

Series on

Environment & Energy



Volume 2

Energy Efficiency in the Data Center

A Guide to the Planning, Modernization
and Operation of Data Centers



■ Publishing Details

Published by:	BITKOM Bundesverband Informationswirtschaft, Telekommunikation und neue Medien e. V. Albrechtstraße 10 A 10117 Berlin-Mitte Tel.: +49.30.27576-0 Fax: +49.30.27576-400 bitkom@bitkom.org www.bitkom.org
Contact:	Holger Skurk Tel.: +49.30.27576-250 h.skurk@bitkom.org
Editor:	Dr. Ralph Hintemann (BITKOM), Dr. Stefanie Pfahl (Federal Ministry for the Environment, Nature Conservation and Nuclear safety) Translation from the German by Caroline Taunt
Assistant Editor:	Christine Faßnacht (BITKOM)
Gestaltung / Layout:	Design Bureau kokliko / Anna Müller-Rosenberger (BITKOM)
Copyright:	BITKOM 2008

The authors of this document have undertaken every effort to ensure the accuracy and correctness of all content herein. However, BITKOM does not assume any legal liability or responsibility for the accuracy, completeness or value of any such information given in this document.



Volume 2

Energy Efficiency in the Data Center

A Guide to the Planning, Modernization
and Operation of Data Centers

Climate protection needs energy efficient data centers

■ A message from Federal Environment Minister Sigmar Gabriel



Sigmar Gabriel, Federal Minister for the Environment, Nature Conservation and Nuclear Safety

In the face of the steady increase in the use of modern information and communication technology and the internet, improving the energy efficiency of our data centers is an important contribution to climate protection. Servers and data centers are among the fastest-growing energy consumers within the sector. Experts estimate that in 2006 the electricity consumption of Germany's roughly 50 000 data centers amounted to 8.67 TWh. That is the equivalent of the annual power output of three medium-sized coal-fired power stations. If no efforts are

made to improve energy efficiency in this area, then data center energy consumption is predicted to rise to 11.1 TWh by 2010. That means that, by 2010, the CO₂ emissions caused by German data centers would rise by approx. 50 per cent as compared with 2001.

The scale of these figures demonstrates the major significance of energy efficient data centers for climate protection.

This is a highly relevant topic for manufacturers and users of IT hardware and software, because dwindling resources mean that cost pressure often directly influences product prices, competitiveness and thus locational and job security, too.

We should make it our common goal to show how these challenges can be transformed into positive synergies for improving resource efficiency, protecting natural resources, lowering costs, generating innovation, saving jobs and creating new ones. New technological solutions and innovations are called for if we are to meet these economic and ecological challenges.

This guide is an important tool in helping us to reach this goal. It points out clearly and concisely the various different practical approaches: cooling strategies for improving data center energy efficiency, for example. It will help you to implement improvement measures quickly and effectively, for the benefit of both the climate and your own pocket.

Content

Welcome	2
Foreword	4
1 Introduction	5
2 Developments in Data Center Energy Requirements	6
2.1 Overview	6
2.2 Energy consumers in the data center	6
2.3 Challenges for the planning and operation of data centers	7
3 Measuring Energy Consumption and Temperatures in the Data Center	8
3.1 Overview	8
3.2 Energy monitoring	9
3.3 Temperature monitoring	9
4 IT Hardware and Software Optimization	10
4.1 Overview	10
4.2 Servers	10
4.3 Storage solutions	12
5 Cooling Optimization	13
5.1 Overview	13
5.2 Air flow in the room and in the rack	13
5.3 Power management	16
5.4 Free cooling	19
5.5 Chilled water temperature	21
5.6 The different types of cooling unit	21
5.7 Different cooling media	22
5.8 Return air temperature	24
6 Optimizing the Energy Supply	25
6.1 Overview	25
6.2 The impact of UPS systems on data center energy consumption	25
6.3 Safe, energy-efficient UPS topology	27
6.4 Redundancy: security and efficiency	28
6.5 Load characteristics, performance category and semiconductor components	30
6.6 Current developments in UPS efficiency	31
6.7 Balancing energy efficiency, security and quality	33
6.8 Rackmount intelligent power strips	33
7 Energy Contracting	36
Endnotes	37
8 Glossary	38
10 Acknowledgements	39

Foreword



Martin Jetter, Member of the BITKOM Executive Committee
General Manager IBM Germany
Chairman of the Board of Directors
IBM Deutschland GmbH

Over the next five years energy consumption will become one of the largest cost factors for many data centers. Borderstep Institute calculations show that data center electricity costs have already more than tripled between 2000 and 2006. Yet investment in energy efficient products and processes often pays off remarkably quickly: payback periods of 2 years and less are the rule here.

However, according to a survey carried out by market researchers Experton, only 7 per cent of German IT decision makers are familiar with the energy requirements of their own IT systems! So it's hardly surprising that data centers are often expensively cooled down

to 18 degrees although they could easily run at up to 26° Celsius. Frequently, it is a cost-accounting problem that lies behind this waste of energy: An organisation's energy consumption often falls under the responsibility of Facility Management. But companies who want to help put Green IT on the map, cut expenditure and save energy whilst protecting the environment at the same time, need to place responsibility for their IT systems' energy consumption where the decisions are made: in the hands of the IT management.

The current climate debate has increasingly brought the subject of Green IT into the public eye. The predicted climate change is threatening our planet and is going to have a dramatic impact on our lifestyles. CO₂ emissions will have to be reduced very rapidly if global warming is to be curbed. While it is true that the ICT sector is responsible for roughly 2% of all CO₂ emissions, it also has to be said that it contributes around 6% of global value added. This sector's energy efficiency is thus three times better than the average for all sectors.

But there is still a lot to be done: For years now, our branch has been addressing the problem of how to reduce global energy consumption even further. Modern environmental protection requires hi-tech solutions. Climate change can only be curbed with the help of innovative ideas and the intelligent application of ICT technologies.

The ICT sector thus embodies the idea of climate protection more than virtually any other sector. But in order to initiate the paradigm shift, we first need to make our own products and solutions more energy efficient. Which is why this guide presents the latest available solutions for reducing costs and increasing energy efficiency.

1 Introduction

Processing power in modern enterprises is continually increasing. Increasing numbers of business processes are being supported by centralized information technology (IT). New improved applications and programme features require high-performance servers. In many cases, comprehensive IT support for business procedures is mandatory if you want to compete successfully in today's global markets. Today, the processing power installed in the data centers of small and medium-sized enterprises equals that which, a few years ago, was only to be found in a handful of major organizations. But this trend - which is essentially a positive one - is also responsible for the high electric power consumption and cooling requirements of today's data centers. This presents new challenges for those involved in the planning, implementation and operation of IT infrastructures.

Nowadays, the application of modern technologies makes it possible to drastically reduce a data center's energy requirements. This offers several advantages:

- The data center's operating costs are considerably reduced. As operating costs account for a high proportion of total costs, any investment in improving a data center's energy efficiency will often start to pay off after only a few months.
- A data center's power and cooling capacity is often limited by the existing infrastructure or by the power provider. In such cases, improving energy efficiency can mean that these limits are never reached, thus obviating the need for additional high investment costs.

- From the point of view of a company's social responsibility, reducing its data center's energy consumption - with the resulting positive effects for the environment - represents an investment in the future.

This guide offers practical advice on the application of modern technologies in the planning, implementation and operation of data centers, in order to considerably improve their energy efficiency¹ whilst simultaneously reducing overall costs: experience has shown that it is possible to achieve savings in a data center's total costs (planning, construction, infrastructure, operation) of 20 per cent or more. These guidelines apply to the planning of new data centers as well as to modernisation projects for existing infrastructures.

2 Developments in Data Center Energy Requirements

■ 2.1 Overview

Information and communication technology (ICT) is an important factor in modern economies: In Germany alone, the Gross Value Added (GVA) of this sector has grown by almost 50 per cent since the mid 90s and now exceeds that of the automobile and mechanical engineering industries. In 2007 the German ICT market had a volume of roughly 148 billion Euros. The sector employs around 800 000 staff and freelance personnel with almost another one million ICT specialists working in other sectors. At the same time, the ICT sector is one of the main driving forces of economic growth. In the applications sectors, investment in ICT leads to increased labour productivity and facilitates significant product and process innovations.

In the past, however, these highly positive developments also lead to a continual rise in ICT energy consumption. In a study² carried out on behalf of the Federal Environmental Ministry, the Borderstep Institute estimates that energy consumption in German data centers more than doubled between 2000 and 2006, rising from 3.98 billion kWh to 8.67 billion kWh. In fact, due to rises in the cost of electric power, data center energy costs have more than trebled, rising from 251 million Euros to 867 million Euros. If this trend continues, electricity costs are likely to double yet again in the next five years.

But the study also reveals that this trend can be reversed through the application of modern technologies, despite the continuing increase in IT performance.

The overall electricity consumption of data centers could be halved within the next five years if state-of-the-art technologies were used. Despite the probable rise in the cost of electric power, electricity expenditure in data centers would still be significantly reduced. (Fig.1).

Electricity Costs
in millions of €




Fig. 1: Development of data center energy costs³

■ 2.2 Energy consumers in the data center

An analysis of Data Center equipment reveals that, on average, only about half the energy consumption is attributable to the IT systems themselves. The additional essential infrastructure such as air conditioning and uninterruptible power supply (UPS) is responsible for the other half. Fig. 2 shows the distribution of energy consumption in the USA between the years 2000 and 2006. On the IT front, the increased energy consumption is mainly attributable to the considerable rise in the number of volume servers. The ever-increasing volume of data means that energy consumption in the storage sector has also risen dramatically. Today it accounts for roughly 15% of the energy demand generated directly by IT systems.

Energy Consumption
in billions kWh/year




Fig. 2: Development of data center energy consumption – relative proportions of individual end-use components⁴

■ 2.3 Challenges for the planning and operation of data centers

The developments we have outlined above present organisations with major challenges when planning and operating their data centers. Power density in data centers has increased dramatically in recent years. The per rack energy consumption alone has tripled over the past three years. It is no longer unusual to find data centers with power densities of 2500 Watts per m² and more. Those responsible for IT and facility management are now facing the following major challenges:

Can the required total power supply be guaranteed?

Electrical power suppliers now sometimes find it impossible to meet the high power demands, particularly in densely populated areas. So what are the options for reducing power consumption?

How can the heat generated in the data center be removed? Are the existing air conditioning systems adequate? How can they be optimally upgraded or modernized to meet today's requirements and thus save additional energy? Can the exhausted heat be utilized to heat the building, for example?

What is the best strategy for dealing with the extreme demands sometimes encountered in individual racks? There are already racks on the market with individual power consumption levels of 30 kW or more. How can hot-spots be avoided or optimally cooled?

Which investment strategy will provide the lowest TCO (Total Cost of Ownership)? It is particularly important that energy costs should be included in the calculation, in addition to the costs for hardware, planning, implementation and management. In some areas, such as volume servers for example, energy costs already equal the cost of the IT equipment itself.

This guide offers you assistance in tackling these questions and challenges.



3 Measuring Energy Consumption and Temperatures in the Data Center

■ 3.1 Overview

If you can't measure it, you can't optimize it. Many data centers have a lot of catching up to do in this department. A large percentage of companies with data centers are unable to specify the size of the IT share in their facility's overall energy consumption. This may be due to the fact that, whilst many IT budgets include costs for planning, purchasing and management of the IT systems, they exclude the energy costs. These are often billed under 'overheads' by the facility management department.

Even the simple act of measuring a data center's total power requirement will often reveal potential for reducing energy consumption and costs. Adequate consideration to the energy consumption factor when equipping a data center could lead to different investment decisions being taken. A slightly higher investment in an energy-efficient cooling system, for example, can start paying off after only a few months.

There are a variety of metrics that can be used to assess the energy efficiency of a data center. Examples are outlined in Table 1. In a top class data center, for example, the ratio of power consumed by the entire facility to that consumed by the IT equipment (Site Infrastructure Energy Efficiency Ratio) could be around 1.5. But there are also data centers in operation with ratios of between two and three, i.e. in addition to the power used by their IT equipment, they are consuming up to twice that amount of power for air-conditioning, UPS etc. With the appropriate software support, it is also possible to continuously calculate and monitor these values in order to uncover further potential for increasing energy efficiency.

Green Grid (USA)⁵ :

- PUE (Power Usage Effectiveness): ratio of a data center's total power consumption to the power being drawn by its IT equipment alone
- DCE (Data Center Efficiency): ratio of power drawn by IT equipment to a data center's total power consumption ($= 1/\text{PUE}$)
- IEP (IT Equipment Power) the power used by the IT equipment to manage, process, store, or route data
- TFP (Total Facility Power) a data center's total energy consumption incl. power for air conditioning (cooling), power delivery components, monitoring, lighting etc.

Uptime Institute (USA)⁶:

- SI-EER (Site Infrastructure Energy Efficiency Ratio) ratio of power consumed by a data center as a whole to that consumed by the IT equipment only
- IT-PEW (IT Productivity per Embedded Watt): value which compares IT productivity (network transactions, storage volume, computing cycles) with the power required for this productivity
- DC-EEP: multiplication of SI-EER and IT-PEW

Table 1: Energy Performance Metrics

But measuring data center energy consumption and temperature can go much further: modern systems make it possible to record and visualize in detail the energy consumption and temperature distribution within a data center.

The latest, flexible techniques allow fast, graphic real-time insights into all levels of energy usage in the data center – right down to individual active or passive components and their application options. Using infrared cameras to record temperatures, for example, will provide a high quality status quo snapshot. However, to ensure that operations management is permanently secure and efficiency-oriented, continuous, comprehensive monitoring is essential, as well as the graphical depiction of historical performance data.

■ 3.2 Energy monitoring

There are two, mutually complementary approaches to monitoring energy consumption in the data center: server monitoring and energy monitoring. Server monitoring records the “production” data and energy monitoring the “supply” data.

Server Monitoring (“Production”)

- low-cost IT monitoring of all relevant load data
- efficient detection of power
- fast and detailed analysis of individual disturbances
- graphic documentation of measures taken (before – after)

Energy Monitoring (“Supply”)

- consolidation of IT and energy management
- integration of additional parameters (meteorological data, etc.)

- common tools for IT, energy, financial management, procurement, etc.
- development of short, middle and long-term energy strategies

■ 3.3 Temperature monitoring

The detailed collection and visualization of temperature distribution data (Figure 3) offers several advantages. Hot-spots can be identified and appropriate measures taken to eliminate them. Temperature distribution within the data center can be optimized by regulating the air volume flow, for example, thus avoiding both unnecessarily low temperatures and dangerously high temperatures. This reduces costs and improves availability. If the temperature distribution is known, any further upgrading of the data center can be realized in an energy-optimized way.




Fig. 3: Temperature distribution in a data center with hot spot (Source: IBM)



4 IT Hardware and Software Optimization

■ 4.1 Overview

The optimization of IT hardware and software is crucial to reducing data center power consumption. Each watt of power saved on the IT side is one watt less which needs cooling or UPS-backup. This translates into a two-fold saving. Detailed advice on IT operations and the appropriate hard and software is beyond the scope of this guide, as its main focus is “data center design and infrastructure”. And so we can only offer a brief overview here of the savings potential which can be realized through IT hard and software.

■ 4.2 Servers

As Figure 2 demonstrates, volume servers account for about 2/3 of the energy consumed by IT hardware. And so it follows that the volume server area harbours great savings potential for reducing overall IT energy consumption. This is augmented by the fact that many of today’s volume servers operate at very low capacities. Average utilization levels of just 10 per cent are not uncommon. But low capacity operation also means poor efficiency levels. Even when in idle mode, a server still requires considerably more than 70 per cent of its maximum performance power.

There are two basic approaches to reducing server energy consumption. One method is to optimize the hardware so that it uses less electric power. The other option is to improve the operation of this hardware so as to increase the average utilization level of the systems. This results in two-fold savings: through lower power consumption and a reduction in hardware. To achieve optimum results, both approaches should be followed simultaneously (see Table 2).

Reduce System Power Consumption

- efficient hardware is highly dependent upon product design
- the right choice of components and software, together with precise dimensioning will reduce energy consumption

Efficient Use of Hardware Resources

- consolidation and virtualization can lead to major reductions in energy and materials consumption
- infrastructure optimization means that hardware can be turned off when not in use

Table 2: Overview of energy saving options for servers

Figure 4 shows a typical distribution of the energy consumption of individual server components. This reveals energy savings potential: small, 2.5” hard disk drives, for example, consume less energy than 3.5” hard drives, hard drives with lower revolution rates use less energy than those with high revolution rates. If hard drive access times and data transfer rates are not critical factors in the planned application, there is often no noticeable difference in system performance.

The choice of appropriate components, right-sized to fit the individual application, has a correlative effect on the other components so that they, too, will have reduced energy/power demands. From an energy efficiency perspective it is preferable to deploy one large memory module than two smaller modules with the same overall capacity. There are now also some highly effective, energy-efficient CPUs, fans and PSUs on the market.

Server air-flow is also an area with high energy savings potential. Servers should be designed to draw air in across their entire front surface to ensure optimum convection cooling. Increasing the size of the air intake vent of each CPU socket means that the ventilation system can operate at higher supply air temperatures or expend less




Fig. 4: Typical breakdown of a server's energy consumption (Source: Fujitsu Siemens Computers)

energy for cooling. With the blade server form factor in particular, it is important that the design of both ventilation system and power supply should be carefully coordinated so that space is saved in the right places. Inappropriate reduction of the size of the enclosure which houses the CPU and RAM will lead to increased heat density, will generate hotspots and require more energy to transport the same volume of air through a smaller cross-sectional area.

Considerable energy savings can thus be achieved by hardware optimization alone, and the potential is often

even greater when systems are used more efficiently: consolidation and virtualization are the keywords here.

The term consolidation describes the process of harmonizing and amalgamating systems, applications, databases or strategies. The aim here is usually to simplify the infrastructure and render it more flexible. The process usually leads to a considerable reduction in energy consumption.

Virtualization means abstraction: logical systems are abstracted from actual physical hardware. In this process resources are not dedicated but are used collectively, which gives greater flexibility as regards availability and means that capacities are better exploited. This can increase system utilization rates considerably and thus save a lot of energy at the same time.

Figure 5 shows an example of server consolidation and virtualization. Overall energy consumption can be drastically reduced – by 50 per cent in this example – whilst still retaining the same levels of system performance and availability, by moving 4 system workloads onto one high-performance system with professional virtualization. This is one of the more conservative examples – appreciably higher savings are possible, depending on the application in question.

Server consolidation saves energy

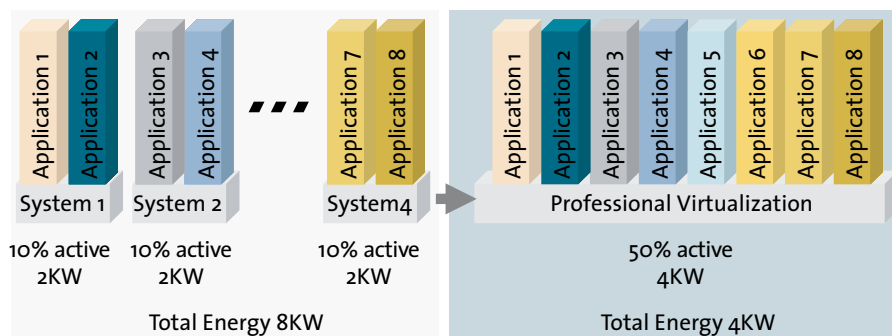


Fig. 5: Energy savings through consolidation and virtualization

Solutions which automatically shut down servers and then restart them also offer great energy savings potential. Many applications make only partial use of servers during a given period – only during office hours, for example, or only Monday to Friday. Servers can be systematically powered down during periods where the computational load is known to be only minimal. Today there are already solutions available for achieving this, especially in conjunction with virtualization technology.

Time management of server activities is another area which harbours opportunities for economizing on energy and costs. Energy demands can temporarily peak when simultaneous server activities overlap, as standard settings are not changed (e.g. certain services are executed on the hour, every hour). Temporal coordination of server loads results in balanced energy demands.

■ 4.3 Storage solutions

Data storage constitutes a relatively small proportion of the overall energy consumption in data centers. Nevertheless, it is worth considering energy efficiency optimization measures in this department, too, for data volume is still growing rapidly, and with it, the storage solution proportion of overall IT energy consumption (see Fig. 2).

One apparently simple method of reducing data storage-related energy consumption is the optimization of data management: unnecessary and obsolete data should be deleted.

In many organizations multimedia files such as mp3 or video files now constitute the majority of all data – even if they are never required for business operations. In addition to this, files are often stored in multiple locations. Precise file-handling procedures supported by high-performance software can drastically reduce the amount of data being stored.

In this context, it is also worth considering the concept of information lifecycle management (ILM). ILM is a storage management concept involving the active management

of information objects throughout their useful life. This is an optimization process, by which a rules engine based on business policies and the assessment of storage hierarchy cost structures, determines the most appropriate storage location for the managed information objects. Only those information objects requiring high availability are stored on expensive, high-performance disks with high energy consumption. ILM helps to save energy because it always selects the optimum – and thus the most energy efficient – storage solution and automatically deletes the information objects at the end of their life-cycle.

Another energy-saving data storage approach is the use of tape systems to archive data. Unlike harddisks, tapes consume no energy until the data is accessed. Energy-optimized storage systems, drives with high storage density and other modern technologies such as de-duplication and automatic powering down of idle hard disk drives can all help to minimize energy consumption.

Table 3 shows the recommended sequence of steps for reducing data storage energy consumption.

1. Optimize data management
■ Delete obsolete data
■ Delete unnecessary data
2. Optimize infrastructure
■ Consolidation
■ Use ILM
■ Use tape systems
3. Optimize devices
■ Apply state of the art technology

Table 3: Recommended steps for reducing energy consumption in data storage environments

5 Cooling Optimization

■ 5.1 Overview

Cooling plays a key role in the design, construction and operation of a data center. This is partially due to the fact that cooling accounts for a large proportion of the energy costs. Depending on local conditions and cooling system design, this proportion generally constitutes at least 20% of total energy costs, but this must be seen as the lower limit. There are also installations in operation where cooling accounts for over 60% of overall power consumption. Designing a cooling system, on the other hand, presents a major challenge, as it usually constitutes a mid to long-term investment which will be in operation for several IT generations. In the following, we point out the various approaches to optimizing data center cooling systems which will give you the lowest possible energy consumption.

■ 5.2 Air flow in the room and in the rack

5.2.1 Hot aisle-cold aisle layout

Traditional data centers use air as a cooling medium in both racks and computer rooms. The heat loads to be dissipated are steadily rising.

The most commonly-used and effective strategy for cooling computer rooms with air is a layout of alternating hot and cold aisles (see Fig. 6). The cooling air is fed under a raised floor and distributed through strategically positioned perforated floor tiles and exhausted via a drop ceiling return air plenum.

The majority of servers produced today draw in the conditioned supply air at the front and expel it again at the back. And so it makes sense to arrange the server racks with perforated front and back doors in such a way as to create hot aisles and cold aisles.

This approach entails positioning the racks face to face, as recommended in VDI 2054 or by ASHRAE Technical Committee 9.9. In this setup, the cooled supply air is distributed into the cold aisle from where it is drawn in by the servers on both sides and then exhausted out of the back of the systems to the hot aisle.




Figure 6: Hot Aisle/Cold Aisle Configuration




Figure 7: Example of air flow in a data center with hot aisle/cold aisle layout and adverse bypassing and recirculation of air

The temperature spread (the difference between return air and supply air) is limited by the maximum permitted exhaust air temperatures of the IT components. In order to meet ever increasing cooling loads, either the intake

air temperature is drastically reduced or the volume of cooling air supplied to the servers by computer room air conditioning (CRAC) units has to be continually increased. This produces extremely high airflow velocities in the room.

When this type of thermal management meets its cooling performance limits, the high air velocities have adverse effects such as air bypassing or recirculation of cooling air (see Fig. 7). Avoiding these phenomena can dramatically reduce operating costs.

The bottleneck in the supply route of conditioned air to the racks is located primarily in the raised floor. The cooling air for the entire room has to be fed through here, despite obstructions to airflow such as cabling or piping for liquid refrigerants. The clearance heights of today's raised floors typically begin at around 300 mm and can reach heights of 2 or 3 metres or more in high-performance data centers.

A hot aisle/cold aisle layout provides the opportunity for improved cable management, both inside the racks and in the underfloor plenum. As far as possible, cabling should be routed through the hot aisles only, to ensure a free flow of air in the cold aisles.

The perforated floor tiles should be installed in the cold aisles only and the CRAC units positioned at the ends of the hot aisles - NOT aligned with the rack rows. A setup where CRAC units are positioned parallel to the rack rows causes a mixing of hot and cold air, which leads to insufficient cooling of the IT equipment in the upper rack sections. And what is more, the lower air intake temperatures of the CRAC units will considerably reduce the overall efficiency of the system.

If planned and implemented with care, a hot aisle/cold aisle layout can dissipate heat loads of up to approximately 5kW per rack. These figures are usually considerably lower in older data centers: between approximately 1 and 2 kW per rack. If the setup is well designed, the cold aisles should be completely filled with cool air, whilst the

warm air passes over the cabinets where it collects and is returned to the CRAC units.

Various different manufacturers offer a range of systems which support or improve cooling efficiency by providing additional cooling capacity in the cold aisles, but not through the underfloor plenum.

5.2.2 Cold aisle and hot aisle containment

The advantages of the hot aisle/cold aisle configuration can be further enhanced by the containment of either the hot or the cold aisle. This prevents the occurrence of hot air short circuits or the mixing of hot exhaust air with cool supply air. Thermal management must always be adjusted to the number of servers.




Fig. 8: Cold Aisle Containment

Care should be taken to ensure that the containment area is fully enclosed. In the case of cold aisle containment, this means sealing off the aisle from above with roof components and to the front and rear with doors, separating hot and cold areas inside the cabinets (unused rack spaces and vertical spaces to the sides of the 19" zone), as well as sealing off the floor in the warm zone (inside the cabinets and in the hot aisles).

This arrangement makes it possible to reduce the vertical temperature gradient from the surface of the raised floor to the tops of the racks from approx. 4 Kelvin to approx. 1 Kelvin. In this way the servers are guaranteed an optimum supply of cool air across the entire air intake surface of the racks.

If the cold or hot aisles are enclosed, CRAC units can be set up parallel to the rack rows (see Fig. 9), as the mixing of warm exhaust air and cold intake air is efficiently prevented.




Figure 9: Positioning of CRAC-units parallel to the racks with cold or hot aisle containment

In a cold aisle containment system where longer runs or multiple rack rows need to be cooled by CRAC units in parallel configuration, it is advisable to carry out precision warm air extraction through a suspended ceiling plenum or ducting system. Hot aisle containment is a further efficiency boosting strategy. But whichever air routing method is adopted, it must always be adapted to the requirements of the individual room situation.

5.2.3 Minimizing obstructions in the underfloor plenum

A raised floor installation is one of the best technical options available for providing a data center with the necessary power and data cabling as well as its vital supply of cooling air. But even the best raised floor will not provide satisfactory results if certain physical laws are not observed.

A standard raised floor is normally installed above the actual base floor (screed, concrete, industrial floor etc.) on a pedestal and stringer understructure. The aim here is to create space for cold air distribution, wiring and cable trays.

The following issues should be addressed to ensure that a raised floor functions as effectively as possible:

- inadequate raised floor height
- incorrect dimensioning of perforated tiles and grilles.
- air leakages: there are gaps at the tile intersections and where the floor meets perimeter wall elements which must be permanently sealed. There should be a minimum air pressure of at least 25Pa in the underfloor plenum to ensure uniform air distribution through the perforated panels.
- constrictions of cross sectional area inside the under-floor plenum
- Cable trays installed across the airflow path

One of the most common thermal management problems is the reduction in cooling capacity which occurs when the underfloor plenum is misused as storage space. It is not unusual to find considerable amounts of cabling and ductwork in the underfloor plenum which impede air circulation and adversely affect the air pressure conditions required for server cooling.

As a result, the air conditioning system has to pump more cold air into the data center in order to achieve the required cooling level. But more cold air means higher air conditioning costs and besides costs there are also security issues to consider – the prevention of so-called hot

spots in the racks for example, and the avoidance of fault sources when altering cable routing.

Ideally, as much as possible of the cross-sectional area in the underfloor plenum should be made available to the air volume flow. Unnecessary obstacles in the underfloor plenum need to be eliminated: surplus cables should be removed and essential cabling laid in such a way that no high crossings occur.

Special cable seals with overlapping brushes are recommended in order to seal openings for cable assemblies and wiring as well as possible. They will guarantee a tight seal, even with very large volumes of cable.

It is advisable to dimension the height of the raised floor adequately, right from the planning stage, to ensure that there will be no bottlenecks if changes need to be made later on, to IT or cabling, for example.

5.2.4 Rack depth extension

Rack depth is a key specification when equipping a data center. In standard configurations of alternating hot and cold aisles rack depth is a prime determinant for the spacing of axes within the data center.

Where heterogeneous server systems are operating within the same rack - which is more likely to be the rule than the exception - IT administrators must opt for a depth-variable solution. A homogenous layout is desirable for energy efficiency reasons.

Up until a few years ago the standard depth for server racks was 1000mm. In recent years racks with depths of 1070mm, 1100mm or 1200 have also come onto the market.

This can be attributed, on the one hand, to the introduction of servers with a rack mount depth in excess of the standard 740mm, and on the other, to a demand for more space behind the installations to accommodate cabling and PDUs.

Optimal and energy-efficient cooling demands intelligent cable management and the deployment of the latest products in the rack. Cabling must not be allowed to impede air flow.

In enclosed cabinet cooling solutions and server cabinets with very high heat loads, cabinets may well have to be even deeper, as additional space is required in front of and behind the 19" installations for the transport of the very high volumes of cooling air (up to 6000 cbm/h in one cabinet). For these applications there are cabinet depths of 1300mm, 1400mm and 1500 available on the market.

■ 5.3 Power management

5.3.1 Power management as an approach to improving energy efficiency

Power management is absolutely essential for an energy-efficient, energy-optimized data center operation.

There are various different HVAC systems available for transporting heat away from the chip and out of the building. They use gaseous or liquid media to do so and include such devices as air circulation units, cold water units, chillers, heat exchangers and condensers. All these systems need to be considered individually and then optimally coordinated.

5.3.2 Power control in the air cycle

As with any technical equipment, precision coolers, too, have many different components where energy savings can be made. The choice of fans, for example, offers great savings potential, for they run 24 hours a day and thus 8760 hours a year. It is this consideration which has led to the rapid breakthrough of EC (electronically commutated) fans. The power consumption of these speed-controlled fans drops drastically when air volume flow is reduced.

When an air handling unit is redundantly configured, major energy savings can be achieved by running all units simultaneously at correspondingly reduced speeds and only when a unit fails, switching the others to rated speed.

The components and systems must be operated as energy-efficiently as possible. This can best be achieved when HVAC units and systems are dimensioned to meet maximum load requirements and can be automatically adjusted to cope with prevailing heat loads. Optimum, energy-efficient power control must encompass both the ventilation and water/refrigerant cycles.

Initially, power control optimizes HVAC technology for static operating conditions. But data center heat loads vary over time, according to installation workloads, with slumps occurring particularly at night and at weekends. In the case of fluctuating loads, it is thus the task of control systems to dynamically adjust the HVAC equipment to meet the varying heat loads.

The trend toward server virtualisation in particular, opens up a whole new set of opportunities for increased efficiency. Workloads can be rapidly migrated to other physical servers and at off-peak times whole cabinets or even rows of cabinets can be cleared of applications and shut down.

HVAC infrastructure should be adjusted accordingly, by shutting down individual CRAC units for example, or throttling back several of these units. Care should be taken to ensure that the minimum air volume of the HVAC units can support the required supply pressure of at least 25Pa in the underfloor plenum.

And so the appropriate deployment and configuration of IT components within the data center plays a vital role in the energy efficient application of cooling technology.

By using intake and exhaust air temperature sensors to systematically monitor power in the racks, it is possible to directly influence both the power consumption and the air volume flow of the cooled cabinets. This process yields

potential HVAC energy savings of between 30 and 60 per cent.

Temperature measurements or electrical energy consumption can both serve as reference variables for power control, as all the electricity is converted into heat after a slight delay. However, this is only valid for data centers where there are no other major energy inputs such as solar radiation, for example. From an energy-efficiency perspective it is therefore preferable to use interior rooms as data center space.

The ASHRAE Technical Committee 9.9 has set an internationally recognized standard for air intake temperatures of electronic equipment in data centers.⁷ According to this standard, the recommended intake temperature for electronic components in the data center lies between 20° and 25°C, with an additional “allowable” range of 15° to 32°C. These recommendations dating from 2004 are currently being revised. It is safe to assume that ASHRAE will raise their recommended temperature levels. Outlet air temperatures and permitted humidity ranges are prescribed by the manufacturers of servers, mass storage devices, switches and other equipment. For reasons of energy efficiency it is advisable to attain the highest possible temperature level, as this directly affects the cooling system’s energy consumption (see section 5.8).

Under ASHRAE guidelines, conventional servers permit a temperature spread of about 15K and blade servers about 28K (between intake and outlet). By fully exploiting this spread it is possible to minimize the volume of cooling air and the energy required for circulating it in the room. When planning a data center, these factors should be verified in advance and incorporated into the design.

5.3.3 Power control in the water cooling/refrigerant cycle

In a water cooling/refrigerant cycle, the collective heat load is conveyed out of the building and usually dissipated into the surrounding air.

Temperature levels in both the cooling cycle and the environment into which the heat load is dissipated, as well as the temperature difference between supply and return piping, are decisive factors for a cooling system's energy efficiency. The higher level is dictated by the temperature of the surrounding environment (ambient temperature), the lower level by the transfer temperature from the air cycle.

The ambient temperature depends on the data center's geographical location as well as on the time of day and season. And the installation location of the heat exchangers is another factor which shouldn't be underrated. Direct sunlight should be avoided at all costs, as this would impair the performance and efficiency of the cooling system.

And so various energy-optimized control strategies for this cycle present themselves.

Firstly, if ambient temperatures are relatively low, a control system can be used to regulate the discharge of waste heat to the environment. In this case, the heat must be temporarily stored. As data centers produce large amounts of heat which cannot be stored indefinitely, the prime option here is to utilize daily temperature fluctuations. The required buffer tank must be adequately dimensioned so that it can store the waste heat during the day and dissipate it to the surrounding air at night.

Secondly, the heat occurring in components can be tapped at the highest possible temperature. The higher the temperature tapped at the ICT equipment, the more efficiently can the HVAC system, particularly free cooling components, be operated. The optimum solution here would be "chip cooling", whereby the heat is removed by means of a liquid heat transfer agent at levels roughly equivalent to the component operating temperatures. This type of chip cooling is already being used in certain special applications, but the technology has not yet been introduced on a wide scale. Before this can happen, servers will have to undergo considerable constructional changes. Furthermore, these solutions would be entirely

determined by the particular CPU and server models and so could not be standardized.

Thirdly, as in the air cycle, power control can also be used to regulate mass flow in the liquid cycle in such a way that the optimal temperature spread is exploited.

5.3.4 Dynamic power control and equipment maintenance

Above and beyond the application of optimized components and systems, the dynamic control of these systems will improve energy efficiency even further.

Whereas the operating conditions at the systems' set point (max. heat load) are relatively constant, power control techniques can change various parameters when the heat load drops.

Power control can throttle back or shut down the HVAC systems to meet the prevailing total heat load.

The following strategies can be particularly recommended here:

- The volume of air in the air cycle can be minimized using speed-controlled fans. (CAUTION: Minimum pressure of 25Pa in the underfloor plenum must be sustained!)
- Subject to the prevailing operating conditions, the feed temperature in the fluid cycle can be set as high as possible.
- By using dynamic control to regulate the free cooling switch point, it is possible to achieve longer periods of cooling without the use of chiller compressors.

This will optimize the main components of the HVAC system's energy consumption:

- The operating energy required by the fans for air circulation is reduced.
- The drive capacities of the chillers' refrigeration units drop because their efficiency increases as the difference between outlet air temperature and feed temperature decreases.

- The periods during which active cooling can be replaced by free cooling are considerably lengthened.

The routine maintenance of all equipment is another important practice for ensuring energy efficient data center operations, and something which is also required by law. Dirty filters in particular cause high energy consumption by increasing the load on fans or on pumps in the water cycle. Power control can only provide optimal utilization if the equipment is properly serviced and maintained.




Fig. 10: Air conditioning units with direct free cooling

■ 5.4 Free cooling

5.4.1 Free cooling as an approach to increasing energy efficiency

Free cooling is particularly suitable for the air conditioning of data centers located in cold and temperate climates. It is important to differentiate between direct free cooling and indirect free cooling. Air conditioning systems for free cooling have been built for over 30 years. In recent years, the design of the systems has been refined in line with technological progress (microprocessor control, speed-controlled components etc.) and this has opened up even greater potential for energy savings.

5.4.2 Direct free cooling

In addition to a cooling cycle, the air conditioners are also equipped with a system of dampers. The damper control system allows outside air to enter the room and indoor air to be transferred back outside or can be switched to a recirculation mode, i.e. the air is circulated within the room being cooled and is cooled via the cooling cycle.

The type of air handling system used depends on external climatic conditions and the required room conditions. Here one has to differentiate between free cooling mode, mixed mode operation (free cooling with additional cooling provided by the cooling cycle) and recirculation mode (cooling via cooling cycle alone).

In direct free cooling, cold outside air is transported into the room. The heat load from the IT equipment is transferred to the air stream and this transports it out of the room. The air current is fed or circulated through an external air filter and a fan unit.

This mode of operation offers considerable energy-savings potential in cold and temperate climates in particular, thanks to annual temperature patterns. If outside ambient temperatures rise, the air handling control system switches to mixed mode so that additional heat can be removed via the cooling cycle.

In mixed mode operation the cooling cycle is initially only operational for short periods when outside temperatures are low. When outside temperatures rise, the cooling cycle runs more frequently until the air handling control system finally shuts off the intake of outside air entirely by means of a system of dampers, and the air handling unit continues in recirculation mode. At high outside ambient temperatures the cooling cycle alone is responsible for extracting the heat from the ICT equipment via the

circulating air stream. In this mode the system also has the highest energy consumption.

Direct free cooling is particularly suited for application in smaller facilities such as telecom containers/shelters. In view of the fact that larger scale facilities with greater heat loads require proportionately larger volumes of outside air to be transported into the room through the building envelope, and that room humidity has to be controlled within narrow tolerances in data center air conditioning, direct free cooling is not normally an option for such installations.

In addition, when operating a system with direct free cooling, the air filters require particular attention. The high volumes of outside air shorten filter lifetimes drastically, especially in locations with high dust levels. They also reduce the efficiency of the air conditioning system and increase operating costs.

5.4.3 Indirect free cooling

With indirect free cooling, unlike the direct version, outside air is not used directly to cool the ICT equipment. All the indirect free cooling systems outlined below work on the principle that, in free cooling mode, a water/glycol solution is cooled by the outside air without using a cooling cycle and can then be used to cool the circulating air stream inside the facility. The interpolated water/glycol solution functions as the coolant and is circulated through a piping system by pumps.

Indirect free cooling offers major energy-savings potential for data center air conditioning. Although the additional components required for indirect free cooling incur higher investment costs for the air conditioning system, this additional outlay is compensated in the medium term by dramatically reduced operating costs.

The savings opportunities offered by indirect free cooling are particularly attractive in cold and temperate climates. Low outside air temperatures make it possible to reduce chiller operation time which, in turn, leads to a reduction

in the energy consumption of the air conditioning system. If outside air temperatures rise, the cooling cycle has to be run more frequently. Depending on the system's design, it automatically shuts down the free cooling function at high outside temperatures and cooling is then provided exclusively via the cooling cycle. This mode of operation is only necessary when outside temperatures are high and then it requires the highest energy input.

The design and construction of indirect free cooling systems varies, sometimes considerably. The systems can be categorized in three basic groups:

- air conditioning units with integrated free cooling function and recooling unit (small and mid-size rooms),
- water cooled units with central chilled water generation and integrated free cooling function in the chillers (mid-size and large rooms),
- water cooled units with central chilled water generation and external free cooling function via recooling unit (large rooms).

5.4.4 Other energy-saving options: free cooling system design

The system design must be geared to the annual temperature pattern of the particular location. All the system components must be carefully configured so that the system can provide the greatest possible energy savings once in operation.

To achieve this it is important to consider both room temperature and room humidity to ensure that the system as a whole can make optimal use of free cooling.

Room temperature

Within a conceptual design framework, room temperature is a crucial factor influencing both the configuration of the system and energy efficiency. Higher room temperatures can be maintained over longer operating periods by indirect free cooling, without the need for additional refrigeration, thus contributing directly to energy savings.

Room humidity

Like room temperature, room humidity also has an effect on energy efficiency. Low relative humidity in combination with a higher room temperature will allow air conditioning units to be run more efficiently. When operated under these conditions, the air conditioning units are able to provide sensitive cooling, i.e. there is no dehumidification.

Excessive relative room humidity combined with a lower room temperature will automatically cause the air conditioning units to produce latent cooling which will inevitably cause dehumidification of the air stream. Extra power is then required to replenish humidity by means of a humidifier to ensure that it does not sink below acceptable levels.

Scalable cooling

Modern air conditioning systems should also be able to react flexibly to partial load conditions within the facility and provide demand oriented cooling using the lowest possible energy input. This can be effected by raising the chilled water temperature, as this will maximise the free cooling operating period still further. Scalability is particularly relevant when a facility is being upgraded in phases or where there are fluctuating heat loads.

An economic feasibility study must be carried out to determine which system is the most suitable. The installation situation, loads and acoustic conditions must also be taken into account here, so that a viable solution is eventually arrived at.

■ 5.5 Chilled water temperature

In large-scale data centers around 90 per cent of the thermal energy is transported in a water/glycol solution (brine). In free cooling systems, brine as a heat transfer medium has a great advantage over direct evaporation systems. When ambient outside air temperatures are low enough, the cold air can be utilized to cool the data center directly.

To fully exploit this effect and reduce the proportion of latent cooling to as near zero as possible, the temperature of the brine in the air conditioning units must be as high as possible, approx. 12°C supply temperature and approx. 18° return temperature. Higher temperatures are even more advantageous from an energy efficiency perspective.

The average room air temperature in data centers corresponds directly with the brine supply temperature. Experience has shown that the energy consumption of an air conditioning system can be reduced by about 15% at an average room temperature of 26°C as opposed to 22°C. The higher temperature of the room air permits higher chilled water supply temperatures, thus increasing the number of free cooling operating hours considerably and substantially raising performance levels (EER).

■ 5.6 The different types of cooling unit

Suitable hardware is required before the above-mentioned cooling strategies can be put into practice. There are a wide range of manufacturers on the market offering air conditioning units of all shapes and sizes. To enable the right choice to be made from an energy point of view, there now follows an examination of a wide variety of different systems.

The types of unit available can basically be divided into comfort and precision air conditioners. In both cases it is possible to use chilled water (CW) and cooling agents (DX = direct expansion) as mediums for removing the thermal energy extracted from the room air. In addition there are also the water-cooled racks and several special, manufacturer-specific solutions.

1. The comfort devices include fan-coil-units (= fan convectors) which have an air-water heat exchanger with integrated blower, and the split unit air conditioners which use a cooling medium and are more commonly used in smaller spaces. But these types of unit have several disadvantages which make them unsuitable for application in computer rooms: The

volume of circulated air is very small, the chilled air is blown in from above, which is highly disadvantageous from a thermodynamic perspective, the split units cause excessive dehumidification, the control systems are only very basic and do not take room humidity into account. What is more, the split units operate inefficiently as far as sensible cooling is concerned⁸. This means that they are neither energy efficient nor suitable for year-round operation.

2. Precision air conditioners were originally developed for use in computer rooms and can be divided into two groups, according to the air-flow routing: down-flow (= air intake from the top, discharge through the base) and upflow (= air intake from bottom front or rear, discharge through the top). These units have microprocessor controls which constantly monitor both temperature and relative humidity and are programmed for energy optimized operation. Precision air conditioners deliver virtually 100 per cent sensible cooling and efficiently apply the energy used for cooling production to reduce room temperature levels. Another strategic point in the design of energy-efficient precision air conditioners is the dimensioning of the heat exchanger surface area⁹. The greater this is, the lower the dehumidification capacity. This can be discerned from the observation of a minimal difference between sensible and overall cooling performance. This difference, also known as latent cooling, extracts humidity from the return air and consumes energy. However, as the relative humidity in data centers needs to be kept within the tolerance limits prescribed by the hardware manufacturers to avoid problems with electrostatic charging, more energy has to be expended to replenish the room air with the extracted humidity.
3. The performance of computer hardware continues to soar, whilst dimensions steadily decrease and developments such as the (blade) server in particular, mean that today's HVAC technology faces the additional challenge of dealing with heat loads of 30 kW and more per rack. The answer is a rack with a built-in

air-water heat exchanger and load-dependent, speed controlled fans. This ensures adequate cooling for all the computers in the rack and it also means that only the computers are cooled, not any other loads such as transmission, lighting etc.

4. In addition to the standard cooling units already described above, there are also some manufacturer-specific solutions which cannot be elaborated upon here (rack-mounted cooling modules which enhance cold-aisle cooling, for example). These solutions certainly have their justifications when used in the specific applications for which they were conceived, but a detailed examination of their energy efficiency characteristics is beyond the scope of this guide.

■ 5.7 Different cooling media

The cooling systems for IT environments can be differentiated according to both the type of unit and the cooling medium it employs. Every type of cooling is based on the physical law which states that energy can never be destroyed, only transported. In computer environments air is used as a primary energy carrier because of its superior distribution qualities. Water, however, can store heat around 3,500 times better than air can, and is thus the better choice of medium for transporting the extracted heat out of the room. Depending on the type of application and local conditions, it may be feasible to use cooling media other than water.

Water as a cooling medium

Water lends itself to use as a cooling medium primarily on account of its easy controllability, ready availability and very moderate expense. In addition, at low outside temperatures it is easy and cost-effective to run a water-based cooling operation which exploits free cooling advantages and reduces energy consumption.

The disadvantages of water within the data center are no longer relevant thanks to the technological progress of the past decades:

- Modern pump systems are infinitely adjustable, depending on the required volume of water and temperature difference and when the pumping capacity is throttled back, energy consumption is automatically reduced as well.
- Through the addition of an antifreeze agent ("glycol"), the freezing point can be lowered to far below 0°C. Nowadays, these agents are also corrosion inhibitors, which makes them doubly useful.
- Today, compared with other cooling mediums, water is much the safer coolant for data center application. This can be accredited to its leak detection potential and/or the fact that the water system can be sealed against electrical systems through the use of conduits with systematic drainage.

Refrigerants as cooling media

Refrigerants are often synthetically produced, chemical substances or compounds. Their low boiling points are taken advantage of in chillers. Of course, natural substances such as water¹⁰, ammonia or carbon dioxide are also used as refrigerants. Ammonia, in particular, is one of the most important refrigerants around, but safety requirements do not permit its use in data centers.

The refrigerant market has undergone a major transformation over the past 10-15 years due to ecological factors such as the destruction of the ozone layer and the greenhouse effect through the emissions of synthetic substances. Air conditioning technology is therefore also attempting to reduce the amount of refrigerant used in systems to an absolute minimum or to switch to natural refrigerants. Here are the most common synthetic refrigerants, all of which belong to the group of alkyl halides: at present, the refrigerant R22 is still in use in many installations. This substance can still be sold for service or maintenance applications up to 31.12.2009, and following that, until 31.12.2014 servicing may still be carried out, but using the recycled product only. However, it is foreseeable that by 2010, the last systems using R22 will have to be decommissioned or at least converted to a different refrigerant with zero ozone depletion potential.

R407C is the most widely used replacement product for R22. It is a blend of three individual components: 23 per cent R32 (difluoromethane CH_2F_2), 25 per cent R125 (pentafluoroethane CHF_2CF_3), 52 per cent R134A (1,1,1,2 tetrafluoroethane CH_2FCF_3). It is relatively easy to replace R22 with R407C. R407C has zero ozone depleting potential, but it does have high global warming potential.

The operational disadvantages of R407C prompted the chemical industry to develop a substitute in the form of R410A. This is also a mixture and consists of 50 per cent R32 (difluoromethane CH_2F_2) and 50 per cent R125 (pentafluoroethane CHF_2CF_3). It is not possible to substitute it for R22 in existing installations however, as cooling cycle pressure levels for R407C are considerably higher. Its global warming potential is slightly lower than that of R407C.

The pure substance R134A is a well-known and widely used refrigerant. However, as its volumetric cooling capacity only amounts to approx. 50 – 70 per cent of that of R22, depending on operating conditions, systems of a comparable size will be correspondingly less efficient. The environmental impact in respect of greenhouse gas emissions is slightly less than that of either R407C or R410A.


Carbon dioxide as a cooling medium

At this point we should like to briefly discuss the use of carbon dioxide (R744). Even if, initially, this should seem somewhat unusual, carbon dioxide does, in fact, constitute an energy efficient alternative cooling medium.

Compared to water, carbon dioxide has a far superior storage capacity, which means that pipe diameters can be smaller and less energy is required for the transportation of the medium from manufacturer to end-user.

The present high investment costs are a disadvantage, which means the larger the installation capacity, the more economic the application.

A further drawback to the use of carbon dioxide as a cooling medium is the need for high operating pressure



levels. But the appropriate air-conditioning components for high-pressure systems are already available, so that carbon dioxide can also be used in the data center.

■ 5.8 Return air temperature

Return air temperature¹¹ plays a crucial role in the energy consumption of air-cooled rooms. The principle here is: the higher the return air temperature, the more efficiently the system will operate.¹² Even so, allowances must be made for the data center's IT equipment to ensure that no damage occurs and the equipment's operating life is not shortened.

A rise in return air temperature automatically involves a rise in the temperature of both supply air and room air.

Research carried out by the Swiss Federal Office of Energy has revealed that between 22°C and 26°C, each degree of temperature increase will result in a 4% energy saving.¹³

Data center airflow must be optimally designed and implemented if the supply air temperature is to be raised. The hot aisle/cold aisle principle, for example, must be strictly adhered to, so that there are no air turbulences which a) prevent the cooling air from reaching its intended destinations and b) cause a drop in the return air temperature.

Otherwise hot air can recirculate into the cold aisle, causing overheating and hot spots. This usually leads to typical countermeasures such as lowering the supply air temperature and/or increasing the air volume flow rate – but such actions are not energy efficient.

6 Optimizing the Energy Supply

■ 6.1 Overview

As far as energy supply optimization is concerned, the uninterruptible power supply (UPS) plays a vital role in reducing the energy consumption of data centers. The UPS is an integral part of a data center's power supply infrastructure. Because of its back-up character, a UPS will not provide any "productive" power conversion during normal operations. As well as efficiency, the question of adequate energy management security is of particular importance for UPS systems. Paragraphs 6.2 to 6.7 deal with the influence of the UPS on a data center's energy consumption. They also highlight various different concepts and specifically address the conflicting objectives of security and energy efficiency.

Paragraph 6.8 presents another approach to power supply optimization: the use of intelligent power strips. These can be remotely controlled via IP networks and allow integrated power monitoring.

■ 6.2 The impact of UPS systems on data center energy consumption

6.2.1 Energy efficient power management with UPS systems

If one sees the UPS as a kind of "insurance", then the electrical energy expended during UPS operation can be compared to the payment of different insurance premiums for different benefits – as in third party liability, third party fire and theft and fully comprehensive cover in the field of automobile insurance. The insured party would rather he never had to claim on his insurance, but on the other hand, he doesn't want to have to face the consequences if the worst comes to the worst.

So the following section will deal with the correlation between efficiency and security. An energy management policy which is geared solely towards saving electricity is definitely falling short of the mark, particularly in the case of UPS.

	UPS Type 1	UPS Type 2	UPS Type 3	UPS Type 4	UPS Type 5
Real power of the load in kilowatts (kW)	100	100	100	100	100
Efficiency of UPS in %	96.0%	94.0%	92.0%	90.0%	88.0%
Input power to UPS in kW	104.2	106.4	108.7	111.1	113.6
Electrical losses from UPS in kW	4.2	6.4	8.7	11.1	13.6
Kilowatt-hours (kWh) per year	36 500	55 915	76 174	97 333	119 455
Power costs in EUR per kWh	0.10€	0.10€	0.10€	0.10€	0.10€
Additional costs due to electrical losses per year (without cooling)	3 650.00€	5 591.49€	7 617.39€	9 733.33€	11 945.45€
CoP (Coefficient of Performance) for cooling (=0.4)	1 460.00€	2 236.60€	3 046.96€	3 893.33€	4 778.18€
kWh per year including cooling	51 100	78 281	106 643	136 267	167 236
Additional costs due to electrical losses per year (including cooling)	5 110.00€	7 828.09€	10 664.35€	13 626.67€	16 723.64€
Additional costs over 10 year operational period	51 100.00€	78 280.85€	106 643.48€	136 266.67€	167 236.36€
Additional costs due to electrical losses over 10 years as compared to Type 1		27 180.85€	55 543.48€	85 166.67€	116 136.36€

CO ₂ emissions caused by UPS electrical losses per year	UPS Type 1	UPS Type 2	UPS Type 3	UPS Type 4	UPS Type 5
1 kWh electric power = 0.6 kg CO ₂ *	30 660 kg	46 969 kg	63 986 kg	81 760 kg	100 342 kg
Difference in CO ₂ emissions per year as compared to Type 1		16 309 kg	33 326 kg	51 100 kg	69 682 kg
Equivalent of driving a Golf-class automobile an annual distance of:		70 290 km	143 635 km	220 241 km	300 329 km

*Source: International Economic Forum on Regenerative Energies (IWR) – <http://www.iwr.de/re/eu/co2/co2.html>

Table 4: Exemplary calculation of UPS power losses and CO₂ emissions

6.2.2 Impact of UPS performance on energy efficiency and CO₂ emissions

Table 4 illustrates the importance of UPS performance. The example shows a data center with a 100 kW load and electricity costs of € 0.10 per kWh with a CoP factor of 0.4 (Coefficient of Performance = a factor representing the additional cooling performance required to cool electrical losses.) Then the CO₂ emissions for the different efficiency levels are calculated and compared with the CO₂ emissions of an automobile.

security requirements. This must be done very early on when supply pathways are being planned. This type of classification is shown in Table 5.

The building's normal power supply (NPS) is fed in either via a direct connection to the national grid (usually up to 300 kW at 400 V) or from the medium voltage grid (up to 52 kV) via distribution transformers up to 2 MVA. For the back-up power supply (back-up power system BPS), one has to differentiate between the emergency power supply (EPS) and the uninterruptible power supply (UPS), depending upon permitted interruption times. This breakdown is shown in Fig. 11.

6.2.3 Power supply in facility management

In order to utilize electric power efficiently, it is important to first classify the electrical loads in terms of their

Type	Example
Normal Power Supply (NPS)	Supplies all systems and loads present in the building
Emergency Power Supply (EPS)	Supplies all emergency systems in hazardous situations e.g.: <ul style="list-style-type: none"> ■ Emergency lighting ■ Fire elevators ■ Fire extinguishing systems
Uninterruptible Power Supply (UPS)	Supplies sensitive loads which must be operated continuously in the event of GPS failure/disturbance, e.g.: <ul style="list-style-type: none"> ■ Servers/computers ■ Communications equipment ■ Control systems ■ Emergency lighting, tunnel lighting

Table 5: Types of Power Supply Feed¹⁴




Fig. 11: Network structure according to load demands¹⁵

Data center air conditioning provides an interesting example for the importance of the uninterruptible power supply. Any interruption in the power supply is basically no problem for air conditioning equipment, as long as the emergency power (EP) generator kicks in on schedule within 10 to 15 seconds. But problems may arise in the EP system if the generator fails to start up. If the air conditioning in the computer room breaks down, the computers, disk arrays and servers could overheat during the shutdown, resulting in severe hardware damage. So it makes sense to have the air conditioning at least partially backed up by the UPS system.

However, a data center's energy efficiency doesn't have to suffer if the following energy management configuration for the air conditioning's UPS-backup is used. Modern UPS equipment offers the option of using a "digital interactive mode" for air conditioning equipment, via digital control and a static bypass switch. This means that the air conditioning system can be powered directly from the AC input, bypassing the defective UPS power electronics. When problems occur in the power supply network, UPS monitoring will ensure that the load is switched to the safe, standby UPS power path.

6.3 Safe, energy-efficient UPS topology

Each component in the power supply chain to the load system creates procurement and operational costs, either through losses or for maintenance and monitoring. As the following presentation assumes an uninterruptible power supply with no breaks (interruption = 0 ms), the supply voltage is converted twice by the power electronics of the UPS. The rectifier converts the incoming alternating current into direct current. This can then charge the battery in the intermediate DC circuit as well as supplying the inverter input with the required power during normal operations. The inverter converts the DC voltage into clean, stable AC output voltage, which is, in turn, completely unaffected by disturbances and fluctuations in the supply voltage. Besides acting as a buffer in the event of power failure, this so-called double conversion UPS also provides optimal power quality which it continually monitors.

Fig. 12 is a schematic showing the components of a typical double conversion UPS with normal operational energy losses.




Fig. 12: Schematic showing a UPS topology and how the components influence efficiency.

The requirements for the individual components of a double conversion UPS depend to a great extent on the application conditions and operational security demands. The input power factor should be about 1 to avoid generator compatibility issues. This can be achieved through power factor correction at the UPS input or by the application of so-called IGBTs (Insulated-Gate Bipolar Transistors) in the UPS rectifier. Possible efficiency losses through power factor correction are compensated for by reduced system perturbations from the IGBT switch. These can affect the input environment.

Thanks to developments in power electronics, it is now possible to operate the IGBT inverter without an output transformer. This does not impair the quality of the output voltage. Apart from avoiding losses in the inverter transformer, it also dispenses with the need for either the galvanic isolation of the battery voltage from the connected load or the use of a transformer in the output network configuration. It is therefore eminently advisable to involve specialists at the early stages of consulting and planning.

■ 6.4 Redundancy: security and efficiency

6.4.1 The conflicting goals of redundancy and energy efficiency

Energy management has to allow for the required levels of equipment reliability as well as technical constraints. Whereas the schematic illustration in Fig. 12 supplies no information on redundancy in the power supply channels, increasing numbers of organisations are engaged in classifying power infrastructure according to availability requirements. The best known example of this is probably the American Uptime Institute's "Tier Classification" system.

Fig. 13 shows a Tier IV configuration calculated to deliver 99.99 percent availability.

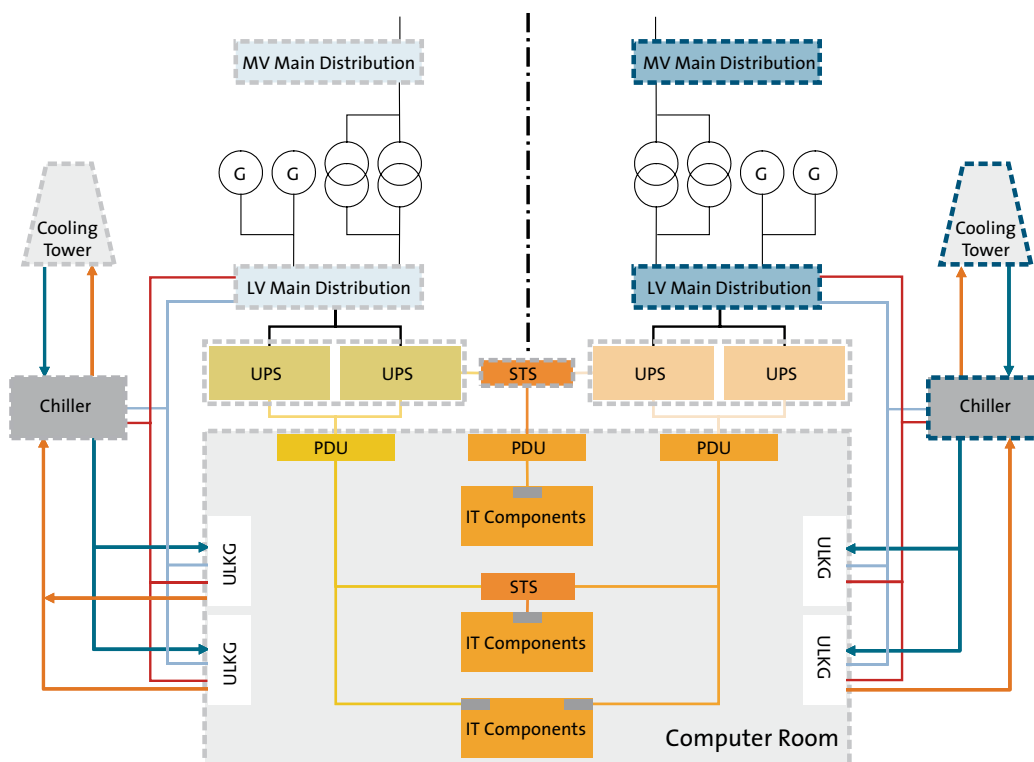


Fig. 13: Infrastructure conforming to Uptime Institute's Tier IV Classification¹⁶

Similar classification systems are offered by IEEE (Availability Classes, Class V is the equivalent of 99.999 percent) and the Harvard Research Group (AEC = Availability Environment Classification). A clear, concise data center perspective on this subject is provided in two BITKOM publications¹⁷: a guide and a planning matrix both entitled “Reliable Data Centers”, which include design suggestions for the individual installation components.

It is also important to strike the right balance between security and efficiency when developing redundancy concepts.

The typical efficiency curve for UPS systems (Fig. 14) shows low efficiency under low load conditions and the best efficiency levels at full load. The main reason for this is that losses occur even when the system is in standby mode and these occur consistently over the entire load range. Proportional losses occur as the load increases. The initial ascent of the efficiency curve should be as steep as possible and then flatten out to a plateau. Some UPS systems operate most efficiently at loads of between 50 and 80 percent.




Fig. 14: Typical UPS efficiency curve (Low mains: 380 V, Nom. mains: 400 V, High mains: 415 V)

When UPS systems are operating in continuous operation mode, it is recommended that they are run at a maximum 80 per cent load, to provide reserve capacity for sudden load variations and instant expansion.

6.4.2 The different UPS configuration concepts in comparison

Single module configuration

Rightsizing is always a critical factor in single module systems, as they are configured to meet future load requirements. They are often oversized for fear that performance limits will be met too soon. But this leads to underutilization and unnecessary losses (see efficiency curve of the selected UPS).

Dual module parallel system for N+1 redundancy (also 1+1)

This is a configuration for individual unit loads of between approx. 20 kVA /kW and 250 kVA /kW, as it offers better availability than single module systems. As with the single module set-up, this system configuration is also sized to allow for future maximum loads, but with the added constriction that each module must stay below 50% load if the redundancy concept is to be properly implemented.

Multiple UPS modules operating in parallel for N+1 redundancy

This set-up is for individual unit loads of between about 20 kVA/kW and 1000 kVA/kW, as it offers better scalability than dual module parallel systems. In this system configuration, an external maintenance bypass switch, sized to meet the desired maximum load and with the corresponding number of UPS connections, makes it possible to connect the required number of UPS modules to meet the momentary load and redundancy level. When load demands increase, one or more UPS modules can be added without compromising normal operations.

This ensures better utilization rates than with 1+1 configuration and the UPS systems can be operated with fewer energy losses.

Modular UPS systems with power modules in parallel mode for N+1 redundancy

This set-up has capacity ratings of between about 10 kVA/kW and 200 kVA/kW per individual module and thus offers better scalability than two or more UPS systems in parallel operation. The system framework of modular

UPS systems is designed to meet the desired maximum final load. UPS capacity can be increased during operation without the need for an external bypass switch with multiple UPS connections. Multiple power modules can be operated in parallel, thus increasing individual module utilization levels and reducing power losses during UPS operation.

Two independently operated USV systems for 2N redundancy

This is a high availability redundancy concept with two separate, independent, fully maintainable power distribution systems (dual bus system). As far as utilization is concerned, this configuration displays identical characteristics to the dual module parallel system for N+1 redundancy (also 1+1).

The following table presents some further redundancy concepts and the possible utilization levels of individual UPS systems and power modules.

Redundancy Level	N	N+1	N+1	N+1	N+1
Configuration	1	1+1	2+1	3+1	4+1
Number of UPS	1	2	3	4	5
% Utilization per UPS	100	50	66	75	80
Redundancy Level	2N	2(N+1)	2(N+1)	2(N+1)	2(N+1)
Configuration	2	2(1+1)	2(2+1)	2(3+1)	2(4+1)
Number of UPS	2	4	6	8	10
% Utilization per UPS	50	25	33	37,5	40

Table 6: Utilization of UPS systems in different redundancy scenarios (permissible maximum capacities in compliance with redundancy classes)

The efficiency statistics provide valuable guidance for UPS system planning as they complement the UPS classifications, but because of the highly complex relationship between technology, operations, security, ambient conditions and disturbance reactions, they should never serve as the sole assessment criteria.

6.5 Load characteristics, performance category and semiconductor components

The influence of the load on efficiency is reflected not only in performance levels, but also in the required voltage waveform. Connected PSUs in servers, PCs or hard disk arrays for example, can cause a phase shift between current and voltage which is denoted by the load power factor. Harmonics can contribute to the non-linear characteristics of the load. A comparison of the influence of non-linear loads and linear (ohmic) loads is shown as a schematic diagram in Fig. 15.




Fig. 15: UPS efficiency as a function of load exploitation for linear and non-linear loads¹⁸

It is also the case for UPS that, assuming the same technology and classification, losses from more powerful units are easier to reduce than those from the small units. The more compact they need to be, the more action needs to be taken as far as heat management is concerned.

Inefficient components mean an increase in waste heat, resulting in the need for additional heat control, possibly involving the implementation of more or larger fans. Another advantage of larger capacity UPS modules is that a three-phase network can supply 400 V as opposed to the single phase AC voltage of 230 V. Furthermore, in the field of UPS technology, there have been continual developments in power semiconductors and control engineering

over the past 40 years. A significant step forward, apart from materials development in metals, semiconductors and synthetics, has been the transition from thyristors to transistors and then to IGBT transistors for inverter power components (Fig. 16).




Fig. 16: UPS efficiency as a function of load exploitation for different electronic components¹⁹

Although higher power densities demand special types of cooling, they also offer considerable advantages in terms of installation size and performance related losses.

Increased semiconductor switching frequency, the introduction of real-time control and finally modern vector control systems have all had a favourable effect on the quality of the load voltage and the reduction of losses.

6.6 Current developments in UPS efficiency

Manufacturers of UPS equipment are cooperating with the EU to produce a directive which will harmonize efficiency assessment and functional requirements for UPS systems. The short-term result is a common code of conduct which sets targets for UPS efficiency levels as from 2008²⁰.

The European norm EN 62040-3 forms the basis for the classification of UPS equipment and the definition of evaluation requirements.

Table 7 shows the planned CoC proposals for the various different UPS classifications for different load and UPS performance categories.

Efficiency in normal mode with non-linear load according to EN62040-3 at:	UPS Output Power ≥ 10 to < 20 kVA			UPS Output Power ≥ 20 to < 40 kVA		
	Class./1/	Class./2/	Class./3/	Class./1/	Class./2/	Class./3/
25% of nominal power	83.0%	88.0%	93.0%	84.0%	88.5%	93.5%
50% of nominal power	89.0%	92.0%	95.0%	89.5%	92.5%	95.5%
75% of nominal power	90.5%	92.5%	95.7%	91.0%	93.0%	96.3%
100% of nominal power	91.0%	92.5%	96.0%	91.5%	93.0%	96.5%

	UPS Output Power ≥ 40 to < 200 kVA			UPS Output Power ≥ 200 kVA		
	Class./1/	Class./2/	Class./3/	Class./1/	Class./2/	Class./3/
25% of nominal power	86.5%	89.0%	94.0%	89.0%	91.5%	95.0%
50% of nominal power	90.5%	93.0%	96.0%	92.0%	94.5%	97.0%
75% of nominal power	92.0%	93.5%	96.7%	93.0%	94.5%	97.7%
100% of nominal power	92.0%	93.5%	97.0%	93.0%	94.5%	98.0%

Notes:

Class./1/: UPS types “VFI – S-...” according to EN62040-3

Class./2/: all UPS types “VI” and “VFI” according to EN62040-3, except “VFI – S-...”

Class./3/: all UPS types “VFD” according to EN62040-3

Table 7: UPS efficiency targets according to the CoC for UPS²¹

The stipulations for additional components with efficiency-reducing effects should also be taken into consideration, as these effects are not always clearly recognizable in the UPS category. Voltage quality and security are critical factors when these components are used in UPS

applications. The list currently specifies additional transformers for UPS modules and filters for the reduction of harmonic distortions.

Normal Mode: Load according to EN62040-3	Additional devices for the reduction of harmonic currents at input better than IEC 61000-2-2, 61000-3-2 and 61000-3-12 <i>The measured losses of these can be deducted from the system losses and shall not exceed the given values per device</i>			
25%	0.6%			
50%	1.0%			
75%	1.6%			
100%	2.5%			
Normal Mode: Load according to EN62040-3	Additional isolation transformer at input or output of the normal mode power path <i>Maximum losses per transformer (in kVA)</i>			
	≥ 10 to < 40	≥ 40 to < 200	≥ 200 to < 500	≥ 500
25% of nominal transformer power	1.5%	1.0%	0.7%	0.5%
50% of nominal transformer power	1.9%	1.5%	1.1%	0.7%
75% of nominal transformer power	2.6%	2.0%	1.7%	1.3%
100% of nominal transformer power	3.6%	3.2%	2.7%	2.0%

Table 8: The influence of additional UPS components on efficiency according to CoC for UPS²²

■ 6.7 Balancing energy efficiency, security and quality

As the previous sections have shown, UPS modules have to satisfy a whole range of different requirements. To conclude this topic, we now present a few considerations on the different UPS component requirements which also affect UPS energy consumption.

For example, in the UPS configuration shown in Fig. 12, various different input filter design alternatives are possible. Normally the aim is for the lowest possible interference with the network installation environment in the vicinity of the UPS. Electric drives and motors, for example, can experience interference from circuit feedback caused by the UPS. Passive or active filters, the use of different rectifiers such as six or twelve-pulse bridge circuits with diode components or IGBT transistor technology can lead to diverging filter effects and to various intrinsic UPS energy losses.

Another UPS parameter is the voltage range within which the rectifier can supply the inverter with power without having to switch to battery operation. If this is more generously configured it can mean that UPS efficiency is reduced during normal operation. But as the batteries do not need to be discharged and charged so often, the overall energy footprint may be more favourable in the long run. The rectifier is normally also utilized for charging the battery and so it should be adequately dimensioned. Great attention should be paid to battery charging and monitoring to ensure that the battery packs are treated with the greatest possible care. This has now become such a complex process that it has not yet been incorporated into the efficiency analysis. A battery management system which is optimally adjusted to the type of battery in operation and which monitors the status of the battery packs is able to give adequate warning should problems occur.

All electronic monitoring and control systems require electrical energy. But does it make sense to do without input level monitoring or fan control in order to save electricity at normal full capacity operation, or perhaps

forgo the weekly battery test which then has to be carried out with a manual meter?

Communication interfaces and equipment displays such as LCD and LED also require energy. Even the EMC components in the UPS which minimize interference within the UPS and in its immediate environment, affect efficiency levels. EMC requirements will be dictated essentially by the UPS installation location. Another way of minimizing energy losses would be to dispense with the electronic bypass switch. But the lack of an electronic bypass would show up negatively in a reliability calculation.²³

There have also been significant developments in the field of inverter technology. Modern power technologies perform at high switching frequencies and provide clean output voltage without the need for an inverter transformer. This would make it possible to improve efficiency by between one and three per cent, but security-conscious users value the inverter transformer because it provides galvanic isolation between the connected load and a DC charge in case of an inverter fault. The so-called transformerless UPS modules have a protection system which becomes operative when DC components are present in the UPS output voltage.

■ 6.8 Rackmount intelligent power strips

As well as the different alternatives for optimizing the energy efficiency of the UPS systems themselves, there is also the option of using “intelligent” power strips, so called Smart Power

Strips (SPS) which save energy on the power supply side. Unlike previous power strips which were only used to accommodate plugs, these SPS feature several characteristic functions: simple Plug and Play, modular design and easy access to different circuits. In addition, these power strips offer remote control over IP networks and an integrated monitoring function. To cater for the stringent requirements of data center operations, special solutions with adapted housing and different current capacities have been developed.



There are two basic construction design types which are typically used in the rack:

- 19" 1U or 2U (8 ports, up to 16 ports),
- vertical power strips which attach to the side of the rack (12 ports, 20 ports)

From a functional point of view the power strips are available in three categories:

- passive power strips for IT environments
- switchable power strips with control input
- intelligent power strips with control input and monitoring functions

Passive power strips for IT environments

With simple models it is important to make sure that the strip has an adequate number of power outlets (for 'Schuko' or rubber connectors), as well as high quality mechanical characteristics. The use of power strips with on-off switches is not to be recommended as the risk of accidental shutoff through mechanical contact is far too great. There is also a model with integrated surge protection available. This helps avoid malfunctions caused by temporary voltage spikes. The main disadvantages of the simple models are the unavoidable start-up peaks (when all loads are switched on simultaneously), the absence of either on/off control or consumption monitoring.

Switchable power strips with control input

A more advanced development of the first category comprises switchable power strips with a control input

(e.g. serial RS232, partially Ethernet) which enables the controlled start-up of devices. The power strips are typically controlled via a KVM switch which also facilitates the grouping and assigning of individual outlets to the connected devices (e.g. servers).

Intelligent power strips with monitoring functions

As well as their switching function, the latest generation of power strips offer power consumption monitoring (overall consumption and individual outlet measurements) to provide you with an accurate overall picture of your data center's energy consumption and efficiency. Here it is important that individual measurements be taken at each power outlet in order to ascertain the exact load distribution in the rack and if necessary to determine the best spatial positioning in the data center for new servers. This can also form a basis for the optimization of cooling and air-conditioning systems.

The measured data can be read out using a monitoring tool via Ethernet and SNMP and recorded as a trend. Following this you then have option of analyzing the individual loads within the framework of a statement of costs. This allows the individual allocation of costs per load (e.g. calculation of energy costs per section or server). Intelligent power strips also offer other functions such as environmental monitoring (temperature, humidity) adjustable thresholds (e.g. power) with alarm functions (notification via e-mail or SNMP).

Category	Application in	Advantages	Disadvantages
Passive power strips	Simple installations	<ul style="list-style-type: none"> ■ cost price 	<ul style="list-style-type: none"> ■ no control features ■ start-up peaks ■ no metering
Switchable power strips	Data centers with server management	<ul style="list-style-type: none"> ■ switch function delays server start-up 	<ul style="list-style-type: none"> ■ no consumption measurement ■ usually only in combination with a KVM switch or console server
Intelligent power strips	Energy efficient data centers	<ul style="list-style-type: none"> ■ allows full control of equipment and energy consumption ■ suitable for standalone applications 	<ul style="list-style-type: none"> ■ cost price

Table 9: Power strip overview

7 Energy Contracting

Cooling and power supply management constitute challenging tasks for data center operators. One way in which the burden on the organisation can be relieved with respect to this task is energy contracting. Purchasing energy via a third party (contractor) can offer various advantages for suitable clients (contractees). For example, clients don't need to invest in the energy system, and this releases funds for investment elsewhere.

The contracting company's services usually include consultancy, planning, financing and the operation of energy systems. In addition they take over the purchase of electric power, as well as gas, oil, district heating or district cooling at favourable market conditions on the basis of negotiated mid and long term supply contracts. The contractor offers professional energy management services and is in a position to guarantee the availability levels required by a modern data center. The contractor is also able to utilize the available potential to reduce energy consumption and operational pollutant emissions. Comprehensive maintenance and repair guarantees are also often included in an energy contractor's service portfolio.

Thanks to its size, experience and specialist competence, the contracting company is able to take measures to exploit the available ecological and economic potential – a task for which the data center operator often has no available human or financial resources.

To sum up we can identify the following potential advantages for contractees:

- Less pressure for facility management and IT departments. The contractor assumes responsibility for coordination and execution of the tasks.
- No investment expenditure: results in a freeing up of financial capital for other projects; increased liquidity.
- Increased supply security.
- Application of efficient and energy-saving technologies.
- Effective energy monitoring.
- Support for internal and external cost allocation.
- Regular maintenance and servicing of the system.
- Cost benefits and synergy effects arise when the contractor handles purchase, planning and installation of the system. These effects can also occur to a certain extent through in-house production of electricity, heat, cooling and the efficient utilization of waste heat.
- In some cases, the value of the property may even be increased by the implementation of state-of-the-art technologies.

Endnotes

1. This guide thus complements the BITKOM guide "Reliable Data Centers" and the matrix "Planungshilfe Betriebssicheres Rechenzentrum" (Planning aid for reliable data centers) which can be downloaded free of charge from the BITKOM website. However, both guides and planning aid are no substitute for expert advice and support from competent consultants and planners.
2. Fichter, Klaus: Zukunftsmarkt energieeffiziente Rechenzentren. Study carried out on behalf of the Federal Environment Ministry (BMU); (Borderstep Institute); 2007
3. Fichter, Klaus: Zukunftsmarkt energieeffiziente Rechenzentren. Study carried out on behalf of the Federal Environment Ministry (BMU); (Borderstep Institute); 2007
4. US EPA (U.S. Environmental Protection Agency, Energy Star Program): Report to Congress on Server and Data Center Energy Efficiency, Public Law 109-431, 2007
5. The Green Grid: GREEN GRID METRICS: Describing Datacenter Power Efficiency. Technical Committee White Paper, February 20, 2007
6. Brill, Kenneth G.: Data Center Energy Efficiency and Productivity, White Paper, The Uptime Institute, 2007
7. ASHRAE, Thermal Guidelines for Data Processing Environments; 2004
8. A unit's sensible cooling capability is its performance in cooling the air without dehumidifying it.
9. Strictly speaking, it is the ratio between air volume and heat exchanger surface area which needs to be optimized. A large heat exchanger surface area requires a sufficient volume of air.
10. Although water is frequently used as a cooling medium, it has not proven itself for application as a coolant in data center environments.
11. Return air temperature is only significant when data centers are air-cooled. In installations where all the server racks are water-cooled, for example, the room air temperature and thus the return air temperature, too, are of no consequence for a data center's energy efficiency. In such cases more attention should be paid to temperatures in the racks' water system.
12. As mentioned above, another advantage of higher return air temperatures is the higher temperature spread which is possible when water is used as a coolant. An increased spread allows higher operational temperatures on the recooling side and this makes free cooling possible, even at higher outdoor temperatures. The longer free cooling is used, the less energy is needed for the operation of compressors.
13. Swiss Federal Office of Energy: Risikofreier Betrieb von klimatisierten EDV-Räumen bei 26° C Raumtemperatur (Risk-free operation of air-conditioned computer rooms at room temperatures of 26° C), Bern 1995
14. Totally Integrated Power: Applikationshandbuch – Grundlagenermittlung und Vorplanung (Application Manual - Basic Data and Preliminary Planning), Siemens AG; 2006
15. Totally Integrated Power: Applikationshandbuch – Grundlagenermittlung und Vorplanung (Application Manual - Basic Data and Preliminary Planning), Siemens AG; 2006
16. Uptime-Institute: www.uptimeinstitute.org
17. Both publications can be downloaded free of charge from the BITKOM website
18. High performance buildings: Data centers uninterruptible power supplies (EPRI-Solutions, Ecos Consulting); 2005
19. Swiss Federal Office of Energy: Program "Electricity" Code of Conduct / Label for Uninterruptible Power Supplies (UPS); 2005
20. Code of Conduct on Energy Efficiency and Quality of AC Uninterruptible Power Systems; 2006
21. Code of Conduct on Energy Efficiency and Quality of AC Uninterruptible Power Systems; 2006
22. Code of Conduct on Energy Efficiency and Quality of AC Uninterruptible Power Systems; 2006
23. White Paper: The classifications define site infrastructure performance (The Uptime Institute Inc.); 2001-2006

8 Glossary

19"-Cabinet	Rack with approx. 40 U, overall height approx. 2 metres, installation width 483 mm, installation height is measured in rack units (U), 1 U = 44.45mm
BPS	back-up power system
Brine	water/glycol mixture
Chiller	water cooled unit
CPU	central processing unit
CRAC	computer room air conditioning
CW	chilled water; air-conditioning units which use cold water
Data Center	server room and /or data center
DX	direct expansion; refrigerant-type air conditioning units
EMC	electromagnetic compatibility
Emission	the release of substances into the atmosphere from a resource or process
EPS	emergency power supply
EU	electric utility
Hot spots	local overheating
I	inverter
ICT	information and communication technology
IGBT	insulated gate bipolar transistor
ILM	information lifecycle management
Immission	the reception of material, such as pollutants, by the environment from any source
IT	information technology
KVM switch	keyboard-video-mouse switch
LV main distribution	low voltage main distribution
Modular	a system constructed of several individual units (sub-assemblies)
MV main distribution	medium voltage main distribution
NPS	normal power supply
Parallel operation	two or more UPS units which combine to supply the connected loads
Power distribution	also main LV distribution panel (LVDP) or PDU (power distribution unit)
Precision air conditioner	air conditioning unit which can maintain both temperature and humidity at constant levels
PSU	power supply unit
R	rectifier
Redundant	multiple configuration for increasing availability (fault tolerance)
Scalable	capable of adapting to demand
SNMP	simple network management protocol
SP	sub panel
SPS	smart ("intelligent") power strips
STS	static transfer switch / static transfer system
Surge protection concept	conventional overvoltage protection includes coarse, medium and fine protection
TCO	total cost of ownership
UPS	uninterruptible power supply

10 Acknowledgements

This guide to “Energy Efficiency in the Data Center” has been produced in consultation with the BITKOM Working Group on “Fail-safe Computer Centres & Infrastructure”.

We would like to express our sincere thanks to all members of the Working Group for their valuable discussions and suggestions, as well as to the following individuals for their contributions and comments:

- Dr. Wolfgang Gnettner (Fujitsu Siemens Computers GmbH)
- Frank Hauser (Raritan Deutschland GmbH)
- Dieter Henze (Rittal GmbH & Co. KG)
- Robert Horn (KKT Kraus Kälte und Klimatechnik GmbH)
- Dr. Siegbert Hopf (Masterguard GmbH)
- Eberhard Knödler (Brach + Moll Kälte und Klimatechnik GmbH)
- Stephan Lang (Weiss Klimatechnik GmbH)
- Achim Pfeiderer (Stulz GmbH)
- Dr. Georg Riegel (deZem GmbH)
- Staffan Reveman (Newave USV Systeme GmbH)
- Christian Richter (Emerson Network Power GmbH)
- Michael Schumacher (APC Deutschland GmbH)
- Judith Wagener (Bull GmbH)
- Manfred Willnecker (Knürr AG)
- Ingolf Wittmann (IBM Deutschland GmbH)
- Ralph Wölpert (Lampertz GmbH & Co. KG)
- Eckhard Wolf (AEG Power Supply Systems GmbH)
- Carsten Zahn (Schnabel AG)

Our very special thanks to Harry Schnabel (Schnabel AG), head of the Working Group.

The German Association for Information Technology, Telecommunications and New Media (BITKOM) represents over 1,200 companies. Its more than 900 direct members generate a sales volume of 135 billion euros annually and employ 700,000 people. They include providers of software, IT and telecommunication services, manufacturers of hardware and consumer electronics as well as digital media enterprises. BITKOM is committed in particular to an improved regulatory framework, a modernized education system and an innovation oriented economic policy.

This translation was kindly sponsored by the following companies:



Rittal - Complete IT Competence



German Association for Information Technology,
Telecommunications and New Media

Albrechtstraße 10 A
10117 Berlin-Mitte
Tel.: +49.30.27576-0
Fax: +49.30.27576-400
bitkom@bitkom.org
www.bitkom.org