Uptime Institute | IT Symposium 2009

LEAN, CLEAN & GREEN

Raritan

Measuring Power in your Data Center

Herman Chan, Director, Power Solutions Business Unit, and Greg More, Senior
Product Marketing Manager

There are many ways power has been allocated, estimated, and sometimes measured in data centers. This white paper will discuss the reasons why data center power should be measured, where power consumption should be measured, what to do with the data once it is gathered, and some of the new technologies available.

Why measure power in data centers?

Suppose your phone rings tomorrow and the CIO is on the line asking, "What are we doing about power consumption in our data centers?" What will you say? Typically, data center managers have not worried about power consumption. Other areas of the company might live without air conditioning after 7 pm, but not the data center. If additional IT-based services were needed, the necessary hardware was purchased and plugged in. This approach is quickly changing as 1) additional power is often not available, 2) the cost of power is becoming a significant cost of operating a data center, and 3) companies are placing a higher value on green initiatives so as to be perceived as good corporate citizens and to prepare for likely environmental regulations.

You can't manage what you don't measure. This is particularly true for power, where rule-of-thumb estimates can turn out to be just plain wrong, leading to unnecessary and sometimes quite substantial costs. Devices that were thought to be consuming very little power may be consuming quite a lot, even while simply sitting idle doing no useful work.

The first step is to baseline current power consumption. Ideally, this will be done in a way that provides useful statistics to be compared over time. Early measurements and estimates may be rough but can be refined over time as the power deployment inside and outside the data center is better understood and as the measurement quality improves.

There are many ways to manage power consumption in a data center but without some baseline measurements it is difficult to know where to start or what efforts will have the greatest impact. Also, without baseline measurements it is impossible to show management past levels of consumption and how you have improved.

Efficiency metrics

An efficiency metric receiving a lot of attention is the power usage effectiveness (PUE). It is the ratio of the total energy used by a data center, including IT equipment, and the energy consumed by the IT equipment only. The total energy includes lighting, cooling and air movement equipment and inefficiencies in electricity distribution within the data center. The IT equipment portion is that equipment which performs computational tasks.

$$PUE = \text{Total Facility Power/IT Equipment Power}$$

A data center that only supplies power to IT equipment would have a PUE = 1.0 because the numerator and denominator would both be IT equipment power. This is obviously not a realistic situation. Even in a lights-out data center, power will be consumed to provide cooling and air movement and there will be electrical distribution inefficiencies.[1]

Corporate average data center efficiency (CADE) takes into account the energy efficiency of facilities, their utilization rates, and the level of server utilization.

$$CADE = \text{(Facility Efficiency) x (IT Asset Efficiency)}$$

- Facility Efficiency = Energy delivered to IT / energy drawn from utilities
- IT Asset Efficiency = Average central processing unit (CPU) utilization across all servers (which is often a small percentage, e.g., 5 percent) until efficiency efforts like virtualization are undertaken.

Where and how to measure

In a data center there are several locations where power can be measured. Moving from the coarsest measurement to the most detailed, the first is the power entering the data center. If the data center is a stand-alone structure this is simply the power feed from the utility. This would be the total power number in the numerator of a PUE calculation.

Very often it's not that easy. The data center may be a floor in a building in which case a submeter for that floor or room should be installed. This submeter would record the total power number provided the data center doesn't share power or building facilities such as cooling equipment. If facilities and power are shared, which is often the case particularly in urban data centers, then work will need to be done to at least get an estimate of the total power consumption of the data center, possibly from several different sources (e.g., the submeter measuring the feed into the data center plus some percentage of the power used by the building cooling equipment).

The next place where power is often measured is at the uninterruptible power supply (UPS). If it only provides power to IT equipment then this data can be used for the

---

[1] Data center infrastructure efficiency (DCiE) is simply the inverse of PUE and won't be covered here.

denominator of a PUE calculation. However, the UPS may also provide power to rack-based cooling equipment.

A third place to measure power is at the rack itself with metered rack power distribution units (PDUs). These figures are generally considered to represent the IT equipment, aggregated to a rack, unless there are fans or rack-side cooling units.

A fourth place to measure power is at the individual outlets of a rack PDU. These intelligent PDUs also typically provide aggregated rack power consumption as well. Monitoring the power at the outlet level ensures that IT equipment power consumption can be identified for a PUE calculation. By providing power information at the individual device level, specific actions can be taken to improve efficiency.

The fifth place to measure power is at the CPU. This gives the purest measurement of what power is actually going into doing purely computational work. In practice, this is not widely used today. In terms of taking actual energy conservation actions, the CPU level is not very useful since, in most cases, an entire device, blade or other piece of IT equipment is what data center staff can change or decommission, not a CPU. The most typical approaches to measuring power consumption in a data center are metered rack PDUs and intelligent rack PDUs that monitor individual outlets.

What to do with the data gathered

Depending on the measurement locations and method of measurement chosen, various energy efficiency initiatives may be undertaken. Individual outlet-level metering is recommended for IT equipment because if provides useful, actionable information.

Monitoring the power consumed at a rack allows data center managers to determine if their original power allocations make sense today. Quite often, power is allocated to IT equipment on the basis of nameplate ratings, which are conservatively high. Even when a percentage, say 70 percent, of nameplate power is used power is often over-allocated. This means more power is going to an IT equipment rack than what will actually be consumed. This "stranded power" could be deployed elsewhere but how do you know you're not leaving the rack vulnerable to running out of power in a peak load situation?

Monitor each individual device at regular intervals, the shorter the better, to ensure that no peak periods are overlooked. With individual device power consumption figures it is possible to set up racks such that equipment

power consumption patterns compliment each other and thus more IT equipment can be supported with the same amount of power. If a rack is close to consuming all the power allocated to it, and therefore at risk of tripping a breaker, having individual IT equipment power consumption data allows IT staff to remove equipment in a logical manner so as to minimize the risk of a breaker tripping while maintaining useful loading levels.

Through tests in its own data center, Raritan determined that rules-of-thumb percentages of nameplate ratings simple don't work. Across 59 servers, 15 had average power consumption of 20 percent or less, 29 had 21 percent to 40 percent, 9 had 41 percent to 60 percent, 4 had 61 percent to 80 percent and 2 had 81 percent or more. Even at peak power consumption, 49 of the servers were 60 percent or less of their nameplate rating. Many data center planners use 70 percent of nameplate, which means there is a lot of stranded power in many data centers.

On the other hand, at peak power consumption 5 of the 59 servers were at 81 percent or more of nameplate and therefore at risk of shutting down. The message is that in terms of power consumption, it is important to know what is going on at the individual device, not some aggregated average which may mask problems both on the high and low side.

Environmental sensors—their impact on power and cooling efficiency

Environmental sensors make an important contribution to power efficiency. It is not uncommon for cooling to consume 30 percent or more of a data center's total power. Vendors provide inlet temperature specifications. As long as the inlet temperature is within the specification the server will perform fine. These specifications are often substantially higher than what is typically provided in data center cold aisles. Thus, the temperature can often be turned up to reduce cooling equipment power consumption.

Temperature sensors should be placed at the bottom, middle and top third of racks on the cool air inlet side. Cooling IT equipment to temperatures lower than required consumes a lot of power without any beneficial effects. Due to a lack of at-the-rack instrumentation, data center managers often overcool to be confident IT equipment won't fail.

New technologies available

Taking an individual snapshot of power consumption at one point in time is not sufficient. IT devices may consume a lot less power at 2 am than they do at 8 am and may hit

peak power consumption at 4 pm on Thursday. Power consumption can also vary by time of year, such as online sales during December.

There are hardware devices that can take snapshots of power consumption at user-defined intervals as often as once every few seconds. Software programs are available to turn these data points into calculations of power usage where the unit of measure is kilowatt hours (kWh). Sophisticated tools can calculate carbon footprints based on energy usage. With actual individual device information data center staff can know the biggest contributors to carbon generation and therefore what needs to be most closely managed.

What to look for

Accurate: As carbon caps, credits and trading are adopted, accuracy becomes important. A +/- 5 percent accuracy, assuming perfect sine waves (which rarely occur in the real world), may be acceptable to determine if a rack is operating with about a 25 percent margin before circuit breakers trip. It is not acceptable when dealing with regulations and carbon credits to be verified and traded on exchanges. Nor is it accurate enough for billing or charge backs.

Open and interoperable: Many data centers have deployed an IT management system. To tie such a system to power measurement look for open standards for integration and interoperability with existing equipment. Ease of use is a key consideration so power management does not become a time-consuming project for already busy IT staff.

Secure: Power is the life blood of data centers. It is important that access to the power management system be secure. Look for systems with high levels of encryption such as 256-bit advanced encryption standard (AES) and the ability to set authentication, authorization, and permissions.

Conclusion

We hope that if the CIO calls you tomorrow and asks, "What are we doing about power consumption in our data centers?" you'll refer to this white paper and outline a plan beginning with a program to gather information to establish some baselines. Collecting data now, and taking a stab at data center metrics such as a PUE calculation, will put you on a path to more efficiently manage power and power costs. And to take calls from CIOs.

About the Author

Herman Chan is Raritan's Director, Power Solutions Business Unit based in Somerset, NJ. This business unit focuses on building data center infrastructure management solutions to solve problems associated with improving uptime, capacity planning and efficiency. Herman has over 15 years of experience in sales and marketing management, and product management for network and data center infrastructure solutions. He has previously held management positions with Fountain Technologies and Panasonic. Herman holds a degree in Economics from Rutgers University and is ITIL certified. Herman can be reached at herman.chan@raritan.com.

Greg More is Senior Product Marketing Manager for Raritan's Power Solutions Business Unit based in Somerset, NJ. Greg has over 15 years of experience in technology marketing and product management. He has previously held positions with Motorola and Integrated Network Corporation. Greg received his MBA degree from Northwestern University's Kellogg School of Management and a BS Industrial Engineering degree, also from Northwestern.

About Raritan

Raritan's energy management solution can be a key component of the Uptime Institute's "Four Metrics that Define Data Center Greenness." The Raritan energy management solution—intelligent rack PDUs, integrated environmental sensors and energy management software—is designed to help data center IT and facilities managers get a handle on efficiently managing power consumption down to the individual device.

The energy management software can output values for current draw, kWh, and carbon footprint. These values can be aggregated and sorted in a variety of ways. For example, kWh usage can be presented by department, current draw can be organized by type of server and carbon footprint can be shown by geographic region, data center, or row or rack of IT equipment.

By measuring actual power consumption at the hardware level, monitoring rack temperatures, and consolidating all the information into easily exported spreadsheets, reports and graphs, data center staff can create accurate baselines and track the effectiveness of green IT efforts. They can easily report the contributions they make to achieve corporate environmental goals such as reducing carbon emissions.

About the Uptime Institute

Uptime Institute is a leading global authority on data centers. Since 1993, it has provided education, consulting, knowledge networks, and expert advisory for data center Facilities and IT organizations interested in maximizing site infrastructure uptime availability. It has pioneered numerous industry innovations, including the Tier Classification System for data center availability, which serves as a de facto industry standard. Site Uptime Network is a private knowledge network with 100 global corporate and government members, mostly at the scale of Fortune 100-sized organizations in North America and EMEA. In 2008, the Institute launched an individual Institute membership program. For the industry as a whole, the Institute certifies data center Tier level and site resiliency, provides site sustainability assessments, and assists data center owners in planning and justifying data center projects. It publishes papers and reports, offers seminars, and produces an annual Green Enterprise IT Symposium, the premier event in the field focused primarily on improving enterprise IT and data center computing energy efficiency. It also sponsors the annual Green Enterprise IT Awards and the Global Green 100 programs. The Institute conducts custom surveys, research and product certifications for industry manufacturers. All Institute published materials are © 2009 Uptime Institute, Inc., and protected by international copyright law, all rights reserved, for all media and all uses. Written permission is required to reproduce all or any portion of the Institute's literature for any purpose. To download the reprint permission request form, uptimeinstitute.org/resources.