**Technische Universität Berlin**
Innovationszentrum Energie

# Electricity Demand Modeling of German Data Centers: Dealing with Uncertainties

Lars Dittmar
Marc Schaefer

**Schriftenreihe** **Nr. 1**
Innovationszentrum Energie

# Autoren

**Marc Schaefer**
Fachgebiet Energieverfahrenstechnik und Umwandlungstechniken regenerativer Energien, TU Berlin
Tel: +49 (30) - 314 25363

**Lars Dittmar**
Fachgebiet Energiesysteme, TU Berlin
Tel: +49 (30) - 314 79123

**ISSN**
**ISBN 978 3 7983 2168 7**

∞ **Gedruckt auf säurefreiem alterungsbeständigem Papier**

**Vertrieb/ Publisher**

Universitätsverlag der TU Berlin
Universitätsbibliothek
Fasanenstr. 88 (im VOLKSWAGEN-Haus), D-10623 Berlin
Tel.: (030)314-76131; Fax.: (030)314-76133
E-Mail: publikationen@ub.tu-berlin.de
http://www.univerlag.tu-berlin.de

# Electricity Demand Modeling of German Data Centers: Dealing with Uncertainties

Lars Dittmar
Marc Schaefer

# Contents

## Abstract

Data centers, i.e., centralized facilities for processing, management, and dissemination of data and information, have become essential to the functioning of contemporary socio-economic systems. Although information technology and data centers provide significant social and economic benefits, concerns about data centers' soaring energy consumption and associated $CO_2$ emissions have emerged in recent years. This concern has been manifested through the public and private initiatives that attempt to promote energy efficiency and "green" technologies in data centers. The rapid emergence of data centers, however, provides great challenges to energy demand and efficiency analysis; even in describing the status quo. The lack of statistical information on the number of data centers, their current floor areas, power densities, and energy use are non-existent, which creates considerable technical uncertainties when estimating data centers' energy use. Policymakers and industrial initiatives that attempt to drive the market for efficiency need to be informed of the role data centers play in current and future energy systems. Thus, robust estimates on current and future trends in data center energy usage are required. This is exactly the aim of this paper. We analyze the nationwide impact of data centers on electricity demand in Germany for the period 1998-2008, including uncertainty analysis.


## 1 Introduction and Background

Modern societies rely increasingly on Data Center services. Companies of virtually any size, from any sector, i.e. government institutions and universities, rely on Data Centers services to facilitate business transactions, administration, information processing, and communications. All of these services depend on IT-infrastructures and Data Centers, which do not only continuously increase performance but can also become increasingly energy intensive. A recent study by the US EPA concluded that the power requirements of Data Centers in the United States more than doubled within six years. The electricity demand of US Data Centers grew from 28.2 TWh in 2000 to 61.4 TWh in 2006. This corresponds to a share of about 1.5 percent of the total electricity demand of the US in 2006. It was also reported (Schäppi, Bellosa et al. 2007) that the electricity demand of Western European (EU15 + Switzerland) Data Centers grew rapidly, ca. 34% during the period of 2003-2006.

Increasing energy prices and concerns about global warming apparently changed the prioritysetting within the Data Center industry. Not only is performance of servers and processors important, but so is also low environmental impact are also of great importance. In recent years, a number of initiatives emerged around the slogan “Green IT”. The “Green Grid” is a global consortium set on increasing the energy efficiency of data centers. The European Code of Conduct for Data Centers and Green IT have become a central topic of several industrial fairs such as CeBit. Nevertheless, the tentative trend appears to be overwhelmed by soaring demand for computing power. Not surprisingly, (EPA 2007a) and (Schäppi, Bellosa et al. 2007) found that with typical business trends, the energy consumption by US and European (EU27) Data Centers will more than double by 2011.

The rapid development of the data center industry provides great challenges to energy demand and efficiency analysis, even in describing the status quo. The lack of statistical information on the number of Data Centers, their current floor areas, power densities, and energy creates considerable technical uncertainties in estimating the data centers’ energy use. Policymakers and industrial initiatives that attempt to drive the market for efficiency need to be informed about the role Data Centers play in modern and future energy systems. Thus, credible estimates on current and future trends in Data Center energy usage are required.

## 2 Definition of Data Centers

There is no universal and distinct definition of the term “Data Center“. In fact, a number of diverse definitions and classifications can be found in literature. Definitions range from the purely qualitative (Cremer, Eichhammer et al. 2003; Fichter 2007) up to classifications that show both qualitative and quantitative elements (Tschudi, Xu et al. 2004; Turner, Seader et al. 2006; EPA 2007a). While, for example, (Cremer, Eichhammer et al. 2003) define data centers in a ”restricted sense“ as server farms that store and process data only for Internet applications, (EPA 2007a) defines it as a place where every room is dedicated to data processing servers, including server closets. And more generally, (Tschudi, Xu et al. 2004) define it as a “facility that contains concentrated equipment to perform one or more of the following functions: store, manage, process, and exchange digital data and information.” (Tschudi, Xu et al. 2004).

In light of these inconsistencies in terminology, it is invaluable to define the term Data Center for this report:

*A Data Center is a self-contained facility (room, building) that accommodates centralized IT equipment (Servers, Networks, and Storage.) and an ancillary site infrastructure (power delivery system and cooling system) to allow a reliable provision of the following services: storing, possessing, or exchanging of digital data and information.*

## 2.1 Eletricity Consumption in Data Centers

Data Centers range in size from small server rooms with dozens of servers and peak power demand in the kilowatt-range up to facilities that host tens of thousands of servers and have peak power requirements in the double-digit megawatt range. Based on a survey conducted amongst 31 German Data Centers, we found that the smallest Data Center had a peak power demand of only 40 kW, but the largest Data Center head one of 13.5 MW, which is comparable to a cement production facility with 1,200 kt annual output (IZE 2009). In spite of vast differences in the magnitude of the electricity consumed by Data Centers, the consumption can basically be traced back to four major sub-systems (cf. Figure 1):

**(1) IT equipment**

Continuous uptime is the crucial objective of Data Center operations as the costs of Data Center downtime can be substantial. Thus, IT equipment in Data Centers is in continuous uptime i.e. 8760 hours at best and operation, respectively. The electrical power consumed by IT equipment can range from 55 percent (typical) to 80 percent (best practice) of total data center electricity consumption.

**(2) Cooling systems**

Virtually all of the electrical power supplied to the IT equipment is converted to heat. This heat must be removed from the Data Center to keep the IT equipment operating properly (ASHRAE 2004). The cooling system accounts for anywhere between 35 percent (typical) to 10 percent (best practice) of total electricity use in Data Centers.

**(3) Power supply systems**

Power supply systems in Data Centers ensure that reliable and high-quality power is supplied to the IT equipment. Electricity is rectified several times from alternating current (AC) to direct current (DC) within the supply chain from the grid to the IT equipment. The power passes through an Uninterruptible Power Supply (UPS), which serves as a backup in case of power disruptions. Subsequently, a Power Distribution Unit (PDU) is connected to feed the different units of IT equipment. The power supply chain accounts for anywhere between 12 percent (typical) to 7 percent (best practice) of the total energy consumption of the facility.

**(4) Other power requirements**

Others power requirements, such as the electricity used for lighting, account only for a small share of the electricity consumed in Data Centers (1 to 4 percent).

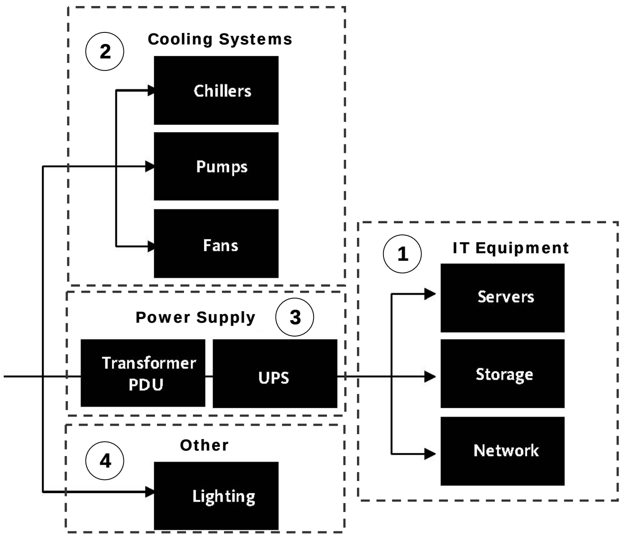


*Figure 1: Electricity Consumptions in a Data Center*

The impact of each subsystem on the overall electricity consumption differ considerably. The different distributions of electricity consumption by the four subsystems is illustrated using two German data centers as an example (Figure 2).

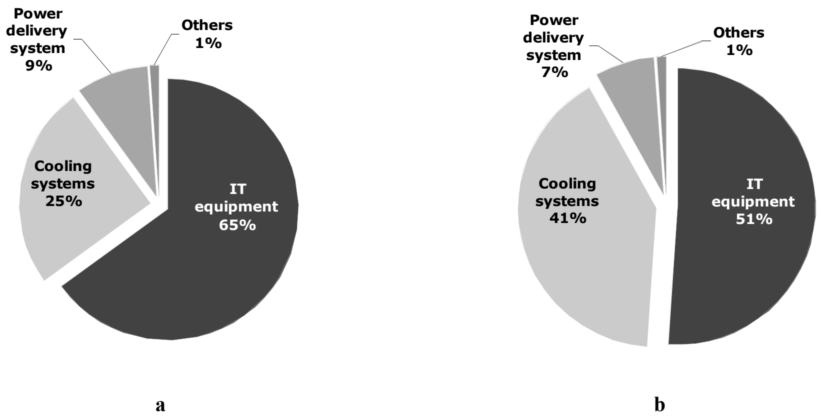


*Figure 2: Eletricity by End-Use for two German Data Centers in 2007 (IZE 2009)*

In the first Data Center (Figure 2a), 2/3 of the electricity is consumed for processing, management, and dissemination of data and information (main subsystem: IT equipment). Only 1/3 of the overall electricity is used for supporting subsystems like cooling and power delivery. In the 2nd Data Center (Figure 2b), by contrast nearly 50 percent of the overall electricity is consumed by supporting subsystems. Depending on the type (e.g., Housing/ Hosting), use (e.g., scientific, public or private businesses) and location (e.g., climate) of the Data Center, the energy consumption of the subsystems can consumption greatly. Therefore, explicit indicators are needed.

## 2.2 Energy Efficiency Indicators of Data Centers

Various energy efficiency indicators for Data Centers have been proposed and discussed in literature. For a detailed discussion of the different indicators, see (Rasmussen 2007). The Green Grid has established Power Usage Effectiveness (PUE)[1] as an indicator, which is currently the most frequently used indicator (GreenGrid 2008a). The PUE is defined as the quotient of the total power used in a Data Center divided by the power requirements of the IT-equipment:

$$\textbf{PUE} = \frac{\text{Total facility Power [kw]}}{\text{IT-Equipment Power [kw]}}$$

Theoretically, the range of the PUE is between 1 and infinity. A PUE of 2 implies that for every watt of IT load there is an additional watt of power required for the operation of the auxiliary site infrastructure.

But PUE lacks crucial information about the efficiency of IT-equipment itself, which could be measured as the amount of data processed per watt IT-load. Moreover, there are currently neither applicable methods of how to measure the actual IT-efficiency in Data Centers nor representative data available for these efficiencies. This report focuses on the energy efficiency of the site infrastructure.

[1]Another indicator, "Data Center infrastructure Efficiency" (DCiE), is also frequently used to evaluate data center efficiency. The DCiE is the reciprocal of the PUE.

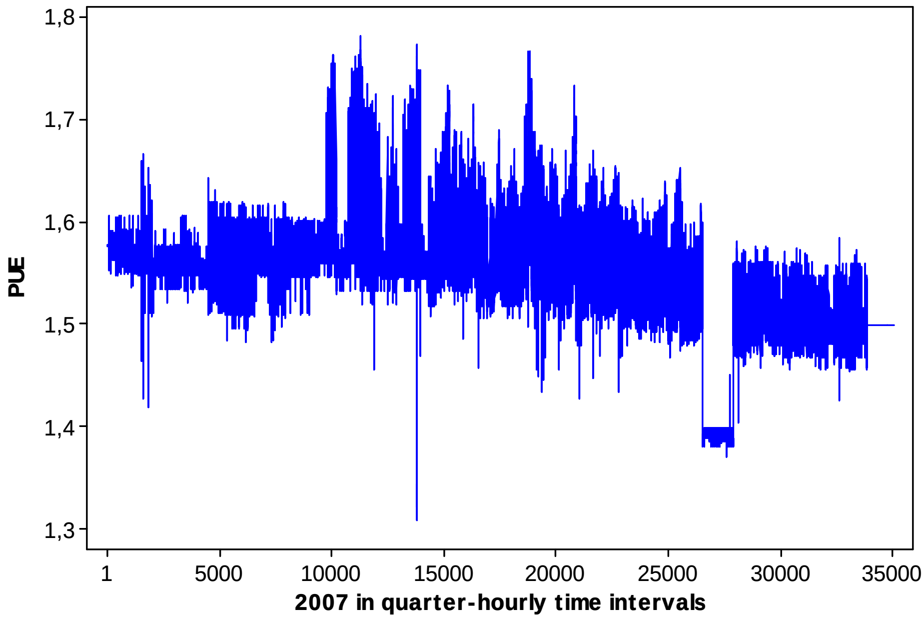


*Figure 3: PUE of a German Data Center in the course of 2007 (IZE 2009)*

Electrical power consumption in Data Centers is not constant and the PUE consequently fluctuates over time. Even under constant IT-Loads, the site infrastructure's power consumption is subject to variation (e.g., changing cooling loads, interval operation of cooling equipment etc. (Rasmussen 2007). Figure 3 illustrates the variation in the PUE of a data center throughout the course of 2007, based on quarter-hourly measurements. Strong fluctuation in PUE reveals the drawback of this static indicator, which is based on only capacity considerations. The minimum PUE is 1.31, the annual average 1.55 and the maximum 1.78.

Currently there are no standards on how to measure the PUE, or more importantly, on which period the PUE should be averaged. From the example presented in Figure 3 the manager of the Data Center could misuse the instantaneous PUE value of 1.3 for marketing purposes. Kenneth Brill of the Uptime institute comments on the marketing PUE relation: "It's important to understand that PUE has become a competitive tool for marketing manipulation and misinformation" (Brill 2008). Besides the risk of "marketing trickeries", a PUE not averaged over a period of one year or at least extrapolated to a period of one year by simulation, will lack the essential information related to the overall efficiency of a Data Center. "The use of a single instantaneous audit-type measurement of DCiE [or PUE] is not effective for benchmarking, it is not effective for trending, and it does not provide actionable insights into the opportunities for efficiency improvement" (Rasmussen 2008). For those reasons, we propose to use annual average PUE values only. The annual average of the

PUE corresponds to the Energy Usage Effectiveness (EUE) (Sijpheer 2008).

$$\textbf{EUE} = \frac{\text{Annual Energy Consumption Data Center [kw]}}{\text{Annual Energy Consumption IT-Equipment [kw]}}$$

## 2.3 Approaches to Model Electricity Consumption by Data Centers

A series of assessments involving data centers' electricity use can be traced back to the internet boom at the turn of the millennium. Several studies have been dedicated to estimating the electricity used by information technologies, including Data Centers for the US (Beck 2001; Mitchell-Jackson, Koomey et al. 2002; Roth, Goldstein et al. 2002), the Netherlands (Hartkamp 2002) , Switzerland (Aebischer, Frischknecht et al. 2003) , and Germany (Cremer, Eichhammer et al. 2003). Recently, the issue of energy consumption in information technology has been revived and has resulted in numerous studies that estimate, amongst other things, the energy consumed by Data Centers (See (EPA 2007a) for the U.S., (Fichter 2007) for Germany, (Clevers and Verweij 2007) for the Netherlands and (Koomey 2008) for worldwide estimates, split up into several regions (the US, Western Europe, Japan, Asia Pacific, and the rest of world)). In quantifying the electricity demand analysis of Data Centers on a national scale, two major methods have emerged. (Cremer, Eichhammer et al. 2003):

**Approach based on Data Center floor areas**

The yearly electricity consumption is calculated by multiplying the floor space [m2] by the power density [W/$m^2$] and the intensity of use [h] (a typical value here is 8760 h for year-round availability). The individual values for floor space are taken from national statistical data. The power density here includes both the IT equipment and the power drain from the site infrastructure. Amongst others, this approach can be found in (Beck 2001; Mitchell-Jackson, Koomey et al. 2002; Koomey, Sezgen et al. 2005; Sijpheer 2008).

**IT-Equipment based approach**

This approach estimates the energy use of Data Centers based on the stock of IT equipment operated in Data Centers and the associated unit energy consumption of the respective types. While earlier assessments did not quantify the electricity used by the site infrastructure (e.g., (Roth, Goldstein et al. 2002) and cp. (Koomey 2008)), recent studies incorporate this issue. (EPA 2007a; Schäppi, Bellosa et al. 2007; Koomey 2008).

Approaches based on Data Center floor areas require the collection and monitoring of a representative sample of Data Center floor space and the corresponding power densities. The corresponding statistical data are not available for Germany or for many other countries. (Koomey, Sezgen et al. 2005) describe the situation in the U.S. and California by saying: "It is surprisingly difficult to obtain market data on the floor area associated with Data Centers." . To the knowledge of the authors, statistical data on Data Center floor spaces and power densities have only been regularly ascertained in the Netherlands (Hartkamp 2002; Clevers and Verweij 2007).

Because of the lack of sound statistical data on floor spaces, the IT-Equipment based approach dominates the electricity demand modeling of Data Centers. Hence, most current international studies are based on this approach. These studies include the work of Jonathan Koomey (Koomey 2007a; Koomey 2007b; Koomey 2008), the work of the United States Environmental Protection Agency (EPA 2007a; EPA 2007b), and the study conducted by (Schäppi, Bellosa et al. 2007) in the framework of Intelligent Energy Europe Program. All aforementioned studies use the number of installed servers as the main indicator for the electricity used in Data Centers, as they are responsible for 70-80 percent of the electricity used by the IT-equipment there. Servers are differentiated into three server classes: volume server, mid-range server, and high-end servers[2]. This classification of servers is based on the convention of the market research institute "International Data Corporation" (IDC). They provided the statistical basis for the number of installed severs for the respective studies. This approach can be summarized as follows:

$$El(t) = \text{PUE} \cdot \text{hours} \cdot \sum_i \text{UPC}_{\text{Server}}(i,t) \cdot N(i,t) \qquad (1)$$

**Where:**

| | |
|---|---|
| $t$ | year |
| $i$ | volume-, mid-range- and high-end server |
| $El(t)$ | total electricity consumption of data centers [kWh] |
| PUE | Power Usage Effetiveness [-] |
| hours | 8760 hours [-] |
| $\text{UPC}(i,t)$ | average power consumption (UPC) of server class $i$ [W] in year $t$ |
| $N(i,t)$ | number of servers type $i$ in year $t$ |

According to equation (1), the total electricity consumption of servers including the Data Centers' site infrastructure, is derived from the product of the average unit power consumption (UPC) values for each server class, the installed base of the respective server classes, the annual intensity of use (8760 hours) and the power usage effectiveness (PUE). Assuming a constant PUE over a year, the PUE is equivalent to the

[2] This classification is based on price classes for servers. Volume-servers costs less than 25,000 U.S. dollars, mid-ranges servers costs up to 500,000 U.S. dollars, and high-end servers cost more than 500,000 U.S. dollars (IDC 2008).

EUE defined in section 2.2.

This approach was actually pioneered by Jonathan Koomey. He was the first to apply the EUE to capture the electricity consumed by servers and the electricity consumed by the associated site infrastructure in the Data Centers (Koomey 2007a). Since servers consume not all, but only 70-80 percent of the electricity used by the IT-equipment in Data Centers, (EPA 2007a) expands on Koomey's approach by including approximate data on the electricity used by storage and network equipment to estimate all the electricity consumed by Data Centers in the US. Additionally, (EPA 2007a) disaggregates and allocates the stock of servers to five Data Center classes: Server Closets, Server Rooms, Localized Data Centers, Mid-Tier Data Centers and Enterprise-Class Data Centers.

# 3 Electricity Demand Modeling of German Data Centers

As a conclusion to the previous discussion, we decided to apply the approach introduced by Koomey to model the electricity use of German Data Centers. However, following (EPA 2007a) we included an additional, but simplified representation of the energy consumption of the network and storage equipment. This was done to estimate the entire quantity of electricity consumed by Data Centers in Germany. Furthermore, we applied a vintage stock model approach to the installed base of servers. The vintage stock approach divides the total stock of servers into vintages, based on their year of commissioning. For each of the vintages and server classes, the specific unit energy consumption of the respective year of commissioning is assigned. The electricity consumption of German Data Centers is the sum over the amount of electricity consumed by the IT equipment (servers, network equipment and storage) multiplied by the Energy Usage Effectiveness corresponding to the electricity use of the site infrastructure (cooling an power delivery system). In the following, we individually describe each of the introduced terms.

## 3.1 Electricity Consumption of Servers

Data on the shipments and stocks of servers were obtained from (IDC 2008) and (TechConsult 2008a). The IDC data include the installed server base and server shipments for the years 1998 to 2008. This data is seperated by server classes: volume, mid-range, and high-end servers. The IDC data set is based on a top-down methodology that uses industry sales statistics to derive the installed base and shipments of the servers. IDC includes all servers except for customized ones, which are usually operated by large Internet companies such as Google. These servers are based on personal computer motherboards and therefore are not officially shipped as servers (c.p. (Koomey 2007a)). (TechConsult 2008a) uses a bottom-up methodology, based on a survey among 3,000 German companies, to ascertain the shipments and installed base of servers. TechConsult's data not only capture servers officially shipped as servers but also customized servers shipped as personal computer motherboards. For the year 2007, the installed server base s measured by TechConsult was 10% higher than the IDC data. The data from TechConsult was not available between the period 1998 to 2006, so we assumed the IDC data underestimated the installed base of servers by 10 percent between 1998 and 2006 and adjusted the analysis accordingly.

Based on the data from server shipments and server stocks, we developed a vintage stock model of the installed server base. The total stock of servers is modeled as a dynamic balance between the annual inflows of new server-vintages and the decommissioning of old server-vintages. The decommissioning of server-vintages is modeled by Weibull survival functions. Figure 4 provides an overview of the simulation results relating to the server stock developments. By inspection, the results from the server stock simulation are generally satisfactory for all of the modeled server classes. The simulation results show an expected service life of 3.6 years for volume servers, 5.7 years for mid-range servers and 6 years for high-end servers. When verified with other studies (Oliner 1993; OECD 2001; Kawamoto, Koomey et al. 2002), these results were consistent.

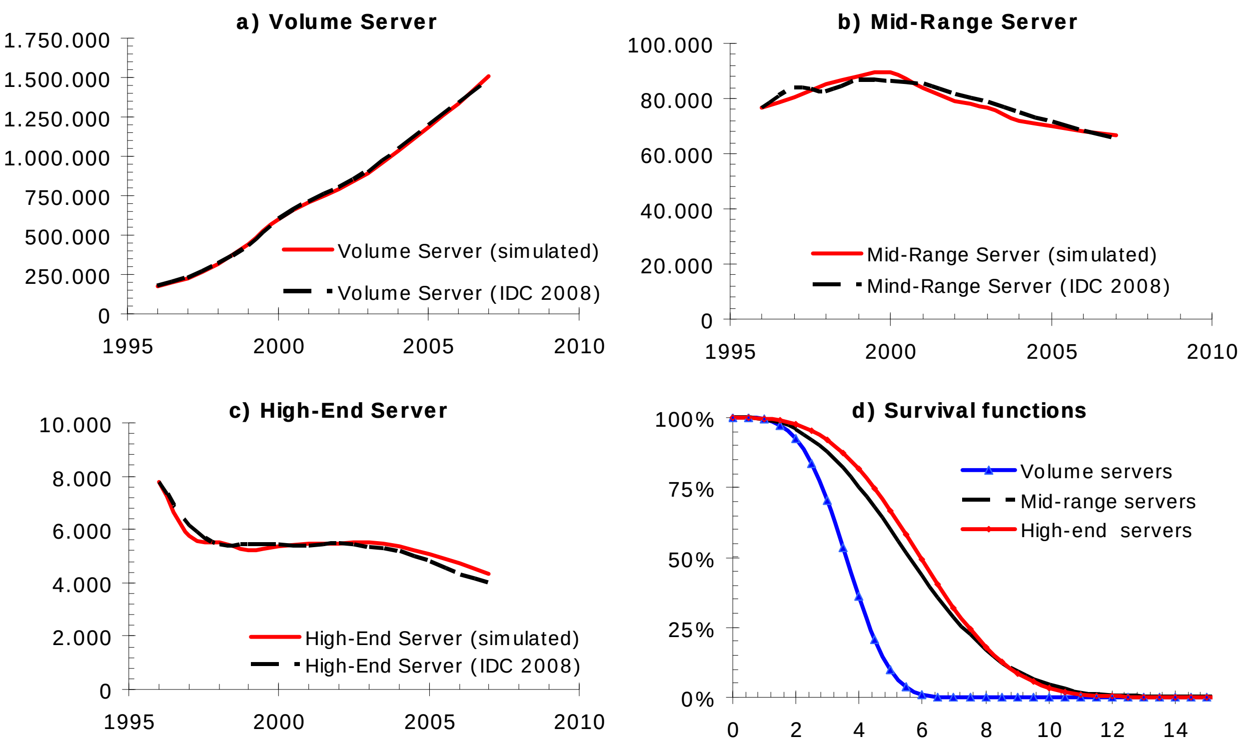


*Figure 4: Simulated installed server base and historic developments according to (IDC 2008; TechConsult 2008a) 1996-2008: a) volume, b) mid-range and c) high-end server. d) Survival function functions of volume, mid-range and high-end servers*

For further modeling, the unit power consumption of each server vintage and each server class has been assigned. The values are given in figure 5. We consulted various sources and found a range of estimates for each unit power consumption value. These ranges are plotted in the figure as error bars and gives the standard deviation from the mean unit power consumption value. For some years, however, there were not enough data points to derive meaningful values for the standard deviation. In those cases, we assumed a standard deviation of +/-15% of the mean unit power consumption, which is the average calculated for the years with sufficient data points.

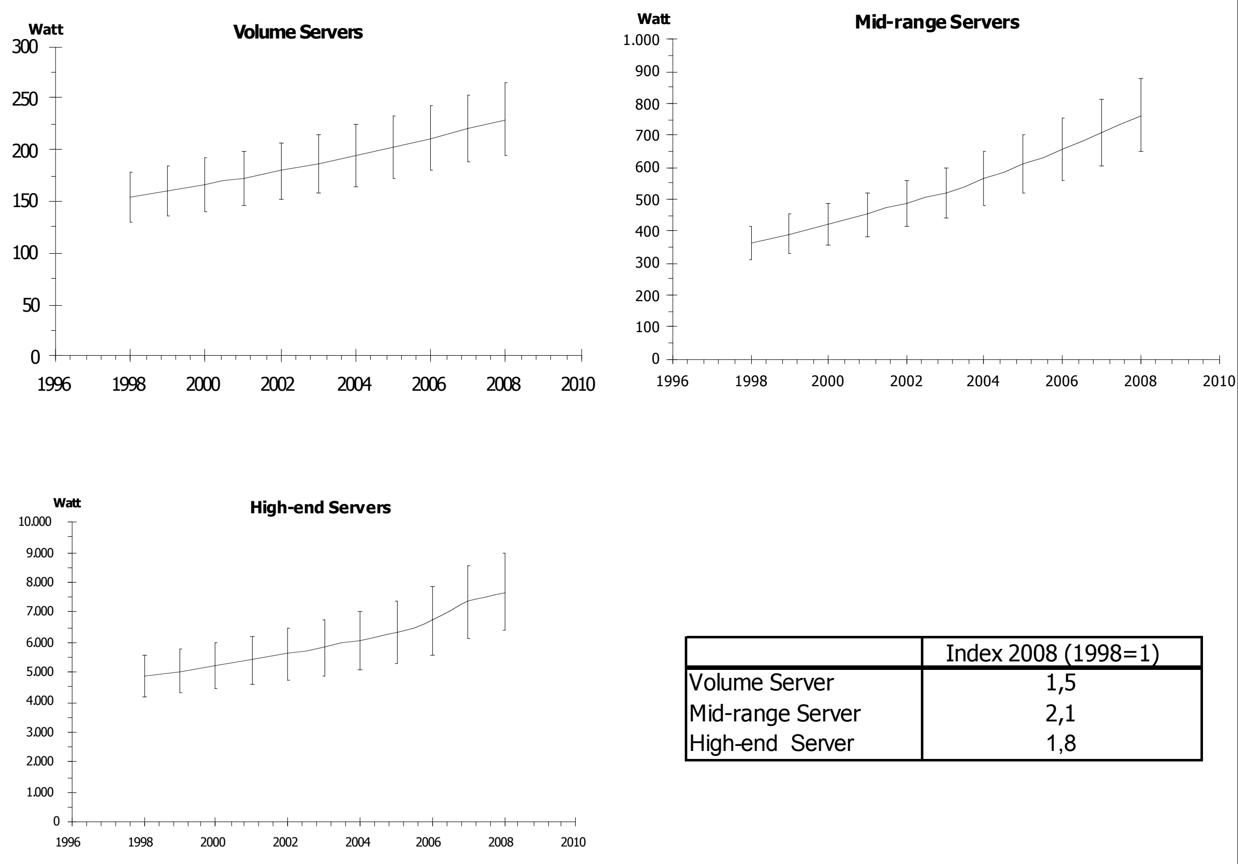


| | Index 2008 (1998=1) |
|---|---|
| Volume Server | 1,5 |
| Mid-range Server | 2,1 |
| High-end Server | 1,8 |

*Figure 5: Development of the unit power consumption of volume, mid-range, and high-end Servers in Germany 1998-2008 (data are based on: (Roth, Goldstein et al. 2002; Cremer, Eichhammer et al. 2003; Roth, Goldstein et al. 2004; Koomey 2007a; Schäppi, Bellosa et al. 2007; SPEC 2008))*

In contrast to intuition (e.g., obtained from unit energy consumptions developments of other appliances) newer server vintages have higher energy consumption than the older ones. In this context, the increase in unit power consumption went along with substantial improvements in the performance per watt. (Belady 2007) reported that the performance per watt of a typical server increased by a factor of 16 between 1999 to

2006. This corresponds to a doubling in performance every two years (Belady 2007).

### 3.2 Electricity Consumption of Network Equipment

The unit energy consumption of network equipment has been derived from US data found in (EPA 2007a). Because of the lack of data on the energy use of network equipment in data centers, (EPA 2007a) derives energy consumption of network equipment based on expert assumptions; the study assumes that volume servers operated in data centers will typically have three network ports, which are connected to network equipment with a power demand of 8 Watts (EPA 2007a; EPA 2007b). Consequently, the average unit power consumption of each volume server increases by 24 watts. Following this approach, we applied this method to our analysis. By considering the uncertainty within such assumptions, we assigned a normal distribution with a standard deviation of 10 percent to this value.

### 3.3 Electricity Consumption of Storage Equipment

We also extracted data on the electricity usage of external storage equipment of Data Centers from (EPA 2007a; EPA 2007b). We first derived the stock of enterprise storage in Germany based on the assumption that the ratio between volume servers and enterprise storage equipment is the same for Germany and the US. Given the amount of storage equipment, the unit power consumption has been assumed constant at 25 watts (again, based on data found in (EPA 2007a)). We also assumed that the unit power consumption is subject to uncertainty; thus, we assigned a normal distribution with a standard deviation of 10 percent.

### 3.4 Energy Usage Effectiveness (EUE)

Currently, there is no comprehensive industry data set for ”Power Usage Effectiveness” (PUE) and ”Energy Usage Effectiveness” (EUE). A common assumption is that the industry-wide average of the PUE is two (EPA 2007a; Koomey 2007a; Schäppi, Bellosa et al. 2007; Koomey 2008). However, this assumption lacks any empirical foundation. Moreover, a number of international studies offer contradictory values for EUE. Table 1 shows the minimum, maximum and the average PUE values from four different studies. None of the studies found average PUE values of two, and all studies show average PUE values below two.

| Country | Minimum | Average | Maximum | Number of Data Centers recorded | Source |
|---|---|---|---|---|---|
| USA | 1.33 | 1.86 | 3.03 | 19 | (Greenberg, Mills et al. 2006) |
| Germany | 1.17 | 1.67 | 1.90 | 13 | (ECO 2008) |
| Netherlands | 1.40 | 1.65 | 1.89 | 17 | (Celvers and Verweij 2007) |
| Germany | 1.32 | 1.74 | 2.77 | 16 | (IZE 2009) |

*Table 1: International comparison of PUE and EUE values for Data Centers*

In some studies, even the maximum PUE is below two. Figure 6 provides the distribution and descriptive statistics of PUE values in the aforementioned studies. Conclusively, this value is subject to a high uncertainty and therefore further research efforts are required.

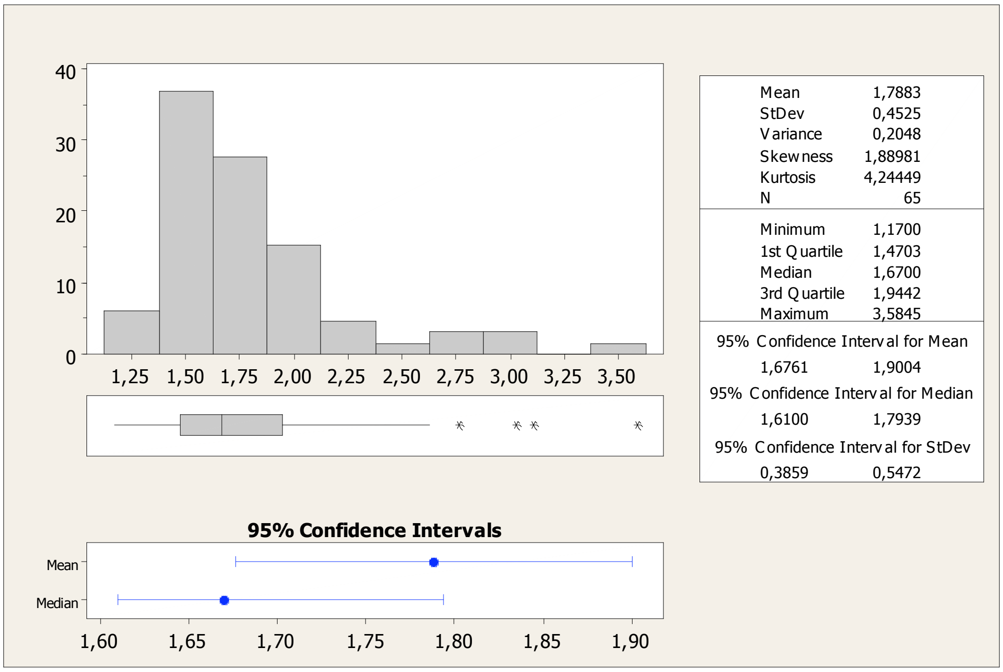


*Figure 6: International comparison of PUE and EUE values for Data Centers*

Aside from the “true” distribution and average EUE value problems, more questions arise. Determining the actual quantity of servers housed in Data Centers (according to the Data Center definition given in chapter 2) is problematic and therefore, the applicability of the EUE multiplier is questionable. As Koomey notes: “The IDC data also include servers that are not housed in Data Centers - the number and location of such servers may affect the appropriate value for estimating power used by cooling and the associated infrastructure“ (Koomey 2007a).

Figure 7 shows the installed server base in Germany in the year 2007 broken down into size ranges of installed server clusters.

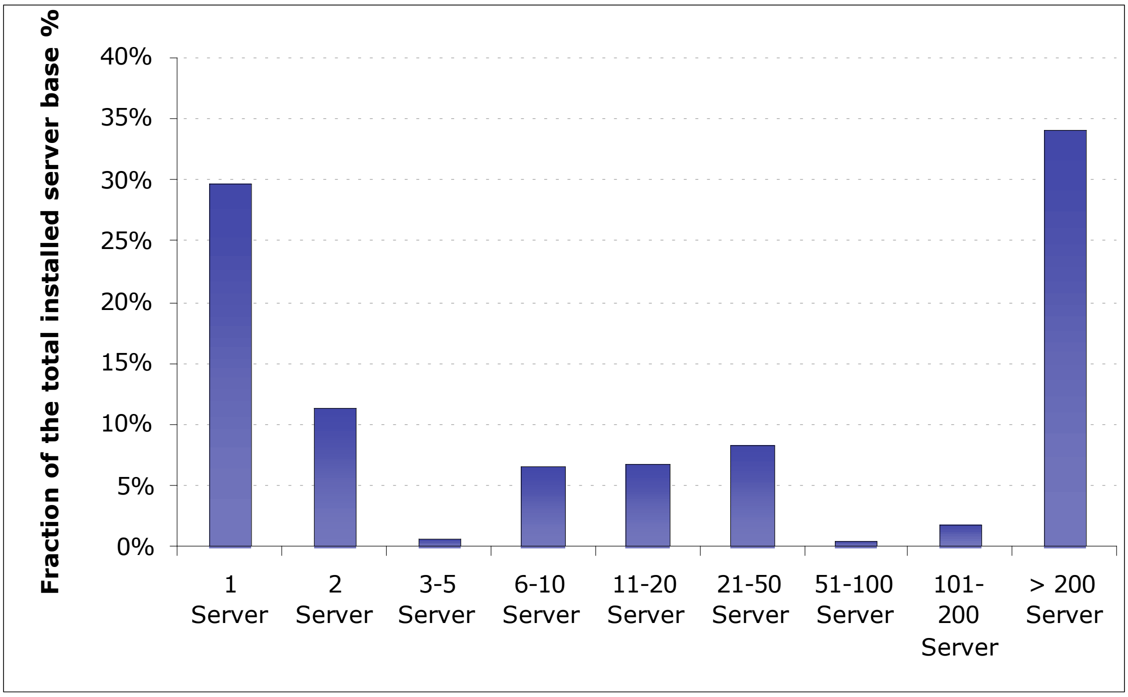


*Figure 7: Breakdown of installed server base by server clusters in Germany 2007 (TechConsult 2008b)*

Figure 7 illustrates that about 30 percent of the total servers installed in Germany are standalone, and 12 percent of the server stock is operated in clusters of only two servers. At the other end of the spectrum, about 34 percent of the servers are located in clusters larger than 200 servers. From these numbers we conclude that most of the standalone servers are not located in Data Centers or non air-conditioned environments. To incorporate this fact into the simulations, we apply a correction factor. This factor reduced the stock of volume servers not housed in Data Centers. As the true number of servers that are not housed in Data Centers is unknown, we apply a uniformly distributed correction factor with a value range of 75 to 100 percent. This implies that 75 to 100 percent of the severs are attributed to Data Centers within the simulations. Given the number of standalone servers provided in Figure 7, this is a conservative estimate.

## 3.5 Summary Modeling Approach

Equation (2) summarizes the main components of the modeling approach applied to the simulation of the electricity demand of German Data Centers.

$$E_{DC}(t) = \begin{bmatrix} \sum_i \sum_j \left[ N_{\text{server}}(i,j,t) \cdot \beta(\text{volume}) \cdot \langle \text{UEC}_{\text{server}}(i,j) \rangle \right] \\ + \text{UEC}_{\text{network}}(t) \cdot \sum_j N_{\text{server}}(\text{volume}, j, t) \\ + \text{N}_{\text{storage}}(t) \cdot \text{UEC}_{\text{storage}}(t) \end{bmatrix} \cdot \text{EUE}(t) \qquad (2)$$

**With:**

| | |
|---|---|
| $t$ | year |
| $i$ | volume-, mid-range- and high-end server |
| $j$ | server vintages (<1998, 1997-2007) |
| $E_{DC}(t)$ | total electricity consumption by Data Centers, year $t$ [kWh] |
| $\beta(\text{volume})$ | correction factor for volume servers not located in Data Centers |
| $N_{\text{server}}(i,j,t)$ | number of servers type $i$, vintage $j$, year $t$ |
| $\text{UEC}_{\text{server}}(i,j,t)$ | unit electricity consumption (UEC) server class $i$, vintage $j$, year $t$ |
| $\text{UEC}_{\text{network}}(t)$ | unit electricity consumption network equipment year $t$ |
| $N_{\text{storage}}(t)$ | number of servers class $i$, vintage $j$, year $t$ |
| $\text{UEC}_{\text{storage}}(t)$ | unit electricity consumption network equipment year $t$ [kWh] |
| EUE($t$) | Energy Usage Effectiveness year $t$ [-] |

The electricity use of German Data Centers:

**Sum over:**

**1** **Total annual electricity consumption of servers:**

Product of the stock of servers $\text{N}_{\text{server}}(i,j,t)$, distinguished by vintage $j$, server class $i$, and a correction factor $\beta$(volume) representing the share of volume servers not housed in Data Centers. Subsequently, the term is multiplied by the respective unit electricity consumption values $\text{UEC}_{\text{server}}\ (i,j,t)$.

**2** **Total annual electricity consumption of network equipment:**

Product of the stock of network equipment derived from the number of volume servers $\sum_j[\text{N}_{\text{server}}(\text{volume},j,t)]$ and the respective unit electricity consumption values $\text{UEC}_{\text{network}}(t)$.

**3** **Total annual electricity consumption of storage equipment:**

Product of the stock of storage equipment $\text{N}_{\text{storage}}(t)$ and the respective unit electricity consumption values $\text{UEC}_{\text{storage}}(t)$.

**Multiplied by:**

**4** **Energy Usage Effetiveness:**

EUE($t$) accounts for the energy use of the site infrastructure.

### 3.6 Dealing with Uncertainties: Monte Carlo Simulations

The input parameters for the estimation of data centers energy use are subject to uncertainty. To quantify the uncertainty within the model results caused by the model input values, we apply Monte Carlo simulation methods. Table 2 summarizes the main assumptions and input value ranges for the simulation.

| | Distributions | Values | Source |
|---|---|---|---|
| EUE | Triangular | 1.5, 1.7 and 2.5 | assumption based on data given in table 1 |
| $\beta$(correction factor volume servers) | Uniform | 75% to 100% | Assumption based on (TechConsult 2008b) |
| Unit energy consumption storage equipment | Normal | 25 Watt +/-10% | (EPA 2007a) |
| Unit energy consumption network equipment | Normal | 8 Watt +/-10% | (EPA 2007a) |
| Unit energy consumption servers | Normal | see Figure 5 | see Figure 5 |

*Table 2: Input parameters of the Monte-Carlo simulation*

### 3.7 Electricity Demand by German Data Centers in the Period 1998-2008

Figure 8 represents the results from the Monte Carlo simulations (10,000 iterations) performed to model the electricity consumption by German Data Centers in the period 1998 through 2008. The plots show the mean value, the standard deviation and the 5 percent / 95 percent percentiles of the simulation results. The results illustrate that the electricity demand of German Data Centers grew by a factor of about five, from 1.7 TWh in 1998 to 8.2 TWh in 2008. This corresponds with a compounded annual growth rate of 17 percent. In the year 2008, Data Centers were responsible for about 1.4 percent of the net electricity demand in Germany. Despite these enormous growth rates, the estimates presented in Figure 8 are considerably lower than previous estimates. In (Cremer, Eichhammer et al. 2003) the electricity consumption by servers in the year 2001 was estimated at 3.94 TWh. This estimate is about 1 TWh (i.e., 66 percent) higher than the results presented here (mean). In (Cremer, Eichhammer et al. 2003), the electricity consumption of the Data Centers' site infrastructures is not accounted for. This deviation, however, can easily be explained by the fact that (Cremer, Eichhammer et al. 2003) did not have access to comprehensive statistics from the installed server base; they instead used a survey from Switzerland and estimated the server stock based on server per employee relations. By using the assumption concerning the unit energy consumption of servers from (Cremer, Eichhammer et al. 2003) and data from (IDC 2008), the energy consumption by the servers decreases from the original 3.94 TWh in (Cremer, Eichhammer et al. 2003) to 1.7 TWh in our calculation. These numbers are within the range of the results from our estimates.

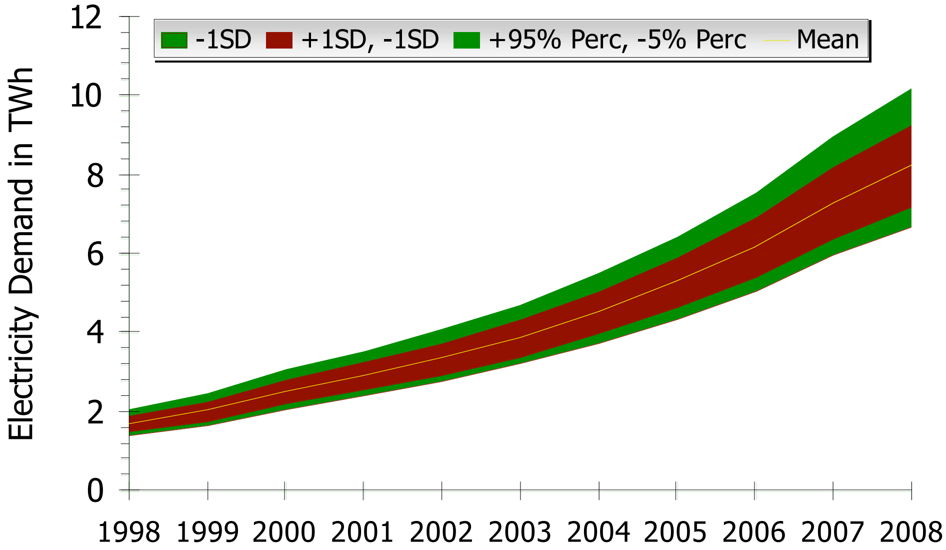


*Figure 8: Electricity demand of German Data Centers from1998-2008*

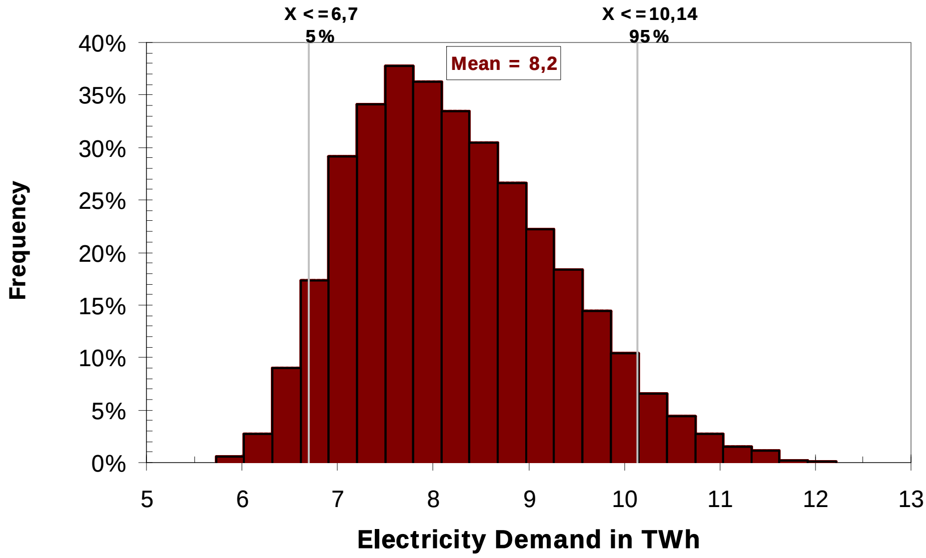


*Figure 9: Distribution of electricity demand in 2008*

Another, more recent study estimates the electricity use of German Data Centers based on the hypothesis that the relative contribution of Data Centers to the total national electricity demand is the same in the US and in Germany. The study uses estimates from (EPA 2007a), which state that data centers contribute 1.5 percent to the total electricity consumption in the US and approximates the electricity demand of German Data Centers as 8.67 TWh in the year 2006 (Fichter 2007). This estimate clearly exceeds the results obtained from our Monte Carlo simulation. Thus, the hypothesis that the relative contribution of Data Centers to the total electricity demand in Germany and the US is identical must be rejected.

Figure 10 shows a Tornado-Diagram. For the year 2008, this illustrates the relative importance of the input variables to the model results based on standardized regression coefficients.

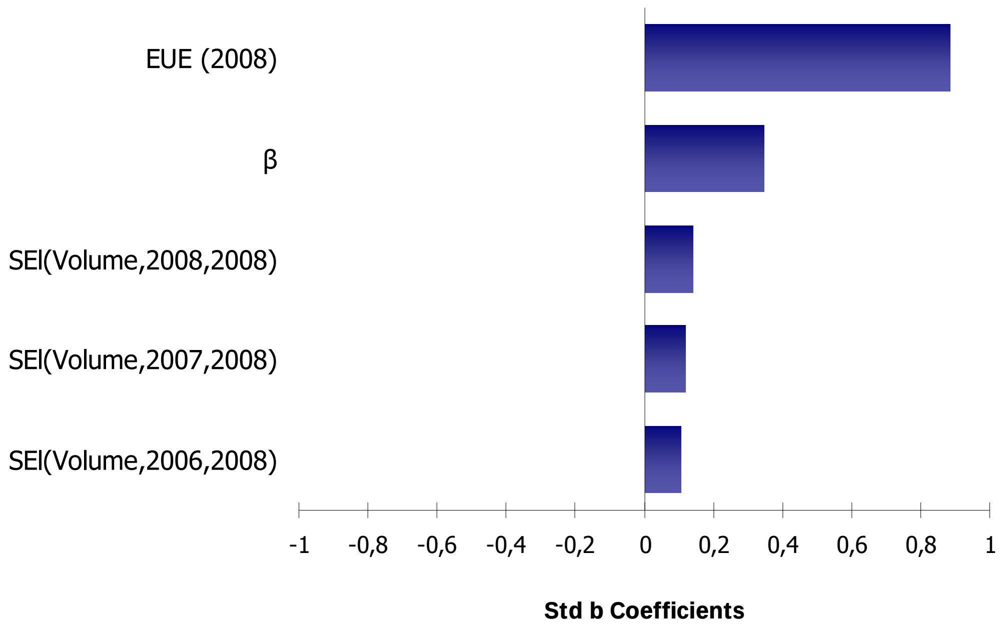


*Figure 10: Tornado-Diagram electricity demand of German Data Centers in 2008*

The tornado diagram illustrates that the EUE has the greatest impact on the modeling results. This agrees with intuition and clearly demonstrates the importance of this indicator. The next important factor is the correction factor beta. This determines how many of the volume servers are actually housed in Data Centers. Furthermore, both the unit energy consumption of the servers shipped in the year 2008 and the unit energy consumption of previous vintages are sensitive to the model results. This is in accordance with our vintage stock approach.

## 4 Closing Remarks

This paper applies a method introduced by Koomey (2007), which has been used in a number of subsequent studies (Koomey 2007a; Koomey 2007b; Schäppi, Bellosa et al. 2007; Koomey 2008) for estimating the electricity demand of Data Centers. The modeling approach has been supplemented by a vintage stock approach and uncertainty analysis. The results consider the most crucial uncertainties associated with this modeling method. We found that the Energy Usage Effectiveness and the amount of servers actually located in Data Centers have a relatively high degree of uncertainty. Furthermore, revealed that previous studies significantly overestimated the electricity demand of German Data Centers. Nonetheless, in the past decade the Data Center industry has experienced enormous growth rates in both electricity consumption and the number of Data Centers. With a compound annual growth rate of 17 percent, electricity demand almost quintupled between 1998 and 2008. Our current research focuses on scenarios projecting possible development pathways for the German Data Center industry to the year 2020.

# 5 Literature


Aebischer, B., R. Frischknecht, et al. (2003). Energy- and Eco-Efficiency of Data Centres Genève.

ASHRAE (2004). Thermal Guidelines for Data Processing Environments. Atlanta, American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

Beck, F. (2001). Energy smart data centers: applying energy efficient design and technology to the digital information sector. Washington DC, REPP Research Report.

Belady, C. (2007). "In the Data Center, Power and Cooling Costs More than the IT Equipment It Supports." Electronics Cooling Magazine **13**(1): 24-27.

Brill, K. G. (2008). "Don't Trust The Marketing Brochures." Forbes.

Clevers, S. H. and R. Verweij (2007). ICT Stroomt Door. Inventariserend onderzoek naar het elektriciteitsverbruik van de ICT-sector & ICT-apparatuur. Den Haag, Tebodin B.V.

Cremer, C., W. Eichhammer, et al. (2003). Energy Consumption of Information and Communication Technology (ICT) in Germany up to 2010. Karlsruhe/Zurich, FhG ISI / CEPE.

EPA (2007a). Report to Congress on Server and Data Center Energy Efficiency, Public Law 109-431. U.S. Environmental Protection Agency - Energy Star Program.

EPA (2007b). Report to Congress on Server and Data Center Energy Efficiency, Public Law 109-431: Appendices. U.S. Environmental Protection Agency - Energy Star Program.

Fichter, K. (2007). Zukunftsmarkt energieeffiziente Rechenzentren. Berlin Borderstep Institut für Innovation und Nachhaltigkeit GmbH.

GreenGrid. (2008). "Green Grid Metrics: Describing Data Center Power Efficiency." from www.thegreengrid.org/gg_content/Green_Grid_Metrics_WP.pdf

Hartkamp, F. (2002 ). Energy Consumption of ICT in the Netherlands, Tebodin B.V. and Novem.

IDC (2008). Installed Server Base Germany 2000-2007. International Data Corporation. Frankfurt.

IZE (2009). Energy and Resource Efficiency of German Data Centres, Innovationszentrum Energie -TU Berlin, (German language).

Kawamoto, K., J. G. Koomey, et al. (2002). "Electricity used by office equipment and network equipment in the US." Energy **27**(3): 255-269.

Koomey, G. K., O. Sezgen, et al. (2005). Estimating Total Power Used by Data Centers in California, LBNL.

Koomey, J. G. (2007). Estimating total power consumption by servers in the U.S. and the world.

Koomey, J. G. (2008). "Worldwide electricity used in data centers." Environmental Research Letters **3**(3): 034008.

Mitchell-Jackson, J., J. G. Koomey, et al. (2002). "National and regional implications of internet data center growth in the US." Resources, Conservation and Recycling **36**(3): 175-185.

Oliner, S. D. (1993). Constant-quality price change, depreciation, and retirement of mainframe computers. Price Measurements and Their Uses. M. F. Ross, M. E. Manser and A. H. Yong. Chicago, University of Chicago Press**:** pp. 19-61.

Rasmussen, N. (2007). Electriccal Efficiency Measurement for Data Centers. White Paper No. 54. APC, American Power Conversion

Roth, K., F. Goldstein, et al. (2002). Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings - Volume I: Energy Consumption Baseline. Washington, DC, Prepared by Arthur D. Little for the U.S. Department of Energy. A.D. Little Reference no. 72895-00. .

Roth, K., F. Goldstein, et al. (2004). Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings--Volume II: Energy Consumption Baseline. Washington, DC, Prepared by Arthur D. Little for the U.S. Department of Energy. A.D. Little Reference no. 72895-00. .

Schäppi, B., F. Bellosa, et al. (2007). Energy efficient servers in Europe: Energy consumption, saving potentials, market barriers and measures Part I: Energy consumption and saving potentials. Energy efficient servers in Europe.

Sijpheer, N. C. (2008). Energy saving at data hotels; Doing more with less. Petten, ECN.

SPEC (2008). SPECpower_ssj2008, Standard Performance Evaluation Corporation.

TechConsult (2008a). Installierte Serverbasis nach Branchen und Mitarbeitergrößenklassen. Kassel.

TechConsult (2008b). Verteilung der Serverbestands 2007 in Deutschland nach Größe der Servercluster. Kassel.

Tschudi, W., T. Xu, et al. (2004). High-Performance Data Centers - A Research Roadmap. Berkeley, Lawrence Berkeley National Laboratory.