Efficient Power Measurement in the Data Centre

Judging the Juice

We spoke to Herman Chan, Director and Greg More, Senior Product Marketing Manager, for Raritan's Power Solutions Business Unit.

NETCOMMS EUROPE: Power consumption wasn't always the first thing on a data centre manager's mind. So, what's changed?

RARITAN: There are a number of factors – additional power is often not available, the price of power is becoming a significant cost in operating a data centre, and, not least importantly, businesses are placing a higher value on green initiatives.

Based on the simple premise that you can't manage what you can't measure, data centres are undertaking steps to measure device-level power consumption. Rule-of-thumb estimates often turn out to be wrong, leading to unnecessary, and sometimes quite substantial, costs – e.g. devices that are thought to be consuming very little power may actually be consuming quite a lot, even while idle.

NE: So where should you start to measure device-level power consumption?

RARITAN: The first step towards measurement 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 as the power deployment inside and outside the data centre is better understood and as the measurement quality improves.

There are many ways to manage power consumption in a data centre, 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 what improvements have been made.

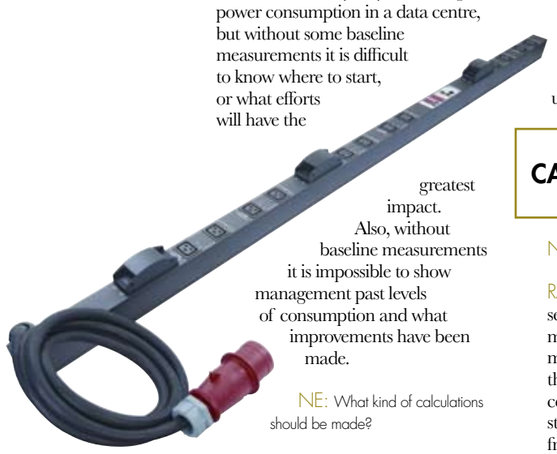

NE: What kind of calculations should be made?

RARITAN: An efficiency metric that receives a lot of attention is the power usage effectiveness (PUE). It is the ratio of the total energy used by a data centre, 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 centre. The IT equipment portion is that equipment which performs computational tasks.

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

A data centre which 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 centre, power will be consumed to provide cooling and air movement, and there will be electrical distribution inefficiencies. DCiE is simply the inverse of PUE.

Corporate average data centre efficiency (CADE) takes into account the energy efficiency of facilities, their utilisation rates and the level of utilisation of servers.
Facility Efficiency = Energy delivered to IT / energy drawn from utilities

IT Asset Efficiency = Average CPU utilization across all servers, often a small percentage such as 5%, until efficiency efforts like virtualisation are undertaken.

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

NE: Where should the power be measured?

RARITAN: In a data centre, 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 centre. If the data centre 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.

Unfortunately, though, it's not often that easy. The data centre 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 centre 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 centres, then work will need to be done to get at least an estimate of the total power consumption of the data centre, possibly from several different sources, such as the submeter measuring the feed into the data centre, plus some percentage of the power used by the building cooling equipment.

The next place where power is often measured is at the UPS. If it only provides power to IT equipment, then this data can be used as an approximation for the denominator of a PUE calculation. However, this is only an approximation, because the power inefficiencies of the UPS itself should not be part of the IT equipment power. 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 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 uniquely 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 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 centre staff can change or decommission, not a CPU. The most typical approaches to measuring power consumption in a data centre are metered rack PDUs and intelligent rack PDUs that monitor individual outlets.

NE: Once you've got this data, how can you use it to make your data centre more environmentally friendly?

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

Monitoring the power consumed at a rack allows data centre 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%, 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 that you are not leaving the rack vulnerable to running out of power in a peak load situation?

Monitoring each individual device at regular intervals, the shorter the better, ensures that no peak periods are overlooked. With individual device power consumption figures, it is possible to set up racks so 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 minimise the risk of a breaker tripping while maintaining useful loading levels.

NE: What have been your findings in your own data centre?

RARITAN: We found that rule-of-thumb percentages of nameplate ratings simply don't work. Across 59 servers, fifteen had average power consumption of 20% or less, twenty-nine had 21% to 40%, nine had 41% to 60%, four had 61% to 80% and two had 81% or more. Even at peak power consumption, 49 of the servers were 60% or less of their nameplate rating. As data centre planners use 70% of nameplate, this means there is a lot of stranded power in many data centres.

On the other hand, at peak power consumption, five of the 59 servers were at 81% or more of nameplate, and therefore at risk of shutting down. In terms of power consumption, therefore, it is important to know what is going on at the individual device, not from an aggregated average which may mask problems, both on the high and low side.

NE: How effective are environmental sensors on power and cooling efficiency in data centres?

RARITAN: Environmental sensors do make an important contribution to power efficiency, as it is common for cooling to consume 30% or more of a data centre's total power. IT equipment 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 centre cold aisles, so the temperature can often be turned up, which leads to less power consumption by the cooling equipment.

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 centre managers often overcool to be confident their IT equipment won't fail.

NE: What new technologies are available to enable better measurement?

RARITAN: 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 also calculate carbon footprints based on energy usage. With actual individual device information, data centre staff can know the biggest contributors to carbon generation, and therefore what needs to be most closely managed.

Taking an individual snapshot of power consumption at one point in time is not always sufficient, as IT devices may consume a lot less power at 2 a.m. than they do at 8 a.m. and may hit peak power consumption at 4 p.m. on Thursday. Power consumption can also vary by time of year, such as online sales during December.

NE: What should data centre managers be looking for in power measurement?

RARITAN: Accuracy. As carbon caps, credits and trading are adopted, accuracy becomes important. +/- 5% 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% margin before circuit breakers trip. It is not, though, 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.

Another thing to consider is that the measurement is open and interoperable. As many data centres deploy an IT management system, to link this to power management, open stands for integration and interoperability with existing equipment are essential. Ease of use is also a key consideration, so that power management does not become a time-consuming project for already busy IT staff.

The power measurement also has to be secure, as power is the life blood of all data centres. Systems with high levels of encryption, such as 256-bit AES and the ability to set authentication, authorisation and permissions, are highly worthwhile.