White Paper

Energy Efficient Infrastructures for Data Centers
a joint White Paper from Fujitsu Siemens Computers and Knürr

Edition July 2007

Pages 16

Worldwide power requirements to operate data centers are growing inexorably and have now reached dimensions that even by comparison with other industry sectors can no longer be side-lined. The constantly growing need for IT support in existing and new business processes is giving rise to ever greater number of servers and storage systems in worldwide operation.

The CO_2 issues associated with this growth are focusing public interest on data center operators. Due to continual increases in energy prices, power is emerging as an ever larger cost factor in internal calculations of overall operating costs. However, the key aspect for many companies is the fact that many data centers are unable to cope with the space, power and air conditioning requirements of the new IT equipment generations. Unless appropriate concepts are adopted, data centers may even compromise business expansion. For example, over the last few years heat loads in data centers and server enclosures have increased to such an extent that classical precision air conditioning at room level no longer functions.

It is therefore vital to quickly identify and implement effective measures for both IT systems and surrounding physical infrastructures in order to sustainably improve energy efficiency in data centers. Many of the components needed to do this are already available, and each component alone provides tangible optimization potential with a sometimes significant impact on the total cost of ownership (TCO).

The following measures make a big contribution on IT level:
- Minimization of power consumption through the use of energy efficient systems
- Usage of cooling optimized systems
- Consolidation onto fewer and higher performant systems (e. g. blade server architectures)
- Improved utilization of IT systems with the help of virtualization technology
- Flexible control of the power consumption of IT systems by means of dynamic IT solutions

The following measures may be adopted on infrastructure level:
- Hot/cold aisle rack arrangements
- Closed, directly cooled racks
- Careful management of the cooling air flow in racks and in the data center
- Constant matching of the cooling system to current operating needs
- Use of low-loss UPS systems
- Right choice of room temperature (every degree Celsius less contributes to higher energy efficiency)
- Right choice of location and insulation of the building/envelope

However, from today's standpoint, a combination of various measures on the IT and infrastructure levels is needed to stop or even reverse the rise in data center energy costs. For this reason, Fujitsu Siemens Computers and Knürr are cooperating on joint concepts to develop integrated solutions that will keep energy costs in data centers under control, today and into the future.

Contents

1 Introduction

1.1 Data Centers at the Center of the Conflict between Classical Demands and Energy Efficiency

The key criteria by which CIOs and IT managers have always been judged are IT performance, functionality and availability coupled with maximum return on investment (ROI). Not surprisingly, these were also the main factors when selecting server and storage systems for deployment in data centers. The associated energy costs were simply the price that had to be paid for added performance. No energy awareness could develop because the data center's electricity bill never landed on the desk of the CIO or of the data center manager. In the meantime, concrete numbers, trends and forecasts by various parties are persuading CIOs to put data center energy efficiency on their agenda.

In this White Paper we intend to identify the main reasons behind these trends and also to put forward concrete proposals for optimizing data center energy efficiency. In most cases, there is no need to build an entirely new data center or to replace a major part of the existing power and air conditioning infrastructure, or such investment can be at least postponed until a later date.

Even though energy efficiency is a topic that is gaining in importance, the priorities adopted when selecting data center IT systems will change only slowly. A primary goal of the solution approaches we describe below is therefore not only to identify potential for energy efficient operation but also to satisfy the classical demands placed on data center infrastructures and to ensure adequate performance, functionality and availability.

1.2 Growing Public Pressure

Enterprises are facing constantly increasing pressure of a financial, environmental and legislative nature to live up to their corporate social responsibility. They are urged to take stronger action not only to cut carbon dioxide emissions or to use recyclable materials but also to address energy efficiency issues when deploying IT technology. This affects the entire spectrum of IT systems. In the past, the focus of all energy saving measures was on end-user products. Recently however, attention has been redirected primarily to data center energy efficiency, starting with IT infrastructures and continuing to the power-supply and cooling infrastructures.

1.2.1 US Environmental Protection Agency Scrutinizes Data Center Power Requirements

According to a study by Jonathan G. Koomey of Stanford University, data centers consumed a total of 45 billion kW hours of electricity or 1.2% of national electricity use in 2005, thus making them one of the USA's biggest energy consumers[5]. This figure equates to the total annual capacity of five 1000-MW power plants. Reason enough for the US government to pass a law in December 2006 to instruct the Environmental Protection Agency (EPA) to determine data center power requirements. The law required the EPA to submit a report on the rapidly growing power needs of private and public data centers to Congress within 180 days. The agency was also requested to make proposals how existing and new programs could be developed to encourage companies and the public sector to make use of energy saving technologies.

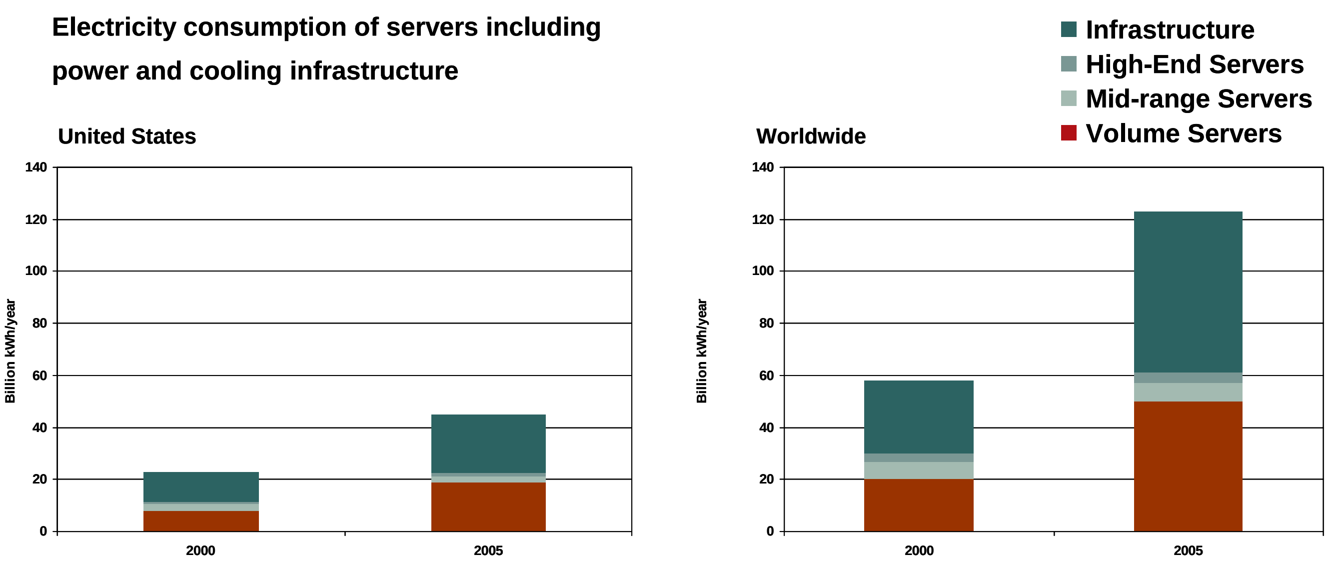

Figure 1: Data centers consume large amounts of electricity - just as much as all color TVs in the USA[5]

1.2.2 Energy Star for Servers

In the view of the EPA there is an urgent need to include servers – in the same way as PCs, refrigerators, DVD players and air conditioning systems – in a classification program to label energy efficient products. In a letter to computer industry representatives the agency stated that it had begun to develop an Energy Star specification for enterprise servers. The Energy Star label for PCs was also officially introduced in the EU in 2002.

1.2.3 BITKOM (German Association for Information Technology, Telecommunications and New Media e. V.)

Also in Europe there are efforts towards labeling of IT products concerning their energy consumption, in order to increase energy efficiency. Fujitsu Siemens Computers submitted a proposal to the German Federal Minister Sigmar Gabriel, which makes it possible to compare the energy consumption of computers, based on a typical annual consumption including the Energy Star specifications. The Federal Minister is very open-minded about the proposal and is ready to start conversations with Fujitsu Siemens Computers and/or the BITKOM about consumption markings on the basis of the existing labels. However, due to the indifferent and dismissive attitude of several American enterprise representatives in the committees, progress in this process was slowed down. From this inititive, Fujitsu Siemens Computers expects a European-wide valid and consumer-friendly labeling obligation of IT systems.

1.3 Increasing Share of Energy Costs in the Overall Costs of Data Centers

The growing financial pressure on IT operators stems from a global rise in power needs that have given rise to a worldwide hike in energy prices. Economists point out not only the immense power requirements of expanding nations such as China and India but also the strong increase in more and more powerful computers. Eurostat has published a study which indicates that between 1 July 2005 and 1 July 2006 the EU-25 aggregated electricity price (in €/kWh) for industrial consumers rose by 16%[2]. The electricity consumption of 45 billion kWh shown in Figure 1 cost US data center operators a total of US$ 2.7 billion in 2005. Worldwide costs are estimated at approx. US$ 7.2 billion[5].

Analysts maintain that power and air conditioning for server and storage systems will, in the medium term, emerge as a dominant theme in data centers. According to IDC, power is already the second largest cost factor in data center operations after maintenance and administration, and this trend is still on the up. As a result, a number of companies operating big data centers have drawn their own conclusions and are building their new data centers in the vicinity of large and, above all, inexpensively priced power stations.

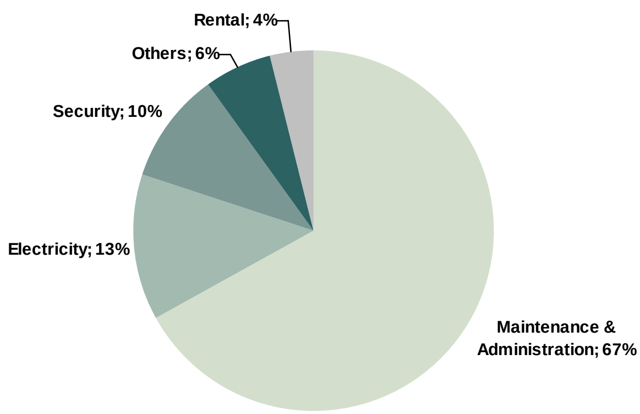

Figure 2: Distribution of operating costs in the data center (source IDC) [10]

1.4 Data Centers at their Capacity Limits – Reasons for Data Center Energy Problems

Analysts predict that data centers will soon reach their capacity limits in terms of power and air conditioning. This may impose tight restrictions on companies who want to expand their business but also need to substantially increase IT support in growing market segments. If there are no changes in the energy efficiency of IT systems and their operation, Gartner forecasts that in 2008 50% of all data centers will no longer be able to meet the power and cooling requirements of their IT systems[1].

1.4.1 Increasing Demand for Greater Performance and the Trend towards More Densely Packaged Systems

Apart from higher energy costs, key causes are the ongoing demand for greater performance and the trend to more densely packaged systems.
The growing use of Internet media and the rise in electronic communication via the Internet are, for many experts, the areas with the greatest future IT support needs. But also the accelerating digitalization of business processes and other influences such as national legislation (e.g. Sarbanes-Oxley Act, Basel II) or disaster recovery measures are dominant, driving factors.

The introduction of newer and even more powerful applications in these usage scenarios often forces IT departments to opt for the latest hardware technology. Servers and also storage systems of the latest generation almost invariably offer greater performance and capacity in less space than comparable predecessors. According to a publication of the "American Society of Heating, Refrigerating, and Air Conditioning Engineers (ASHRAE)[3]" on the development of power density in IT systems, the power density of servers has risen by a factor of 10 in the last 10 years and will continue to rise although at a slower pace.

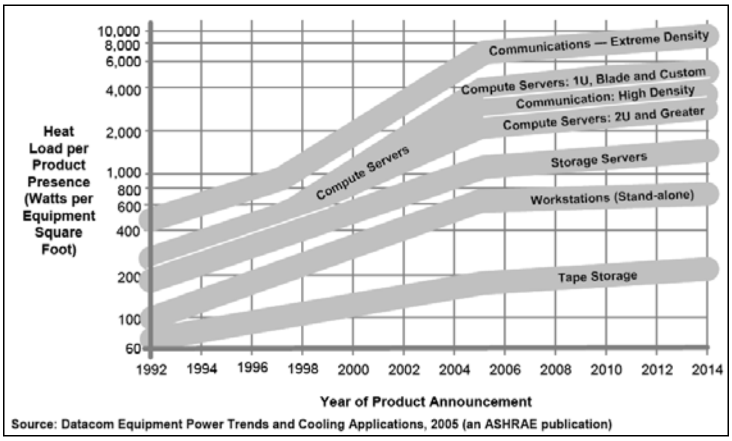

Figure 3: Power density of servers has increased by a factor of 10 in the past 10 years[1]

According to Gartner, the power consumption of a rack fully equipped with blade servers will increase from 15 kW (2005) to approaching 52 kW (in 2011), specifically due to the dramatic rise in power density as a result of the disproportionately high growth rates of blade servers[1]. However, higher power requirements caused by increased server density is just one side of the coin. Data center operators must also provide additional power for the higher demands placed on the data center infrastructure in the form of larger air conditioning systems, fans and power supplies.

Growing power requirements as a result of higher power densities in data center IT systems and constantly rising energy prices are producing an energy costs spiral. In many companies this means that even today energy costs to operate and cool systems are higher than procurement costs. According to the Uptime Institute, the 3-year costs for powering and cooling a server are already between 1 and 1.5 times the procurement costs. A projection into the year 2012 suggests a costs ratio of 22:1 in the worst-case scenario and a ratio of 3:1 in the best-case scenario; nevertheless, the latter is still double the current ratio[4].

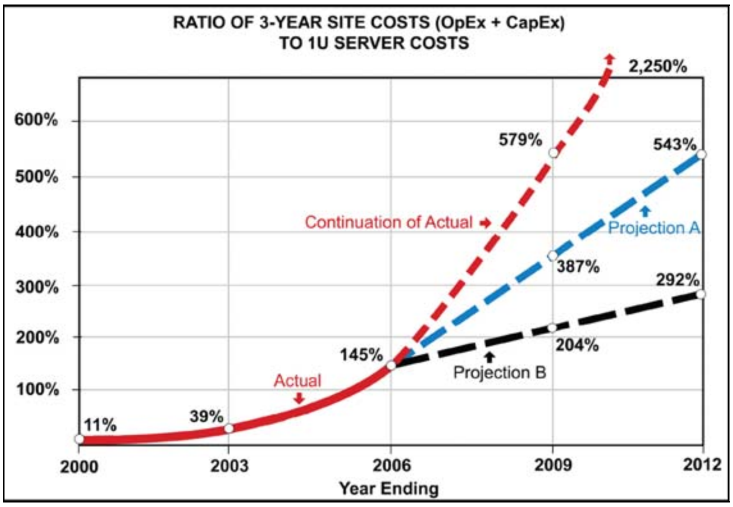

Figure 4: Infrastructure costs for powering and cooling data centers are rising faster than the procurement costs for new servers[4]

1.4.2 Power Wastage due to Inefficient Components and Inefficient Operation

A large part of the energy used in IT operation is not translated into computing power but is dissipated as heat. The past tendency to over-dimension IT systems or to use them for a single application only has resulted in systems that are extremely underutilized, at least in the standard server sector. An idling server still consumes about 60 to 70% of the energy compared to a server in normal production operation.

2 Solutions to Boost Energy Efficiency in IT Infrastructures

Below we discuss a number of measures to effectively counter the energy problems faced by data centers. In this context, it is of course appropriate to take a brief look at the client systems on which most energy saving measures have focused so far. There, the use of energy saving components and of energy management functions has a long tradition. Whereas measures in the client field are naturally restricted to system and component level, the challenge faced in data centers calls for a combination of measures on different levels.

Certainly, a first step is to deploy more energy efficient hardware technology on system level. Server and storage systems based on the latest component technology are much more energy efficient and deliver higher performance with the same or even reduced energy consumption. By using energy efficient memory, CPU, power supplies and fans together with an energy optimized system design, vendors can contribute substantially to the reduction of energy consumption.

Even greater optimization potential when running rather static IT infrastructures is provided by technologies that reduce the total amount of IT systems in operation, thus delivering substantial energy savings on the infrastructure level. Here, switching to multi-core servers or using virtualization solutions can significantly contribute to a consolidation of IT systems. However, virtualization of IT resources is not the only means of achieving more efficient IT operation in static environments by ensuring enhanced resource utilization and consolidation. Virtualization technology is a key component when it comes to building dynamic IT infrastructures.
As a scalable computing basis, virtualization technology is the ideal platform for implementing server pool concepts which, in conjunction with automation technology, guarantee greater flexibility in the control of IT resources as compared with static environments and therefore support more energy efficient IT operation. A higher level of integration between application and IT resource management also permits substantially improved matching of available IT resources to application needs.

The integration of IT resource management into management tools of the cooling infrastructure represents a further step on the path to full optimization of the energy efficiency of data center infrastructures. This enables IT systems to signal changes in cooling needs. By means of timely information slow-reacting cooling systems are able to align cooling capacity to match anticipated demand in good time. This has the advantage that considerable amounts of electricity needed for cooling system operation can be saved because the systems no longer have to run constantly at maximum cooling capacity.

At the same time, many of the potential measures call for little investment as compared with the upgrading of data center infrastructures or the construction of a completely new and modern data center. Ultimately, it is a matter of ensuring that companies can continue to use their existing data center capacity and are able to postpone heavy capital investment in new power and cooling infrastructures for as long as possible.

2.1 Focus on Industry Standard Servers

The greatest potential energy savings are to be made not only in power and cooling infrastructures but also, without doubt, in the area of volume servers (according to IDC, per-system costs below $25,000), simply due to their market share; industry standard servers based on Intel and AMD technology account for the major share in this category. More than 40% of the energy used for server systems worldwide in 2005 was consumed by this server class. With an average yearly increase of 20%, volume servers are also posting the biggest energy consumption growth rates. Between 2000 and 2005 consumption has more than doubled[5]. And the strong annual growth rates of close to 40% for blade servers are also adding to the problems associated with volume servers due to the trend towards greater power density in data centers[1].

A further reason for the concentration on volume servers is certainly that this segment has the greatest technological backlog demand. Most of the servers in the mid-range ($25,000 - $500,000) and high-end (> $500,000) sector already run with high utilization rates. In addition, virtualization technology has been state-of-the-art for many years in this server sector. Utilization rates of up to 90% can be achieved by combining hardware-based and software-based virtualization.

Consumer	2000		2005		Consumption CAGR
	Energy consumption (billion kWh)	%	Energy consumption (billion kWh)	%	
Volume servers	20	34	50	41	20
Mid-range servers	7	12	7	6	0
High-end servers	3	4	4	3	5
Infrastructure	28	50	62	50	17
Total	58		123		17

Figure 5: Server and infrastructure energy consumption worldwide. The greatest energy saving potential is in the area of power and cooling infrastructures and volume servers[5]

2.2 Reduced Energy Consumption through the Use of Energy Efficient Components

2.2.1 Structure of an Energy Efficient Server

Fujitsu Siemens Computers has a "PRIMERGY Power Configurator" tool to determine energy consumption figures for servers. The tool enables energy consumption to be calculated for a specific component configuration level. Figures 6 and 7 show a comparison between a PRIMERGY RX300 S3 system with a standard configuration level and an energy optimized variant with identical performance data. Consumption is reduced by one third through the use of low-voltage processors, a smaller hard disk (2.5" instead of 3.5"), only one power supply and four 2-GB RAM instead of eight 1-GB modules. This is offset by an additional 220 euros of capital investment for the server. The electricity costs for operation are reduced from an average 304 to 210 euros per year. Taking into account the fact that approximately the same energy savings can be made in the cooling infrastructure, the extra procurement costs for the energy optimized system are amortized after about one year.

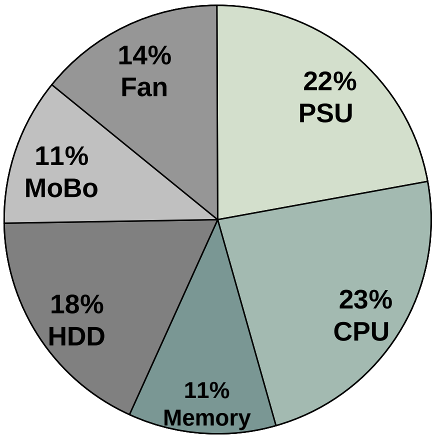

Figure 6: Standard system
E5320 CPU, 8x 1GB RAM, 6x HDD 3.5" 15K

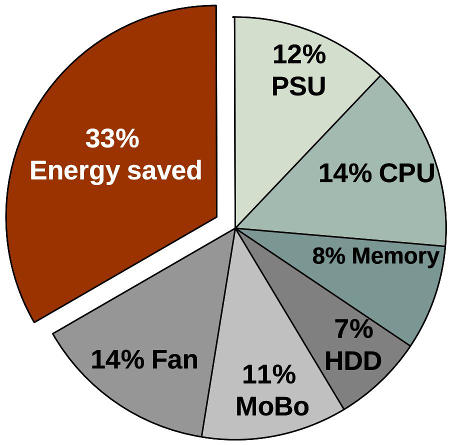

Figure 7: Energy optimized system
L5320 CPU, 4x 2GB RAM, 6x HDD 2.5" 10K

2.2.2 The Difference between US and European Vendors

Despite the fact, that there are still no official comparison and bench mark tests available, already today the possibility exists to compare the different energy values of different server classes on the basis of vendor data. These comparisons show the following results. Servers of a large US vendor consume on average 21% more energy, than comparable servers of Fujitsu Siemens Computers. This means extra costs of approximately 500 Euro within a life cycle of a volume server of at least three years. In all likelihood, the reason for this is caused by inefficiencies of the used power supplies, cooling and design of the systems. The average value is based on a comparison of 55 different configurations. These configurations were selected according to a random sample. The table shows the additional energy consumption for some selected server classes compared to those from Fujitsu Siemens Computers (status quo July 2007).

1 HE – 1 Socket Intel Server	+24 %
1 HE – 2 Socket Intel Server	+35 %
2 HE – 2 Socket Intel Server	+11 %
2 HE – 2 Socket AMD Server	+20 %
Blade – 2 Socket Intel Server	+13 %

For many years, European and Japanese enterprises, private and political organizations are very sensitive to topics around energy savings and environmental friendliness. Therefore, Fujitsu Siemens Computer took over a pioneering role in relation to these requirements. This becomes clear also in this comparison.

2.2.3 World's Most Energy Efficient Server: PRIMERGY TX120

The PRIMERGY TX120 model by Fujitsu Siemens Computers is currently the most energy efficient server in the marketplace. Fully equipped this system consumes only 163 W. The low cooling requirements in conjunction with high-quality fans permit an extremely low noise level of 28 dB in idle mode and of only 31 dB during operation.

2.2.4 Multi-Core Processors: More Performance per Watt

The use of multi-core processor technology is certainly one of the most effective strategies to deliver more computing power with the same or reduced energy consumption. Significant optimization of the performance per watt ratio enables IT managers to increase packaging density without coming up against power supply limits.

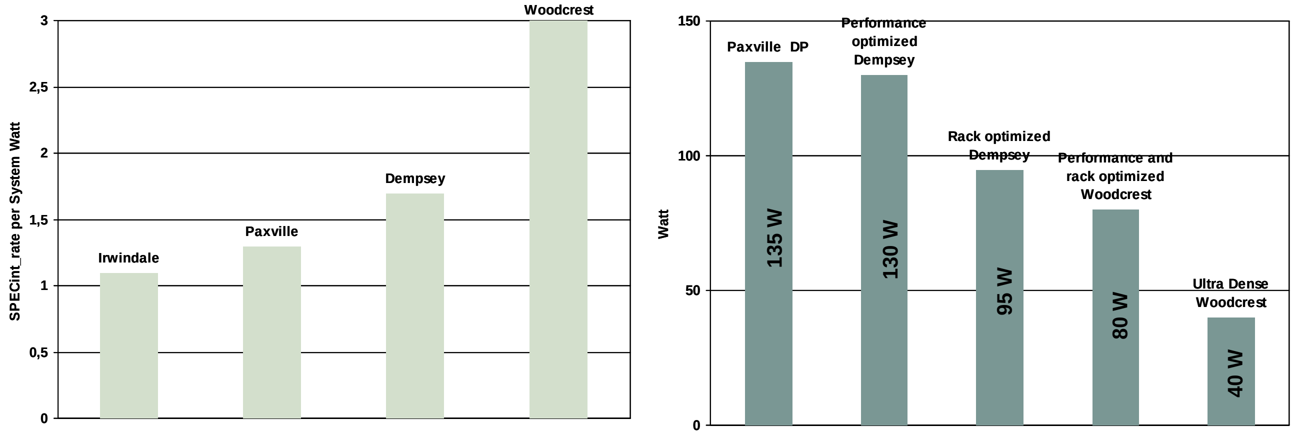

Figure 8: Constantly increasing computing power with the same or reduced energy consumption

2.2.5 Massive Increase in Energy Efficiency Provided by PRIMERGY RX300 S3 Quad-Core Servers

Measurements by Fujitsu Siemens Computers to compare a single-core PRIMERGY RX300 S2 server (launched in 2004) and a quad-core PRIMERGY RX300 S3 server have revealed that the SPECint_rate_base2000 performance benchmark increases 7.6 fold for the same power consumption when moving from the RX300 S2 to the RX300 S3 model. In other words, given the same electricity consumption many more users can be served, and to achieve the same performance the RX300 S3 needs only 13% of the power draw of an RX300 S2.

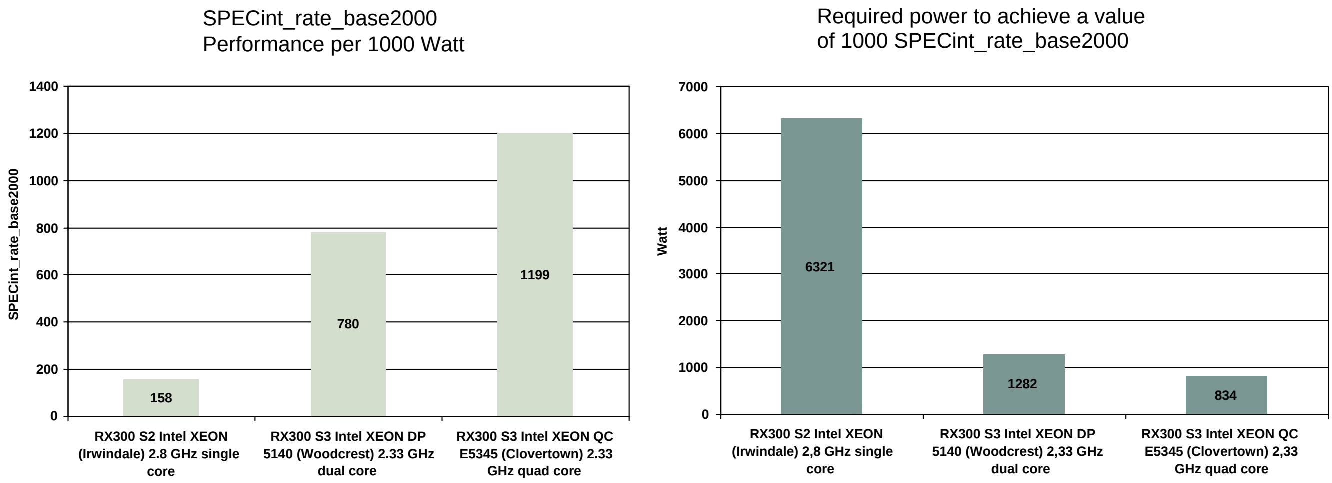

Figure 9: 7.6-fold performance boost for the same power draw or 87% less electricity for the same number of users

In a terminal service scenario, this resulted in concrete energy savings of more than 60% and a cost savings potential of 8.460 euros per year to serve the same number of users. In this scenario, a PRIMERGY RX300 single-core server was able to serve 127 users whereas the RX300 S3 quad-core model served 300 users. Consequently, only eight quad-core servers instead of 20 single-core servers were needed to comply with Service Level Agreements for 2500 users.

2.3 Optimization Potential in the Operation of a Static IT Infrastructure

2.3.1 Systems with Higher Utilization Rates Use Energy More Efficiently

The widespread practice in many data centers of deploying a server for a single specific task is a growing problem, not only in terms of procurement and management costs but also from an energy efficiency point of view. The principle of assigning a separate server to each application does prevent software conflicts and is a pragmatic way of increasing system reliability but it means that many such servers are severely underutilized and "twiddle their thumbs" for long periods. Average utilization rates on standards-based servers of between 5% and 10% are nothing out of the ordinary.

As the difference in energy consumption between a server under normal load and a server under reduced load is very minor, this kind of architecture represents a massive waste of energy and gives rise to unnecessary costs for non-productive systems.
Servers running in idle mode still consume a large amount of energy. Measurements on a typical database server show that the server consumes 550 W under peak load and 400 W during normal production operation but still needs 350 W in idle mode. What's more, additional energy is wasted to cool servers that are not running productively but are simply generating hot air.

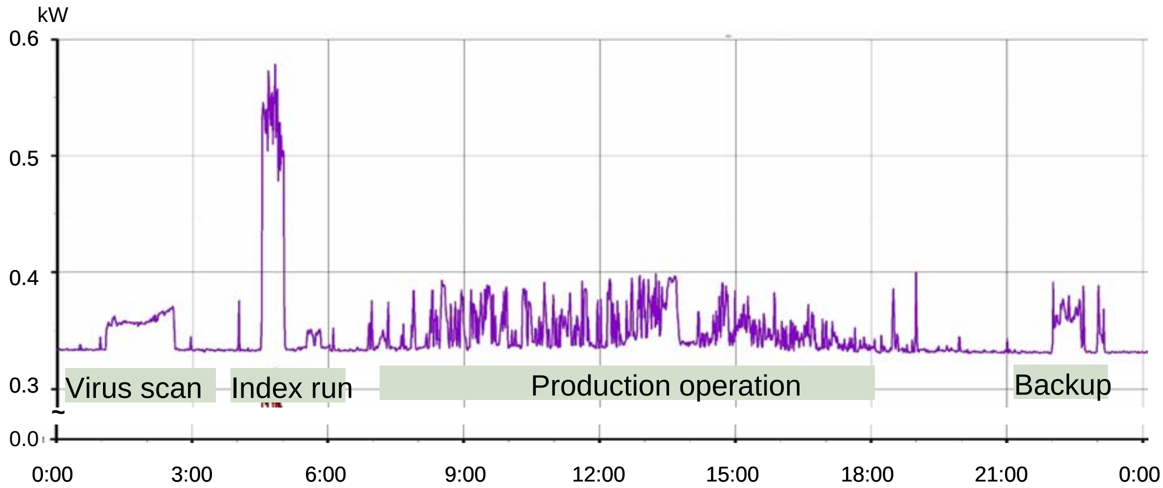

Figure 10: A day in the life of a database server

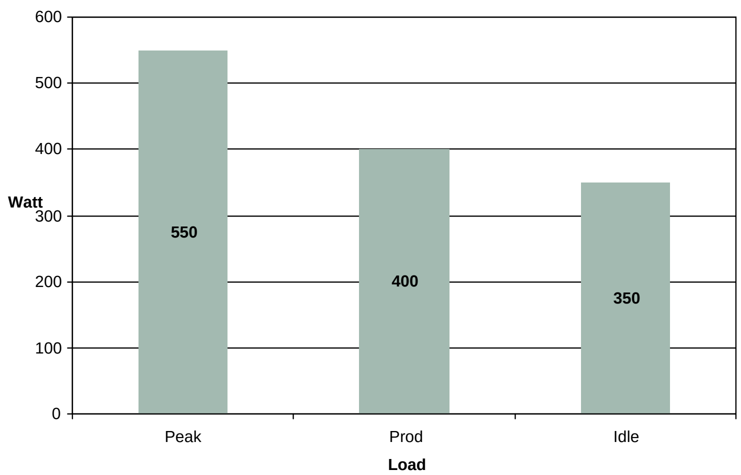

Figure 11: A server in idle mode consumes almost as much electricity as during production operation

2.3.2 Virtualization Reduces Energy Consumption

In the search for a way out of this dilemma many companies are opting for the enormous potential offered by virtualization technology of all kinds to increase system utilization.

A server which has a higher workload because it serves numerous virtual instances makes much more efficient use of the electricity supplied to it. For example, power supplies are much more efficient at high load than in idle mode. VMware is today's most popular server virtualization solution for standards-based servers. There are other software solutions such as Xen (Open Source) and Virtual Server from Microsoft. Fujitsu Siemens Computers even offers a combination of software-based and hardware-supported virtualization concepts in the mid-range and high-end sector on its PRIMERGY BladeFrame, BS2000/OSD, PRIMEQUEST, PRIMEPOWER and SPARC enterprise platforms. On mainframe platforms utilization rates in excess of 90% can be achieved so that they are among the most efficient platforms of all. Gartner estimates that in the next three to five years, virtualization technology will also boost utilization rates on standards-based servers from today's figure of between 10% and 20% to at least 50% to 60%[6].

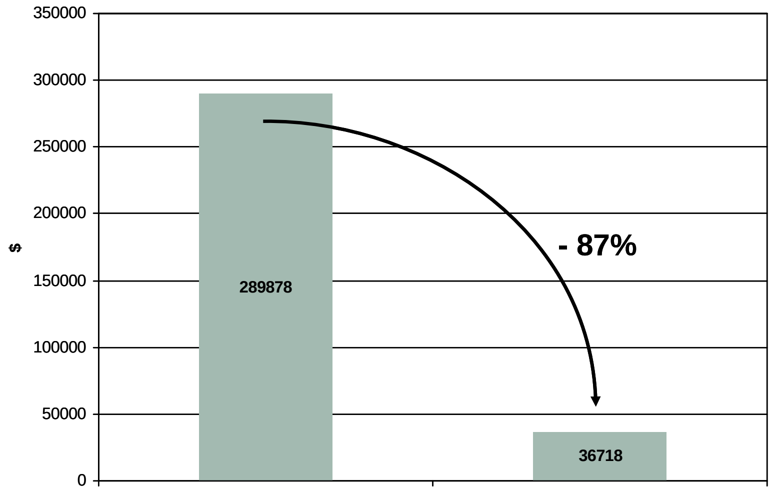

Figure 12: Yearly energy savings potential due to server virtualization with VMware for a consolidation rate of 12.5 : 1

2.4 Transition from a Static to a Dynamic IT Environment

Beyond the use of energy efficient components and the virtualization of servers, a further step to even greater energy efficiency is the deployment of dynamic infrastructures. By making IT operation more flexible, IT resources can be better matched to current needs. Wasted energy consumption by systems not currently needed, can be reduced even further. By reference to a terminal server environment, Figure 13 illustrates how use of the latest quad-core processor technology in conjunction with management tools that support more flexible operation can lower energy costs by 74.5%.

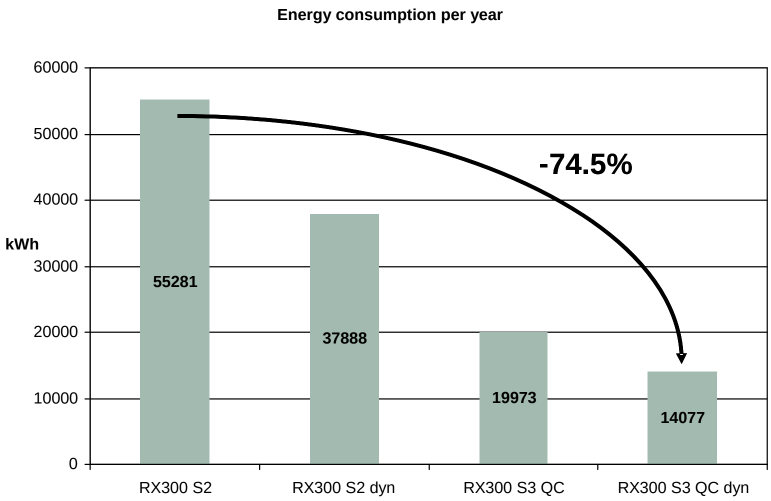

Figure 13: Energy savings through a combination of the latest processor technology and dynamic operation in a terminal server environment

2.4.1 A Step Ahead: The Dynamic Data Center Strategy of Fujitsu Siemens Computers

Migration of many smaller application islands, each consisting of production, backup and test systems, to an architecture based on a pool of similar IT resources substantially reduces the number of servers needed and thus makes a valuable contribution to large energy savings. Again, the use of virtualization technology ensures that services are not tied to a particular server so that each service can run on any server in the pool.

2.4.2 Dual Benefits for Customers – Dynamic Data Center Solutions Combine Blade Server and Virtualization Technology

Blade servers and virtualization technology target many common areas that make a significant contribution to enhanced energy efficiency in IT environments. Blade server architectures permit the shared use of components such as power supplies and fans and therefore deliver energy saving benefits on system level. Utilization of the individual components is improved and fewer components are needed to support IT operation.

Similarly, the deployment of virtualization technology on an infrastructure level results in fewer and better utilized systems. By combining both technologies in an integrated solution, customers are given the unique opportunity to profit doubly from a single implementation.

2.4.3 Flexible Energy Consumption in IT Systems through the Use of Dynamic IT Solutions

By integrating the management tools required to monitor the IT infrastructure into the management environment of application vendors, IT systems respond better to application resource needs and therefore allocate IT resources in a more flexible way.

- When peak load situations necessitate server farm scaling or if it is necessary to react to error situations, management tools ensure that additional servers are activated only when they are really needed and do not waste energy running in idle mode.
- IT administrators are also provided with support to optimize IT resource utilization. Servers not currently needed, can be switched off when workloads drop. In such scenarios, automation technology simplifies the implementation of energy efficient day/night operation.
- Dynamic Data Center solutions also feature an energy efficient N:1 redundancy concept. In this case, all that is needed is a standby server in a server farm. If a server fails, the standby server is activated automatically. If permitted by the Service Level Agreement, the standby server need not even be switched on so that no energy is wasted.
- Shared storage systems in a NAS/SAN environment ensure better utilization of disk capacity and therefore support energy efficient operation.

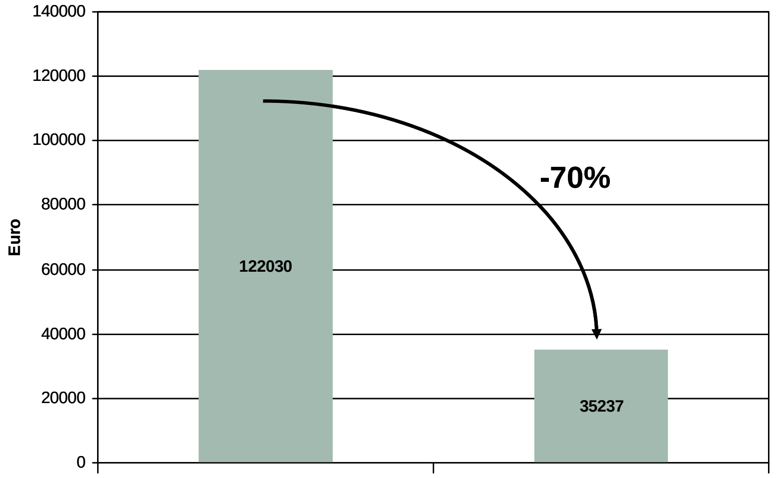

Figure 14: Energy cost savings through the use of the "FlexFrame for Oracle" Dynamic Data Center solution in a concrete customer scenario

3 Solutions to Boost Energy Efficiency in Cooling Infrastructures

Despite the considerably improved characteristics of servers with regard to power consumption, total power requirements to operate servers in data centers have risen dramatically over the past few years. The "power guzzlers" are not only servers and other components but also and in particular the physical infrastructure. The same amount of energy is consumed to supply power and cooling as to run the entire IT equipment. And all the electrical energy supplied is ultimately dissipated to the outside as heat. This means that each watt saved in terms of IT equipment counts double because one watt is also saved in terms of the infrastructure. It also means that there is an urgent need for action to improve infrastructure energy efficiency.

Most technologies and solutions needed for this purpose already exist but must be implemented in a systematic, coordinated way. However, there is enormous optimization potential into the future.

3.1 Optimization of Components and Systems

3.1.1 Classical Cooling Solutions

At present most data centers are air-cooled. Servers and other IT components are generally mounted in racks so that they draw in cool air at the front and expel warm air at the back. Provided that sufficient cooling air is available and heat dissipation by air through the computer room works, even high-performance computers such as blade servers can be air-cooled successfully. Today's server racks are designed so that the perforation quality of the doors (> 80% open perforation pattern) presents no appreciable flow resistance.
According to ASHRAE, as well as satisfying relative humidity requirements the air intake temperature for IT equipment should be between 20 and 25°C[7]; however, most servers from Fujitsu Siemens Computers today support a much wider temperature range between 5 and 35°C, without any drawback on the MTBF (Mean Time Between Failure).

The maximum cooling performance per rack is limited technically by the ability of the surrounding room to supply sufficient cooling air to the front of the racks and to direct the warm air at the back of the racks out of the room and to the circulating air cooling units usually located at the outer edge of the data center.

Figure 15 shows the best cooling setup for the cabinets in a cold aisle/hot aisle arrangement[8]. The cooling units force the cooling air under the raised floor, distribute it there and blow it through perforated tiles into the cold aisles on the air intake side of the cabinets.

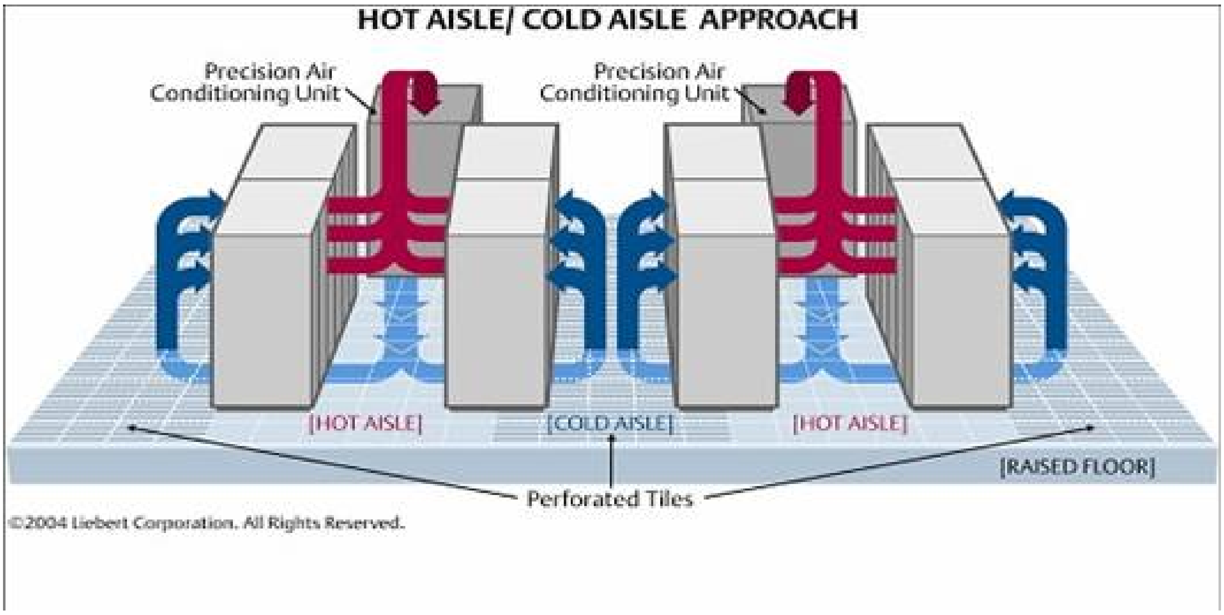

Figure 15: Cold aisle / hot aisle arrangement for optimal cooling

Other cabinet arrangements, some without raised floors, are often used, particularly in small but mainly in older data centers. However, the optimization approaches described here can also be used.

Assuming that cold aisle/hot aisle arrangements are carefully planned and constructed, approx. 3 to 5 kW cooling capacity per rack can be achieved. In older data centers these values are usually much lower at about 1 to 2 kW per rack. The cold aisles are then completely filled with cool air (operators generally want temperatures around 20°C) and the warm air is collected above the cabinets and returned to the cooling equipment.

If cooling requirements are higher, the flow of cooling air supplied is no longer sufficient and some warm air is drawn back over the cabinets to the cool side. Consequently, air temperature near the top of the cabinets may be unacceptably high.

To lower the temperature, the outlet temperature is often set to lower values in the range between about 15 and 18°C. This improves the situation in the area at the top of cabinets but consumes a lot of extra power.

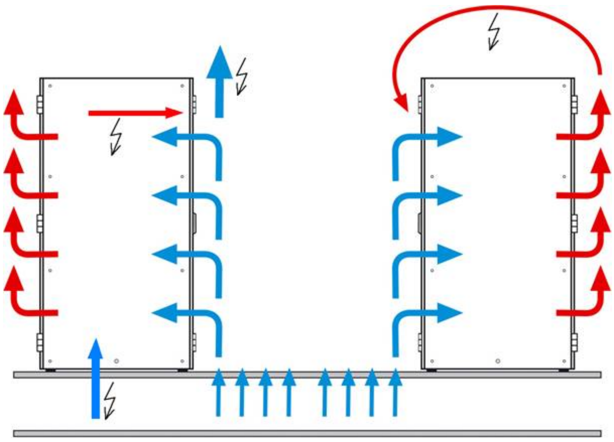

Figure 16: Warm air is drawn across the cabinets to the cold side

The reason for the extra cooling power needed lies in the cooling system outside the room.

Figure 17 shows a typical cooling system setup outside the computer room. In the circulating air equipment there is a heat exchanger that is supplied with cold water or refrigerant by chillers located outside the building. In the chillers a compression refrigeration machine raises the temperature of the cooling medium (about 10°C) to such an extent that it can be cooled again using ambient air. Even in central Europe this should function perfectly well, even at outside temperatures of up to 40°C; in other words, the cooling medium must be heated to approximately 45 – 50°C in order to ensure effective heat dissipation.
At lower ambient temperatures (below about 8°C in our example) the coolant can be cooled directly by the outside air without the need for a chiller ("free cooling").

Compression refrigeration machine drives consume relatively large amounts of electricity – usually about 30 – 40% of cooling capacity. The lower the desired temperature of the cooling medium, the higher is the actual amount of electricity needed.

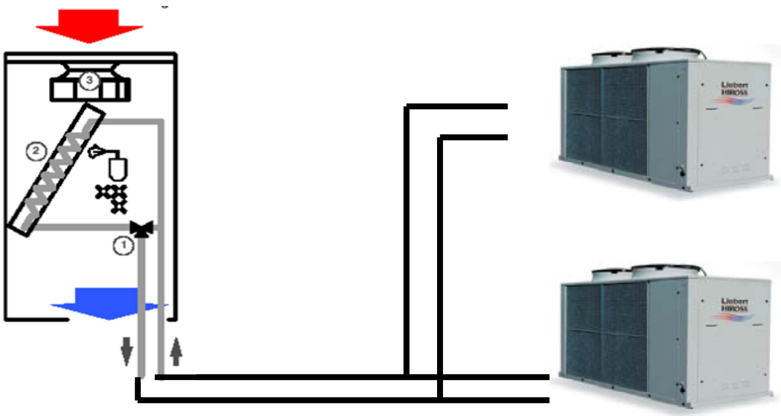

Figure 17: Typical cooling system setup outside the computer room

The power requirements to drive the cooling systems are classified as follows for the individual components.

High
Drive power for the cooling equipment in the chillers; the higher the desired temperature of the cooling medium and the lower the outside temperature, the lower the power requirements. The efficiency of the chillers rises with every degree of temperature, and the potential free-cooling periods (where there is no need to operate the chillers) become disproportionately longer.
What's also important are optimal component efficiency (compressor, electric drive) and energy-optimized control.

Medium
Drive power for the fans (in the circulating air equipment and in the chillers); the power is proportional to the air volume flow (m³/h) and to the difference between the air intake pressure and the pressure (Pa) in the raised floor.
Optimization of fan flow, the drive motors (preferably EC motors) and energy-optimized control also play a major role.

Low
Drive power of the fluid pumps in the coolant system; the power requirements are almost negligible. Nevertheless, the individual components, piping system (low flow resistance) and control system should also be optimized and pressure-controlled pumps must be used.

Consideration must also be given to the humidification and dehumidification of the room air, an aspect which occasionally plays a very significant role. However, due to the complexity of this issue, we won't deal with it in depth here.

Figure 18 shows a typical distribution of power needs for a well planned data center. The IT components need about 50% of the power, the cooling system about 25%, and the air circulating fans about 12%. The share needed by the power supply (UPS system, transformers) is relatively low at approx. 10% but can also be substantially optimized.

This produces the following clear energy optimization priorities:

1. IT components and system
2. Cooling system optimization by keeping all temperature levels as high as possible
3. Air circulation optimization by preventing cooling air losses. The recirculated air volume should be kept as low as possible.

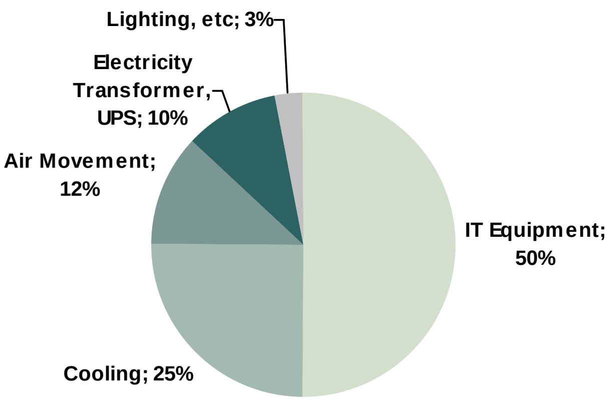

Figure 18: Distribution of power needs for a carefully planned data center

It should also be ensured that the entire cooling infrastructure is not over-dimensioned. The actual electrical power consumed by the servers – and therefore the cooling load – is approx. 30 – 40% below the values specified on the rating plates.

3.1.2 Cooling Solutions for Higher Cooling Performance

Due to the enormous growth in computing power over the last few years and despite the greatly improved energy efficiency of servers (see section 2.2), the packaging density of electronic components has risen in parallel to such an extent that the heat load of individual cabinets has increased from about 1-2 kW (2000) to between 10 and 20 kW or even more (2007).
These heat loads exceed the capacity of classical precision air conditioning equipment (see section 3.1.1). Once the desired cooling capacity tops 5 kW per rack, classical air conditioning units are generally no longer able to dissipate the heat produced in the racks. A number of approaches with open cooling architecture (enclosed aisles, extra cooling devices near or in the racks) can be adopted to prevent hot spots in data centers. These function well but are not discussed here. But intelligent solutions are also required at cabinet level.

3.1.2.1 The Right Choice of Rack Counts: PRIMECENTER Racks from Fujitsu Siemens Computers and Knürr

At cabinet level it is crucial to ensure minimum resistance to air flow. Air-cooled PRIMECENTER racks with low air-resistance doors, increased cabinet width (700 mm instead of 600 mm) and structured cabling deliver maximum capacities of up to 10 kW if the air flow characteristics in the computer room are suitably optimized. They are therefore able to dissipate the waste heat of fully equipped server racks to the outside.

3.1.2.2 PRIMECENTER Liquid Cooling Racks with Cool-safe™ Technology / Knürr CoolTherm™

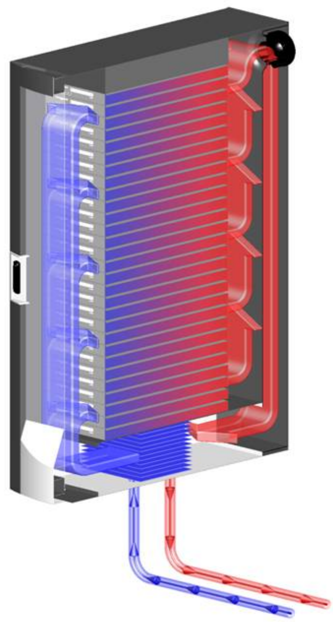

Only closed, directly cooled racks function reliably and economically at very high cooling capacities of above 20 kW per rack. However, due to their favorable characteristics, products of this kind are also used for cooling capacities upwards of 10 kW.
Knürr CoolTherm™ and PRIMECENTER Liquid Cooling Racks with Cool-safe™ technology transfer almost all heat generated into the cooling water. Consequently, minimum heat is dissipated into the room and a substantial burden is removed from the data center's air conditioning infrastructure. This means that even future server generations can be operated reliably and that the capacities of conventional air cooling systems can be exceeded.

The greater efficiency of cooling infrastructures that use water (water carries heat 3500 times better than air) permits higher packaging densities not only within the water-cooled PRIMECENTER Liquid Cooling Racks; data centers too can be filled with racks and servers so that they are "heat-neutral", regardless of the cooling capacity. Packaging density then depends only on the geometric properties in the data center and the capacity of the chillers. PRIMECENTER Liquid Cooling Racks permit space savings greater than a factor of 3 (this equates to a tripling of data center rack space) and offer substantial savings potential.

The picture on the right shows a cabinet of this kind in which an air-water heat exchanger is incorporated right at the bottom to rule out any risk of leakage onto the IT components.
The rack is closed on the air side and the cooling air is routed in a closed circuit at a fixed nominal temperature in front of the air intake of the IT components. As a result of the strict separation of the cold and hot sides, the cooling air temperature is the same from top to bottom and, due to the circulation principle, there are no heat nests, even at maximum cooling capacities.
The strict end-to-end routing of the cooling air, the short air paths and the built-in fans deliver very high cooling capacities. The fans themselves are power-optimized (fan wheels, EC motors), and the heat exchanger is generously dimensioned in terms of heat transfer and flow resistance. The cooling capacity and the fan speed are regulated steplessly so that the exact cooling capacity currently needed is made available.
In this concept the air circulation power requirements are already energy-optimized so that no special measures are needed outside the cabinets.

As concerns cooling system optimization, there is considerable potential as the cooling water intake temperature is much higher than in the "classical" solution.

3.2 Static System Optimization

It is primarily the careful routing of cooling air that offers optimization potential in the classical solution. The sealing of all leaks in the raised floor, the careful selection and positioning of the floor tiles and the separation of cold and warm air within the racks are the basis for power-optimized operation.
It is essential to carefully balance the permissible heat loads (and therefore the IT equipment levels) in each cabinet in order to prevent the backflow of hot air. No circulating air cooling system can be designed to ensure that exactly the same flow of cooling air is available in each cabinet or that hot air is extracted equally well at all points. In other words, each cabinet has its own individual power budget which must be taken into account when equipping the cabinet with IT components.

This ensures that only the amount of air that is actually needed is circulated. In addition, the fans in the circulating air cooling devices must be speed-controlled by means of reference temperature sensors in the room. This helps minimize air circulation power requirements.

Added potential is provided by the option of expelling air from the raised floor at higher temperatures. A temperature of between 20 – 22°C should be sufficient. As illustrated above, this would allow the entire cooling system to run on considerably less electrical power (better chiller efficiency, longer free-cooling periods).

In a closed system (water-cooled server cabinets) air circulation is already optimized for all operating conditions. The possible intake temperatures are higher than in optimized classical cooling solutions. Still more potential can be exploited by equipping cabinets so that they remain below the maximum cooling load. For example, in a 25-kW cabinet fitted to a level of approx. 15 kW, a cooling water temperature of approx. 20°C is sufficient to keep the internal cooling air temperature constant at 25°C over the entire height of the cabinet.

Depending on which servers are used, the cooling air temperature can be increased still further. A continuous temperature of 25°C generally presents no problems. And some servers are designed for still higher continuous temperatures up to around 30°C. It is therefore conceivable, at least in closed cabinets, to fully dispense with chillers in the cooling system and to use wet cooling towers the whole year round. This would totally eliminate the major power guzzlers.

It is a particularly beneficial that high power density (due to high electronic packaging density) coupled with systematic server utilization (by means of virtualization) gives rise to especially energy efficient cooling solutions. As compared with conventional cooling solutions and depending on the specific situation, up to 40% energy savings can be made for the cooling system and for air circulation.

3.3 Dynamic System Optimization

As not only current power needs – and therefore the heat to be dissipated – but also the outside temperature and humidity fluctuate continuously, there are still great savings to be made in constantly matching the system to current operating conditions. The objective is to keep the cooling water temperature as high as possible and the circulated air flow as low as possible. There is hardly anything to be gained by adjusting cooling medium flow as the extra capital investment often doesn't pay for itself.
The above purpose is achieved mainly by evaluating the temperature sensors located in the room or in the cabinets in the physical infrastructure, possibly also by measuring the current power consumption of the servers and occasionally by measuring pressure differences and flow rates. Dynamic physical infrastructures achieve special significance when used in dynamic IT environments. Only in this way can potential energy savings be fully exploited.
Building on this development in IT and infrastructure technologies, dissipated heat is set to play an important role in the future; to heat offices for example. Although dissipated heat cannot be used directly even given higher cooling water temperatures of about 20°C, from this level on the temperature can be raised to the requisite 50 - 55°C using a heat pump with a very good performance rating.

4 Data Center Quality Services

4.1 More than 20 Years of Experience in Data Center Construction

The Data Center Quality Services of Fujitsu Siemens Computers IT Product Services are targeted at all data center areas (air conditioning, power supplies, fire prevention/protection, water supply, security facilities) in the form of a modular service portfolio. Numerous modules are available which, when used in combination, support a holistic approach to the implementation of data center infrastructures that are safe and economical. Services cover not only consulting, risk analysis, conceptual and solution design but also project management and custom solution implementation.

Highest priority is given to air conditioning and electrical engineering aspects in the design of energy efficient solutions. Potential energy savings are examined and calculated when selecting air conditioning and cooling systems and when designing energy transport and conversion.

With more than 20 years of experience in data center construction, Fujitsu Siemens Computers is the professional partner of choice for data center operators and offers all products, solutions and services from a single source.

5 Outlook

Further optimization potential in terms of IT and infrastructure is there for the taking. Many reputable companies have therefore recently joined forces in open, non-profit organizations such as "Green Grid" and the "Climate Savers Initiative" to develop new ways of saving power in data centers. "Green Grid" focuses on optimizing data center infrastructures and developing "best practices". For example, it is intended to establish key figures and measuring methods for power consumption and waste heat generation. By contrast, the "Climate Savers Initiative" concentrates more on the energy efficiency of individual devices (PCs and servers). The goal of this initiative is to halve PC power consumption by 2010 and therefore to reduce carbon dioxide emissions by 54 million tons per year. At the same time, the "SPEC Power-Performance Committee" is currently in the process of developing a first benchmark to evaluate the energy efficiency of servers. This would, for the first time, enable fair and consistent statements to be made on server energy consumption.

Significant improvements will be necessary on all levels to achieve the ambitious targets of these initiatives. For example, on component level all vendors are working on optimizing the energy consumption of their products, above all for low to medium utilization rates. In future, management tools will be able to measure and display energy consumption. Energy consumption will be used as an additional key metric in the management of IT environments so that IT systems can be better matched to actual power needs.

What's more, in future the intelligent interaction of IT systems with physical infrastructures will generate further energy savings in the double-digit percentage range. This will call for close cooperation between IT and infrastructure manufacturers, as successfully practiced by Fujitsu Siemens Computers and Knürr for many years.

6 List of References

1. "Meeting the Data Center Power and Cooling Challenge", Gartner, 19 December 2006
2. "Strompreise für private Haushalte und industrielle Verbraucher zum 1. Januar 2006", Eurostat - Umwelt und Energie — Statistik kurz gefasst, 11 November 2006
3. "Datacom Equipment Power Trends and Cooling Applications", ASHRAE, 2005
4. "Data Center Energy Efficiency and Productivity", The Uptime Institute, 2007
5. "Estimating total power consumption by servers in the U.S: and the world", Jonathan G. Koomey, February 2007
6. "Important Power, Cooling and Green IT Concerns", Gartner, 23 January 2007
7. "Thermal Guidelines for Data Processing Environments", ASHRAE, 2004
8. "Design Considerations for Datacom Equipment", ASHRAE, 2005
9. "Liquid Cooling Guidelines for Datacom Equipment Centers", ASHRAE, 2006
10. "Managing the Data Center of the Future – Aufbau und Betrieb von Service-orientierten Infrastrukturen", Thomas Meyer, IDC's Dynamic IT Conference, Munich, July 2007

Fujitsu Siemens Computers
Holger Kordowske
Infrastructure Products Marketing

Knürr
Dr. Peter Koch
SVP Products Technology & Quality

Fujitsu Siemens Computers Internet:
www.fujitsu-siemens.de/products/standard_servers

Knürr Internet:
www.knuerr.com/whitepapers