LINEAGE POWER®

The Efficiency Challenge

Balancing Total Cost of Ownership with Real Estate Demands

written by:

Mike Wagner
VP & GM Embedded Power AC-DC

The Efficiency Challenge

Balancing Total Cost of Ownership with Real Estate Demands

With the digital expansion over the past few decades, new requirements for data centers to store, transmit, and manage data of all sorts are reaching a critical level. Companies are feeling the crunch as data centers are filling up quickly with IT equipment, real estate remains at a premium, and power demands and energy prices are skyrocketing. Two primary issues are confronting data centers today. First is the total cost of ownership — which consists of the IT equipment as well as the data center infrastructure used to support operations. The infrastructure includes the power distribution, cooling, cabling, and several other components. Second is the fact that data center real estate is at a premium. Each inch of space within a rack enclosure needs to perform at maximum efficiency vis-à-vis computing performance, power, and cooling.

In this article, we will discuss the path to improving efficiency as it applies to AC-DC power supplies (AC-DC front ends) within server equipment. Although the front end is only a portion of the data center power train, the AC-DC front end power supply remains a key element in the efficiency challenge.

As a note to the reader, there is a significant amount of activity centered on alternative architectures (in an effort to reduce power train losses and improve overall efficiency) such as 380 VDC or 48 DC architectures. One should be aware that there are always tradeoffs to any approach. In the case of a high-voltage bus such as the 380 VDC, there is a lack of building and maintenance code standards nor connectors or circuit breakers in place today for such a system. For the purpose of this article, we will discuss the existing AC to 12 volt architecture at the server level and what can be done from a power-supply perspective to improve its efficiency performance — resulting in better overall system level efficiency.

Increasing Server Density

The number of servers per rack enclosure is increasing immensely and to help with the lack of real estate in data centers, servers are becoming smaller with increased processing capabilities. These servers are also becoming more power dense. In 1996, the average rack enclosure housed 7 servers. This number is expected to reach 20 servers per rack by 2010, according to research analyst firm IDC. New and high-density blade servers are gaining ground quickly, and as data center power density is increasing by approximately 15 percent per year, blades will be a significant portion of all servers installed (see Figure 1). At the same time, because blade servers pack so much computing power into a small space, the power and cooling challenges intensify. Therefore, on one hand, these dense IT solutions are helping alleviate data center real estate issues, yet simultaneously, they are also contributing to the power and cooling challenges.

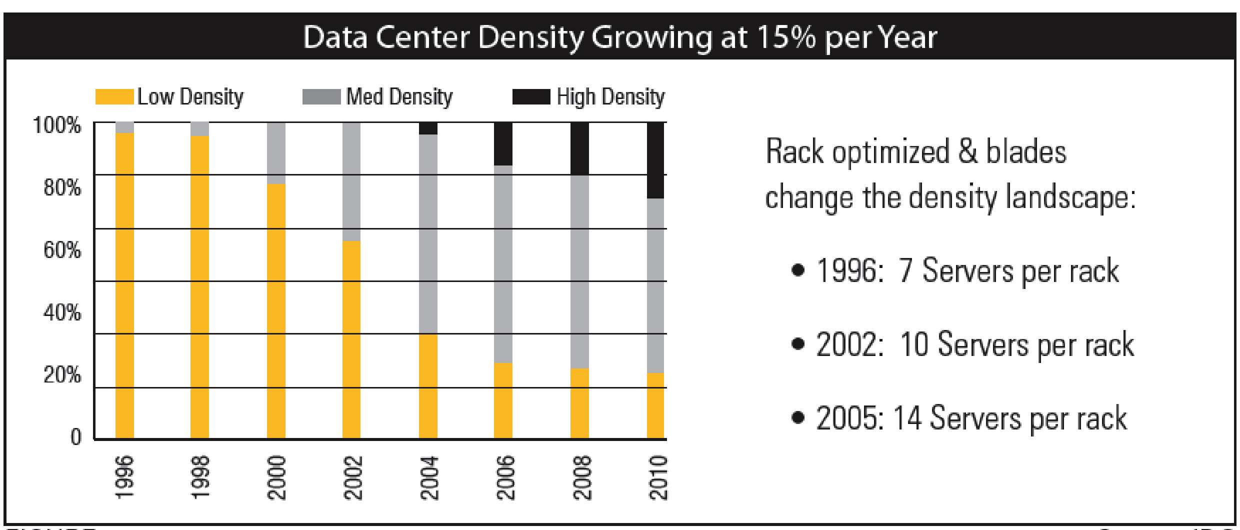

FIGURE 1 Source: IDC

Power and Cooling Costs Skyrocket

As mentioned, the change in server and power density helps with the data center issue of space and data-processing capabilities, but it also contributes to rising costs in several areas such as power and cooling (see Figure 2).

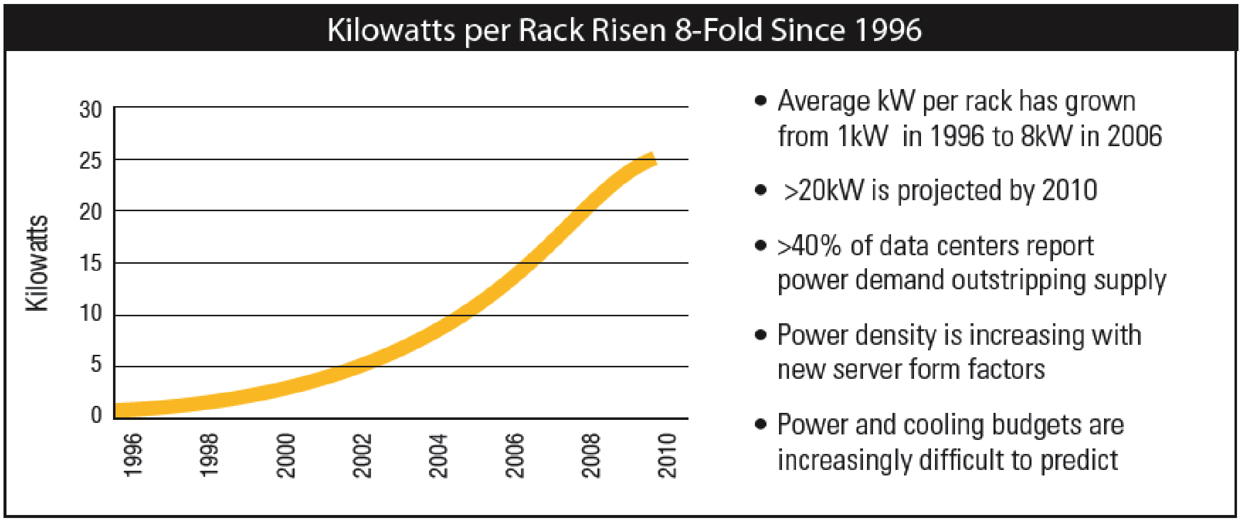

FIGURE 2 Source: IDC

Conservative figures estimate that energy costs exceed over one-half of data center maintenance costs. "The real cost, however, is not just in the power being used but in the costs of the infrastructure equipment - generators, UPSs, PDUs, cabling and cooling systems. For the highest level of redundancy and reliability - a Tier 4 data center - for every kilowatt used for processing, some $22,000 is spent on power and cooling infrastructure" (Robb, 2007).

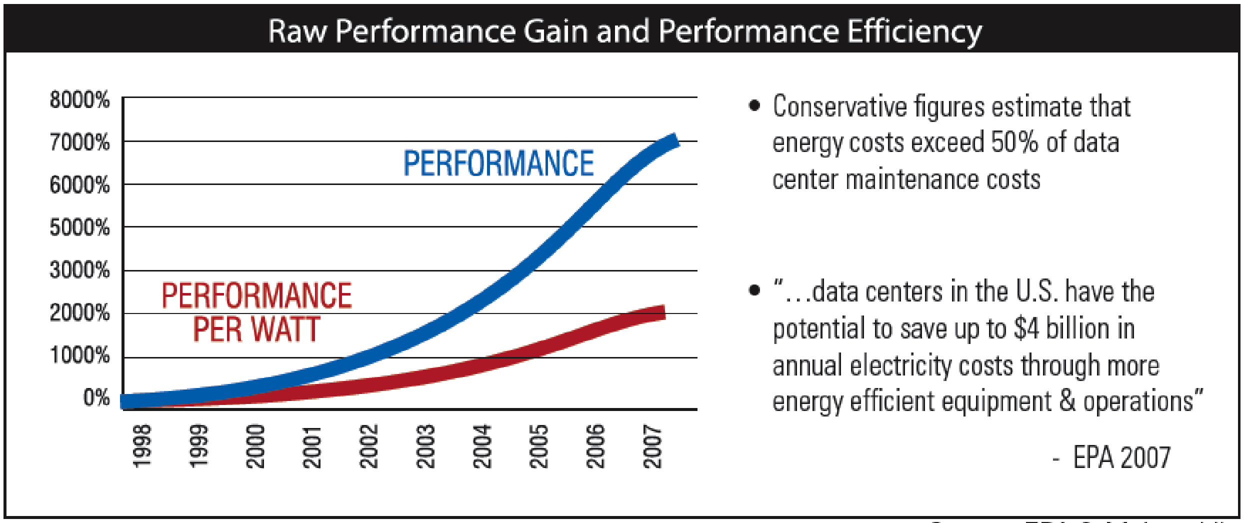

FIGURE 3 Source: EPA & McLaughlin

Smaller and more powerful has been the mantra for design engineers, but while the raw server computing performance has increased significantly, the performance per watt has been slower (see Figure 3) (McLaughlin, 2007b). Most of the industry data describes lack-luster efficiencies for multiple-output, PC-grade power supplies. For distributed power architectures which implement front ends and rectifiers, the power supplies exhibit much better performance gains. Nonetheless, the theme is that the demand for power and cooling is at an all-time high. Experts agree that what needs to be done - to decrease our carbon footprint and reduce our energy costs - is to improve data-center efficiency. Industry efforts such as The Green Grid (www.thegreengrid.com) and new EPA initiatives are ringing the bell for change.

Data Center Power and Consumption

Figure 4 represents the current power train in data centers:

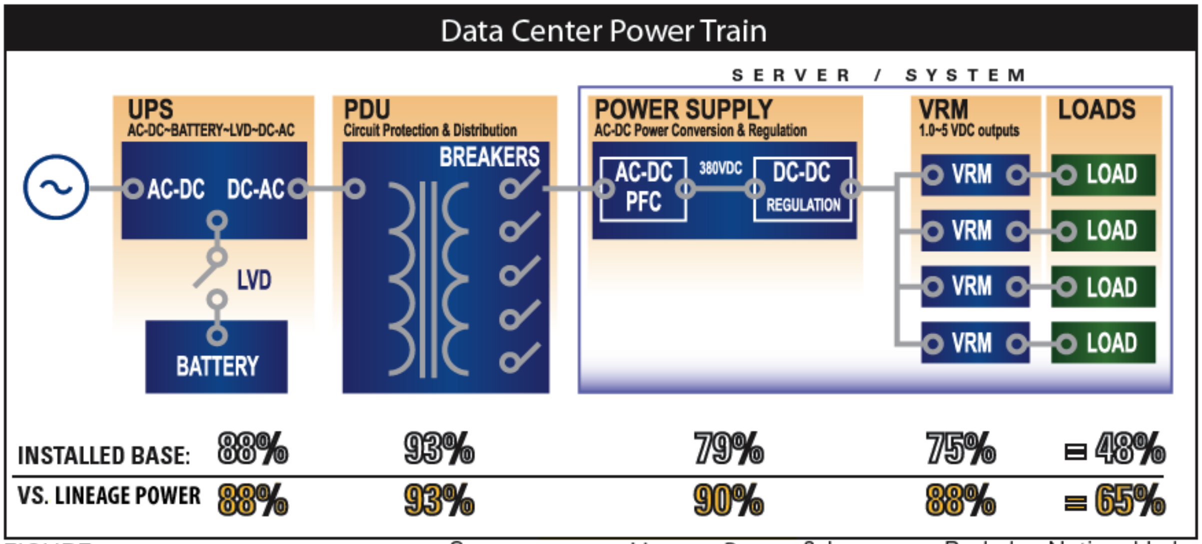

FIGURE 4 Source: Lineage Power & Lawrence Berkeley National Labs

This AC power distribution architecture is the most common architecture, provides the lowest risk in the supply base of IT equipment, and has an aggressive cost base for equipment as well. This power architecture also has some associated weaknesses which include low system efficiency, low power densities, and a higher number of components affecting mean time between failures.

The following are the efficiency ranges for the important elements of the typical AC power train as seen in "existing" data centers (Lawrence Berkeley National Labs, 2007 and Lineage Power):

- UPSs – 85-92%
- Power Distribution – 98-99%
- Power Supply – 80% (Note 1)
- DC-DC & VRM – 78-85%

There are numerous studies with varying efficiency ranges(Note 2), yet no matter which statistics are used, it is clear that these figures leave ample room for improvement in several power train areas - particularly the switch-mode power supply area. Lineage Power, being an international manufacturer of power supplies, is a key contributor with advanced technologies that will greatly improve performance efficiencies as illustrated later in this paper.

Current Power Solutions

As the industry has matured, so has the power supply industry's understanding of IT equipment's actual performance models. For many years now, original equipment manufacturers have requested power supplies that demonstrate their optimum efficiency at or near full load. What followed was a power supply efficiency curve optimized for peak efficiencies at 75 - 85% of full load, while at the same time sacrificing light load performance — as opposed to today's computing demands which is in fact requiring "light load" efficiency performance (See Figure 5).

Over time, what the industry has demonstrated is that servers typically run at a mere 20% load because they are operating in redundancy. For example, servers which run redundantly and in parallel, bring the maximum load down to 50%. And of that 50% load, more than half is often left under utilized.

Conversely, power supplies in the installed base were designed for optimal efficiency under full load. Per Lawrence Berkeley National Labs study, typical power supplies that offer 79% efficiency (maximum), would yield only 63% efficiency under experienced 20% load conditions (typical).

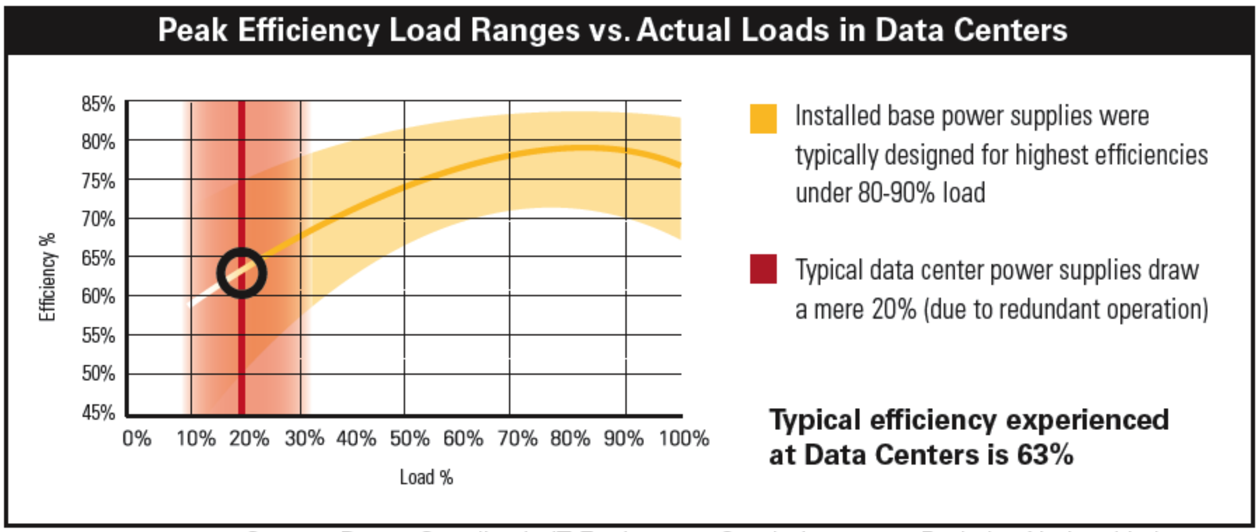

FIGURE 5 Source: Power Supplies in IT Equipment, Graph, Lawrence Berkeley National Labs, 2007

Maximizing Power Efficiency and CAR1212 / CAR2512 Overview

Given these data center efficiency challenges, Lineage Power sees its role for improving the dissipation losses and performance efficiencies of power supply units. Over the years we have found that there are certain key elements that lead to inefficiencies. Some of these inefficiencies lend themselves to improvements via creative thinking and innovative circuitry—while others are highly dependent upon component and materials advancement subject to market availability. The technical challenge in line with server-industry trends is twofold. One objective is to increase the overall efficiency of the front end while the other is targeted at specifically improving mid-to light-load performance efficiency given typical server load operating levels.

Efficiency Challenges Within the Power Supply

The most productive way to analyze efficiency improvement within the front end is to examine the dissipation losses by mapping areas slated for improvement. In a general sense, the power-train can be broken into sections including the PFC stage, bias and fan control, primary DC-DC topology/stage and finally the output section. The efficiency of these stages is multiplicative. Hence, concentrating on any of the areas mentioned above will have a direct impact on the overall front end efficiency. The areas with the largest power dissipation will of course have the most impact on the entire efficiency equation.

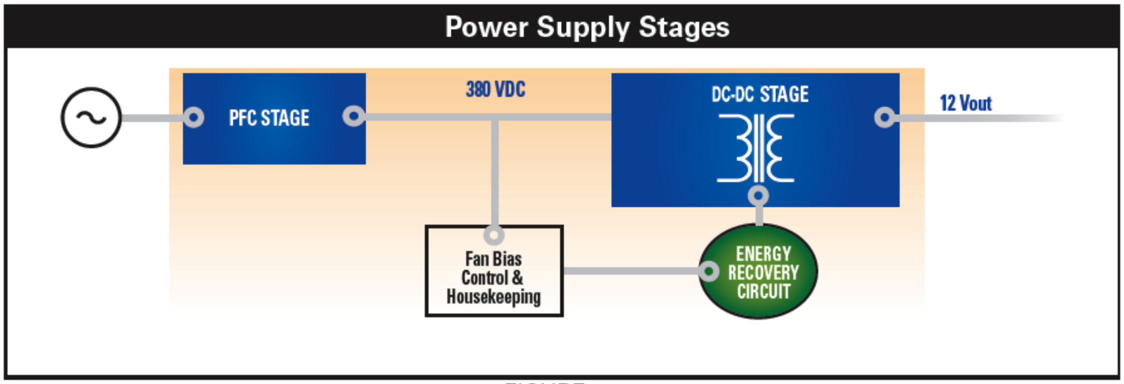

FIGURE 6

Today at Lineage Power, the PFC stage is already running at 97.5% and is not considered "low-hanging fruit" relative to other areas. The biggest impact will come from improvements made to the DC-DC step-down stage along with the bias and fan control circuitry. Today, on our 92% efficient front end, the DC-DC stage accounts for an efficiency of 94%. In order to improve upon this performance, it is required to categorize the losses by type so that each one can be addressed independently. Much of our successful effort to get us to an overall front end efficiency of 92% has been focused on switching losses. This we addressed through the use of zero voltage switching (ZVS) techniques using a full bridge topology which serves as the best platform to reduce losses during the switching transitions of the MOSFETs. Lineage Power also employed a patent-pending circuit, called a lossless snubber, which acts to recover the energy associated with reverse current of the output synchronous rectifiers normally dissipated as heat. Approximately 80% of this energy is recovered and reused back in the circuit as recycled energy.

Coupling our approach above with a technique called synchronous rectification (replacing Schottky diodes with low RDson MOSFETs) we maximized our performance in this area. In both cases of primary and secondary side, we used best in class MOSFETs with very low RDson resistance enabling us to minimize the conduction losses even further. In addition, optimizing the DC-DC transformer turns ratio reduced the primary current and secondary voltage, minimizing losses and improving efficiency performance.

The bottom line to these techniques is a 12 volt front end with 92% efficiency at a high power density of 27W3 as portrayed in our CAR2512 series.

Further improvement in this area will depend upon the availability of new components and materials that are not available for purchase on the market today. To get a jump start on leading edge technology, Lineage Power has partnered with strategic suppliers so that new design challenges are shared as early as possible in the development process. Exchange of technology and product road maps is on-going and will lead to new technological improvements.

Increasing Efficiencies Under Light Loads at the Data Center

Having addressed the first objective of increasing the overall efficiency of the front end, the second area is to address mid- to light-load performance efficiencies. As mentioned earlier, servers tend to draw power at levels much lower than previously anticipated. Servers normally operate in redundant mode, sitting idle until an application drives demand for power. The average server draws between 10-30% of the available power from the front end just to support its functions. Consequently, the performance of a typical front end, which is optimized for full load, is relatively poor. Lineage Power has recognized this fact and addressed this need by providing technology which targets typical operating ranges and at the same time supports full power reliably.

There are several approaches that can be taken to address the load balancing issue within the front end. These include:

- Lowering bias/fan supply power draw levels
- Optimizing the power supply hold up time versus transformer turns ratio
- Reducing PFC switching frequency as the load decreases
- Using improved material (MPP) for the main output choke to reduce core losses at light load
- Balancing RDson (primary MOSFET resistance) versus capacitance for optimal total loss solution
- Optimizing Synch FETs when load is below 30% to conserve bias and housekeeping power

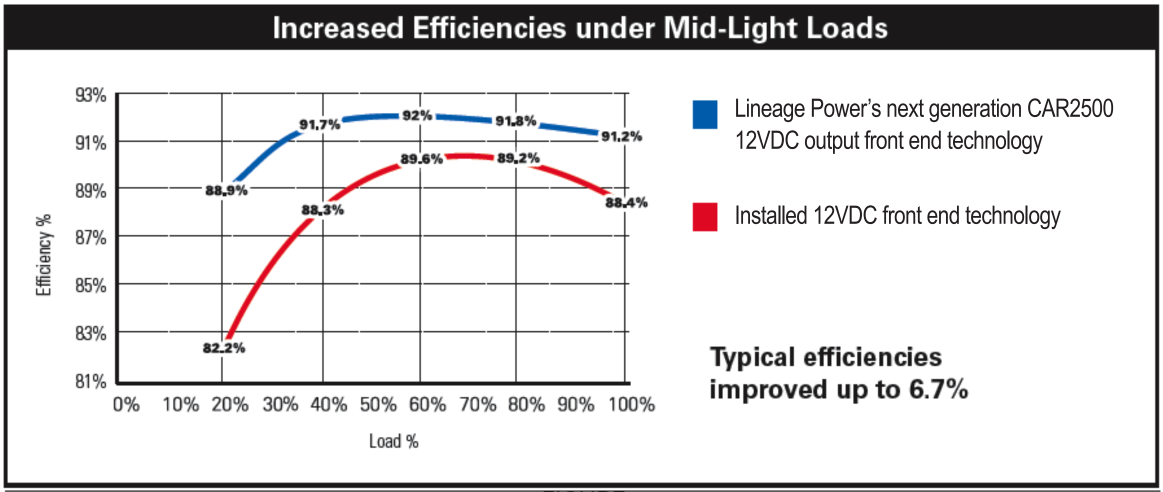

FIGURE 7

Total Cost OF OWNERSHIP REDUCTION

To calculate total cost of ownership (TCO) savings via power supply improvements, we can use data from a current 80 Plus® Server Project as a baseline for typical efficiencies in existing data centers.

The 80 PLUS® Server Project is furthering a Silver (75-87 percent efficiency) and Gold Standard (77-90 percent efficiency) for power supply efficiencies. At these levels, 80 PLUS calculates the following server savings over the lifetime of server equipment, which is estimated at 4-6 years, depending on type of server.

Using this data, one can demonstrate how to predict the following standards on the data center TCO in relationship to server equipment. To calculate TCO, these alternative scenarios provide a representative overview.

Server & Data Efficiency Savings

	Server Class	Lifetime Energy Savings per Server (kWh)
Silver Standard	Volume	480 - 1290
	Mid-Range	1932 - 3366
	High-End	21716 - 36000
Gold Standard	Volume	740 - 1800
	Mid-Range	2680 - 4710
	High-End	30336 - 50742

* Energy cost are 8.8 cents per kwH (per EPA)
* Data Center PSU Loaded at 20% of rated power

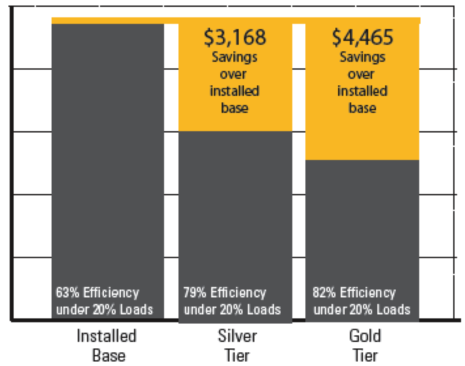

Life Time Savings of one High End Server (6 years)

FIGURE 8

Source: 80 Plus® Organization & EPA

The EPA reported an average commercial electricity rate of 8.8 cents/kWh in the US for 2006. Using these rates, one can see that the annual savings from the installed base efficiencies (as reported by 80 PLUS) to either the Silver or Gold Standard are substantial. Over the lifetime of a single high-end server (6 years), the Silver Standard leads to $3,168 in savings (36,000kWh * $.088 = $3168) and the Gold Standard delivers $4,465 (50,742kWh x $.088 = $4,465). Estimated savings will be even greater utilizing Lineage Power's next generation CAR2500/1200 designs.

With enterprise-class data centers housing hundreds to thousands of servers (EPA, 2007), the reduction of the TCO translates to large sums of money. For 1000 servers, the Silver Standard reduces costs by $3,168,000 and the Gold Standard by $4,465,000.

Summary

Based on the data that we have analyzed, it is clear that inherent inefficiencies are currently built into each component of the data center AC power train. These inefficiencies lead to greater server power consumption and recurring electricity costs that substantially outweigh the original purchase price of the IT equipment. Furthermore, data shows that such inefficiencies are adversely affecting the overall power usage of the data center because of the need for cooling equipment, thus creating an environment with some of the lowest power efficiency levels and some of the highest carbon footprints of any commercial facility.

Our analysis shows, that in terms of efficiency, an area for improvement within the power train is the AC-DC front end. Lineage Power has concentrated on these products to increase power densities and true operational efficiencies. This approach allows data centers to leverage existing power architectures while experiencing substantial reduction in their TCO through reduced power consumption, physical size and cooling costs. Lineage Power's next generation CAR2512 Technology has been used as an example to demonstrate the significant gains that can be achieved in a data center by addressing just one component of the AC power train.

NOTES

Note 1: Lawrence Berkeley National Labs (LBNL) shows a range of 68-72 percent for power supply unit efficiencies, and it is important to note here that these statistics — as reported by LBNL — reflect multiple-output, PC-grade power supplies rather than previous generation higher-efficiency front ends. Lineage Power's experience shows that efficiency levels are much higher at 80% or greater.

Note 2: From Lineage Power's professional experience, we believe the PDU efficiency number might be too high and the power supply number and DC-DC number too low. Therefore, we have included what we see as more representative numbers corresponding to today's information alongside the LBNL numbers.

References

- 80Plus.org 2007. 80Plus Server Project. http://www.80plus.org
- EfficientPowerSupplies.org. 2007. Power Supplies: A Hidden Opportunity for Energy Savings. Natural Resource Defense Council (NRDC) report prepared by Ecos Consulting, May 22, 2002 (also available from EPA Energy Star Website) http://www.efficientpowersupplies.org/efficiency_opportunities.html
- EPA. August 2, 2007. Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431. U.S. Environmental Protection Agency ENERGY STAR Program. http://yosemite.epa.gov/opa/admpress.nsf/0/4be8c9799fbceb028525732c0053e1d5?OpenDocument
- Hewlett-Packard Company. August 2007. Budget going up in a puff of hot air. http://h71028.www7.hp.com/ERC/cache/49205-0-0-0-121.aspx?bodycontentparams=437510-0-0-0-121&ERL=true
- Intel - Griffith, Brian. September 2004. Enabling High Efficiency Power Supplies for Servers: Update on industry, government and utility initiatives. http://hightech.lbl.gov/events.html
- Kiosk, William. March 2007. Taking the Long View: A systematic approach to data center designs. Member of the Facilities Panel for the US Defense Science Board Energy Strategy Task Force. Energy & Power Management magazine.
- Lawrence Berkeley National Labs – Tschudi, William. 2007. Data Center Power Distribution. http://hightech.lbl.gov/events.html
- Lawrence Berkeley National Labs – Tschudi, William. 2005. Energy Efficient High Tech Buildings: Can anything be done to improve data center and cleanroom energy efficiency? http://hightech.lbl.gov/events.html
- McLaughlin, Patrick. May 2007. Metrics and Efficiency: The future of data centers. Cabling Installation & Maintenance magazine. http://cim.pennnet.com/articles/article_display.cfm?Section=ARCHI&C=feat&ARTICLE_ID=290948&KEYWORDS=metrics%20and%20efficiency&p=27
- McLaughlin, May 2007b. Reported quote from HP's Christian Belady during his keynote address at the Data Center World conference in March 2007. Cabling Installation & Maintenance magazine. http://cim.pennnet.com/articles/article_display.cfm?Section=ARCHI&C=feat&ARTICLE_ID=290948&KEYWORDS=metrics%20and%20efficiency&p=27
- Pacific Gas & Electric (PG&E). 2007. PG&E Joins The Green Grid Expands Industry-Leading Financial Incentive Programs for Energy Efficient Computing Equipment. August 7, 2007 http://www.pge.com/news/news_releases/q3_2007/070807.html]
- Robb, Drew. August 27, 2007. Green IT Changes Data Center Power Demands. Computerworld. http://www.itworldcanada.com/a/News/54a2d90e-aeb8-4abf-83b5-b6cc86bdd16d.html
- Uptime Institute. 2007. Data Center Efficiency and Productivity.

ADDITIONAL RESOURCES

- 7 x 24 Exchange http://www.7x24exchange.org
- 80 Plus Organization: www.80plus.org
- Data Center Calendar: http://hightech.lbl.gov/dc-powering/calendar.html
- Data Center Documents: http://hightech.lbl.gov/dc-powering/documents.html
- EPRI Solutions: www.eprisolutions.com
- High Tech Events: http://hightech.lbl.gov/events.html
- The Green Grid: www.thegreengrid.org

Contact Us

For more information, call Lineage Power toll free at **888-LINEAGE (888-546-3243)**, or +1 972 244 9288 and visit us on the Web at **lineagepower.com**

LINEAGE POWER®