feasibility reasons it was not possible to review large amounts of results and had to tighten the query by using a time range. Now, we successfully have defined these queries and composed a classification tree, we can adjust the time range of which results are published. Can we discover any trends in scientific literature regarding green cloud computing? The results of the same queries, but within a time range "up to 2014" instead of "from 2014", are shown in figure 3.2.

Figure 3.2 shows that the distribution of scientific literature of the 4 categories has not changed that much comparing to the publications before 2014. The relative distribution is actually identical as shown in figure 3.1.

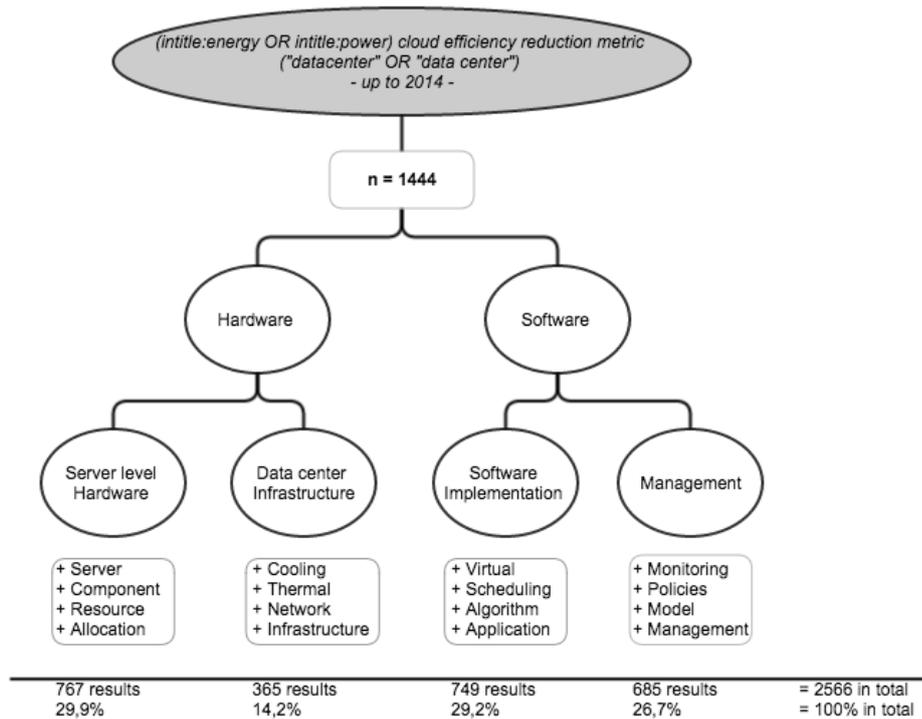


Figure 3.2: Classification tree up to 2014

3.3.2 A more specific approach

Although these aforementioned trees give a rough overview of the accents in scientific literature, there is some room for improvement. Next, another method is used to classify the results of the Google query. Because Google does not offer a variety of options to handle the amount of results, a solution is found elsewhere. Some logic is used to specify the query without losing results. With a simple equation using boolean operators you can prove that:

$$(1) \quad (A \cap B) + (A \cap \neg B) = A.$$

Which is also visualised in a Venn diagram in the following figure:

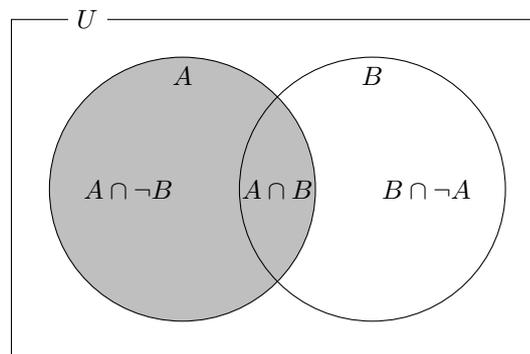


Figure 3.3: Visualisation of equation 1 in the form of a Venn diagram

This Venn diagram is a visual expression of the equation mentioned above. $(A \cap B)$ means all items which belong to subset A and subset B. This is the overlapping part of circle A en B in the middle. $(A \cap \neg B)$ means all items which belong to subset A and do not belong to subset B. This the left part of the left circle in the Venn diagram. Those two parts together are equal to

the whole (grey) circle which defines subset A.

When considering the primary query equal to "A" and a newly introduced keyword to "B", it is possible to split the results into two subsets based on the appearance or lack of keyword "B". As described above, "A" is equal to the primary query and has therefore 475 results. When adding keyword "B", Google will only return articles of this set which contain keyword "B". After that, when adding "not B", Google then only returns the articles which does not contain keyword "B". These two subsets combined is equal to the results of the primary query. In this way you can build a tree by including and excluding certain keywords from the subset. An article can only appear once in this tree. Since every keyword has two versions, the "B" and "not B", using all 16 keywords previously mentioned will result in a tree with $2^{16} = 65536$ leaves. A tree that large is not desirable. Therefore, another approach is used. Near the root of the tree keywords which roughly split the amount of results are preferred. By doing so, the tree will be better balanced. With this in mind, the following tree is built:

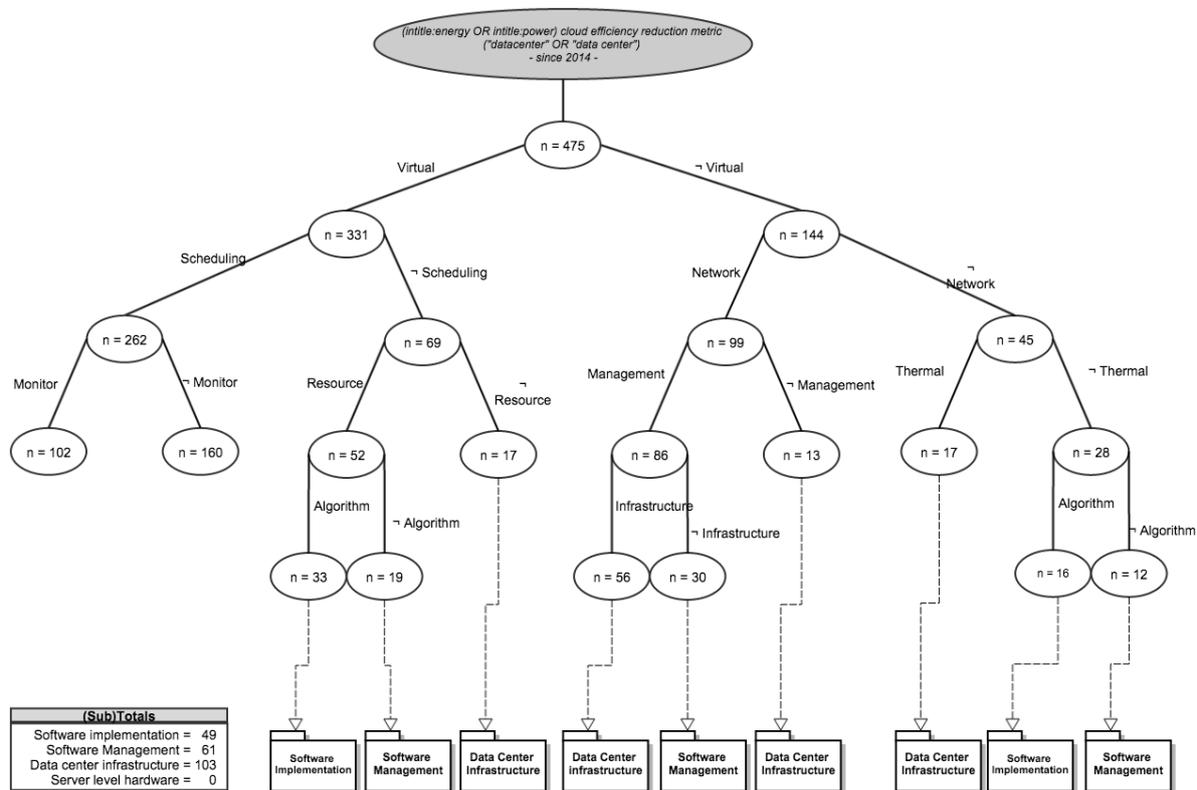


Figure 3.4: More specific classification tree based on the appearance or absence of keywords using the same query as in figure 3.1

This tree requires more interpretation and experience to build. The results are split by using the 16 keywords mentioned earlier. When dividing articles based on the occurrence or lack of a certain keyword, a subset is created. The results of this subset are then divided by using another keyword. The keywords are chosen from the 16 keywords mentioned earlier. The keyword which is used to split the results at each node is chosen to the best of my abilities. Also *WriteWords*, is used to give insight of which keywords are most used in each particular subset. This process continues until a subset is made where the subject of the articles included reflects to one of the four categories for the biggest part. Besides, no article is counted twice. When adding all leaves, the total amount of articles is equal to the number of results in the root.

As shown in figure 3.4, The left-hand side of the tree is unfinished. This is because it was impossible to specify these subsets with all possible keywords what would lead into a subset which reflects on of the four categories. In this subset, all keywords occur or do not occur. This could be because these articles use all the 16 keywords mentioned earlier. Because of this, it is not possible to split this subset based on one keyword. This classifying method has come to an end.

In the bottom left corner a table of (sub)totals is represented. These numbers represent the amount of articles of which it was possible to allocate to a specific category. Although these numbers are unfinished it is hard to draw any conclusions. The only think you can notice is that the category "Server level hardware" has no articles assigned to it. This could be because terms like server and allocation for example are used in multiple contexts.

This problem could be solved by counting the number of occurrences of a certain keyword. When Keyword "A" occurs more times than "B", it is more likely that this article belongs to "A"'s subset. This way of weighing the importance of keywords based on their occurrence is not possible with Google Scholar.

Due to feasibility reasons of this thesis, one category will be chosen in the next chapter to handle in greater depth. Then is becomes possible to answer the next research questions.

Literature review

At the beginning of this thesis, we discuss the Best practices report from the EU. After that, we gave a brief overview to give insight in the distribution of subjects in scientific literature. Due to feasibility reasons, it was not possible to cover all subcategories. Therefore, we chose one category to examine in detail. As shown in figure 2.2, "Management" and "Monitoring" are subjects of the EU best practices document that have some room for improvement which also has a relatively high value of influence. Therefore the category "Management" as shown in figure 3.1 is chosen to handle in greater depth. This should cover "Management" as well as "Monitoring" since "Monitoring" is an important part of "Managing". In the previous chapter To do this in a fair manner, a Systematic Literature Review (SLR) is used. As shown in figure 3.1, this category contains 220 articles. The following paragraphs show how SLR is implemented.

4.1 Research question

The literature study consists of several parts, which is done systematically to make the review more objective. First of all, a research question must be defined. Which answers are you looking for? The research question which will be answered in this chapter is:

RQ 2: *What do we know so far regarding data center management and monitoring?*

RQ 3: *And in what way and how could these findings help improving the EU's Best Practices?*

4.2 Methodology

Next is the search process. This process describes the used strategy to search research literature. After that, the results must be filtered. This is done by setting some inclusion and exclusion criteria. The remaining articles are used to extract the important data. This is called data collection [13].

Just like the literature overview, Google Scholar is used to find literature. This data source allows you to compose queries with several properties and dependencies. For example you can use parenthesis, AND / OR operators, -intitle, "-" (Results must NOT include any words) etc. These advanced search methods are used in the queries.

Mendeley ¹ is used to manage the Google Scholar results. This is an online reference (and PDF) manager. Thanks to Mendeley's Web Importer it is possible to save all Google Scholar results on the current page. Unfortunately, Google Scholar does not allow you to export all

¹<https://www.mendeley.com/>

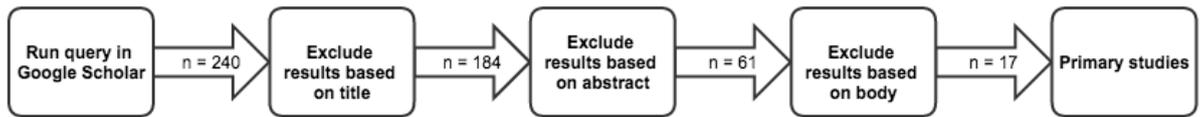


Figure 4.1: Flowchart of the search process to define the primary studies. "n" represents the amount of results of that particular stage of the process.

results at once. Not even with a self made script or third party software. ². Last, ATLAS.ti ³ is used in the last stage of the SLR, which will be discussed later on.

As mentioned earlier, the SLR is based on the category "Management". This resulted therefore in the following query:

(intitle:energy OR intitle:power) cloud efficiency reduction metric ("data center" OR "data center") +(Monitor AND Policy AND Model AND Management)

This query returned 240 ⁴ results. This amount of results falls within the range of results Google Scholar can offer.

4.2.1 Inclusion and Exclusion criteria

When selecting the useful articles, you have to set some inclusion and exclusion criteria. Which criteria must an article satisfy to be useful? First of all, the resulting articles must be in English because of the feasibility reasons of this research. Second, the energy efficiency must be considered. There are articles for example, which explain how certain solutions must be implemented, or provide new (scheduling) algorithms which theoretically could be more efficient. For this research, we are only interested in applicable solutions or methods. Third, a data center must be the infrastructure of the provided cloud service. Since the focus lies on policies, models and monitoring, the management part must be considered. Finally, the full body must be found with the license of the University of Amsterdam. Articles, which are not accessible with this license, will be passed.

4.2.2 Search process

The query returned 240 results instead of the 220 results as shown in figure 3.1. Although the query is identical, the number of results changed somewhere. In the next section the inclusion and exclusion criteria are shown. Based on these criteria articles will remain or will be removed of this subset. This is initially done by title. The remaining articles in the subset are read again, but now not just the title, but also the abstract. After that, the full body is read to check whether the article is useful or not. In figure 4.1 is shown how many articles reached the new subset.

The residual 17 articles are called primary studies. These articles will be discussed thoroughly in the following chapter. The bibliography of these 17 primary studies is listed in Appendix A.

²<https://scholar.google.com/intl/en/scholar/help.html#export>

³<http://atlasti.com/>

⁴Last visited june 2015

Research findings

This chapter will present the results of the systematic literature review as described in the previous chapter. First, an overview is given of the primary studies. Next, a more detailed view is made.

5.1 Overview of the findings

All primary studies are about data center management, but there are some differences and obviously similarities between those articles. The first partition made has to do with the accentuation of the article. There are roughly two groups: *Implementation* and *Goals*. When Managing a data center, you have to take several goals into account. A goal could be remaining a Quality of Service (QoS) or a service Level Agreement (SLA) while using less energy. When a goal is defined, an implementation or strategy is required to accomplish that goal. In order to know if a goal is achieved, metrics must be defined. These metrics give insight in the extend to which a goal is accomplished, or to adjust an implementation when necessary. Furthermore, metrics are used to help defining a goal. This makes a goal more concrete and manageable.

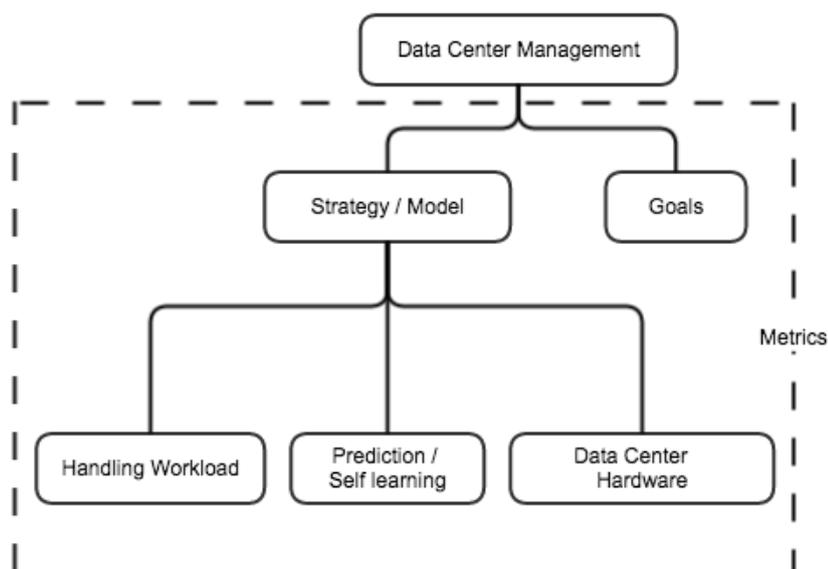


Figure 5.1: Visualization of the subject distribution of the content of the primary studies regarding management.

Implementation includes three subgroups: *Data center hardware*, *Prediction / Self Learning* and *Handling workload*. In other words: How do you efficiently manage data center hardware? How can a data center manage itself? And how should a data center manage different workloads? Data center hardware includes resource allocation, but also Dynamic Voltage and Frequency Scheduler (DVFS) for example. As long as the main focus lies on managing hardware. When the accent lies on handling the workload by using scheduling algorithms for example, the article fits in the second category: handling workload. The third category is about the anticipation or prediction of workloads. Self learning is a key factor in here. At last, there is a category called metrics. As mentioned above, metrics are necessary to be able to manage a data center, regardless of the subject. In Table 5.1 is shown which articles cover which subjects.

As shown in Table 5.1, some articles cover multiple subjects and occur therefore in multiple categories.

5.1.1 Handling workload

When selecting the articles regarding data center management, the subject *workload* or *load balancing* occurs frequently. Here, everything will be discussed regarding workload handling. We start with the subjects related to handling different types of starting positions. What to do when a data center has to manage a certain workload? First, Dynamic Voltage (Frequency) Scaling (DVFS or DVS) and Adaptive Link Rate (ALR) were found [14, 15]. DVFS controls the usage of the CPU based on the system load. When there is not much to do, the energy state of the CPU is then being adjusted. ALR is somewhat similar to DVFS, but it is applied to a network instead of a CPU. ALR is also trying to reduce the capacity and therefore the consumption when there is not much to do. This workload depends on the time, season and geographical location of where the cloud service is used and located [16]. Another thing mentioned is server state scheduling [16]. This method looks a lot like DVFS, but then it schedules the state of a server. When a server is not used, it will be put into sleep or it hibernates that server.

Second, are subjects regarding handling the workload itself. The first subject mentioned is Geographical Load Balancing (GLB). This method divides the load over multiple data centers. An other option mentioned is classifying incoming workload based on the importance of that load of jobs or the kind of jobs involved [17]. An other option is to schedule all jobs in such a way that a maximum amount of green energy is used [18].

At last, there are subjects regarding the managing part of a data center. Here a trade-off is made between power and performance [15]. The goal is to use as little energy as possible, but without violating SLA's or QoS or other Business Level Objectives (BLO) [19, 20]. Since virtual machines are often used in data centers, the VM Manager has an important role. The VMM manages VM migration if necessary and VM Allocation [17].

5.1.2 Data center Hardware

Articles about data center hardware can be classified into 4 subclasses, being *Cooling related*, *PDU related*, *Energy measurement* and *redundancy*. First, in cooling related articles the term Computer Room Air Conditioner (CRAC) and Computer Room Air Handler (CRAH) occur often. CRAC works practically the same as an air conditioner at home, but then big enough to handle ha data center. CRAH works a little different. Now, chilled water is used to cool the air. Since cooling uses a lot of energy, it is more energy efficient to balance the cooling according to the workload.

Subject	Description	Occurrence	Percentage	Primary studies
Handling workload	Covers all workload balancing, scheduling, VM placement, Resource Allocation subjects.	11	26.2%	[Kalaitzoglou et al., 2014], [Kong and Liu, 2014], [Lent, 2015], [Subirats and Guitart, 2015], [Elijorde and Lee, 2015], [Norouzi and Bauer, 2015], [Zhou and Jiang, 2014], [Procaccianti, 2015], [Chen et al., 2014], [Noureddine, 2014]
Data Center Hardware	Covers all server level hardware, data center infrastructure, communication, cooling etc.	10	23.8%	[Kong and Liu, 2014], [Lent, 2015], [Zhou and Jiang, 2014], [Uchekukwu et al., 2014], [Procaccianti, 2015], [Vitali et al., 2015], [Arianyan et al., 2015], [Chen et al., 2014], [Capozzoli et al., 2014], [Noureddine, 2014]
Prediction / Self Learning	Covers also workload or cooling for example, but with the accent on self learning, adapting, anticipation, prediction.	2	4.8%	[Subirats and Guitart, 2015], [Procaccianti, 2015]
Metrics	"A quantitative measure of the degree to which a system, component, or process possesses a given attribute" [Vitali et al., 2015]	16	38.1%	[Kong and Liu, 2014], [Capozzoli et al., 2014], [Javaid, 2014], [Subirats and Guitart, 2015], [Arianyan et al., 2015], [Fiandrino et al., 2015], [Vitali et al., 2015], [Uddin et al., 2014], [Uchekukwu et al., 2014], [Norouzi and Bauer, 2015], [Zhou and Jiang, 2014], [Lent, 2015], [Kalaitzoglou et al., 2014], [Noureddine, 2014], [Procaccianti, 2015], [Elijorde and Lee, 2015],
Goal	Covers all subjects about the purpose of certain implementations. Think of Quality of Service (QoS), Service level Agreement (SLA) etc.	3	7.1%	[Javaid, 2014], [Norouzi and Bauer, 2015], [Arianyan et al., 2015]

Table 5.1: Percentage of appearances of primary studies divided per category .

Next, there are articles about Power Distribution Units (PDU) [21]. These PDUs manage the voltage levels through the entire data center. They also Transform the voltage levels suitable for servers. Unfortunately, there is a constant (or proportional to the workload) energy loss in this process.

Third, there are articles about Energy measuring. This can be done at different levels. These articles often mention a energy metering framework like JouleUnit [20], PTEC [22], PowerScope [23], JouleMeter [23] etc. These frameworks has to receive information from the hardware to be able to calculate the energy use of a process. To do this in the most precise manner, power metering devices must be installed. The software components which read the power usage from the metering devices are called energy collectors. When these power metering devices are not available, recent studies showed [14] that the server's power consumption can be estimated by using a linear relation with the power usage of the CPU. This, off course, is because the CPU consumes relatively much power in a data center.

Finally, are articles about (hardware)redundancy. When servers are not used and therefore redundant, the server must be turned off or hibernated to prevent unnecessary energy consumption [16, 24]. This can be done by migrating VMs to the least servers possible. The redundant servers can then be turned off of hibernated.

5.1.3 Prediction and Self Learning

This subsection is not the biggest one, but is important since there is less or no human interaction needed when managing a data center. Therefore this is covered in a separate subsection. Prediction can be done at several levels. First, there is cooling [22]. Here a model is made from the previous energy usage and can then choose the most energy efficient to use for the future. This is an ongoing process of predicting the load based on the workload until now and is all done in order to minimize the the energy consumption of the CRAC or CRAH units. This can also be used to forecast VM migration or CPU usage [19, 20]. Based on the previous and current workload, it is possible to estimate the future workload. This is not only beneficial for cooling a data center, but also for the energy usage of VMs, CPUs or other data center components.

Second, prediction is used for the performance or response time of applications [25, 15]. The response time is useful to make sure the cloud service hits it's SLA or BLO for example. This prediction is then used to minimize the energy consumption without violating policies, agreements or goals.

5.1.4 Metrics

In almost every article (16/17) metrics are used. Unfortunately, these metrics are almost all only relevant to a specific situation or implementation. There are only a few metrics widely used. The first one mentioned is Power Usage Effectiveness (PUE) [26, 18, 19, 21, 27]. This metric gives insight into the amount of energy used by IT equipment in comparison with the total energy used by the data center. The closer the value of the ($PUE = TotalFacilityPower / ITEquipmentPower$) gets to 1.0, the better it is. Another term used to measure this is Data Center Infrastructure Efficiency (DCIE). This metric is the same as PUE, but formatted in a percentage of energy used by IT equipment against the total energy usage of the data center. Modern data centers have a PUE around 1.1 - 1.2 [18]. Another frequently used metric is the Rack Cooling Index (RCI). As the name suggests, it provides insight in the rack cooling efficiency according to the thermal guidelines. There are too many metrics to mention and explain here. For example, there are already 17 metrics mentioned only about communication in one article [28]. The important thing about metrics is that it provides an index. With this index you are able to compare it with other cloud services or with your own cloud service through time.

5.1.5 Goals

When managing a data center, certain goals has always have be thought of. A goal could be using less energy, or lowering the PUE of a data center. Besides the technical related goals, there are also goals which are related to the users of the cloud service. The aforementioned terms like BLO, SLO or SLA are important. Goals are important to manage a data center or any other company in the commercial industry. Depending on the goals that have been set, several metrics are needed. These metrics can be metrics which are widely accepted, but you can also define your own metrics. As long as you can keep track of the relevant processes to accomplish your goal.

5.2 Comparison with EU best practices

In what way can these results help to improve or tighten the EU Best Practices? First, the Management and monitoring subjects provided by the EU will be handled in more detail. After that, the comparison can be made and the EU Best Practices could be improved.

5.2.1 Relevant Information from the EU best practices

In Table 5.2, the relevant best practices are shown up to subsection detail. The descriptions are left out due to the large size of the table when displaying them all.

5.2.2 Improving EU best practices

When comparing the relevant EU best practices with the findings mentioned in section 5.1 there are several things to say.

Section of the EU best practices	Best Practice	Expected	value
3.1.1	Group involvement	Entire Data center	5
3.2.1	Consider the embodied energy in devices	Entire Data center	3
3.3.3	Lean provisioning of power and cooling for a maximum of 18 months of data floor capacity	New build or retrofit	3
3.3.4	Design to maximise the part load efficiency once provisioned	New build or retrofit	3
5.2.1	Scalable or modular installation and use of cooling equipment	Optional	3
5.2.2	Shut down unnecessary cooling equipment	Optional	3
5.2.3	Review of cooling before IT equipment changes	Entire Data center	2
5.2.4	Review of cooling	Entire Data center	2
5.2.5	Review CRAC / AHU Settings	Optional	3
5.2.6	Dynamic control of building cooling	Optional	3
5.2.7	Effective regular maintenance of cooling plant	Entire Data center	2
6.2.1	Reduce enginegenerator heater temperature setpoint	Optional	2
6.2.2	Power Factor Correction	Optional	3
9.1.1	Incoming energy consumption meter	Entire Data center	4
9.1.2	IT Energy consumption meter	Entire Data center	4
9.1.3	Room level metering of supply air temperature and humidity	Optional	2
9.1.4	CRAC / AHU unit level metering of supply or return air temperature	Optional	3
9.1.5	PDU level metering of IT Energy consumption	Optional	3
9.1.6	PDU level metering of Mechanical and Electrical energy consumption	Optional	3
9.1.7	Row or Rack level metering of temperature	Optional	2
9.1.8	IT Device level metering of temperature	Optional	4
9.1.9	IT Device level metering of energy consumption	Optional	4
9.2.1	Periodic manual readings	Entire Data center	3
9.2.2	Automated daily readings	Optional	4
9.2.3	Automated hourly readings	Optional	4
9.2.4	Achieved economised cooling hours	New build or retrofit	4
9.3.1	Written Report	Entire Data center	4
9.3.2	Energy and environmental reporting console	Optional	4
9.3.3	Integrated IT / M&E energy and environmental reporting console	Optional	4
9.3.4	Achieved economised cooling hours	New build or retrofit	4
9.4.1	Server Utilisation	Optional	3
9.4.2	Network Utilisation	Optional	3
9.4.3	Storage Utilisation	Optional	3
11.1	Air flow / Delta T	New IT Equipment	3
11.3	Further development of software efficiency definitions	Optional	3

Table 5.2: Relevant Best Practices organized by the corresponding section in the EU Best Practices document.

Handling workload

In the EU best practices, section 5.2.1, the EU mentions handling cooling equipment. This is in line with the findings of the SLR of this thesis. In the EU best practices, section 5.2.6, real

time cooling scheduling is even suggested. Another best practice is also about workload, being section 9.2.4 of the EU best practices document. Here they suggest to schedule the cooling ours in the most economical way on a years basis. What the EU did not mention is Dynamic Voltage Frequency Scaling. This subject showed up in several primary studies [18, 15]. With the DVFS it is possible to manage the CPU state depending on the workload. If the EU mentions managing cooling equipment based on the workload, why can't they mention DVFS which is also based on workload. Furthermore, VMs are not mentioned at all. Subjects like VM Migration, VM Manager and VM Allocation remains therefore uncovered [17, 20]. Data centers use virtualization to handle their big amount of computing power. An energy efficient VM manager which can migrate and allocate VMs is thus a good addition. Finally, geographical load balancing is not covered. This is not useful for small cloud services, but load balancing between data centers has a lot of potential [18]. This could be added in chapter 11 of the EU best practices ("11 Items under Consideration").

Data center hardware

On the hardware subject there is not that much to add. In section 9.1.1 and 9.1.2 of the EU best practices, the installation of metering equipment is mentioned. This corresponds to the findings of the SLR. Without these meters it is hard to measure the energy consumption of different components. Section 5.2.2 of the EU best practices, "Shut down unnecessary cooling equipment", is also confirmed by the findings of the SLR [16, 24].

Metrics

Metrics are briefly mentioned by the EU. The best-known metric PUE is mentioned, but no hard numbers are given. According to [Kong and Liu, 2014] modern data centers by the year 2014 have a PUE between 1.1 and 1.2. In the United States though, the average PUE was 1.7 in 2014. This number is based on self-reported PUEs [29]. Maybe it is possible to provide a scale of PUE values from 1.0 to 2.0, where 1.7-1.8 is acceptable and 1.1-1.2 is really good. Apart from the scale, at least the PUE values should be mentioned in the same section in the best practices document.

There are lots and lots of metrics and it is not possible to mention them all, but some simple, applicable metrics can be useful. Metrics like Rack Cooling Index (RCI) or Data Center Cooling System Efficiency (DCCSE) [28] give some guidance and insight in the energy efficiency of cooling in a data center. RCI for example, measures the over and under temperatures. Over and under temperatures are temperatures which are above maximum or below minimum recommended. A RCI of 100% means that there are no over and under temperatures. A $RCI < 90\%$ is considered as a poor performance. DCCSE is the ratio between the average cooling system power consumption and load of the data center [28]. Since cooling is such a big subject within the EU best practices, some of these metrics should be added. At last, reports are mentioned often by the EU. PUE is a metric which they require to mention in such a report. This could of course be tightened by using Annual Component Consumption (ACC), Relative Idle Consumption (RIC) or Component Consumption per Unit of Work (CCUW). ACC is the annual energy consumption per application component measured in kWh [14]. RIC is the percentage of annual idle energy consumption to total per application component [14]. CCUW is the average energy consumption (in kWh) of each software component per unit of work delivered [14]. These metrics could make such a report more useful.

Prediction and Self Learning

Prediction and Self Learning do not appear in the document. Nevertheless, there are some findings of the SLR which could be used. Prediction of CRAC or CRAH [22] has a lot of potential. Since cooling is a big topic in the best practices document, this should be mentioned whether in chapter 11 ("Items under Consideration") or in the sections regarding cooling. In section 4.2.4 of the EU best practices "New deployment of IT strategies" of the best practices document is mentioned that you should "select efficient software." According to [Procaccianti, 2015] when developing software or outsourcing this development, self-adaptation is more suitable

when building software from scratch. Section 4.2.4 of the EU best practices should therefore mention this to show that this is one of the possibilities. Although it could be too early to incorporate into this best practices document, but the SLR showed that there are opportunities using this.

Conclusion & Discussion

To answer the first research question, we have to provide an overview of the accents in scientific literature with respect to green clouds. Google scholar is used to find scientific literature. Because we have to provide an overview, the search query can not be too specific. After trying several queries, they returned too many results and therefore some adjustments had to be made. The adjustment with the biggest influence was setting a time range. When setting this time range from 2014 until now, the amount of results became feasible. With this subset a classification tree is composed. In order to make a classification tree, some classes or categories are created, being *Server level hardware*, *Data Center Infrastructure*, *Software Implementation* and *(Data center) Management*. Besides that, there are criteria upon which the classification is based. The criteria consist out of four keywords that all most occur in order to classify a result into that category. Each category has 4 keywords which are representative for that category.

After successfully building the first tree, we could use the same query and classification to compose another tree. This time another time range is used. We chose to adjust this from "until 2014" to "up to 2014". When the second tree was finished, the first thing noticed was that the subject distribution did not change that much. Moreover, the relative distribution was identical "until 2014" compared to "up to 2014". Consequently, the aspects in scientific literature did not change that much.

With this method, an article can occur in several categories. Therefore we tried to compose a more specific classification tree, where each article can only appear in one category. This is done by using some extra logic in the query. Based on the appearance or absence of certain keywords, new subsets are made. Once a subset consists out of articles with the same subject, this subset is assigned to a specific category. Unfortunately, some articles had identical keywords appearing, but did not belong to the same category. With Google Scholar, it is not possible to check for keywords in a certain context or define a threshold of how much a keyword must occur. We were not able to finish this attempt due to this limitation of Google Scholar.

This could be an interesting point of future research. Exporting the articles, although Google blocks that as well, into an arbitrary content manager could offer a lot more possibilities than in this setup.

After the literature overview, one category is chosen to study in more detail. This category is data center management. By using a SLR, recent information is presented regarding this subject. The findings are divided into some categories as well, being *Data center hardware*, *Prediction / Self Learning*, *metrics*, *Goals* and *Handling workload*. The findings are discussed for every category separately.

Finally, with the findings we tried to improve, or tighten best practices proposed by the EU. Several suggestions are proposed to adjust or tighten some EU best practices.

When we want to define the energy efficiency of a data center or compare it with other data centers, metrics are necessary. The problem is that there are many metrics, which are not all applicable to every data center. Defining standards of which metrics should and could be used

by every data center would be a big step forward. Then we can compare the results of those metrics of different data centers. There are some metrics widely used in the world. Power Usage Effectiveness (PUE) is one of them. According to [18] the PUE of a modern data center should be between 1.1 and 1.2 in the year 2014. In the United States though, the average PUE was measured at 1.7 in 2014. An idea could be to make a scale of PUE values which goes from 1.0 to 2.0, where the range 1.1 to 1.2 is considered as really good, and 1.7 to 1.8 is considered as minimal. To complete this, more research is required to build the best possible scale. Another metric called Rack Cooling Index (RCI) is also an applicable metric where some hard numbers can be provided. A RCI of 100% means that there are no over and under temperatures. A RCI $< 90\%$ is considered as a poor performance. Just like the PUE, a scale could here be provided as well. These kind of metrics could easily be compared between data centers. Maybe a standard can be defined of which metrics every data center should use to report its energy efficiency.

For future research it could be interesting to review literature of other subjects than data center management, as well as expanding the amount of applicable metrics. This way, we can help improving the energy efficiency of data centers and make cloud computing more eco-friendly.

Bibliography

- [1] Dennis Abts and Bob Felderman. A guided tour through data-center networking. *Queue*, page 10, 2012.
- [2] Jonathan Koomey. Growth in data center electricity use 2005 to 2010. *A report by Analytical Press, completed at the request of The New York Times*, 2011.
- [3] Archana Venkatraman. Global census shows datacentre power demand grew 63% in 2012. <http://www.computerweekly.com/news/2240164589/Datacentre-power-demand-grew-63-in-2012-Global-datacentre-census>, 2012, (accessed May 1, 2015).
- [4] Krishna Kant. Data center evolution: A tutorial on state of the art, issues, and challenges. *Computer Networks*, pages 2939–2965, 2009.
- [5] Greening the cloud project. <http://www.greeningthecloud.nl>, 2015.
- [6] Gunnar Schomaker, Stefan Janacek, and Daniel Schlitt. The energy demand of data centers. In *ICT Innovations for Sustainability*, pages 113–124. Springer, 2015.
- [7] Co2 emissions from fuel combustion iea statistics international energy agency highlights. <http://www.iea.org/publications/freepublications/publication/CO2EmissionsFromFuelCombustionHighlights2013.pdf>, 2013, (accessed 30-03-2015).
- [8] Lizhe Wang and Samee U Khan. Review of performance metrics for green data centers: a taxonomy study. *The Journal of Supercomputing*, pages 639–656, 2013.
- [9] Brandon Heller, Srinivasan Seetharaman, Priya Mahadevan, Yiannis Yiakoumis, Puneet Sharma, Sujata Banerjee, and Nick McKeown. Elastictree: Saving energy in data center networks. In *NSDI*, pages 249–264, 2010.
- [10] Dzmitry Kliazovich, Pascal Bouvry, and Samee Ullah Khan. Greencloud: a packet-level simulator of energy-aware cloud computing data centers. *The Journal of Supercomputing*, pages 1263–1283, 2012.
- [11] Íñigo Goiri, Josep Ll Berral, J Oriol Fitó, Ferran Julià, Ramon Nou, Jordi Guitart, Ricard Gavaldà, and Jordi Torres. Energy-efficient and multifaceted resource management for profit-driven virtualized data centers. *Future Generation Computer Systems*, pages 718–731, 2012.
- [12] European Commission. Best practice guidelines v5 1 1r - europa. http://iet.jrc.ec.europa.eu/energyefficiency/sites/energyefficiency/files/files/documents/ICT_CoC/2014_best_practice_guidelines_v5_1_1r.pdf, 2014 (accessed March 21, 2015).
- [13] Barbara Kitchenham, O Pearl Brereton, David Budgen, Mark Turner, John Bailey, and Stephen Linkman. Systematic literature reviews in software engineering—a systematic literature review. *Information and software technology*, pages 7–15, 2009.

- [14] G Kalaitzoglou, M Bruntink, and J Visser. A Practical Model for Evaluating the Energy Efficiency of Software Applications. *ICT for Sustainability 2014* (. . . , 2014.
- [15] X Zhou and C J Jiang. Autonomic Performance and Power Control on Virtualized Servers: Survey, Practices, and Trends. *Journal of Computer Science and Technology*, 2014.
- [16] Ricardo Lent. Analysis of an energy proportional data center. *Ad Hoc Networks*, 25:554–564, 2015.
- [17] F Elijorde and J Lee. Attaining Reliability and Energy Efficiency in Cloud Data Centers Through Workload Profiling and SLA-Aware VM Assignment. *International Journal of Advances in Soft Computing & . . .*, 2015.
- [18] F Kong and X Liu. A Survey on Green-Energy-Aware Power Management for Datacenters. *ACM Computing Surveys (CSUR)*, 2014.
- [19] J Subirats and J Guitart. Assessing and forecasting energy efficiency on Cloud computing platforms. *Future Generation Computer Systems*, 2015.
- [20] G Procaccianti. Energy-Efficient Software. 2015.
- [21] A Uchechukwu, K Li, and Y Shen. Energy Consumption in Cloud Computing Data Centers. *International Journal of Cloud . . .*, 2014.
- [22] J Chen, R Tan, G Xing, and X Wang. PTEC: A system for predictive thermal and energy control in data centers. *Real-Time Systems . . .*, 2014.
- [23] A Nouredine. Towards a better understanding of the energy consumption of software systems. 2014.
- [24] M Vitali, B Pernici, and U M O’Reilly. Learning a goal-oriented model for energy efficient adaptive applications in data centers. *Information Sciences*, 2015.
- [25] F Norouzi and M Bauer. Autonomic Management for Energy Efficient Data Centers. *CLOUD COMPUTING 2015*, 2015.
- [26] M A Javaid. A Strategic Model for Adopting Energy Efficient Cloud Computing Infrastructure for Sustainable Environment. *Available at SSRN 2389084*, 2014.
- [27] M Uddin, R Alsaqour, A Shah, and T Saba. Power Usage Effectiveness Metrics to Measure Efficiency and Performance of Data Centers. *Appl. Math*, 2014.
- [28] C Fiandrino, D Kliazovich, P Bouvry, and A Zomaya. Performance and energy efficiency metrics for communication systems of cloud computing data centers. 2015.
- [29] Uptime Institute. 2014 data center industry survey. <http://journal.uptimeinstitute.com/2014-data-center-industry-survey/>, 2014 (accessed March 29, 2015).

Primary studies SLR

E Arianyan, H Taheri, and S Sharifian. Novel energy and SLA efficient resource management heuristics for consolidation of virtual machines in cloud data centers. *Computers & Electrical Engineering*, 2015. URL <http://www.sciencedirect.com/science/article/pii/S004579061500155X>.

A Capozzoli, M Chinnici, M Perino, and G Serale. Review on performance metrics for energy efficiency in data center: The role of thermal management. *Energy Efficient Data ...*, 2014. URL http://link.springer.com/chapter/10.1007/978-3-319-15786-3_9.

J Chen, R Tan, G Xing, and X Wang. PTEC: A system for predictive thermal and energy control in data centers. *Real-Time Systems ...*, 2014. URL http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=7010489.

F Eljorde and J Lee. Attaining Reliability and Energy Efficiency in Cloud Data Centers Through Workload Profiling and SLA-Aware VM Assignment. *International Journal of Advances in Soft Computing & ...*, 2015. URL <http://home.ijasca.com/data/documents/4.Frank-Eljorde-et-al.pdf>.

C Fiandrino, D Kliazovich, P Bouvry, and A Zomaya. Performance and energy efficiency metrics for communication systems of cloud computing data centers. 2015. URL http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=7090996.

M A Javaid. A Strategic Model for Adopting Energy Efficient Cloud Computing Infrastructure for Sustainable Environment. Available at SSRN 2389084, 2014. URL http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2389084.

G Kalaitzoglou, M Bruntink, and J Visser. A Practical Model for Evaluating the Energy Efficiency of Software Applications. *ICT for Sustainability 2014 (...)*, 2014. URL http://www.researchgate.net/profile/Magiel_Bruntink/publication/265301532_A_Practical_Model_for_Evaluating_the_Energy_Efficiency_of_Software_Applications/links/540850320cf2bba34c268a3a.pdf.

F Kong and X Liu. A Survey on Green-Energy-Aware Power Management for data centers. *ACM Computing Surveys (CSUR)*, 2014. URL <http://dl.acm.org/citation.cfm?id=2642708>.

Ricardo Lent. Analysis of an energy proportional data center. *Ad Hoc Networks*, 25 554–564, 2015. ISSN 1570-8705. doi 10.1016/j.adhoc.2014.11.001. URL <http://dx.doi.org/10.1016/j.adhoc.2014.11.001>.

F Norouzi and M Bauer. Autonomic Management for Energy Efficient Data Centers. *CLOUD*

COMPUTING 2015, 2015. URL http://www.researchgate.net/profile/Carlos_Westphall/publication/275463143_CLOUD_COMPUTING_2015_-_The_Sixth_International_Conference_on_Cloud_Computing_GRIDs_and_Virtualization/links/553cf5300cf2c415bb0d02a3.pdf#page=154.

A Nouredine. Towards a better understanding of the energy consumption of software systems. 2014. URL <http://hal.univ-lille3.fr/tel-00961346/>.

G Procaccianti. Energy-Efficient Software. 2015. URL http://www.researchgate.net/profile/Giuseppe_Procaccianti/publication/275348368_Energy-Efficient_Software/links/553a06070cf226.pdf.

J Subirats and J Guitart. Assessing and forecasting energy efficiency on Cloud computing platforms. *Future Generation Computer Systems*, 2015. URL <http://www.sciencedirect.com/science/article/pii/S0167739X14002428>.

A Uchekukwu, K Li, and Y Shen. Energy Consumption in Cloud Computing Data Centers. *International Journal of Cloud ...*, 2014. URL http://www.researchgate.net/profile/Uchekukwu_Awada/publication/263580831_Energy_Consumption_in_Cloud_Computing_Data_Centers/links/00b7d53b561293aa31000000.pdf.

M Uddin, R Alsaqour, A Shah, and T Saba. Power Usage Effectiveness Metrics to Measure Efficiency and Performance of Data Centers. *Appl. Math*, 2014. URL <http://qpl.naturalspublishing.com/files/published/62y855naf6685f.pdf>.

M Vitali, B Pernici, and U M O'Reilly. Learning a goal-oriented model for energy efficient adaptive applications in data centers. *Information Sciences*, 2015. URL <http://www.sciencedirect.com/science/article/pii/S0020025515000614>.

X Zhou and C J Jiang. Autonomic Performance and Power Control on Virtualized Servers Survey, Practices, and Trends. *Journal of Computer Science and Technology*, 2014. URL <http://link.springer.com/article/10.1007/s11390-014-1455-4>.