

How to Achieve Lower PUE in Your Data Center

Sean Evanuk, EIT
Mitsubishi Electric Power Products, Inc.
UPS Division
Application Engineer
2/18/15

What is PUE and Why Should You Care?

Power usage effectiveness (PUE) is a metric defined as the ratio of the total facility energy to the IT equipment energy.

$$PUE = \frac{\text{Total Facility Energy}}{\text{IT Equipment Energy}}$$

This data center metric was developed by the Green Grid Association (TGG) in 2007 (The Green Grid, 2007). The refinement and widespread adoption of the metric has led TGG to consolidate their previously published content related to PUE into an all-encompassing white paper (The Green Grid, 2012). The TGG white paper #49, "PUE: A Comprehensive Examination of the Metric", is an invaluable resource for any data center wanting to report and use PUE.

In short, the PUE describes how well a data center delivers energy to IT load and supports it. PUE does not, nor does it try to, measure the computational efficiency of the IT load. It is best used as a tool to identify opportunities to improve operational efficiency within the data center.

Knowing where power is utilized will help determine where the most gains can be made. Chris Malone, Thermal Technologies Architect at Google, said at the 2011 Uptime Symposium, "We have built this infrastructure up using best practices which are so simple that everybody should employ them and see good results" (uptimeinstitute, 2011). This article will explore some simple techniques that can be used to improve PUE.

Where is the Power Being Used in the Data Center?

Let's first begin with a look at how the energy use is distributed inside an example data center. For this example, it will be assumed the PUE was measured at 2.0. A data center with a PUE of 2.0 would have 50% of its energy being consumed by IT equipment and 50% consumed by the overhead associated with running the data center and supporting the IT load. In other words, every watt used for IT load will require an additional watt of overhead. Figure 1 depicts a hypothetical breakdown of the total data center energy.

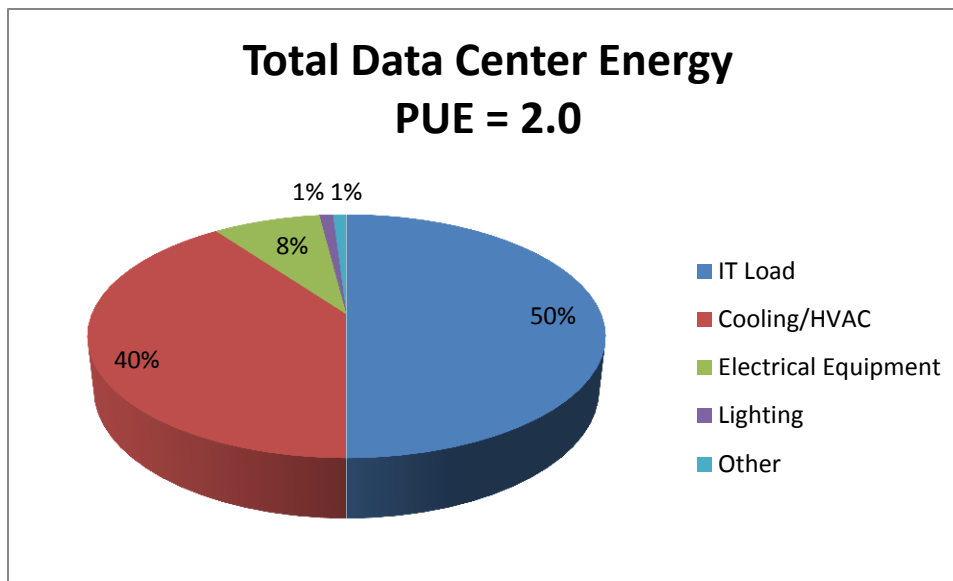


Figure 1: Total Data Center Energy Breakdown (U.S. Environmental Protection Agency, 2007)

It can be seen that the majority of the power being used to support the IT load in this instance comes from cooling and HVAC. The rest of the energy is being used by electrical equipment, lighting, and other sources such as fire suppression systems and security. This means the biggest gains can be made by decreasing the energy used for cooling.

This is the philosophy Google has taken. Again at the Uptime Institute Symposium, Chris Malone said, “Fix cooling first. It’s the biggest term in your overhead so it offers the most opportunity for improvement” (uptimeinstitute, 2011).

After cooling, the power distribution infrastructure is the second best area to look for optimizations. With proper measurement of PUE, the least efficient areas of the data center can be targeted. Proper tracking of PUE also allows optimizations to be evaluated. Data center operation cannot improve without knowing the current state of the data center. This requires taking specific measurements throughout the data center. This leads to the next section of what is needed to measure PUE.

Looking at Your Current PUE

TGG defines a basic (Level 1), intermediate (Level 2), and advanced (Level 3) measurement for PUE (The Green Grid, 2012). All three levels require the total facility energy to be measured at the utility input, although the different levels recommend additional measurement points. The different levels also differ in the required measurement for the IT load.

The level 1 measurement will provide a basic indication of the PUE by means of measuring the IT load at the Uninterruptible Power Supply (UPS) outputs. The level 2 measurement will provide a more accurate PUE figure by measuring the IT load from the Power Distribution Unit (PDU) outputs. The advanced level 3 measurement provides the most accurate PUE figure by measuring the IT load at the input of the IT equipment. This can be accomplished through metered rack PDUs.

Additional metering locations can provide more insight into your data center’s energy usage. The data collected can be fed into computer models that will allow a data center operator to predict the effects

of various changes. Beneficial locations for additional monitoring would be on the outputs of automatic transfer switches (ATS), inputs of UPS, and the inputs of mechanical equipment. Branch circuit monitoring can also provide power usage at the rack level. Measuring at these different locations not only provides additional granularity for insight but can be used as inputs for a computational fluid dynamics (CFD) model.

A CFD model can provide an accurate picture of the cooling capabilities of a data center by using numerical methods and algorithms. There are numerous software packages that provide user interfaces for creating CFD models specifically for data centers. Use of this software provides helpful information for upgrading infrastructure in an existing data center. This includes determining where hot spots have been created, where hot/cold air is mixing, and how to right size the cooling.

Accurately performing a CFD study requires not only power usage information but information on the equipment installed and taking airflow measurements on all perforated tiles. After a CFD study is performed, identifying improvements for cooling and airflow management becomes an easier task. Improvements can be achieved using simple best practices, addressed next.

Cooling Down Your Energy Usage by Warming Up the Data Center?

There are some basic strategies that can be deployed for a typical raised floor data center. If a CFD study has been done, problematic areas can be targeted. The fundamental goals are to prevent the mixing of hot and cold air and to optimize cooling throughout the data center. This can be accomplished using several simple solutions.

- 1) Use a hot/cold row layout
- 2) Use Blanking panels and grommets
- 3) Eliminate under floor obstructions.
- 4) Use hot or cold aisle containment
- 5) Increase inlet temperature to the racks
- 6) Make use of economizers and free cooling

A common hot/cold aisle layout in a data center arranges the rows of racks in an alternating hot row cold row configuration. With the rows of racks placed in a back to back configuration, the hot air exhaust of the equipment is separated from the cold air used by the equipment. With the hot/cold aisle layout it is important to keep all perforated tiles in the cold row. This maintains the separation of the hot and cold air.

To achieve maximum effectiveness of using hot and cold rows, it is imperative to use blanking panels in cabinets and to seal off any holes in the data center floor. The use of brushed grommets can be used for cable cutouts to seal any holes in the raised floor. It is also important to eliminate under floor obstructions. Any under floor cable tray should be routed in the hot rows and not routed directly in front of the Computer Room Air Conditioner (CRAC) units.

CFD studies have shown the end of aisles and the top of cabinets are the most prevalent location for mixing of hot and cold air. This mixing can be eliminated by using a containment system. There is a choice of hot and cold aisle containment. There are arguments for both types of containment systems and the best choice will depend on the data center. One advantage cold aisle containment can offer is the ability to do partial containment. A partial containment system would create barriers at the end of the rows but not use any barrier at the top of the racks. This will still lead to considerable prevention of mixing hot and cold air.

Another recent trend seen amongst data centers is the increase of the inlet temperature to the racks. In 2008, the American Society of Heating, Refrigerating and Air-Conditioning (ASHRAE) expanded the recommended temperature range for IT equipment (ASHRAE TC 9.9, 2011). The maximum recommended operating temperature was increased from 77 degrees F to 80 degrees F. Raising the required equipment inlet temperature will allow the supply temperature of the CRAC units to be increased and easier use of economizers. The use of outside air by economizers allows for significant reductions in HVAC energy.

Deliver Efficiency with UPS

The electrical infrastructure is the second largest overhead associated with the data center. The distribution system can be optimized by minimizing the resistance and power losses. This can be achieved by making runs as short as possible, bringing higher voltages to the equipment, and using more efficient pieces of equipment. One particular piece of equipment that has seen significant efficiency increases in recent years is the uninterruptible power supply.

The UPS has continued to evolve and has become more efficient over the years. Old UPS installations use transformers and large filters due to 6 pulse and 12 pulse converter sections. These units have a typical efficiency of 90% to 95%.

New technology has allowed UPS systems to become more efficient and create less heat. This all equates to energy savings within the data center and lower PUE. Making use of a transformer-less UPS with a three level IGBT converter can allow for efficiencies up to 97%. Not only are the losses decreased in the power electronics, the elimination of transformers increases efficiencies and reduces cooling requirements.

The latest advancement is the use of silicon carbide (SiC) transistors which will allow for operational efficiencies up to 98%. Upgrading the UPS in the data center can have a major impact on the PUE.

Don't Forget About the Small Things

Efficient lighting can play a large part in lowering the PUE even though it makes up only a small percentage of the total load. LED bulbs, motion sensors, and lighting control are key components for optimizing the lighting. The goal is to have the lights off when the light is not needed. A simple approach can be implemented using motion based occupancy sensors to turn the lights in the data center.

An advanced technique that could be implemented is one that provides illumination only where the technician is working. As the technician moves around the data center, the lighting can dynamically change. This is accomplished via a sensor network tied into a central control system. In addition to controlling the lighting, the centralized control system can also collect data that can provide insight into energy usage.

LED lighting will use less energy than the common fluorescent lighting typically found in a data center. Using LED lighting with a central power-conversion module can also save on energy by eliminating electrical losses and decreasing heat. This will have the added benefit of reducing the HVAC load.

Beyond PUE and Current Trends

PUE is an effective metric for determining how efficient the infrastructure is being utilized. In some cases increasing efficiency in the data center might actually raise the PUE. An example is replacing computing equipment with more efficient equipment but the cooling is not changed to reflect the lower amount of energy being used.

So the overhead energy stays roughly the same while the IT load decreases. This can result in the PUE increasing even though the data center is consuming less energy. The PUE though does show that the data center energy is not being used effectively.

There are also trends in the industry to make use of containerized systems. A container would house all required cooling and power. The container would be optimized to ensure a very low PUE. All the best practices can be isolated and used in the containerized system.

Machine learning and neural networks are also being used to create a smarter data center. Since the data center is a complex system, machine learning allows one to see how a simple change can impact the rest of the data center. This will allow for temporary change of the infrastructure to allow efficiencies to be maintained during a planned or unplanned outage.

The use of best practices can allow data centers to efficiently use the infrastructure. Large companies such as Google are leading downward trends of PUE citing a 12 month trailing PUE of 1.12. It might be unrealistic for a colocation provider to reach the same PUE levels of Google but implementing some or all of the aforementioned techniques can help any data center lower their PUE.

Works Cited

ASHRAE TC 9.9. (2011). *2011 Thermal Guidelines for Data Processing Environments - Expanded Data Center Classes and Usage Guidance*. Retrieved March 4, 2015

The Green Grid. (2007, February 20). *Green Grid Metrics: Describing DataCenter Power Efficiency*. Retrieved February 18, 2015, from www.thegreengrid.org: <http://www.thegreengrid.org/en/Global/Content/white-papers/Green-Grid-Metrics>

The Green Grid. (2012). *PUE: A Comprehensive Examination of the Metric*. Retrieved February 18, 2015, from www.thegreengrid.org: <http://www.thegreengrid.org/en/Global/Content/white-papers/WP49-PUEAComprehensiveExaminationoftheMetric>

U.S. Environmental Protection Agency. (2007, August 2). *Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431*. Retrieved February 24, 2015, from www.energystar.gov: http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf

uptimeinstitute. (2011, June 1). *Google Data Center Efficiency Lessons that apply to all Data Centers*. Retrieved February 20, 2015, from <https://www.youtube.com/watch?v=0m-yRYEMZVY>