

Characterization of Data Center Energy Performance

● David F. Snelling ● C. Sven van den Berghe

This paper presents an intuitive, two-parameter metric for fully describing the energy efficiency of data centers (DCs). The metric accurately characterizes the energy performance of a DC from when it is first commissioned through to full capacity and thus can be used to predict future performance and inform deployment policy. The metric also describes the theoretical ideal performance of DCs and can therefore be used to compare DCs of different sizes at different stages of deployment or in different phases of design and development. Application of this metric to two Fujitsu DCs, 600 kW and 3 MW in terms of IT power, demonstrated that it is accurate with respect to both simulation and measurement results.

1. Introduction

The efficiency with which data centers (DCs) use their energy supply is a significant issue for cost and environmental reasons. Collaborative efforts to increase efficiency within the DC industry have been facilitated by a consensus on an easily understood and easily measured metric for the efficiency of a DC and on a method for tracking its change over time.

In an agreement reached between the U.S., the EU, and Japan in February 2010, power usage effectiveness (PUE) was adopted as the cornerstone of a strategy for harmonizing DC metrics aimed at energy efficiency. The overall strategy for promoting the use of harmonized metrics to effectively achieve energy efficiency in DCs is built around two concepts, as paraphrased below from the press release.¹⁾

- 1) Measure the actual information technology (IT) work output of the DC compared to actual energy consumption of the DC as a whole. The two aspects of this metric can and need to be developed independently.
- a) IT efficiency: Measure the IT work output compared to the energy consumed by the IT

equipment.

- b) DC efficiency: Measure the DC infrastructure efficiency, i.e., PUE.
- 2) Measure the use of renewable energy technologies and the re-use of energy to reduce carbon consumption.

The harmonization effort resulted in the selection of PUE as the metric to be used for international discussions and comparisons regarding the DC efficiency metric. The details of how PUE should be calculated²⁾ were laid out by the Green Grid (TGG) consortium.³⁾ In particular, annual measurements of energy input to the DC (kWh_{IN}) and energy output from the power distribution units (PDUs) to the IT equipment (kWh_{IT}) are used to compute PUE (kWh_{IN} / kWh_{IT}). For reasons that will become apparent, we will use the reciprocal of PUE, DC infrastructure efficiency (DCiE), in the remainder of this paper.

The use of PUE as a metric to derive DC efficiency is limited by a number of weaknesses (as discussed in Section 2) and by its focus on the current operational environment of the DC. An improved metric would enable discussion of the

inherent efficiency of a DC's design and how the DC is operating in relation to its potential. We derive such a metric in Section 3 and discuss its application to actual DCs in Section 4.

2. DCiE weaknesses

There are a number of known weaknesses in the measurement and use of DCiE in isolation for measuring DC efficiency. Before we can discuss them, however, we need to establish a basic understanding of the definition of DCiE.

DCiE is the ratio of the energy consumed by the IT equipment to the energy used by the DC facility as a whole, ideally measured over a year. Thus, the more efficient the DC infrastructure, the closer to 1.0 is DCiE. Mechanical and electrical (M&E) equipment, which are not included in the measurement of IT, include transformers, uninterruptible power supply (UPS) systems, power transmission and distribution facilities, air conditioning, chillers, water pumps, and to a lesser extent lighting and fire suppression systems.

2.1 Measurement point

The practicalities of measuring IT power consumption mean that the point at which the energy used by the IT equipment is measured can vary. For example, one could assume that all energy coming from the UPS systems was dedicated to the IT equipment. While this might be true, it does not take into account non-IT energy lost in the PDUs.

The two-parameter metric proposed here assumes that the measurement point is at the PDU output although it does not rely on this assumption.

2.2 Aggregation

While the DCiE metric can be computed continuously or aggregated over any period of time (e.g., a day, a week, a month, a year), only the longer term aggregations reflect the true capabilities of a DC. Any DC will be more efficient

if the measurement is done for an hour in the middle of a cold winter's night than for an hour at high noon in summertime. The harmonization effort resulted in the recommendation of using yearly aggregations of the DCiE measurements to ensure that the effects of most climate conditions are incorporated into the measurements.

The metric proposed here follows this guideline but can be applied on any time scale.

2.3 Life cycle

Since DCiE is a unitless ratio, it is of little use in assessing the actual amount of energy used by a DC. Economies of scale can greatly affect the efficiency of DCs, as can factors associated with the life cycle stage of the DC. For example, a new, state-of-the-art DC that is only fractionally utilized will typically seem inefficient according to DCiE measurements. This is due to the base load imposed by the M&E equipment. In a fully utilized DC, the base load will be a much smaller fraction of the total load, so the DCiE will be dominated by proportional loads.

The commercial and regulatory effect of these life cycle differences creates a significant problem for DC operators. Customers are already seeking information on the energy efficiency of service providers—including DCs. DCiE can easily be used to misrepresent the efficiency of a DC, or conversely a new DC can easily be undervalued by customers simply because its utilization is not yet near its maximum capacity.

The proposed metric is aimed at ameliorating this problem and thus provides a way to assess DCs independently of their life cycle stage.

2.4 Other anomalies and problems

In general there are a number of anomalies and problems that arise as a result of the traditional DCiE approach to measuring DC efficiency.

- **QoS:** Higher tier^{note 1)} DCs typically have a lower DCiE due to the higher level of redundancy in their M&E systems.
- **False economy:** A major program to upgrade the IT equipment and technology in a DC (e.g., to reduce energy costs and CO₂ output) can result in a lower DCiE simply because it reduces the relative contribution of the proportional aspects of DCiE in favor of the fixed M&E load.
- **Erroneous predictions:** Using DCiE to predict behavior or plan provisioning is risky. For example, inaccurate predictions of the rate at which DCiE will improve as a DC is provisioned can lead to greater than expected energy bills.
- **Hypothetical DCs:** Informed discussion about hypothetical DCs is nearly impossible since DCiE is only meaningful when a DC is under a specific provisioning load. However, in terms of long-term strategic planning, DC operators need to be able to assess DC technology in hypothetical settings.
- **Data availability:** Frequently, access to measured data for a given DC is almost impossible to obtain. Most DCs operate in very static configurations, so any data that is available does not provide insight into trends over time or between configurations.
- **Complex alternatives:** Alternative metrics to DCiE, such as DC performance per energy (DPPE)⁴⁾ from the Green IT Promotion Council,⁵⁾ provide more information but do so at the cost of complexity. For example, DPPE is a composite metric based on four independent terms linked by an arbitrary weighting scheme.

In the remainder of this paper, we describe our two-parameter metric aimed at addressing these issues. Our primary focus

note 1) DC “classes” were defined by the Uptime Institute⁶⁾ to describe reliability and availability. Higher tier DCs are more reliable.

is on its life cycle independence, that is, the ability to use it to characterize a DC’s energy efficiency independently of its provisioning load. DCs operate over a number of years, and, while instantaneous measurements of efficiency are needed, there are also cases in which understanding the potential of a DC is also important. For example, highly modular DCs become more efficient sooner in their provisioning cycle, so we need to understand and quantify this in making provisioning decisions. Large DCs are in general more efficient but only once fully loaded, so the full value of a large DC is only realized over its full life cycle.

3. Metric derivation

From the forgoing discussion, two factors stand out as playing a key role in the characterization of DC energy performance: 1) DCs are most energy efficient when fully provisioned, and 2) the rate at which a DC increases in efficiency as it is provisioned significantly affects its full life cycle efficiency.

Figure 1 depicts the DCiE of two hypothetical DCs as their level of provisioning increases. DC ‘A’ is slightly more efficient when fully provisioned than DC ‘B’; however, DC ‘B’ achieves a reasonable level of efficiency sooner in the provisioning cycle than DC ‘A’.

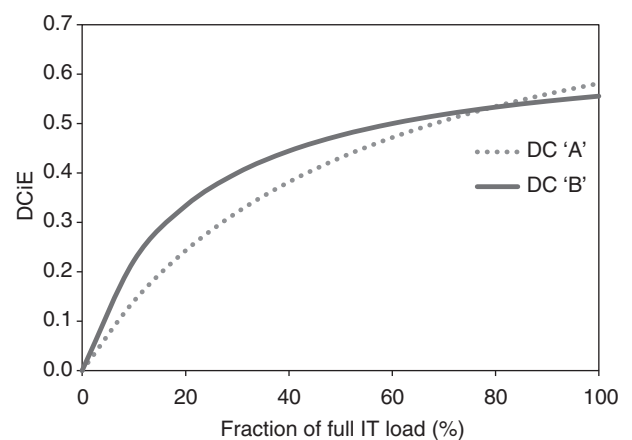


Figure 1 Example DCiE plot for two hypothetical DCs.

If we compared the DCiE of these DCs only in their fully loaded condition, we would conclude that DC ‘A’ was the more efficient. If we made this comparison when they were only partly provisioned, our conclusion would depend on the relative provisioning levels. This simple example clearly highlights the importance of the two key factors listed above.

Our proposed two-parameter metric is aimed at addressing both factors. The metric’s parameters are defined as follows.^{note 2)}

- DCiE_∞: The theoretical, asymptotic maximum efficiency of a DC.
- P_{1/2}: The IT load at which the DC performs at half its maximum theoretical efficiency, i.e., DCiE_∞. P_{1/2} is expressed in kW but may also be represented as a percentage of the maximum designed IT load of the DC.

To understand the formulation of these parameters, we present an example set of measurements from a Fujitsu DC (DC1).^{note 3)} **Figure 2** is a plot of the IT load versus the total site load, both measured in kW.

The apparent linearity suggests that the M&E load increases linearly with the IT load. DCs operating in a non-linear regime will tend to have reduced efficiency at high load, so DC design capacities tend to be specified so that operations stay within the linear part. This effect is detectable in Figure 2 only at the highest load point. Therefore, assuming that the apparent linear relationship above holds in general, we




Figure 2
Plot of average annual IT energy load against average annual total site load.

can describe the average annual IT load (L_{IT}) in terms of the average annual site load (L_{DC}).

$$L_{IT} = \alpha + \beta \times L_{DC} \tag{1}$$

Then, letting L_{ME} be the M&E load, where $L_{DC} = L_{IT} + L_{ME}$, we can define DCiE:

$$\begin{aligned} DCiE &= L_{IT}/L_{DC} \\ &= L_{IT}/(L_{IT} + L_{ME}) \\ &= \beta \times (L_{IT}/(L_{IT} - \alpha)). \end{aligned} \tag{2}$$

The theoretical maximum DCiE for a DC, as L_{IT} tends to infinity, is therefore

$$DCiE_{\infty} = \beta. \tag{3}$$

Substituting equation (1) into equation (2) shows that a DC reaches half this efficiency ($\beta/2$) when the IT load equals $-\alpha$. That is,

$$P_{1/2} = -\alpha. \tag{4}$$

Therefore, α and β characterize the DC and are the slope and intercept of the IT load plotted against total site load, as shown in Figure 2. This then forms the basis for a model for DCiE as a function of IT load:

$$DCiE(L_{IT}) = DCiE_{\infty} \times (L_{IT}/(L_{IT} + P_{1/2})). \tag{5}$$

4. Metric application

Because, as mentioned above, large DCs generally become more efficient as they become more fully loaded, the DCiE_∞ parameter reflects the efficiency of the DC if it were “infinitely

note 2) This approach was inspired by the work of Professor Roger Hockney in the 1980’s, when he applied a similar approach to describing the performance of vector supercomputers. The parameters defined by Hockney were ‘r-infinity’ (r_{∞}), which is the theoretical peak performance of a machine in terms of floating point operations per second, and ‘n-half’ ($n_{1/2}$), the length of the computational vector required to reach half that performance.⁷⁾

note 3) In this paper we make use of actual data from two Fujitsu DCs. In the interest of security and confidentiality, we refer to them simply as DC1 and DC2.

loaded.” As such, it represents an idealized implementation of the DC. Most large DCs will approach this figure as they are utilized to their maximum designed IT capacity, i.e., IT_{max} .

The $P_{1/2}$ parameter represents the point in the DC provisioning life cycle at which the DC achieves half its maximum ideal efficiency. Comparisons between DCs can be made by normalizing $P_{1/2}$ so that it is a percentage of IT_{max} . In general, DCs with a smaller $P_{1/2}$ are more modular and achieve greater efficiency sooner in their provisioning life cycle.

Note that, in some cases, DCs with a lower $DCiE_{\infty}$ will actually be more energy efficient if assessed across their full life cycle if they have a small $P_{1/2}$. However, a small $P_{1/2}$ and a large $DCiE_{\infty}$ tend to go together since well-designed modular DCs also tend to be efficient at full load.

4.1 Measurement

We recommend that users of this metric follow the guideline set out by The Green Grid for measuring PUE.²⁾ In short, IT and site energy consumption should be measured over a year to generate average values (in kW), and the site energy should be measured at the point of hand-off from the utility provider to the DC and should incorporate all DC support functions including DC offices.^{note 4)}

IT energy consumption should be measured at the output from the PDUs to the IT equipment, or even at the IT racks. The intention is to exclude all non-IT related load values from the IT load value.

In general, these guidelines are difficult to implement in practice primarily due to changes in the IT load in a DC over the course of a year.

Also, the calculations of α and β are ideally based on a least squares fit of a number of data points over a range of IT loads. In practice, these

points may be difficult to obtain, so the use of a simulator tuned to match the DC in question is recommended for generating annual IT and site energy figures for a variety of DC loads. One such simulation tool is provided to members of the Data Center Specialist’s Group of the BCS.⁸⁾ Fujitsu Laboratories of Europe has developed a similar tool for use within the Fujitsu Group companies.

4.2 Use cases

At Fujitsu Laboratories of Europe, we have applied this model to a number of DCs. **Figures 3 and 4** show simulated $DCiE$ curves for two of them.

It is clear from this data that DC1 is significantly more efficient and reaches half its ideal efficiency much earlier in the provisioning cycle than DC2, i.e., 19% maximum load for DC1 versus 29% for DC2. However, in absolute terms, DC2 is more efficient than DC1 for small IT loads. This is because DC1, having been designed for a very large maximum capacity, has a larger base load than DC2. To get the most value from a DC like DC1, it is important to implement IT services in the DC as quickly as possible so as to take advantage of the long-term greater efficiency inherent in its design.

4.3 Validation

As part of commissioning DC1, Fujitsu ran a full load test in which the DC was provisioned with a simulated IT load (using electric space heaters) from 100 kW to 3 MW in 20 steps. For each step, the DC was allowed to stabilize, and the IT load and total site load were both recorded. This data provides the basis for computing the instantaneous $DCiE$ for each level of provisioning.^{note 5)} As shown in **Figure 5**, the simulated data closely agree with the measured

note 4) Details on how to deal with shared use facilities and non-electrical energy sources, e.g., natural gas and chilled water, are available online.²⁾

note 5) Note that this data cannot be used to predict the average performance for the year as all measurements were taken on a single day, meaning that the seasonal effects were not taken into account.




Figure 3 DCiE plotted against fraction of maximum IT capacity for two Fujitsu DCs.

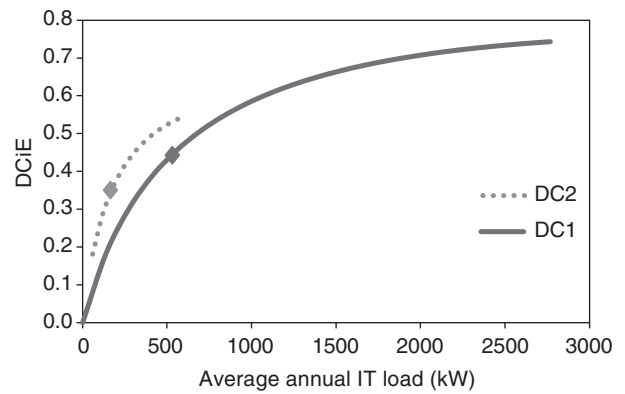


Figure 4 DCiE plotted against actual IT load (in kW) for two Fujitsu DCs.

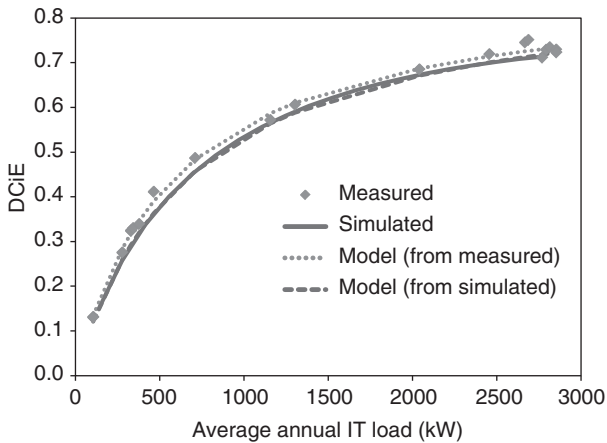


Figure 5 Measured and simulated DCiE for DC1 and results of models based on measured and simulated data against average annual IT load.

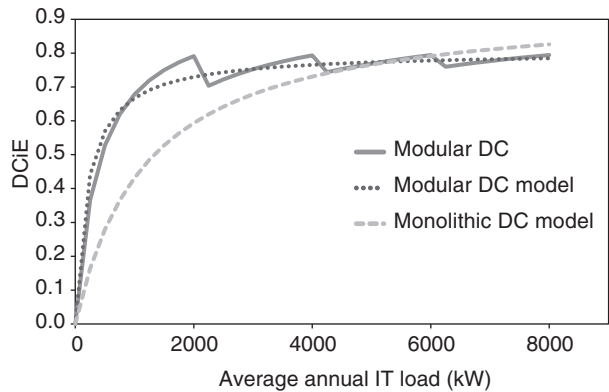


Figure 6 Comparison DCiE data for modular and monolithic DCs.

data. Moreover, the data obtained using a model based on the measured $DCiE_{\infty}$ and $P_{1/2}$ and using one based on the simulated $DCiE_{\infty}$ and $P_{1/2}$ closely agree with the measured data. The Pearson correlation coefficients were 0.999 between the measured data and the data from the model based on the measured $DCiE_{\infty}$ and $P_{1/2}$ and 0.998 between the measured data and the data from the model based on the simulated $DCiE_{\infty}$ and $P_{1/2}$.

4.4 Modular DCs

A modular DC design is generally thought to be more energy efficient. In **Figure 6**, we plot

DCiE against the average annual IT load for a modular DC design (showing the provisioning stages), a modular design based on our model, and a monolithic DC. While the values at full capacity are virtually the same, the modular DC would be more efficient early in its provisioning life cycle. This indicates that the focus should be on designing small modules with the emphasis on maximizing $DCiE_{\infty}$.

5. Conclusion

In this paper, we have presented a simple two-parameter metric that can be used to

characterize a DC regardless of where it is in its operational life cycle. While direct annual measurements of DC power usage effectiveness are important, having a model that can be used to predict the performance of the DC over its entire life cycle is a valuable tool for DC operators.

There are a number of other ways in which the proposed metric could be applied. It could be integrated across a changing IT workload profile and thereby provide an indication of the overall expected energy consumption of the DC for the given profile. This would provide very useful guidance in capacity planning and budget management. Moreover, the various components of the two main load types (IT and M&E) could possibly be decomposed into sub-components. That is, IT could be decomposed into server, storage, and networking, and M&E could be decomposed into cooling, electrical, and emergency components. This would provide a deeper understanding of a DC's load sources.



David F. Snelling
Fujitsu Laboratories of Europe Limited
Mr. Snelling is engaged in research covering the design, management, and environmental effects of large-scale, complex distributed systems.

References

- 1) Harmonizing Global Metrics for Data Center Energy Efficiency - The United States of America, European Union and Japan Reach Agreement on Guiding Principles for Data Center Energy Efficiency Metrics (February 2, 2010).
http://www.energystar.gov/ia/partners/prod_development/downloads/Harmonizing_Global_Metrics_for_Data_Center_Energy_Efficiency.pdf
- 2) The Green Grid: Recommendations For Measuring and Reporting Overall Data Center Efficiency Version 2 - Measuring PUE for Data Centers (May 2011).
<http://thegreengrid.org/en/Global/Content/Reports/ingandReportingOverallDataCenterEfficiencyVersion2>
- 3) The Green Grid.
<http://thegreengrid.org>
- 4) Green IT Promotion Council: Concept of New Metrics for Data Center Energy Efficiency, February 2010.
http://www.greenit-pc.jp/topics/release/pdf/dppe_e_20100315.pdf
- 5) Green IT Promotion Council.
<http://www.greenit-pc.jp/e>
- 6) The Uptime Institute.
<http://www.uptimeinstitute.com/publications#Tier-Classification>
- 7) R. W. Hockney: Characterization of parallel computers and algorithms. *Computer Physics Communications*, Vol. 26, No. 3-4, pp. 285-291 (1982).
- 8) BCS Data Centre Specialist Group.
<http://dcsg.bcs.org>



C. Sven van den Berghe
Fujitsu Laboratories of Europe Limited
Mr. van den Berghe is engaged in the development of an energy-flow simulator for DCs.