Determining PUE in an Imperfect World

Dr. Jay's Technical Tip for Measuring PUE in Dedicated and Mixed-Use Facilities

Modius, Inc.
71 Stevenson Street, Suite 400
San Francisco, CA 94105
www.modius.com
(888) 323-0066

modius

Dr. Jay's Technical Tip
Determining PUE in an Imperfect World

I am routinely surprised at how difficult it can be to determine the total energy consumption for many data centers. Stand-alone data centers can at least look at the monthly bill from the utility, but as The Green Grid[1] points out when discussing PUE metrics, continuous monitoring is preferred whenever possible. Moreover, measurement in an environment where power and cooling infrastructure is shared with non-data center facilities can be more even more complex.

A common situation is a data center in a mixed environment, such as a corporate campus or an office tower, in which the chilled-water system—often referred to as the "mechanical yard"—is a shared resource. As difficult as it sometimes can be to set up continuous power monitoring for a stand-alone data center, it is considerably trickier when the mechanical yard is shared. Again, it's simple in principle, but often surprisingly painful in practice.

One way to address this problem is to use The Green Grid's partial PUE, or pPUE. While the number should not be used as a comparison against other data centers, it provides a metric to use for tracking improvements within a single data center.

This isn't always a satisfactory approach, however. Given that there is a mechanical yard, it's pretty much guaranteed to be a major component of the overall non-IT power overhead. Using a pPUE of the remaining system and not measuring, or at least estimating, the mechanical yard's contribution masks both the overall impact of the data center *and* the impact of any efficiency improvements you make.

There are a number of ways to incorporate the mechanical yard in the PUE calculations. Full instrumentation is always nice to have, but most of us have to fall back on approximations. To be meaningful, the approximations should go beyond a simple fixed percentage or constant value for unknown quantities. For instance, the approximations should account for variation in time, including seasonality associated with weather and business cycles. Fundamentally, you want to know how much energy the mechanical yard consumes and what portion of the cooling load is allocated to the data center.

The Perfect World

In an ideal situation, you have sub-metering for the mechanical yard's power—i.e., chillers, cooling towers, and all associated pumps and fans. It's not unusual to have a single distribution point where measurement can be made, or perhaps even a dedicated automatic transfer switch (ATS). Then for the ideal solution, all you need is sub-metering of the chilled-water going into the data center. The calculus is explained as such:

The heat load, h, of any fluid cooling system can be calculated from the temperature change, ΔT, and the overall flow rate, q: $h = Cq\Delta T$, where C is a constant that depends on the type of fluid and the units used. As much as I dislike non-metric units, it is easy to remember that C=500 when temperature is in °F and flow rate is in gal/min, giving heat load in BTU/h. (Please don't tell my physics instructors I used

1 http://thegreengrid.org/

BTUs in public.) Regardless of units, the total power to allocate to your data center overhead is $P_{dc}=P_{mech}\ (h_{dc}/h_{mech})$. Since what matters is the ratio, the constant C cancels out and you have $P_{dc}=P_{mech}\ (q\Delta T_{dc}/q\Delta T_{mech})$.

You are pretty much guaranteed to have the overall temperature and flow data for the main chilled-water loop in the BMS system already, so you have $q\Delta T_{mech}$. You are much less likely to have the same data for just the pipes going in and out of your data center. If you do, hurrah, you're in The Perfect World, and you're probably already monitoring your full PUE and didn't need to read this article at all.

Perfect and You Don't Even Know It

Don't forget to check the information from your floor-level cooling equipment as well. Some of them do measure and report their own chilled-water statistics, in which case no additional instrumentation is needed. In the interest of brand neutrality, I won't go into specific names and models in this article, but feel free to contact me[2] with questions about the information available from different equipment.

Perfect Retrofit

If you're not already sub-metered, but you have access to a straight stretch of pipe at least a couple feet long, then consider installing an ultrasonic flow meter. You'll need to strap a transmitter and a receiver to the pipe, under the insulation, typically at least a foot apart along the pipe. No need to stop the flow or interrupt operation in any way. Either inflow or outflow is fine. If they're not the same, get a mop; you have other more pressing problems. Focus on leak detection, not energy monitoring.

If the pipe is metal, then place surface temperature sensors directly on the outside of the inflow and outflow pipes, and insulate them well from the outside air. Might not be the exact same temperature as the water, but you can get very close, and you're really most concerned about the temperature difference anyway. For non-metal pipes, you will have to insert probes into the water flow. You might have available access ports, if you're lucky.

PUE in an Imperfect World

But, what about some of the options available for the large population of data centers that don't have perfect instrumentation, and can't afford the time and/or money to purchase and install it right now?

Let's revisit the example of the most common shared resource—the chilled water system—from a common campus or building mechanical yard. We looked at the simple way to allocate a portion of the power consumed by the mechanical equipment to the overall power consumed by the data center.

The approach there assumed perfect sub-metering of both the power and chilled water, for both the data center and the mechanical yard. It's a lovely situation if you have it or can afford to quickly achieve it, but not terribly common out in the hard, cold world. Thus, we must turn to estimates and approximations.

Of course, any approximations made will degrade the ability to compare PUEs across facilities—already a tricky task. The primary goal is to provide a metric to measure improvement. Here are a few scenarios that fall short of the ideal, but will give you something to work with:

[2] jay.hartley@modius.com

- *Can't measure data-center heat load, but have good electrical sub-metering.* Use electrical power as a substitute for cooling load. Every watt going in ends up as heat, and there usually aren't too many people in the space routinely. Works best if you're also measuring the power to all other non-data-center cooled space. The ratio of the two will get you close to the ratio of their cooling loads. If there are people in a space routinely, add 1 kWh of load per head per 8-hr day of light office work.

- *Water temperature is easy, but can't install a flow meter.* Many CRAHs control their cooling power through a variable valve. Reported "Cooling Load" is actually the percentage opening of the valve. Get the valve characteristics curve from the manufacturer. Your monitoring system can then convert the cooling load to an estimated flow. Add up the flows from all CRAHs to get the total.

- *Have the heat loads, but don't know the mechanical yard's electrical power.* Use a clamp-on hand meter to take some spot measurements. From this you can calculate a Coefficient of Performance (COP) for the mechanical yard, i.e., the power consumed per cooling power delivered. Try to measure it at a couple of different load levels, as the real COP will depend on the % load.

- *I've got no information about the mechanical yard.* Not true. The control system knows the overall load on the mechanical yard. It knows which pumps are on, how many compressor stages are operating, and whether the cooling-tower fan is running. If you have variable-speed drives, it knows what speed they're running. You should be able to get from the manufacturer at least a nominal COP curve for the tower and chiller and nominal power curves for pumps and fans. Somebody had all these numbers when they designed the system, after all.

Whatever number you come up with, perform a sanity check against the Departm ent o f E nergy's DCPro online tool[3]. Are you in the ballpark? Heads up, DCPro will ask you many questions about your facility that you may or may not be prepared to answer. For that reason alone, it's an excellent exercise.

It's interesting to note that even the Perfect World of absolute instrumentation can expose some unexpected inter-dependencies. Since the efficiency of the mechanical yard depends on its overall load level, the value of the data-center PUE can be affected by the load level in the rest of the facility. During off hours, when the overall load drops in the office space, the data center will have a larger share of the chilled-water resource. The chiller and/or cooling-tower efficiency will decline at the same time. The resulting increase in instantaneous data center PUE does not reflect a sudden problem in the data center's operations; though it might suggest overall efficiency improvements in the control strategy.

PUE is a very simple metric, just a ratio of two power measurements, but depending on your specific facility configuration and level of instrumentation, it can be remarkably tricky to "get it right." Thus, we are confronted with the ever-expanding array of tier levels and partial alternative measurements. Relatively small incremental investments can steadily improve the quality of your estimates. When reporting to management, don't hide the fact that you are providing an estimated value. You'll only buy yourself more grief later when the reported PUE changes significantly due to an improvement in the calculation itself, instead of any real operational changes.

[3] http://www1.eere.energy.gov/industry/saveenergynow/pdfs/dc_pro-energy_profiler.pdf

The trade-off in coming to a reasonable overall PUE is between investing in instrumentation and investing in a bit of research about your equipment and the associated estimation calculations. In either case, studying the resulting number as it varies over the hours, days, and seasons can provide excellent insight into the operational behavior of your data center.

Continuous monitoring of total power is critical to managing the capacity and efficiency of a data center. Whether your concern is PUE, carbon footprint, or simply reducing the energy bill, the monthly report from the utility won't always provide the information you need to identify specific opportunities for improvement. Even smart meters might not be granular enough to identify short-term surges, and won't allow you to correlate the data with that from other equipment in your facility. It's hard to justify skimping on one or two meters for a new data center. Even in a retrofit situation, consider a dedicated meter as an early step in your efficiency efforts.

For more information about how you can overcome challenges to monitoring and uncovering the true state of your data center (including server rooms, branch offices and IDF closets), watch this demo[4] of OpenData software by Modius.

[4] https://modiusevents.webex.com/modiusevents/lsr.php?AT=pb&SP=EC&rID=2621382&rKey=c1e99b02a6f51f0e