

Accurately Measuring Datacenter Power Efficiency

Timely and accurate measurements with SunSM Eco Services



Summary

- Accurately measuring power consumption in today's datacenters is difficult because datacenters are in constant change. This makes it a challenge to pin down the impact of energy-efficiency improvements.
- Workloads fluctuate depending on human activity, adding to the challenge of obtaining accurate power-consumption metrics.
- SunSM Eco Services measures at a higher level rather than attempting to measure each individual device's power consumption. This helps build a statistical model that can be used to predict and track the impact of energy-efficiency changes.
- Sun's approach gives timely, cost-effective measurements that can be used to motivate efficiency improvements, which can save companies both power and real estate expenses.



Nothing makes an argument more compelling than the cold, hard facts. There is also no better way to motivate an organization toward increased datacenter power efficiency than with accurate, aggregated energy-consumption data. Solid data does more than convince an organization that efficiency can cut expenses — a baseline measurement gives a starting place and point of comparison for datacenter-efficiency improvements.

The problem with standard measurement techniques is that they are time consuming and costly to produce good results. Power measurements must be taken over a period of time long enough to incorporate the changes in power consumption caused by periodic workload fluctuations. When this is considered, along with the fact that datacenters are constantly changing entities, measurements can be inaccurate and out of date by the time they have been gathered.

So how does an IT organization measure datacenter power consumption accurately enough to reflect efficiency improvements, yet fast enough so that the information is timely and actionable? Sun has a solution, one that SunSM Eco Services can implement in your datacenter.

Using solid metrics

The most useful datacenter power-efficiency metric is the ratio of IT equipment load to total utility power:

$$\text{Efficiency} = \frac{\text{IT Load}}{\text{Total Facility Power}}$$

In a typical datacenter, this ratio is 0.3, meaning that only about 30 percent of the total utility draw actually powers servers. Another third of the power is used to remove the heat that servers generate. And roughly another third is consumed by power conversions in uninterruptible power supplies (UPS), power distribution units (PDUs), and miscellaneous power draws ranging from humidification to spot cooling systems.

Reducing the numerator, for example, by replacing old, energy-inefficient servers with UltraSPARC® T1 processor-powered servers results in a lower IT load that is amplified by lower cooling costs and increased efficiency. Reducing the denominator, for example, by increasing cooling-system efficiency, also results in a higher efficiency number. In addition to the value of the ratio, the absolute values of the numerator and denominator can be used to calculate total energy costs.

Measuring power in the datacenter power hierarchy

From a power standpoint, a datacenter is a hierarchy of electrical devices. Utility power feeds one or more UPS, which in turn feed a greater number of PDUs, each of which then feed multiple racks.

Measuring power consumption lower in the hierarchy gives more accurate measurements, but it takes longer to obtain results. The more devices you measure, the more expensive the measurement and the more dynamic the target becomes. There is only one power meter for the facility, but there are tens of PDUs and thousands of servers.

The key to effective datacenter power measurement is finding a timely and cost-effective measurement technique. Typical best practices use special equipment to measure each server and even the volume of chilled water circulating through the cooling infrastructure. These low-level measurements can be used to create an energy model for the datacenter that aggregates the low-level data into higher-level statistics.

In contrast, measuring power consumption higher in the datacenter power hierarchy yields fast, inexpensive, and accurate results so long as the data is adjusted for changes in the datacenter configuration during and after the measurements are made.

Human-generated workloads

One factor complicating energy-consumption measurement is the cyclical nature of workloads imposed by human beings. When measuring power consumption for servers responding to cyclical workload variations, measurements must be taken across enough cycles so that the workload and its variability are well understood. Measurements taken without understanding the cyclical nature of the workload can have variability greater than the amount of energy savings obtained in any one conservation step, thus obscuring the benefits of the improvements.

Workloads can vary by the hour of the day, day of the week, and week of the quarter. Retailers typically have much higher IT workloads during the holiday season. Some organizations book more sales at the end of the month, even more at the end of the quarter, and the most at the end of the fiscal year. A retailer's power consumption in January compared to December might erroneously suggest that some dramatic savings had taken place when the source of the change is normal business fluctuations.

Experience at Sun

Sun uses Sun Ray thin clients throughout its workplaces in order to save the power and administrative costs of desktop workstations. These appliances display a user's desktop, with all of the real work taking place on a server.

The average daily workload plot shown in green illustrates a ramp-up of work in the morning, a peak just before lunch, a dip while employees typically have lunch, and another workload peak when people return from lunch. This workload varies by the day of the week, with Tuesday being the busiest day and Friday being the slowest. The red lines illustrate the variability in these measurements and serve as a reminder that if you try to measure a change in energy efficiency of only a few percent, it's likely to be overshadowed by the range of error in the measurement itself.

Measuring efficiency improvement is further complicated by the fact that energy-saving technology is likely to be phased in over a number of months to a year, at the same time that numerous other business-related changes take place, for example, deploying new servers to handle a growing business. Consider aggregating multiple, small improvements so that their impact can be measured more easily.

Sun's measurement approach

Sun's approach for measuring IT power consumption in existing datacenters begins with aggregate measurements at the PDU level, where software or manual methods can be used to track power consumption for all IT devices over significant periods of time. Accurate measurements can be made, for example, by monitoring PDU power consumption over multiple business cycles.

Measuring at a higher level in the datacenter power hierarchy eliminates the high cost of taking point measurements at each server and moving measurement equipment from device to device. In a short amount of time, it can yield the IT-load-to-facility-power ratio that indicates efficiency.

Creating a statistical model

With measurements taken at the PDU level, a statistical model of IT load can be created by correlating each PDU's power consumption with the devices that they power. This yields a system of linear equations that can be solved to obtain average power-consumption data per server type.

Having statistics on average power consumption per server type rather than for individual server configurations provides an adequate level of detail that allows actionable data to be generated quickly and at low cost. Using average data rather than data for individual servers evens out the variability between configuration and workload differences for each server type.

Averages allow close estimates of future power consumption to be made without the pitfalls of using a smaller sample size to predict future energy use.

Measuring the impact of changes

The beauty of the statistical model is that it can be used to generate new power-consumption estimates without remeasuring each PDU over multiple business cycles. Datacenters change on a daily basis, much faster than the business cycles over which power must be measured. The statistical model, however, can be used to generate overall power-consumption estimates based on any moment's snapshot of datacenter assets.

Given power-consumption averages for each server type, overall power-consumption statistics can be generated from the datacenter asset list. Simply multiply the consumption of each server type by the number of servers of each type; the result is the overall IT load.

This momentary estimate of IT load can be divided by facility power to obtain a datacenter-efficiency metric that accurately tracks the moment-by-moment changes in datacenter configuration. The ratio of IT load to facility power factors out the impact of growth and contraction due to business cycles, helping organizations to focus on improving efficiency even as overall business — and power consumption — may be expanding.

Calculating confidence intervals

The statistical model, in fact, any measurement approach, requires determining a confidence interval for the data. A confidence interval can give immediate insight into whether, for example, a 3-percent reduction in IT load can actually be detected using the statistical model or the measurement technique. In the case of many datacenter power measurements, changes that have small impacts on power consumption are difficult to measure given the rate of change of datacenter assets and workload-induced power-consumption fluctuations.

Measuring datacenters of the future

The method of measuring IT load at the UPS level is an excellent starting point for building a statistical power model and is appropriate for existing datacenters and those without more sophisticated monitoring capabilities. Future datacenters with fine-grained monitoring capabilities present new opportunities and challenges.

At Sun, these datacenters are in operation today. In order to consolidate its IT infrastructure and its real estate holdings, Sun created a "pod" design that could be replicated in its sites worldwide. A pod is a self-contained group of racks and/or benches that optimize power,

cooling, and cabling efficiencies to facilitate rapid and simple replication throughout the datacenter. Sun's energy-efficient datacenter design emphasizes flexibility and modularity, incorporating dramatic improvements in power and cooling infrastructure. The design has the ability to accommodate high-density infrastructure today and is ready to handle even higher densities in the future. The result of Sun's effort with its Santa Clara, California, datacenter is an 80-percent reduction in datacenter-related real estate holdings and more than a 60-percent reduction in datacenter facility power consumption.

Granular, in-rack power monitoring

Sun has designed its datacenters with in-rack power distribution units that can provide minute-by-minute power-consumption data on a per-rack and even a per-socket basis. With accurate power-consumption data available for each device in the datacenter, it can be tempting to model power consumption on a per-server basis, when average power consumption per server type is more useful for predicting future datacenter power consumption.

Driving the statistical model

Per-server power-consumption data can be used to drive the statistical model with data averaged across server types or averaged across server types qualified by major differences in configuration, including amounts of memory, disk drives, and peripherals. The important factor to consider is workload characteristics, including the effect of human-induced workloads. Even though datacenters can provide minute-by-minute power-consumption data, it must be measured over a significant enough time to achieve the accuracy required to observe incremental improvements in efficiency.

Putting Sun's methodology to work

Corporate attention spans can be short, and the best way to get attention and motivate efficiency improvements is to be able to generate accurate, up-to-date power statistics on a moment's notice.

Taking accurate measurements takes time, however, and a statistical model built from a set of baseline measurements provides a means of closely estimating efficiency improvements even while a datacenter is in a state of constant change.

Sun Eco Services can help you develop a statistical model for your datacenter, empowering you to generate timely and cost-effective datacenter power-efficiency measures.

Learn More

For more information on Sun Eco Services and energy-efficient datacenters, contact your Sun sales representative or visit sun.com/eco.