

Energy Efficiency Indicators for Data Centers

Bernard Aebischer, Centre for Energy Policy and Economics (CEPE), Swiss Federal Institute of Technology (ETHZ), Zurich, Switzerland

Huston Eubank, Rocky Mountain Institute (RMI), Snowmass, CO, USA

William Tschudi, Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA, USA

Abstract

Three complementary and synergistic projects dealing with indicators of energy efficiency in data centers are presented in this paper.

A recent comprehensive data center Charette¹ conducted by RMI identified concepts that could lead to data centers that would cost less to build, work better, and save astonishing amounts of electricity. The participants recommended detailed metrics that could be useful in evaluating and influencing energy efficiency. An independent organization that would develop, demonstrate, and implement comprehensive and generally applicable cost, performance and energy intensity indicators was also proposed.

LBNL has benchmarked energy consumption in a number of operating data centers. The metrics that were used allow energy efficiency comparison of various system and component designs and may provide indication of overall performance. A large sample of measured performance will be useful in identifying attainable good performance based upon current practice and in setting future performance goals.

For public policy applications, energy efficiency indicators must be robust and easy to determine. One indicator for data centers, the ratio of IT equipment electricity use to the total was proposed for use in evaluating construction permits and continuous energy efficiency monitoring. The outcomes of the LBNL and RMI projects may lead to the use of a detailed system of energy efficiency indicators including data center systems and components by providing credible target values that can be referenced in standards and point the way to even more efficient practices.

1. Introduction and problem setting

A commonly used measure of the energy performance of commercial buildings is the average annual energy consumption per unit of floor area. In order not to compare apples and oranges one often corrects energy consumption for climate differences and classifies buildings according to their HVAC (e.g. fully air-conditioned building, building with mechanical ventilation, building with no cooling and no mechanical ventilation) and/or according to the economic sector or building type (e.g. office building, hospital, school). For all these building configurations, energy and electricity consumption are dominated by HVAC and lighting and the mean electric load intensity is typically of the order of 10-100 W per square meter.

In some parts of some buildings energy consumption is dominated by specific processes such as cooking, food storage, washing, or data processing (computer rooms). Mean electric loads may locally

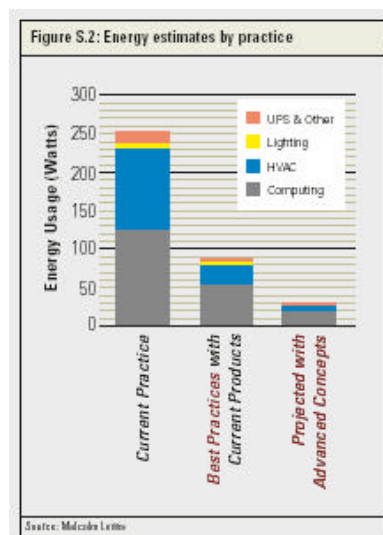
¹ The term "charrette" is adopted from the storied practice of Ecole des Beaux Arts architectural students in nineteenth century Paris who reputedly could be seen still drawing their projects until the last minute as they were carried on the "cart" or en charrette on the way to the design jury. In its modern-day adaptation, charrette refers to an intensive design workshop involving people working together under compressed deadlines.

be of the order of several hundred watts per square meter. If these high electric-load zones represent a substantial fraction of building energy consumption, then the measure of energy consumption per unit of floor area is no longer a good measure of energy efficiency of the building. One may therefore use another indicator or process metric, such as energy per meal or per guest for a restaurant, or exclude this process-induced energy consumption in the calculation of energy per unit of floor area and consider it in a separate approach.

In this paper we present three projects looking at the question of how to measure energy efficiency in data centers. A first project is a report (Aebischer et al., 2003) on behalf of the government of Geneva making a proposition how to implement a voluntary agreement between Authorities, NGOs and operators of data centers into the administrative process of construction permits and in the planned continuous monitoring of energy efficiency by all large energy users.

A comprehensive design workshop of data centers that cost less to build, work better and save astonishing amounts of electricity is the second example (RMI, 2003). On invitation by the Rocky Mountain Institute (RMI), three dozen experts in all relevant topics of data center business and engineering came together and sketched in much detail an innovative data center, consuming 70% to 90% less energy than a typical modern data center (figure 1) for the same computing performance. A detailed description of the outcomes of the design charrette is given in another paper presented at this conference (Eubank et al., 2004).

Figure 1. Power demand of a 1-U box for computing in a data center environment today, best practice and with advanced concepts (RMI, 2003, p. 14)



Ongoing research activities at Lawrence Berkeley National Laboratory (LBNL) (<http://datacenters.lbl.gov/>) is the third example. LBNL has benchmarked energy use in fourteen data centers and is in the process of benchmarking a number of additional operating data centers. This work quantifies the energy consumed in data centers by breaking it into its end use. The data obtained is for IT processing loads and key building systems and is obtained through direct measurement in most cases. The benchmarks reveal the current energy intensity of IT equipment, and the relative performance of building infrastructure systems. The metrics allow comparison of various types of systems and their relative efficiencies. The better performing data centers and their systems will then be analyzed to determine the practices that lead to better performance. A summary of some of the benchmarking results and energy efficiency improvement ideas are presented.

The common topic

One common topic of these three studies is the question of how to measure energy efficiency. This aspect is looked at from three different points of view, the planning process, the scientific analysis aimed at understanding the differences in the use of energy in existing data centers and the implementation in a policy framework.

Different actors and different purposes need different indicators to measure efficiency or performance. A useful classification (framework for analysis) is the life cycle of a data center – i.e.: the planning, design, construction permitting, operation, renewing and dismantling phases - and the respectively relevant actors or stakeholders: the developer, planner, designer, operator or client of a data center and the utility, the government and in some cases NGOs on the other side.

There are more synergies between the three projects:

- The canton of Geneva introduced limits and target values for indicators of energy efficiency, but the empirical base for the specific values is poor. The workshop where RMI gathered many experts and practitioners in that field was a excellent opportunity to discuss these values and the outcomes can be used to improve and refine the requirements in the future. The empirical data collected and the thorough analysis done by LBNL may become the starting point for a new and better-funded regulation.
- The RMI workshop and the LBNL analysis – by their nature – propose many more and alternative indicators to the one proposed in the Geneva study. Detailed measures of energy efficiency at the level of components and sub-systems are needed as an alternative to the Geneva approach and could become mandatory in the case where targets at the aggregate level are not reached. The detailed regulation serves as "stick" in the process where the carrot is a voluntary target at a rather aggregate level.

2. Different indicators with specific characteristics for different purposes proposed in the framework of the RMI-charrette

In a recent comprehensive design workshop of data centers that cost less to build, work better and save astonishing amounts of electricity, recommendations regarding benchmarking is a key output (RMI, 2003, 6.4): "Gathering and benchmarking operating data about computing facilities and data centers is essential, and is a key recommendation of the charrette. Feedback on costs is essential both for operations (short run) and planning (long run) of data flow and processing capacity. The data must be globally available, transparent, and translatable across boundaries...". Some of the recommended metrics include the computational output, e.g. kW per unit of computational output; others consider the efficiency of the basic infrastructure, e.g. by the ratio of electrical computer equipment load to the total building/data center load; and others look at more technical aspects, like air changes per hour in computer room or power transformer efficiency. Some of the recommended metrics can be calculated for well-documented data centers; others are not yet in use. An independent organization is proposed to work – among many other aspects – on the development, demonstration and implementation of a comprehensive and generally applicable set of indicators to measure cost, performance and energy intensity of data centers.

"One metric for comparing the efficiency of data centers proposed at the charrette is total utility power delivered to the facility divided by the net power that goes directly into computing equipment. (RMI, 2003, S. 15) This metric is commonly known to specialists as the "delivery factor". The next figure shows this delivery factor for the three scenarios of figure 1. The inverse metrics – energy going into the equipment divided by total energy – is the indicator C1 discussed in more detail in section 4.

Figure 2. Total power delivered to the data center divided by the net power that goes directly into computing equipment (source: (RMI, 2003, p. 15)

Figure S.4		
Total Power / Computing Power =	With Current Improvements in Computing	Holding Computer Power Constant
Current Practice:	1.97	1.97
Best Practice with Current Products:	1.54	0.38
Projected with Advanced Concepts:	1.36	0.13

Other recommended metrics include:

- Metric of computational output—kW per unit of computational output;
- kW per rack equivalent—allows tracking of “packing factor”;
- UPS efficiency or losses—ratio of total kW in to UPS power output, kW of HVAC/kW of UPS;
- Plug-process load W — W/ft^2 nameplate energy labeling for peak, end use, idle, power supply efficiency;
- Total kW demand per kW provided to the servers (a measure of parasitic power demand) or to all IT equipment, or the ratio of electrical computer equipment load to the total building or data center load (this would be a measure of the infrastructural energy efficiency);
- Cooling—kW/ton, ft^2/ton , unit of cooling per unit of data-processing;
- Air recirculation— cfm/ft^2 , W/cfm , air changes per hour in computer room;
- Power transformer efficiency—percent efficient;
- Lighting— W/ft^2 as used (net of any control savings); and
- Effective air infiltration or leakage area effect.

(RMI, 2003, p. 76)

On the level of the IT equipment (servers), the experts discussed the interest and the feasibility to extend the Energy Star approach to these equipment. It became clear that "several things are needed before Energy Star ratings can be applied to servers and data centers. These include:

- numerical, quantifiable statements about energy usage;
- good metrics (flops/W, calculations/cycle, etc.); and
- good baselines (find a good model somewhere).

As a first step, consider creating an Energy Star rating that focuses only on power supply and fan efficiency. This avoids the problem of defining performance metrics because these measures are independent of the processor that is used in the server." (RMI, 2003, p. 80)

At this conference, the RMI-charrette is described in detail in a separate paper: “High Performance Data Centers” (Eubank et al., 2004).

3. Analysis of energy benchmarking leads to energy indicators – ongoing work at LBNL

3.1 Introduction

Although electrical power demand in data centers is high, there currently is little information concerning energy efficiency issues. This lack of energy efficiency focus is true for the “processing” of information (e.g. transactions/W) and the building infrastructure. There is little energy benchmark data available to provide information to highlight what can be achieved in the design of new systems.

The benchmarking reported here involves a strategy to obtain energy end use breakdown and collect information on the operating efficiency of infrastructure systems commonly in use in data centers. The main prerequisite to developing meaningful energy efficiency indicators applicable for data centers is a large enough sample of measured energy performance. Lawrence Berkeley National Laboratory (LBNL) has benchmarked energy use in fourteen data centers and is in the process of benchmarking up to ten additional operating data centers. This work quantifies the energy consumed in data centers by breaking it into its end use. The data obtained is for IT processing loads and key building systems and is obtained through direct measurement in most cases. The benchmarks reveal the current energy intensity of IT equipment, and the relative performance of building infrastructure systems. The metrics allow comparison of various types of systems and their relative efficiencies. The better performing data centers and their systems will then be analyzed to determine the practices that lead to better performance. A summary of some of the benchmarking results and energy efficiency improvement ideas are presented.

Data centers are prevalent in a wide range of industries, universities, and government facilities. Energy demand in these facilities is thought to be growing due to computing technology changes and

IT professionals seek to maximize computing per square foot of data center. In the past, many methods have been used to estimate and quantify energy intensity. As a result, there has been considerable confusion over IT equipment's electrical use (energy intensity) in current configurations, and possible increases as IT equipment evolves in the future. Research aimed at understanding the present electrical intensity, total facility end-use, and key facility systems' performance was undertaken with a short term goal of identifying the better performing systems and how that performance was obtained. A review of the summary benchmark results is performed to identify more efficient systems and practices, and the benchmarks can also discover operational problems. This supports the ultimate goal of identifying and/or developing energy efficiency improvements in these buildings.

Metrics used in this study allow comparison of the current power density of computing equipment (in terms of kW/sq ft), and provide indicators of the efficiency of key facility systems (e.g. kW/ton of chilled water). One metric was included that may provide insight into the efficiency of the heating, ventilating and air conditioning (HVAC) system in performing its role of cooling the computing equipment.

The data collected can be used to understand current performance, to establish design and operational criteria for new projects, to identify current best practices in design and operation, and suggest ways to improve reliability. This information may help identify performance indicators that can be determined for use in public policy decisions. A review of the current best practices is likely to identify gaps where additional research is needed to achieve a new level of improved energy efficiency.

Definitions

Acronym	Definition
HVAC	Heating, Ventilation, and Air Conditioning
UPS	Uninterruptible Power Supply
CRAC unit	Computer Room Air Conditioning unit
IT	Information Technology

3.2 Background

Energy demand of today's IT equipment and for systems removing the heat they produce is high compared to ordinary commercial buildings. Although data centers contain various types of computing equipment, building systems in data centers usually have similar characteristics and can account for more than 50% of the total energy use in a data center facility. Unfortunately, many of these systems are oversized which causes them to operate inefficiently at partial load conditions. This may be caused by many factors such as provision for future growth, misunderstanding true IT equipment loads, liability for undersized systems, etc. A segment of the data center market where space for IT equipment is leased in critical facilities even uses excess infrastructure as a selling point, resulting in a trend to further oversize electrical and HVAC systems.

Planning for the future, whether at the utility level, facility level, or computing equipment level has been a challenge due to uncertainty in the heat intensity of future IT equipment and the overall amount of computing that will be required. To improve this situation, a good starting point is to better understand the current operating conditions.

Previously, little publicly available energy benchmark data existed for data centers. Confidentiality of facility operating information is important to a majority of data center operators and this has hindered the dissemination of reliable building energy performance. Data center operators typically track whole building energy use and the energy used by the computing equipment. What has been lacking is a consistent method of obtaining and comparing electric power density (Watts/ft² or m²), energy end use, computational work produced, and efficiency comparisons of key facility systems (HVAC, UPS, lighting, etc.).

Electrical requirements for data center infrastructure systems are dominated by HVAC systems however other losses such losses in uninterruptible power supply equipment can also be significant.

Typically chilled water production in a large central plant provides the cooling for air systems that move large quantities of air throughout the data center. Current technology typically relies on air cooling of IT equipment using energy intensive HVAC systems often using specialized computer room air-conditioning (CRAC) units, however other large air handlers may be used. Typically, environmental conditions are tightly controlled in order to protect the IT equipment and this leads to additional energy use. Overall electrical demand is relatively constant, reflecting the constant power consumption - and resulting heat production - within IT equipment.

3.3 Performance indicators identified through benchmarking

Performance Indicators can be found by examining a number of the metrics that can be measured in an operating data center. Performance metrics can be very useful in comparing data centers and their facility systems. The metrics allow energy efficiency comparison of widely varying data centers regardless of the design, and the types of computing equipment. These metrics illustrate measured electric intensity, which is useful to trend overall load growth and to predict future needs. They also provide insight into how efficiently the building systems were designed and are operating. Energy use and systems operational information was obtained primarily on chillers, UPS systems, and CRAC units. This data was obtained by connecting power sensors to the host electrical panels, or by direct readout from equipment meters, if in existence. Additional operational data, such as flow, and temperature measurements were obtained from existing facility management systems to the extent they were available and were supplemented by direct measurement if not readily available.

Fourteen data centers in eleven facilities (where three facilities had two data centers each) were included in the initial study performed by LBNL. Previous case studies and summary benchmarking data are available through the LBNL website: <http://datacenters.lbl.gov>

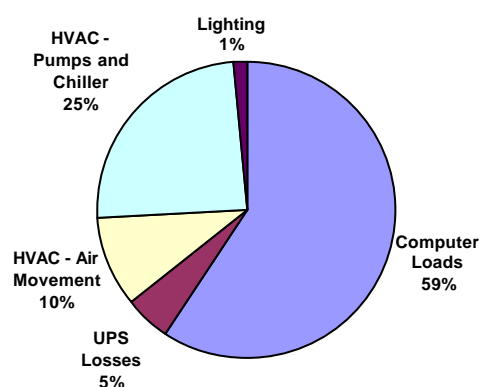
To develop a more robust data set, on-going work is targeting an additional six to ten data centers. Once this additional information is available, the relative performance of a data center facility and its systems will be studied. This information will be used to focus on the better performing systems and identify the practices that enabled them to achieve higher performance. The information will also be valuable to form a baseline so that heat intensity trends can be evaluated. In addition, a mechanism for self-evaluation is planned through development of a "self-benchmarking protocol", that would allow a data center owner/operator to compare his data center's performance to a larger sampling of data centers. This information should improve the ability to predict future power requirements and size systems more efficiently.

Appendices I and II summarize the metrics and other information obtained through this study.

3.4 Benchmark Results

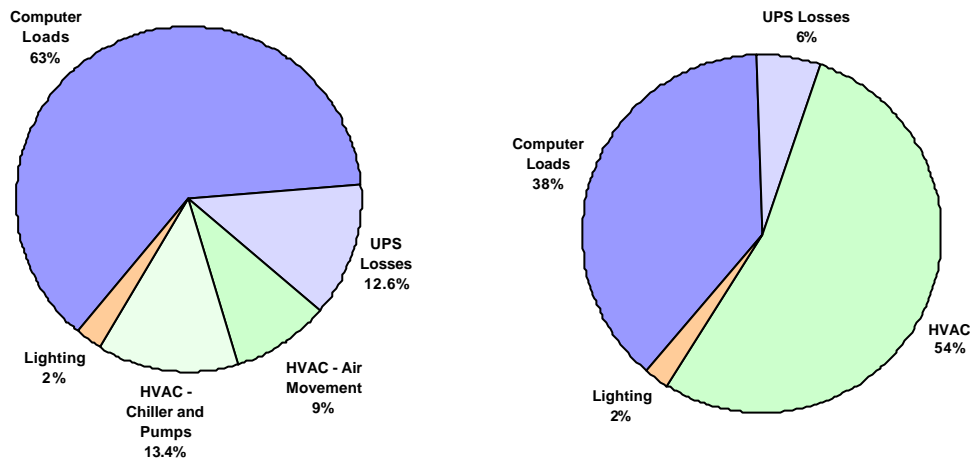
End use energy breakdowns are obtained for each of the data centers. Figure 3 illustrates typical energy end-use information that is provided by the benchmark measurements.

Figure 3. Representative energy end use breakdown



Typically, the end uses that were measured consisted of the IT equipment (fed from UPS systems and power distribution units), UPS system losses, HVAC – chilled water plant, HVAC – computer room air conditioners, and lighting. The relative percentages of each of these systems varied according to the computing load intensity and the efficiency of the infrastructure systems. For example, the percentage of power delivered to the computing equipment varied between 33% and 73%. This indicates, in part, that the efficiency of the infrastructure systems varies considerably. Similarly, the other end use components varied substantially as shown in the examples in figure 4.

Figure 4. Benchmarking examples



The data center on the left utilized a highly efficient system that was thoughtfully designed using best practices with better than standard HVAC components and controls. The data center on the right utilized traditional computer room air conditioners. These benchmarks begin to provide some insight into how to achieve better performance.

Indicators of other systems' performance were also evident. UPS systems, lighting, and other systems are compared. These benchmarks can lead to surprising revelations. In one facility, benchmarking discovered that the entire cooling for the IT equipment was being handled through the make up air (house) system, yet all of the computer room air conditioners were operating moving air throughout the data center thereby creating unnecessary fan heat that added to the cooling load.

HVAC Systems

By focusing on the various HVAC systems and their components, the benchmark data reveals that energy use can vary by factors of 3 or more for systems that serve essentially the same purpose. A possible indicator of cooling efficiency is being investigated. By comparing the energy used for cooling the data center (i.e., the HVAC power in kW) to the energy used for the IT equipment (in kW), an indicator of HVAC system performance is obtained. A lower value may indicate that the system is more energy efficient because proportionally less HVAC energy is required to cool the energy consumed by the IT equipment. This metric is defined as follows:

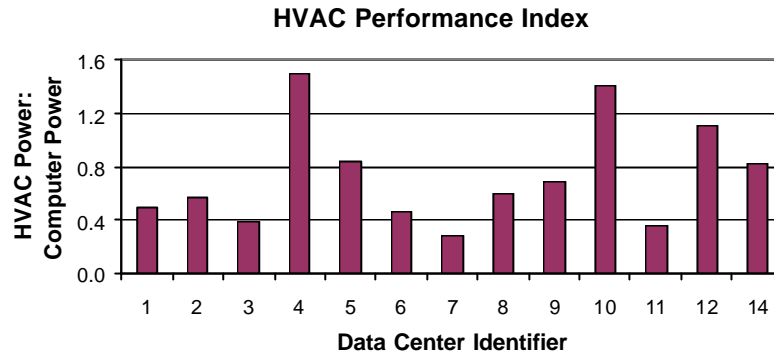
$$HVACperformanceindex(\%) = \frac{kW_{HVAC}}{kW_{UPSOutput}}$$

Another way of looking at this indicator has been proposed by the Uptime Institute, a leading data center industry association. This organization proposes a ratio of the total infrastructure systems' energy use (primarily HVAC, but including other systems and losses) to total energy use. This indicator could provide insight into the relative efficiency of HVAC systems, but could also signal efficiency issues with other systems.

Many different HVAC system designs have been studied. Figure 5 shows a comparison for 13 system configurations measured to date in this study. This information suggests that there is wide variation in

the performance of the systems due to system design. This wide variation underscores the need to understand the features and principles of the more efficient systems. LBNL plans to investigate the attributes of the better performing systems and document the “best practices” that are observed. This investigation should provide valuable insight into the validity of the HVAC performance index as an indicator of performance.

Figure 5. Relative data center HVAC system performance



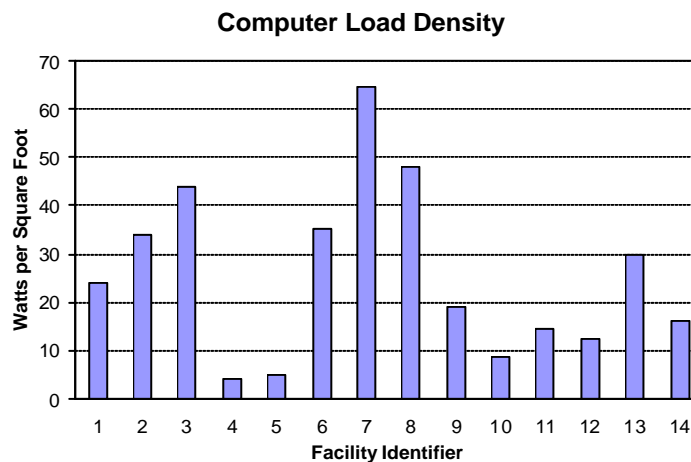
Further investigation into the efficiency of HVAC systems includes examining chilled water and computer room air handling systems. For the chilled water plant, a traditional efficiency metric, kW/Ton, is used. Chiller, pumping, and fan energy (for cooling towers) and the corresponding tons of chilled water produced were obtained. Wide variations in efficiency (dominated by the chiller efficiency) were observed.

Computer room air conditioning energy was similarly determined. Here, it is more difficult to obtain accurate airflow measurements – typically delivered by many air handlers into common underfloor areas or a network of ducting. As a result, an efficiency metric such as kW/cfm is not generally obtainable. We instead rely on the comparison of overall HVAC performance. This is an area where further research could pinpoint additional efficiency issues.

Computing Loads

The heat produced through electrical consumption in IT equipment is typically removed by air-cooling. The electric power density (in terms of W/sq.ft.) can vary significantly from data center to data center depending upon the amount and type of IT equipment employed. Figure 6 shows some results of the measured electric power density of the IT equipment alone. In the calculation to determine this metric, the Uptime Institute’s definition of “electrically active” floor area is used in the denominator. This effectively excludes areas such as walkways or storage spaces, which are more likely to have electric power density similar to commercial office buildings.

Figure 6. Benchmarked computer load densities (1 Watt per square foot = 10 W per square meter)



Uncertainties in predicting the computing equipment loads make HVAC system sizing a challenge. Measured electric power density is typically quite less than what is specified during design. This occurs for several reasons: facility designers may use name plate data (which typically leads to loads that are several times greater than actual); future equipment power densities are uncertain; unnecessary or compounding conservatisms are applied; computing equipment (and load) typically is added gradually over time as business needs dictate; etc. In addition, reliability strategies require that multiple, redundant equipment be available. This introduces significant inefficiency for standby or part load operation. While the computing load varies for each application, measured data from facilities with similar computing missions will help “right-size” the cooling equipment. Infrastructure systems are often more efficient when operated near their full design load. This study found computing loads at all facilities below 65 W/sq ft, which is well below their design values and the intensities frequently quoted. Use of benchmark data can lead to better prediction of design loads and better build-out strategies. Designing systems and components in closer alignment with actual operating loads will also lead to more efficient operation.

3.5 Summary

Energy benchmarking results can be used as indicators of overall data center performance as well as individual system efficiency. Measured energy use determined by a benchmarking program can provide a baseline for tracking energy performance over time. It can be used to better predict future needs leading to more efficient sizing of supporting facility systems. Benchmarking can also be used to prioritize where resources need to be applied to achieve improvements in energy efficiency.

Use of the benchmarking metrics also provides a mechanism for comparison of the efficiency of facility systems and components between different data centers. This is possible even though the system design and configuration may be completely different. By analyzing the results, better designs can be identified. Large apparent variations in system or component energy use may signify design, installation, operational, or maintenance problems. Examining poorly performing systems could identify operational or maintenance problems or could lead to methods to correct inefficiencies originally built into the facility. For data center designers, access to actual operational data will highlight better practices and lead to new creative energy efficient designs and operating strategies.

Additional insight on important data center efficiency topics may be found in LBNL’s research roadmap for High Performance Data Centers and is available through the LBNL website. On-going research will provide additional benchmark results, develop a “best practices” summary, develop a self-benchmarking protocol, investigate power supply efficiency in IT equipment, investigate UPS system efficiency opportunity, and demonstrate new or emerging technologies.

4. Energy efficiency indicator proposed to be used in a policy process – outcomes of a study on behalf of the Canton of Geneva

4.1 Why Geneva is concerned about data centers

In the last couple of years ICT -companies promoted facilities and services where servers, routers and the like are operated in data centers or “server hotels”. One major feature of these facilities is their particularly high electric load from 200 to more than 1000W/m². In 1999/2000 the Authority of the Canton of Geneva faced several applications of construction-permits for data centers, with the potential to increase considerably the electricity consumption of the Canton of Geneva and therefore compromising the energy-policy of the Government aiming to reduce electricity consumption in the Canton.

In (Cremer et al., 2003) we did a compilation of estimates of energy consumption by data centers in different countries (table 1). It is of the order of a few tenths of percent of total electricity consumption and is increasing fast. These values are means for an entire country. Locally, the share in total electricity demand of data centers may be much higher. In Geneva, the electricity demand of the new data centers planned in 1999/2000 was of the order of 10 to 20% of the total electricity demand of the Canton!

Table 1. Electricity demand of data centers in fraction of total electricity demand of the country and in electricity per inhabitant (Source: Cremer et al. (2003))

Country	Year	Fraction of total electricity demand, in %	Energy demand by capita, kWh/cap	Evaluation method	Source
CH	1988	2.2%	143	Equipment	Spreng/Aebischer 1993
CH	1999	0.8%	57	Equipment, computer-networks	Aebischer et al. 2002d
USA	2000	0.3%	58	Equipment	Roth et al. 2002
USA	2000	0.6%	123	Equipment, computer-networks, phone-network	Roth et al. 2002
USA	1998?	0.6%	100	Equipment	Kawamoto et al. 2000
USA	2000	0.1%	20	Floor area	Mitchell-Jackson 2001
USA	2000	0.1%	15	Floor area	Beck 2001
NL	2000	0.6%		Floor area	Hartkamp 2002

4.2 Monitoring of data centers in Switzerland in the 1990s

In Switzerland, energy consumption of 10 computer centers – mainly of large banking and insurance companies - was surveyed for several years (Jund, 1996). Unfortunately only two annual power-load measurements (one in summer and one in winter) were performed and most of the detailed information is confidential. Nevertheless, some interesting information, like the repartition of energy (table 2) could be derived from the data (Aebischer et al. 2003). It is compared to more recent measurements.

Table 2. Electricity end use in % of total electricity demand of data/computer centers (Source: Cremer et al. (2003), LBNL (2003))

	Aebischer 1992	Aebischer et al. 2002 min.	Aebischer et al. 2002 max.	Mitchell-Jackson 2001	LBNL 2003
ICT-equipment	42 %	43 %	63 %	49 %	52%
Losses in UPS	6 %				
Losses in transferring el. current	4 %	10 %	10 %	11 %	9.9
Cooling	17 %	30 %	20 %	14 %	19%
Ventillation	28 %	15 %	5 %	23 %	16%
Lighting	2 %	2 %	2 %	3 %	3%

These banking and insurance companies developed an indicator of energy efficiency for their computer centers. It is defined by the ratio of the electricity used by the computer system itself to the electricity consumption of the entire computer center, i.e. computer system and infrastructure needed to operate the computer system.

4.3 Energy efficiency indicators proposed to be used by the Canton of Geneva

A study commissioned by the Authority of the Canton of Geneva (Aebischer et al., 2003) explores how to measure the energy efficiency of data centers and how to integrate it in a policy framework aimed at promoting energy efficiency as well in the phase of planning and construction as in the phase of operation of the energy-consuming facilities.

Ideally, energy efficiency of a data center should be measured in terms of energy consumption per unit of service delivered to the customer. However, there exists no commonly agreed method to measure the service provided by a data center. Even if a standardized method of measuring the service did exist, its value would probably vary so fast – due to technological progress – that it would not be possible to define any reference value necessary to evaluate efficiency of a data center. Using the floor area as a reference for the energy consumption does not make much sense in a data center. Indeed, the same electricity consuming equipment may either be dispersed over a large area – leading to a low value of electricity per unit of floor area – or densely packed together in fully-

equipped racks resulting in a high value of energy per m², with no difference in the specific electricity consumption per unit of energy service. As an alternative measure of energy efficiency in data centers we propose to use the CEE concept described in the next section.

To be useful in a policy process, an indicator must give indications of potential improvements. By choosing smart technical solutions the electricity consumption of a typical central infrastructure may be substantially reduced. In a non-optimized data center C1 is of the order of 50% to 60%, to be compared to the efficiency of 75% in the optimized data center (in table 3, C1 = ICT equipment / Total).

Table 3. Characteristics of infrastructure and relative electricity consumption of the main end uses for an optimized, a conventional and an inefficient layout and operation of the infrastructure in data centers in Geneva (Source: Aebischer et al. (2003)) (Weather conditions in Geneva: a few days below 0° Celsius in Winter and a few days with over 30° Celsius during day-time and over 20° during night-time in summer).

	optimised infrastructure ³⁾	conventional infrastructure ³⁾	inefficient infrastructure ³⁾	(Mitchell-Jackson 2001)
shares based on:	kWh/a	kWh/a	kWh/a	kW
free-cooling	yes	yes	no	unknown
computer room temperature	26°C	22°C	22°C	unknown
cold-water temperature	13/19°C	11/17°C	6/12°C	unknown
COP chillers	4	2.5	2.5	unknown
supply air temperature	14°C	12°C	12°C	unknown
pressure loss in CRAC	350Pa	500Pa	900Pa	unknown
fan efficiency	65%	60%	55%	unknown
ICT equipment	74.20%	62.20%	50.30%	48.50%
HVAC	14.80%	21.80%	27.70%	36.9% ¹⁾
Light	2.00%	3.00%	4.00%	3.40%
Power distribution unit	2.00%	4.00%	5.00%	²⁾
UPS	5.00%	7.00%	10.00%	²⁾
Others	2.00%	2.00%	3.00%	11.20%
Total	100.00%	100.00%	100.00%	100.00%

1): 13.6% for central chiller plant, 23.3% for fans, CRAC (computer room air conditioning) units, AHUs (air handling units).
2): Included in "Others".
3): Simulation by Altenburger (2001)


4.4 The CEE concept

The Coefficient of Energy Efficiency (CEE) expresses the ratio of the electricity consumed by processors, hard disks and the like (u, so-called "useful electricity consumption", see figure 7) divided by the electricity purchased from the utility or produced on-site (T).

$$CEE = C1 * C2. = U/T * u/t = u/T$$

C1 is a measure of the energy efficiency of the infrastructure's design and operation. The coefficient C2 is a measure of the losses on the level of the specific "production apparatus" itself and represents – in the case of data centers – a rather ambitious and innovative approach. A detailed description of CEE can be found in Appendix 3. The technical feasibility of the CEE-concept was explored by investigating whether the two components of CEE, i.e. C1 (measure of efficiency of the infrastructure) and C2 (measure of the efficiency of the equipment), can be determined and whether potentials and strategies to reduce the losses can be identified.

Figure 7. Scheme of a colocation site and its electricity consuming parts. Source: Aebischer et al. (2003)



The use of C2 in the Geneva policy process had to be abandoned, because no way could be found to measure the efficiency of the computing equipment in a sufficiently accurate way. Energy policy measures to foster energy efficiency of ICT equipment is rather limited on a local and even national level.

On the other hand, measuring C1 is possible and widely done. The same metric was used by some computer centers in Switzerland in the 1990s (section 4.2). A measuring concept to determine C1 is shown in Appendix 4. In most of the existing data centers it can be implemented with little (new computer centers) or reasonable (existing computer centers) investments, but in some centers, e.g. smaller computer rooms that share some infrastructure (e.g. production of chilled water) with other users, substantial investments would be needed.

The values for C1 to be used in the construction permitting process and in the follow-up monitoring process were mainly derived from simulation calculations. As starting values we recommended the following target values²: C1 > 0.65 to be reached by a new data center in the construction-permission procedure and in the follow-up monitoring process, and C1 > 0.55 for existing data centers (Aebischer et al., 2003).

4.5 Policies to favor energy-efficiency in data centers (and other important energy consumers)

Aebischer et al. (2003) discussed scenarios (S2-S4) that are all based on target values or mandatory values (standards) for C1 and/or C2. These values are used in the process to get construction and operation permits and in the later phase of monitoring energy consumption. A first scenario (S1) describes the legal procedure effective today.

- The scenario “procedural permission (S1)” is a description of the approach defined in a recent revision of the legislation in the Canton of Geneva. It postulates the involvement of the promoter of a new building in a four-stage permission-procedure (preliminary authorization, final authorization and operation permission), which includes the elaboration of an energy concept.
- The scenario “voluntary agreement on target values (S2)” basically envisages an energy self-regulation. Both decision and control are de facto delegated to data centers or to other large consumers. It is inspired from the two approaches “Energimodell Zürich” und “Energimodell Schweiz” summarized in (Aebischer et al., 2003).
- The scenario “formal authorization based on mandatory values (S3)” corresponds to a rather traditional understanding of authorization procedures, such as building permits. It proposes to

² Valid for temperate weather conditions (see Table 3)

make energy decisions according to the procedures already in place for construction regulation. In other words, formal authorization increases the public sector's power while it is only sporadically wielded by it or is delegated to private actors (for example the data centers themselves, indirect control by environmental NGOs).

- The scenario "integrated control (S4)" goes a step further in the direction of strict legal and administrative monitoring of the conditions for authorization. It increases the role of the state, since it implies rigorous state control and sanctions to ensure that conditionally authorized installations and equipment then actually conform to set standards. Here, private economic and/or environmental actors no longer play a role.

The CEE index can be used in all the scenarios as an indicator of the target value and/or as a monitoring index. Its status and its usefulness, however, vary from case to case. In the second scenario, its final value can be negotiated before it becomes a target value. In scenarios 3 and 4, its value is measured for each case and is then compared to the target value set in the regulations. It becomes a criterion for granting authorization, and afterwards, monitoring the CEE index allows measurement of how well the fixed objectives are attained.

The described scenarios in their present state do not contain any instruments for economic incentives (e.g. incentive tax or subsidy, a system of financial reward and punishment according to the energy efficiency of the data centers, electricity tax). It is however perfectly possible to include such instruments and to combine them with the authorization procedure which lies at the heart of the proposed scenarios.

5. Conclusions – outlook

There is no single best indicator. The appropriateness and usefulness of an indicator always depends on the application of the indicator. Data center energy efficiency indicators may be classified in different ways.

In the planning and construction phase of a data center many different professions are involved: architects; civil, HVAC, and electrical engineers; computer scientists; and many others. All these specialists have their standard ways of working and their specific indicators. The kind of indicators they use are in general technical characteristics of components, equipment or sub-systems of data centers. Working together and considering the important aspects of other disciplines was the big challenge of the RMI-charrette. The outcome – a proposal for an integrated planning approach – shows the way to data centers that cost less to build, work better and save astonishing amounts of electricity.

More than 50 recommendations filling almost 50 pages is a lot of information. RMI is well aware that "A significant amount of education is required for the creation of more efficient data centers" (RMI, 2003, p. 83). Design workshops in other parts of the US and the world are envisaged. But a broad and fast diffusion needs many other initiatives. Some important aspects are treated in a planning aids manual on energy efficiency in HVAC for data centers, which is in preparation on behalf of the Swiss Federal Office of Energy.

For applications in a policy process, energy efficiency indicators must be robust and easy to determine. A pragmatic solution like the one using C1 as an efficiency indicator is the most promising way to get a tool that can realistically be hoped to be used in near future in a policy process of target setting and monitoring.

The voluntary target setting approach proposed for Geneva is inspired by today's practice in Zurich for large energy consumers. If participants in that program do not reach the targets they have to fulfill a multitude of technical specifications for components and systems. Such detailed specifications are needed to implement a similar scheme for data centers.

The ongoing work at LBNL is essential for different applications:

- measured data and benchmarking is needed to define target values and/or standards which are credible
- detailed metrics for components and sub-systems are needed in order to make the voluntary target value for aggregated indicators attractive
- better performing systems can be studied to determine how they achieved better efficiency
- standard benchmarking protocols being developed will facilitate getting a large sample of benchmark data
- research into efficiency opportunities in IT equipment power supplies and UPS systems has the potential to dramatically reduce overall energy consumption

The activities at LBNL could be a starting point for the strong recommendation made at the RMI-workshop regarding the establishment of an independent organization to work – among many other aspects – on the development, demonstration and implementation of a comprehensive and generally applicable set of indicators to measure cost, performance and energy intensity of data centers.

References

Aebischer, B., R. Frischknecht, Ch. Genoud, A. Huser, F. Varone, 2003: Energy- and Eco-Efficiency of Data Centres. Report commissioned by the Canton of Geneva, Geneva, Switzerland
http://www.cepe.ethz.ch/download/projects/data_centres_final_report_05012003.pdf

Aebischer, B. et al., 2002. CO₂-Reduktionspotential Erdgas. Projektphase 1: Referenzszenario. Zürich, 16. April

Aebischer, B., 1992. Zentrale Rechenanlagen. In "Strom rationell nutzen. RAVEL Handbuch". vdf-Verlag Zürich, Mai (ISBN 3 7281 1830 3)

Altenburger, A., 2001: Electricity consumption of HVAC equipment in data centres. Personal communication (12.12.2001). Amstein + Walthert, Zürich

Cremer, C., W. Eichhammer, M. Friedewald, P. Georgieff, S. Rieth-Hoerst, B. Schlomann, P. Zoch; B. Aebischer, A. Huser, 2003: Energy Consumption of Information and Communication Technology (ICT) in Germany up to 2010. Karlsruhe/Zürich <http://www.cepe.ethz.ch/download/projects/INFO-KOM/ISI%2BCEPE-ICT-english.pdf>

Eubank, H. et al., 2004: High Performance Data Centers. International Conference on Improving Energy Efficiency in Commercial Buildings, 19 - 20 April 2004, Frankfurt, Germany

RMI, 2003: Design Recommendations for High-Performance Data Centers. Report of the Integrated Design Charrette, 2-5 February 2003. Rocky Mountain Institute, Snowmass, CO
<http://www.rmi.org/store/pid2424.php>

Jund, E., 1996: Energieeffizienz in Rechenzentren. Kennzahlen von 14 Rechenzentren in der Schweiz. SBG LIEG/LIHA, Zürich

LBNL, 2003: Data Centers Case Studies. Case Study Reports 1-10 <http://datacenters.lbl.gov>

LBNL websites:

<http://datacenters.lbl.gov>

<http://enduse.lbl.gov/Projects/InfoTech.html>

Mitchell-Jackson, J. D., 2001: Energy needs in an Internet economy: A closer look at data centers. University of California, Berkeley

Spreng, D., B. Aebischer, 1990: Computer als Stromverbraucher. Schweizer Ingenieur und Architekt Nr. 50, p. 1459-1463.

Appendix I. Data Center Metrics

Whole Building Electrical Power:	kW
Load Intensity: Data Center floor area Total load density Computing load density HVAC load density	square feet (sq ft) W/ sq ft W/ sq ft W/ sq ft
HVAC: Chiller plant Chiller Efficiency Chilled Water Plant Efficiency Chiller load Data Center Load HVAC air systems CRAC unit fan, and humidity control energy Central air handling fan power Air handler fan efficiency (where possible to obtain) External temperature and humidity	kW/Ton kW/Ton Tons Tons Cubic Feet per Minute per kW (CFM/kW) °F, %
Data Center Electrical power demand: UPS Loss Computer load (from UPS Power) HVAC - chilled water plant (if central plant exists) HVAC - central air handling, and/ or CRAC Unit energy Lighting	kW kW kW kW kW
Design Data: Design basis for Computer load Design basis for Chilled Water, air side HVAC, and UPS Systems	kW/sq ft Temperature, Humidity, Flow rate, % Efficiency, Total load, etc.

Appendix II. Additional Data Center Information

Features and System Descriptions	Example Descriptions
HVAC	Central water-cooled chilled water plant, central air handling system with VAV control
	Distributed air-cooled CRAC units
	Air-cooled chillers with CRAC units suppl. air under floor
	Central air handlers use outside air economizers
Variable-speed-drives	Centrifugal chiller with VFD
	Primary/ Secondary with VFD
	Central air handler with VFD
Electrical Distribution	N+1 UPS's
	N+1 at the PDU level
	Backup power generators
Control Strategies	Multiple cooling towers operated in parallel
	Minimum number of chillers operated
	CRAC units in empty areas turned off.
	Humidity control disabled on CRAC units
	VAV system with duct static pressure of 0.75"
	Chilled water set point fixed at 50 °F
	Condenser water set point fixed at 70 °F
Chiller kW/Ton monitored continually	
Air side economizers used on	
Temperature and Humidity Set points	Return air temperature maintained at 70 °F ± 5 °F
	Supply air temp. central air handlers maintained at 50 °F
	Relative humidity maintained at 50 % ± 10%
Redundancy/Reliability	N+1 at UPS level
	N+1 at PDU level
Estimate of Occupancy	Data center is 40% full - physical capacity
	Operating at 30% of UPS capacity

Appendix III. CEE (definition)

The Coefficient of Energy Efficiency (CEE) measures the efficiency of electricity consumption of a data center. CEE expresses the ratio of the electricity consumed by processors, hard disks and the like (u, so-called "useful electricity consumption", see figure 7 in the section 4) divided by the electricity purchased from the utility or produced on-site (T). But, it does not measure the energy efficiency of processing, storing and transmitting information depending on the hardware, but also the software used.

$$CEE = u / T = C1 * C2$$

with

$$C1 = U / T = U / (U + CO+VE+ME+UP+TR+LI+OT)$$

and

$$C2 = u / t = u / (u + co+ve+me+up+tr+ot)$$

Abbreviations:

- U = t total el. power load / consumption of equipment
- CO electricity used for production of cold, refrigeration
- VE electricity used for ventilation, evacuation of heat by air
- ME electricity used for other mechanical work, e.g. pumps
- UP electricity used for uninterruptible power supply
- TR electricity used for transformation and correction of electric. power
- LI electricity used for lighting
- OT electricity used for others (miscellaneous consumers and losses)

Depending on the measurement concept and due to its inferior relevance "lighting" (LI) is often including in "others" (OT).

- u "useful" el. power load / consumption of equipment
- co electricity used for production of cold, refrigeration not included in C1
- ve electricity used for ventilation, evacuation of heat by air not included in C1
- me electricity used for other mechanical work, e.g. pumps not included in C1
- up electricity used for uninterruptible power supply not included in C1
- tr electricity used for transformation and correction of electric. power not included in C1
- ot electricity used for others not included in C1

Appendix IV. Measurement concept to determine the components of C1 in a data center (developed by A. Huser, encontrol GmbH)

