New Technology Demonstration Program

FEMP

Federal Energy Management Program

Technology Focus

Power Supply Options for Data Centers

William Tschudi
Charles Williams

Lawrence Berkeley National Laboratory
Berkeley, CA

LBNL – 63339

Revised July 2007

The work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Federal Energy Management Program, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Power Supply Options for Data Centers

Executive Summary

For many federal facilities, the fastest growing end-use of electric energy is found in concentrations of computing capacity commonly known as data centers. For these users, the critical importance of information processing to their agency mission will present a serious challenge to meeting the aggressive new energy efficiency goals in Executive Order 13423. Federal energy managers can find a variety of methods for reducing energy intensity, in both design and operations, for these high-technology facilities at <http://hightech.lbl.gov/datacenters.html>. This report summarizes a recent demonstration of one such technique – configuring power supply systems for data centers so that they use DC (direct current) power throughout, eliminating the conventional practice of multiple conversions from utility-supplied AC (alternating current) to DC and back again at every stage of the power supply system. This eliminates both the power loss and heat generated by each such conversion (which drives air conditioning energy use).

The demonstration suggests that direct powering, coupled with selection of high-efficiency power-supply components, can result in as much as 30% improvement in power conversion and distribution to IT equipment as well as overall facility level efficiency, when compared to a typical AC-powered data center. Data do not exist for more than the broadest ballpark estimate of energy use in federal centers, which suggests an order-of-magnitude figure of as much as 6 TWh/year. While data processing equipment decisions are made on the basis of many criteria other than energy efficiency, and no systematic effort was made in this demonstration to estimate the cost-effectiveness of retrofit of power supplies, this demonstration does suggest that federal data centers could potentially be using as much as 1.8 TWh/year less energy.

Background

The importance of high-speed data processing and communication to modern society and economy can scarcely be exaggerated. Thomas Friedman, in The World is Flat (2005: Farrar, Straus & Giroux) argues that they have wrought a more profound revolution change in communication and trade than did the Gutenberg printing press – and have changed the world permanently in far less time. The proliferation of PCs, PDAs and cellular communications, the ubiquity of the internet and the accelerating development of software that allows all of these technologies to interact on a common communication framework have “flattened” our world, in Friedman’s parlance, by literally erasing the significance of geography and national boundaries in the conduct of commerce and trade. Together with the emergence of the leadership and management practices to take advantages of these “flatteners”, and the collapse of political and economic barriers that accompanied the end of the cold war, what Friedman calls the “triple convergence” has literally “changed everything”. Friedman’s observations are supported by the words of the CEO of Hewlett Packard, that the world is entering “an era in which technology will literally transform every aspect of business, every aspect of life and every aspect of society.”

These transformative effects of information technology have had no less impact on our Federal government than on other elements of society. Modern governance would be unthinkable without high-speed data-processing capacity.

A few examples suffice to illustrate:

- Many Americans may feel that they produce very few sets of numbers and calculations more complex than their annual 1040 tax return. In 2005 over half of returns filed were submitted to IRS over the internet, and all 135 million are reviewed in IRS data processing centers.

- Weather prediction in the U.S. is the envy of the world; more than ten of the installations on the list of the world's 500 largest supercomputers are operated by or for the National Oceanic and Atmospheric Administration (NOAA). A handful of others are operated by the Navy for oceanography or meteorology.

- The sophisticated and photogenic warfighting technologies of the U.S. Armed Forces – smart bombs, cruise missiles, supersonic stealth fighters, etc – are obviously computer-based, and the bases that deploy them and the defense contractors that produce them account for another significant fraction of supercomputers in America.

These examples demonstrate the increasing importance to the Federal government of processing prodigious quantities of data efficiently and accurately and/or enormous volumes of communications simultaneously in real time. These functions are typically performed in computer processing facilities often referred to as data centers. A variety of other terms are often employed – server farm, computer center, internet hotel, co-location facility are among the many descriptive terms to describe large complexes of multiple computers.

Figure 1: Data Center

For this report, we will use the term data center, and employ the functional definition of data centers used by Lawrence Berkeley Laboratory in the energy efficiency research roadmap for data centers that LBNL prepared for the California Energy Commission:

“Generally, we use the term data center to be a facility that contains equipment to perform one or more of the following functions:

- Store, manage, process, and exchange digital data and information. This includes corporate data centers, educational facilities, etc.
- Provide application services or management for various types of data processing, such as web hosting, Internet, intranet, telecommunication, and information technology.

We do not consider spaces to be data centers that primarily house office computers including individual workstations, servers associated with workstations, or small server rooms. Generally, the data centers we include are designed to accommodate the unique needs of energy intensive computing equipment along with specially designed infrastructure to accommodate high electrical power consumption, redundant supporting equipment, and the heat dissipated in the process.”[1]

Two of the phrases in this quote, “…heat dissipated in the process” and “redundant supporting equipment” are the keys to the energy intensity of data centers. In recent decades, users of personal computers, cell phones, electronic organizers or other forms of transistor-based telecommunications and data processing devices have observed increases in speed and capacity that can only be measured in orders of magnitude. These advances have been made possible primarily by increasing the number of transistors on the silicon-chip integrated circuits that drive these devices. Gordon Moore, the cofounder of Intel, made an observation in 1965 that became known as Moore’s Law, that “the number of transistors per square inch on integrated circuits had doubled every year since the integrated circuit was invented. Although this industry rule-of-thumb has since been modified to a doubling every eighteen months, it is nonetheless expected to remain valid into the 2020’s for at least another two decades.”

The movement of electrons through these chips, with the transistors determining where they go or don’t go, is of course what drives binary data processing systems and telecommunications. More transistors means more electrons moving, which generates resistance in even the most conductive circuits, and therefore more heat. The relationship tends to be linear; a doubling of transistors generally means a doubling of the amount of heat generated. As more and more capacity can be built into smaller and smaller units, the heat is both concentrated and intensified. For this reason, data centers that further concentrate hundreds of processors into a single space need cooling capacity that can be 50 times the equivalent of conventional office spaces.

Another basic fact that the average user of data processing equipment knows, too often as a result of unfortunate experience, is that if they lose power through grid failure, disconnection, battery failure or any other reason, work in process can be lost, and perhaps even all information stored in memory lost as well. The greater the volume of data processed by or electronic communication transmitted through a given data center, the more catastrophically costly and completely unacceptable such a loss of power is likely to be. This is the reason for the common reference to data processing centers being built to a "six nines" (99.9999%) standard of availability. The typical configuration

---

[1] See http://hightech.lbl.gov/documents/DataCenters_Roadmap_Final.pdf; p. 4.

of data center power supply to achieve that level of reliability – the bulk of the “redundant supporting equipment” noted above -- entails both power losses and increased production of heat – making power supplies the second major source of extremely intensive energy use in data centers.

Current Practice – Electrical distribution in data centers

Electrical power flow from entry into the data center through to the chip level where processing takes place is characterized with numerous changes between alternating (AC) and direct current (DC) involving many different voltages. Typically, the electrical service to the building arrives as 480 Volt AC and undergoes many conversions before serving the end loads for the electrical IT components which require various DC voltages. Depending upon the equipment selected the conversion losses can easily account for 20-30% or more, of the total facility electrical consumption.

This can be shown schematically as in figure 1.

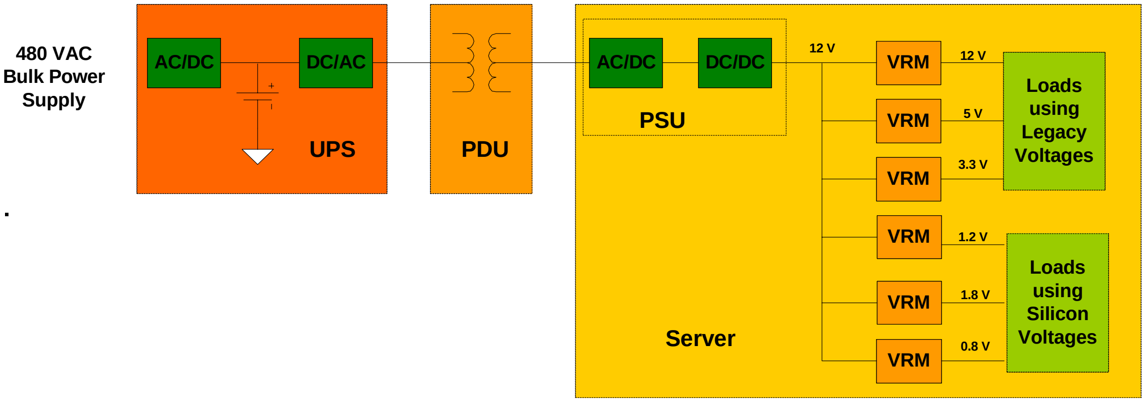

Figure 1
Courtesy of Intel Corporation

High voltage AC (480volts) enters the data center and is converted to DC in an uninterruptible power supply (UPS) in order to charge batteries that are used to backup the power supplied by the utility. Then the power is converted back to AC. A further AC to AC conversion occurs in a device called a power distribution unit (PDU). Here voltages are reduced to the voltage input into the IT equipment.- typically 208 volts.

Once inside the server a conversion from AC to DC occurs along with power factor correction followed by further multiple DC to DC voltages conversions. A byproduct of each power conversion is heat which must then be removed through the HVAC system – requiring additional energy. The end to end efficiency of a traditional system from entry to the building through to the end processor may be on the order of 60%.

Studies have shown that there is a wide range of energy efficiency in both UPS systems and the power supplies inside servers. Figure 2 below, illustrates the efficiency ranges that were measured in benchmark testing[2].

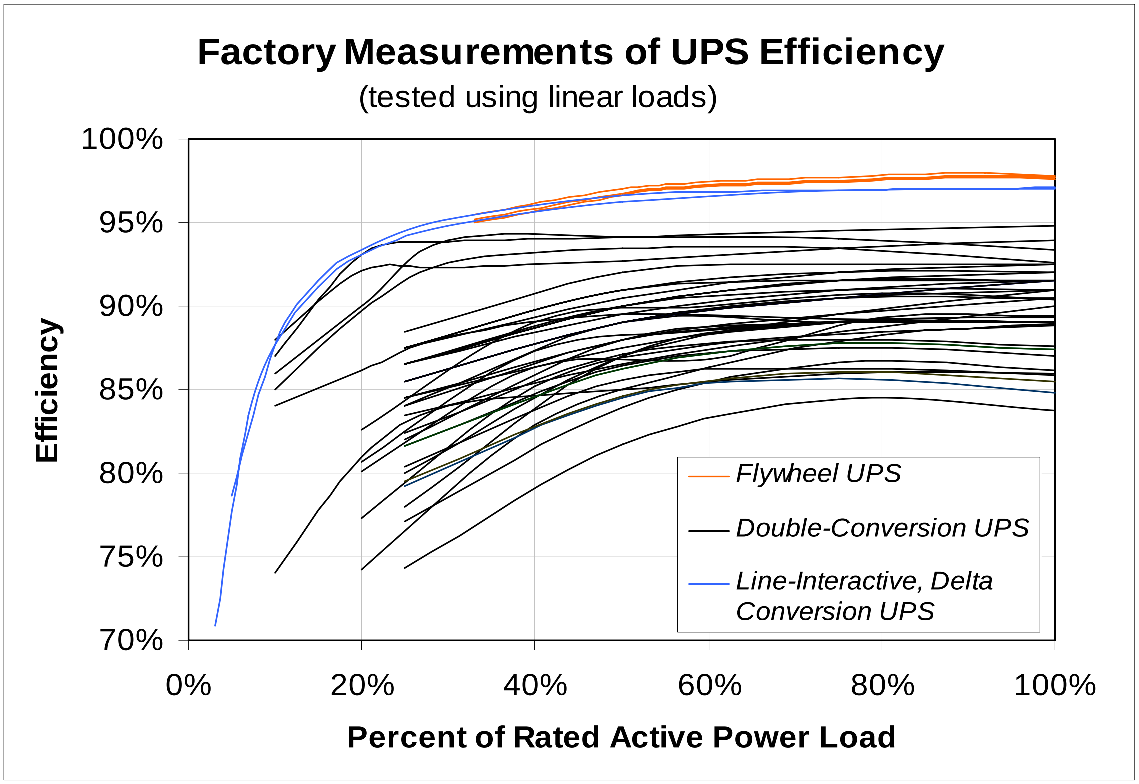

Figure 2

Notice that efficiencies are typically lower for lightly loaded UPS systems. For redundancy purposes most UPS systems are loaded at less than 40% of their rated load. Frequently systems are oversized to allow for future growth or load uncertainty, which creates further inefficiency. In today's data centers it is not uncommon to consume 15% or more of the power within the UPS system serving the IT equipment.

Similarly, the efficiencies of power supplies used in IT equipment vary considerably and have a similar efficiency drop off at lower loadings. Figure 3 shows efficiencies of benchmarked, low-end servers.[3]

---

[2] Study conducted by Ecos Consulting and EPRI Solutions for LBNL - see: http://hightech.lbl.gov/ups.html

[3] Study conducted by Ecos Consulting and EPRI Solutions for LBNL - see: http://hightech.lbl.gov/psupplies.html

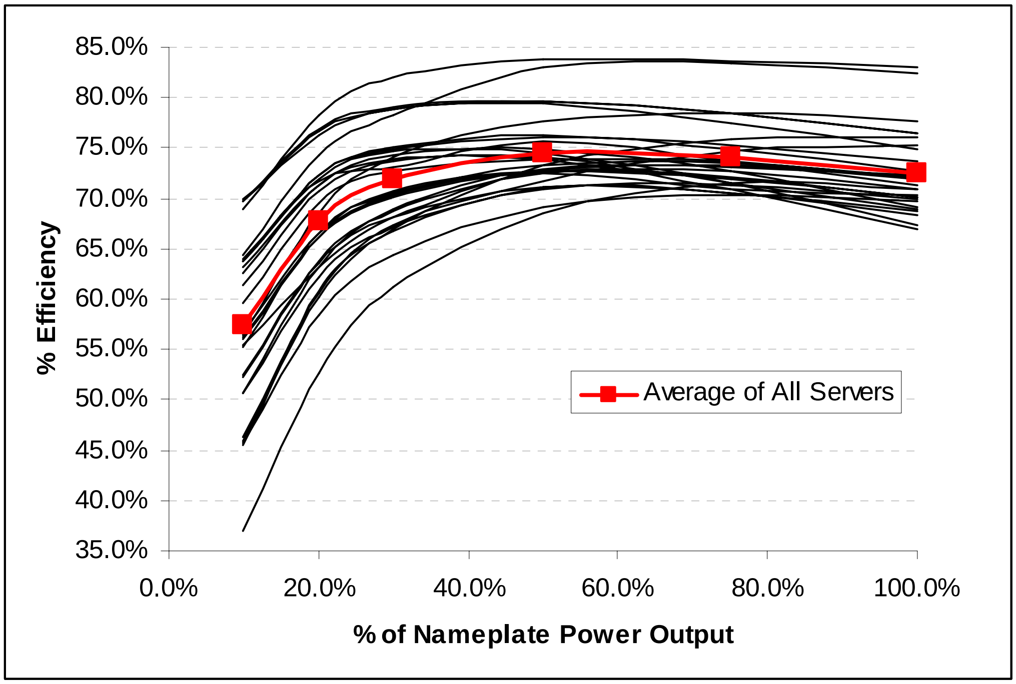

Figure 3

For this group of server power supplies, a typical operating point would be at about 70% efficiency. This means that 30% of the power intended for the computing operations was lost in the power supply. This, coupled with the losses in the UPS system accounts for a large percentage of the power used in data centers - both for the direct losses and for the HVAC required for cooling. Selecting efficient components or eliminating some of the conversion losses in these two areas represents a huge potential for most data centers. Combined energy savings for UPS and power supply efficiency improvement could easily reach 20% in some cases simply by selecting more efficient components.

Smaller losses also occur in the power distribution units (PDUs) which are essentially transformers which condition the AC power being delivered to the servers. These losses are relatively small; however, eliminating these from the distribution path would result in some savings. Relatively small losses in the power cabling also contribute to the overall efficiency.

Benchmarking of HVAC systems shows a similar wide variation of performance however the industry typically assumes a 1:1 correlation between IT equipment loads and building loads. Figure 4 shows that the ratio of power that finally reaches the IT equipment can vary considerably which indicates that some centers are more efficiently delivering the necessary power and cooling. Assuming a 1:1 correlation, a 20% energy saving through the power distribution to IT equipment, would translate into a 20% *Facility Level* energy savings on average.

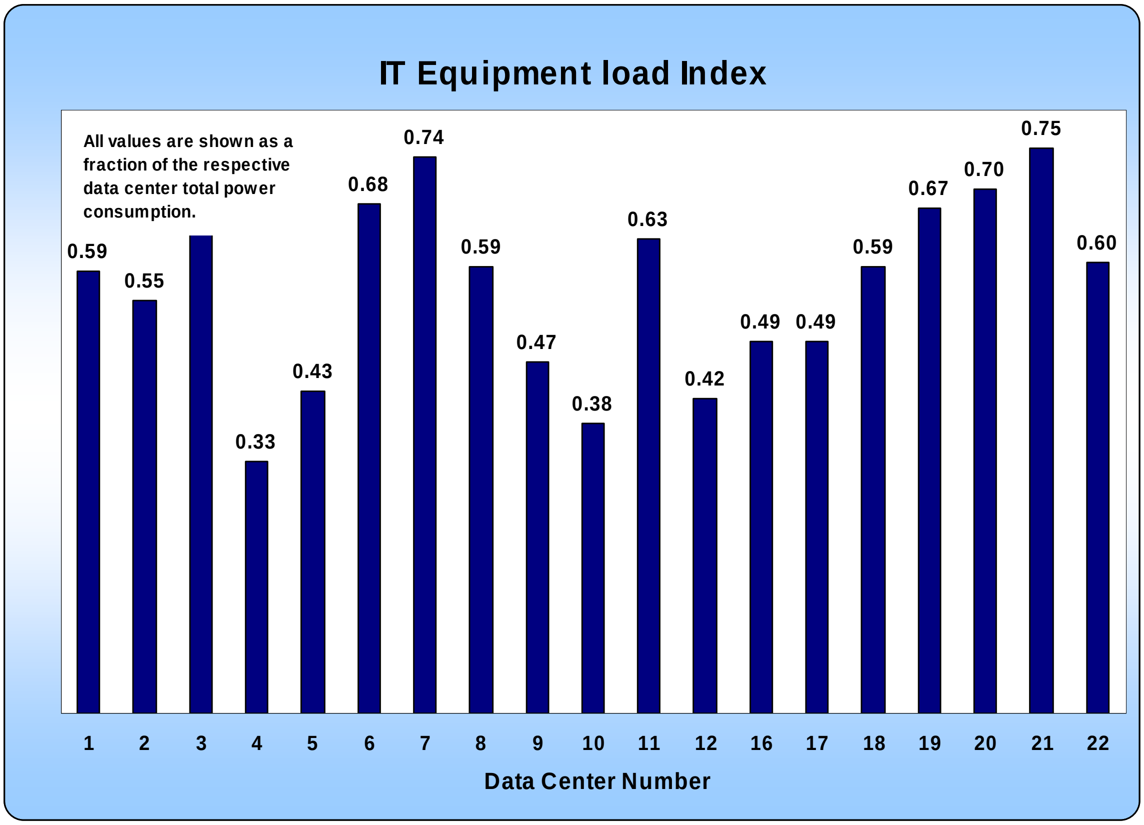

Figure 4

High Voltage DC distribution

But could the efficiencies go even higher? A different approach in the distribution of power in a data center involves completely eliminating some of the conversion steps. A proof of concept demonstration was established at Sun Microsystems' Newark, CA Campus to show how such a system could be assembled with commercially available equipment. Over 20 firms participated along with LBNL in providing equipment or services necessary to assemble the demonstration. In this scheme, DC power is supplied directly to the IT equipment as shown in figure 2:

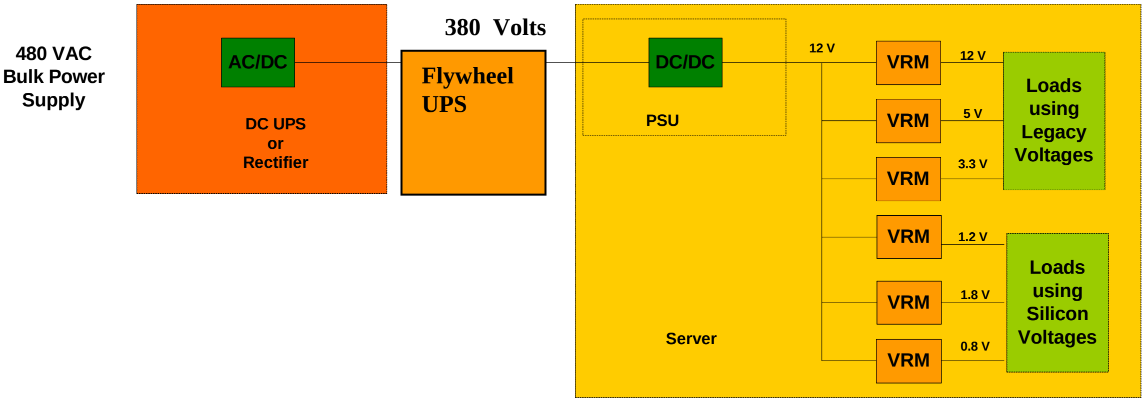

In this scheme, the power entering the data center is converted once from 480 Volts AC to 380 Volts DC. Backup capability is provided by a highly efficient rotary UPS system (98% efficient) or it could be provided through traditional battery systems. In any case, 380V DC power is then delivered directly into the computers. Multiple conversions are eliminated resulting in improved energy performance

In this proof of concept demonstration, a traditional AC distribution system was assembled next to the DC system for a side by side comparison. Figure 5 shows the side-by-side demonstration in operation at Sun Microsystems' Newark, CA facility. For additional technical detail on the Data Center DC Power Demonstration project, see "DC Power for Improved Data Center Efficiency"; My Ton, Brian Fortenbery and William Tschudi, Lawrence Berkeley Laboratory, January 2007 at http://hightech.lbl.gov/dc-powering and additional research into energy efficiency for data centers at http://hightech.lbl.gov/datacenters.

Figure 5

Identical servers were provided in both systems with the exception that the DC system servers were modified to directly accept 380 Volt DC by by-passing the traditional AC-DC conversion (and power factor correction) thereby eliminating the first conversion inside the computer. The AC distribution system contained a very efficient UPS system yet the DC system was able to achieve energy savings of 10-14% over the traditional system. If the DC system were compared to "average" or "poor" distribution system energy savings in excess of 30% could be obtained.

UL rated, DC power distribution equipment is commercially available as was shown in the proof of concept demonstration. This fact alone resolves the question raised by many about safety when using a DC distribution system. DC systems are widely used in many applications including transit systems, military, elevators, power transmission, etc. For use in data center applications, there could be additional benefits of eliminating many of the power quality issues when using AC power, resulting in improved reliability and up time. In addition since the DC system will contain fewer components, there would be improved reliability due to fewer potential points of failure. The conversion from DC to AC in a traditional UPS, the power distribution unit (PDU), and the first stage conversion inside the servers could all be eliminated with this scheme.

While DC power has been shown to offer significant energy savings, there is also a potential for first cost savings. If this technology were to be deployed on a wide scale, the capital cost of equipment would likely be less than the AC counterpart since there are fewer devices, and reliability could be improved because there are fewer points of potential failure.

Federal Data Center Facilities

The federal government owns and operates a large segment of the data center market. High Performance Computing centers are numerous within the federal sector, as are computer centers supporting virtually every aspect of the federal government. Of the top 500 supercomputers in the world, approximately 20% are owned and operated by the federal government.[4] These are among the most energy intensive centers in the market often requiring 10 megawatts or more of electrical power. Other centers support the wide range of government activities such as the IRS, flight control, Social Security, Defense, etc.

Federal Data Center Power Usage

No listing of data centers owned and operated by or for the federal government has been compiled, much less an assessment of their power usage. Many of the larger centers are “Defense Classified”[4] ; a great number of smaller data centers will be included in larger office facilities and neither their existence as data center not their power consumption will appear separately in federal facility inventories or energy use reports. With some simplifying assumptions, however, we can develop an order-of-magnitude estimate.

A 2004 study by EPRI Solutions and Ecos Consulting estimated that servers in the U.S. accounted for 14.6 TWh of electric use annually, and UPS systems for 7.1 TWh. If federal facilities represent 20% of the data center market in the U.S. (perhaps

---

[4] Of the 500 largest supercomputers in the world, at least 90 are designated as operated by federal agencies, USDOE National Laboratories, or public-private consortia that rely on Federal funding. There may be 10-12 more Federal facilities among the U.S. entries in the Top 500 list. The list has a dozen or so non-specific entries like “Government” or “Defense Classified”).

conservative, given that 20% of the supercomputers in the entire world are operated for the U.S. government), then federal data centers may represent 4TWh/year of electric usage. Moreover, bench-marking of a dozen data centers in California found that HVAC power consumption (primarily cooling of all the heat generated by servers and power supplies) averaged more than a third of total data center power consumption. Given this, then federal data centers electric usage may be approximately 6TWh/year. Data processing equipment decisions are made on the basis of many criteria other than energy efficiency. But if 30% efficiency gains are attainable through use of efficient DC power supply systems, then this ballpark estimate suggests that federal data centers could potentially be using as much as 1.8 TWh/year less energy.