Guidance for Calculation of Efficiency (PUE) in Data Centers

White Paper 158

Revision 3

by Victor Avelar

Executive summary

Before data center infrastructure efficiency can be benchmarked using PUE or other metrics, there must be agreement on exactly what power consumptions constitute IT loads, what consumptions constitute physical infrastructure, and what loads should not be counted. Unfortunately, commonly published efficiency data is not computed using a standard methodology, and the same data center will have different efficiency ratings when different methodologies are applied. This paper explains the problem and describes a standardized method for classifying data center loads for efficiency calculations.

RATE THIS PAPER ★ ★ ★ ★ ★

APC by Schneider Electric White Papers are now part of the Schneider Electric white paper library produced by Schneider Electric's **Data Center Science Center**
DCSC@Schneider-Electric.com

Schneider Electric

Introduction

The benefits of determining data center infrastructure efficiency as part of an effective energy management plan are widely recognized. The standard metrics of Power Usage Effectiveness (PUE) and its reciprocal Data Center Infrastructure Efficiency[1] (DCIE) have emerged as recognized standards.

In the course of providing energy audits for data centers, Schneider Electric has identified a number of practical problems with establishing the PUE of specific data centers. Here are some of the issues that cause difficulty in classifying power-consuming subsystems as IT loads, physical infrastructure, or neither:

- Devices are encountered in data centers that draw power, but how (or if) their power data should be counted in the efficiency calculations is unclear
- Certain data center subsystems are not present in some data centers (for example, outdoor lighting or NOC)
- Some subsystems support a mixed-use facility and are shared with other non-data-center functions (for example, cooling towers and chiller plants), so the fraction of power attributable to the data center cannot be directly measured
- Some subsystems are extremely impractical or costly to instrument for power consumption (for example, PDUs due to number of output connections, or switchgear)
- Some practical power measurement points include loads that are unrelated to the data center but cannot be separated during measurement

One or more of the above practical problems are frequently present in typical data centers, and are almost always present in data centers that exist in a shared use facility like an office building. Since most data center operators who attempt to determine PUE will encounter one or more of the above problems, a standard way to deal with them should be defined. That is the purpose of this paper.

This paper defines a standard approach to collecting data from data centers and showing how to use it to calculate PUE, with a focus on what to do with data that is confusing or incomplete.

Users can always make up their own approach to collecting and processing energy data, but data center efficiency cannot be benchmarked without a standard method. It is important to note that various standard bodies such as the Green Grid and others recognize the problems outlined here and are working on solutions in the form of guidelines and standards. One such guideline is available in a publication from ASHRAE and The Green Grid titled, *PUE: A Comprehensive Examination of the Metric*[2]. It is anticipated that standards will emerge related to many of the issues described in this paper. When such standards exist, this paper will immediately be updated as appropriate to describe and align with such standards. Therefore, before applying the methods described here, visit whitepapers.apc.com to make sure you are using the most current revision of this paper.

[1] In this paper, PUE will be used as the efficiency metric. All discussions equally apply to DCiE

[2] http://www.thegreengrid.org/en/Global/Content/Books/Book1-PUEAComprehensiveExaminationoftheMetric

A three-part methodology

To overcome these problems, the following three-part methodology is outlined in this paper:

1. Establish a standard **to categorize data center subsystems** as either (a) IT load, (b) physical infrastructure, or (c) not included in the calculation
2. If a subsystem's power consumption cannot be directly measured because it is **shared with non-data-center loads**, estimate the power using a standardized methodology for that subsystem type
3. If a subsystem's power cannot be directly measured because of **technical barriers to measurement**, estimate the power using a standardized methodology for that subsystem type

Each of the three parts of the method is described in turn.

Part 1: Standard categorization of data center subsystems

To perform data center efficiency calculations, it is critical to clearly define which power constitutes **IT load**, and which power is **physical infrastructure**. While the designation of servers and storage devices as IT load is obvious, and the designation of UPS and cooling/air-handling units as physical infrastructure is obvious, there are many power-consuming subsystems in the data center that have a less obvious place in the efficiency calculation, including personnel spaces, switchgear, and network operations center (**Figure 1**).
These subsystems consume power. If they are not uniformly classified for all data centers, it is not possible to directly compare computed efficiency results across different data centers. There are ongoing efforts by customers, governmental bodies, and utility providers to benchmark data center efficiency, so it is essential to establish standardized guidelines for what is considered to be **IT load**, what is considered to be **physical infrastructure** (sometimes called site infrastructure), and what is **not included** in data center efficiency calculations. In addition, there are often practical problems with obtaining key data required to calculate data center efficiency.
There are both theoretical and practical considerations in assigning loads to the above three categories for the purpose of data center efficiency calculation. Rational people could draw different conclusions about whether different types of loads should be counted in the PUE metric at all, and if counted, whether they are IT loads or not. The following practical guidelines are applied to construct the categorization system in this paper:

- If managing the energy use of the subsystem is a realistic and desirable outcome of data center energy management, then the system load should be included in the PUE metric.
- If it is typically not practical to separate the measurement of a specific load during a measurement process, then it should not be prescribed that the specific load be separated for computation of the metric.
- A comprehensive list of the types of devices found in a data center should be standardized, with a clear assignment to "IT load", "physical Infrastructure", or "not Included", for the purpose of data center efficiency calculation.

The following tables are a preliminary attempt to establish a list of common subsystems and their assignment to one of the three categories, namely "IT load", "Physical Infrastructure", or "Not Included", for the purpose of data center efficiency calculation. These tables are consistent with the ASHRAE and The Green Grid publication titled, *PUE: A Comprehensive Examination of the Metric*. **Tables 1-4c** provide a list of subsystems that should be accounted for in PUE and those that should be excluded. These subsystems are grouped into four separate tables of equipment types: IT, power, cooling, and other.

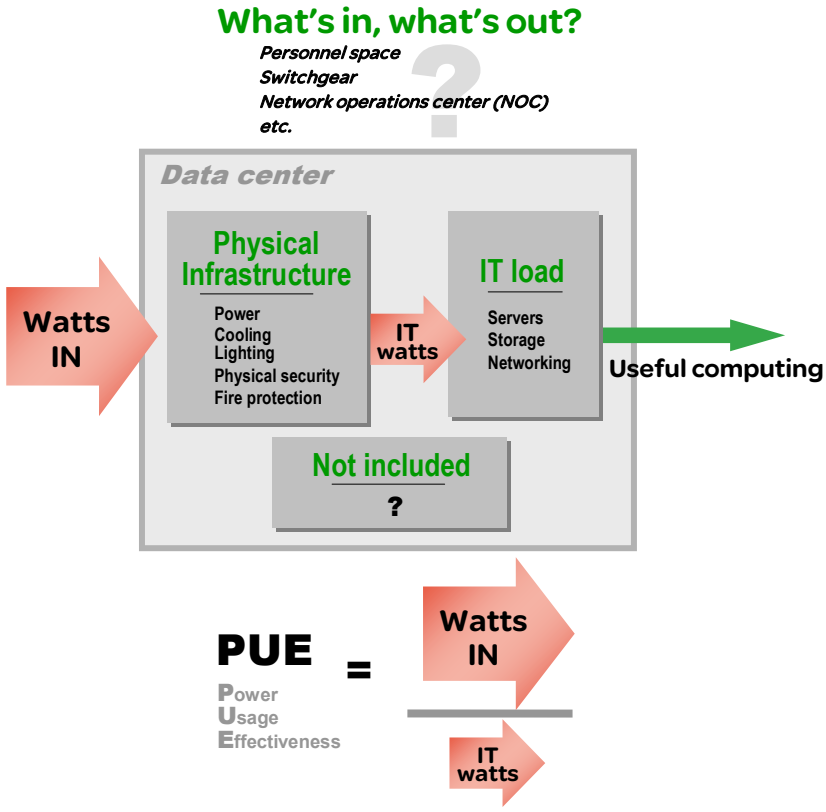

Figure 1

Uniform classification of power-consuming subsystems is essential for benchmarking data center efficiency

In **Table 1**, various IT loads are identified and categorized in the IT load category. Some types of IT equipment are obviously categorized as such, but other types are not always categorized consistently. For example, there is typically some confusion about whether disaster recovery IT loads or network operations center (NOC) loads should be categorized as IT loads within the data center.

Almost all data centers have some form of network operations center, which may be as small as a desk in the computer room, or it may be a large room with dozens of workstations and walls of monitors. In some cases, an organization may have a large central NOC that monitors a number of data centers, in which case one data center may be burdened with a large NOC while another similar data center has only a small NOC. There are arguments that favor inclusion of the NOC in PUE calculations, and arguments to exclude the NOC. However, using the principle that it is desirable to manage energy of data center related loads such as a NOC, the NOC should be included in energy reporting, and therefore in the PUE. Furthermore, NOC related energy use such as power and cooling, may be quite difficult to separate out from other data center related energy use. This again suggests that it is practical to include the NOC in energy reporting and PUE calculations.

Table 1

Proposed categorization of data center IT subsystems
*(refer to **Figure 1** for context of categories)*

	Data center subsystem	Power category: IT load	Power category: Physical infrastructure	Power category: Not included	Issues
IT	Servers	√			
	Storage equipment	√			
	Networking gear	√			Some have argued this should not be included, but it is very difficult to separate it.
	KVM and monitors	√			Some have argued this should not be included because it is not really IT equipment, but it is very difficult to separate it.
	Disaster recovery IT loads	√			The purpose of listing disaster recovery (which is a type of data center application) is to cover the possibility that some data centers have DR on the same site. Ideally DR is handled completely separate from a main data center facility and should be managed and measured as an independent data center. Consequently, if DR is on the same site and has shared resources of physical infrastructure then it must be included and is just another business process that is supported within the main DC.
	IT equipment in network operations center (NOC)	√			This load is primarily computers and displays for data center control. It could be argued to include it in any of the 3 categories, but in some cases it is not located at the same location as the data center, making assignment difficult. General assessment to include in the IT consumption when it is shared within same facility and sharing grey space resources. Could be a free standing facility as a NOC, which would drive a new category of data center and should be included as IT load with separate rating for PUE baseline.

Table 2 lists power subsystems that can be found in a typical data center and how they are assigned in the PUE calculation. While most of these assignments are obvious, an important value of this chart is as a checklist to make sure that all the appropriate data has been found and integrated into the PUE calculation. In particular, many reported measurements of PUE incorrectly assign PDUs and static transfer switch (STS) units as part of the IT load during PUE calculations, which can result in large errors, especially in partially loaded data centers. While it can be argued that it is difficult to measure the losses of these devices, it is quite possible to estimate them with a high degree of accuracy. Therefore it is never valid to omit them from efficiency calculations.

Other common mistakes made in PUE calculations are omissions of switchgear and automatic transfer switch (ATS) equipment. In many cases, it is not practical or cost effective to measure the losses of these devices, but it is always practical to estimate them. Although such estimates have some error, the fact that these subsystems are a small percentage of the overall system energy use (around 1%) causes the estimate to contribute only a very small error to the final calculation of energy use or PUE.

Alternate energy generation systems such as wind turbines and solar panels operate in parallel with data centers and their energy contribution should not be incorporated into PUE because it masks true data center performance and distorts the purpose of the metric.

Similarly, PUE does not give "credit" for using the waste heat from a data center to heat an office building. The Green Grid metric, ERE[3], should be used for this purpose.

Table 2

Proposed categorization of data center power subsystems
*(refer to **Figure 1** for context of categories)*

	Data center subsystem	Power category: IT load	Power category: Physical infrastructure	Power category: Not included	Issues
Physical infrastructure - power	**Switchgear & panel boards** (outdoor or indoor)		√		Many sites have no way to determine the switchgear power, so it would need to be estimated if included. Energy, space and resources are consumed to provide and operate this equipment and are an essential part of all data centers.
	Automatic transfer switches (ATS)		√		
	Generators (block heater, jacket water heater, strip heaters, generator controls, generator battery charger)		√		
	Uninterruptible power supplies (UPS)		√		
	Static transfer switches (STS)		√		These are clearly part of the power path, but some have argued that this should be considered part of the IT load because it is difficult to separate. In general, estimation is typically required.
	Power distribution units (PDU)		√		
	Disaster recovery power systems		√		Following the same logic above, if DR can be isolated, then it is not included in the main facility PUE, otherwise it must be included.
	Alternate energy systems (PV, wind, etc.)			√	These systems almost always operate in parallel with the data center, and are simply co-located with the data center. Their inclusion in PUE calculations would mask the true performance of the data center.

Table 3 lists various cooling subsystems that should be included in the PUE calculations. Note that some of these systems are often inadvertently omitted when calculating data center efficiency, such as make-up air systems, central humidifiers, and unit heaters, resulting in an overestimate of data center efficiency.

[3] http://www.thegreengrid.org/en/Global/Content/white-papers/ERE

Table 3

Proposed categorization of data center cooling subsystems
*(refer to **Figure 1** for context of categories)*

	Data center subsystem	Power category: IT load	Power category: Physical infrastructure	Power category: Not included	Issues
Physical infrastructure - cooling	**Chillers**		√		
	Chilled water pumps (primary, secondary, tertiary)		√		
	Condenser water pumps		√		
	Cooling tower (fans, basin heaters)		√		Some have argued that this is difficult to separate out in mixed-use facilities. In these cases, a method of estimation of the assigned power to the data center is required.
	Water treatment (sand filter pumps & injectors)		√		Water treatment includes water filtration and chemical systems.
	Pipe freeze protection (heat tape, heat trace, strip heaters)		√		
	Air compressors		√		Air compressors oftentimes support pneumatic valves in chilled water plants.
	Centralized humidifiers		√		
	CRAH / CRAC (fans, reheat coils, humidification)		√		CRAH - computer room air handlers CRAC - computer room air conditioners In addition to the IT space of the data center, these units are sometimes used to condition the mechanical and electrical rooms which should also be accounted for.
	Condensate pumps		√		
	Make-up air / fresh air system power		√		Some have claimed this is difficult to measure and have argued not to include it.
	Unit heaters		√		Small heaters (usually ceiling-mounted) are sometimes used in mechanical and electrical rooms during the winter months.
	Condensers		√		(Condensers are used only for air-cooled CRAC units.)
	Dry coolers		√		
	Cooling for NOC		√		
	Disaster recovery cooling systems		√		
	Well pumps		√		Much water can be consumed by cooling systems. Some data centers pump their own water but some purchase water as a utility that is pumped by remote utility pumps. Water as a resource requires energy to cool, store, and move it.

Table 4

Proposed categorization of data center other subsystems
*(refer to **Figure 1** for context of categories)*

	Data center subsystem	Power category: IT load	Power category: Physical infrastructure	Power category: Not included	Issues
Physical infrastructure - other	Lights in data center's "white space"		√		Some have argued that only power path and cooling are to be included in efficiency calculations.
	Lights in mechanical & electrical rooms		√		This is easier to quantify in purpose built data centers but in multi-use buildings, a proportional share of the mechanical and electrical space lighting must be accounted for.
	Other plant controls (fire, dampers, physical security, BMS HVAC, PLCs)		√		HVAC - heating, ventilation, and air conditioning PLC - programmable logic controller BMS - building management system Includes servers used to control/manage the data center (i.e. BMS, security, fire)
	Outdoor lights			√	These are highly variable by site and many data centers are within multi- purpose facilities so it is very difficult to assign load to data center.
	Personnel office loads			√	Some have argued that there is no clear standard for deciding what personnel space is assigned as data center infrastructure, so it should not be included. Others say it is difficult to separate out. This space varies greatly and is best driven to an efficient consumption model though US Green Building Council (USGBC).
	Lights in data center's personnel areas			√	This is not included and this space varies greatly in mixed use facilities. It is best optimized thru the best practices of organizations like the USGBC.

Table 4 lists lighting and other devices that are typically found within data centers but do not appear in the previous tables. In general, lighting is included in data center energy calculations, but outdoor lighting is excluded due to the high degree of variation between sites. If outdoor lights are present and the energy use cannot be easily removed from the data center loads during measurement, they can be removed by subtracting their estimated values.

Note that personnel / office loads are specifically excluded from the data center efficiency calculation. This is because it is extremely difficult to assign personnel loads to the data center in many situations because in the vast majority of data center facilities there are non-data-center personnel functions in the same building, and the amount of load associated with non-data center personnel varies over a very wide range from site to site. It can be challenging to separate out data center loads from personnel office loads in many mixed-use facilities, but it does not need to be expensive or complex to do this. The next sections address this type of issue.

Part 2: Estimation of shared resources

Correctly and consistently classifying data center loads, as described in the prior section, is an important part of determining the PUE for a data center. However, as explained in the introduction, some devices that consume power and are associated with a data center are shared with other uses. For example, a data center may share a chiller plant with an adjacent office building, or the data center UPS may also provide power to a call center. Even an exact measurement of the energy use of such a shared device is unhelpful to the data center efficiency calculation, because the losses of that device associated with loads other than the data center should not be included in the PUE.

A common approach taken when a device is shared is to simply omit the device from the PUE calculations. This can result in major errors, especially if the device is a major energy user such as a chiller. Such an approach invalidates the PUE calculation for benchmarking purposes. A better strategy is to estimate (or indirectly measure) the fraction of the losses of the shared device that are associated with the data center and then use these losses in the PUE calculations. This approach can yield surprisingly accurate results.

Consider the case of a chiller plant that is shared with other non-data-center loads. The goal is to measure or estimate the fraction of the power that the chiller plant consumes that is related to the data center. Here are three alternative approaches to determining the energy use of a shared chiller plant:

1. Measure/estimate the thermal load on the chiller using known electrical losses of all other data center loads, measure/estimate the chiller efficiency performance, then use this information to calculate the electrical power that the chiller is using for data center loads
2. Measure/estimate the fractional split of the thermal load between the data center and other loads (using water temp, pressure, pump settings, etc), measure the chiller input power, then allocate the fraction of the chiller power to the data center according to the fractional split
3. Shut off the non-data center loads on the chiller, then measure the chiller to determine the chiller power associated with the data center[4]

These indirect measurements and estimations are typically made during an expert data center energy audit, but can be attempted by sophisticated data center operators. Once the technique is established for a specific data center, it is easy to re-use it over time for efficiency trending.

Similar techniques can be used for other types of shared resources. These estimation and modeling techniques can be formalized in the software management tools of the data center to provide continuous efficiency reporting. For example, the InfraStruXure Central data center management suite from Schneider Electric offers a continually increasing number of options to handle shared resources in its energy management tools.

[4] This allocation is not exact, and typically slightly overstates the losses assigned to the data center, because some losses in the chiller plant are fixed and do not stop when the non-data-center loads are switched off.

Part 3: Estimation of devices impractical to measure

Every power-consuming device in a data center can have its energy use measured. However, it can be complex, expensive, and impractical to measure energy use of some devices. In many cases, indirect measurement and estimation of devices can allow determination of the PUE in a practical and cost-effective manner.

Consider the case of a power distribution unit (PDU), which is clearly a power device. In a partially loaded data center, the losses in PDUs can be in excess of 10% of the IT load, with a significant effect on the PUE calculation. Yet most data center operators omit PDU losses in PUE calculations because they are considered so difficult to determine, causing a serious error in the PUE calculation.

There are several reasons why the losses of a PDU are very difficult to measure directly:

- Instrumentation within PDUs never directly provides loss information
- Input and output instrumentation in PDUs typically does not provide watt values, only VA or amps
- Input and output instrumentation in PDUs is not accurate enough to allow subtraction of output from input to determine losses
- There are a large number of outputs on a PDU that might need to be summed to obtain the output power

Fortunately, the losses in a PDU are quite deterministic, in that they can be directly calculated from the IT load if the characteristics of the PDU are provided. The losses of a PDU can therefore be estimated with a high degree of precision if the load is known in either watts, amps, or VA. In fact, estimating the losses in this way is typically MORE accurate than using built-in PDU instrumentation[5].

Once PDU losses are estimated, they are subtracted from the UPS output metering to obtain the IT load, and they are counted as part of the infrastructure load in determining the PUE. This simple method greatly improves the PUE calculation when compared to ignoring the PDU losses. This function can be built into the software tools used during an energy audit, or the function can be built directly into real-time energy management software, as it is in the InfraStruXure Central energy management software system.

These estimation techniques are described in more detail in other white papers listed at the end of this paper.

[5] PDU loss measurement errors are amplified because they are the difference between two large numbers. The errors are so great that it commonly causes PDU efficiency determined by direct measurement to exceed 100%, which is obviously impossible.

Conclusion

PUE is defined as the ratio of two numbers, **data center input power** over **IT load** power. While it at first appears to be a problem of simply obtaining two measurements and taking their ratio, it rarely is this simple in production data centers.

Data centers are often part of a building that has multiple purposes, so it can be impossible to find either a single point to measure the total data center input power or to measure the IT load. This means that the appropriate measurement points to capture all data center energy use need to be identified, and typically the power use data from different subsystems needs to be combined to find the total data center input power. To further complicate the matter, some of these power measurements may be impractical to obtain, or the devices may be shared with non-data center applications.

This paper has provided lists of the various data center subsystems that need to be included in energy use analysis, along with how the energy use of these subsystems should be incorpo-rated into PUE calculations. In addition, the paper has described practical approach-es to determining the energy use of shared devices and devices that are impractical to measure, in order to allow for PUE calculations.

Data center operators need to understand that determining PUE does not require extensive, expensive instrumentation because many losses in a data center can be very effectively estimated by indirect measurement and estimation.

While these methods appear difficult for typical users to implement, low cost, easy to use software tools are available to simplify the problem and allow any data center, large or small, new or retrofit, to have an effective real-time energy management system.

RATE THIS PAPER ★★★★★

About the author

Victor Avelar is a Senior Research Analyst at Schneider Electric. He is responsible for data center design and operations research, and consults with clients on risk assessment and design practices to optimize the availability and efficiency of their data center environments. Victor holds a Bachelor's degree in Mechanical Engineering from Rensselaer Polytechnic Institute and an MBA from Babson College. He is a member of AFCOM and the American Society for Quality.

Resources

Electrical Efficiency Measurement for Data Centers
White Paper 154

Electrical Efficiency Modeling for Data Centers
White Paper 113

Implementing Energy Efficient Data Centers
White Paper 114

Allocating Data Center Energy Costs and Carbon to IT users
White Paper 114

Browse all white papers
whitepapers.apc.com

Browse all 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 at
www.apc.com/support/contact/index.cfm