Federal Ministry
for Economic Affairs
and Energy

IT2Green.de

Energy-Efficient ICT in Practice

Planning and Implementation of GreenIT Measures in Data Centres and the Office

Imprint

Publisher
Federal Ministry for Economic Affairs and Energy (BMWi)
Public Relations
10115 Berlin
Germany
www.bmwi.de

Text and editing
Fraunhofer Institute for Reliability and Microintegration (IZM)
Fraunhofer Institute for Systems and Innovation Research (ISI)
VDI/VDE Innovation + Technik GmbH

Design and production
VDI/VDE Innovation + Technik GmbH

As of
November 2014

Printing
Silber Druck oHG, Niestetal

Picture credits
Vishnukumar/Maxim_Kazim, Fotolia.com, VDI/VDE-IT (Title);
VDI/VDE-IT (p. 4, 6); Tiberius Gracchus, Fotolia.com (p. 8);
Iceteaimages, Fotolia.com (p. 11); Eisenhans, Fotolia.com (p. 14);
Axel Springer SE (p 17); Universität Paderborn (p. 19);
Sashkin, Fotolia.com (p. 20); .shock, Fotolia.com (p. 22);
Minerva Studio, Fotolia.com (p. 30), OFFIS p. 33, 34);
yanlev, Fotolia.com (p. 36); agsandrew, Fotolia.com (p. 40);
gena96, Fotolia.com (p. 44)

Content

Continuous improvement energy-efficient ICT in offices workplace concepts knowing technology and applications Thin Clients Cloud Power Management user acceptance presets legislative regulations monitor mobile devices load profiles power analysis usage analysis data protection user-driven IT-resources Client Probe energy-optimised WLAN turn off on low load control by Resource-on-Demand-algorithms energy-efficient ICT in data centres hardware procurement configuration life cycle assessment miniaturisation modern hardware holistic optimisation of energy and resource use components system configuration thermal anagement IT2 Green measurement of energy consumption data collection beforehand air-conditioning power supply detailed planning troubleshooting plausibility checks IPMI-values key performance indicators evaluation of energy efficiency Green application systems load forecasts load and power management methods energy-oriented migration to the Cloud energetic control of applications TOSCA Green Policies sustainable technology design energy management system ISO 50001 holistic environmental assessment Green Software methodical obstacles limitation of data traffic programming languages modular software efficient utilisation of memory

1. Introduction

GreenIT – A balancing act

GreenIT, or green computing, refers to the practice of designing and operating IT systems used in telecommunications, data centres and office environments in a way that conserves resources and makes efficient use of energy. Rather than dealing with the energy savings that can be achieved through ICT ("green through IT"), this guide looks at ways to reduce the energy needed to actually run the ICT systems themselves ("green in IT"). Rising electricity costs, increasingly strict environmental legislation and customers' demands for green products are the main reasons why GreenIT has become a hot topic in recent years. At the same time, there is a growing awareness in many companies, public authorities and other organisations of the need to conserve resources and energy in business operations, and of the cost-saving potential such action brings.

Energy efficiency and quality of service must be addressed and implemented collectively.

Many strategies to increase energy efficiency in data centres and the office environment are already known, with the literature, internet blogs, conferences, and service providers offering comprehensive advice. Simple measures can usually be implemented very quickly. Greater energy savings achieved through better system design, sizing and utilisation and the more dynamic use of IT systems, on the other hand, often require more complex planning and involve higher initial costs. Further to this, there are conflicting aims when it comes to using IT. As the principle approach to saving energy consists of the needs-based provision of IT resources, which is mostly only achieved by the complete or partial shutdown of active IT elements, availability often suffers. Compliance with the agreed quality of service is essential for IT administrators, however. The IT system should be "stable" with no "loss of convenience" for the users.

In short: energy efficiency and quality of service must be addressed and implemented collectively.

Setting up an IT system that saves energy and conserves resources will always be a balancing act for IT administrators in business enterprises and public authorities if clear priorities are not set. For this reason, it is important to take a "whole-of-the-system" approach to GreenIT and identify and articulate the positive business aspects. Furthermore, it should be noted that businesses choosing to green their IT will need to take a very close look at their own IT infrastructures and business processes, which itself offers many advantages.

Benefiting from GreenIT

GreenIT drives innovation and learning processes in business enterprises and quickly demonstrates the importance of cross-departmental collaboration. This fact is central to the sustainability of GreenIT measures. Many workers can use this opportunity to make a name for themselves in the company, boosting their image both in-house and towards the public. This guide is primarily geared towards IT administrators in small and medium-sized enterprises, public authorities and other institutions that independently run ICT equipment in offices and IT systems in server rooms and data centres. However, decision-makers and users who define and demand IT resource requirements can also benefit from this guide as it contains key principles and practical fundamental knowledge for the energy-efficient design and use of IT systems. The guide aims to provide reasons for implementing green computing in businesses, and encourage individuals and businesses alike to continuously work towards the goal of saving energy and conserving resources.

GreenIT is like exercising – you need to keep at it to stay fit.

IT administrators face two central challenges, not only in the context of GreenIT. First of all, the IT resources available must be provided to the users efficiently and at the defined level of quality. To this end, continuous improvement processes must be established and implemented successfully. Secondly, as IT resource needs change dynamically, they must be analysed regularly to make important planning decisions over the medium term.

When developing new IT strategies, the priority is to make a good forecast of IT resource needs and compare these needs against new technical possibilities. In this connection, the persistently high level of technological dynamism coupled with an intensity of use that continues to grow makes it

The IT2Green projects present their solutions for energy-efficient ICT at CeBIT 2014 in Hannover.

necessary to continuously monitor the state of existing IT systems and periodically analyse user needs and new technological options.

> The IT2Green projects clearly demonstrate how difficult robust power consumption measurements can be.

General principles

This guide presents a fundamental methodology framework on the basis of which GreenIT projects can be implemented in businesses and public authorities. The content and practical examples are primarily derived from the research results of the IT2Green programme. Rather than attempting to provide the technical details of every strategy and action to improve energy efficiency, this brochure seeks to explain the causal relationships between the development and use of technology and the consumption of energy and resources. Furthermore, this guide aims to provide a solid basis for measuring load and power consumption in data centres and the office environment.

Questions concerning our own energy efficiency, what technologies and measures deliver which particular savings, and the best approach to check the continuity of these effects, in particular, often prove difficult to answer. However, it is precisely these topics and questions that are central to the success of green information and communication technology in practice.

Structure of the guide

This guide consists of two parts:

- "Energy-efficient ICT in the office" (p. 9)
- "Energy-efficient ICT in data centres" (p. 21)

The first part deals with the energy-efficient operation of end user devices in office environments and general workplace concepts, looks at configurations for energy-saving power management and an energy-efficient WLAN, and presents power demand analyses of IT resources on the basis of real user needs.

The second part of the guide discusses energy-related issues in data centres, including technology procurement, configurations and integrated energy and resource optimisation. Ways to measure energy consumption and energy efficiency indicators are also explained. Furthermore, this section also presents strategies for the efficient control of data centres via a green control station and solutions for migration to the cloud.

Gateway to specific information

This brochure can only touch a limited number of topics and illustrate a few sample solutions. In reality the situation is far more extensive and diverse. Therefore, this guide cannot claim to be a single reference manual in itself, but is rather a gateway to more specific information. Details on more in-depth literature, interesting points of contact and organisations are provided at many points in the brochure. In addition to the references given in the text, a bibliography and a list of helpful websites dealing with GreenIT are also provided.

Continuous improvement process

Given that IT in businesses is constantly changing at a faster pace and increasingly higher demands are put on the availability of IT, GreenIT should be understood as a continuous improvement process. With his PDCA cycle, American scientist William Edwards Deming (1900–1993) formulated the basis for continuous improvement processes. Standing for "Plan – Do – Check – Act", PDCA describes a policy of action which has the broadest take-up in quality management and environmental management today. For the purpose of implementing GreenIT, the concept of this methodological approach is adopted and transferred to a practice-related project structure.

1. Status quo: In the first step, the actual situation must be captured in the form of data and measurements, analysed and presented in a transparent manner. Several questions need to be answered in this connection: Which IT systems are used, who uses them, how are they used, and what are the associated energy requirements? What causal relationships or particular characteristics can be identified?

2. Conceptual design: The second step involves clearly defining the requirements which the IT system should meet in a specific time frame. Energy-efficiency goals must be weighed against quality of service at this point. The challenge lies in identifying technical parameters and product options in line with individual user needs.

3. Implementation: Fundamental decisions must be made at this stage of the process. What is important here is to integrate all the involved parties and the different sub-sections into a goal-oriented process. GreenIT measures can only be planned and implemented successfully by an interdisciplinary team. Support from all sides must be guaranteed. In this respect it is helpful to quantify the business-specific and individual benefits of GreenIT as precisely as possible.

4.Usage optimisation: A continuous process of checks and improvements begins with the implementation of GreenIT measures. Conditions change in practical operations and individual problems crop up. Developments brought about by software and applications must be constantly tracked and integrated into a monitoring strategy. The steady increase in the capacity utilisation and/or needs-based use of existing IT resources and the customisation of a functional energy management system are the primary objectives of the last step.

2. Energy-efficient ICT in the office

For a number of years, ICT in the office has been undergoing a substantial change that also concerns energy and resource consumption. Computing technology is becoming more compact and mobile, a fact which has a positive effect on the development of resource-conserving technologies. However, while the energy efficiency of individual devices is increasing, software applications such as simulations or graphic programmes are becoming more compute-intensive and the number of devices in use is on the up and up. For example more and more workstations are fitted with multiple monitors, driving up total power consumption despite the lower energy demand of the individual devices. Added to this, many products have a relatively short life span, with the result that a product's environmental impact depends more on the material and energy used in the production phase. Therefore, efficiency gains in the utilisation phase of short-lived products lose value in relative terms.

The continuous availability of high-bandwidth internet makes it possible to transfer computing operations to the cloud. The end user devices can therefore be slimmer (thin clients) and also use less energy. Whether this actually has a positive impact on the overall environmental footprint is primarily contingent upon whether the data centres use load-adaptive network technology, which is rarely the case in practice to date.

The information below provides a brief overview of the next few sections, that present important aspects of energy efficiency at work in more detail.

Workplace concepts

The "Workplace concepts" section presents technical trends that create the framework for energy-saving ICT workplace designs. The focus here is on the practical implementation of new strategies and the improvement in practice of heterogeneous ICT systems that have evolved over time. Particular attention is given here to interaction with the ICT user in the office. Investigating, analysing and optimising the interaction of users with the IT resources available to them is central to boosting energy efficiency. Sample methods and approaches are presented using specific results from the IT2Green projects.

Power management and user acceptance

Many new end user devices already offer comprehensive power management options. The "Power management and user acceptance" section describes challenges posed when using effective power management settings. The overriding priority here is to find an appropriate balance between user convenience and saving energy. In addition, holistic sensor and control systems are presented that use presence detection to control both the ICT systems and the extended building infrastructure, such as the heating and lighting. These end-to-end software solutions can tap additional potential for energy conservation. To be economically viable, however, this would require far more detailed planning, and a minimum number of users in most cases.

Power demand analysis

As the user and the user's role define what IT resources are needed in the office, the user's conduct also has a bearing on energy efficiency. The "Power demand analysis" section looks at the detailed measurement and analysis of power consumption rates related to end users. In this connection, it is important to never lose sight of data protection during any energy monitoring activity. The fundamental requirements that need to be considered in this respect are described in this chapter.

User controls IT resources

In optimisation processes, it is very difficult to determine the needs of the user, the perceived speed and quality of the IT services used, and such information is taken into consideration all too seldom as a control parameter. The "GreenIT Cockpit" project has found a way to automatically record control parameters and use them to optimise the energy and response times of the data centre.

Energy-optimised WLAN

In addition to "visible" equipment in the office, the WLAN also offers certain potential for optimisation. One example of successful optimisation is presented by the University of Paderborn as part of the GreenPAD" IT2Green project". An example of measurement and analysis in practice is provided here. Hard facts and data present real-scale systems and potential for energy savings, serving as an excellent starting point for appropriate measures.

2.1. Workplace concepts

Knowing the technology and the applications

The amount of energy office ICT actually needs depends heavily on the application-related ICT strategy and the use of appropriately configured products. The more precise information is known about the application requirements and the more these can be translated to software and hardware specifications, the better these requirements can be matched to the products available on the market. In other words, it is important to be fully aware of user needs and technical possibilities, which is only possible through continuous monitoring.

ICT equipment on the market is becoming ever more specialised. The fact that performance data for computing power, storage capacity and connectivity are just as varied as the applications themselves is testimony to this. Businesses would therefore be advised to take software needs and software options as the basis for developing an ICT strategy, and define the fundamental architecture and performance parameters using the current possibilities afforded by new technologies and products.

Modern equipment is energy efficient

ICT equipment nowadays has a very high energy efficiency. Product-related environmental legislation, energy labelling, and the implementing measures of the European Ecodesign Directive have put a certain amount of pressure on manufacturers. However, the technical problem of potential increases in energy density and the practical need for mobility prove to be more important for improving energy-efficiency.

The miniaturisation of semiconductor components has been an ongoing trend for decades. The number of transistors per unit of area roughly doubles every two years. To ensure that the energy density, and therefore the waste heat, does not increase at the same rate, the voltage is reduced. Accordingly the progressive increase in performance on the one hand is accompanied by constant or lower power needs on the other hand. The table below shows that the average power consumption of computers in sleep and idle mode dropped by around 20 percent over a two-year period.

Energy Star Programme		Sleep (W)		Idle (W)	
Product	Class	2010/2011	2012/2013	2010/2011	2012/2013
Desktop PC	B	2,1	1,9	36,1	31,8
	C	2,5	1,8	46,9	37,9
Notebooks	B	1,3	0,7	13,3	8,2
	C	1,9	1,8	23,3	18,8

The demand for mobility is the second factor driving energy efficiency. With battery capacity limited and technological advances in the area comparatively slow, sophisticated power-saving technologies are needed to deliver functional, sophisticated mobile devices. Many innovations from laptop and mobile phone development have become standard, increasing the energy efficiency of both computers and monitors – particularly from the change to LED backlighting.

Thanks to this positive trend in the development of end user devices, businesses purchasing and using ICT systems can continue to benefit from large optimisation potentials. A product's specific energy usage should always be considered when making a purchase. This applies to both the general power consumption of the hardware and the choice of 0-Watt power supply units. At the same time, when purchasing new hardware it is also important to bear in mind the amount of computing power that is needed by the particular users and whether there is any excess capacity.

After almost three decades of continuous performance enhancement in conventional PCs (fat clients), modern technologies and products use the powerful network options for thinner cloud-based workplace concepts (thin clients). In between, there is a multitude of hybrid options that take into consideration workplace-specific requirements, such as a high level of mobility or security, for instance.

The more accurate the knowledge of the users' software requirements, the better these requirements can be matched to the IT products available on the market.

Thin or not to(o) thin

A detailed needs analysis is required to determine which workplace concept and hardware are appropriate. In addition to taking stock of the hardware currently installed, the following questions should also be answered:

- Which users and how many users need computing capacity and how much is needed altogether?
- Which users are mobile workers and to what extent?
- Is server-based computing already used for individual applications, such as CPU-intensive simulations?

Based on the demands determined, a thin client solution can be worthwhile provided that individual irregular applications require a high level of computing capacity and there are few mobile workers. Then, efficiency gains can be expected compared to when workers are permanently equipped with high-performance computers. The energy demand of individual workstations is far lower for thin clients. However, the resources saved are sometimes needed elsewhere. Apart from server operation, it is also important to ensure appropriate infrastructure and sufficient network capacity to minimise the likelihood of failure from overload.

If, on the other hand, workers tend to use common Office applications, mini-PCs can also offer energy advantages over standard desktops or server-based thin clients. Where employees are often on the road and need to be mobile – taking frequent business trips, for instance – the best alternative could be to provide workers with laptops. Run with a docking station and external monitor, laptops are far more efficient than conventional desktop PCs. Apart from focussing solely on energy usage, however, other factors also need to be considered, such as costs, maintainability, user-friendliness and security aspects. In addition, it is also essential to clarify the complexity of a migration to another workplace concept and the advantages and disadvantages this would entail.

Data security is a critical aspect when making the change to thin clients. Given that resources are shared, businesses run the risk of unauthorised access to larger volumes of data than would be possible in configurations with autonomous computers if inadequate security measures are in place. At the same time, a thin client environment minimises the risk of security mistakes by users negatively impacting server operation as safety-related settings generally cannot be modified by users on well-monitored servers.

> If many server-based applications are already used, the switch to thin clients is often easier on account of the existing network and hardware infrastructure and can deliver substantial savings.

In addition to data security, availability is also an important factor that must be considered when using thin clients. If individual computers fail – be that thin clients or desktops – this does not pose a problem for staff's work. If the network is unavailable due to technical complications, it would still often be possible to continue working on local computers. Thin clients, on the other hand, lose their entire functionality in such situations as they lack their own dedicated hardware.

The Cloud

Businesses running their own servers incur high fixed costs and the associated risk is also high. As an alternative, it is possible to transfer server operations to a cloud. The services rented in this way can also work out cheaper during ongoing operation as the costs of excess capacity and emergency power generators are shared among the users. Effects of scale also mean that cloud solutions can offer advantages from an energy perspective.[1] By choosing the right size data centre or dedicated solutions, availability is ensured and there is even protection against system failure.

However, these solutions also present a conflict of interest with data protection, as the data could be stored in regions not covered by the necessary data protection regulations.

Following the usual procedure to send data packages in the cheapest way possible, such regions can also act as transit states. If the data are poorly encrypted or unencrypted when transmitted, this opens the door to data espionage. Apart from considerable damage to the company's public image, this can also result in compensation claims for which provisions must be made or additional insurance policies concluded. The substantial financial loss that might occur must be compared against the savings made. Therefore, more and more consideration is given to the topic of legal certainty in technology programmes, such as "Trusted Cloud".[2]

2.2. Power management and user acceptance

When it comes to reducing the energy consumption of ICT equipment, the choice of general workplace concept and the devices to be used, as described above, is the most important factor. Building on this, numerous specific optimisations can then be made during the implementation stage. Modern ICT devices have extensive, ready-integrated power management options. The default power management settings should be checked during implementation and changed where necessary to suit the specific requirements of the workplace.

General and individual default settings

As a general rule, the default settings for standby mode on PCs, laptops and monitors should be as short as possible to fully avail of the energy saving potential. The very fast reactivation times of just a few seconds from standby to operating mode means that user acceptance is not an issue. Problems can occur, however, if security settings require long reactivation and log-in times (e.g. by co-ordinating security with the network), resulting in delays in the workflow. In such cases it is important to first of all determine whether the security checks after every reactivation are necessary or whether they could be pared down to once every 24 hours, for instance. This approach could fully tap the potential of power management without holding up users. If this is not possible, administrators should sit down with the users to agree a slightly longer time before the system switches to standby mode.

The optimum solution is a combination of time-controlled power management, a standby switch on the keyboard and awareness and understanding among staff of the significance of energy use and standby modes.

1 Eric Masanet, Arman Shehabi, Lavanya Ramakrishnan, Jiaqi Liang, Xiaohui Ma, Benjamin Walker, Valerie Hendrix, Pradeep Mantha: The Energy Efficiency Potential of Cloud-based Software: A U.S. Case Study, Lawrence Berkeley National Laboratory, 2013, http://crd.lbl.gov/assets/pubs_presos/ACS/cloud_efficiency_study.pdf

2 "Trusted Cloud" technology programme of the Federal Ministry for Economic Affairs and Energy: www.trusted-cloud.de

During implementation, it is also essential to check whether there are special user groups for which the preset power management settings and the resulting changeover to the standby mode would cancel an active application. Typical examples here would be programmes with database applications or simulations. It is advisable to test the compatibility of the power management setting of the workstation computer with software and IT architectures used in the firm.

Standby switches can also be used to rapidly set workstations to the energy saving mode. On leaving the workstation, the user can press the button on the PC or keyboard to quickly and easily set the computer to standby, and simply needs to press the same button again to resume work at the PC later on. As they are easier to reach, switches on a keyboard are used more frequently than switches on a PC. This approach has the advantage that no time is lost until automatic standby, but it does require the active intervention of the user.

Therefore the optimum solution is a combination of time-controlled power management, a standby switch on the keyboard and awareness and understanding among staff of the significance of energy use and standby modes. When configuring the power management default settings, there should always be agreement on which energy saving mode should be used.

Legal regulations regarding sleep mode

In the field of PCs, energy saving modes are classified according to the Standard Advanced Configuration and Power Interface (ACPI). ACPI system level S3 – suspend to RAM – is the common standby mode used. Nowadays, reactivating the system from this mode is only a matter of seconds.

In accordance with the new EU regulation 617/2013 implementing Directive 2009/125/EC of the European Parliament with regard to ecodesign requirements for computers, the standby electric power demand for desktop computers and notebook computers may not exceed 5 Watt and 3 Watt respectively with effect from 1 July 2014. An additional allowance of 0.7 Watt can be applied for a wake-on-LAN (WOL) function. Desktop computers and integrated desktop computers with an idle state power demand less than or equal to 10 Watt are not required to have a discrete system sleep mode. Power demand in the off mode may not exceed 1 Watt for all products. There are no legal specifications regarding default power management settings but a setting in the range of 10 to 20 minutes is usually recommended.

Monitors

Strict power management requirements should also apply to monitors in addition to workstation computers. While modern LED-LCD and OLED displays have very good power efficiency ratings, the total power demand of monitors is quite significant given that screens are getting bigger and bigger and the growing trend to have a second monitor installed. As monitors can be reactivated immediately, their power management settings can be very strict. Short intervals of between 2 and 10 minutes should normally apply for monitors. Having the screen dim briefly before going to standby instead of abruptly turning black can increase user acceptance for monitor standby mode, as this gives users a response window and maintains convenience of use.

Mobile devices

In addition to direct energy savings, power management on laptops can also contribute greatly to battery life. For this reason, it is important to check whether the charging settings can be changed to help prolong the lifetime of the battery. To this end, very short charging cycles should be avoided and the battery should not always be kept fully charged. In addition, power management options can be configured to specify when battery charging should begin (e. g. only when battery is less than 80 percent charged). The needs of the user must also be taken into consideration here.

Load profiles – as different as night and day

To optimise power management, it can be useful to analyse the load profiles of different (types of) employees. Load profiles can differ significantly:

- Periods of absence often vary greatly between workers and managerial staff (e. g. many meetings)
- Different areas of responsibility sometimes demand different hardware or computing capacity with different load profiles and power management requirements

A precise analysis of the times of usage enables considerable energy savings.

Adaptive systems that are based on an extended infrastructure (presence detection; ICT systems, lighting and heating controlled by the same software) can also tap additional potential for saving energy but require detailed planning and a minimum number of specific users in most cases for them to be economically viable. This would require a detailed analysis of the power demand, an example of which is presented in the next section.

2.3. Power demand analysis

While modern devices already have dedicated power management systems, not all potential to save energy is utilised to the full. This can be attributable both to user acceptance and to the organisation of working times. The reactivation time required can be a criterion for opting against the economical use of resources. And although the reactivation time of modern computers is very short, even a tiny delay can often be objectionable. If the effort needed to reactivate a system is too big compared to the brief break from work, the power saving mode is generally not enabled and time-controlled power management is not triggered (at this stage), with the result that the computer remains switched on.

Frequent breaks between active use

The IT2Green project "pinta" investigated how long and often users were absent. The absent times of twelve workers from different departments were studied in a trial of 139 days. Ten of the workers had average absent times of more than 20 percent of the gross work time. Two absences of less than 5 minutes, one up to 15 minutes and two more than 15 minutes were measured on average. The last two measurements also included the lunch break. The time-controlled power management system of the workstations was not triggered in the shorter periods of absence, resulting in a marked loss of efficiency.

Readings from other projects also show frequent breaks between active phases with very short load peaks. Most of the time the workstations are active but the level of capacity utilisation is very low (idle mode). Only during longer breaks and during the night does the pre-configured power management system set the devices to a sleep mode with a power consumption usually far less than 5 Watt.

Therefore the "pinta" project focuses on using sensors to detect the presence of workers, and the associated need for local IT devices, and co-ordinating the computer-controlled shutdown of the devices.

Motion detectors at key positions in the office determine the position of the users. Outside the office space, on the other hand, an app queries the position and direction of movement using the inbuilt technology of smart phones and transmits the results by WLAN to the pinta system. Solutions with Bluetooth transmitters are also possible. The combination of internal office technology and mobile technology makes it possible to set the computers to sleep mode or shut them down entirely depending on the situation.

If the worker leaves the building this is registered by the system. The system can also switch off the lights and heating, if necessary, and switch them back on again when the worker returns. This control system runs in the background and is managed by the pinta platform, resulting in far bigger potential savings and providing incentive well beyond the parameters of IT. In addition, it is also possible to create schedules and profiles for assembly rooms in which facilities such as heating, air conditioning, lighting or network connectivity only start when they are actually needed.

With regard to power demand, energy savings of up to 36 percent have been achieved compared against computers with standard configurations. Even compared with computers with optimised power management systems and an average time-to-standby of 5 or 15 minutes respectively for the monitor and computer, substantial savings can be made by linking power management to lighting and heating control.

Further to this, personalised settings can be made using an app. In an open-plan office this will, of course, need to be coordinated with other users. Thanks to administrators and the fact that all devices are interconnected, coordination can be better effectuated via a system than through verbal communication.

From a business point of view, the costs and benefits will also need to be examined in addition to the aspect of user friendliness. Motion sensors are already installed in many modern office buildings to reduce lighting bills, and connection to an internal WLAN network is also the norm. If businesses prefer not to use smartphones, Bluetooth transmitters present a low-cost alternative as only the infrastructure for the Bluetooth transmitter needs to be developed.

These results were obtained in a test operation where users had to lock their PCs by way of validation each time they left their workplace. The data compiled in this way using the "Windows Context Detection" software were then evaluated, compared with the results of the sensor system, and profiles created for different types of worker to determine daily influencing factors. The savings to be made here are fundamentally contingent upon the type of worker. The longer a worker is away from his/her workplace, the lower the potential saving in terms of ICT and lighting but the higher with regard to cooling and heating costs.

pinta – Pervasive Energy via Internet-Based Telecommunication Services

Project Management:

Jelena Mitic, Siemens AG, Munich

jelena.mitic@siemens.com
www.pinta-it2green.de

Analysing usage and upholding data protection

The previous examples illustrate that detailed investigation and analysis of user behaviour can be very useful, and indeed necessary, to correctly scale and configure optimum power management, which may also be automated depending on the circumstances. In Germany, however, businesses are forbidden by law to process personal data, which would include the motion and attendance profiles of workers. This should be taken into account when planning any measures. According to German data protection law, however, a ban with an authorisation option applies. This means that justifiable exceptions require either a legal regulation or the permission of the party concerned.

"On the one hand we must ensure the careful use of the world's finite resources. On the other hand it is just as important to uphold the civil rights of a society that itself is the basic requirement in finding the optimum balance between these two principles."

Martin Rost of the Independent Centre for Data Protection in the State of Schleswig-Holstein (ULD) on reconciling data protection and GreenIT

The current debate on data protection addresses six protection goals[3]:

- Safeguarding availability seeks to ensure that a process is available in time and can be used properly. A key protective measure is the redundant design of data stock, systems and processes.
- Safeguarding confidentiality seeks to ensure that only authorised persons can access processes and data. This is achieved by sealing off systems and processes, partitioning roles, and encrypting data and communication.
- Safeguarding integrity seeks to ensure that processes remain intact, traceable and complete. This means that hash values are compared in data, actual/target states are compared in processes, and hardware and software are audited by external bodies.
- Safeguarding transparency seeks to ensure that the processing of personal data can be verified, checked and evaluated without unreasonable effort. For this reason, the components of a process must be documented completely and procedures recorded.
- Safeguarding intervenability seeks to ensure that parties concerned can effectively assert their rights and that organisations must perform controlled change management. To this end, parties concerned must be granted immediate access to their data and processes, and organisations must integrate data protection processes into their process frameworks, such as ITIL or CoBIT, and their IT security management system.
- Safeguarding unlinkability seeks to ensure that personal data cannot be gathered, processed or used for a purpose other than that originally intended. If full unlinkability cannot be achieved, it should be realised to the extent that linking would require disproportionate efforts to establish such linkage. This is accomplished through anonymisation and pseudonymisation mechanisms.

Therefore it is important to consider data protection and pseudonymisation and/or anonymisation as early as the data collection stage. At the same time, the works council should be brought on board at an early stage in all the planned and implemented measures that concern the collection of user data. It must be ensured that the data are used solely for the designated purposes (optimisation of IT and energy efficiency) and are not analysed for any other purpose.

2.4. User controls IT resources

Quality of experience as a parameter of orientation

The user, the user's applications and the speed of service the user perceives should be parameters when deciding on the size and scale of IT systems. When calculating the IT hardware requirements, it is important to also factor in the applications which the hardware will be running in the future to keep the number of appreciable delays to a minimum. Different waiting times can certainly be tolerated for the various applications, so the maximum possible computing speed (milliseconds) does not always need to be achieved.

If these aspects are considered in the design and control of the hardware, particularly the server and network infrastructure, an optimum hardware design can be achieved without the hardware infrastructure being oversized or undersized for the job. For this, the response times of a data centre can additionally be optimised for individual applications. This measure of the user's experience is known as the "Quality of Experience":

In business practice, the QoE of IT systems can often only be gauged by asking the end user about his/her experience. This must be performed separately for individual applications, requiring much time and effort. To now determine the individual response times during operation, user behaviour must be evaluated. In contrast to the process for controlling the operating modes of the end user devices, not only the presence of the user but also the user's specific behaviour at the PC (when does the user use which particular application and for how long) must be determined here. Such an approach, however, would result in an in-depth "surveillance" of the user, which is at odds with the data protection requirements discussed in the previous section, an issue which could only be resolved by consensus with the employees and works councils.

Quality of Experience (QoE) is a measure of the customer's subjective experience with a service. Generally speaking QoE looks at an offering from the standpoint of the customer.

[3] Martin Rost of the Independent Centre for Data Protection in the State of Schleswig-Holstein (ULD) on reconciling data protection and GreenIT, http://www.it2green.de/de/1223.php

Client probe

The GreenIT Cockpit project took a different approach to determining the QoE: In the place of a user, a "workstation click probe" continuously measures a few (5-8) quality features of an IT service as a QoE value (this IT service is a communication service in the "GreenIT Cockpit" project). To this end, application-specific load generators are installed and measured. This end-to-end measurement from the workstation automatically reconstructs the relevant use of an IT service from a user's perspective and presents the QoE achieved in numerical form, e. g. as a time or quantity measurement. The click probe is always an additional office computer. No work is performed on this computer in order to avoid legal restrictions.

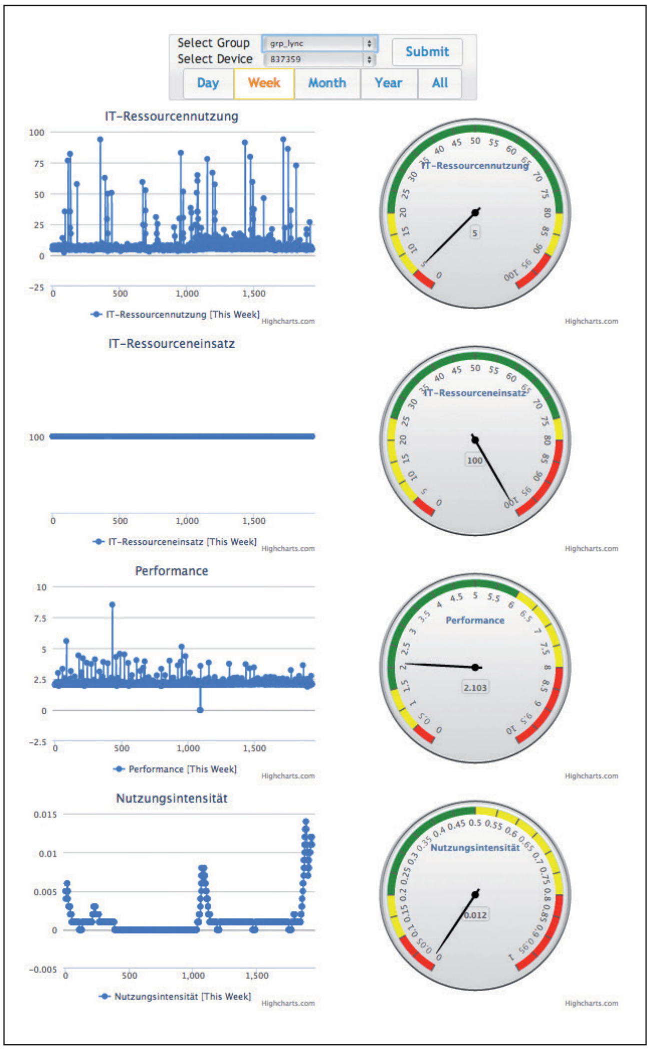

The relevant uses of the IT services from a business perspective are indicated on a dashboard. Source: Axel Springer SE

The QoE data gathered at fixed intervals are stored in a central location and displayed on a dashboard together with the data simultaneously available on the energy consumption of the IT service, system utilisation and the current quantity of relevant uses of this IT service from a business perspective (see figure). The information is displayed as a current value (dial gauge) and over time with a selectable display period.

Using quality assurance methods (such as the quality control card), the collection of data is monitored for specific trends so that the IT service causing the trend can be regulated before warning or error thresholds are exceeded. The regulation focuses on optimising the energy consumption of the IT system that delivers the particular IT service. The dial gauges have coloured markings at either end of the scale (corresponding to "too good" or "too bad"), representing the threshold values. The click probes were implemented on the basis of the Sikulix software for Windows-based and OS X-based workstations. Data collection and web-based displays are implemented with MySQL and JavaScript.

Such application-specific load generators are possible for a wide variety of applications and can be adapted to the specific place of application and associated critical business processes. The advantage of such a client probe over standard load generators is that instead of generating an unspecified load in the data centre, the client probe generates precisely the same load as the real application on the machine that also processes the service. The real QoE of the user can therefore be reconstructed and used as a control parameter to resolve the conflict between conserving energy and achieving the fastest possible response times.

GreenIT Cockpit – Business-Oriented Management Cockpit for Energy-Efficient ICT Throughout an Organisation

Project Management:
Thomas Leitert, TimeKontor AG, Berlin
thomas.leitert@timekontor.de
www.greenit-cockpit.de

2.5. Energy-optimised WLAN

Shutdown on low load?

While the use of WLAN is widespread in home environments, more and more businesses and institutions are also offering and using WLAN on their premises. As is frequently the case in IT, WLAN infrastructures too are often oversized in order to prevent any bottlenecks should issues arise. Over broader areas, this can certainly result in significant energy demand. According to a study conducted by Steven Lanzisera, Bruce Nordman and Richard E. Brown, the WLAN infrastructure accounts for roughly 5 percent of the electricity consumed by the network infrastructure. [4]

To save energy, it would seem obvious to shut down sections of the unused infrastructure and only restart them if necessary. However, WLAN, in particular, should be available to the user any time, without delay and with the full quality of service. The GreenPad project at the University of Paderborn analysed this conflict in greater detail and developed ways for optimisation. [5]

The user traffic and the data volumes of the individual access points were measured to gain an understanding of the utilisation of the networks. This data volume is a prerequisite for identifying under-utilised access points and associated approaches for load-adaptive operation.

At the University of Paderborn, large differences were found between weekdays and weekends and fluctuations between day and night that offered latitude for optimisation. These quite typical load profiles can also be expected to the same or similar extent in business and public authorities. In addition, the analysis revealed that on average 65 percent of the access points are not used. By completely shutting down all of the access points at times of low load (11 p.m. to 7 a.m.), the University of Paderborn could, in theory, save 25 percent of the energy used by the WLAN infrastructure. The service would not be available during this period, however. While this would be a very simple solution, it would not be very user-friendly.

Controlling with resource-on-demand algorithms

As witnessed with server operation, energy savings can be achieved not only by complete service shutdown but also by using load-dependent operation.

> The individual access points must be treated as an entire infrastructure to be optimised as a whole, and not as separate components.

To this end, the GreenPad project studied and appraised a variety of resource-on-demand algorithms for access point control:

- Green clustering
- IP green clustering
- Green star clustering (newly developed)

The distance between the individual access points and the presence of obstacles, such as walls, are central to deciding whether access points can be shut down without losing network coverage. While it would have been possible to evaluate this by hand, this would have been too time-consuming. The project therefore decided to use data that can be evaluated automatically, such as the RSSI readings of the access points. By leveraging these data an algorithm can then determine which access point should be shut down and when.

To quantify the potential savings achieved, the power de-

> The result revealed that energy savings of 15 percent could be achieved at the University of Paderborn through load-dependent operation while still maintaining the quality of service (compared with 25 percent savings in event of full service shutdown).

[4] Steven Lanzisera, Bruce Nordman, and Richard E. Brown: Data Network Equipment Energy Use and Savings Potential in Buildings, Lawrence Berkeley National Laboratory, ACEEE Summer Study on Energy Efficiency in Buildings, 2010, http://www.aceee.org/files/proceedings/2010/data/papers/2195.pdf

[5] Maximilian Boehner, Sebastian Porombka, Gudrun Oevel: WLAN energy efficiency, University of Paderborn

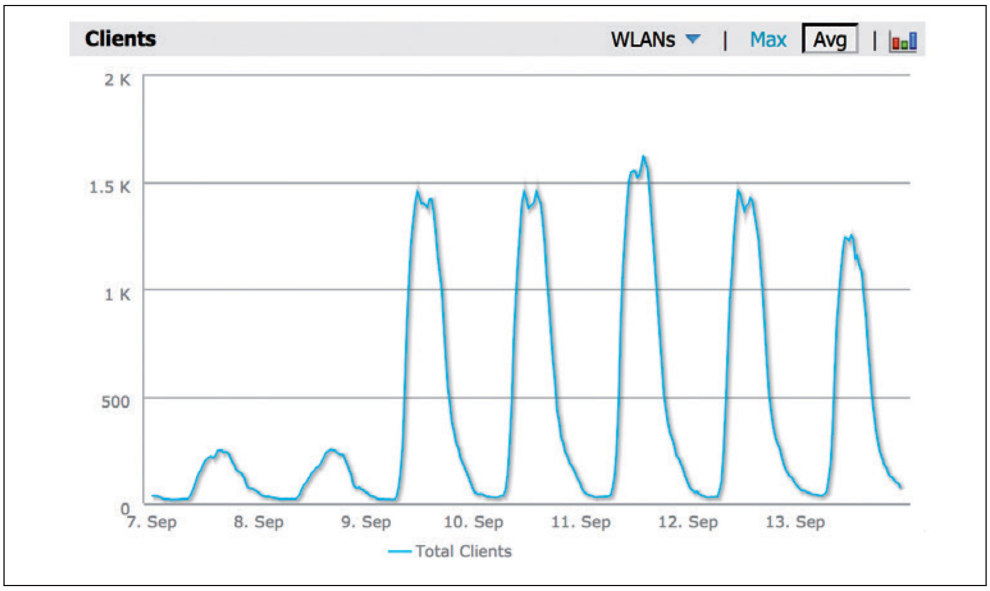

Energy savings with WLAN can be achieved by load-dependent operation. Source: University of Paderborn

mand of the access points was recorded beforehand based on manufacturer specifications (maximum demand) and on individual measurements (idle demand and load-dependency). This was based on a combination of the evaluated algorithms (IP green clustering during high-fluctuation periods, green star clustering during low-fluctuation periods).

The strategies adopted to optimise the power consumption of the WLAN infrastructure can be easily transferred to other WLAN infrastructures and could also be integrated directly into future WLAN technologies.

GreenPAD – Energy-Optimised ICT for Regional Economic and Knowledge Clusters

Project Management:
Thomas Dyroff, unilab AG, Paderborn
thomas.dyroff@unilab.de
www.green-pad.de


Power
Active
Power
Active
Power
Active
Power
Active
Power
Active
Power
Active
Power
Active
Power
Active

3. Energy-efficient ICT in the data centre

Rising energy costs in production often prove to be the motivating factor behind greater energy efficiency and GreenIT in the data centre. Once a business sets itself the goal of saving energy, a challenging process unfolds in planning the introduction of active measures. A wide range of information can be found on the Internet and many commercial offerings for technical solutions are available on the market to help businesses green their data centres.

The process of improving energy and resource efficiency in the data centre should start with the IT equipment and work its way outwards from there. This can even culminate in the collaborative interaction of multiple data centres in a closed cloud. The possibilities for optimisation are many and varied, affecting both hardware and software aspects.

"Much has happened in recent years; particularly when I think about what the industry was up to just a few years ago. Planning was primarily centred on security and availability criteria. Energy efficiency was of hardly any concern. Nowadays only very few data centre operators could afford to adopt such an attitude."

Ulrich Terrahe, dc-ce RZ-Beratung, Frankfurt am Main

Technology procurement and configuration

In the first stage, substantial wins can be achieved with even the simplest, low-cost measures. This includes activating existing power management options or improving how cables are bundled to optimise cooling in the rack. Climate control and heat management offer enormous potential for optimisation. In this connection, local conditions and individual possibilities for cost-cutting solutions must be examined.

Operation at high temperatures also presents considerable potential to save electricity over the entire system as the substantial cooling requirements can be pared down. Such a measure is only permissible on condition that all the technological elements, including the servers, storage, network technology and power distribution system, withstand the higher operating temperatures and the cooling system reacts quickly to variations.

Ultimately consideration must also be given to the fact that the energy efficiency of IT systems is still steadily improving in parallel with increased system performance. Therefore, regular hardware upgrades provide great potential for optimisation. An important prerequisite for efficiency in this connection is the right size and scale of the IT system and the right choice of products for the desired performance and quality of service.

Energy and resource optimisation

While simple measures can often deliver substantial energy savings, far higher energy savings can mostly only be achieved by the targeted increase in capacity utilisation and the load-adaptive operation of all elements in the data centre. More complex interactions and correlations must be established to boost energy efficiency even further. The appropriate sizing and consolidation of ICT resources in the data centre is a fundamental goal.

In this context, virtualisation – i.e. the largely automated utilisation of servers, storage systems and network technology – plays an important role. However, successful virtualisation also requires a high degree of interoperability between elements in the data centre that have often been controlled separately up until now. It is evident that thermal management is becoming increasingly complex with virtualisation.

Measuring energy demand and applying KPIs

A thorough analysis of heating and energy-related conditions, across all sub-sections, forms the basis of any successful optimisation strategy. Nuanced and continuous measurements of power consumption, load profiles, utilisation and thermal conditions are indispensable and create the framework for delivering a data centre that is both performance-oriented and energy-efficient.

In this connection, many options are available to record in detail the load profiles and power demands of ICT and infrastructure systems using existing sensory systems and software options or by installing additional measuring technology. The creation of a continuous, adequately proportioned and informative load and energy monitoring system proved to be a major challenge in many projects.

Action to improve the energy efficiency of a data centre starts with the IT equipment.

This section provides examples of factors that must be taken into consideration during measurements and simulations, and what key performance indicators (KPIs) and metrics are available.

Green control station and migration to the cloud

The load transfer or migration of applications not only within the data centre itself but also to the cloud – i. e. a group of several collaborating data centres – only delivers energy savings if there is a high level of automation and system component interoperability. With such a format, favourable local conditions – such as the availability of wind or solar energy – can be exploited for green data centre operation.

Time-dependent load transfer to a remote data centre can make a great deal of sense as it enables the effective use of renewable energy, which is intermittent and subject to local variations.

If a very dynamic, load-adaptive operation of ICT and infrastructure systems is desired, this often results both in higher initial costs for the hardware and software and the danger that delays or even errors can be caused by less flexible or badly synchronised systems. Reliable, available and energy-efficient data centre concepts can only be realised by meticulous planning, testing and continuous optimisation.

The research results of the IT2Green technology programme illustrated below cannot fully present the "what" and "how" of all the measures possible today. Instead the section focuses more on the "why", with specific approaches, tips and pointers provided for the strategic further development of energy efficiency in the data centre.

3.1. Technology procurement and configuration

Usage dominates in energy demand

The average use phase of a server is four years, depending on the server type and application. Even in the case of simple servers with just one or two processors, the electrical energy consumed over this period accounts for over 80 percent of the server's total environmental footprint over its lifecycle, with external cooling needs not yet factored in.

This is the finding of scientific studies that analysed the product carbon footprint (PCF) of servers over the entire product lifecycle.[6]

Conscious product procurement and the energy-efficient use of servers is all the more important, a fact which, incidentally, is also true for the data storage systems and network technology in the data centre. Only with an increasing degree of integration – such as current microserver concepts like the Moonshot system[7] – will it be possible to significantly reduce the energy-related environmental impact. Future optimisation measures will then centre on the content of recyclable material in the ICT and therefore resource efficiency.

In summary, the electrical power demand of ICT and therefore the procurement of energy-efficient products and the energy-optimised operation of technology are the primary environmental objectives for data centres right now and will remain so over the next few years. IT managers are aided by the continued dynamism in the development of technology to create increasingly efficient and smarter hardware.

Continued miniaturisation

The energy efficiency of ICT doubles roughly every 18 to 24 months. This was the conclusion reached in 2011 by a team of engineers from Intel and Microsoft led by scientist Jonathan Koomey from Stanford University.[8]

The team showed the causal relationship between the energy efficiency of ICT and Moore's law, which describes the paradigm of miniaturisation in the semi-conductor industry and predicts the doubling of transistor counts per integrated circuit in a constant time cycle.

Given that transistor structures nowadays are just a few atomic layers thick, the technological challenges of gradually increasing the integration density, and therefore computing power, are considerable. Potential increases in energy density, and the resulting additional waste heat produced locally during operation, pose a serious problem.

Apart from the use of new materials, the electrical voltage is reduced more and more by technological means. The consistent development of multi-core architectures for processors also promotes energy efficiency.

The aim of these technologies is to deliver better system utilisation and the load-adaptive operation of individual computing cores. At low workloads, therefore, better low-power modes can be achieved. The semi-conductor industry will still be able to continue this pattern of miniaturisation and energy efficiency in the next generations of chip development.

Leading manufacturers emphasise this point, making reference to the possibilities afforded by 3D integration and new, high-integration cooling concepts at the chip and PCB level, which are far from exhausted.

Modern hardware: the basis for energy efficiency

The power consumption of the central processing unit (CPU) primarily determines the power demand of the servers in a data centre. With a rising trend in both the number of cores and the average number of sockets per device unit, careful attention must be paid to appropriate server configuration and economical operation despite theoretical improvements in energy efficiency.

The growing demand for computing power and technical progress in the development of energy demand are two opposing processes, but increasing energy demand from the growing demand for computing power tends to predominate.

6 Markus Stutz: Carbon Footprint of a Typical Rack Server from Dell, 2011, http://i.dell.com/sites/content/corporate/corp-comm/en/Documents/dell-server-carbon-footprint-whitepaper.pdf

7 Technical Whitepaper: HP Moonshot System, http://h20565.www2.hp.com/portal/site/hpsc/template.BINARYPORTLET/public/kb/docDisplay/resource.process/?spf_p.tpst=kbDocDisplay_ws_BI&spf_p.rid_kbDocDisplay=docDisplayResURL&javax.portlet.begCacheTok=com.vignette.cachetoken&spf_p.rst_kbDocDisplay=wsrp-resourceState%3DdocId%253Demr_na-c03728406-3%257CdocLocale%253D&javax.portlet.endCacheTok=com.vignette.cachetoken

8 Jonathan G. Koomey: Growth in Data Center Electricity Use 2005 to 2010, Stanford University, 2011

Against this backdrop, the recommendation could be made to replace servers regularly in order to benefit from the potential efficiency gains of new products. However, this recommendation also requires new servers to be sized correctly for the intended application and server capacity to be properly utilised.

Measurements show that CPU utilisation scales quite closely with the total power consumption of the device. The minimum and maximum thermal design power values (TDP) of the central processing unit are simple indicators for the range of power consumption of a device, wherein currently between 40 and 60 percent of the power consumed by a server is attributable to the CPU, depending on the particular configuration. This value varies depending on the memory and redundancy specifications.

Energy efficiency benchmarks

In addition to the TDP values, more complex benchmarks that relate the performance to the power also provide important information to make a good purchase decision. The Standard Performance Evaluation Corporation (SPEC) currently provides two relevant benchmarks for servers. The older SPECpower_ssj2008 benchmark has been used by many producers since the end of 2007.

This benchmark reports power consumption over the entire load range. Specific test data for hundreds of server types can be accessed in a public database.[9] On the surface, the test results are very informative as they report power consumption for different load levels. However, experts are of the opinion that the benchmark's load profile is not very realistic and therefore a new SPEC benchmark has been trialled since late 2013.

This benchmark, known as the Server Efficiency Rating Tool (SERT), was developed in close collaboration with the ENERGY STAR® international energy label and the US Environmental Protection Agency (EPA) ENERGY STAR for Computer Servers programme. SERT seeks to simulate the load generation of different applications and therefore deliver a better representation of real usage profiles and energy efficiency. The environmental authorities and the businesses involved have high hopes for the SERT standard as the desire is to define precise energy efficiency requirements with SERT, which would come into effect as part of new procurement directives.

Just as SERT was specifically developed for the server specifications of the ENERGY STAR® programme, the EPA commissioned the Storage Networking Industry Association (SNIA) to develop the EmeraldTM Power Standard, which contains test specifications and criteria for a standardised and reproducible energy efficiency rating for data storage systems. The first tests with this data storage benchmark also got underway in 2013 and are set to form the basis for award criteria of the ENERGY STAR® programme.

Apart from the ENERGY STAR® international energy label, environmental requirements of servers, data storage systems and network technology are also currently being examined in a preparatory study for the European Directive for Ecodesign Requirements (ErP Directive 2009/125/EC). Analysing a broader subject matter, this study is due to release its final report at the start of 2015.

The study addresses both energy consumption in the utilisation phase and material and resource requirements over the entire product lifecycle, thereby examining environmental requirements for IT equipment in a data centre from a "whole-of-product" approach for the first time ever.[10]

9 SPEC Power: SPECpower_ssj2008 Results, http://www.spec.org/power_ssj2008/results/

10 ErP Preparatory Study on Enterprise Servers (Lot 9), www.ecodesign-servers.eu

3.2. Integrated energy and resource optimisation

Components and system configuration

Ongoing modularisation at the equipment level offers big scope for energy savings. The key here is that the individual elements have energy-saving properties, are interoperable and can be well matched. This demands highly efficient and appropriately sized power units that can offer an efficiency rating of over 90 percent today even under partial load.

While high-efficiency power supply units (PSU) do not solve the problem of frequent under-utilisation, they are a cornerstone of a low-power system.

In practice, many PSUs are, unfortunately, still not equipped to meet the growing demands of more frequent load changes, with low partial-load and energy-saving modes on the one hand and brief power peaks on the other. Some PSUs neither support modern energy-saving modes nor higher ambient temperatures.

Both aspects are a key prerequisite, however, for dynamic, low-energy data centre operation. The specifications of the power supply units must be examined accordingly. The same is true for all elements of the power supply system in the data centre. Reduced conversion loss is a building block of energy efficiency.

The topic of direct D/C power supply is often raised in this connection. Relevant systems have been available on the market in recent years and energy savings can be made, in theory. However, direct current requires an integrated data centre design and appropriately trained staff, which is usually not an option for small and medium-sized data centres.

Thermal management

As processors and other microelectronic components give off the electrical power consumed almost entirely again as heat, functional and effective heat management is another basic building block of energy efficiency. In this connection, cooling should be designed from the ICT out, i.e. starting from the ICT or rack and working its way outwards to the room or building. The real climate control needs and flow conditions must be determined through appropriate measurements and calculations (see below).

The cooling systems in many data centres are oversized for the data centre needs. Generally speaking, the reliable and energy-efficient operation of ICT is subject to defined climatic conditions, including temperature and humidity.

Many ICT vendors are currently developing products that allow operation at higher inlet temperatures of up to 35°C. These products are often somewhat more expensive. The electronics and mechanical components and the overall thermal design are designed for higher temperatures or temperature changes. In this way, the product vendors are supporting data centre operators in their endeavour to significantly reduce cooling energy by allowing higher room temperatures.

In recent years, the ASHRAE guidelines for permissible temperature and humidity limits in a data centre have also changed accordingly. Regulation of the inlet temperature value must also take the overall behaviour in the rack into consideration. It should also be noted that electrical fuses in the power distribution system at rack level, for example, must be designed for higher temperatures.

Data centre operators are advised to test out higher room temperatures and capture interaction between heat and energy in this connection. Ensuring adequate ventilation and avoiding thermal hotspots and short-circuits have priority and require specific analysis with appropriate sensory technology.

In conjunction with computational fluid dynamics tools (CFD), temperature and infrared sensors installed throughout the data centre can depict temperature and air flows as 3D heat maps.

Not only can this technology be used to make immediate adjustments but also to analyse new load scenarios. Modern data centre infrastructure management tools (DCIM) offer functions for embedding CFD analyses.

These tools deliver valuable information and pointers for energy optimisation, particularly in the event of large-scale conversion or technology upgrade projects.

With increasingly dynamic, load-adaptive operation, temperature changes will be more frequent and have larger amplitudes.

Software-based load control in the form of virtualisation can also result in large thermal differences in the rack with a negative impact on air conditioning requirements. It is precisely these factors – i.e. dynamic load adjustment and larger temperature changes – which can result in system failure if the hardware is not properly designed to handle such conditions.

Advanced strategies

While simple measures can often deliver substantial energy savings, more complex correlations and interactions must be identified and created to boost energy efficiency even further.

Increasing server capacity utilisation with the aim of consolidating ICT resources is a fundamental strategy to improve the use of energy and resources.

> The software-based, largely automated control of the utilisation of servers, storage and network technology must be extended to the entire data centre infrastructure in order to produce high energy savings.

The virtualisation of all IT elements must be synchronised with thermal management. Less-than-optimum temperature states will occur again and again in practice – in the rack for instance – because load distribution does not take the position of the servers in the rack into account.

Another approach centres on more load-adaptive operation where both the ICT resources and the climate control and power supply infrastructure in the data centre are automatically adapted to the immediate needs of IT production. Measures in this context range from BIOS settings for servers to the configuration of the climate control system and appropriate threshold settings.

> Predictive resource planning for IT production is a prerequisite for effective capacity utilisation. One approach involves the evaluation of historical load patterns and application profiles.

A third approach looks at optimisation over the entire data centre where applications are specifically moved between physically separate data centres, for instance, to benefit from temporary advantages of the site, such as free capacity or the local availability of green electricity.

To conclude it can be said that significantly higher energy savings are usually only achieved by specifically increasing capacity utilisation and by the load-adaptive operation of all elements in the data centre. These measures require far more dedication and commitment and initial costs are higher.

Despite the many open options, data centre operators are nevertheless reluctant to systematically implement GreenIT measures for a number of reasons that are easy to understand. They are particularly concerned about the availability of their services and the economic viability of the system as a whole. Active measures to save energy can certainly involve a certain reduction in availability considering that individual ICT resources are mostly set to sleep mode or even switched off in order to effectuate energy savings.

Therefore it is important to examine whether defined delays or other availability restrictions brought about by measures to save energy are acceptable for the user or client.

3.3. Measuring power consumption

Collect data before planning

When planning and implementing GreenIT measures, the current state of affairs should be recorded and evaluated by collecting data that encompass all the disciplines and sub-sections in the data centre.

> Initially containing energy-related information from product data sheets, the data collection process can extend to reading out existing sensor data – via IPMI for instance – to recently introduced power energy measurements of all the elements in the data centre that use electricity. Thermal imaging can also meaningfully complement the measurement data.

In practice it has been proved time and again that dedicated measurements offer several advantages when taking stock of the current situation. For example, businesses can themselves determine and adjust the granularity of the measurement as required. The consolidation or synchronisation of the measurement results is also often easier as usually only one software programme is used, rather than having to work with several logs. When reading out the data of different sensors, there is also often the problem that measuring intervals cannot be modified and the accuracy of the measurement cannot be verified.

ICT, climate control and power supply

For a robust analysis of existing energy efficiency, it is necessary to record both the power consumption over time and the capacity utilisation profiles (CPU, RAM, I/Os) for the servers installed. As storage equipment and network technology are becoming increasingly important in the data centre, these systems must also be included in the measurement in addition to the servers.

Furthermore, climate control should be measured in a similar, synchronised manner. Aside from the electrical power consumption of all the ICT, climate control and power supply elements, additional data are also needed for a good analysis. Therefore, the input and output temperature of the servers, the fan speed, the flow speeds, the thermal load and the outside temperatures, if applicable, should be recorded. With regard to power supply, the level of efficiency or power loss over the entire load spectrum is of interest for the analysis. As power supply units, in particular, tend to fail frequently, these events should also be documented.

Implementation requires detailed planning

Before any measurement, the entire inventory in the data centre must be examined closely, including the cable lists and circuit diagrams. A top-down approach should be followed here, as it helps to capture measuring technology that is already installed (sensors) and therefore potential sources of data.

In the following section, detailed planning focuses on determining the measurement parameters and the measuring technology best suited to the task. To measure the energy in the data centre, products are available from professional service providers that contain hardware and software elements, not only helping to significantly improve data collection but also aiding data interpretation in particular.

> Important decisions must be made regardless of whether measurements are taken independently or the measurement task is assigned to a service provider.

The sensor types and precise measuring points must be identified. Then the measurement intervals, the data volume and the consolidation and evaluation of the measured values must be planned. Spot measurements on static consumers reduce the measurement effort. The detailed documentation of the measures implemented and the measurement results is very important for the subsequent analysis process.

Troubleshooting and plausibility checks

It is absolutely essential to check the plausibility of the measurement results at an early stage as the process of pinpointing sources of error is a time-consuming task that is anything but simple. This was the opinion of all the IT2Green projects that performed complex measurements in the data centre.

It is advisable to check the first measurement results after a few days on the basis of information from product datasheets, SPEC benchmarks or similar sources. The projects also experienced a series of technical errors. For example, measuring devices (e.g. transformers) were installed incorrectly or were incorrectly assigned in the analysis tools.

Other sources of error included faulty or incorrectly calibrated measuring devices. Hidden errors are often only revealed by clearly identified changes to the operation, e.g. new technology installed or a test run performed. Against this backdrop, businesses would be well advised to schedule plenty of resources and time for the localisation of faults.

Interdependence between ICT and infrastructure

The data gathered should be evaluated periodically with a range of comparative analyses. The first step of the analysis should look at the power consumption or energy demand of individual devices in relation to capacity utilisation.

The primary focus here is on cyclic changes and metastable load profiles. Attention must be paid to how the devices respond during high or low utilisation. This also applies for virtual machines (VMs), VM behaviour under load and the capacity utilisation of physical machines. The power consumption of different types of server (rack, blade, multinode etc.) should also be compared if they run the same services or applications.

The analysis becomes even more interesting if the interdependence between individual ICT systems, applications and infrastructure elements is studied.

For example the study could look at how load-adaptive the cooling system in the data centre already is. If other energy-saving options are to be tested, a test environment can also be set up outside production where load generators simulate different server and VM configurations using the measured data from production, and dynamic load adjustment can be evaluated.

Summary

Load and energy measurements in the data centre support optimisation measures and provide hard facts for planning decisions to be taken by the executive board, the IT department and facility management. As a first step, it is advisable to take stock of the current situation in a comprehensive and consistent manner using professional tools. When planning the measuring points, it is important to take a top-down approach and ensure existing technology is also incorporated into the process. Building on the experience gathered, a leaner monitoring system can be set up later on as part of integrated data centre management (DCIM solution).

The measurement results should be checked and queried constantly in order to identify potential sources of error early on. The analysis should focus on causal relationships between capacity utilisation, power demand and infrastructure response.

IPMI values and external power measurement

As part of the IT2Green project "GreenPAD", an external measuring instrument was used to measure the power consumption of three different server systems over the entire load range, and the results were then compared with values obtained via IPMI.

It was found that the values read out via the Intelligent Platform Management Interface (IPMI) largely matched the measuring results, suggesting that the energy consumption of the servers can be very accurately captured with the internal IPMI values and that there is no absolute necessity for an additional external measurement for invoicing or energy analysis purposes, for instance. Among other things, IPMI comprises specifications for the measurement of hardware states using sensors. With IPMI it is possible to read out power consumption values over the entire load range between idle-load and full-load in increments of 4 Watt. The project took measurements on rack servers and on a blade system and compared them with the values obtained through IPMI.

In the case of the rack server, the power consumption of the entire functional computer was measured, including the local hard drive, fans and also PSU loss. The energy use of the network switch infrastructure, which is entirely provided externally, is not included in the measured values. In the measurement, partial load scenarios were created gradually by putting load on only a section of the processor core using the Linux "Stress" testing programme. As the synthetic load for complete processor capacity utilisation, the Linpack benchmark was used in parallel in multiple virtual machines with high overbooking of the processor cores physically available. On comparing the real measurement and IPMI-based values, only minor discrepancies were found in all load ranges with an average error of under 3 percent.

In the case of the blade servers, a total of 18 server blades are mounted in an enclosure with shared active infrastructure components. In the configuration selected, these constitute six shared power units, fans, network switches to connect the server blades and two redundant management blades. The IPMI values for the individual server blades cannot take the power demand of this shared infrastructure into account. In the production environment, many different applications run on a complete blade server. Therefore billing information must be recorded at the blade level.

From the measurement results it can be concluded that when billing on the basis of the blade energy use values obtained by IPMI it is necessary to factor in both the inaccuracy of the measured values and a relatively high surcharge, in percentage terms, for infrastructure consumption values that cannot be allocated directly. If entire blades are shut down in times of low load, the complete infrastructure must continue operating. Measurements here show that in this load range (i.e. below idle with a total of 1,750 Watt), the measurement errors of the internal sensor systems for total power consumption increase significantly.

In conjunction with a deterioration in the level of efficiency of the power units, the actual power consumption measured externally is considerably higher than the value displayed based on internal measurements. Therefore, in the overall energy balance the infrastructure overhead increases disproportionately. For application scenarios where entire computing nodes are shut down, the blade servers are, therefore, not as well suited as an alternative configuration with nodes that can be shut down individually (rack servers).

Compared directly with the measurement results for the rack server, the measurement results for the blade server indicate higher power consumption and lower application performance particularly under full load.

The poorer application performance might be attributable, in part, to the effects of scales of the 18-blade configuration, the fact that the measurement was taken earlier, and the less elaborate optimisation of the more complex system. Oracle's slower Java software was evidently used to perform the measurements. The original hope that energy efficiency could be significantly improved by using blade servers – which are more complex – compared to using comparable individual servers cannot be confirmed by the measurement results we have obtained to date and by the manufacturer's benchmarks.

GreenPAD – Energy-Optimised ICT for Regional Economic and Knowledge Clusters

Project Management:
Thomas Dyroff, unilab AG, Paderborn
thomas.dyroff@unilab.de
www.green-pad.de

3.4. KPIs at a glance

Assessing energy efficiency

Many data centre managers now use key performance indicators (KPIs) to assess the energy efficiency of and in data centres. As a wide range of KPIs have been released for data centres in recent years, the question is (i) how robust and informative are the individual comparative values, and (ii) for what areas of the data centre do they provide sound information (IT systems, power supply, cooling system etc.).

The most well-known metric for GreenIT in data centres is power usage effectiveness (PUE), developed by the Green Grid Association. As with many of the common KPIs, the PUE measures the energy efficiency of the infrastructure in relation to the IT load present. There are metrics that map the efficiency of the entire data centre or only particular portions of the data centre (e.g. the partial Power Usage Effectiveness (pPUE) can be determined for the cooling system or an individual server room). Other metrics intentionally only focus on the infrastructure. The Self-Benchmarking Guide of the Lawrence Berkeley National Laboratory (LBNL) contains 14 metrics for data centre operators to assess the building infrastructure. This also reveals the biggest problem with current metrics:

The efficiency of the IT load itself is not taken into consideration in the PUE and in most other metrics.

While with its Four Metrics system the Uptime Institute published a scientific approach that contains "useful work done" (H-POM) as the most important variable, no exact definition of "useful data centre work" has been provided so far. The areas of software (e.g. server virtualisation) and IT hardware (e.g. CPU, power supply units, fans) can be captured through IT capacity utilisation (DH-UR, DH-UE).

The same applies for the Corporate Average Data Center Efficiency (CADE) KPI, also developed by the Uptime Institute. The most important value for the CADE is "IT energy efficiency" (IT-EE). However, up to now it has not been possible to measure this value, so the Uptime Institute recommends the value be estimated. Therefore the results of the CADE depend heavily on the quality of the estimates and comparability is difficult.

While there are many KPIs for data centres, the level of information they deliver regarding energy efficiency differs greatly.

In contrast to the other technical approaches, the Cost of Services model presented by the British Computer Society (BCS) has a financial orientation. The idea is to present the total data centre energy efficiency as a ratio of data centre power demand and the value add of IT in the business. The costs for power demand should be transferred to the IT equipment, the software and ultimately the IT services. The last step, which could, in turn, be used to assess the IT itself is not described, unfortunately.

The examples given are just some of the metrics and KPIs for climate control and power supply of IT components, also demonstrating how difficult it is to assess the efficiency of IT.

The most important vehicle in all efforts surrounding GreenIT in the data centre is the use of energy-efficient IT components, for which a host of metrics is now also available.

When it comes to assessing server energy efficiency, the first name that must be mentioned is the Standard Performance Corporation (SPEC), which seeks to develop standards to measure the performance of servers. The most well-known standard is the SPECpower_ssj2008 benchmark. This determines the maximum number of computing operations possible per second and then measures the power demand of the server for different load levels. One disadvantage of the SPECPower_ssj2008 is that it only tests a special workload that puts particular "Stress" on the CPU and RAM but hardly puts any load on the data storage system or network interfaces, for instance. To map the energy efficiency of a server for a variety of application scenarios, SPEC developed the Server Efficiency Rating Tool (SERT). As part of the benchmark, various workloads (servlets) are executed in multiple load levels and measured.

With regard to the energy efficiency of storage systems, the Storage Performance Council has developed the SPC-1/E, SPC-1C/E, SPC-2/E and SPC-2C/E benchmarks. The SPC-1/E and SPC-1C/E benchmarks generally determine the ratio between the number of randomly generated read and write operations and the specific power demand for multiple load levels ranging from idle to full load.

Approaches to improve the energy efficiency of data storage facilities are being developed under the Green Storage Initiative (GSI). The most important GreenIT approach here is the SNIA (Storage Networking Industry Association) Emerald Power Efficiency Measurement Specification. This determines the energy efficiency of storage units in standby mode and at medium system power consumption in gigabytes per Watt, and calculates the energy efficiency of systems in active mode from the number or data volume of read and write operations and the average rate of power consumption.

To measure the energy efficiency of network components, the Alliance for Telecommunications Industry Solutions (ATIS) has defined the Telecommunications Energy Efficiency Ratio (TEER). At different load levels, the TEER quantifies the energy demand of network components that are categorised by device class or their position in the network (access, high speed access, distribution/aggregation, core).

The Energy Consumption Rating Initiative (ECRI) has developed the Energy Consumption Ratio (ECR) also to assess the energy consumption of network equipment depending on the data throughput rate. Here, the energy demand and data throughput are captured at full load, variable capacity utilisation and in the extended idle mode, and the average energy efficiency of the individual device is determined. Like ATIS, the ECRI also defines different device classes.

However, these metrics alone cannot determine whether the energy-efficient IT hardware in the data centre is actually used efficiently.

Conclusion: By using a clever combination of the metrics available, the energy efficiency can be well represented for the data centre building technology and for the ICT to some extent.

Many institutes are exploring so-called data centre productivity metrics to depict the actual level of IT efficiency. The aim here is to put the costs (energy, financial costs) in relation to the benefit (IT service):

- CS: Distribution of the data centre (operating) costs to the sum of the IT services
- Green Grid and SPEC: Datacenter Performance Efficiency (DCPE)
- Uptime Institute: IT Productivity per Embedded Watt (IT-PEW)
- TU Berlin and Ecoinstitute: Carbon footprint for specific IT applications

There is therefore still a great need for further research. This is being addressed in projects such as the IT2Green incentive programme.

A detailed analysis of the existing energy efficiency metrics for data centres is available in the GreenIT dossier produced by IT2Green (http://www.it2green.de/de/1365.php, in German only).

Links to metrics and organisations

- The Green Grid Association: www.thegreengrid.org
- Uptime Institute: www.uptimeinstitute.com
- Lawrence Berkeley National Laboratory: http://hightech.lbl.gov/dctraining/about.html
- British Computer Society: www.bcs.org
- GreenIT Promotion Council: http://home.jeita.or.jp/greenit-pc/e/
- Standard Performance Evaluation Corporation: www.spec.org
- Storage Performance Council: www.storageperformance.org/home
- Storage Networking Industry Association: www.snia.org
- Alliance for Telecommunications Industry Solutions: www.atis.org
- Energy Consumption Rating Initiative: www.ecrinitiative.org
- American Society of Heating Refrigerating and Air Conditioning Engineers: www.ashrae.org
- Centre for Energy Policies and Economics: www.cepe.ethz.ch
- Transaction Processing Performance Council: www.tpc.org

3.5. Data centre with green control station

Load forecasts for effective virtualisation

Virtualisation, and consolidation measures building on virtualisation, can reduce the number of servers needed and the associated energy demand. However, if computing loads fluctuate greatly it is often not possible to permanently reduce the number of physical servers.

In such cases, cross-server control coupled with the temporary shutdown of individual hardware elements can nevertheless still offer potential for savings.

This requires knowledge of the resources needed by the applications running on the servers. Users specify the quality of service, including associated time parameters, via service level agreements (SLAs).

The aim of cross-server control is to adapt the active servers to the current, actual need for computing capacity. One approach is to dynamically re-arrange the applications at runtime to better utilise individual physical machines.

Server power models that depend on capacity utilisation provide information about the expected power demand. The system resources studied comprise:

- CPU time in percent
- RAM use in bytes
- Storage I/O in bytes

Load and power management processes

The AC4DC research project developed a load and power management (LPM) system for servers and services that adapts the number of active servers to the actual resource needs of the services at any time. To use the process, the servers to be managed must run in a virtualised environment and the services must be encapsulated in virtual machines (VM). Using a technique called live migration, which allows VMs to be transferred between servers on the fly, the services can be managed dynamically and entire servers can be consolidated.

The graphic on page 33 provides an example of the principle procedure for an environment with three servers. In the initial arrangement at time t1, the capacity utilisation of the servers is already quite high (level of VMs) so optimisation is not possible. However if resource needs drop (VM 6 at t2 for instance), the servers can be consolidated by migrating the VMs on server A to the other two servers. Server A is now currently not needed and can be set to a low-power mode. Given that servers at idle still require a certain percentage of their maximum power consumption (modern servers: approx. 20 percent; older servers: over 50 percent[11]), much energy can be saved with this approach.

The load and power management system must also react to growing resource demands, however, as illustrated in time t3 in the example. If the resource needs of VM 6 increase again, VM 1 must be migrated back to server A to prevent any resource bottlenecks. However, considering it takes some time to boot a server and migrate a VM, it would not be possible to prevent this bottleneck with a purely reactive process. An important requirement of LPM can be identified from this example:

The resource needs of the VMs must be known in advance, through forecasts for example.

The LPM system developed in the project therefore has a component for predicting the resource needs of services. To this end, the periods where patterns are repeated are identified in the historical load data and used to forecast future resource requirements. The more historical data available, the better the result. However, only services with regular behaviour patterns can be reliably modelled with this predictive method. This involves business applications whose use correlates largely with working times. Sporadic applications, on the other hand, must be considered separately by assuming a static, statistical (maximum) capacity utilisation rate for them.

The graphic on page 34 illustrates the other components and interfaces of the AC4DC load and power management system. As already explained, the observed (historical) resource needs of the VMs are needed to model future requirements. The controller then uses the forecasts to plan any possible optimisation in the future. To this end, the controller also takes into account load-dependent power models of the server hardware and the externally agreed service level agreements (SLAs) to select the servers to be used and determine the target level of utilisation. The control commands required for optimisation (migrations, server deactivation and reactivation) are collected

[11] Results of the SPECpower_ssj2008 benchmarks at http://www.spec.org/power_ssj2008/results/

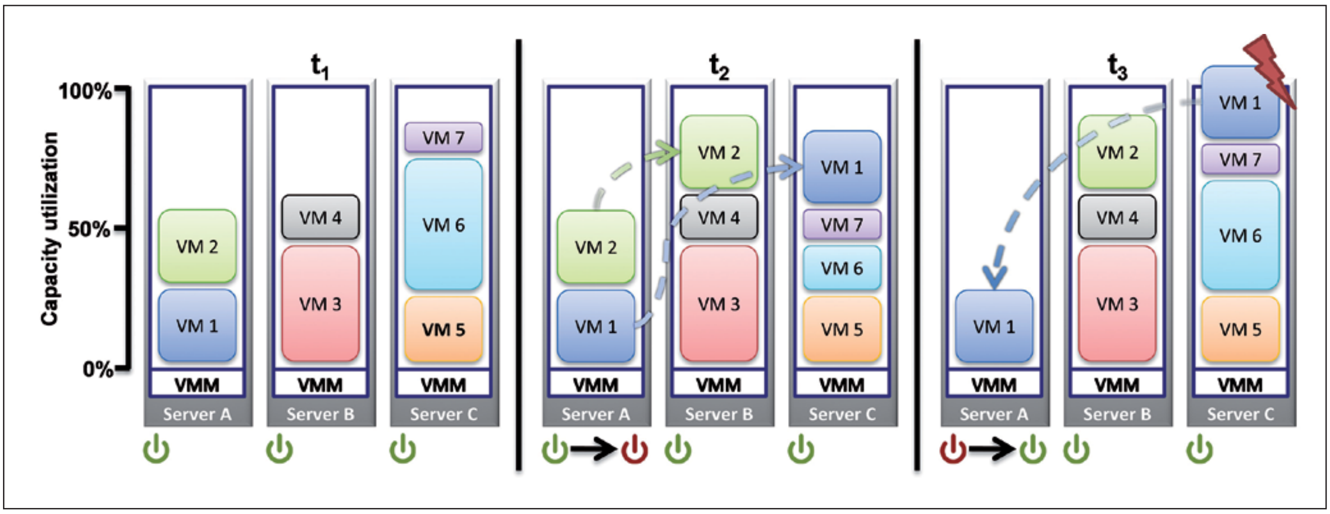

Principle of dynamic load distribution to save energy in virtualised environments. Source: OFFIS

in a queue along with the time stamp of the planned execution and then processed. The times of the actions to be performed are selected such that bottlenecks are not possible provided that the forecasts are correct. This means that when demand for resources increases in the future, the servers are rebooted in time and the VMs concerned are migrated before a bottleneck can occur.

Simulations suggest potential savings of 30 to 55 percent of the server hardware's energy demand.

Measurements taken in a test environment consisting of eight blade servers running 62 VMs were able to confirm the simulation results. Here is was possible to save roughly 50 percent of the energy required to operate the servers.

AC4DC – Adaptive Computing for Green Data Centres

Project Management:
Bernd Hanstein, Rittal GmbH & Co. KG, Herborn
hanstein.b@rittal.de
www.ac4dc.de

3.6. Migration to the cloud

Energy-oriented migration to the cloud

The transfer (migration) of applications and services to an external data centre (cloud) can make sense from a power perspective if renewable energy is available at another site, for instance, or energy-optimised operation is ensured in the data centre. When migrating services to the cloud, IT administrators usually need to answer four main questions:

- What should be migrated to the cloud?
- How should the migration be performed?
- What effect will this have on power consumption?
- How cost-effective is the migration?

The "MIGRATE!" IT2Green project examined the migration of data centre applications to the cloud from these standpoints and worked on effective solutions.

Applications are currently migrated by hand – a considerably time-consuming task. A software-based, automated process is a far more practical option.

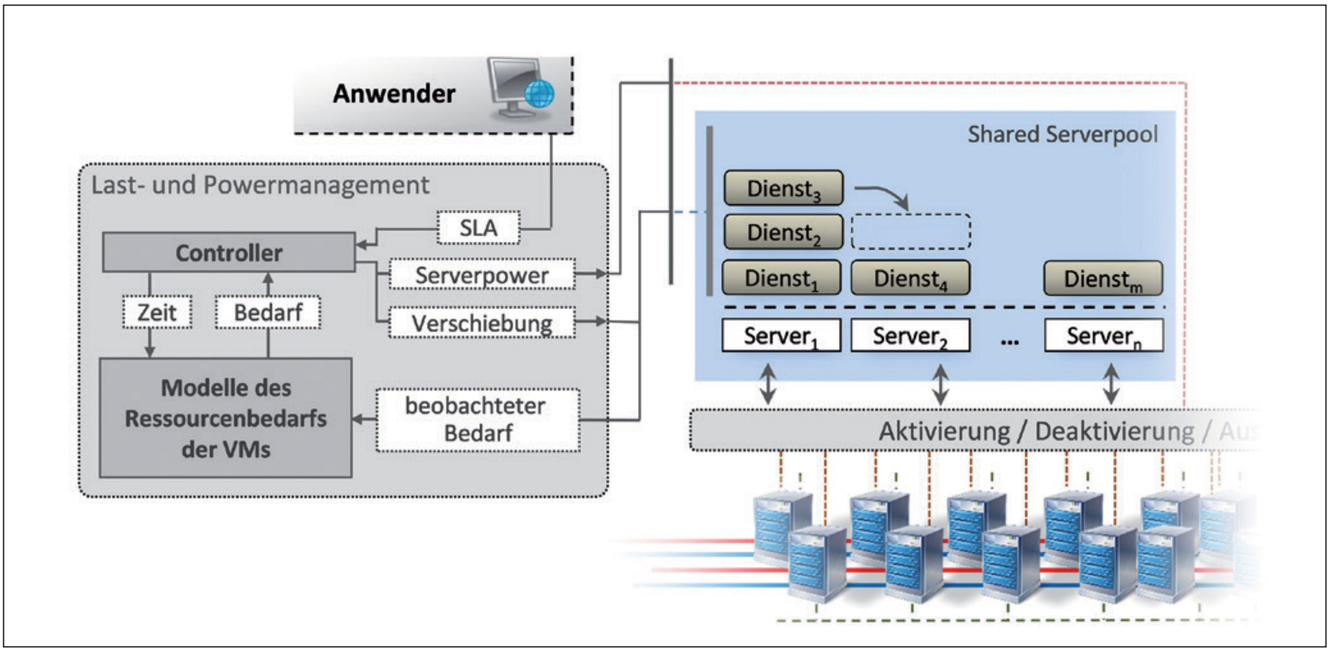

Outline architecture of "AC4DC" load and power management. Source: OFFIS

Control with TOSCA & green policies

If a user infrastructure is migrated to a cloud environment that is more energy-efficient, the cloud provider should also be able to control the application from an energy perspective. For example, the provision of resources should be adapted to anticipated or actual usage patterns.

Energy-related control requires the close technical integration of the application and cloud environment. However, this integration should not lock in the user, as this could result in high switchover costs if the user changes provider later on.

> TOSCA[12] is a new IT industry standard to resolve conflict between close integration and portability.

TOSCA stands for Topology and Orchestration Specification for Cloud Applications and is being developed by the OASIS standardisation organisation with the participation of leading cloud technology providers.[13] The basic idea is to describe applications as cloud services in formal models in such a way that these models can be executed directly by cloud providers for deployment and control in the cloud environment. The description of a cloud service is provided in a TOSCA service template and comprises the topology – i.e. the components – their requirements and relationships as well as plans for deployment and control. TOSCA is vendor-independent and does not make any specifications as regards to particular components, requirements or plans.

The plans contained in a TOSCA service template are very interesting for GreenIT: As the plans define the management of an application in the cloud, they can also be used for energy control. With this in mind, the "MIGRATE!" IT2Green project developed sample green policies and incorporated them as a meta requirement in the TOSCA standardisation system.

A green policy can refer to the entire cloud infrastructure or just individual components. Examples include PUE specifications, compliance with data centre certificates, use

[12] OASIS: Topology and Ochestration – Specification for Cloud Applications – Version 1.0, OASIS Standard, 2013, http://docs.oasis-open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.pdf

[13] OASIS: https://www.oasis-open.org/committees/tosca

of green electricity, server efficiency categories, power consumption of servers, time-related provisioning of resources (times during day/night, tools, seasonal differences) and flexible quality of service levels.

What is important is that green policies are not defined by TOSCA. Instead users can define the policy content individually but in a standardised language.

How is a TOSCA model executed automatically?

Leading providers of cloud management systems involved in the TOSCA standard have already announced products or launched them on the market for this purpose. Model execution is the job of what is known as a TOSCA container, which deploys and controls via the cloud environment hypervisor. In addition to industrial solutions, an open source development by the name of OpenTOSCA[14] is also available.

Cloud migration with a TOSCA container involves the following steps:

1. The user describes the cloud service, along with the green policies, as a TOSCA service template.

2. Using the TOSCA container, the cloud provider performs deployment automatically in a cloud environment while taking the green policies into account.

3. The user uses the cloud service.

4. The cloud environment delivers operating data that are relevant for green policies back to the cloud management system.

5. By comparing operating data and green policies, the cloud provider controls the cloud environment via the TOSCA container.

With these standardised steps, users are able to automatically migrate their applications to the cloud and also define and specify individual green policies at the same time.

MIGRATE! – Models, Methods and Tools for the Migration to Cloud-Based, Energy-Optimised User Infrastructures and their Management

Project Management:
Prof. Dr. Stefan Kirn, Universitat Hohenheim, Stuttgart
stefan.kirn@uni-hohenheim.de
www.migrate-it2green.de

[14] OpenTOSCA: http://demo.opentosca.org

4. Outlook

Sustainable technology design

GreenIT – the design of information and communication technology so that it conserves energy and resources – is likely to grow rather than decline in importance in the years ahead. This is due to the continued increase in energy prices and a range of current market trends:

- The implementation of interconnected sensors in all kinds of objects, and the resulting high volumes of data generated (smart factory, smart building, smart home)
- The smart capture and analysis of these ever expanding sources of data (big data or smart data)
- Automated processing of data including the automatic triggering of actions (ambient intelligence)
- The almost entirely IT-based management of business processes, requiring more and more real-time synchronisation (e-business)
- High performance computing (HPC) to support simulations in research and development, situation analyses and forecasts for the business community and policy makers
- The next generation of smart networks and the opportunities they afford for adaptive adjustment to altered environmental conditions (software defined networks)
- The convergence of stationary and mobile applications and use patterns including Internet use (convergent ICT, cloud)
- Non-volatile memories that are fast and hold large data volumes
- Strict data protection requirements that generate large volumes of data due to coding and multiple duplication

Growing demand and the continued dynamic development of technologies pose a considerable challenge particularly for small and medium-sized enterprises, public authorities and universities. Consolidating and harmonising different targets, such as availability, quality of service and energy efficiency as well as migratability and the handling of legacy systems, requires additional analytical capabilities and a complex data base.

Even today holistic management systems for greater automated control and monitoring of ICT and facility infrastructures or specifically for the data centre sector are becoming firmly established.

However, there are still considerable interface issues in practice.

Energy management system according to ISO 50001

Even compared with other countries, there has been a marked increase in the use of energy management systems in German businesses in recent years. The use of systems certified to ISO 50001 is a prerequisite in Germany if businesses are to qualify for concessions on energy and electricity tax. Prompted by this legislative framework, businesses in the manufacturing industry, in particular, are leading the way in the implementation of such systems.

The deployment of energy management systems requires extensive knowledge of energy usage in the business, thereby also enabling the constant optimisation of in-company energy use. In this respect, the use of methods and tools based on sophisticated IT solutions that are integrated into existing structures provides an opportunity to tap the potential of energy management to the full. Appropriately designed ICT makes it possible to capture detailed energy use data, and identify and leverage energy savings potential within the framework of an energy management system. The integration of different applications using information technology provides a wide range of opportunities.

The involvement of everyone concerned, from staff to management, is key to the success of such a system.

The smart use of suitable ICT can render energy management far more transparent. The results achieved can be documented and presented for specific target audiences to secure acceptance and the continued success of the system. The current policy framework at both the national and European level suggests that there will be greater uptake in the use of energy management systems, so a growing market for appropriate solutions can be expected.

Overall environmental impact assessment

Apart from energy efficiency, more and more attention will also be paid to the overall environmental impact assessment, which not only examines the use of ICT hardware but also the production and disposal of the hardware. While early product replacement can appear to make sense from an energy-efficiency perspective – particularly in the case of short-lived devices – from an overall environmental point of view the opposite is often true, such that the longer use of hardware should be given precedence over early product replacement. In addition to the carbon footprint, the use of critical raw materials, in particular, will become increasingly relevant in the overall environmental impact assessment. There are two reasons for this:

1. There is a growing demand for critical raw materials, such as indium, gallium and rare earth elements, in all high-tech sectors.
2. The efficient recovery of these materials, which are often used in tiny amounts in products, is not possible with today's recycling technologies.

"Green" software also offers additional potential but is as yet in its infancy. The following section examines the approaches that already exist, and the challenges presented by the implementation of green software.

Green software

Apart from addressing the improvement of ICT hardware through the use of more efficient components, for example, up to now GreenIT has also primarily looked at the energy-optimised control of such devices. In end user devices, this is reflected in the form of integrated power management – now a standard feature. In data centres, it primarily takes the form of virtualisation, as well as other solutions for load-adaptive operation as described in the previous section.

Up to now, the programming of applications in a way that conserves energy has not been addressed at all or only in a very selective manner. Known as green software, the aim is to optimise resource consumption by using more efficient software solutions. Creating a conscious awareness for resource efficiency as early as the programming stage is at the forefront of green software. Due to the over-availability of computing and storage power, common programmme today do not communicate with the hardware efficiently. The focus is on the functionality of the programme alone. Either hardware resources are reserved for programme functions that are not required or computing load is generated which is not even needed to run the programme functions the user wants. The programming process itself and the programming language used play a central role here.

Methodological challenges

From a methodological standpoint, green software presents a challenge. Software can never be considered as a separate entity. Instead it is always part of a complex system, be that as part of a device-network-data centre framework, or on an end user device with the specific configuration of an operating system and software used in parallel. Other components of the system interact with the software and make it difficult to make general statements. The most important trends and approaches are described in the following section.

Limiting data traffic

Audio-visual entertainment is the biggest factor driving data traffic and the associated demand for resources. Reducing the volume of data to be transmitted can therefore be a goal of optimisation. A number of aspects need to be taken into consideration in this respect, however: Users have different demands when it comes to the quality of image and video files. To adapt resource requirements to user needs, users are increasingly offered a choice of resolution for the image and video files. In the case of embedded graphics, a resolution that suits the application should be selected, not the best possible resolution.

File compression can be used to reduce the data stream. The reduced stream must be weighed up against the greater computing effort needed to code and decode the files. The networks used for the transmission also play a key role. If the comparatively inefficient mobile communications network is used to send the data, weak file compression carries more weight than when the more efficient LAN or WLAN networks are used. Therefore, in the case of mobile applications, in particular, a software should not dictate the use of the mobile communications network if WLAN is available. Apart from reducing file sizes, it would also be advisable to check how frequently a (mobile) end application needs to communicate with a data centre – be that to search for updates, query status messages or send log files.[15] Efficiently programmed apps also still offer potential for optimisation here.

[15] David Bicknell: Green Tech – Putting GreenIT, sustainability and lean thinking into context: 8 ways to make your software applications more energy efficient, 2012, http://www.computerweekly.com/blogs/greentech/2012/02/making-your-software-applications-more-energy-efficient.html

Programming language and modular software

Enormous potential for optimisation is offered by the improvement of programming languages. Web-based applications, or apps, make increasing use of dynamic programming languages that still communicate with the hardware resources inefficiently. In particular, sufficient support is not yet lent to the parallel use of processors in the form of multithreading. On account of limited battery capacity, these programmes are already programmed with resource efficiency on the end user device in mind, but currently cannot exploit the performance potential of stationary multi-core systems. Furthermore, it must be examined whether the software in the entire system also works efficiently. While the apps addressed are often optimised for operation on mobile devices, they shift a share of the computing operations required to servers in data centres where they work more or less efficiently.

According to studies conducted by BITKOM, half of the functions of standard software programmes are never used by the average user.[16] The modular approach of mobile applications can serve as a model so that users only select features that they actually use. As users will probably choose a maximum number of functions if in doubt, it is necessary to render the features more transparent. Furthermore, the software must make it easy for the user to make settings that conserve resources. It must be noted that updates, in particular, often reset personal settings.

In this respect, attention should be paid to ensuring better version compatibility with regard to customised programmes.

Efficient use of storage

In light of low hardware prices, economic interest in making efficient use of storage space is currently relatively low. Technically, however, this area does offer potential in the form of file compression and specific archiving or deletion of file duplicates (deduplication). If different types of hardware are installed, software can automatically identify the most efficient storage medium to use so that data rarely used in data centres are stored, for instance, on tape systems while frequently fluctuating data are saved to flash memories. Freeing up resources that are not needed increases user performance and can reduce the obsolescence of hardware already installed as this hardware will not need to be replaced so quickly because of higher performance requirements.

Conclusion

While the implementation of green software is still in its infancy, the fundamental approaches to implementation are already there and can be expanded upon and developed further. This is driven by modern-day environmental awareness on the one hand and purely economic considerations on the other. More and more attention will be given to resource efficiency in new software products in the years ahead. Modularity of the programmes and awareness of future software requirements when choosing the programming languages are among the key factors setting the course for green software. By compiling specific programming guidelines and collecting best practice examples, the topic of green software can and should be actively encouraged and promoted.

[16] Ralph Hintemann: 4th Annual Convention of the GreenIT Science Forum: Green Software: Potential for Boosting Resource Efficiency in IT – Specific Approaches to Reduce Energy and Resource Consumption by Software, Borderstep Institute

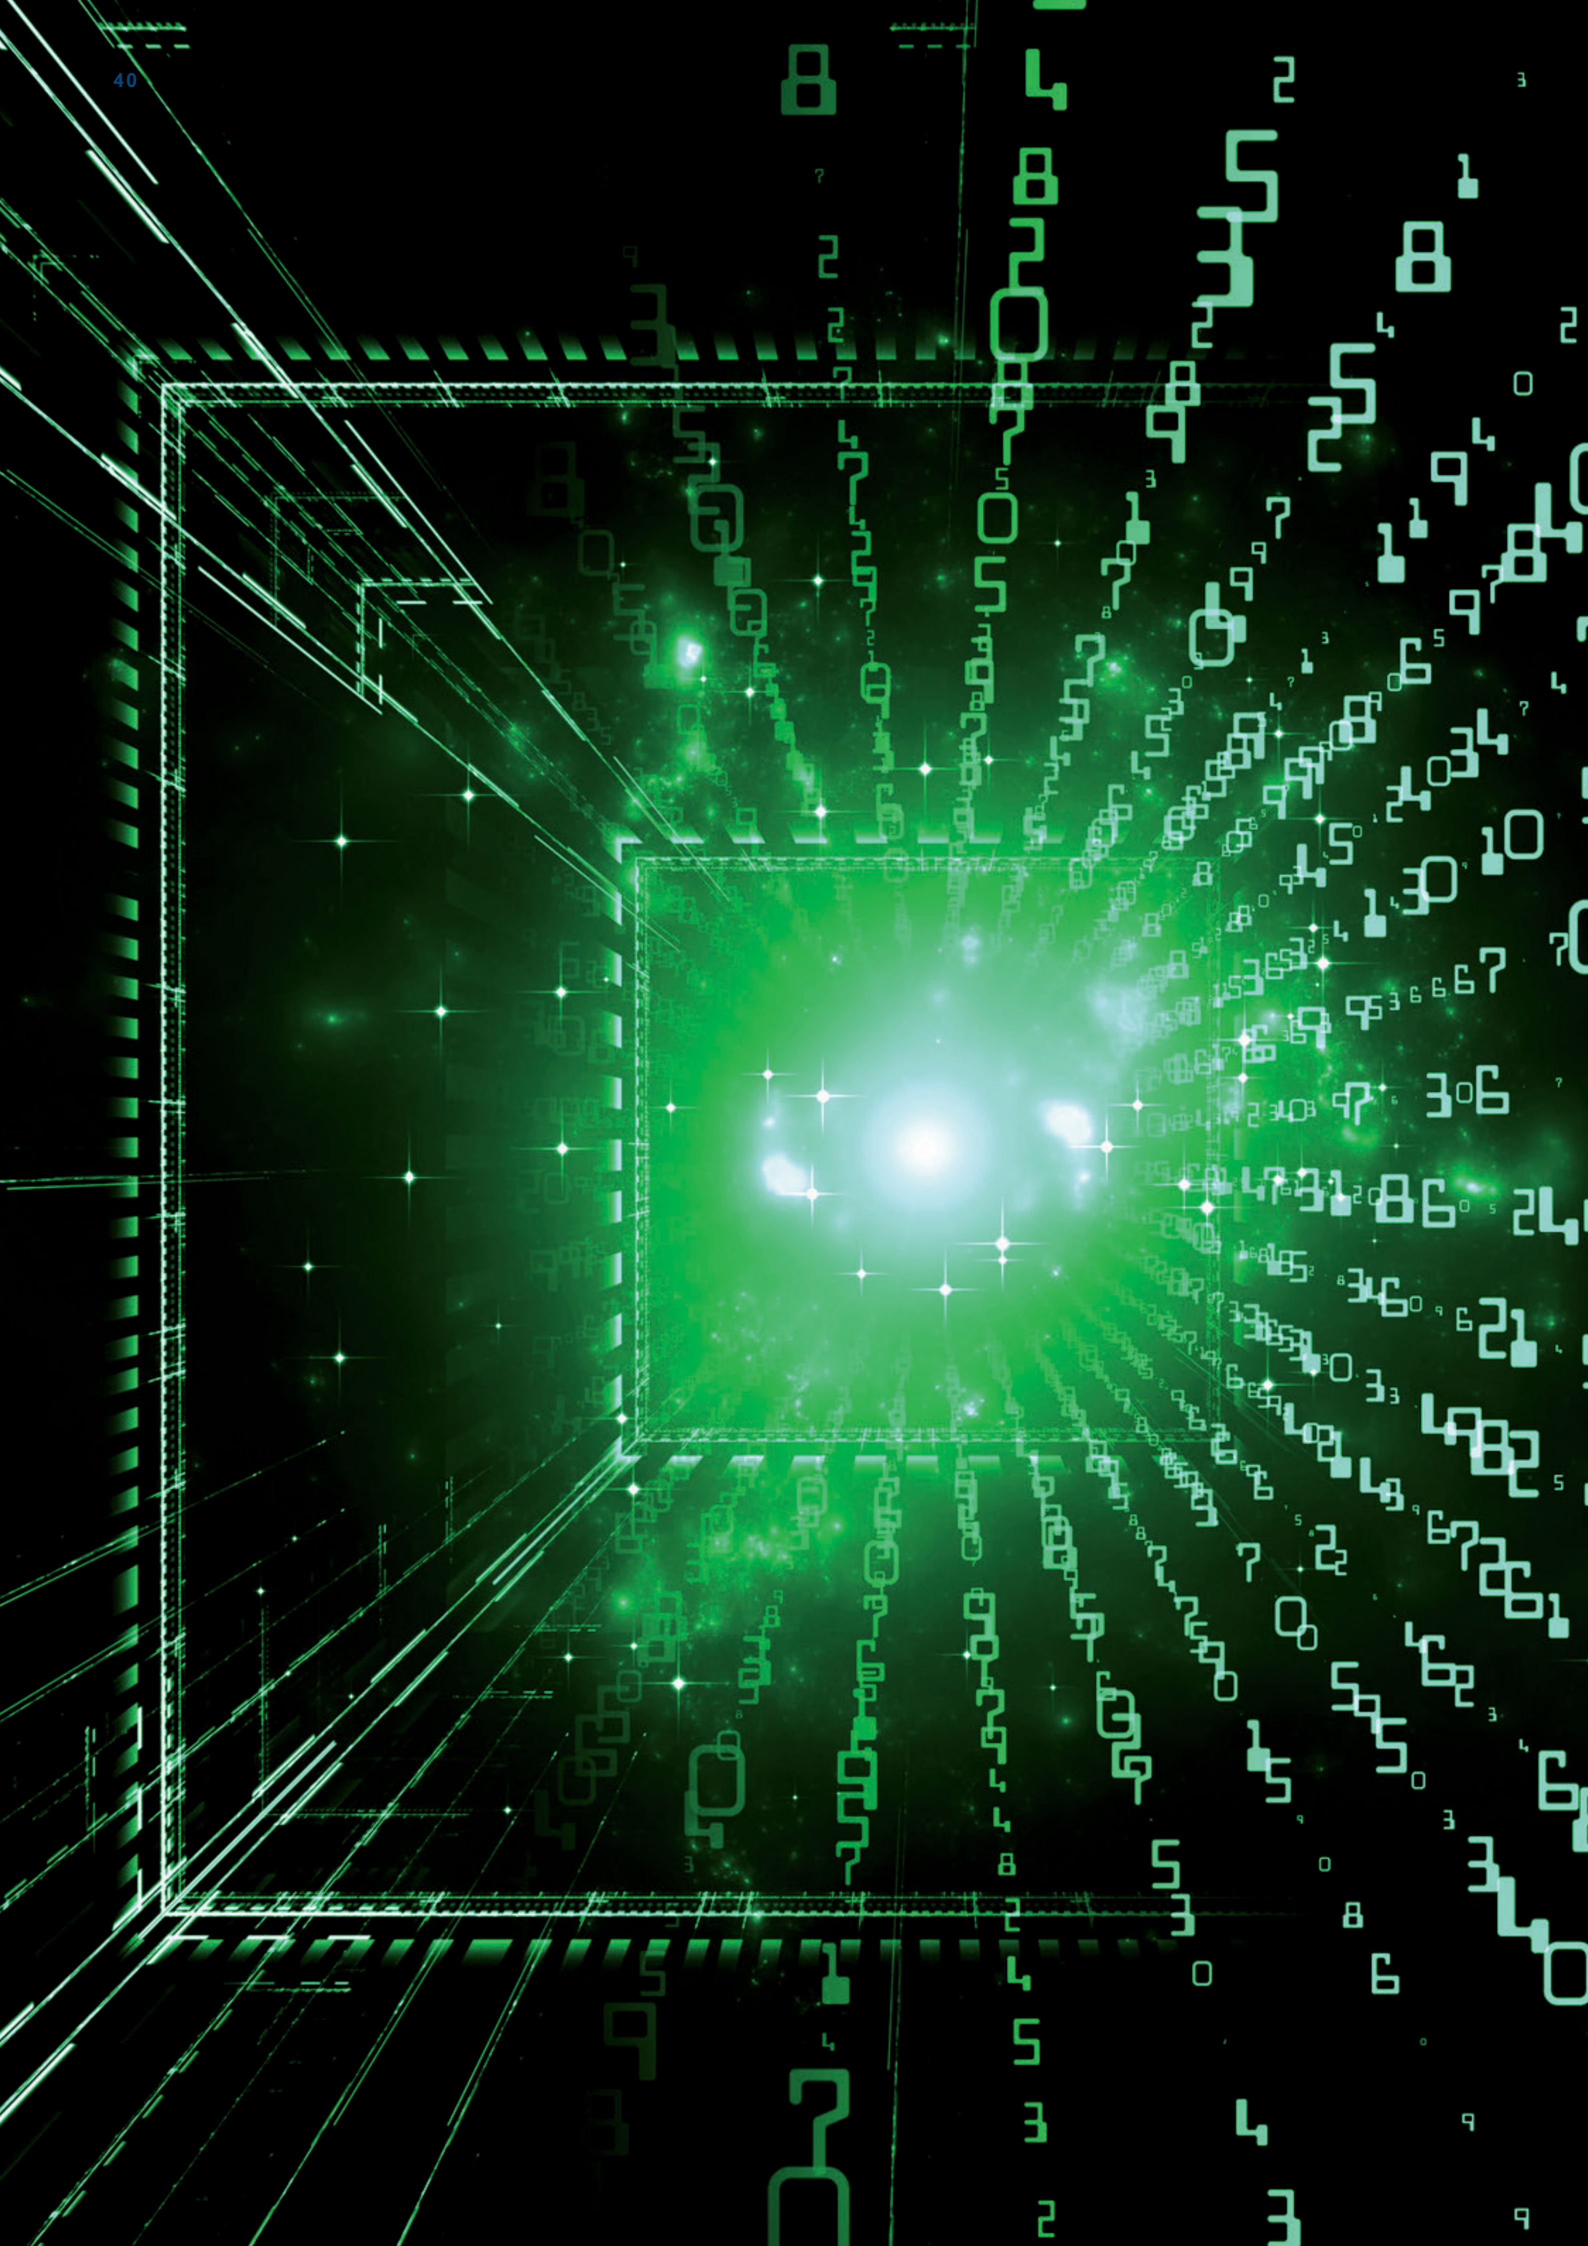

5. IT2Green technology programme

Energy-efficient ICT for SMEs, administration and the home

This brochure was created as part of the technology programme "IT2Green – Energy-efficient ICT for SMEs, administration and the home" of the Federal Ministry for Economic Affairs and Energy and builds on the research results and findings of the projects involved. The IT2Green technology programme funds a total of ten collaborative projects that investigate information and communication technologies in different areas of application with the aim of optimising the energy efficiency of ICT systems. The work of the project therefore does not look at the energy savings that are possible through the use of ICT but rather at the reduction of the power needed to run the ICT systems themselves.

Improving the energy efficiency in the ICT sector is a complex task. On a technical level, this encompasses the development of new hardware and software along with the smart coordination and control of these elements in a product system. At the same, economic and legal framework conditions, in particular, are crucial factors for the overall optimisation of ICT systems which can often constitute a real limiting aspect. The IT2Green collaborative projects pursue integrated, exemplary solutions. They demonstrate new principles and systemic approaches to energy optimisation in the field of ICT, including telecommunications, data centres and end user devices. Customer-defined performance parameters and quality of service (QoS) have a significant bearing on this process and constitute primary framework conditions.

One overarching optimisation approach is, for example, the needs-based provision of IT resources. This means that:

> ICT systems should only be available when they are needed and should automatically switch over to low-power operating modes when not required.

This technical approach requires a broader view of the system. The research projects focus on the smart interconnection of existing sensor technology to actively switch user systems on and off, and also on predictive tools that gather and evaluate a variety of data and use them to forecast demand.

Another approach is the targeted increase in the utilisation of existing IT resources. In contrast to existing technologies such as virtualisation, the IT2Green research projects centre on complex load transfer processes, monitoring and management solutions that consider and control a variety of technical and economic factors. As, in practice, no ICT system resembles the other, the optimum operating point of a data centre or a telecommunications system, for instance, needs to be individually readjusted at regular intervals. The solutions in the IT2Green technology programme realise this advanced system view and, in the process, deploy load transfer concepts that factor in locally favourable conditions. For example, the migration of IT applications between physically separate data centres is examined and tested.

> The goal of the research projects is to improve energy efficiency as a whole.

Correctly measuring any improvement in energy efficiency is a considerable task. The research projects address this issue by developing relevant metrics, for example. The extensive measurements taken in data centres and office buildings have also demonstrated that the software-based consolidation and analysis of measurement data from various sub-sections, such as the climate control and power supply infrastructure or ICT devices, involves much manual work and offers considerable room for improvement.

Finally, the research projects focus on the specific conditions and needs of ICT users in SMEs, administration and the home. Interdisciplinary research teams and regional networks were formed to develop and test energy-efficient ICT systems in realistic settings.

> Close cooperation between research, industry and users enables model projects in which transferable solutions are developed for broad-based use.

Three thematic clusters

IT2Green comprises ten funded projects grouped in three thematic clusters: "Telecommunication Networks", "Data Centres and Clouds" and "Monitoring and Management". With these different focal topics, the research projects concentrate on the specific conditions and needs of ICT users in SMEs, administration and the home to develop and trial practical and realistic solutions for energy-efficient ICT systems.

Cluster: Telecommunication Networks
The "Telecommunication Networks" cluster studies and puts diverse approaches for the load-adaptive operation of telecommunications technology on trial. The central goal is to switch individual network elements to energy-saving modes when low data traffic is expected, without the customer experiencing any loss in the quality of service. Coordinated energy management must be implemented at multiple network levels to this end.

Projects:
- ComGreen (www.communicate-green.de)
- DESI (www.desi-it2green.de)
- IntelliSpektrum (www.intellispektrum.de)

Cluster: Data Centres and Clouds
In the "Data Centres and Clouds" cluster, methods and technologies are developed that make it possible to increase the utilisation of IT systems in an individual data centre or in data centres working as a group. This is intended to allow energy-related locational advantages to be better utilised at any given time.

Projects:
- AC4DC (www.ac4dc.com)
- GGC-Lab (www.ggc-lab.de)
- GreenPAD (www.green-pad.de)
- MIGRATE! (www.migrate-it2green.de)

Cluster: Monitoring and Management
The "Monitoring and Management" thematic cluster develops smart control systems for the needs-based provisioning of IT resources in the office workplace. The aim is to only run the ICT devices at full power if the user is present and to set the IT resources automatically to sleep mode when the user leaves the office workplace.

Projects:
- Adaptive Sense (www.adaptive-sense.de)
- GreenIT Cockpit (www.greenit-cockpit.de)
- pinta (www.pinta-it2green.de)

Accompanying research

To successfully realise the programme's goals, the research projects are supported by the accompanying scientific research conducted at the Fraunhofer Institute for Reliability and Microintegration (IZM) and the Fraunhofer Institute for Systems and Innovation Research (ISI). Responsible for communication is VDI/VDE Innovation + Technik GmbH.

The central task of accompanying research is the scientific support and evaluation of the ten funded research projects. Furthermore, accompanying research also seeks to closely examine the project research results in order to derive generally applicable, transferrable knowledge from them. With regard to cross-cutting issues, three interdisciplinary working groups were set up. Led by accompanying research, overarching issues of the projects were discussed in these groups.

This guide is therefore based on the results of the research projects and the findings of the working groups.

6. Contact and information

Project management

German Aerospace Centre (DLR)
Project Management Agency
Convergent ICT/Multimedia on behalf of the Federal
Ministry for Economic Affairs and Energy
Linder Höhe
51147 Cologne
Germany
www.pt-multimedia.de
it2green@dlr.de

Accompanying research

Fraunhofer Institute for Reliability and
Microintegration (IZM)
Gustav-Meyer-Allee 25
13355 Berlin
Germany
www.izm.fraunhofer.de
it2green@izm.fraunhofer.de

Fraunhofer Institute for
Systems and Innovation Research (ISI)
Breslauer Straße 48
76139 Karlsruhe
Germany
www.isi.fraunhofer.de
it2green@isi.fraunhofer.de

VDI/VDE Innovation + Technik GmbH
Steinplatz 1
10623 Berlin
Germany
www.vdivde-it.de
it2green@vdivde-it.de

www.it2green.de

7. References

Energy-efficient ICT for SMEs, administration and the home

David Bicknell: Green Tech – Putting GreenIT, sustainability and lean thinking into context: 8 ways to make your software applications more energy efficient, 2012, http://www.computerweekly.com/blogs/greentech/2012/02/making-your-software-applications-more-energy-efficient.html

Eric Masanet, Arman Shehabi, Lavanya Ramakrishnan, Jiaqi Liang, Xiaohui Ma, Benjamin Walker, Valerie Hendrix, Pradeep Mantha: The Energy Efficiency Potential of Cloud-based Software: A U.S. Case Study, Lawrence Berkeley National Laboratory, 2013, http://crd.lbl.gov/assets/pubs_presos/ACS/cloud_efficiency_study.pdf

Results of the SPECpower_ssj2008-Benchmarks: http://www.spec.org/power_ssj2008/results/

ErP Preparatory Study on Enterprise Servers (Lot 9), www.ecodesign-servers.eu

Jonathan G. Koomey: Growth in Data Center Electricity Use 2005 to 2010, Stanford University, 2011

Markus Stutz: Carbon Footprint of a Typical Rack Server from Dell, 2011, http://i.dell.com/sites/content/corporate/corp-comm/en/Documents/dell-server-carbon-footprint-whitepaper.pdf

Martin Rost, Unabhängiges Landeszentrum für Datenschutz Schleswig-Holstein (ULD), http://www.it2green.de/de/1223.php

Maximilian Boehner, Sebastian Porombka, Gudrun Oevel: Energieeffizienz im WLAN, Universität Paderborn

OASIS: Topology and Ochestration – Specification for Cloud Applications – Version 1.0, OASIS Standard, 2013, http://docs.oasis-open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.pdf

OASIS: https://www.oasis-open.org/committees/tosca

OpenTOSCA: http://demo.opentosca.org

Ralph Hintemann: 4. Jahrestagung Wissenschaftsforum GreenIT: Grüne Software: Potenziale zur Steigerung der Ressourceneffizienz in der IT – Konkrete Ansatzpunkte zur Reduktion von Energie- und Ressourcenverbrauch durch Software, Borderstep Institut

SPEC Power: SPECpower_ssj2008 Results, http://www.spec.org/power_ssj2008/results/

Steven Lanzisera, Bruce Nordman, and Richard E. Brown: Data Network Equipment Energy Use and Savings Potential in Buildings, Lawrence Berkeley National Laboratory, ACEEE Summer Study on Energy Efficiency in Buildings, 2010, http://www.aceee.org/files/proceedings/2010/data/papers/2195.pdf

Technical Whitepaper: HP Moonshot System, http://h20565.www2.hp.com/portal/site/hpsc/template.BINARYPORTLET/public/kb/docDisplay/resource.process/?spf_p.tpst=kbDocDisplay_ws_BI&spf_p.rid_kbDocDisplay=docDisplayResURL&javax.portlet.begCacheTok=com.vignette.cachetoken&spf_p.rst_kbDocDisplay=wsrp-resourceState%3DdocId%253Demr_na-c03728406-3%257CdocLocale%253D&javax.portlet.endCacheTok=com.vignette.cachetoken

"Trusted Cloud" (Federal Ministry for Economic Affairs and Energy): www.trusted-cloud.de

Links

- IT2Green: Green-IT Dossier (compilation of information on the topic of GreenIT), http://www.it2green.de/de/1365.php
- Blauer Engel Rechenzentren: http://www.blauer-engel.de/de/produkte_marken/vergabegrundlage.php?id=226
- Energy Star (US): www.energystar.gov
- Energy Star (EU): www.eu-energystar.org/de
- JRC Code of Conduct for Energy-efficient Data Centres: http://iet.jrc.ec.europa.eu/energyefficiency/ict-codes-conduct/data-centres-energy-efficiency
- Standard Performance Evaluation Corporation: www.spec.org

www.bmwi.de