White Paper
eco Datacenter Expert Group

Guide to Power Density and Load Determination of Servers, Data Cabinets and Data Centers

Ulrich Terrahe, Marc Wilkens

5th June 2012

datacenter.eco.de

Association of the German Internet Industry

WE ARE SHAPING THE INTERNET.

eco

Content

About the Authors

Ulrich Terrahe,
CEO, DC-CE RZ-Beratung GmbH

Ulrich Terrahe, Dipl.-Wirtsch-Ing (FH), studied at the Technical Universities in Gießen and Berlin. His first professional engagements were in Lennox Industries, England, Kraftwerksunion, Rudolph Otto Meier and Raab Karcher Wärmertechnik.

Since 1997 he has worked exclusively in the planning of data centers. Initially, he worked for Schnabel AG (from 2000 to 2007 as a Board Member), and after this as the owner of his own business, dc-ce RZ-Beratung. He is a specialist in the data center industry and offers consultation and planning services for data center infrastructure.

Mr. Terrahe presents papers on a variety of data center themes at national and international conferences. In 2007 in London he won the European Datacenter Award in the category "Future Thinking and Design Concepts" and has been a member of the jury ever since. Today, he is the organizer of the event "future thinking" and the "German Data Center Prize", a platform for innovative development in the industry.

http://dc-ce.de

Marc Wilkens,
Research Associate, TU Berlin, FG IKM

Marc Wilkens is an expert for energy efficiency in data centers. He developed the process for data center benchmarking at the TU Berlin, and an extensive coding system which forms the basis for the development of sustainable energy-saving potential.
http://blauerengel.ikm.tu-berlin.de/

From March 2010 to April 2011 Marc Wilkens was deeply involved as a Technical Consultant in drawing up the foundations for awarding the "Blauer Engel für energiebewussten Rechenzentrumbetrieb" ("Blue Angel for energy-conscious data center operation") under the leadership of the Federal Environmental Agency. He has led the new Competence Group "Sustainable IT Management" at the eco-Association since 2011.

Marc Wilkens graduated from his degree in Industrial Engineering at the TU Berlin in autumn 2006. From 2008 to 2010 he was Managing Director of "IFE-RE" at the Innovation Center: Energy (IZE) at the TU in Berlin. Under his leadership, the "Concept Study on Energy and Resource Efficiency in the Operation of Data Centers", among other studies, was produced.

Marc Wilkens has been Project Leader for the TU Berlin in the Federal Ministry for Economics and Technology project "Government Green Cloud Laboratory" since June 2011. Apart from this, he is working on his doctorate on the theme "Energy Efficiency in Data Centers" – his study deals with the development and testing of new parameters for the representation of the overall efficiency in data centers.
http://www.it2green.de/de/778.php

His experience in a multitude of data center projects stands him in good stead as a freelance Consulting Engineer. In this work he draws up independent expert reports on the energy efficiency for the planning and operation of data centers, and supports his clients in the practical implementation of the measurement concepts.
http://wilkens-esb.de/

Foreword

Although there have already been data centers, data center planners and data center service providers for many years, it was with the liberalization of Telecommunications in 1996 that the basis for today's data center landscape was first laid out. The unrestrained imagination of the dotcom boom led to numerous data centers worldwide being built, which were only to be filled years later. 9/11 also provided its sad contribution to a new understanding of data centers and security against outages, and increased the demands for resilient design. Not to be forgotten, Moore's law has retained its validity for more than 50 years, and IT hardware manufacturers love to adorn themselves with the concept "bigger, better, further".

In the early days of our just 16-year-old industry, only moderate power and cooling densities of between 500 W/m2 and 800 W/m2 were used as the basis of design. Today the norm is 1,000 W/m2 and often the expansion to 1,500 W/m2 is planned for. These changed requirements are also clearly visible in requests for proposals. Ten years ago, data center capacities were expressed in square meters – today the kW output and the power density are the priorities. The corresponding areas can easily be calculated. This shows how, in just one decade, our industry has completely turned around.

The producers of IT hardware and software have had a not insignificant role to play in this change. The capabilities of an off-the-shelf smartphone from today exceed those of a standard server from 1996. Where 10 to 15 years ago we had to fight for space for the servers, for the last few years we have been seeing ever longer and wider racks appearing, so that there is sufficient space for the power and data cables. From 4 U space servers to 2 U space servers, pizza boxes and blade servers, miniaturization has entered the world of the modern data center, and technologies such as water-cooled racks have been developed to provide high density cooling.

Thanks to PowerPoint, pdf, jpg, YouTube and countless billion emails – just to name some examples – the demand for memory capacity has also risen dramatically. And it's not just that the requirement for storage units is increasing almost exponentially, but also that their weight – at around 1.5 t/m2 – is creating a small challenge for almost every data center. Added to that, the power consumption increases accordingly with each fully-equipped rack of servers.

More supply and more variety has led inevitably to an extremely diversified data center landscape with very different specifications for redundancy, security against outage, power and cooling density. As a result the almost unified design found 15 years ago is a thing of the past. This means that builders and planners need to carry out exacting preparation with regard to the client-specific requirements for today and for the next 10 to 15 years, as this is the length of life a data center is designed for.

Unfortunately, experience has shown that the power density per square meter or per rack is often greatly overestimated by the client. This may be as a result of caution, or of relying too heavily on manufacturer specifications and peak current. Along with the other challenges, realizing these high power densities and cooling capacities results inevitably in an over dimensioned solution which is too expensive for all parties.

For these reasons, an expert group was established in January 2012 with the goal of defining realistic power and cooling densities for the varying requirements in the modern world. The participants of this expert group were drawn from the areas data center operation and planning, server and software manufacture, and also research and education, and they presented and discussed their knowledge.

The initial conclusions of this expert group, which is only the beginning of many further discussions, is recorded in this White Paper. We don't need a crystal ball to realize that these results will need to be regularly adapted to the new circumstances over the next few years. But they still provide us with a reliable basis for further activities.

Happy reading!
Dr. Béla Waldhauser
Leader, Competence group Data Center Infrastructure

Target Group

The White Paper is addressed at planners, developers, operators and owners of data centers and server rooms for large and medium-sized companies in the manufacturing industry and commercial activities such as banking and insurance, whose data centers are essential for their business operations. The data centers have a heterogeneous application structure, with applications like exchange server, web, print, fileserver, CRM systems, accounting software, database, backup, and archiving systems.

The IT services mostly use standard, commercially available hardware, such as one, two or more U space servers, or sometimes equivalent blade server technology. Further, it is assumed that the appropriate network technology and, if applicable, telecommunications technology is an integrated feature.

The guideline figures introduced here for the power density in data centers are not intended for highly specialized data centers providing services such as colocation and webhosting, or high-performance computing.

Presentation of the Problem

One of the most important constructional criteria for a data center is the power supply to be provided. What electrical capacity is needed for the data center today and in the future, and can be calculated for a lifetime of 10 to 15 years?
As there are often no measurements to use as a basis, and the ongoing development of computer technology is hard to predict, miscalculations are common.

The power requirements are discussed within the market range of 2kW to 20kW per rack. What makes this more difficult is that at specialist conferences as well as in product presentations, only the high-performance solutions are discussed; which then give the impression that inevitably more and more supply is required to the same limited space.

The power density miscalculations have an impact on all areas of the data center.

The result of excess provision is that more space is used for ICT and support infrastructure than is actually required. The technical facilities such as transformers, UPS, diesel and switch plants, batteries, climate control and air-conditioning systems are correspondingly larger.

As a rough estimate, every kW overestimated costs 3,000 to 5,000 euros, for a data center with a n+1 redundancy strategy.

A server room with 5 data cabinets, where the provisioning for each cabinet is 10kW but the requirement is actually only 5 kW, would be 25kW over-provisioned, and would entail between Euro 75,000 and Euro 125,000 of unnecessary investment.

Alternatively, the miscalculation can lead to the under-provisioning of the infrastructure, so that the data center reaches its capacity limit too quickly, and expansion is no longer possible. This variant of miscalculation is, however, less common.

Objective

Until now there has been very little information available to provide independent and vendor-neutral help to those responsible for the construction or modification of data centers, in order to minimize the risk of a miscalculation in the capacity.

This White Paper provides a rough guide to the correct capacity for the ICT equipment in a data center.

In principle, making use of a configuration program to analyze energy requirements of the existing IT and the energy requirements of future IT is recommended, in order to calculate the actual requirements.

Independent assessments for "best and worst-case scenarios" should then be incorporated.

The data used as a basis for this guide

The conclusions of this White Paper are based on research, calculations and measurements, as well as on field experience. The starting point was an expert group which met on the 18th January 2012 in Darmstadt to analyze and discuss the existing and future power requirements for data centers. The experience of data center operators and planners, and server and software producers, along with research findings from the OFFIS Institute and the Technical University Berlin were considered. The capacity of a variety of servers with different loads was measured, as well as being calculated using the configuration programs supplied by the manufacturers. Furthermore, a study from the TU (Datacenter Benchmarking), in which the data from 74 data centers was collated, was taken into account. On the basis of this information, "Best Practice Scenarios" were developed, which condensed the results into a widely applicable trend forecast.

Rack Categories

The White Paper differentiates four distinct rack categories for the target group.

A value for each rack category is determined. It can be assumed that more than 80% of the currently operational racks corresponding to the respective category have values lying under the levels of the applicable conditions for that category, and therefore also lower than the calculated power supply. These values are only applicable for racks used for the housing of ICT equipment which is less than four years old. Older ICT components are often less energy efficient and have not been taken into account in these calculations. Older ICT equipment is unlikely to be placed in new data centers or modified existing data centers.

Category 1

Rack Category 1, "Normal" is a standard cabinet with 42 Units, an average level of occupancy of (active) ICT components of 70%, and a heterogeneous server structure composed of one U space, two U spaces, or more. The proportion of physical servers on which virtual machines run (degree of virtualization) is less than 25%, and the mean load of the server is 20% of maximum. The expert group agreed that, within the above-named target group, more than 80% of the racks with a low degree of virtualization have values lower than these configuration characteristics.

Category 2

Rack Category 2, "Virtual", is a standard rack with an average level of occupancy of (active) ICT components of 70%, and a heterogeneous server structure composed of one U space, two U spaces, or more. The rack occupancy is therefore identical with Rack Category 1 (see above). However, in this category the degree of virtualization of the servers allowed for is over 50%, whereby the mean load of the servers rises to 60% of maximum. According to the discussions of the expert group, within the above-named target group, more than 80% of the racks with a high degree of virtualization have values lower than these configuration characteristics.

Category 3

In Rack Category 3, "Blade Normal", blade centers are utilized which essentially function as a replacement for the standard server. It assumes that there are no highly specialized applications in operation. The prime example for Category 3 is two blade centers and three 1 U servers, with a level of occupancy of 50%. The maximum load for the servers is set at 60% of maximum.

Category 4

In Rack Category 4, "Blades with higher demands", blade centers are also used, but these are used for highly specialized applications. It is assumed that 4 blades are used, that the level of occupancy is at 70%, and that the maximum load of the servers is 75% of maximum.

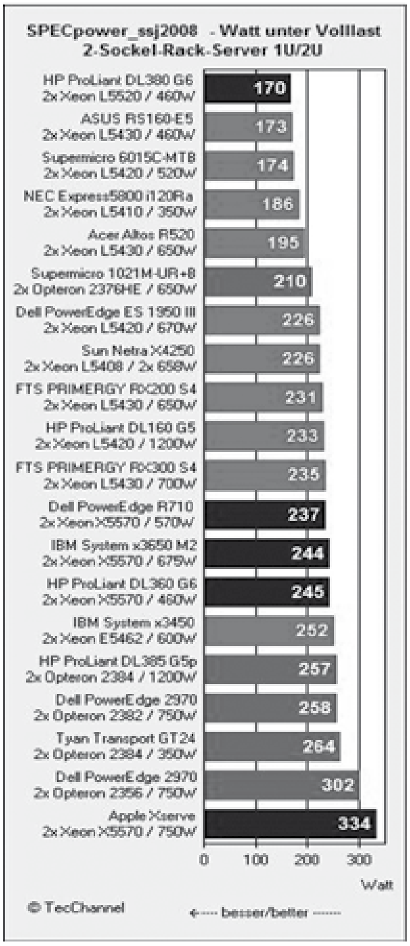

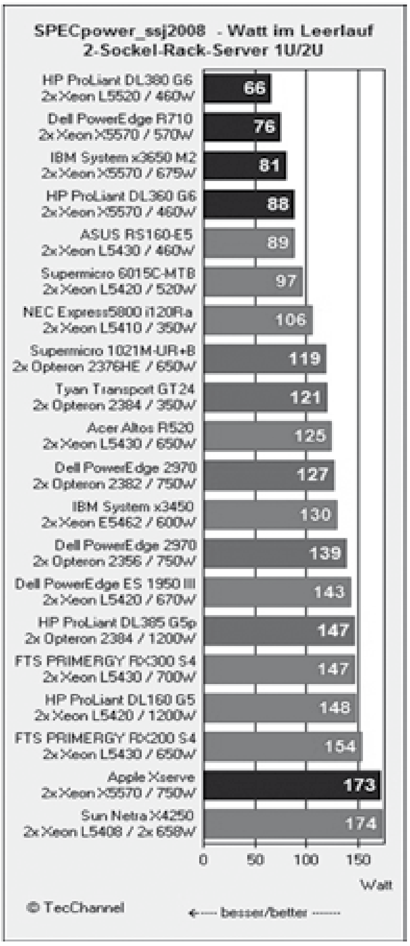

Table 1

Power consumption 1/2U Servers in max. and idle mode

(Source: TecChannel, 2009)

Power consumption per rack

(according to SPEC, determined through the TU Berlin)

Table 2

Load	1 U Server	2 U Server	4 U Server	Bladecenter
100%	200 W	280 W	680 W	2950 W
60%	145 W	220 W	560 W	2000 W
20%	100 W	160 W	450 W	1500 W

Analysis of Server Load

The basis for the calculation of power supply is the power consumption of the server. Here we differentiate between the mains power supply capacity, the maximum power consumption and the power in partial load operation. In order to calculate the total consumption of the data center, under no circumstances should the capacity of the power supply units of the servers be simply added together. This leads to over capacity, with the affects described above.

It is highly recommended that appropriately specialized experts validate the calculations to safeguard the servers.

The actual power consumption of the server can be obtained through measurements. Alternatively, it is possible to get power consumption data for most servers (and server components) from the independent Institute SPEC (Standard Performance Evaluation Corporation), which can instead be used for the calculated values.

Table 1 shows the standard 1U/2U servers from TecChannel with their consumption under maximum load and in idle mode. These loads lie between 170 and 334 watts in full operation, and between 66 and 174 watts when idle.

In table 2, the SPEC data on various servers of a range of technologies or unit sizes is evaluated for the years 2008 to 2011. From this, average consumption can be calculated for comparable technology from various manufacturers. These averages were calculated for different server loads (100%, 60%, 20%). For the planning of new server rooms or data centers, consumption and performance values which are worse than these average values should not be used on priciple because the efficiency of new server technology has generally improved in recent months.

From this data analysis it can be seen that the load behaviors of standard servers for the aforementioned target group are quite similar, which suggests that the assessment in this White Paper has a good level of accuracy. Nevertheless, it should be mentioned that there can be a deviation of 100% between a server under low load and a server under high load, so that a detailed analysis of the ICT components is always sensible.

For the initial needs analysis for the planning of IT loads in data centers, the following guide is nonetheless a good basis.

Example Calculation per Category

Following are four typical examples for ascertaining the power densities for racks with different assemblies of ICT components.

Example Rack Category 1 “Normal”:

For the calculation of an example rack corresponding to the above-defined Category 1, a 42 U rack was equipped with six 1 U servers, seven 2 U servers, two 4 U servers, and two switches with a total of 30 active ports. This corresponds to a rack occupancy with active components of approximately 70%, or 30 U.

For the above-defined server load of 20%, the relevant SPEC server types result in the following mean power consumption:

Consumption calculation per rack for Rack Category 1

Source: TU Berlin, FG IKM (2012) nach SPEC-Serverdaten von 2008-2011

Table 3

Occupancy		Average Load					
Number	Server Type	20% (Rack Category 1)		60% (Rack Category 2)		100% (theor. max. load)	
		1 Server	Total	1 Server	Total	1 Server	Total
6	1 U Server	100 W	600 W	145 W	870 W	200 W	1200 W
7	2 U Server	160 W	1120 W	220 W	1540 W	280 W	1960 W
2	4 U Server	450 W	900 W	560 W	1120 W	680 W	1360 W
2	Switch (30 active ports)	180 W	360 W	180 W	360 W	180 W	360 W
Consumption Value per Rack			2980 W		3890 W		4880 W

Category 1 therefore requires a power density of almost 3kW per rack.

For the power consumption of more recent switch models, one can assume that the energy requirements increase with the number of ports; the power requirements per switch is currently not or only minimally dependent on the load on the ports.

Example Rack Category 2 "Virtual":

For the calculation of an example rack corresponding to the above-defined Category 2, a 42 U rack was equipped with six 1 U servers, seven 2 U servers, two 4 U servers, and two switches with a total of 30 active ports.

The rack occupancy is therefore identical with Rack Category 1 (see above). However, in this category it is assumed that, as a result of the virtualization of the servers, the mean load of the servers rises to 60%. This results in the following mean power consumption:

Consumption calculation per rack for Rack Category 2

Source: TU Berlin, FG IKM (2012) nach SPEC-Serverdaten von 2008-2011

Table 4

Occupancy		Average Load					
Number	Server Type	20% (Rack Category 1)		60% (Rack Category 2)		100% (theor. max. load)	
		1 Server	Total	1 Server	Total	1 Server	Total
6	1 U Server	100 W	600 W	145 W	870 W	200 W	1200 W
7	2 U Server	160 W	1120 W	220 W	1540 W	280 W	1960 W
2	4 U Server	450 W	900 W	560 W	1120 W	680 W	1360 W
2	Switch (30 active ports)	180 W	360 W	180 W	360 W	180 W	360 W
Consumption Value per Rack			2980 W		3890 W		4880 W

The power density for Category 2 stands at almost 4 kW per rack for the described rack occupancy and server load.

Example Rack Category 3 "Blade Normal":

For the calculation of an example rack corresponding to the above-defined Category 3, a 42 U rack was equipped with two blade centers, three 1 U servers, three 2 U servers, and two switches with a total of 16 active ports. This corresponds to a rack occupancy of 24 U; the rack is therefore close to 60% equipped with active components.

Without the operation of highly specialized applications (e.g. batch jobs in high-performance computing), the utilization of blade technology in the racks results in a server load of maximum 60%.

This results in the following mean power consumption for the relevant SPEC server types:

Consumption calculation per rack for Rack Category 3

Source: TU Berlin, FG IKM (2012) nach SPEC-Serverdaten von 2008-2011

Occupancy		Average Load					
Number	Server Type	20%		60% (Rack Category 3)		100% (theor. max. load)	
		1 Server	Total	1 Server	Total	1 Server	Total
3	1 U Server	100 W	300 W	145 W	435 W	200 W	600 W
3	2 U Server	160 W	480 W	220 W	660 W	280 W	840 W
2	Blade Center	1500 W	3000 W	2000 W	4000 W	2950 W	5900 W
2	Switch (16 active ports)	150 W	300 W	150 W	300 W	150 W	300 W
Consumption Value per Rack			4080 W		5395 W		7640 W

Table 5

The power density for Category 3 stands at almost 6 kW per rack for the above-described rack occupancy.

Example Rack Category 4 “Blades with higher demands”:

For the calculation of an example rack corresponding to the above-defined Category 4, a 42 U rack was equipped with four blade centers and two switches with a total of 30 active ports. This corresponds to a rack occupancy of 30 U; the rack is therefore close to 70% equipped with active components.

It can be assumed that for this intensive operation of blade technology there is precise scheduling of server loads (e.g. using scheduled batch jobs). The server load therefore rises to 75%, and results in the following power consumption:

Consumption calculation per rack for Rack Category 4

Source: TU Berlin, FG IKM (2012) nach SPEC-Serverdaten von 2008-2011

Occupancy		Average Load					
Number	Server Type	20%		75% (Rack Category 4)		100% (theor. max. load)	
		1 Server	Total	1 Server	Total	1 Server	Total
4	Blade Center	1500 W	6000 W	2400 W	9600 W	2950 W	11800 W
2	Switch (30 active ports)	180 W	360 W	180 W	360 W	180 W	360 W
Consumption Value per Rack			6360 W		9960 W		12160 W

Table 6

The power density for Category 4 stands at a maximum of 10 kW per rack for the described rack occupancy and server load.

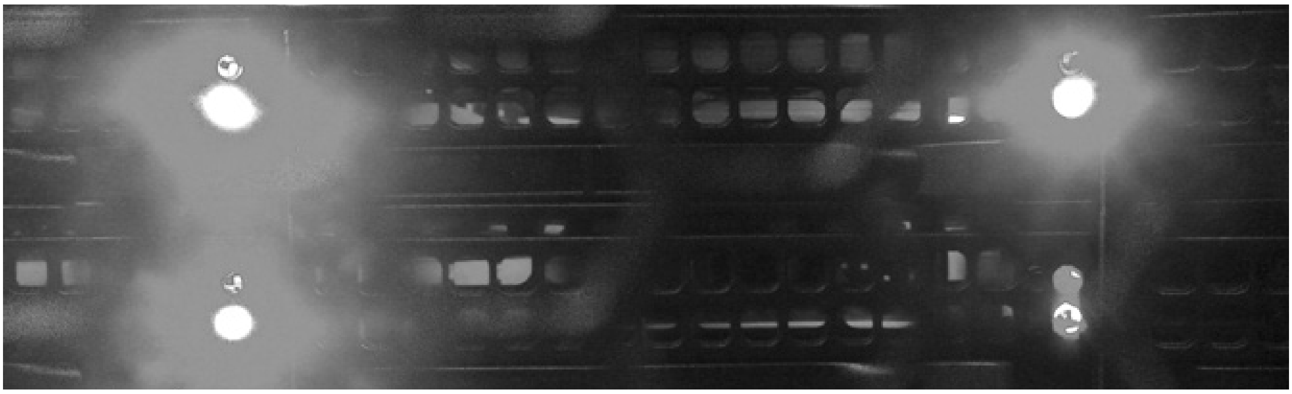

Power Density in the Field

Results from Data Center Benchmarking at the TU Berlin

The above-specified assumptions and calculations were confirmed through the results of a survey conducted by TU Berlin on the energy efficiency of data centers.

The TU Berlin developed Data Center Benchmarking (DCB) through the subject areas Information and Communication Management and Energy Process Engineering. Since 2009 data on energy requirements and data center infrastructure has been regularly collated there. So far, 74 data centers have taken part in DCB. The analysis of the data shows that the power density in data centers is also dependent on the purpose of operation (see assumptions on rack categories above). The results of the DCB demonstrate clearly the differences for the Categories 1 and 2, as well as for 3 and 4. Data to the Rack Categories 3 and 4 have not yet been separately collated in the DCB.

Rack Category 1:

In data centers where the ICT predominantly corresponds to Rack Category 1, the servers are as a rule not, or only very minimally, virtualized. In the DCB, only 2 Category 1 data centers have minimal virtualization; the other 15 data centers operate no virtual machines.

Power Density Rack-Category 1

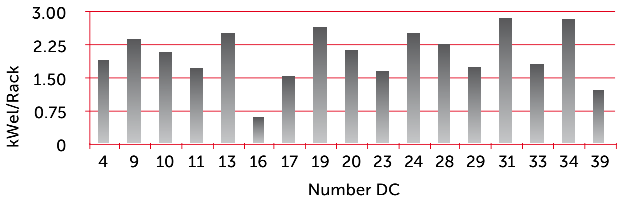

Source: Data Center Benchmarking, TU Berlin, April 2012

The result of the DCB shows that the power density in data centers where predominantly Category 1 racks are installed stands at an average of almost 2.9 kW per rack. Thus, the results of the DCB are on average approx. 1 kW lower than the example calculations (see above).

Rack Category 2:

For data centers with a high level of virtualization (see Rack Category 2), the power consumption of the servers increases considerably. In the DCB, data centers built using predominantly Category 2 racks have 2 out of 3 servers virtualized.

Power Density Rack-Category 2

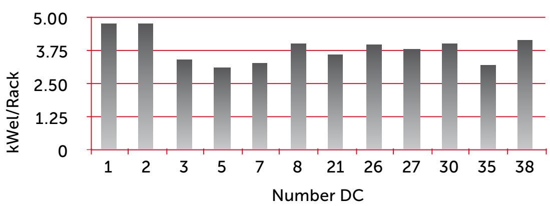

Source: Data Center Benchmarking, TU Berlin, April 2012

The power density, at 3.9 kW, is on average 1 kW per rack higher than the Category 1 racks without virtualization; the results of the DCB are for this category in line with the example calculations. The maximum power density in the DCB is almost 4.8 kW per rack – in this data center, individual servers are operated with up to 17 virtual machines, which in turn suggests a special application/utilization in this data center.

Note: A direct connection between the higher power density (installed IT capacity) and the use of virtualization technology (load on the server) cannot be deduced so far from the results of the DCB. However, the practical experience of the expert group confirm this connection.

Rack Categories 3 and 4:

In the DCB, a range of data centers have been included that predominantly use blade center and mainframe technology. The majority of these data centers are operated in the scientific area, for high-performance computing. Here there are very high loads on the servers, due to the exclusive use of the servers for batch jobs and the use of parallel computer operations. The power densities are therefore comparable between Categories 3 and 4.

Power Density Rack-Category 3 and 4

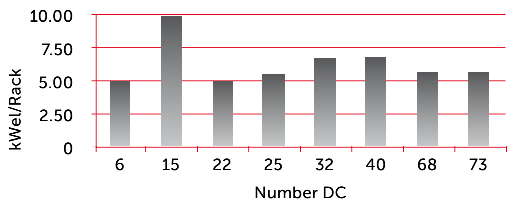

Source: Data Center Benchmarking, TU Berlin, April 2012

In the DCB, the maximum power density of 10kW was measured in a data center involved in HPC, with dedicated mainframe technology. However, the mean power density remains at 6 kW, substantially lower than this maximum consumption, and also below the consumption of 10kW per rack calculated above. The average value for power density here is admittedly less conclusive, as the Categories 3 and 4 were examined together.

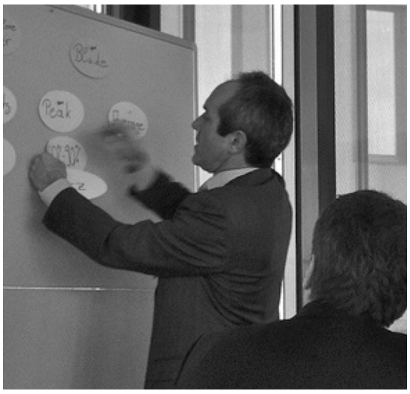

Conclusion

This Whitepaper deals with determining the power density of data centers; this is designed to provide a realistic orientation for the construction and dimensioning of data centers with the necessary technical infrastructure.

This paper is aimed at data center planners, developers and operators who do not operate data centers as a core business, run no specialized applications (gaming, webhosting, etc.) and who work in a heterogeneous server landscape.

The calculated results should serve as an initial approximate benchmark, in order to correctly estimate the required total power capacity. In particular, this should help to avoid over-provisioning. With these results, we hope to motivate readers to seriously consider calculating power density for their own specific requirements.

Within the scope of this study, four categories were crystallized, showing different uses of racks. The actual borders between the individual categories are often fluid, particularly in the context of renovating a data center. It should be remembered that a change from one category to another results as a rule in space-saving (fewer servers) and therefore not in an increase in total load.

The values designated in Table 7 summarize the calculations, analyses and the field reports given above.

Guide to Power Density and Load Determination of Servers, Data Cabinets and Data Centers

Rack Category	Conditions	Power Density per Rack
Category 1 "Normal"	Heterogenous server structure (1 U server, 2 U server, more U server)	< 3 kW
	Up to 70% rack occupancy with ICT components	
	Low degree of virtualization (< 25%)	
	Server load approx. 20%	
Category 2 "Virtual"	See Rack Category 1	< 4 kW
	See Rack Category 2	
	Higher degree of virtualization (> 50%)	
	Server load approx. 60%	
Category 3 "Blade Normal"	Use of blade technology	< 6 kW
	No specialized applications	
	Up to 60% rack occupancy with active components	
	Server load approx. 60%	
Category 4 "Blade with Higher Demands"	Majority blade technology	< 10 kW
	With specialized applications	
	Up to 70% rack occupancy with active components	
	Server load approx. 75%	

Table 7

[1] "Degree of virtualization" means the number of virtual servers in relationship to the number of physical servers in the whole data center.

The process for an initial calculation of power density requirement for a data center should be the classification of the planned racks into one of the categories listed above and calculating the sum of all racks in order to define the total electrical ICT consumption. As a rule, this result will still lie above the actual loads, as the assumptions made in this white paper are oriented toward the upper load limit that was determined for each target group.

The analysis also shows that despite an increase in performance of servers in the last few years, the energy requirements of the new server generation have decreased. From today's perspective, there is no sign of this trend changing. Rather, the indication is that the power requirement of a server in idle mode will continue to sink, which will make the load behavior considerably more dynamic.

Regardless of this, it must be taken into consideration that the processing demand on a server is defined according to the applications which run on it. Very often, much too powerful and "energy-hungry" servers are employed for applications that do not require it.

Overall, it is clear that an examination of servers and their load-dependent consumption should be an essential component of the planning of data centers.

About eco

eco has been the Association of the German Internet Industry for more than 15 years, and represents the industry's interests in politics and international committees.

With more than 750 member organisations, eco is shaping the Internet: eco develops markets, fosters technologies and forms framework conditions. In the Competence Network, eco deals with questions of infrastructure, legal and regulatory issues, innovative applications and the use of content.

A focus of the association is to unify the various perspectives of the industry with those of politics and to foster communication between all market participants. eco organizes many public events, in order to bring those involved together and to create tranparency in the debates.

Detailled information on eco can be found on the association website at

http://international.eco.de.

Your Notes

Your Notes

Association of the German Internet Industry

eco – Association of the German Internet Industry
Lichtstraße 43h, 50825 Cologne, Germany
Phone +49 (0)221/70 00 48 – 0 , Fax +49 (0)221/70 00 48 – 111
info@eco.de, international.eco.de

WE ARE SHAPING THE INTERNET.

eco