

INDUSTRY INITIATIVES, SUCH AS THE GREEN GRID AND CLIMATE SAVERS, HAVE SET THE AMBITIOUS GOAL OF HALVING DATA-CENTER ENERGY USAGE BY 2011. TO ACHIEVE THIS OBJECTIVE, THE INDUSTRY MUST SHIFT TO THE USE OF DIGITAL-POWER CONTROL AND MANAGEMENT AT THE AC/DC- AND DC/DC-POWER-SUPPLY LEVEL. A SHIFT TO DC-BUS-POWER DISTRIBUTION WITHIN THE DATA-CENTER INFRASTRUCTURE ALSO SHOWS PROMISE.



Eltek Valere's Flatpack2 HE 48/2000 rectifier module offers 96% typical power-conversion efficiency, which the company claims is the highest efficiency level available.

# PUSH FOR POWER EFFICIENCY

## FORCES CHANGES IN SERVER-CENTER HARDWARE AND SOFTWARE

BY MARGERY CONNER • TECHNICAL EDITOR

Advances in server-processing power have coupled with increasing demands on workloads to drive up data-center power consumption in the United States. Currently, data centers, or server “farms,” account for 1.5% of the total US electricity bill, and that figure should rise to nearly 3% by 2011 (Reference 1). In addition, data-center-electricity costs now outpace hardware costs over the four-year life of a server. Purchasers of servers, who previously pinched pennies when specifying server hardware, now are open to spending a little more for hardware if a more sophisticated and efficient power-supply design will yield lower power bills. The need to reduce electricity costs and the desire to responsibly use energy have made power efficiency at the hardware, software, and infrastructure levels a top priority for data centers. And, because data centers are among the largest end applications for power supplies, their needs will influence the agenda for designers and

manufacturers of power supplies.

The US government's EPA (Environmental Protection Agency) encourages data-center efficiency through efforts such as Energy Star, which has just released the first draft of its “Program Requirements for Computer Servers” (Reference 2). Industry groups such

as the Green Grid and Climate Savers are developing guidelines for best practices, promulgating new technologies, and generally shining the light on the need for greater power efficiency in businesses.

Until recently, the data-center industry focused on leveraging facility space rather than power usage. As servers got smaller in size and cost, their physical facility's price and maintenance remained the same or increased. Thus, it made sense to cram as many servers as possible into each facility. However, a by-product of operating a server is heat: For every watt a server processor consumes, it wastes another half-watt in cooling; ac/dc- and dc/dc-power conversion and regulation cost another half-watt. As long as energy costs were relatively low, it made sense to solve

the cooling problem by using air conditioners. As energy costs soared, however, this approach became economically and environmentally impractical.

This concern for power efficiency is relatively recent in the data-center world, according to Tom Darby, who manages Texas Instruments' data centers and is the company's representative on the Green Grid. As an example, he points out that IT (information-technology)-industry trade shows once focused on increases in server-processing speeds and capability. In the past year, however, the trade shows have shifted their focus onto power efficiency—to the extent that an EPA director gave a recent keynote speech. “Green” technology may be a nice-to-have feature in many parts of the US economy. For the data-center sector, however, cost savings due to energy efficiency has become a dominant force, affecting hardware design and even placement of the data center itself. For example, one of the key reasons for Google's decision to start a data center in The Dalles, OR, was the community's proximity to cheap hydroelectric power.

The power-supply community has for several years been discussing a move to

### AT A GLANCE

- ▣ Data-center efficiency now consumes 1.5% of the total US electrical power.
- ▣ Data-center operators are willing to pay a premium for power-efficient hardware.
- ▣ Although it holds promise, dc-power distribution requires a major change in the power infrastructure.

digital power. In particular, digital control of the PWM (pulse-width-modulated) control loop has struggled to catch on in the power-supply industry, especially the “silver-box” ac/dc-power supply that powers every rack of servers in a data center, because of the price sensitivity of these supplies. Even 50 cents' worth of additional parts is enough to put a silver-box supply out of the competition, and digital-power controllers can add a couple of dollars. In addition, an analog-control loop tuned for a specific load can perform as well as a digital-control loop. Combine greater cost, brand-new technology that's alien to analog designers, and no clear technological advantage; it's no wonder that

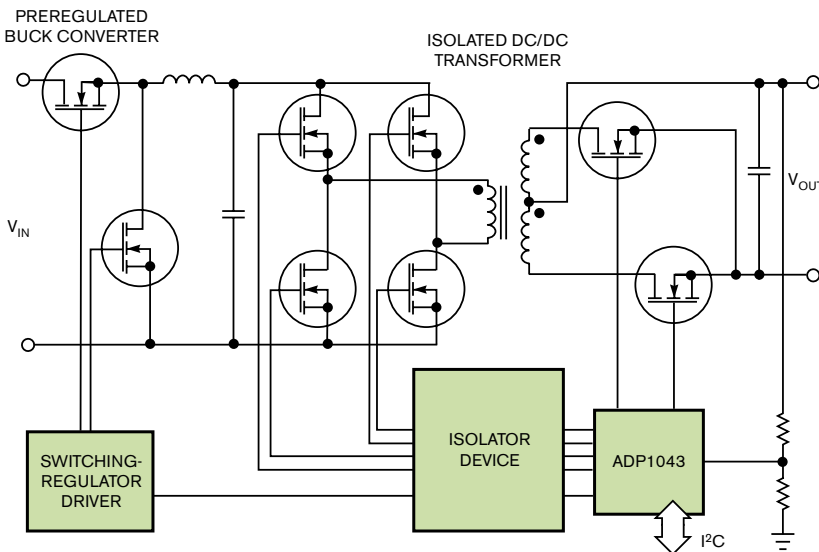
digital power has not taken off in the silver-box-supply market.

However, the quest for improved efficiency will drive the adoption of digital power in server applications because the power supply must be able to maintain a high efficiency over not just one load, but a range of loads. Analog loops can't provide a high efficiency over loads that can range from 20 to 100% of full load.

Servers must respond over such a widely varying load because the vital nature of a data center's workload requires redundant power supplies to serve as backups for each rack of servers. Both of the redundant supplies operate at 50% of the load under normal conditions, with the load varying from 100 to 10% or less depending on the workload and power-supply status. Achieving efficiency over such a wide load swing is difficult with an analog-PWM-loop supply, but a digitally closed loop can tune itself for changes in the load and make efficient operation possible over the load range.

Just a few years ago, the cost savings due to increased power efficiency wouldn't have been enough to offset the higher cost for parts. Now, IC vendors are eyeing the greater complexity as a way to justify higher IC costs. For example, Analog Devices recently introduced its first digital-power-controller IC and pointed out that it supports a complicated design such as that in **Figure 1** (**Reference 3**). The circuit increases a converter's efficiency by as much as 1%—enough to get the attention of power-supply designers. Analog Devices officials hope that the ability to support the new two-stage design will be a plus for the chip.

Digital power also confers the ability to communicate with individual supplies both within and among server racks. Power-supply-system management requires intelligence within the power supply to self-monitor heat, operating time, and load response. Digital-power ICs offer a range of power-management options, including a host-controlled system, such as a PMBus-(power-management-bus)-based controller or Power-One's Z-One digital-power-management system. These power-communication buses require a separate intelligent host that can query supplies about their operating time, heat, and fan response, and they can also sequence supplies down



**Figure 1** Digital-power controllers, such as Analog Devices' ADP1043, can simplify the incorporation of complex new-circuit designs within server power supplies. This two-stage dc/dc-controller design must drive and control seven FETs but can result in a power-supply efficiency increase of 1%. The design takes the input voltage from the power-factor-correction output bulk capacitor and buck-converts it down to a lower regulated voltage. The isolation stage now runs as a high-efficiency dc/dc converter running in an open loop.

or up. The controllers can also store information about the power system to schedule maintenance before a failure. Not all systems require such elaborate control. For example, some digital-loop power controllers can also communicate between controllers in a simpler no-host-control scheme that may provide enough control for some systems. For example, Zilker Labs' ZL2006 and 2004 controllers allow communication between chips without a host or a communication bus.

In addition to increasing power-supply efficiency, data centers can optimize system power efficiency through software virtualization of hardware servers. With virtualization, a server can support several jobs at once, rather than dedicating one physical server to a job regardless of whether the job requires the server's entire processing bandwidth. In effect, virtualization subdivides the physical server into multiple virtual machines until no excess processing capacity remains on the server. With the advent of server virtualization, a system host requires no dedicated server: If a hosting system requires only some of a hardware server's bandwidth, the rest of the bandwidth can host one or more separate systems, depending on the hardware's capacity and the bandwidth needs of the hosting system.

According to Kevin Bross, modular-systems architect at Intel, a data center responding to a recent virtualization survey reported that virtualization had allowed a 10-to-1 reduction in the number of servers and a 2-to-1 reduction in the total power usage because the remaining servers were receiving heavier usage.

Although using digital power and virtualization are two relatively straightforward power-reduction tools, the third

## FOR MORE INFORMATION

**Analog Devices**  
[www.analog.com](http://www.analog.com)

**Climate Savers**  
[www.climatesavers.computing.org](http://www.climatesavers.computing.org)

**Eltek Valere**  
[www.eltekvalere.com](http://www.eltekvalere.com)

**Energy Star**  
[www.energystar.gov](http://www.energystar.gov)

**Green Grid**  
[www.thegreengrid.org](http://www.thegreengrid.org)

**Intel**  
[www.intel.com](http://www.intel.com)

**Power-One**  
[www.power-one.com](http://www.power-one.com)

**Texas Instruments**  
[www.ti.com](http://www.ti.com)

**Validus**  
[www.validusdc.com](http://www.validusdc.com)

**Vicor**  
[www.vicr.com](http://www.vicr.com)

**Zilker Labs**  
[www.zilkerlabs.com](http://www.zilkerlabs.com)

tool, using a dc-power-distribution bus, requires changing a basic part of the data-center infrastructure. Current methods involve twice converting data-center power to dc: This approach converts power as it enters the center to power the UPS (uninterruptible power supply) system and then converts the power back to ac for distribution through the building to the rows of server racks. This ac-to-dc-to-ac conversion shuffle causes the loss of approximately 20% of a data center's power.

One of the first companies to announce an infrastructure-level dc-voltage plan is Validus DC Systems. The Validus system's basic architecture calls for the conversion of ac power to -575V dc where it enters the building at a power-quality-module air-cooled rectifier. The -575V dc routes through the facility through a power-distribution module, a scalable "switchboard" that integrates power distribution from multiple system inputs, such as the utility, a UPS, or battery banks. Also at this point, alternative-energy sources, such as solar- or wind-generated power, can enter the power-distribution scheme. The power switchboards integrate and distribute the high-voltage dc to the power-converter units for each row of server racks and convert -575V-dc to -54 V-dc power, distributing as much as 120 kW to the server row.

Ron Croce, chief operating officer at Validus, says that the company has received an increasing number of inquiries from data-center operators interested in power distribution for alternative energy, which is often dc. Validus' switchboard approach lends itself to paralleling the output of photovoltaic systems to boost the dc power to usable distribution levels, he says. In addition, few data-center operators would want to rely

### MORE AT EDN.COM

⊕ For more on digital-power control see, [www.edn.com/article/CA6335298](http://www.edn.com/article/CA6335298).

⊕ For more on the Energy Star server-efficiency draft, go to [www.edn.com/blog/1470000147/post/1340021934.html](http://www.edn.com/blog/1470000147/post/1340021934.html).

⊕ Go to [www.edn.com/080320df](http://www.edn.com/080320df) and click on Feedback Loop to post a comment on this article.

at least initially on a purely alternative-energy source. Validus' approach allows the integration of several power sources, allowing a data center to use, for example, dc-photovoltaic power during the day and switch to ac-grid power at night. Over a five-year period, the total cost of ownership, including equipment, installation, maintenance, and efficiency saving, yields a 50% savings over a conventional ac approach, according to Croce.

Intel's Bross points out that using a dc-voltage bus for power distribution is tried-and-true method within the telecom industry, in which  $-48\text{V}$  dc is the norm. Pragmatically, Bross suggests that, rather than getting into unending discussions about the practicality of dc power within the data-center industry, engineers can make extensive evaluations with currently available telecom equipment. He goes further than just pointing out dc power's practicality, however. He also offers a sophisticated online calculator that allows facility engineers to perform side-by-side evaluations of the efficiency of ac-power-system power versus that of  $-48\text{V}$ -dc power systems. The calculator does not rely on generalized assumptions about the efficiency of a UPS or the gauge of a distribution cable. It instead allows you to plug in your own numbers for every piece of the complicated power-distribution and -conversion puzzle. Check out the ac-versus-dc-power calculator for yourself, and you'll see the difference dc-power distribution makes: It's a winner (**Reference 4**). For example, Eltek Valere's Flatpack2 HE 48/2000 rectifier module offers 96% typical power-conversion efficiency, which the company claims is the highest efficiency level available. Prices begin at \$450.

Steve Oliver, vice president of marketing and sales for Vicor's VI Chip, agrees that  $-48\text{V}$  dc holds promise for energy-efficient power routing and points out that a move to dc power also brings up the opportunity to rethink power distribution within server racks and servers. Oliver says that, when converting from an intermediate 12 or 9.6V rail to the load, designers often use a synchronous-buck converter. With this topology's duty cycle, the higher the input voltage relative to the output voltage, the worse the circuit's efficiency.

For example, a synchronous-buck-rectifier circuit that drops 12V down to 6V has an efficient 1-to-1 duty cycle. But a circuit that drops 12V down to 1V requires a less efficient 1-to-12 duty cycle; it also stresses the switching and control FETs. Oliver sees a move away from synchronous-buck converters and toward circuits and devices that can efficiently make the drop from 48V to less than 1V with little or no loss in efficiency. Vicor's VI chips can make the drop from 48V to, for example, 0.8V in one stage, he claims.

Although servers do need higher voltage rails to power devices such as disk drives, 90% of a large system's load powers the processor with 0.9V and the memory stack with approximately 1.2V. Oliver suggests a hybrid power architecture that includes a main power unit on each server going from 48V down to the low processor and memory voltages and a much smaller supply handling vestigial higher voltage rails. **EDN**

## REFERENCES

- 1 "Report to Congress on Server and Data Center Energy," US Environmental Protection Agency, Energy Star Program, [www.energystar.gov/ia/partners/prod\\_development/downloads/EPA\\_Datacenter\\_Report\\_Congress\\_Final1.pdf](http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf).
- 2 "Energy Star Program Requirements for Computer Servers, Draft 1," [www.energystar.gov/ia/partners/prod\\_development/new\\_specs/downloads/Draft1\\_Server\\_Spec.pdf](http://www.energystar.gov/ia/partners/prod_development/new_specs/downloads/Draft1_Server_Spec.pdf).
- 3 Jang, Yungtaek, Milan M Jovanovic, and Yuri Panov, "Multi-Phase Buck Converters with Extended Duty Cycle," Delta Products Corp, 2006, [www.deltartp.com/dpel/dpelconferencepapers/2.4\\_10348.pdf](http://www.deltartp.com/dpel/dpelconferencepapers/2.4_10348.pdf).
- 4 "A Facility-Level Analysis of Efficiency between AC and  $-48\text{VDC}$  Servers," Intel, 2007 [http://innovate.intel.com/facility\\_operating\\_expenses\\_calculator\\_energy.html](http://innovate.intel.com/facility_operating_expenses_calculator_energy.html).

You can reach  
Technical Editor  
**Margery Conner**  
at 1-805-461-8242  
and [mconner@connerbase.com](mailto:mconner@connerbase.com).

