

Unlock Data-Center Power Capacity with Software Defined Power

[Electronic Design](#)

[Mark Adams](#)

Fri, 2017-01-13 12:43



The booming demand for cloud computing and data services will only accelerate as the number of conventional computer-based users is rapidly dwarfed by the multitude of connected “things” that the veritable Internet of Things (IoT) threatens to bring about. Challenges will emerge in trying to scale up the capacity of networks and data centers to support the 50 billion sensors, controls, and other “things” that Cisco forecasts will require access to the internet by 2020. Perhaps more challenging will be addressing the constraints of data-center power capacity.

New infrastructure will inevitably be required, but building new data centers takes time and considerable investment, and “more of the same” isn’t necessarily the answer. Neither is simply adding more processing and storage capacity to existing data centers. Why? Because one of the most constrained resources is power, and data-center power capacity often is the first to become exhausted.

Related

[5 Methods For Improving Data Center Efficiency](#)

[What Trends Emerged at APEC 2016?](#)

[Software Monitors Data Center Components over the Cloud](#)

The solution to this problem isn’t just providing more power capacity, as this is expensive and time-consuming. Plus, just adding more power supplies assumes the local utility can support the additional load, which for a typical data center is forecast to grow from 2015 levels of around 48 GW to 96 GW by 2021.

So what’s the answer? Well, firstly, two key factors limit how much of the installed capacity of a traditional power architecture is actually available to use. Historically, data centers have had to ensure high availability to cope with mission-critical processing workloads. This need has been met through supply redundancy, which often leaves power supplies idle. Then there’s the way power is partitioned between server racks, which may result in supplies operating well below their load capacity.

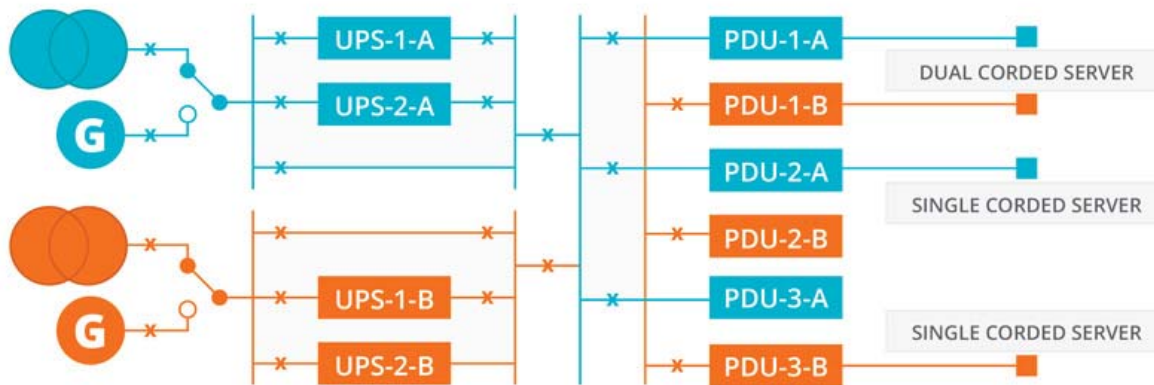
Power provisioning also needs to allow for fluctuations in the highly dynamic processing loads experienced by data-center servers. Light processing loads typically translate to lightly loaded power supplies, which will always operate less efficiently than fully loaded supplies.

What’s needed here is a method to even out power-supply loads, perhaps by redistributing processing tasks to other servers or by pausing non-time-critical tasks or rescheduling them to quieter times of day. Other methods can address demand fluctuations by using battery-power storage to meet peak demands without impacting the load presented to the utility supply.

Intelligence Quotient

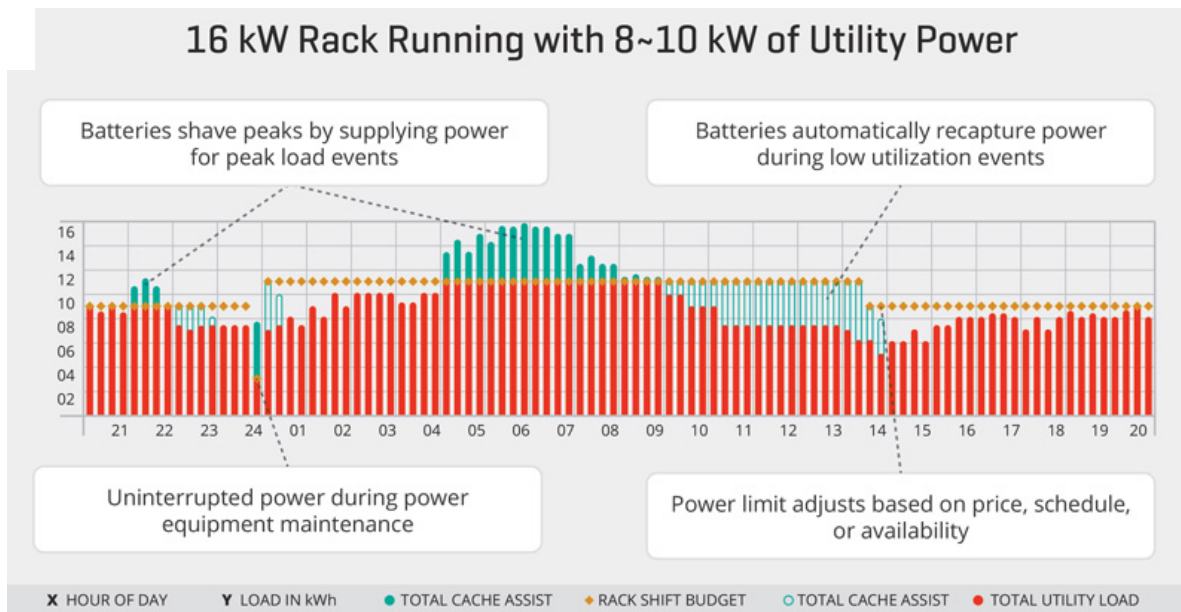
What characterizes all these potential solutions is the need for greater intelligence, not just in the management of the data center's processing operations, but in the way power is managed. One potential solution is the Software Defined Power (SDP) that might unlock the underutilized power capacity available within existing systems. To understand the benefits it can offer, let's consider some examples:

- A tier 3 or tier 4 data center, which has to provide 100% redundancy for mission-critical tasks, uses a 2N power architecture (*Fig. 1*), where essentially every element in the power-supply path is duplicated. However, the reality is that not all tasks in such a data center are mission-critical. If, for example, 30% of the servers are handling test or development tasks, then potentially half the power provisioned for them isn't really required (i.e., 15% of the data center's total power capacity could be redeployed elsewhere).



- Server power-supply capacity is conventionally designed to handle peak CPU utilization. However, depending on the task being undertaken, that utilization and the consequent power demand can vary considerably. Google found that the average-to-peak power ratio for servers handling web mail is 89.9%, while web search activity has a lower ratio of 72.7%. This means that provisioning power for web search could result in surplus capacity, by up to 17%.

- At the server rack level, where power is being drawn from the utility supply, loads can vary predictably throughout a day or change very rapidly. Longer duration peaks in demand can be met with generating sets to supplement a constrained utility supply. However, for more dynamic situations, battery storage is a better solution—the batteries store energy during periods of low demand in order to supply power during peak load events. This is typically referred to as “peak-shaving.” The example in *Figure 2* shows how a 16-kW rack can be operated with a more even 8 to 10 kW of utility power.



Implementing Software Defined Power in data centers and similar IT systems doesn't need to be difficult. At its simplest, SDP uses feedback to control the distributed power architecture found in data server racks, dynamically responding to changing loads and adjusting intermediate and point-of-load voltages to maximize the efficiency of every supply in the chain. On a broader front, SDP becomes part of a higher-level hardware and software solution that embraces techniques such as peak-shaving and dynamic power sourcing.

One such solution might be the Intelligent Control of Energy (ICE), which is offered by Virtual Power Systems in partnership with [CUI](#) as a complete power-management capability that can be deployed in both existing and new data-center installations. By releasing redundant capacity from non-mission critical systems and using its power switching and Li-ion battery storage modules for load distribution and maximizing utilization, ICE claims to deliver up to 50% reduction in the total cost of ownership of server power installations.

To learn more, visit <http://www.cui.com/sdp-infrastructure-solutions>.



Looking for parts? Go to sourceesb.com

Source URL: <http://electronicdesign.com/power/unlock-data-center-power-capacity-software-defined-power>