

More Data in the Data Center

The true challenge is to increase MIPS without burning up

BY INTERNATIONAL RECTIFIER

Exponential growth in e-commerce, Internet communication, and entertainment has led to an ever-increasing demand for more servers and data storage. The move of these communication paths from a real-time media environment to a demand-driven, IP packet-based environment creates a tremendous opportunity for energy-efficient infrastructure equipment for the data center. As central offices give way to network switches, commerce becomes server-driven, and media programming becomes packetized, it becomes necessary to re-think power management strategies. While many companies seek to be “green” and reduce the environmental impact of electricity use, sufficient space and power or cooling capacity required to expand their businesses must also be developed.

Data centers accounted for approximately 1.5 percent of total US electricity consumption in 2006 — more than 60 billion kilowatt-hours (kWh) (see figure 1) — at a cost of \$4.5 billion, while the cost to power and cool servers worldwide surpassed \$26 billion.

With this sum growing at four times the rate of server spending, information technology professionals must consider the total cost of ownership of the data center in their purchasing decisions of infrastructure equipment, such as servers, storage area networks, routers, and switches. The cost to operate this equipment over its lifetime is now about three times larger than the original hardware purchase. As a result, aggressive targets to cut energy consumption have been set by the leading technology solutions providers.

The traditional method of measuring the performance of the data center by throughput (MIPS: millions of instructions per second) or performance density (MIPS/square

foot) alone is no longer sufficient. The key metric moving forward to measure efficiency is based on performance efficiency, or MIPS per watt (W). Leading data center equipment providers have set aggressive goals to increase MIPS/W by a factor of ten by 2010, while reducing overall power consumption by 20 percent.

This new figure-of-merit (FOM) presents its own challenge as increasing MIPS has traditionally led to a proportionate increase in dissipated power. Successfully developing and executing strategies to decouple MIPS from watts will allow more processing power to be crammed into smaller enclosures, thus reducing electricity bills and building infrastructure costs.

There are two main sources of power consumption in the

Rising energy use

Electricity use for data storage, ranging from small office servers to large data centers, more than doubled from 2000 to 2006, and is expected to double again in the next five years.

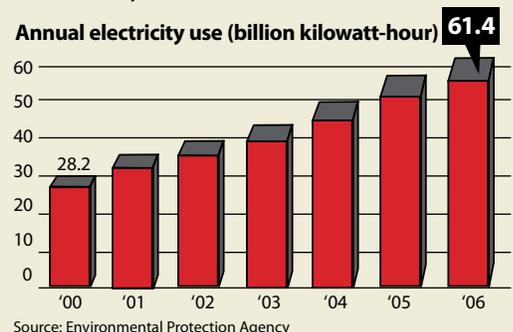


Figure 1: Data storage electricity consumption

data center. Servers comprise the first of these sources, and the infrastructure required to cool and protect the second. The energy usage from each is about equal, and they are directly related. Therefore, for every dollar saved in server energy consumption, an additional dollar can be saved in infrastructure energy costs.

Server energy consumption consists of three main elements. The electronic loads such as microprocessors and memory banks consume 60 to 70 percent of energy, while power supplies consume 25 to 30 percent, and cooling uses five percent. While there is significant advancement in reducing the load's power profile, by, for example, the introduction of multi-core efficient processors and virtualization technology, much can be gained by adopting a holistic approach to system power management in order to significantly reduce the consumption of all three of these major energy consumers.

A Holistic Approach

New smart power management system solutions (see figure 2) include the co-design of several critical components of the power supplies that are integrated into the platform. The key elements of the power system are highly efficient and dense power stages, advanced highly responsive power controllers, digital interfaces for programmability and diagnostics, accurate power monitors, system controllers, and sequencing.

Advanced Power Stages

Advanced power stages reduce power loss in power supplies by up to one third compared to traditional designs. New products take advantage of industry-leading MOSFET technology that achieves significantly higher power density and advanced packaging that exhibits nearly zero package resistance and inductance with lowest industry thermal impedance (see figure 3). Compared to standard plastic discrete packages,

As central offices give way to network switches, commerce becomes server-driven.

the metal can construction of these benchmark MOSFETs enables dual-sided cooling to effectively double the current handling capacity of high frequency dc-dc buck converters. This dramatically cuts energy losses while shrinking the design footprint of the circuit board. Optimized driver ICs co-designed with these MOSFETs deliver benchmark efficiency over heavy and light-loads.

By using efficient power devices coupled with innovative control schemes, it is possible to obtain an optimal combination of efficiency and electrical performance. Combined, these technology improvements can improve



Figure 2: Integrated power management platform

server efficiency by about five to six percent with increased density over time.

Advanced Power Systems

Advanced power systems can reduce power dissipation even further. High power loads in the system, such as microprocessors and memory banks, have a very unpredictable power profile due to rapid changes in their required performance and function. Under severe requirements, these loads can exceed their thermal limits and require thermal throttling, or stepping back in performance, to allow the silicon and the package to cool. Once sufficiently cooled, the load must crank back up again. This creates a highly inefficient stop/start, or thermal and power cycle.

Permitting high-performance silicon to thermal and power cycle is both an enormous waste of energy and performance. Employing a holistic approach to system power management eliminates thermal throttling. By dynamically monitoring instantaneous power, recording its trends over time, and understanding the thermal impedance of the load, the power system can accurately predict thermals in the system at any point. With this information, the power system can then alter the load's electrical characteristics (like dynamically changing its core voltage or reducing the digital silicon's clock speed) to limit its power and establish the correct cooling conditions in advance through energy-efficient advanced variable speed motion control, to guarantee that the load never leaves its required thermal envelope, optimizing its throughput, and hence its performance. This can eliminate up to 15 to 20 percent of server total power dissipation.

Application Example

The modularity, low price, and size of blade servers makes them an ideal way to expand capacity as needed within the already crowded confines of the data center. So the trend is to add racks containing high-density blades.

The problem is the large amount of heat that such racks

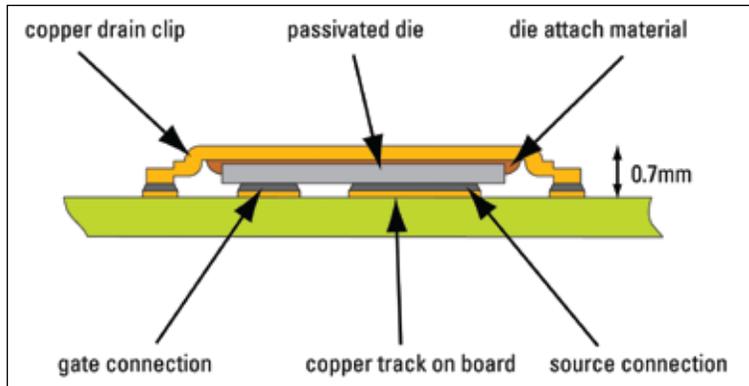


Figure 3: Advanced power MOSFET packaging technology

generate. The latest blades may have up to four processors per board, and their power requirements are significant. This creates thermal problems that limit the ability of users to take full advantage of the computing density blades offer. In practice, data centers often leave slots empty to provide more cooling and keep systems within thermal specifications. A mainstream dual processor blade can consume between 600 and 1000 W on its own, while the data center will dissipate an equal amount in infrastructure and cooling losses. Assuming total data center power consumption for each blade at 1.6 kW equates to over 14,000 kWh and almost \$1300 per blade per year in operating costs.

These two methods can help reduce the total power consumed in the blade, reducing the cooling requirements, and thus allowing greater blade density in the rack. The first method is to employ highly efficient on-board power supplies. Approximately 80 percent of the power drawn by the blade is consumed through the on-board power supplies, so the efficiency of the power supplies has a large impact on the system efficiency. Much of this power is consumed through the microprocessors and memory. For example, a typical high-

performance microprocessor will operate at 130 amps at 1.1 volts, or 146 W. Today it is typical for on-board power supplies to have about 80 percent efficiency, or 20 percent losses.

With advanced power control and conversion technology, such as International Rectifier's XPhase scalable multi-phase architecture and DirectFET MOSFETs, it is possible to increase system efficiency to over 88 percent. By reducing the power supply's power loss by 40 percent (from 20 percent losses to 12 percent losses) about 900 kWh and \$82 per year can be saved.

Accurate real-time power monitoring ICs from International Rectifier can be part of an advanced power system that reduces dynamic power loss in the blade, an additional 2100 kWh and \$191 per year can be saved.

In total, by employing these advanced power management techniques, 3000 kWh and \$273 per year can be saved (over 21 percent). The cost of employing these technologies is significantly less than the savings generated. These approaches also offer secondary benefits. Board level fans can run slower, thus saving more energy, as there is less heat to distribute and acoustic noise can also be reduced.

Summary

Over the next three years, up to 25 percent of server and full data center heat load can be saved with the adoption of optimized power management systems that incorporate advanced power stages, accurate and dynamic power monitoring, and high-performance power controllers. This new approach will allow data center equipment vendors to meet their performance objectives on time and enable a new generation of high-performance and cost effective products aimed at servicing the insatiable performance needs of the ever-changing media environment. ■