

Overview

The Power Usage Effectiveness (PUE) metric has become the de facto standard for measuring data center energy efficiency. PUE compares the total power going into a data center with the amount of power used to power IT equipment (servers, storage, and network). There is increasing pressure being exerted on data center managers to take measures to reduce the PUE. Unfortunately, the proper usage of PUE is often misunderstood and, by focusing solely on this single metric, it may mean data center managers are missing out on other opportunities to affect sustained reductions in energy use.

What is PUE?

The Power Usage Effectiveness (PUE) metric was introduced by the Green Grid, an association of IT professionals focused on increasing the energy efficiency of data centers. In the white paper *Green Grid Data Center Power Efficiency Metrics: PUE and DCiE* (Belady, Rawson, Pflueger, & Cader, 2007), the authors lay out the case for the introduction of metrics to measure energy efficiency in the data center.

The Green Grid believes that several metrics can help IT organizations better understand and improve the energy efficiency of their existing datacenters, as well as help them make smarter decisions on new datacenter deployments. In addition, these metrics provide a dependable way to measure their results against comparable IT organizations.

Without proper metrics in place, it is difficult to determine the effectiveness of changes made to improve data center energy efficiency.

There is a great deal of truth in the adage “You can’t manage what you can’t measure”. In order to manage energy efficiency in the data center, it is imperative to have metrics in place to measure the impact of changes. There were two primary metrics introduced, PUE and DCE (Data Center Efficiency). The latter was later changed to DCiE (Data Center Infrastructure Efficiency). Both metrics measure the same two parameters, the total power into the data center and the IT equipment power.

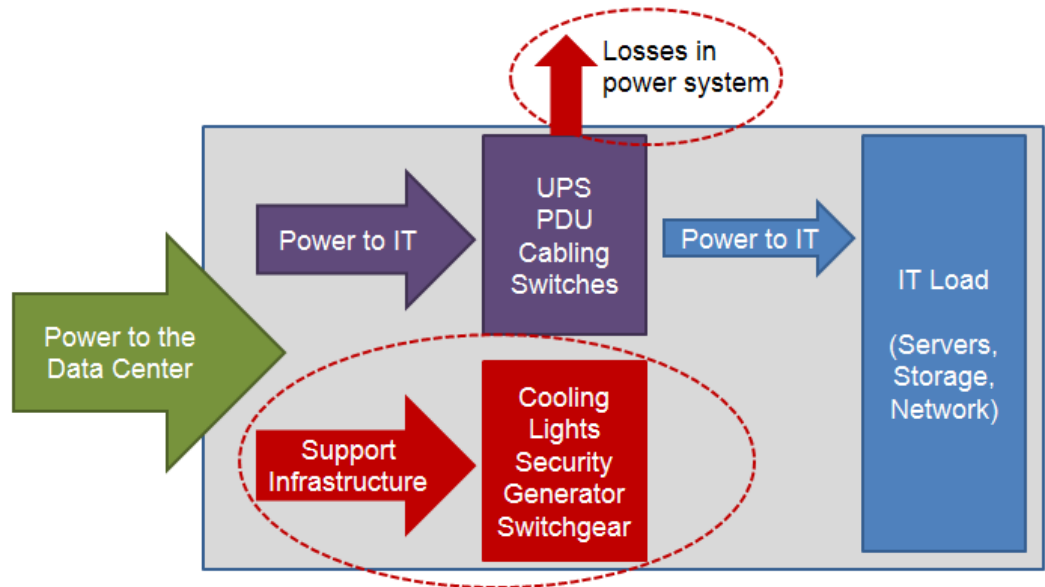
$$PUE = \frac{\text{Total Power into Data Center}}{\text{IT Equipment Power}}$$

$$DCiE = \frac{\text{IT Equipment Power}}{\text{Total Power into Data Center}}$$

While both metrics had their supporters, PUE became the standard metric.

A PUE value of 1 would represent the optimal data center efficiency. In practical terms, a PUE value of 1 means that all power going into the data center is being used to power IT equipment. Anything above a value of 1 means there is data center overhead required to support the IT load.

What components make up this overhead? Let's look at where the power going into a data center is consumed.



Reducing losses in the power system and the support infrastructure will improve data center energy efficiency.

Ideally, we would like all power entering the data center to be used to power the IT load (servers, storage and network). This would result in a PUE value of 1. Realistically, however, some of this power must be diverted to support cooling, lighting and other support infrastructure. Some of the remaining power is consumed due to losses in the power system. The remaining power then goes to service the IT load.

Let's look at an example to see how PUE is calculated. If the power entering the data center (measured at the utility meter) is 100 kW and the power consumed by the IT load (measured at the output of the UPS) is 50 kW, we would calculate PUE as follows:

$$PUE = \frac{100}{50} = 2.0$$

A PUE value of 2.0 is fairly typical for a data center. This means that for every watt required to power a server, we actually consume 2 watts of power. It is important to remember that we are paying for the power entering the data center, so every watt of overhead represents an additional cost. Reducing this overhead will reduce our overall operating costs for the data center.

If we want to improve data center energy efficiency, there are two areas in which we can affect change. If we can reduce the power going to the support infrastructure or reduce losses in the power system, more of the power entering the data center will make it to the IT load. This will improve our energy efficiency and reduce our PUE.

Are There Drawbacks to Using PUE?

PUE is a great tool for the facilities side of the data center. It allows facility engineers to measure the impact of changes they make to the infrastructure, things like raising the data center temperature, upgrading to a higher efficiency UPS, increasing voltage to the rack and so on. PUE must be used with care, however. It must be understood that IT changes can have a dramatic impact on PUE.

Under pressure to reduce costs, and in some cases to try to match the reported PUE from other companies, data center managers are being pushed to significantly reduce their PUE value. Unfortunately, this is not always the right approach. The drive to reduce PUE at all costs can actually have a negative impact. If data center managers focus only on reducing PUE, they may inadvertently use more energy and increase data center costs.

Let's run through an example on how this can happen. Suppose we have a data center which has input power of 100 kW, 50kW of which is being used to power IT equipment. As previously illustrated, this would give us an initial PUE value of 2.0.

$$PUE = \frac{100}{50} = 2.0$$

Suppose we now decide to virtualize a number of servers. In fact, we are so successful with virtualization that we are able to reduce the power to IT equipment by 25 kW and the overall power to our data center by the same amount. What will happen to our PUE?

$$PUE(\text{after virtualization}) = \frac{75}{25} = 3.0$$

Wait a minute – isn't a higher PUE value something we want to avoid? Not necessarily. It is important to understand what can cause the PUE to increase or decrease. While it may seem counterintuitive, any reduction in IT load without an equivalent reduction in infrastructure load will actually result in a higher PUE.

This becomes easier to understand if we break PUE down into its components:

$$PUE = \frac{IT\ Load + Infrastructure\ Load}{IT\ Load} = 1 + \frac{Infrastructure\ Load}{IT\ Load}$$

When the IT load is reduced, $\frac{Infrastructure\ Load}{IT\ Load}$ will always increase, resulting in an increase in the PUE. Inversely, increasing the IT load will always decrease the PUE.

So, if our PUE has gone up, does this mean the data center is now less energy efficient? No, the data center is now *more* energy efficient. We are accomplishing the

Decreasing the overall energy usage in the data center may result in a HIGHER value for your PUE!

same work while using less energy at less cost. To illustrate this, let's calculate the annual energy usage and cost both before and after virtualization.

Before virtualization:

$$\text{Annual Energy Use} = 100 \text{ kW} * 8,760 \text{ hrs/yr} = 876,000 \text{ kWh}$$

$$\text{Annual Electric Cost} = 876,000 \text{ kWh} * \$0.10/\text{kwh} = \$ 87,600$$

After virtualization:


$$\text{Annual Energy Use} = 75 \text{ kW} * 8,760 \text{ hrs/yr} = 657,000 \text{ kWh}$$

$$\text{Annual Electric Cost} = 657,000 \text{ kWh} * \$0.10/\text{kwh} = \$ 65,700$$

The virtualized data center is clearly more energy efficient. In fact, the data center can probably be made even more energy efficient if the support infrastructure is now reduced to more closely match the reduced IT load.


PUE by itself is a meaningless number if we don't know how to use it to measure the results of changes in the data center. Since we know that virtualization will likely increase the PUE, should we avoid it? Of course not! It is important, however, to note when the virtualization took place when we examine our PUE over time. In addition to tracking PUE, we must also track any changes in the infrastructure or IT load so we can correlate the changes to the PUE value.

When tracking PUE it is also important to track IT load and infrastructure changes so we know WHY the PUE changed.



It must be understood that there are many other factors which may impact PUE. There will always be tradeoffs between availability and energy efficiency. Redundancy, for example, will increase PUE. Data center equipment – from cooling equipment to UPSs to server power supplies – will run more efficiently when more heavily loaded.

Consider the following graph of server power supply efficiency versus the load on the power supply (Ecos and EPRI, 2008). The power supply operates more efficiently as the load increases, with the highest efficiency attained with loads of 40% and higher.



Similarly shaped curves can be found for UPS, computer room air conditioning and other support infrastructure.

Redundancy, while improving availability, will reduce the load across multiple systems. As the load is reduced, energy efficiency will also be reduced. Virtualization and consolidation, while reducing the overall energy usage, will actually increase the PUE unless the power and cooling infrastructure are downsized to align with the IT load. Raising the server inlet temperature may reduce the PUE, but the overall energy usage may actually increase if the increased power required for server fans is greater than the cooling savings.


The bottom line is that PUE, while an important piece of the energy efficiency puzzle, is just that – one piece of the puzzle. PUE is only one component of a comprehensive energy management program which must consider both the IT and facility sides of the house.

PUE is just one component of a comprehensive energy management program.

What Should I Be Measuring Besides PUE?

PUE is best used for tracking the impact of changes made to the data center infrastructure. It is less useful for tracking the improvements resulting from reducing the energy consumption of IT equipment.

Let's return for a moment to the diagram of the power usage in the data center.



“Reductions in energy consumption at the IT equipment level have the greatest impact on overall consumption because they cascade across all supporting systems”.

While it is important to reduce losses in the power system and the power used for the support infrastructure, we need to realize that the bulk of the power consumption in the data center goes to the IT load itself. If we can reduce the IT load, we will reduce the overall power required for the data center.

In fact, reducing the IT load has a compounding effect, as it will also reduce the losses in the power system and the power required for the support infrastructure. In the Liebert white paper *Energy Logic: Reducing Data Center Energy Consumption by Creating Savings that Cascade Across Systems* (Emerson Network Power, 2009), the authors call this the “cascade effect” and state that “reductions in energy consumption at the IT equipment level have the greatest impact on overall consumption because they cascade across all supporting systems.”

Let's see how this cascade effect works. If one watt can be saved at the IT load, this will also reduce losses in the server power supply (AC to DC conversion), reduce losses in the power distribution (PDU transformers, losses in the wiring itself), reduce power losses in the UPS, reduce the amount of cooling required and, finally, reduce power losses in the building transformer and switchgear. The end result of the cascade effect is that saving one watt at the IT load may actually result in two or more watts of overall energy savings.

We must keep in mind that the real goal, whether from an environmental or cost standpoint, is to reduce overall energy usage. Since powering the IT load is such a large portion of the overall electricity cost in a data center, reduction of the IT load must be a primary consideration in any energy efficiency initiative. In order to measure our success in reducing the IT load, power usage at the IT device level must become a critical metric.

How Much Does it Cost to Run a Server?

So how does PUE translate into cost in the data center? The annual cost to run a server can be calculated with the following formula:

$$\text{Server Power Cost} = \frac{\text{Average Power (W)}}{1000} * 8760 \frac{\text{hours}}{\text{year}} * \text{PUE} * \frac{\text{Cost}}{\text{kWh}}$$

The annual cost of powering a 400 watt server can be \$800 or more.

The annual cost to power a server drawing 400 watts in a data center with a PUE of 2.25 and a utility cost of \$0.10/kWh would be calculated as follows:

$$\text{Server Power Cost} = \frac{400}{1000} * 8760 * 2.25 * \$0.10 = \$788.40$$

Note that this \$788.40 only includes the power cost. It does not include such items as software licenses and other operational costs.

It becomes clear that there are two ways to reduce the server power cost: (1) reduce the average server power, or (2) reduce the PUE. Much has been written about reducing the PUE and there are a number of tools from various vendors which will measure the infrastructure power usage. What is less talked about is how to reduce the IT power usage. There are also fewer tools to measure the IT power usage, particularly down to the device level.

What Can I Do to Reduce the IT Load?

Since powering the IT load is such a large portion of the overall electricity cost in a data center, reduction of the IT load must be a primary consideration in any energy efficiency initiative. There are a number of ways to reduce the IT load including the following:

- Decommission or repurpose servers which are no longer in use
- Power down servers when not in use
- Enable power management
- Replace inefficient servers
- Virtualize or consolidate servers

It is estimated that 10-15% of data center servers are “ghost servers”.

These servers can consume 70-85% of the power of a server running at 100% CPU usage while producing no useful output.

Decommission or Repurpose Servers

In the *Server Energy and Efficiency Report 2009* (Alliance to Save Energy, 1E, 2009) the authors commissioned Kelton Research to conduct a survey of IT professionals responsible for server operations at large global organizations. Based on the results of the survey, it was estimated there are 4.7 million servers running 24/7 worldwide which are not doing anything useful. These are sometimes called “ghost servers”. Since “ghost servers” consume 70-85% of the power of a server running at 100% CPU load but produce no useful output, it is important to identify these servers and to either decommission or repurpose them.

Data center managers have always struggled with how to identify unused or lightly-used servers. The typical method is to use CPU utilization as a measure of whether or not a server is being actively used. Unfortunately, this doesn’t always tell the whole story. A server may appear to be busy when it is actually only performing secondary or tertiary processing not related directly to the primary services of the server. In defining primary services, the authors of *The Green Grid Data Center Compute Efficiency Metric: DCcE* (Blackburn, Azevedo, Ortiz, Tipley, & Van Den Berghe, 2010) provide a great example.


As a generic example, the primary service of an e-mail server is to provide e-mail. This same server may also provide monitoring services, backup services, antivirus services, etc., but those are secondary, tertiary, and similar types of service. If the e-mail server stops being accessed for e-mail, the monitoring, backup, and antivirus services may no longer be necessary, but the server may still continue to provide them. So from a CPU-utilization standpoint, the unused server may appear to be busy, but that may only be secondary or tertiary processing. Hence, CPU utilization is not a precise enough measure.

The authors introduce the Server Compute Efficiency (ScE) metric as a means to determine whether or not a server is being used for primary services. The ScE metric measures CPU usage, disk and network I/O, incoming session-based connection requests and interactive logins to determine if the server is providing primary services. When this data is plotted over time, a clearer picture of server usage begins to form.

The ScE metric can provide data center managers with the ability to determine which servers are providing primary services as well as provide some insights into which servers may be good candidates for virtualization or consolidation.

The following graph shows when each ScE component (CPU, Disk I/O, etc.) is active over the course of a day. A higher stacked bar indicates that more components were active at the time of the snapshot. An empty space indicates time periods where no primary services were being performed.

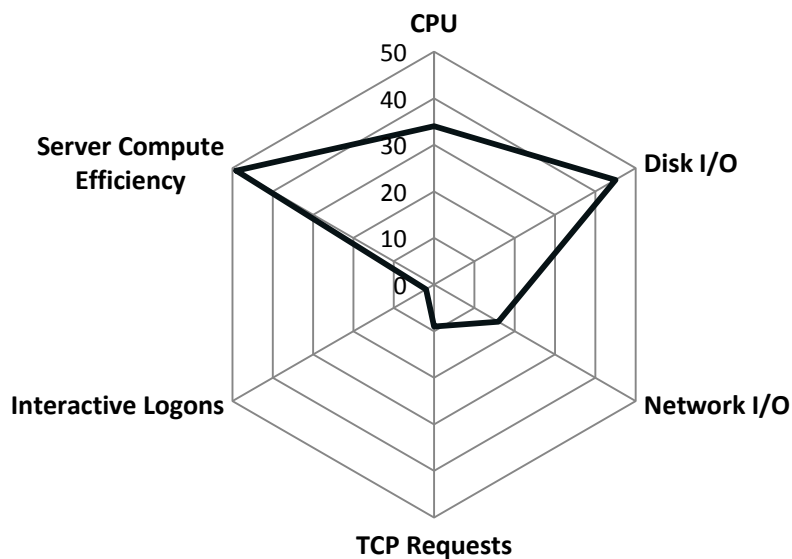
Server Compute Efficiency (ScE) By Component



Server Compute Efficiency (ScE) provides a metric for measuring server activity.

Another way of looking at this data is to compute the percentage of time the server was engaged in providing each of the ScE components. The following radar chart shows the server has an overall server compute efficiency of 50%, meaning that the server was providing primary services 50% of the time. We can also see the percentages of activity by CPU, Disk I/O, Network I/O, TCP request and interactive logons, providing us a much more detailed view of the server activity than simply looking at CPU.

ScE Components (% Active)



We power down office computers when not in use, so why don't we do the same with servers?

Power Down Servers When Not in Use

We power down personal computers when we go home at night, so why don't we do the same with servers when they aren't being used? While the majority of servers in data centers may be utilized around the clock, there are some servers which may only be used during certain parts of the day or week. When time periods of inactivity can be identified, the servers should be turned off during these times. If it has been determined that a server is no longer in use, turn it off while waiting to decommission or repurpose it.

Enable Power Management

Many processors and operating systems provide the ability to reduce a server's power consumption during periods of low utilization. In the book Energy Efficiency for Information Technology (Minas & Ellison, 2009) the authors explain that altering the P-state of the CPU can save as much as 75 percent of the full-speed CPU power.

Altering the CPU's P-state can reduce a server's power consumption at low utilization without limiting any of the performance needed at peak demand levels. The switch between P-states is dynamically controlled by the operating system and occurs in micro-seconds, causing no perceptible performance degradation.

Servers respond very quickly to P-state changes so the processor power and speed closely follow the workload. As a result, heat does not build up unnecessarily, which also provides for greater power savings from reduced cooling loads.

By enabling Demand-Based Switching (DBS), significant savings can be realized in the data center.

Enabling power management on 500 servers could save \$100,000 or more in annual energy costs!

	Typical CPU Utilization		
	15%	30%	45%
System Power with DBS Off	258 W	291 W	316 W
System Power with DBS On	201 W	220 W	240 W
DBS Power Savings Per System	57 W	71 W	76 W
Energy Cost Savings Per System*	\$148	\$186	\$200
Energy Cost Savings Per 500 Systems*	\$74,000	\$93,000	\$100,000

* Annual cost: assumes \$0.10/kWh and cooling costs double that of platform power

Source: Intel, 2008

Replace Inefficient Servers

Once a server has been purchased, there is sometimes a tendency to consider it a “sunk cost”. What is not taken into account in this argument, however, is the ongoing operational costs, including power, cooling, software licensing and so on. With the cost to power a server now exceeding the purchase cost, the ongoing operational costs become a prime factor in determining when to retire an aging server.

A new multi-core server may replace as many as 15 single-core servers, saving as much as 93% of the power usage. In addition to the power savings, software licensing and other maintenance costs can also be considerably reduced. Additional savings include a reduction in data center cooling costs and the potential to reclaim valuable rack space.

Virtualize or Consolidate Servers

There are a number of excellent reasons to virtualize or consolidate servers. From a business continuity viewpoint, virtual machines can be isolated from physical system failures to increase system availability. In addition, parallel virtual environments allow for an easier transition to a backup facility.

From an energy efficiency viewpoint, virtualization provides a number of opportunities for energy savings. Virtual machines provide much more granular control over workloads. Virtual machines can be moved to additional servers as demand increases, while unused or lightly-used servers can be managed to minimize power usage. Overall, virtualization can increase server CPU usage by 40-60%. As the CPU usage is increased, the energy efficiency of the server power supply will also increase. Virtualization and consolidation will have the effect of decreasing the total number of servers or, at the least, deferring the purchase of new servers.

By looking beyond PUE, data center managers are more likely to find other good opportunities to reduce overall energy usage.

Summary

While the Power Usage Effectiveness (PUE) metric can provide valuable information for measuring data center energy efficiency, it should be considered as only one component in a comprehensive energy management program. While there is increasing pressure being placed on data center managers to reduce the PUE, doing so without a full understanding of power usage in the data center might actually have a detrimental effect.

By considering other data center metrics such as energy usage at the IT device level and server compute efficiency, data center managers are more likely to find other excellent opportunities to affect sustained reductions in energy usage.