Electrical Efficiency Measurement for Data Centers

White Paper 154

Revision 2

by Neil Rasmussen

> Executive summary

Data center electrical efficiency is rarely planned or managed. The unfortunate result is that most data centers waste substantial amounts of electricity. Today it is both possible and prudent to plan, measure, and improve data center efficiency. In addition to reducing electrical consumption, efficiency improvements can gain users higher IT power densities and the ability to install more IT equipment in a given installation. This paper explains how data center efficiency can be measured, evaluated, and modeled, including a comparison of the benefits of periodic assessment vs. continuous monitoring.

Contents
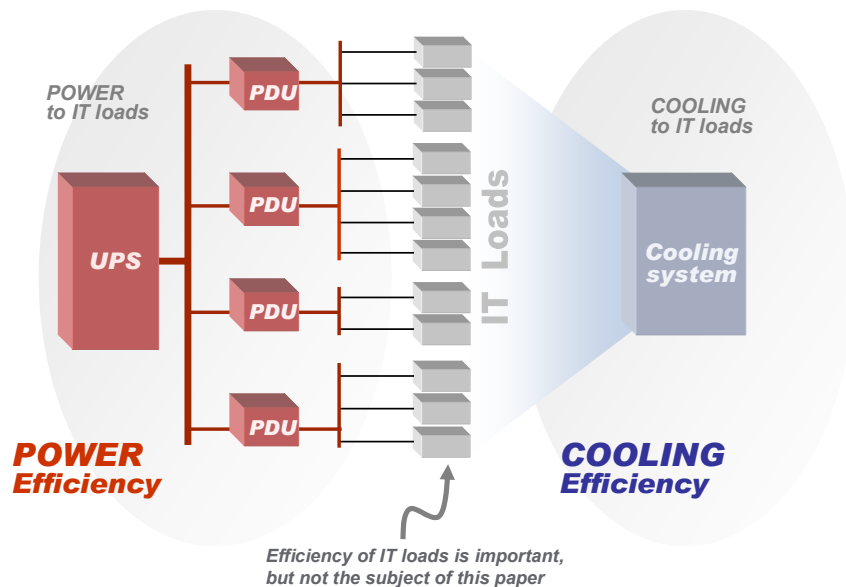
Click on a section to jump to it

Schneider Electric

Introduction

In today's environment, it is good public policy and good business to consider the options to control data center energy consumption. A 1 MW high-availability data center can consume $20,000,000 of electricity over its lifetime. Recent articles suggest that for some customers the cost of electricity is greater than the cost of IT hardware. Many companies are beginning to consider the carbon consumption of their ongoing operations and realizing that data centers are significant contributors to the environmental burden of business and industry. Research by Schneider Electric and others is showing that efficiency varies widely across similar data centers, and – even more significant – the actual efficiencies of real installations are well below the practical achievable best-in-class values.

In this paper, we will focus on the power consumption and efficiency of data center power and cooling infrastructure, which consumes 50% or more of the electrical power in a typical installation. The power consumption and efficiency of the actual IT equipment supported by the power and cooling infrastructure, while critically important to overall energy use, is not the subject of this paper (**Figure 1**).

Figure 1

*Efficiency of **power and cooling** is the subject of this paper*

Efficiency of IT loads is important, but not the subject of this paper

You can't control what you can't measure. The first problem users face is not having the data regarding the performance of their facilities. In addition, even if they do have data, they don't know how to evaluate it. They don't know how to compare or benchmark their performance to the efficiency of other similar facilities. They don't even know how to compare the actual data to the as-designed performance they should expect from their facilities.

The majority of data center operators don't have efficiency data – and even if they did, they wouldn't know how to act on it. Solving these problems requires:

- A standard language for **describing** data center efficiency
- A standard method for **measuring** data center efficiency
- A standard method for **specifying** data center efficiency
- A standard method for **analyzing** data center efficiency and determining the contributors to inefficiency
- **Benchmark data** relating to efficiency
- Efficiency modeling **tools** to assess proposed improvements or alternative designs
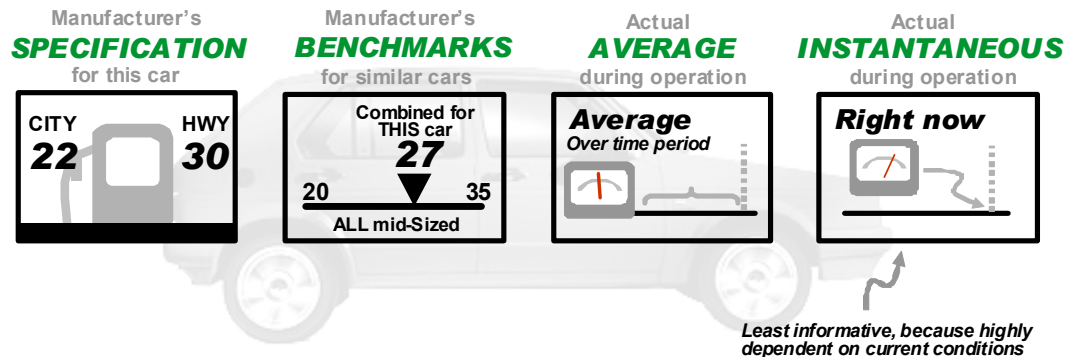
Fortunately, there are no fundamental practical or technical barriers to achieving these objectives. In fact, this paper asserts that data center operators should consider the management of efficiency as a top priority for both existing and new data centers. This paper addresses all of the above issues except the third one, **specifying** data center efficiency.

A closer look at data center efficiency

Comparing data center electrical efficiency with automobile fuel economy (miles per gallon or kilometers per liter) is a very helpful way to think about the subject. When we discuss automobile fuel economy with regard to a specific car, we talk about four basic concepts (**Figure 2**):

- The manufacturer's fuel economy **specification** for the car (e.g., "30 MPG highway / 22 MPG city")
- The fuel economy **benchmarks** of similar vehicles
- The **average** fuel economy we would experience with the car
- The **instantaneous** fuel economy we are experiencing right now

These same four key concepts of efficiency apply to the data center, and it is very important to articulate and separate them in our discussions.

Figure 2

Four measurements of efficiency, using automobile fuel economy as an example

Of all of the measures of automotive fuel economy, the instantaneous fuel economy we are experiencing right now is the least useful number, because it is highly dependent on current driving conditions (uphill or downhill, accelerating, standing at a traffic signal, coasting, pulling a trailer, etc). Similarly, taking a single efficiency measurement of a data center is an inadequate description of data center efficiency. **It is very important to understand that simply taking a "measurement" of the efficiency of a specific data center does not allow us to draw any conclusions or take any action, without additional information.**

The automobile analogy teaches us that efficiency measurement is only meaningful when placed in a framework of expected specification, industry benchmarks, defined operating conditions, and time averaging.
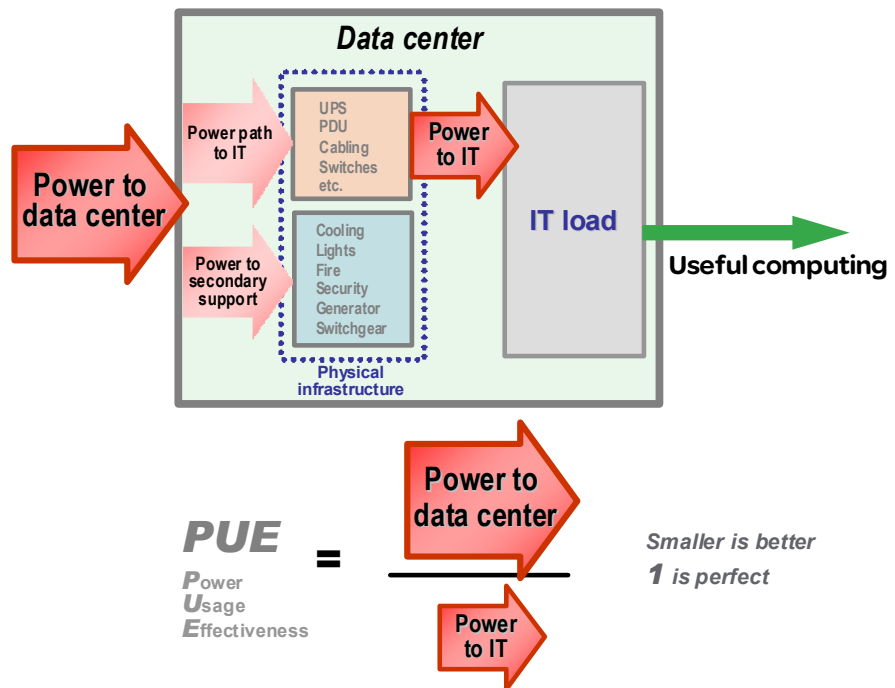
The PUE metric

In this paper, we will use Power Usage Effectiveness (PUE) as the metric for measuring data center physical infrastructure efficiency. PUE is expressed as the ratio of total data center input power to IT load power. Lower PUE means less energy "overhead" is consumed in powering the IT load, which means higher electrical efficiency for the data center. A PUE of 1 would indicate no energy overhead and perfect efficiency, which would mean the only power needed to *support* the IT load is the power actually *consumed* by the IT load.

If the data center were perfectly efficient, all the power supplied to the data center would reach the IT loads – this is the ideal case. In the real world there are a number of ways that electrical energy is consumed by devices other than the IT loads – for example, transformers, UPS, distribution wiring, fans, air conditioners, pumps, humidifiers, and lighting. Some of these devices, such as UPS and transformers, are in the power path that supplies power to the IT loads. Other devices, such as cooling and lighting, provide secondary support and protection, but do not supply power to the IT loads. (**Figure 3**).

Figure 3

PUE quantifies the power "overhead" that is consumed in supporting the IT load

$$PUE = \frac{\text{Power to data center}}{\text{Power to IT}}$$

Power Usage Effectiveness
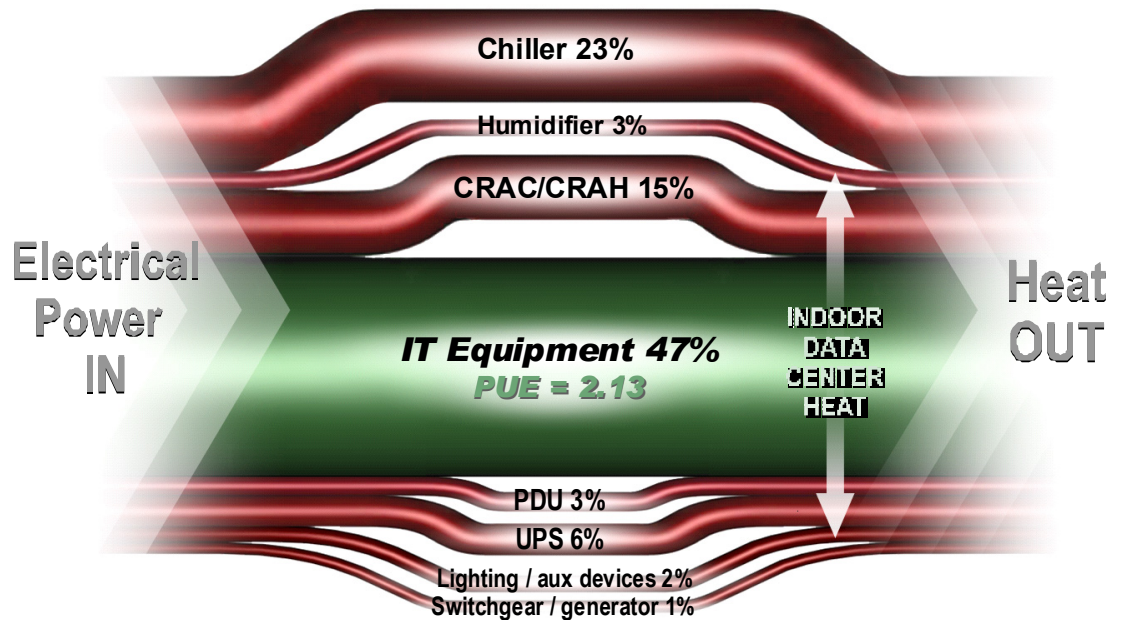
Smaller is better
1 is perfect

Note in **Figure 3** that the only power that "counts" in data centers – as far as PUE is concerned – is the power that actually makes it to the IT loads. All other power, considered waste, or "loss," in this analysis, includes heat from device inefficiency in the power path plus ALL power to the secondary infrastructure – all of which is fair game for efficiency improvements, either from better device design or completely new technologies such as "free cooling." For more about this distinction between "useful" power and "other" power (waste) power see White Paper 113, *Electrical Efficiency Modeling for Data Centers*.

Where does the "other" power go?

If the "useful" power of the data center (the PUE denominator) is the power delivered to the IT loads, where does the rest of the power go? **Figure 4** shows where power flows in a typical data center. Note that virtually all the electrical power feeding the data center ultimately ends up as heat.

The data center represented in this figure is a typical high-availability, dual power path data center with N+1 CRAC units, operating at a typical load of 30% design capacity.

Figure 4

Power flow in a typical data center

Chiller 23%
Humidifier 3%
CRAC/CRAH 15%
Electrical Power IN
IT Equipment 47%
PUE = 2.13
INDOOR DATA CENTER HEAT
Heat OUT
PDU 3%
UPS 6%
Lighting / aux devices 2%
Switchgear / generator 1%

> What about IT load efficiency?

Providing power and cooling for the IT loads is the "work" of the data center's physical infrastructure, with success measured by how little energy is consumed in the process of delivering that power and cooling – this is what PUE measures.

Overall efficiency of the entire data center will also depend on the efficiency of the IT loads in providing THEIR useful work with minimal power consumption and minimal heat generation (loss). Overall data center efficiency will be a combination of physical infrastructure efficiency (the subject of this paper) and IT efficiency (useful computing per IT watt consumed).

Improvements in the power management of computing devices, as well as improvements in software efficiency such as virtualization, are key elements in increasing the useful work generated per IT watt consumed. These important improvements complement and multiply any improvements in PUE, but are not the subject of this paper.

Any strategy to reduce the environmental impact of data centers must focus on improving both PUE **and** IT efficiency.

Note that less than half the electrical power feeding this data center is actually delivered to the IT loads. With 47% of the input power delivered to the IT load, the data center of **Figure 4** is operating at a PUE of 2.13.

Efficiency data for devices

Manufacturers provide efficiency data for power and cooling equipment. For power equipment, efficiency is typically expressed in percent. Cooling equipment efficiency is expressed in a number of ways, typically coefficient of performance (COP) for heat pumps – the ratio of heat (kW) removed to electrical power (kW) consumed, KW/ton for chillers, and EER (energy efficiency ratio) for rooftop systems. These cooling metrics are very similar and can be converted from one to another. However, these metrics do not provide complete efficiency data because their published values are calculated using a single point at standard temperature and humidity. The actual efficiency of the equipment will vary from the published value due to conditions experienced during a year. The effect of outdoor conditions on efficiency is covered in more detail later in this paper.

Lighting power is simply expressed as a simple constant loss – a 60-watt light bulb is 0% efficient when considering data center efficiency, because it generates 60 watts of heat and provides no power to the IT load.

Averaging efficiency measurements

Many cooling-related devices in data centers cycle on and off over time – humidifiers, some types of CRAC units, condensate pumps, some types of chillers, dry-cooler fans, and fresh air ventilation fans. This causes the instantaneous measurement of data center **power consumption** to vary, even if the IT load is constant. Therefore, these cycling behaviors cause the instantaneous **efficiency** to vary over time.

The energy consumption of a data center over a period of time is computed using the **average** of the data center efficiency over that period. Therefore, when we speak about data center physical infrastructure efficiency, or PUE, we really are interested in the average efficiency over a period of time. The cycling periods of most data center devices that have cycling behavior is on the order of tens of minutes, so the shortest averaging time that gives meaningful information is approximately one hour. **In this paper, when we speak of the instantaneous PUE we actually mean the value based on 1-hour averaging**.

There are good reasons to also consider other averaging periods for PUE. For example, due to seasonal variation of efficiency the *annual average value* is useful for energy cost estimation. For different applications, daily, weekly, and monthly averaging are also useful. We will refer to these averages as the daily, weekly, monthly, and annual PUE. Most popular literature on data center efficiency, whether using PUE or some other measure, does not identify or quantify the averaging period. This creates a significant ambiguity. An inspection of the common use of efficiency measurements suggests that annual PUE is the intended metric, because it is most easily related to average energy consumption over the data center lifetime.

Unfortunately PUE, measured instantaneously, will generally NOT be equal to the annual PUE, or for that matter, to the daily, weekly, or monthly PUE. At any instant in time, a data center has a single number for its instantaneous efficiency measurement. This number is the result of the inherent power consumption characteristics of the power, cooling, and lighting systems and the value of the IT load at that moment. However, the efficiency will vary over time because the IT load varies over time, the outdoor conditions vary over time, and the cooling mode of the data center can change over time. All of these factors affect the data center efficiency and the result is that data center efficiency is constantly changing. Any particular one-time measurement of data center efficiency is just a "snapshot" and cannot be used to predict future performance. **Single measurements of data center efficiency are inherently inaccurate and cannot be used as a basis for benchmarking or efficiency management.**

In order to determine how to measure data center efficiency, we need to understand how conditions affect efficiency.

> " Most popular literature on data center efficiency does not identify or quantify the averaging period. "

Factors influencing efficiency

Several variables in the data center environment have a significant impact on efficiency. These include the effect of the IT load, the outdoor conditions, and user configurations and settings for equipment.

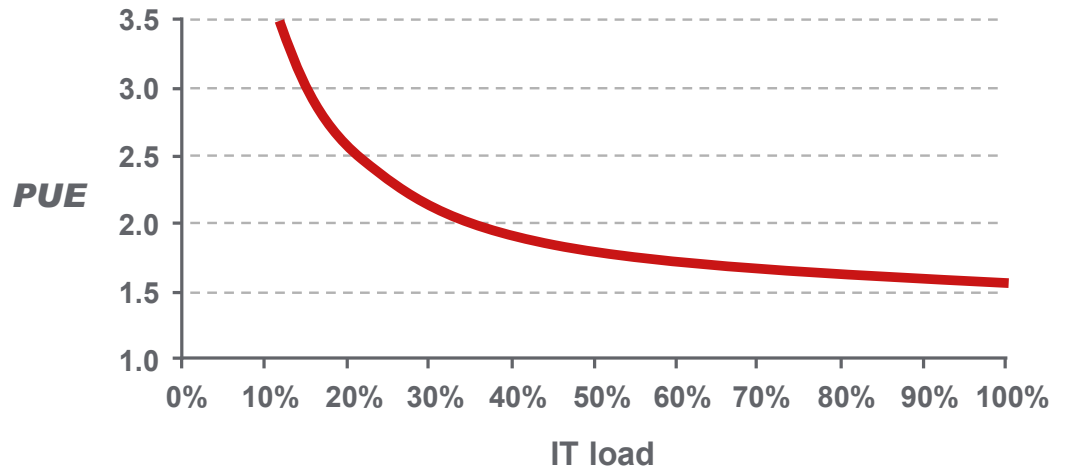### Effect of IT load on efficiency

The fact that conditions of a data center change over time cause the efficiency of the data center to change over time. One of the most important conditions that changes over time is the IT load. Power management features in newer generation IT equipment can cause the IT load to vary moment-to-moment, while the removal and addition of IT equipment by the data center operator causes longer term changes in the IT load. **Figure 5** shows how the efficiency of a typical data center varies with the IT load.

Figure 5

PUE is always a function of IT load – efficiency degrades at lower loads

A fixed installation of power and cooling equipment in a data center has a PUE curve of the form shown in **Figure 5**, with lower efficiency at lighter loads. From this curve, we can see that efficiency varies dramatically with the load.

If the IT load were to vary over the course of a day due to varying IT compute workloads, the graph shows that we can expect a variation of the instantaneous PUE over the course of the day. This means the instantaneous PUE would not equal the daily PUE. If weekend IT workloads are different from the workday IT workloads then we would find that a daily PUE measurement would typically not match the weekly PUE. This has important implications in how we measure efficiency and how we interpret the measurements, as will be discussed later in this paper.
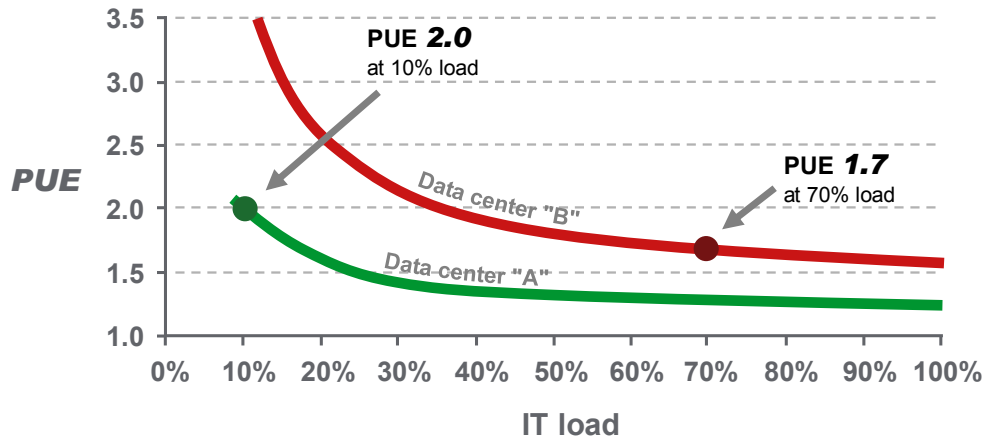
The variation of PUE with load has another important effect on how efficiency data is interpreted. Consider the case of two data centers to be compared. Suppose data center **A** has a PUE measurement of 2.0, while data center **B** has a PUE measurement of 1.7. Data center B appears to have the better efficiency measurement, so it seems reasonable to assume it is a "greener" data center of fundamentally superior design. However, consider a more detailed view of these two data centers as shown in **Figure 6**.

Figure 6 shows that data center **A**, even though it has a worse PUE at the moment it was measured, has a superior overall efficiency curve than data center **B**. From a specification standpoint, data center **A** probably uses much higher efficiency power and cooling equipment and an optimized airflow design. Nevertheless, because its percentage IT load is lower than data center **B**, it operates at a more inefficient point on its efficiency curve.

Which design is better in **Figure 6**? An engineer would say the inherent design efficiency of data center "A" is superior. However, a business person would say that the sizing of the data center was poorly planned, which negated all the technical efficiency gains. The decision as to which "design" is better depends on whether you consider appropriate sizing to be part of the design.

Figure 6

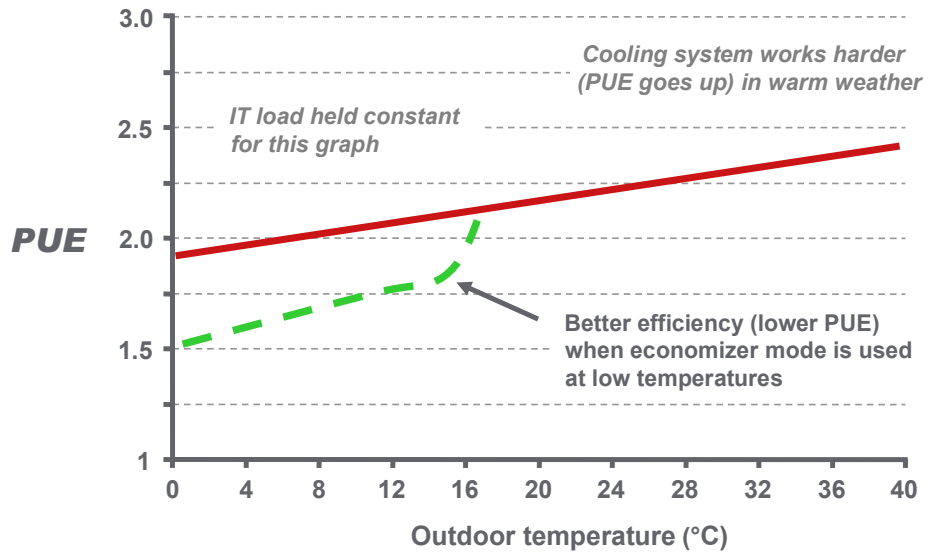*Comparison of PUE curves for two different data centers*

Effect of outdoor conditions on efficiency

Outdoor conditions are another factor that varies with time and affects data center efficiency. While various factors such as sunlight, humidity, and wind speed can affect efficiency, the most important variable is the outdoor temperature. **Figure 7** shows how the efficiency of a typical data center varies with outdoor air temperature.[1]

Figure 7

PUE as a function of outdoor temperature

The efficiency of a typical data center declines as temperature increases because heat rejection systems consume more power when processing the data center heat, and because outdoor heat infiltration into the data center becomes an additional heat load that must be processed. The dotted curve of **Figure 7** shows how the efficiency might improve at lower temperatures if the cooling system has an "economizer" operating mode. Clearly the efficiency varies with outdoor temperature, especially if the data center is equipped with economizer cooling modes.

---

[1] Note that an appropriate choice of the x-axis may be the so-called "wet-bulb" temperature or "dew point temperature," instead of the ambient temperature, depending on the type of heat rejection system used.

Effect of user configuration and settings on efficiency

There are a variety of actions users can take to affect the PUE. Users can change temperature or humidity set points, move or add vented floor tiles, fail to clean air filters, etc. These effects are highly variable and depend on the exact design of the power and cooling systems. For the purposes of this paper's analysis, these user-dependent "settings" will be treated as attributes of the data center design, and not as variable conditions like IT load and weather. When the user changes these "settings" – for example, moves a vented floor tile, changes a filter, or changes a temperature set point – the data center design is considered to have been changed and new measurements are required.
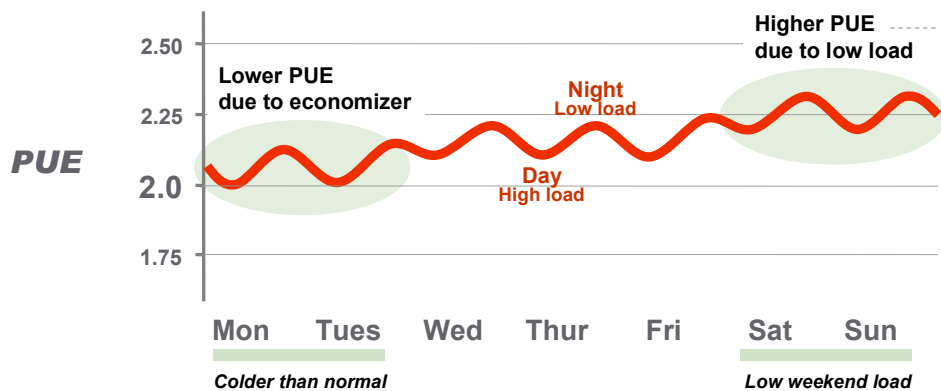
Combined effect of varying conditions

In an actual data center, we have explained that the efficiency changes as IT load and temperature change. **Figure 8** shows how the efficiency of a data center might vary during the course of a week as a result of three different effects:

- Daily variation caused by daily swings in the IT load and outdoor temperature
- Day-to-day variation caused by the weather
- IT load reduction on weekends

Figure 8

Variation of instantaneous PUE over the course of a week

While the day-to-day variation is not dramatic, it seriously degrades the usefulness of a specific, one-time measurement of PUE. Even if a specific efficiency measurement is made with extreme precision, its accuracy in predicting electrical bills, measuring the effect of efficiency improvement activities, and performing trending analysis is poor, because of the expected day-to-day variation in efficiency. The use of a single instantaneous audit-type measurement of PUE is not effective for benchmarking, is not effective for trending, and does not provide actionable insights into the opportunities for efficiency improvement. Fortunately, there is a way to use measurements to achieve all of these objectives, but it requires a new way of thinking about the purpose of measurement.

Efficiency modeling of data centers

A mathematical model that accurately represents the workings of a specific data center, and accepts as inputs the IT load, outdoor weather statistics, etc., could be used effectively in a data center energy management program. Unlike the measurement of an actual operating data center, which provides only data for the conditions at the time of measurement, a model can provide data for any input conditions fed to it. For example, a model could provide the expected value of the efficiency of a data center at full load, even when the actual IT load is a small fraction of the rated load. Models of two different data centers can be fed the same

input conditions, allowing meaningful comparison. Furthermore, a model can even be created before a data center is even built, allowing the prediction of performance in advance. In an operating data center, we cannot simply change the weather or the IT load in order to normalize, or standardize, the conditions. When we take a single "snapshot" measurement, we have no choice: we get the current IT load and the current weather. An effective model allows us to normalize the measurements under different conditions and establish the fundamental operating performance capabilities of the data center.

Here is a list of some of the benefits that a mathematical model provides that **cannot be obtained** by simply measuring and recording data center efficiencies:

- Accurately predict the efficiency performance of a proposed data center design in advance
- Estimate with high accuracy the efficiency performance of an operating data center for conditions that are impractical to measure, such as for different IT loads, or for industry benchmark conditions
- Estimate with high accuracy the efficiency performance of a data center for situations with only partial information, where it is not practical to measure power consumption of all of the circuits
- Estimate with high accuracy the efficiency of the data center for various outdoor conditions, and allow estimation of efficiency averages over time
- Identify and quantify the contributions specific devices in the data center power, cooling, and lighting systems are making to the inefficiency of the data center
- Identify data center subsystems that are operating outside of their expected efficiency parameters
- Benchmark and compare the models of different data centers

While a single measurement of data center efficiency at an operating point is informative, it is not actionable without the capabilities listed above. For this reason, the mathematical model is the key to creating a process and system for efficiency management. It is the model that allows understanding of the causes of inefficiency; **therefore, the purpose of data center efficiency measurement is to establish the parameters of the efficiency model**.

In summary, a key purpose of measurement of data center efficiency is to obtain information that contributes to the creation of an accurate model for that data center. **It is the model that provides actionable information about data center efficiency, not the measurement**.

Practical models of data center efficiency

Although the benefits of data center efficiency modeling are valuable, there remains a question as to whether developing and operating a data center efficiency model for a specific data center is feasible and practical. The answer to this question depends on the degree of precision required from the model. Our research finds that it is very practical to create and operate models of sufficient accuracy to obtain the objectives of efficiency management outlined in this paper. In fact, we believe developing and operating efficiency models will be a standard practice in data centers in the near future, and that such models will be built into the standard software tools used to manage data centers.

An example demand / loss diagram for a practical data center efficiency model is shown in **Figure 9**.

Figure 9

A data center efficiency model showing demands on devices and resulting losses

This diagram shows how the data center infrastructure power consumption (losses) are determined by the **IT load** and **outdoor temperature/humidity** conditions, giving rise to energy demands within the power, cooling, and lighting systems. Each type of device (UPS, CRAH, etc.) has demands placed on it (i.e., electrical demands, thermal demands) and generates power consumption (loss) according to the device configuration and inherent efficiency characteristics. *System* configuration elements that impact the efficiency model include such things as types of devices powered by the UPS, type of cooling plant, and the use of economizer modes. The model includes statistical weather data for the site. A detailed description of the theory, construction, and use of a data center efficiency model are described in White Paper 113, *Electrical Efficiency Modeling for Data Centers*.

An efficiency model for a data center can be created for an existing data center, or it can be created before a data center is even constructed, if the design and the characteristics of the

power, cooling, and lighting devices are known. If the model accurately represents the design, the data it provides will be similarly accurate. While the electrical performance of some types of devices, such as lighting, UPS, and transformers are very consistent and predictable, there are many uncertainties regarding the as-built performance of devices such as pumps and air conditioners that cause the model to lose accuracy. This is where measurement can help.
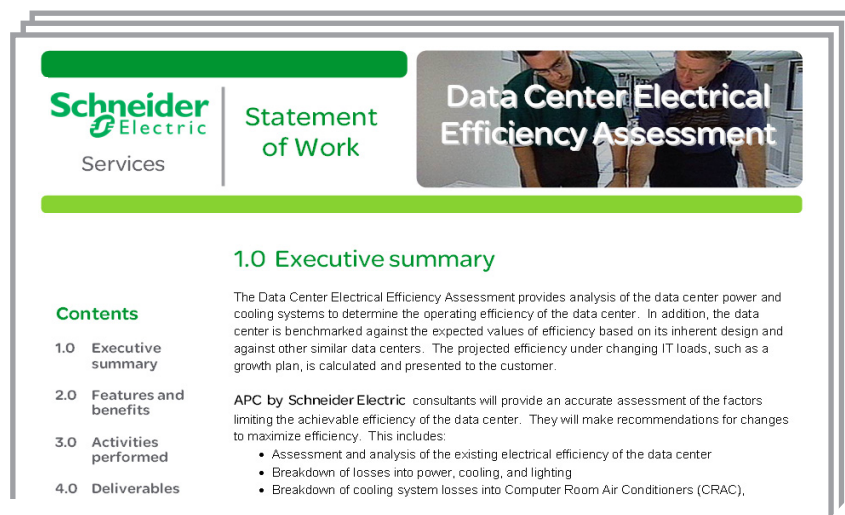
Preparing an efficiency measurement plan

The preceding discussions explain the limited usefulness of relying solely on periodic measurements of data center efficiency, due to poor repeatability and lack of guidance. Periodic measurements, however, can be part of an overall management strategy that includes both initial and ongoing measurements. These two types of measurements have differing objectives:

1. **Initial measurements** to calibrate the data center efficiency model, establish "as-is" and "should-be" performance, and identify potential efficiency improvement opportunities. Initial measurements typically require measurements on individual power and cooling subsystems, in addition to an overall efficiency measurement.

2. **Ongoing measurements** to compare against the model to provide alerts of unexpected inefficiencies and to quantify improvements. Ongoing measurements can be made either by periodic sampling or by continuous instrumentation.

Initial measurements

The first time the efficiency of a data center is measured it should be part of an overall efficiency assessment by experts. In addition to making an efficiency measurement, a data center efficiency assessment typically provides an analysis of the as-built configuration and recommendations regarding efficiency improvement. Ideally, an assessment should provide a mathematical model of the data center as one of its deliverables. **Figure 10** shows an example of a statement of work for a data center efficiency assessment.

Figure 10

Example of a statement of work for an electrical efficiency assessment (cover page)

The data collected during an efficiency assessment should include measurements of enough subsystems to allow a mathematical efficiency model to be calibrated.

Once a data center model is calibrated using initial efficiency measurements, it can be immediately used to establish benchmark performance, such as expected average efficiency

over a year, or to establish the efficiency at other reference IT loads and outdoor conditions, such as those established for industry comparative benchmarking.

Ongoing measurements

After data center efficiency has been measured and the efficiency model has been calibrated, ongoing measurements should be made to quantify any efficiency improvements and to provide warnings regarding any undesirable loss of efficiency. When the measured value of efficiency is found to change from the original measured value, we have already shown that this can be due to variations in the IT load or in the weather. The model can correct for the effects of IT load and weather and reveal whether any efficiency variations found are due to these effects or are caused by actual changes in the underlying data center infrastructure.

Ongoing measurements of efficiency can be performed periodically such as twice per year (ideally once in the summer and once in the winter), or a continuous real-time measurement system can be installed.

The next sections of this paper explain why real-time measurement can be expected to be a standard feature in new data centers within a few years.

Identification of electrical circuits to be measured

To measure the data center efficiency at a particular operating point, we must measure the total input power to the data center and the total IT load. If the facility is a dedicated data center, the input power can be measured at the utility mains connection to the building. If the load were a single giant IT load device, then the IT load power would be a single measurement of the power at the device electrical connection. Only two measurements would be needed in this hypothetical case. **Unfortunately, this ideal situation is never achieved**. Most data centers are part of multi-use buildings with other loads besides the data center, and all data centers are made up of collections of IT devices – possibly thousands – many with separate electrical circuits.

To ensure an exact measurement, all devices that draw power from the mains to make up the total data center load need to be instrumented separately from other non-data-center loads and their input power summed. Furthermore, to obtain the total IT load, all IT devices need to be separately instrumented and their input power summed. In a 1 MW data center this could require thousands of simultaneous power measurements, which is a technical and practical challenge. For this reason, some data center operators have concluded that efficiency measurement is impractical. Fortunately, **complicated instrumentation and time consuming measurement are not necessary** because a small number of consolidated measurement points combined with mathematical models of the system can be demonstrated to provide efficiency measurements of sufficient accuracy.[2]

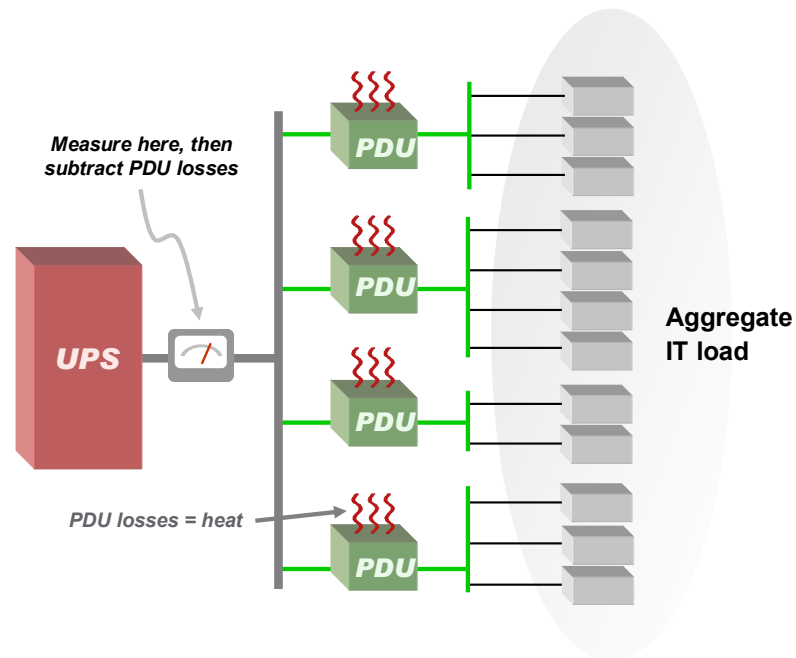### Measuring the aggregate IT load

As an example of how actual measurement can be dramatically simplified using the model, consider the case of measurement of the IT load. **Figure 11** shows the typical situation of a large number of IT loads being provided power from power distribution units (PDUs) supplied from a UPS.

---

[2] A complete description of this approach is not within the scope of this paper.

Figure 11

Measuring the aggregate IT load

To obtain the actual aggregate IT load, it may appear necessary to measure and sum all of the branch circuits supplying every IT load, which could be thousands of measurements. If instead we measure the power at the output of the UPS, we have dramatically simplified the problem to a single measurement, but we have introduced an error because at this location we are measuring the sum of the IT loads plus the power consumed by the PDUs (the PDU losses). In addition, in some data centers other devices such as air handlers may be powered from the UPS. The error caused by using the UPS output power as representative of the actual IT load can be in the range of 2% to 25% depending on the data center, which is quite significant. However, PDUs can be very accurately characterized in the model, and the model can compute the PDU losses to a high degree of accuracy given the UPS output power. The model can subtract the PDU losses from the UPS output power to obtain the IT load with no significant error. In this way, the model allows us to cut potentially thousands of measurements down to a single measurement.

The technique of using the model to gain a large amount of information from a small number of measurements can be used for one-time measurements, and it also can be used as part of a continuous permanent efficiency monitoring system, dramatically reducing the cost of continuous efficiency monitoring.

Ongoing vs. initial investment

The number of power measurement points for ongoing efficiency measurements is typically less than are used for initial measurements. Ongoing measurements need to focus on circuits that are most likely to change over time and are the largest contributors to power consumption. For example, there are some devices such as transformers where it would be valuable to measure them initially, but whose ongoing power contribution can be computed with high accuracy by the model. In general, ongoing measurements require about half the number of initial measurement points. This suggests an instrumentation strategy where the ongoing measurement points are monitored by permanently installed power metering, while the additional points needed for initial measurement are measured using portable instrumentation.

The determination of which circuits should be measured initially and on an ongoing basis is a key part of a professional data center efficiency assessment.

Instrumentation for efficiency measurement

Data center efficiency is a calculated number based on power measurements. **Figure 12** shows examples of the types of equipment used for measurement.
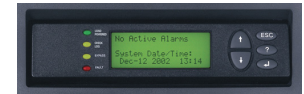
Portable power measurement equipment (Fluke 435)

Permanently installed power measurement equipment (Power Logic PM700)

Permanently installed built-in power measurement functions in cooling and power equipment (APC Symmetra UPS)

Data centers can be permanently instrumented for efficiency, or efficiency can be audited periodically using portable instrumentation. In either case, the power circuits to be measured must be identified as described in the previous section. This paper has explained that it is not necessary to measure the power flows in all of the thousands of circuits in a data center. Very accurate computations of efficiency can be made by measuring the power flow in a small subset of the power circuits. The efficiency measurement strategy consists of the following elements:

- Deciding on permanent vs. periodic/portable measurements
- Identifying the appropriate measurement points
- Establishing a system for reporting efficiency data

Permanent vs. portable instruments

Efficiency measurement instrumentation can be permanently installed or measurements can periodically be made by qualified service personnel with portable instruments. The preferred solution depends on a number of factors, with benefits to each approach.

Permanent instrumentation provides continuous real-time efficiency data and provides significant advantages. One important advantage of permanent instrumentation, compared to periodic audits, is that there are no people attempting to make measurements on live power circuits with portable instruments – any time people are attempting to make measurements on live circuits there is a risk of error and downtime. The disadvantage of permanent instrumentation is the greater first cost, especially for retrofit installations. In most cases, permanent instrumentation can be installed without disruption of the critical IT load.

Periodic audits using portable instrumentation have a lower first cost, and are particularly well suited for existing data centers near end-of-life.

For new data centers, permanent instrumentation should be specified. Some power and cooling devices, such as UPS, may already have built-in power measurement capability. If the accuracy of this built-in capability is sufficient, built-in instrumentation saves the need to

use any additional measurement instrumentation on the affected circuit. Efficiency management software should be able to capture power measurements from power and cooling equipment with built-in power metering instrumentation.

Identifying the appropriate measurement points

Whether the instrumentation is permanent or portable, the appropriate power measurement points must be selected. This paper has made it clear that it is **not necessary** to measure all of the thousands of circuits in the data center. The objective is to measure the fewest circuits that will provide the data necessary to construct a model with requisite accuracy. For existing data center, this problem may be compounded by practical matters related to how the building was wired. Furthermore, some circuits or devices, such as cooling towers, may be intentionally shared with loads outside of the data center to be measured, making it impossible to directly measure the consumption related to the data center. Therefore the instrumentation measurement point selection may need to be customized for a specific installation.

Data collection during efficiency measurement

Many data center operators will attempt to make efficiency measurements, or retain professionals to make such measurements. This paper has shown that the typical one-time measurements that are made are not effective at achieving the goals of energy management. Ideally, data center efficiency assessments will be performed as part of an energy management program and follow the principles outlined in this paper. However, one-time measurements still can be useful as long as they are provided with other supplemental data. In order for a later analysis to use one-time measurement data, the following information must be included with the efficiency measurement:

- Date of the measurement
- Time of day of the measurement
- Averaging interval used
- Measurement points and measurement equipment used at each measurement point
- Full load rating of the data center (maximum IT load)
- Actual operating load at the time of the measurement (current IT load)
- Outdoor temperature and humidity
- Data center redundancy levels
- Type of heat rejection used (dry cooler, cooling tower, free air, packaged chiller)
- Economizer type installed (if any)
- Status of the economizer at the time of measurement (engaged, disengaged)

If this data is available it allows later efficiency analysis and modeling that is otherwise not possible. Therefore, data center operators who choose not to use trained specialists to make efficiency measurements should always ensure that the persons performing the measurement collect and record the above information at the time of any efficiency measurement.

Establishing a reporting system

The previous section shows that electrical efficiency is not directly measured by the instrumentation, and that additional calculations and analysis are required. Furthermore, establishing the overall efficiency curve for the data center power and cooling infrastructure requires that the measurements be used in conjunction with the efficiency model for the data center.

For periodic audits, the efficiency curve for the data center is calculated using the principles described in this paper, and an audit report including this result is generated.

For permanent instrumentation, a real-time calculation of efficiency is possible and preferred. This function should be provided by the power and cooling capacity management system, which already contains the topological model for the data center. When properly implemented, electrical efficiency trends can be reported, and alerts generated based on out-of-bounds conditions. Furthermore, an effective system would provide the ability to diagnose the sources of inefficiency and suggest corrective action.

Efficiency data should be reported at least annually, but can be reported continuously if an energy management system is implemented. The following is a list of some of the key data that should be included in a data center efficiency report:

- Actual annual PUE performance, measured or extrapolated from the model
- Expected annual PUE under industry-standard conditions, from the model
- Design value for annual PUE under industry-standard conditions, if available
- Annual average IT load (% of rated capacity)
- Breakdown of power consumption by major data center subsystem
- Comparison of actual annual PUE to other data centers of similar classification (worst case, typical, best case)
- Comparison of expected annual PUE under industry-standard conditions to other data centers of similar classification (worst case, typical, best case)
- Total annual electrical consumption (kW-hr)
- Estimated annual electrical cost
- Trending of actual annual PUE, annual electrical consumption, electrical costs

Conclusion

The measurement of data center efficiency is for the purpose of managing the use of electricity. Taking action to control electrical usage requires knowledge of

- The sources of inefficiency
- Opportunities for improvement
- The expected benefits of efficiency improvements

Making a simple measurement of data center efficiency is interesting, but does none of these things. Measurements are only useful when used in conjunction with models. For this reason, modeling is a critical aspect of efficiency management and the data needs of the model establish the requirements for the measurement of power flows within the data center.

To manage data center efficiency we need to make initial power measurements to establish the model and baseline conditions, then periodically or continuously monitor power flows on an ongoing basis to provide information about efficiency trending, inefficient conditions, and opportunities for improvement.

The analogy with automobile fuel economy is helpful in understanding the types of information we need in order to interpret data center efficiency data, and is illustrated in **Table 1**.

Automobile fuel economy		Data center infrastructure efficiency
The manufacturer's fuel economy specification of our automobile (i.e. 26 MPG highway, 22 MPG city)	→	The data center designer's efficiency specification of our data center (i.e. 55% PUE at 75% IT load and 25 deg C outdoor temperature)
The economy benchmarks of similar vehicles	→	The efficiency benchmarks of similar data centers
The average fuel economy we are actually experiencing	→	The average PUE we are actually experiencing (daily, weekly, monthly, annual)
The real-time fuel economy we are experiencing right now	→	The real-time PUE we are experiencing right now

There are thousands of power flows within a typical data center. Analysis shows that it is not necessary to measure all of these flows in order to measure and manage efficiency. When combined with appropriate modeling and information about the power and cooling devices, it is possible to create an efficiency management system of high accuracy with only a small number of measurements.

To manage efficiency, users have the option of permanently instrumenting the appropriate power flows for real-time efficiency metering, or they can perform periodic audits of power flows using portable instruments. For new data centers, permanent instrumentation is appropriate. For existing data centers, a mix of limited permanent instrumentation and periodic audits is recommended.

✎ About the author

Neil Rasmussen is a Senior VP of Innovation for Schneider Electric. He establishes the technology direction for the world's largest R&D budget devoted to power, cooling, and rack infrastructure for critical networks.

Neil holds 19 patents related to high-efficiency and high-density data center power and cooling infrastructure, and has published over 50 white papers related to power and cooling systems, many published in more than 10 languages, most recently with a focus on the improvement of energy efficiency. He is an internationally recognized keynote speaker on the subject of high-efficiency data centers. Neil is currently working to advance the science of high-efficiency, high-density, scalable data center infrastructure solutions and is a principal architect of the APC InfraStruXure system.

Prior to founding APC in 1981, Neil received his bachelors and masters degrees from MIT in electrical engineering, where he did his thesis on the analysis of a 200MW power supply for a tokamak fusion reactor. From 1979 to 1981 he worked at MIT Lincoln Laboratories on flywheel energy storage systems and solar electric power systems.

Resources

Click on icon to link to resource

Electrical Efficiency Modeling for Data Centers
White Paper 113

White Paper Library
whitepapers.apc.com

TradeOff Tools™
tools.apc.com

Contact us

For feedback and comments about the content of this white paper:

Data Center Science Center
DCSC@Schneider-Electric.com

If you are a customer and have questions specific to your data center project:

Contact your **Schneider Electric** representative