High Performance Buildings: Data Centers

Server Power Supplies

December, 2005

My Ton
Ecos Consulting
Portland, Oregon

Brian Fortenbury
EPRI Solutions
Knoxville, TN

The work described in this report was performed as part of Lawrence Berkeley National Laboratory's (LBNL) "High-Performance High-Tech Buildings" project focused on improving energy efficiency in Laboratory, Cleanroom, and Data Center facilities. This project was sponsored by the California Energy Commission's Public Interest Energy Research (PIER) Industrial program with technical leadership by LBNL. Ecos Consulting and EPRI Solutions performed the research reported herein collaboratively under subcontract to LBNL and with input from industry advisors and UPS manufacturers.

Legal Notice

Table of Contents

1 Executive Summary

As part of the California Energy Commission's PIER (Public Interest Energy Research) initiative on efficient data centers, Ecos Consulting and EPRI Solutions were tasked with investigating and characterizing server power supplies as well as the power supplies used by other devices in data centers.[1] The goals for this project include an analysis of power supply efficiencies, which included the development of an accepted test protocol for server power supplies, lab and field testing of a broad range of server power supplies and documenting the results. Other objectives of the project include the wide circulation of efficiency findings to the industry through the CEC, PIER, Lawrence Berkeley National Laboratories (LBNL), and other industry and efficiency venues, such as SSI – the Server System Infrastructure group, the PSMA – the Power Sources Manufacturers Association, the American Society of Heating, Refrigeration and Air Conditioning Engineers – ASHRAE, and www.efficientpowersupplies.org. Using the efficiency findings and market data, we also estimated the overall energy consumption of servers in the United States, especially in the State of California, as well as the potential savings from the use of more efficient units. Finally, we also worked with industry groups to press the case for more efficient power supplies.

1.1 Current Power Supply Efficiencies

1.1.1 Laboratory Testing

Our test results showed that most server power supplies' efficiency at converting AC to DC typically peaks at loads between 50-60% and drops off dramatically at loads under 30%. The tested power supplies have efficiencies in the 70-75% range (at 50% load). The most efficient power supply tested demonstrate that significantly more efficient server power supply designs do exist on the market compared to the average performers, and this may represent significant energy savings potential for data centers or enterprise computing where servers are operated on a "24/7" basis. Our findings also show that server power supply designs with poor efficiency are still available and can result in unnecessary power consumption and excess heat generation while in use.

1.1.2 Field Testing

To investigate the relationship between server operation and power supply loading, we developed testing protocols and collected field data on servers from a number of facilities, including LBNL (scientific computing), EPRI Solutions (commercial/high tech facility), and Ecos Consulting (general business applications). Our finding from these field tests shows that even highly utilized machines like the LBNL servers do not fully utilize the capacity of their power supplies. It also suggests that server power supplies are oversized for the actual requirements of the machines in which they are being used.

Most servers power supplies measured were operating somewhere between about 20% and 50% of their rated load all of the time, which is the exact portion of power supply curves at which efficiency begins to dramatically decrease. None of the servers tested ever exceeded 50% of its rated output. The combination of low efficiencies and oversizing of server power supplies can have dramatic effects on the net energy consumption of those power supplies. An oversized, inefficient power supply would often waste two to three times as much net AC power[2] to meet that load as a properly sized, efficient one.

---

[1] At the most basic level, power supplies are used to convert alternating current to direct current needed by the processors in servers as well as other electronic components.
[2] Here net AC power is defined as the difference between the AC power input of the power supply and its DC power output.

1.1.3 CPU Utilization

In addition to measuring power consumption, we also extracted CPU utilization data from each server's log for a 24-hour period during field testing.[3] A CPU utilization curve for each server was constructed representing the percent of time the server's CPU utilization exceeded a given percent. It is difficult to draw any conclusions from such a small sample, but it appears that for the most part power supply loading has very low correlation with "server activity."

The servers we measured do not exhibit any activity-based power management, whereby power consumption is reduced when the server is operating below full capacity. Activity-based power management attempts to scale processor power consumption to the demand for processing resources placed on the CPU. The absence of this technology in the servers that we measured might help to explain the lack of any correlation between the utilization curve of the CPU and the load duration curves of the power supplies.

1.1.4 Proposed Power Supply Efficiency Standards

Based on the efficiency levels documented in existing equipment and the field studies conducted, we closely coordinated with SSI and Intel[4] to propose changes to the SSI specifications for higher power supply efficiencies.

At the Intel Developer Forum in March 2005, in a joint presentation, Intel announced new proposed SSI Industry specifications to encourage more efficient power supplies. The new specifications include testing conditions which align with ATX12V, as well as required and recommended efficiency levels at 20%, 50%, and 100% loading. SSI is also now supporting a power supply technology known as Power Supply Management Interface (PSMI). This new industry standard provides a basic internal mechanism for server power supplies to report power consumption and efficiency data directly to the server's motherboard with 5% to 10% accuracy. The technology could conceivably be used for a number of purposes. Power and efficiency data could be logged on the server's hard drive or reported to HVAC equipment so that fans and cooling equipment could scale their output to the heat output of IT equipment.

1.2 Energy Saving Potentials

We estimated the energy savings potential of more efficient server power supplies in three steps:

1. We estimated the current AEC (Annual Energy Consumption) of servers in the United States using the basic methodology developed in the ADL/US DOE report.[5]
2. We determined the percentage of current server energy use that could be saved due to more efficient power supplies, based on the new SSI recommended specifications.
3. By applying the percentages established in #2 to the AEC estimates established in #1, we were able to estimate the overall energy savings potential of the new recommended efficiency levels.

For California-specific estimates, we used a range of 10% to 15% of estimated US results.

1.2.1 Annual Energy Consumption of Servers

Starting from the basic methodology established in the ADL/US DOE study, we constructed a revised estimate of the annual energy consumption (AEC) of servers in the United States.

[3] CPU utilization describes how busy a server's processor or processors are; it tells us the percent of the CPU's processing capacity, or throughput, that is being used to perform tasks.
[4] Brian Griffith, Intel Power Server Architect, EPG (SSI Coordinator).
[5] Roth, Kurt W., Fred Goldstein, and Jonathan Kleinman (Arthur D. Little). "Office Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings. Volume I: Energy Consumption Baseline." Prepared for US DOE, Building Technologies Program, January 2002. NTIS Number: PB2002-101438.

Figure ES 1 shows our revised estimate of server AEC for 2004. The revised AEC is 14.6 TWh, which represents a 45% increase over the ADL/US DOE estimate of 10.1 TWh for 2000. This increase is attributable to growth in both the number and average power draw of low-end servers, which comprise the majority of units. The estimate for California is in the range of 1.5 TWh to 2.2 TWh.

Figure ES1. Server Annual Energy Consumption (AEC)

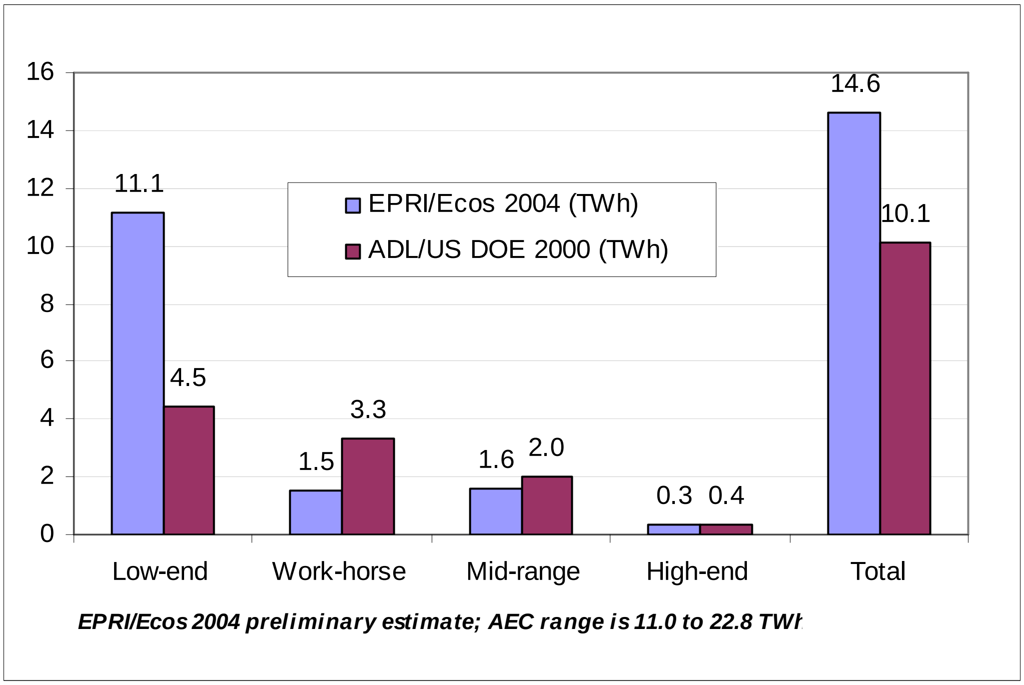

1.2.1 Energy Saving Potentials – Servers

Based on our overall estimate of server AEC, we estimate the amount of electricity that can be saved due to more efficient power supplies would be approximately 1.5TWh. We also project a high-efficiency case, where power supply efficiency is 83% and electricity savings of 2.3 TWh are realized. An estimate for the potential savings for California is between 10% and 15% of the overall US savings potential, or between 230 GWh and 345 GWh.

Table ES1. Energy Savings Potentials

Server Category	Number of Servers (Millions)	Annual Energy Consumption in TWh @ 70% PS Efficiency (Current)	Annual Energy Consumption in TWh @ 78% PS Efficiency (New SSI)	Annual Energy Consumption in TWh @ 83% PS Efficiency (High Adoption)
Low-end	6,587,061	11.1	10.0	9.4
Work-horse	506,470	1.5	1.4	1.8
Mid-range	151,678	1.6	1.5	1.4
High-end	14,730	0.3	0.3	0.4
Total	**7,259,939**	**14.6**	**13.1**	**12.3**

Improved power supply efficiency will also yield compounded savings from reduced air conditioning loads and UPS losses not included here. We also note that the potential for the infrastructure load reduction (i.e. improved overall efficiency and reduced cooling load) is on the same order of magnitude as the power supply efficiency improvement.

1.2.2 Annual Energy Consumption – Other Data Center Devices

We also investigated the range of other devices used in data centers, and arrived at a first order estimate of their energy consumption. Our estimates used a similar methodology to the one employed by ADL/US DOE (2002). This methodology, however, does not account for existing power supply efficiency, nor does it account for the fact that devices normally operate at a fraction of their stated output rating. Therefore, revised our initial AEC estimates to reflect an average of 30% loading and 70% power supply efficiency. Table ES2 presents the "adjusted" AEC.

Table ES2. AEC, adjusted for 70% PS Efficiency and 30% Loading

Segment	AEC (TWh)	Adjusted AEC (TWh)	% of Adjusted AEC
Routers	1.3	0.6	8%
Switches - LAN	9.0	3.9	59%
Switches - WAN	0.4	0.2	3%
Hubs	1.3	0.6	8%
Storage Devices	3.3	1.4	22%
Total	15.3	6.6	100%

Note: AEC is adjusted by a factor of .43 =.3/.7 to account for loading and PS efficiency.

Together, these "other devices" consume a little less than half of the energy that servers do. LAN switches contribute the most to AEC simply due to their large numbers. Remember that, as with servers, not all of the energy consumed goes to power the device itself. Some energy is "lost" or consumed by the power supply itself; the amount of energy consumed depends on the efficiency of the power supply. In the typical existing case, where power supply efficiency is 70%, 30% or about 2.0 TWh is consumed by the power supply (i.e., energy losses)

1.2.3 Energy Savings Potential – Other Data Center Devices

As with servers, energy savings can be achieved by increasing the efficiency (or reducing the energy losses) of the power supply of other devices. These high-level estimates show that there is significant energy savings potential from improving efficiency of power supplies in devices other than servers. While the savings potential is not as large as in servers in absolute terms, the extension of efforts to these "other' devices would be relatively straightforward. Much of the groundwork for improving power supply efficiency and developing specifications for servers, desktops, etc. has already been done or is in process. This work could serve as a starting point for efforts targeting other devices.

The main challenge to encouraging efficiency in the power supplies of other IT equipment like routers and switches is that there is currently no industry body like SSI coordinating standards and efficiency improvements in these products. Power supply designs for this type of equipment are usually customized for a particular product and may vary from manufacturer to manufacturer, whereas in servers, only a few distinct form factors for power supplies exist. The diversity of power supply designs in equipment like routers and switches may complicate efforts to uniformly encourage efficiency improvements.

1.3 Conclusions and Recommendations

This project served well as a first foray into this particular area of high-tech buildings and data centers, and has provided a good overview of server supplies utilization and current efficiency levels. Because of the broad scope, however, we were only able to identify areas for further investigations, but not able to carry out in-depth investigations of any particular server category or application. Below are a number of observations and recommendations based on the findings of this study:

Server Loading: While we got fairly consistent results from our small sample indicating most server power supplies are not fully taxed in everyday use, it would be difficult to project what the loading would be for servers across all data centers. We recommend a more rigorous sampling of servers and their various tasks as a logical next step.

Server Replacement Cycles: We have not investigated this issue, or the associated issue of what type of market penetration can be expected as these new systems go into the field if new SSI standards are adopted. We currently do not have such information, other than anecdotal – this is another area that we would recommend for additional research

Retrofit Market for Server Power Supplies: We have also not investigated whether or not there exists a retrofit market for server power supplies. Similar to the replacement issue, this is another area that we would recommend for additional research.

Implications of Blade Servers: This server class is currently experiencing high growth, and does not require a power supply to be associated with it. While this category has the potential to increase efficiency due to the fact that many blades can share one power supply, the efficiency of blade power supplies has not been adequately considered.

Continue Work with the SSI Group: Given that the SSI had already published some new standards, this one area that we can best support the industry by helping to promote continuous energy efficiency improvement industry wide. Our recommendation is that the project team continue the relationships that we developed.

More, Better Industry Data: We were not able to obtain regional data on product shipments, thus we were only able to approximate the California market based on population estimates. The server locations that we tested are representative of locations found in CA, and two locations out of the three that were tested employ servers that belong in the "low-end" category. This is another area where additional research is recommended.

Documenting Efficient Server Power Supply Designs: We have not investigated any power supply design changes under the assumption that the industry will pursue their own configurations that meet new proposed standards. It is worth noting that here are a variety of approaches that can be taken to design power supplies for higher efficiency. For example, from the PIER sponsored power supply primer, we find that the most lossy subsystems include:

- The switching element
- Control IC
- Transformer
- Output rectifier

Losses can further be addressed with:

- More efficient Power Factor Correction chip
- Active Clamp on the main transformer to replace lossy RCD clamp

To illustrate some of these approaches, the PIER funded Efficiency Challenge design competition in 2004 yielded around a dozen designs all using different techniques to achieve greater than 80% efficiency. Some of those techniques included:

- Optimized selection of control IC using a variable off-time technique to lower losses.
- Flyback transformer optimized to reduce leakage inductance and winding resistance
- Operation mode optimized to strike a balance between switching loss and conduction loss
- The MOS switch is carefully selected to reduce switching and conduction losses
- Output rectifier is also carefully selected to reduce switching and conduction losses
- Output rectifier can use synchronous rectification with low loss MOS switches
- Output rectifier can use Silicon Carbide diodes to lower losses
- Synchronous rectifier on the input

- Litz wire used in flyback transformer design
- One method called burst-mode control halts all switching for light loads

A manufacturer can use a few of these methods to achieve improved levels of efficiency; the methods chosen may depend on cost, availability, or designer preference. But in conclusion, the industry has a wide variety of choices available and will likely develop even more. There will not likely be any stifling of creativity as a result of these standards.

2 Existing Power Supply Efficiency

This task's objective was to document the existing efficiencies of server power supplies. To do this we:

- Assessed the power supply topologies in servers.
- Developed the test protocol for testing internal power supplies.
- Conducted laboratory testing of different power supply configurations and loadings.

2.1 Server Power Supply Topologies

Servers come in all shapes and sizes, as do the power supplies that they contain. The Server System Infrastructure[6] (SSI) group led by Intel has identified several distinct categories of server power supplies (Table 1).

Table 1. SSI Power Supply Product Categories

Category	Description	Form Factor
EPS1U	Entry-Level Power Supply	1U
EPS2U	Entry-Level Power Supply	2U
ERP2U	Entry Redundant Power Supply	2U
EPS12V	Entry Non-Redundant Power Supply	PS/2
ERP12V	Entry Redundant Power Supply	Pedestal servers
EPS	Entry Power Supply	PS/2
TPS	Thin (low profile) Power Supply	Low-profile servers
MPS	Midrange Power Supply	Midrange chassis, fits 3 or 4 across within a 19" rack mount system
DPS 2.0	Distributed Power Supply	For system using Distributed Power System (DPS) architecture – delivering 48VDC bulk power. 2.74W x 4.86H x 12.8L (inches)

Based on industry market research (See Section 3.2) and consultation with Intel, SSI members and other industry partners, we determined that the dominant technology in the marketplace is the multi-output front-end ac-dc power supply. These power supplies are commonly found in 1U, 2U, and pedestal servers[7], which account for over 90% of the server market and are found in both offices and data centers alike.

Consequently, we selected three corresponding SSI power supply categories—**EPS1U, EPS2U, and EPS12V**—as the focus of this study. These topologies were chosen mainly because they represent the most common designs used in rack and pedestal servers today. They are also non-redundant power supply designs, meaning that they contain only one path for AC electricity to be converted and

[6] SSI, www.sssiforum.org, is an industry initiative intended to provide ready to use design specifications for common server hardware elements (chassis, power supplies, and racks) to promote and support server industry growth. The initiative is comprised of "Promoters" such as Dell, Intel (SSI's founder), IBM, and Silicon Graphics, and "Adopters" -- companies that have agreed to adopt and utilize final SSI specifications in product design and manufacturing. A complete list of Adopters can be found on the SSI website: http://www.ssiforum.org/membership.aspx.

[7] The "U" in 1U and 2U servers denotes that the server is 1 or 2 units (1.75 or 3.5 inches) in height.

delivered to the server's various components (although there may be more than one unit installed in each server box).

A straightforward test method has been developed to test the efficiency of non-redundant internal power supplies, but this method does not yet address redundant designs, in which there may be two or more paths for AC electricity to flow to the server.[8] For this reason it was most sensible to focus on the above categories for the purposes of this study. A standard test procedure for redundant supplies is needed, but that it would likely have different loading guidelines than non-redundant supplies, as redundant units may spend more time at lower loading levels.

2.2 Test Protocol for Power Supplies

SSI has developed design guides, which include recommended efficiency specifications, for many of its designated power supply categories.[9] Current versions of these specifications are available on SSI's website at http://www.ssiforum.org/specifications.aspx.

Early versions of the SSI specifications did include power supply efficiency recommendations, but only focused on *full load* conditions. Yet, as field tests demonstrate (See Section 2.2), many server power supplies operate at less than full load a majority of the time and at a significantly lower efficiency than they would at full load. Reasons for this might include low "traffic" at a particular points in time (e.g., overnight), oversized power supplies, or (although not the primary focus of this study) the use of multiple power supply units to insure up-time in the event of a failure.

To better document the efficiency of server power supplies at these lower loads, EPRI Solutions and Ecos Consulting worked with the SSI group and other industry stakeholders, including the PSMA – the Power Sources Manufacturers Association to develop a standardized test method for measuring server power supply efficiency across a *range of loads*. Such procedures allow the development of standardized efficiency curves, so purchasers can size power supplies properly and reasonably estimate the efficiencies they will achieve in operation.

EPRI Solutions previously developed a standardized test procedure for internal desktop computer power supplies. (See **Proposed Test Protocol For Calculating The Energy Efficiency of Internal Ac-Dc Power Supplies, Revision 4.0** this document is currently available at:

http://www.efficientpowersupplies.org/pages/Generalized_Internal_Power_Supply_Efficiency_Test_Protocol_R4.pdf.

Building upon this work, and in consultation with Intel and SSI, EPRI Solutions developed a new test protocol for calculating the energy efficiency of internal ac-dc power supplies typically used in computer servers. This document is available at:

http://eetd.lbl.gov/ea/mills/ht/Documents/PS/Server_Draft_PSTestProtocol.pdf.

Detailed loading guidelines were added for different types of server power supplies. A draft test protocol was circulated to interested parties (e.g., all SSI and project advisors) and was also distributed at various industry conferences and meetings (e.g., Intel's Fall IT Symposium and the ASHRAE's Technical Committee 9.9 Annual Meeting in 2004). Based on comments received, the test protocol was revised a second time.

Following is a key excerpt from the protocol:

> *4.3 Power Supply Loading*
>
> *The efficiency of the power supply shall be measured at 20%, 50%, and 100% of rated current. In addition to these three load conditions, other loading conditions may be identified*

[8] The redundant units can be configured in a number of ways, thus creating different load levels on the PS depending on the configurations.
[9] Efficiency in this case is specifically defined as measured input power to the power supply minus measured output by the power supply, and divided by the measured input.

that are relevant to the manufacturer and user of the power supply. Procedures for loading power supplies are described in detail in Section 6.1.1 below[10]*.*

Testing at a load condition below 25% load condition should be guided by IEC 60180 Ed 1.0, Measurement of Standby Power, which establishes the measurement methods for low-power-mode operation of an appliance.

2.3 Laboratory Testing and Results

2.3.1 Power Supplies Tested

Having developed the test protocol, 28 power supplies were obtained from manufacturers and other industry partners, such as the Power Sources Manufacturers Association (PSMA) (Table 2).Some of these power supplies were purchased; others were borrowed and later returned to the manufacturer.

S. No	SSI Type	Model No.	Rated Power
			Watts
1	TPS1U	DPS-200PB-118	203
2	EPS1U	DPS-350-PB	350
3	EPS1U	DPS-350-PB	350
4	MPS	DPS-450-CB 1	450
5	MPS	DPS-450-CB A	375
6	TPS1U	DPS-125FB	125
7	EPS1U	ENH-0620	200
8	EPS1U	P1G-6300 P	300
9	EPS12V	FSP550-60PLN	550
10	EPS12V	ENS-0246B	460
11	EPS1U	FSP460-631U	460
12	EPS12V	FSP460-60PFN	460
13	EPS1U	ENH-0635A	350
14	EPS2U	P2G6460P	460
15	EPS1U	TC1U35	350
16	EPS1U	TC1U40	400
17	EPS1U	FSP350-601U	350
18	EPS2U	TC2U35	350
19	EPS2U	TC2U40	400
20	EPS12V	API4FS06	550
21	EPS12V	API4FS06	550
22	EPS1U	AP13FS43	500
23	EPS12V	PSM6600P	600
24	EPS1U	M1G6500P	500
25	EPS1U	M1G6500P	500
26	EPS12V	HP2-6500P	400
27	EPS12V	HG2-6400P	500
28	DPS	DPS-400-GB-1	400

Table 2. List of Server Power Supplies Tested

[10] Loading guidelines are needed to ensure consistency when measuring the efficiency of server power supplies with multiple outputs. Guidelines currently do not exist, but may be added in the future. Reference: http://www.ssiforum.org/html/adoptedspecs.asp, for a specification that defines a non-redundant power supply that supports an entry-level server.

Each of the power supplies was tested according to the procedures outlined in the *Proposed Test Protocol for Calculating the Energy Efficiency of Internal AC-DC Server Power Supplies, Review Draft Revision 2.0.* Figure 1 provides a photo of the laboratory testing facility.

Figure 1. Laboratory Test Setup

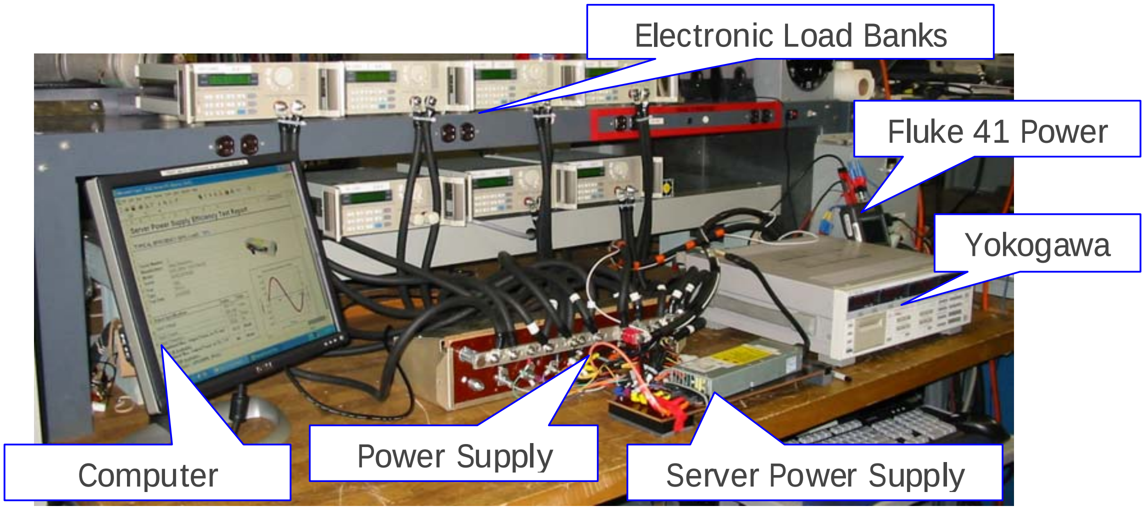

2.3.2 Efficiency Test Results

Figure 2 summarizes the results of the power supply tests.[11] Each black line on this chart represents the measured "efficiency curve" of an individual server power supply: the red line represents the average of all power supplies tested. The largely consistent shape of these curves shows that the measured efficiency typically peaks at loads between 50-60% and drops off dramatically at loads under 30%. Most power supplies are grouped closely together, with efficiencies in the 70-75% range (at 50% load). A few outlying power supplies show relatively high or low efficiency (i.e., the efficiency curves at the top and bottom of the chart). The most efficient power supply tested demonstrate that significantly more efficient server power supply designs do exist on the market compared to the average performers, and this may represent significant energy savings potential for data centers or enterprise computing where servers are operated on a "24/7" basis.

Unfortunately, our findings also show that server power supply designs with poor efficiency are still available and are causing unnecessary power consumption and excess heat generation in data centers., The least efficient power supply measured was a 300W EPS1U unit, which was 52.7% efficient at 20% load, and only 37.0% efficient at 10% load. This means that at 20% load, where many server power supplies are typically loaded, this unit would consume roughly half of the electricity flowing through it simply through inefficiencies in the design.

[11] Individual test reports for each power supply are available at: http://eetd.lbl.gov/ea/mills/ht/Documents/PS/Server_Draft_PSTestProtocol.pdf.

Figure 2. Measured Power Supply Efficiencies

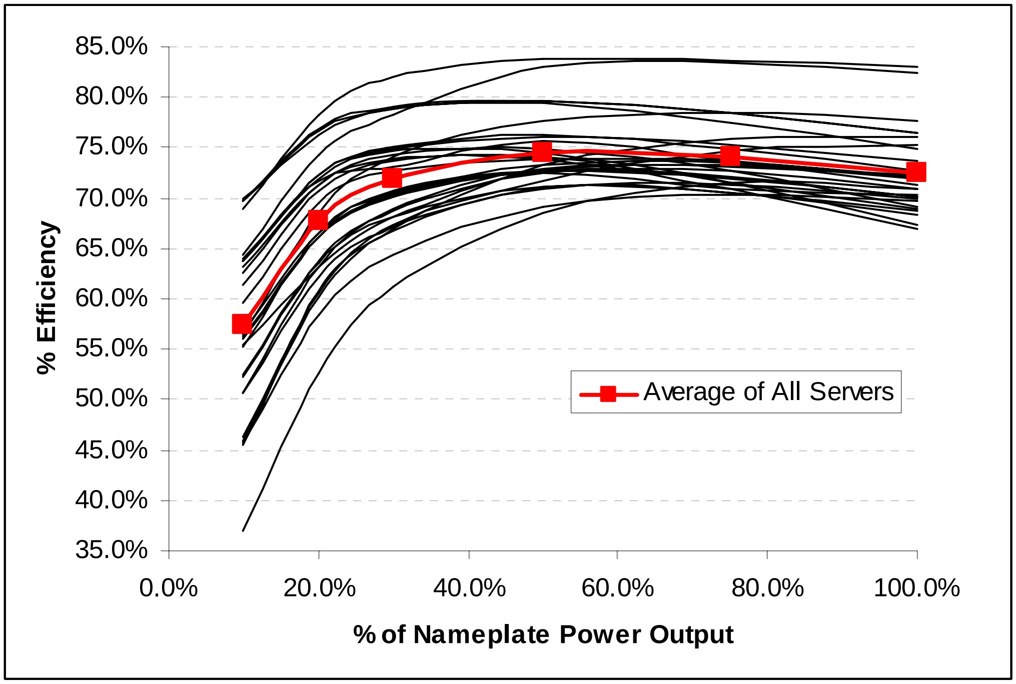

Table 3 shows that the average measured efficiency of all 28 power supplies was only 57.4% at 10%. Peak average measured efficiency for all power supplies was 74.6% at 50% loading.

Table 3. Average Measured Power Supply Efficiency (by SSI Type)

SSI Type	# of PS Tested	Loading %					
		10%	20%	30%	50%	75%	100%
EPS12V	8	60.4%	69.8%	73.3%	75.5%	74.5%	72.6%
EPS1U	12	55.7%	66.3%	70.9%	73.8%	73.8%	72.3%
EPS2U	3	62.9%	71.7%	74.5%	75.0%	72.6%	69.6%
TPS1U	2	53.6%	65.9%	71.6%	75.3%	75.8%	75.0%
MPS	2	48.2%	60.3%	65.8%	70.1%	70.9%	69.9%
DPS	1	64.5%	74.4%	78.2%	82.9%	83.5%	82.4%
Best-in-class	1	69.0%	78.3%	82.1%	83.9%	83.6%	83.0%
Worst-in-class	1	37.0%	52.7%	61.1%	68.6%	71.5%	70.9%
All Power Supplies	28	57.4%	67.7%	71.9%	74.6%	74.2%	72.5%

Note: Efficiency values shown for “All Server PS” correspond to red line in Figure 2.

While the data set is not large enough to state definitively that one power supply form factor is automatically more efficient than another, we can draw the following broad conclusions:

- Power factor correction is much more common in server power supplies than desktop power supplies, though it is becoming increasingly prevalent in high end desktop units
- Surprisingly, average server power supply efficiencies tend to be lower than average desktop power supply efficiencies in our data sets. Part of this is likely due to the recent response by

desktop manufacturers to the market opportunity presented by ENERGY STAR and efforts to push for efficiencies above ENERGY STAR levels (such as 80 Plus)[12] though both also extend to servers.

- The consequences of power supply oversizing remain significant, since power supply efficiency curves are rarely flat. Power supplies designed to operate routinely at 10 to 30% load are generally running about 12 to 15 percentage points less efficiently than if they were sized to operate routinely at 50 to 75% load.

2.3.3 Manufacturer-Reported Efficiency Data

In addition to conducting laboratory tests, we surveyed the manufacturer specification or datasheets for information about power supply efficiency. We discovered that many manufacturers do not report power supply efficiency, and those that do typically only report it at full load (i.e., 100%). Of more than 100 datasheets surveyed, only 71 reported power supply efficiency data. Table 4 provides a summary of the data collected.

Table 4. Manufacturer-Reported Efficiencies (by SSI Type)

SSI Type	# of Power Supplies Surveyed	Average Efficiency (at full load)
EPS 12V	17	67.7%
EPS 1U	26	70.7%
EPS 2U	13	70.1%
ERP12V	0	n/a
ERP2U	0	n/a
EPS	0	n/a
TPS	4	69.8%
MPS	4	65.0%
DPS	7	81.1%
All Power Supplies	71	69.6%

Figure 3 shows the average manufacturer-reported efficiencies by SSI category and how they compare to our laboratory test results at 100% loading. The manufacturer-reported data match quite closely with the power supply test results, likely indicating similarity in test method (at least at full load, which is the simplest way to test a power supply). The similarity also shows that manufacturers are not over-reporting the efficiency of their products. In fact, for all SSI Types except EPS 2U, the test results actually yielded slightly higher average efficiencies than the manufacturer-reported data. In the case of EPS 2U, the difference between the average manufacturer reported efficiency and tested efficiency was negligible. For all power supply categories, average manufacturer-reported efficiency at 100% load was 69.6% versus 72.5% from the tests results.

[12] 80 Plus is a program to encourage manufacturers to voluntary adopt more energy-efficient power supplies using utility rebates. See www.80plus.org for more details.

Figure 3. Manufacturer-Reported Efficiencies and Test Results at 100% Loading (by SSI Category)

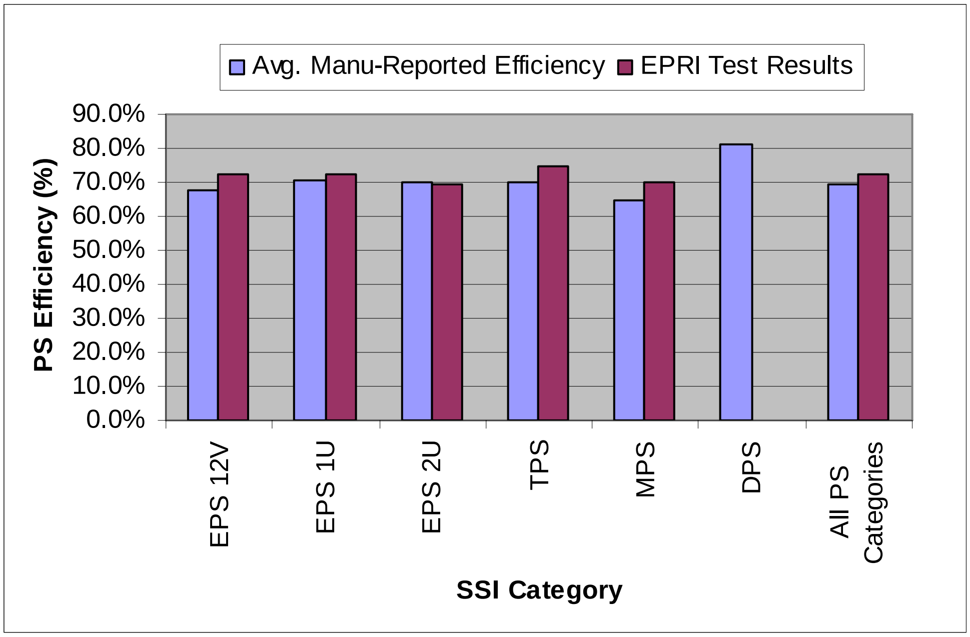

Note: EPRI Solutions did not test any DPS power supplies.

Nevertheless, as Figure 2 (Measured Power Supply Efficiency) makes clear, 100% loading is only one part of the story. Data centers virtually never operate server power supplies at 100% load, so their efficiency at that level is academically interesting, but has no relevance to total cost of ownership. It is critical to expand the consideration of server power supply efficiency to loads below 100% when studying overall energy use and opportunities for energy savings. The next section of this report (Section 2, Field Testing) reinforces this point by examining the typical loading conditions of operating servers.

3 Field Testing

This task's objective was to conduct field testing to determine achievable energy savings from more efficient power supplies in operating server installations. To do this we:

- Developed field-testing protocols.
- Conducted field tests at EPRI Solutions, Ecos Consulting, and LBNL facilities.

3.1 Field Testing Protocol

We developed a simple three-step test protocol/procedure for conducting *in situ* field tests to measure server power consumption.[13]

1. Before taking any measurements, some basic information (Table 5) was recorded about each server's function, technical specifications, power supply size and configuration:

[13] In some high-tech environments (e.g. datacenters) this protocol was not applicable and a customized approach was used to gather data directly from the equipment's log.

Table 5. Server and Power Supply Information

Server Information	Power Supply Information
Make	Number of power supplies
Model	Number of AC input cords
Application (e.g. print, file, e-mail)	Rating of each power supply (watts)
Processor	Power supply configuration (n, n+1)
Memory	Cross reference to SSI specification
Hard Drive	
Operating System	
Power Management enabled/disabled	
Other/Miscellaneous	

2. Each server was measured using a “WattsUp? Pro” power meter that recorded power consumption (in watts) for a continuous period of 15 to 144 hours at data intervals ranging from four seconds to 34 minutes.

3. The data collected from each server were then rank ordered and divided into nine percentiles (1%, 5%, 10%, 25%, 50%, 75%, 90%, 95%, and 99%). In this way, we could determine what level of power was being drawn by the server for what percent of the time.

3.2 EPRI, Ecos, and LBNL Field Tests

Initial field tests were conducted in EPRI Solutions and Ecos Consulting facilities located in Knoxville, Tennessee, and Portland, Oregon, respectively. Both EPRI Solutions and Ecos can be described as “small office” environments—approximately 50 people and 10,000 sq. feet each. The servers located at each of these facilities handle various routine business functions (e.g., print and file server, web server, Microsoft Exchange server, etc.), all of them operate on a 24/7 basis.

Field tests were also conducted at an LBNL data center located in Berkeley, California. The LBNL servers, although similar to some of the others tested in terms of power supply topology and configuration, are custom servers used primarily for distributed scientific computing. As such, these machines operate at close to 100% processor utilization all of the time and have a significantly different usage profile than the office servers tested at Ecos and EPRI Solutions.

3.2.1 Servers Tested

A total of 18 servers—six at EPRI Solutions, nine at Ecos, and three at LBNL—were measured using the field-testing protocol described above. Table 6 lists the servers that were tested and some information about their function and power supply configuration. Power management was not enabled on any of the servers tested.

Table 6. Servers Tested

Server Function	Processor Information	# of Power Supplies	# of AC Cords	Rating of Each Power Supply (W)	Redundancy Config	Cross Ref to SSI
Print& File Server	Single Intel Pentium 3 733MHz	3	3	320	N+2	ERP12V
Accounting Server	Dual XEON Processors @ 2GHz	2	1	500	N+1	ERP12V
Web & SQL Server	P3 ~ 700 MHz	1	1	330	N	EPS
Terminal Server	P3 933MHz	3	1	300	N+2	ERP12V
Exchange Server	Single Intel Pentium 3 - 933 MHz	2	2	330	N+1	ERP12V
PQ Remote Monitoring Server	Dual XEON (P3) 550MHz	3	1	275	N+1	ERP12V
Telephony Server	Intel PIII 700Mhz	2	2	300	N+1	ERP2U
Antivirus Server	Intel Dual PIII 745Mhz	1	1	150	N	EPS1U
Accounting/ADP Payroll Server	Intel Dual PIII 745Mhz	1	1	150	N	EPS1U
GoldMine Sync Server	Intel PIII 1.26Ghz	1	1	200	N	EPS1U
Print & File Server, Domain Controller	Intel Xeon 2Ghz	2	2	350	N+1	ERP2U
MS SQL Server	Intel Xeon 700Mhz	2	2	270	N+1	ERP2U
Web Application Server	Intel PIII 1.26Ghz	1	1	200	N	EPS1U
Mail Server	Intel P450Mhz	1	1	60	N	EPS12V
Mail Server	Intel P450Mhz	1	1	60	N	EPS12V
Scientific Computing	AMD Opteron 1.4 GHz	1	1	350	N	EPS
Scientific Computing	AMD Opteron 1.4 GHz	1	1	350	N	EPS
Scientific Computing	AMD Opteron 1.4 GHz	1	1	350	N	EPS

3.2.2 Server and Power Supply Loading Curves

A load duration curve for each server was constructed from the percentile data representing the percent of time the server power exceeded a given value in watts. We derived a single server load duration curve by averaging these results for the 18 different servers (Figure 4, blue line). Because of the relatively large difference in function, we also constructed separate single server load duration curves for the group of three LBNL servers and for the 15 EPRI Solutions and Ecos server. (Number of servers is shown in parentheses.) As Figure 4 shows, the LBNL servers consumed about 50 watts more, on average, than the EPRI and Ecos servers. As mentioned above, these machines are closer to a scientific mainframe than typical office servers, and this difference in intended use may explain the additional power consumption; however, it may be too early to draw any hard conclusions about this difference in load duration curves. A number of factors beyond the computers' application, including power supply rating, redundancy configuration, processor type, and server utilization could be responsible for these results.

Figure 4. Measured Load Duration of Servers

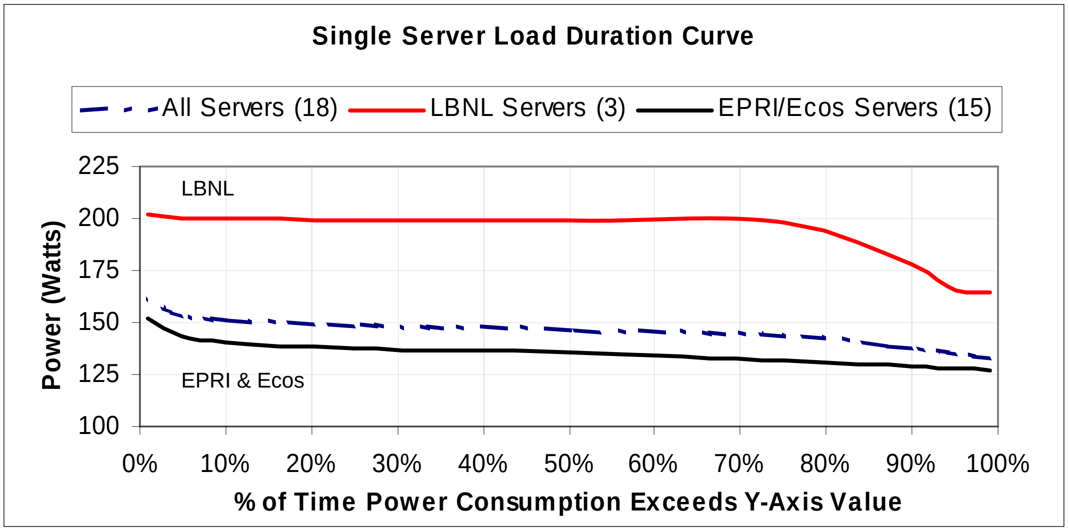

Source: EPRI/Ecos field measurements.

Figure 4 shows that the median (50th percentile) power consumption of a single server is 145.7 watts.[14] By annualizing the server power data, we can estimate total median server power consumption for a small office environment, like EPRI Solutions or Ecos, depending on the number of servers found. For example, Table 7 shows that, a five-server office would consume, on average, 6,384 kWh per year.[15]

Table 7. Estimated Annual Energy Consumption of Servers

# of Servers	Power Draw (W)	kWh per Year
1	145.7	1,277
2	291.5	2,553
3	437.2	3,830
4	583.0	5,107
5	728.7	6,384

[14] This is close to the 125W used to estimate server energy consumption in Roth, Kurt W., Fred Goldstein, and Jonathan Kleinman (Arthur D. Little). "Office Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings. Volume I: Energy Consumption Baseline." Prepared for US DOE, Building Technologies Program, January 2002. NTIS Number: PB2002-101438.

[15] While not included in this calculation, EPRI-Solutions found that the median power consumption for all the other network equipment (e.g., router, network interface card, etc.) needed to support their six-server IT infrastructure was about 150W, or 1,314 kWh per year—slightly more than adding an additional server.

Using the load duration data, the known power supply rating and configuration of each server, and an average of power supply efficiency, we were also able to generate a power supply loading curve showing the percent of time the server exceeded a given percentage of the power supply rating in watts, using the following formula:[16]

$$\%\ of\ Rated\ PS\ Watts = \frac{Watts * 70\%}{PSRating * \#ofPS}.$$

Figure 5 shows the average power supply loading for servers with different redundancy configurations. From our small sample it appears that server power supplies for small office applications are normally loaded in the 15%-35% range. This highlights the importance of light load power supply efficiency, or the need to better match power supplies with their load. Light load efficiency is especially important for redundant systems. The median power supply loading for non-redundant servers (N configuration) is just over 30%; whereas, for redundant power supply configurations—N+1 and N+2—the median loading was only 18.7% and 15.3%, respectively.

Figure 5. Power Supply Loading and Redundancy

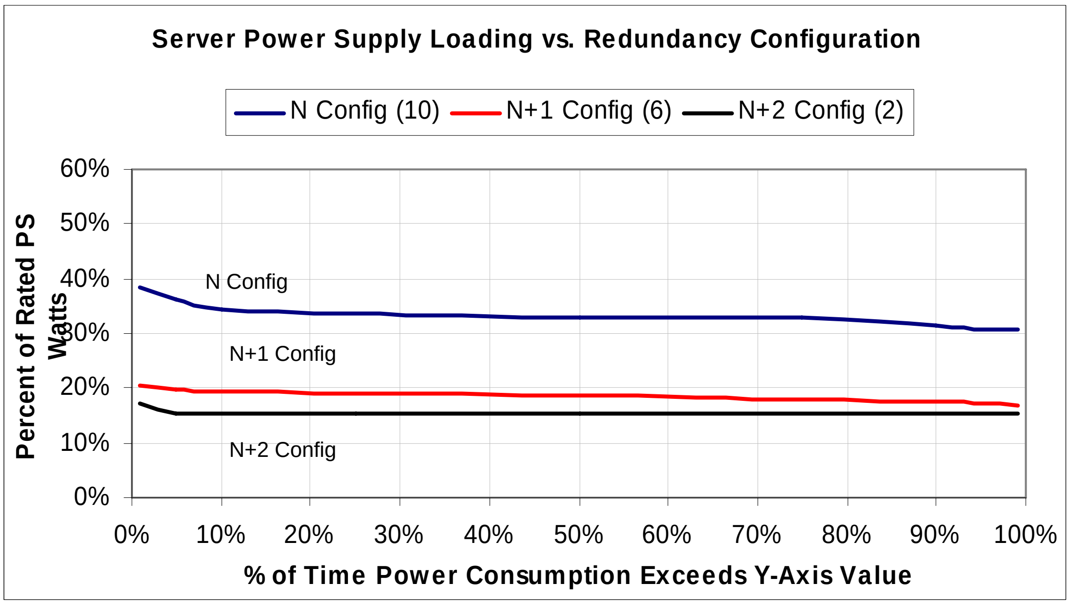

Source: EPRI/Ecos field measurements.

Figure 5 also shows that the loading curves are relatively flat, indicating not only that the server is drawing power or operating 100% of the time, but that there is little variation in the percent loading of its power supply during that time.

To investigate whether server function played any role in power supply loading, we compared the three LBNL servers to seven EPRI/Ecos servers with the same power supply redundancy configuration (N). Figure 6 shows that the LBNL servers were at the upper end, but were not substantially different from the other servers in terms of power supply loading. This suggests that even highly utilized machines like the LBNL servers/mainframes do not fully utilize the capacity of their power supplies. It also suggests that server power supplies are oversized for the actual requirements of the machines in which they are being used. Most servers measured were operating somewhere between about 20% and 50% of their rated load all of the time, which is the exact portion

[16] This formula multiplies the watts recorded by the power meter by the power supply efficiency, 70% or 0.7, to account for 30% energy loss, which is "consumed" by the power supply. It is only an approximation, as power supply efficiency varies with load.

of power supply curves at which efficiency begins to dramatically decrease (see the previous section for more on this topic). None of the servers tested ever exceeded 50% of its rated output.

Figure 6. Loading of Server Power Supplies

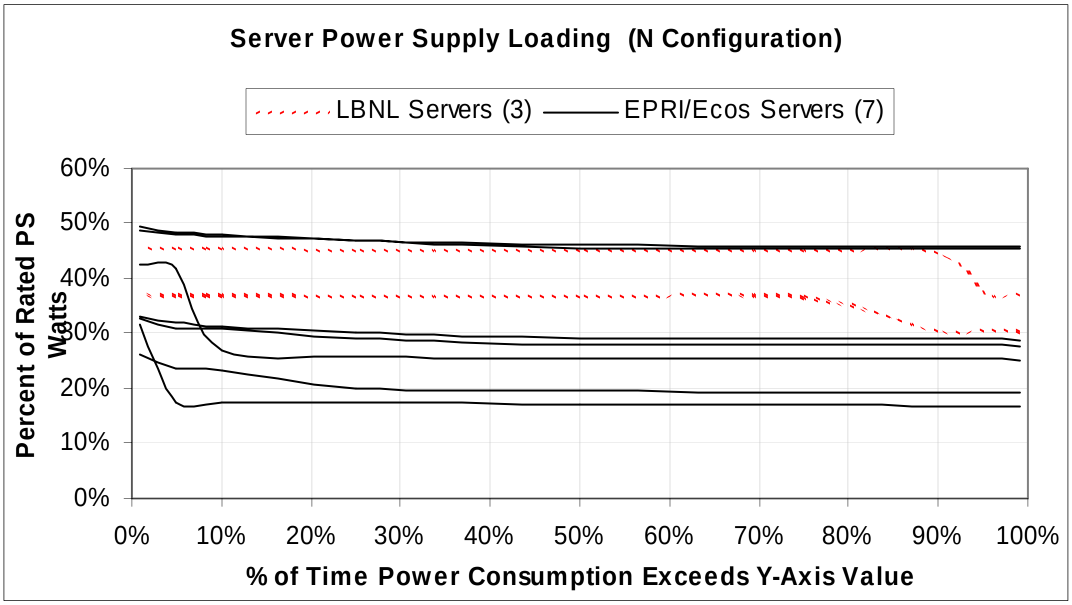

Source: EPRI/Ecos field measurements.

The combination of low efficiencies in, and oversizing of server power supplies can have dramatic effects on the net energy consumption of those power supplies. As Figure 6 illustrates, a dual Xeon processor-based server might typically require 100 to 300 watts DC from its power supply. An oversized, inefficient power supply would often waste two to three times as much net AC power[17] to meet that load as a properly sized, efficient one.

3.2.3 CPU Utilization

In addition to measuring power consumption, we also extracted CPU utilization data from each server's log for a 24-hour period.[18] CPU utilization describes how busy a server's processor or processors are; it tells us the percent of the CPU's processing capacity, or throughput, that is being used to perform tasks.

The utilization data were divided into percentiles, using the same methodology as described for the power data. A CPU utilization curve for each server[19] was constructed representing the percent of time the server's CPU utilization exceeded a given percent.

For comparison, Figure 7 shows the CPU utilization of the LBNL servers (data center/scientific computing) was very different from the Ecos servers (office environment). Average CPU utilization is a remarkably low 5 to 6%.[20] In fact, more than 95%% of the time, these servers operate at less than 30% utilization. In contrast, the LBNL servers spend a majority of their time operating at, or near, full capacity. It is difficult to draw any conclusions from such a small sample, but it appears that for the

[17] Here net AC power is defined as the difference between the AC power input of the power supply and its DC power output.
[18] Note: the time period for which CPU utilization data was collected was not necessarily the same as the time period for which power consumption was measured.
[19] CPU utilization data was collected for 7 Ecos servers, 5 EPRI Solutions servers, and three LBNL servers; three of the Ecos servers were not included.
[20] This median number is somewhat misleading. It could be (and is likely) the case that an organization with multiple servers relies heavily on one or two servers which have very high utilization, but the rest are performing secondary or backup functions and are idle much of the time. By averaging all servers together, we have masked this phenomenon.

most part power supply loading (which was relatively flat, see Figures 5 and 6) has very low correlation with "server activity."

Figure 7. Server Utilization

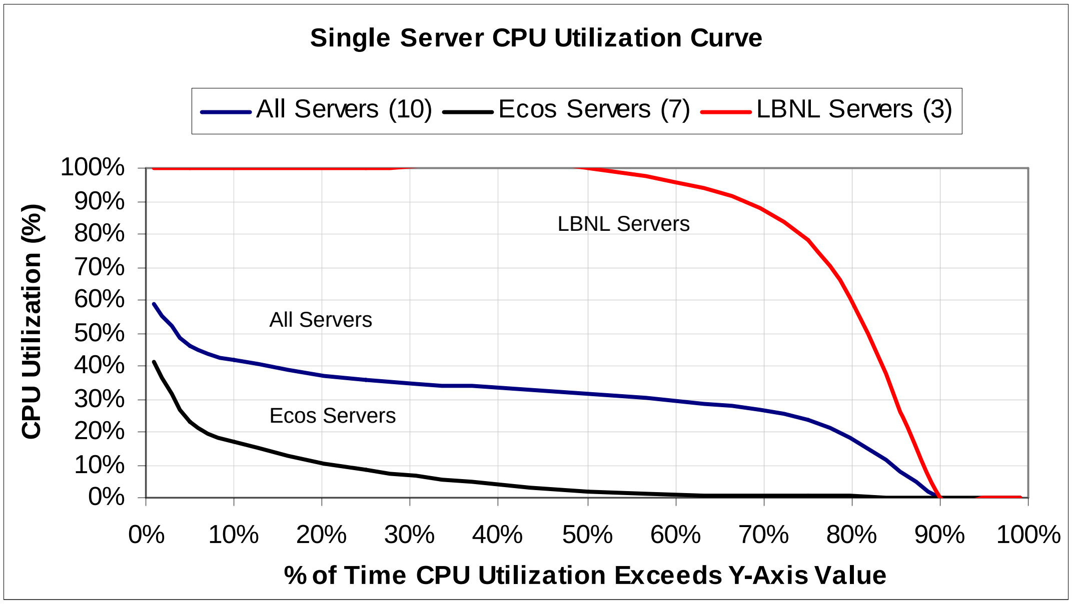

Source: EPRI/Ecos field measurements.

The servers we measured do not exhibit any activity-based power management or demand based switching (DBS), as Intel calls it, whereby power consumption is reduced when the server is operating below full capacity. Activity-based power management attempts to scale processor power consumption to the demand for processing resources placed on the CPU. The absence of this technology in the servers that we measured might help to explain the lack of any correlation between the utilization curve of the CPU and the load duration curves of the power supplies. Figure 8 provides a side-by-side comparison of processor utilization and power consumption in a server, demonstrating the poor correlation between power use and processor use.

Figure 8. AC Power Input vs. % CPU Time

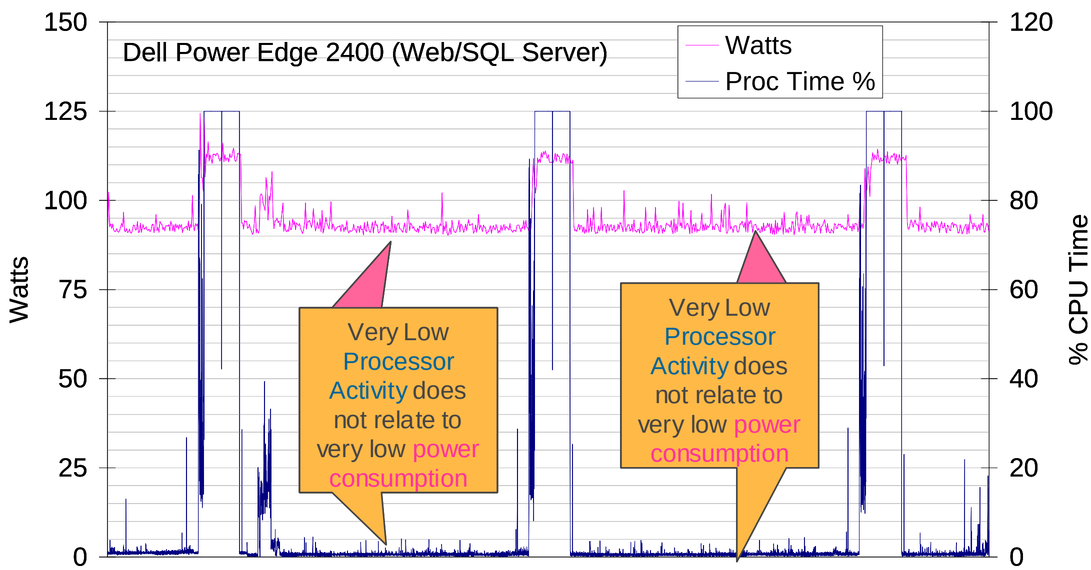

Source: EPRI field measurements.

4 Proposed Power Supply Efficiency Specifications

This task's objective was to recommend new efficiency levels to SSI for consideration in its future specifications. To do this we:

- Coordinated with SSI on recommended new efficiency levels in power supply guidelines.
- Performed market research to assess the market penetration of different server configuration and power supply topologies.
- Estimated the energy savings potential by combining market research with the proposed efficiency levels.

4.1 Coordination with SSI

Based on the efficiency levels documented in existing equipment and the field studies conducted, we closely coordinated with SSI and Intel[21] to propose changes and to the SSI specifications based on the *Proposed Test Protocol.*

At the Intel Developer Forum in March 2005, in a joint presentation, Intel announced new proposed SSI Industry specifications to encourage more efficient power supplies. The new specifications include testing conditions which align with ATX12V, as well as required and recommended efficiency levels at 20%, 50%, and 100% loading. SSI is also now supporting a power supply technology known as Power Supply Management Interface (PSMI). This new industry standard provides a basic internal mechanism for server power supplies to report power consumption and efficiency data directly to the server's motherboard with 5% to 10% accuracy. The technology could conceivably be used for a number of purposes. Power and efficiency data could be logged on the server's hard drive or reported to HVAC equipment so that fans and cooling equipment could scale their output to the heat output of IT equipment.

Table 8 provides a comparison of the old and new SSI specifications with regards to efficiency requirements. Establishing high minimum recommended levels at 20% and 50% loading is particularly critical for improving energy efficiency, since field-testing (Section 2) showed that servers operate in this range most of the time.

[21] Brian Griffith, Intel Power Server Architect, EPG (SSI Coordinator).

Table 8. Comparison of Old and New SSI Specifications

SSI Category		Old Specification	New Specification		
EPS12V		**Rev 2.1**	**Rev 2.8**		
	Loading	**100%**	**100%**	**50%**	**20%**
	Required Efficiency		70%	72%	65%
	Recommended Efficiency	68-72%	77%	80%	75%
ERP12V		**Rev 1.0**	**Draft Rev 1.4**		
	Loading	**100%**	**100%**	**50%**	**20%**
	Required Efficiency		70%	72%	65%
	Recommended Efficiency	68-70%	77%	80%	75%
EPS1U		**Rev 2.1**	**Rev 2.9**		
	Loading	**100%**	**100%**	**50%**	**20%**
	Required Efficiency		70%	72%	65%
	Recommended Efficiency	65-75%	80%	83%	78%
EPS2U		**Rev 2.1**			
	Loading	**100%**	**100%**	**50%**	**20%**
	Required Efficiency				
	Recommended Efficiency	68-72%			
ERP2U		**Rev 2.0**	**Rev 2.2**		
	Loading	**100%**	**100%**	**50%**	**20%**
	Required Efficiency		70%	72%	65%
	Recommended Efficiency	70-82%)	80%	83%	78%

Note: Recommended minimum efficiency in old SSI specification depended upon PS rating (watts).

4.2 Market Penetration of Servers and Power Supplies

We performed market research, through literature review and other industry sources, to assess the market penetration of different server configurations and their power supply types.

This research was performed in two stages:

1. First, we gathered market data on servers to determine which types of configurations and manufacturers had the largest market penetration.
2. Next, we gathered information on the power supplies contained in those units.
3. With respect to California-specific data, this information was not available on any consistent basis, as industry data are not reported for any one state. Thus, we have not attempted to breakout any CA-specific estimates of product shipment or server population. We have incorporated approximations for California where specific estimates may make sense, such as energy savings projection.

4.2.1 Servers

The server market can be sliced many different ways (e.g., by shipments, revenue, geographic region, class, manufacturer, price range, operating system, processor, U-rating,[22] etc.) IDC's Worldwide Quarterly Server Tracker and Forecast tools are probably the most comprehensive data source for statistics on server shipments. The Tracker provides actual quarterly data for more than 15

[22] The "U" rating is based on the height of the server: A 1U server is 1.75-in. tall.

different data categories (including those mentioned above). The Forecast provides fewer data categories, but projects future unit shipments and revenues.[23]

Such data can be purchased from IDC at a substantial cost. However, we felt we could obtain a fairly accurate picture of the U.S. server market relying on limited data from IDC and various other quantitative and anecdotal sources all of which seemed to corroborate.

Worldwide Server Shipments

IDC reported actual 2003 unit shipments of servers of 5.3 million (Table 9), almost a 19% increase over 2002 shipments of 4.4 million.[24]

Table 9. Worldwide Server Market, 1999-2003

Worldwide Server Factory Revenue ($M) and Units Shipped					
	2003	2002	2001	2000	1999
	$M	$M	$M	$M	$M
WW Total	$46,131	$44,649	$50,496	$61,675	$57,708
Units	5,281,231	4,442,690	4,276,119	4,369,840	3,761,141

Source: IDC, 2004 Release as shown in *The Business for Storage Networks. Chapter 1, Industry Landscape: Storage Costs and Consumption*. Cisco Systems, Sept 3, 2004. Available at:
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns516/c1272/cdccont_0900aecd80257124.pdf

IDC expects worldwide server shipments to grow steadily through 2008 (Figure 9).

Figure 9. Historical and Forecast Server Shipments, 2000-2008

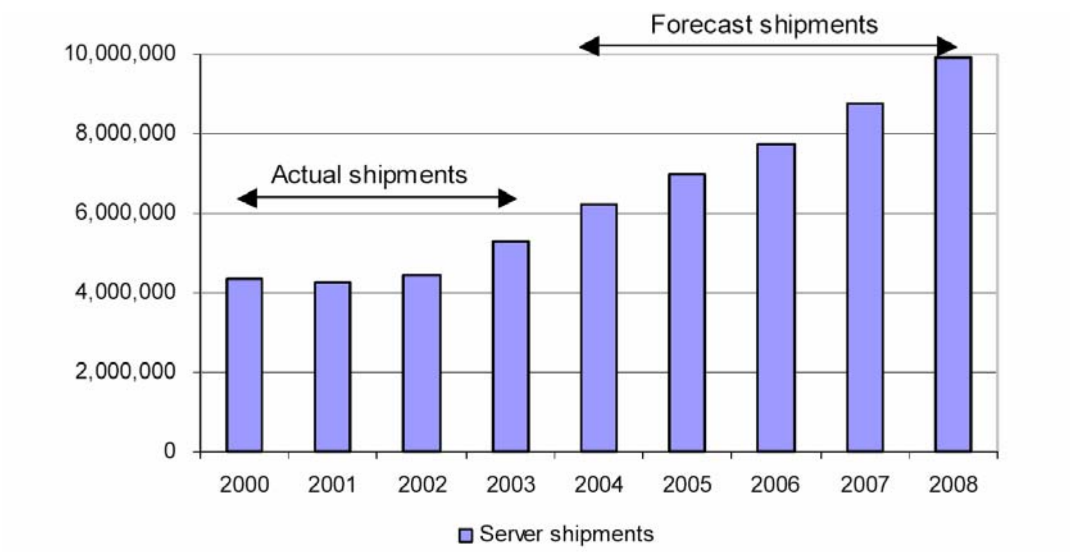

Source: IDC, 2004 as shown in Humpreys, John and Jessica Yang, IDC White Paper sponsored by Rackable Systems. Server Innovations: Examining DC Power as an Alternative for Increasing Data Center Efficiency and Reliability, August 2004

[23] http://www.idc.com./getdoc.jsp?containerId=IDC_P348 and http://www.idc.com./getdoc.jsp?containerId=IDC_P5514.
[24] Gartner Group, another well-know data source, reports similar 2003 shipments of 5.4 million.

U.S. Server Shipments

The United States represents about 41% of the worldwide market in terms of shipments (Table 10). On a population basis, the State of California would account for about 10% to 15% of the overall US shipment of products.

Table 10. US Percentage of Worldwide Server Shipments

Year	Worldwide Shipments (Units)	US Shipments (Units)	US %
2000	4,327,511	1,882,184	43%
2001	4,425,977	1,712,614	39%
2002	4,610,328	1,949,361	42%
2003	5,469,016	2,256,918	41%
2004	6,700,000	2,747,000	41% est.
Average			**41%**

Source: Gartner Group/Dataqwest, with some interpolation by Ecos.
2004 sales from http://uk.news.yahoo.com/050225/221/fd8j9.html

The dominant manufacturers in the U.S. market are Hewlett-Packard, Dell, and IBM, and to a lesser extent Sun Microsystems and Gateway. These five companies represent over 70% of the U.S. market in terms of units shipped (Table 11).

Table 11. U.S. Server Shipments (Units), 2001-2002

Company	2002 Shipments	2002 Market Share (%)	2001 Shipments	2001 Market Share (%)	Growth (%)
Hewlett-Packard	506,589	26	483,938	28.3	4.7
Dell	487,984	25	399,236	23.3	22.2
IBM	225,315	11.6	217,171	12.7	3.8
Sun Microsystems	143,753	7.4	124,056	7.2	15.9
Gateway	18,000	0.9	25,803	1.5	-30.2
Others	567,720	29.1	462,410	27	22.8
Total	1,949,361	100	1,712,614	100	13.8

Note: The data for Hewlett-Packard and Compaq have been combined.
Source: Gartner Dataquest (January 2003)

Volume Servers

Many top manufacturers, such as Dell, emphasize volume or low-end servers. This emphasis parallels recent server sales trends, which reveal that severs priced under $25,000 are the market's bright spot.[25] Lower-end units are becoming ever more capable. "An entry-level box now can do what it took a four-way to do," according to Meta Group analyst Carl Greiner.[26] Manufacturers like these servers not only because their low price point helps boost volumes of sales with small to medium sized business, but also because these servers are increasingly turning up in more high-tech environments, like data centers.

A number of manufacturers are relying on the entry-level server market to drive their overall server business and boost volume sales. For example, in September 2004, Sun Microsystems indicated that the company is renewing its commitment to the low-end server market and increasing its “portfolio of

[25] http://www.newsfactor.com/story.xhtml?story_id=21934
[26] http://www.newsfactor.com/story.xhtml?story_id=21934

offerings in the x86 segment."[27] Volume or low-end servers are thought to be the primary driver of market growth, as indicated by the following market reports:[28]

- In 2002, Gartner Group reported that the U.S. server market experienced a strong "resurgence," with server shipments increasing by 13.8 percent to a total of 1.9 million units. Much of this growth was attributed to low-end server sales.[29]
- In the third quarter of 2003, Vernon Turner, group vice president of IDC's Worldwide Server Group indicated, "Volume servers [those priced less than $25,000] are generating most of the positive momentum in the worldwide server market. This shows that the IT community has embraced volume server deployments as a mainstream technology to meet a wide range of data-processing requirements and to support a wide variety of computing workloads."[30]
- IDC data for the second quarter of 2004 indicated 22.7% year-over-year unit shipment growth for servers, reflecting "strong unit growth in the volume server segment."[31]

Volume or low-end servers also represent over 90% of the market.

- Roth et al. (2002) showed that from 1998 to 2000 low-end servers represented between 88% and 91% of the U.S. market in terms of units shipped (Table 12).

Table 12. U.S. Server Shipments by Class, 1998-2000

Year	Low-end	Work-horse	Mid-range	High-end	Total	Low-end%
1998	1,082,180	104,776	37,813	2,852	1,227,621	88%
1999	1,367,839	119,641	40,340	2,663	1,530,483	89%
2000	1,615,126	121,097	41,314	2,510	1,780,047	91%

Source: ADL/US DOE 2002.

- Gartner Group reported that in the last quarter of 2003, inexpensive dual processor servers containing 32-bit processors from Intel or Advanced Micro Devices accounted for 1.45 million or 91% or of the 1.59 million servers shipped worldwide.[32] IDC data for the same quarter matched closely, reporting that unit shipments of basic x86 servers grew at 23 percent to nearly 1.4 million servers worldwide, with factory revenues growing at 15 percent to $5.5 billion.[33]
- IDC also indicated that single and dual processor capacity servers represented 91% of all server unit shipments in 2003. IDC expects two-processor capacity server unit shipments to account for more than 6.2 million units in 2008, nearly double the 3.5 millions units shipped in 2003.[34]
- In 2004, x86 servers dominated shipments, with 91%, and at current growth rates, the X86 server business will account for over 50% of all server (dollar) sales in 2005.[35]

Rack-Optimized vs. Pedestal Servers

Presently, there is a fairly even split between rack-optimized and non-rack optimized (pedestal) servers, with blade servers accounting for a small fraction of the market (Figure 10). In the future, the number of pedestal servers is expected to remain stable, while the number of rack-optimized, and especially blade servers, grows.[36]

[27] http://www.expresscomputeronline.com/20040906/coverstory01.shtml
[28] http://nwc.serverpipeline.com/midrange/46200036
[29] http://www.serverwatch.com/news/article.php/1575051
[30] http://www.prnewswire.com/cgi-bin/stories.pl?ACCT=105&STORY=/www/story/11-26-2003/0002065610
[31] http://www.tekrati.com/T2/Analyst_Research/ResearchAnnouncementsDetails.asp?Newsid=3460
[32] http://news.com.com/2100-1010_3-5149716.html?tag=nefd_top
[33] Bekker, Scott. "IDC: Volume Servers Shaking Up Server Market." ENTNews, March 3, 2004. http://www.entmag.com/news/article.asp?EditorialsID=6151
[34] http://www.internetnews.com/ent-news/print.php/3389211
[35] http://uk.news.yahoo.com/050225/221/fd8j9.html (Yahoo News, UK, Feb 25, 2005)
[36] Note: Blade servers were excluded from our study, as these servers are still relatively new and **currently** represent only a small portion of the market. Nevertheless, blade server technology should be closely watched as it represents an emerging trend in IT and datacenter applications. (See Blade Servers.)

Figure 10. Historical and Forecast Adoption of Pedestal, Rack-Optimized and Blade Servers, 2000-2008

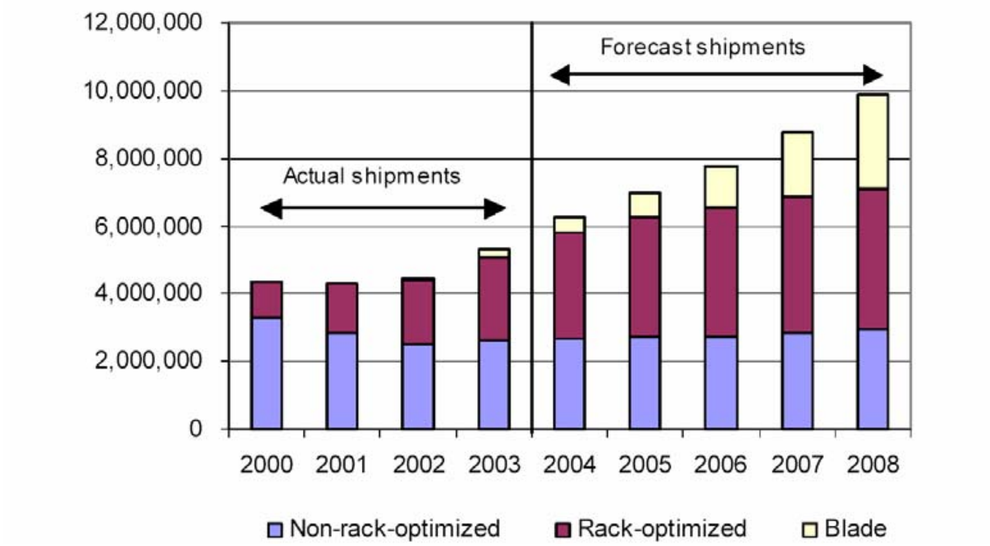

Source: IDC, 2004 as shown in Humpreys, John and Jessica Yang, IDC White Paper sponsored by Rackable Systems. Server Innovations: Examining DC Power as an Alternative for Increasing Data Center Efficiency and Reliability, August 2004

IDC tracks the U-rating[37] of servers. Table 13 shows a break down of the servers shipped in the U.S. in the forth quarter of 2001 by U-rating. 2U, 1U, and 0U (tower/pedestal) servers each account for a little over 20% of the market.

Table 13. U.S. Server Shipments by U-Rating, Q4 2001

Server U-rating	Units (Q401)	%
2	213,558	23.05%
1	209,882	22.65%
0	199,938	21.58%
5	138,040	14.90%
4	79,066	8.53%
7	27,760	3.00%
6	9,610	1.04%
All Others	48,598	5.25%
Total	**926,452**	**100.00%**

Source: IDC, Worldwide Quarterly Server Tracker, demo – additional data can be purchased from IDC.

Historically, large "mainframe" computers dominated the datacenter. Now, even though some high-end models like IBM's zSeries and the HP Superdome are selling quite well, there is a trend toward using smaller servers, including blades.[38] At a recent APC event in Portland, Oregon, one presenter indicated that 5 years ago, the average server in a data center was a 5U, 1 processor machine, whereas, today, the average server is 1.5U, dual processor with at least 2 gigabytes of memory.[39]

[37] The "U" rating is based on the height of the server: A 1U server is 1.75-in. tall.
[38] http://www.cioupdate.com/trends/article.php/3289721
[39] APC Presentation, Multnomah Athletic Club, Portland, OR, July 13, 2004.

Blade Servers

Many analysts believe that blade servers are poised to transform the server industry. According to an article entitled “The Data Center of the Future” by Drew Robb:

> Blade servers offer companies low cost scalability, since it is easy to assign a batch of these servers to a particular application, rather than having to buy a more expensive server which then remains underutilized. Since blade servers stuff a dozen or more servers into a single box, they drastically cut the infrastructure costs for racks, cabling and cooling. Then there is the ease of support. When one goes down it is a simple act to swap out a server card and let the system automatically rebuild.[40]

Roughly 185,000 blade servers were sold in 2003.[41] However, IDC predicts blade servers will account for one out of every four servers sold by 2007.[42]

In fact, IDC forecasts that blade server sales will reach over $1 billion in 2004 on the strength of Dell's new blade offerings and interest from small and medium-sized businesses.[43] The market is fluid, and many manufacturers are revamping their blade product lines.

Of the large vendors, HP was the first to offer (Proliant) blade servers. However, second-quarter 2004 numbers from IDC, reveal that HP is now in second place, shipping 32% of units, while IBM leads the market with 44% of unit shipments.[44] IBM plans to expand its blade offerings and began shipping a new 7U, 14 blade BladeCenter Express Chassis at the end of November 2004. Dell also recently moved forward with its most significant modular server—the PowerEdge 1855, which supports up to 10 servers in a 7U chassis and is geared for companies that rely on large Web farms and high-performance computing clusters.[45]

While Bruce Kornfeld, Director of Enterprise Marketing at Dell, believes that customers are looking at blades for “great server technology, better density and savings in power and cooling to ease management,” many are still “not willing to pay a premium for blade servers over traditional 1U and 2U rack servers."[46] The implications of blade server growth for the power supply market are significant. It will become increasingly common in the future to see one or two high wattage (1000 watts or more) blade server power supplies powering up to six blades, rather than a larger number of smaller wattage power supplies powering individual servers.[47]

4.2.2 Power Supplies

While market data on server shipments are readily available, data on the power supplies contained within those servers are more difficult to find. We sampled some popular server models to get an idea of the size, shape, and configuration of the power supplies they contain. Table 14 provides the power supply information listed for 17 different popular new product offerings.

[40] http://www.cioupdate.com/trends/article.php/3289721
[41] http://www.itfacts.biz/index.php?id=P1517
[42] http://www.serverwatch.com/news/article.php/3411471
[43] http://www.serverwatch.com/news/article.php/3440321
[44] http://www.serverwatch.com/news/article.php/3411471
[45] http://www.serverwatch.com/news/article.php/3436331
[46] http://www.serverwatch.com/news/article.php/3436331
[47] Jeremiah P. Bryant, “AC-DC Power Supply Growth Variation in China and North America,” presented at the Applied Power Electronics Conference (APEC) 2005, Darnell Group, March 6-10, 2005.

Table 14. Power Supply Specifications in Selected Popular Server Models

Make	Sever Model	Form Factor	Power Supply Description
Dell	Poweredge 750	5U Tower or Rack	650W non-redundant or 675W hot-plug redundant power 110/220 Volts
Dell	Poweredge 1750	1U	Optional, hot plug redundant 325 W power supplies, 110/220 Volts, Optional -48V DC power supplies
Dell	Poweredge 1800	5U Tower or Rack	650W non-redundant or 675W hot-plug redundant power 110/220 Volts
Dell	Poweredge 1850	1U	550W, optional hot-plug redundant power, 110/220 Volts
Dell	Poweredge 2800	5U Tower or Rack	930W, optional hot-plug redundant power, 110/220 Volts
Dell	Poweredge 2850	2U	700W, optional hot-plug redundant power, 110/220 Volts
HP	Proliant DL360G4	1U	Optional hot plug redundant power supply (460 W)
HP	Proliant DL380G4	2U	575-Watt CE Mark Compliant Optional Hot Plug AC Redundant Power Supply
HP	Proliant DL585	2U	Optional hot plug redundant power supply (800 W)
HP	Proliant ML110	5U Tower or Rack	350W power supply
HP	Proliant ML350G4	5U Tower or Rack	Most models: 725W, optional hot-plug redundant, NHP SCSI models: 460W, non hot plug supply
IBM	eServer 326	1U	411W 1 std/1 max
IBM	Xseries 336	1U	Hot-swappable 585W power supplies, 1 std/2 max
IBM	Xseries 346	2U	Hot-swappable 625W power supplies, 1 std/2 max
IBM	Xseries 306	1U	300W 1 std/1 max, 110 or 220 volt universal auto sensing
Sun	SunFire V210	1U	320W power supply
Sun	SunFire V240	2U	One required, two for redundancy (hot-swappable) with separate power cords (400W)

Source: PC Magazine and technical specifications available at www.dell.com, www.ibm.com, www.hp.com.

This information tracked closely with the power supply types and configurations that we observed in our field testing (See Sections 2.2 and 2.3), and reinforced the decision—made in consultation with Intel, SSI members and other industry partners—to focus this study on multi-output front-end AC-DC power supplies commonly found in 1U, 2U, and pedestal servers (non-redundant PS).

4.3 Energy Savings Potential of Proposed Efficiency Levels

We estimated the energy savings potential of more efficient server power supplies in three steps:

1. We estimated the current AEC (Annual Energy Consumption) of servers in the United States using the basic methodology developed in the ADL/US DOE report "Office Energy

Consumption by Office and Telecommunications Equipment in Commercial Buildings. Volume I: Energy Consumption Baseline."[48]

2. We determined the percentage of current server energy use that could be saved due to more efficient power supplies, based on the new SSI recommended specifications.
3. By applying the percentages established in #2 to the AEC estimates established in #1, we were able to estimate the overall energy savings potential of the new recommended efficiency levels.

For California-specific estimates, we used a range of 10% to 15% of estimated US results.

4.3.1 Annual Energy Consumption of Servers

Starting from the basic methodology established in the ADL/US DOE study, we constructed a revised estimate of the annual energy consumption (AEC) of servers in the United States.

To construct the revised 2004 EPRI/Ecos estimate, we:

- Updated the server stock numbers by adding four additional years of market data (2001-2004) on numbers of units shipped.

Table 15. Estimated Server Units

Server Categories	2000 ADL/US DOE Stock (Units Millions)	2004 EPRI/Ecos Stock (Units Millions)	% Change
Low-end	4,065,145	6,587,061	62%
Work-horse	577,960	506,470	-12%
Mid-range	185,195	151,678	-18%
High-end	16,549	14,730	-11%
Total	**4,844,849**	**7,259,939**	**50%**

- Revised the power draw for low-end and workhorse servers based on a sample of current data gathered from manufacturers, as well as information gathered during laboratory and field tests about typical power supply efficiency and server loading.

Table 16. Estimated PS Characteristics

Server Categories	2000 ADL/US DOE		2004 EPRI/Ecos	
	Max Power Draw (Watts)	Avg Power Draw (Watts)	Max Power Draw (Watts)	Avg Power Draw (Watts)
Low-end	250	125	450	193
Work-horse	1,300	650	800	343
Mid-range	2,450	1,225	2,450	1,225
High-end	5,040	2,520	5,040	2,520

Notes: For 2000 ADL/US DOE average power draw was calculated as 50% of maximum. For 2004 EPRI/Ecos, maximum power draw for Low-end and Workhorse was revised based on sample of manufacturer data. Average power draw was also revised to reflect 30% server loading, and 70% power supply efficiency.

Figure 11 shows our revised estimate of server AEC for 2004. The revised AEC is 14.6 TWh, which represents a 45% increase over the ADL/US DOE estimate of 10.1 TWh. This increase is attributable to growth in both the number and average power draw of low-end servers, which comprise the majority of units. The estimate for California is in the range of 1.5 TWh to 2.2 TWh.

[48] Roth, Kurt W., Fred Goldstein, and Jonathan Kleinman (Arthur D. Little). "Office Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings. Volume I: Energy Consumption Baseline." Prepared for US DOE, Building Technologies Program, January 2002. NTIS Number: PB2002-101438.

Figure 11. Server Annual Energy Consumption (AEC)

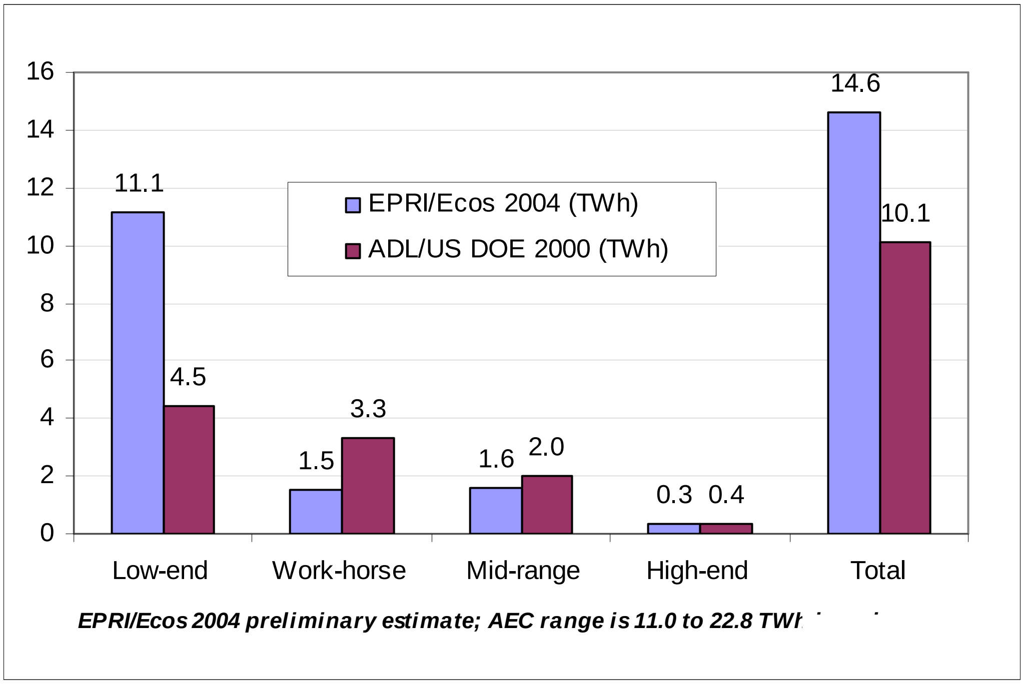

4.3.2 Existing vs. Proposed Power Supply Efficiency Levels

Table 17 compares the existing and recommended efficiencies at different loading points for the EPS 12V and EPS 1U power supplies. The SSI Recommended (2005) efficiency levels, recently endorsed by ENERGY STAR in its first draft specification revision for servers, yield significant efficiency increases, depending upon loading.[49] Likewise, a utility-funded efficiency program currently being operated by Ecos Consulting for desktop and server power supplies, 80 Plus, improves existing low end server power supply efficiencies by at least 14 to 31% at the 20% load condition where server power supplies often operate.[50]

Table 17. Comparison of Existing and Proposed Efficiency

Specification		20% Load	50% Load	100% Load
EPS12V (pedestal)	Existing Server	70%	76%	73%
	SSI Required (2005 proposed)	65%	72%	70%
	SSI Recommended (2005 proposed)	75%	80%	77%
	80 Plus	80%	80%	80%
EPS1U (rack)	Existing Server	61%	71%	72%
	SSI Required (2005 proposed)	65%	72%	70%
	SSI Recommended (2005 proposed)	78%	83%	80%
	80 Plus	80%	80%	80%

[49] See http://www.energystar.gov/index.cfm?c=revisions.computer_spec.
[50] See www.80plus.org for additional information.

4.3.3 Energy Savings Estimates

Table 18 provides an example to demonstrate the energy use of single EPS1U server power supply under various specification scenarios. As the table shows, implementing the SSI Recommended (2005) specifications would result in more than a 50% decrease in power supply electricity use. (Note that this estimate in energy use reduction is based on both utilization data as well as efficiency data showing that power supplies typically spend most of their time at lower load levels, where they are less efficient).

Table 18. Power Supply Energy Use for Single 425W EPS 1U Server

				Power Supply Electricity Use per year	
Specification	**20% Load**	**50% Load**	**100% Load**	**kWh**	**$**
Typical Server	61%	71%	72%	556	$56
SSI Required (2005 proposed)	65%	72%	70%	521	$52
SSI Recommended (2005 proposed)	78%	83%	80%	274	$27
80 Plus	80%	80%	80%	216	$22

Note: Assumes loading in the 20-50% range for a 425W EPS1U power supply and $0.10 per kWh.

Power supply efficiency measures the ratio of output power to input power. Thus, a 70% efficient power supply has energy losses of 30% and must draw more power than actually needed to run a server. Improving power supply efficiency reduces energy losses or the energy "consumed" by the power supply itself. *For example, replacing a 70% efficient power supply with an 85% efficient power supply cuts losses in half (from 30% to 15%)*.

Table 19 shows the percent energy savings for (a) the power supply and (b) the whole server for different efficiency gains.

Table 19.

(a) Reduction in Net Power Supply Energy Use

		New Power Supply Efficiency							
		60%	**65%**	**70%**	**75%**	**78%**	**80%**	**83%**	**85%**
Existing Power Supply Efficiency	**60%**	0%	19%	36%	50%	58%	63%	69%	74%
	65%	-24%	0%	20%	38%	48%	54%	62%	67%
	70%	-56%	-26%	0%	22%	34%	42%	52%	59%
	75%	-100%	-62%	-29%	0%	15%	25%	39%	47%
	78%	-136%	-91%	-52%	-18%	0%	11%	27%	37%
	80%	-167%	-115%	-71%	-33%	-13%	0%	18%	29%
	83%	-225%	-163%	-109%	-63%	-38%	-22%	0%	14%
	85%	-278%	-205%	-143%	-89%	-60%	-42%	-16%	0%

(b) Reduction in Server Energy Use

Existing Power Supply Efficiency	New Power Supply Efficiency							
	60%	65%	70%	75%	78%	80%	83%	85%
60%	0%	8%	14%	20%	23%	25%	28%	29%
65%	-8%	0%	7%	13%	17%	19%	22%	24%
70%	-17%	-8%	0%	7%	10%	13%	16%	18%
75%	-25%	-15%	-7%	0%	4%	6%	10%	12%
78%	-30%	-20%	-11%	-4%	0%	2%	6%	8%
80%	-33%	-23%	-14%	-7%	-3%	0%	4%	6%
83%	-38%	-28%	-19%	-11%	-6%	-4%	0%	2%
85%	-42%	-31%	-21%	-13%	-9%	-6%	-2%	0%

Assuming a baseline efficiency of 70%, Tables 19 (a) and (b) indicate a 34% reduction in power supply electricity use, and a 10% reduction in overall server energy use. Our comparison of existing and proposed efficiency (Table 20) yields, on average, a power supply efficiency increase of 8%.[51]

Based on our overall estimate of server AEC (Section 3.3.1), we estimate the amount of electricity saved due to more efficient power supplies would be approximately 1.5TWh. Figure 12 also shows a high-efficiency case, where power supply efficiency is 83% and electricity savings of 2.3 TWh are realized. An estimate for the potential savings for California is between 10% and 15% of the overall US savings potential, or between 230 GWh and 345 GWh.

Improved power supply efficiency will also yield compounded savings from reduced air conditioning loads and UPS losses not included here. We also note that the potential for the infrastructure load reduction (i.e. improved overall efficiency and reduced cooling load) is on the same order of magnitude as the power supply efficiency improvement.

Table 20. Energy Savings Potentials

Server Category	Number of Servers (Millions)	Annual Energy Consumption in TWh @ 70% PS Efficiency	Annual Energy Consumption in TWh @ 78% PS Efficiency	Annual Energy Consumption in TWh @ 83% PS Efficiency
Low-end	6,587,061	11.1	10.0	9.4
Work-horse	506,470	1.5	1.4	1.8
Mid-range	151,678	1.6	1.5	1.4
High-end	14,730	0.3	0.3	0.4
Total	**7,259,939**	**14.6**	**13.1**	**12.3**

[51] The 8% efficiency gain was derived by assuming that servers operate, on average, at 30% loading (i.e., 20% two-thirds of time and 50% one/third of time) and then calculating a weighted average of efficiency increases at 20% and 50% loading for EPS12V and EPS1U. This also implicitly assumes an even split between pedestal and rack servers, which appears consistent for server sales in 2004 and earlier.

Figure 12. Energy Savings Potential of PS in Servers

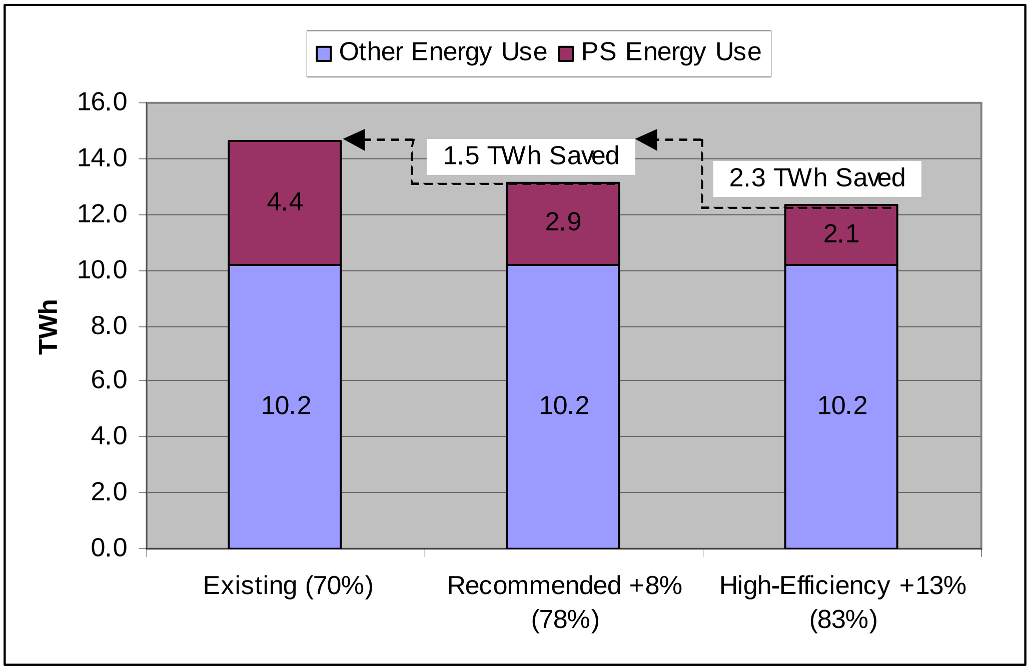

5 Power Supply Energy Savings in Other Devices

This task's objective was to assess other power supply savings opportunities in data centers. To do this we:

- Identified relevant topologies and configurations of "other" power supplies (e.g., routers, switches, hubs, etc.)
- Compiled information about the existing efficiency of power supplies in these devices through testing and manufacturer data.
- Estimated the market penetration of these devices and energy savings potential of more efficient power supplies.

5.1 Power Supply Topologies in Other Devices

Servers are certainly not the only devices in data centers that have potential to yield energy savings through improved power supply efficiency. Other devices such as routers, switches, hubs, and data storage units also contain power supplies. While the saving potential per unit is lower in these components, the sheer number of these devices warrants investigation.

As cited in Aebischer (2003), Mitchell-Johnson (2001) found that the composition of equipment at a typical data center in the United States was 60% servers, 18% switches, 9% disks and 8% routers.[52]

Using this information, a rough estimate of the number of devices in use can be made from our server stock estimates (Section 3.3.1). We estimated that there are approximately 7.2 million servers in use

[52] Aebischer, Bernard, Energy- and Eco-Efficiency of Data Centres, 2003, p. 14. Available at: http://www.cepe.ch/

in the United States. Table 21 shows how this would translate into numbers of other devices if all servers and other data handling equipment were found only in data centers.

Table 21. Other of Devices found in Data Centers

	Stock (units)	% of Equipment found in Data Centers
Servers	7,259,939	60%
Switches	2,177,982	18%
Disks	1,088,991	9%
Routers	967,992	8%
Other	604,995	5%
Total Devices	**12,099,898**	**100%**

Note: Switches represents the number of devices, not number of ports.

We know, however, that servers and other data handling equipment are routinely used in office environments as well, so we conducted more in-depth market research on four broad categories of "other" devices—routers, hubs, switches, and storage devices.

5.1.1 Routers

A router is a device that forwards data packets along networks. Routers are located at gateways, the places where two or more networks connect. Commonly, routers are used to connect two LANs or WANs or a LAN and an ISP's network.[53]

Cisco dominates the router market, claiming a 90 percent market share. Cisco's 2600 router is one of the most popular models on the market and now has added features such as Voice over Internet Protocol (VoIP).[54] Rivals Juniper Networks and 3Com have launched recent efforts to attack Cisco's dominant share. Juniper has started to take away market share with new systems in the terabit range.[55]

Routers in use today are typically stand-alone rack-mounted boxes, ranging in size from 1U to full-rack systems.[56] Some of the more common router deployments (and their power supply types) used by service providers, data centers, and enterprise networks are described in Table 22.

[53] http://www.webopedia.com/TERM/r/router.html
[54] http://news.com.com/Juniper+to+invite+Cisco+to+%27Pepsi%27+challenge/2100-1037_3-5171594.html
[55] http://techrepublic.com.com/5100-22_11-5363222.html
[56] ADL/US DOE (2002) p. 68.

Table 22. Description of Common Router Deployments

Router Class	Site	# of Routers	Examples	Power Supply
Service Provider	Telecommunications or Data Centers	Several dozen routers/switches or more in one facility.	Cisco 6513/7613, Cisco 12000, 12400, Juniper M160, M320, T320. Foundry or Extreme switches.	Mainly dual power supplies, sometimes 3 per system (two active, one in backup). Choice of AC or DC. Large telco would be DC.
Service Provider (Core)	Telecommunications	2 or 4 multiple chassis working together.	Juniper T640, Avici TSR, Cisco CRS-1 *Note: These are rate – a pair of T320s (above) would be more common.*	
Mid-range	Data Centers		Juniper M20, M40e, Cisco 7500, 7600/6500 systems.	Usually redundant dual power supplies AC or DC.
Small	Customer premise or enterprise network	Switches also used	Cisco 2500/2600	Some accommodate dual power supplies, but many do not. Some provide for a DC option but AC is most common.
Very Small	Residential DSL or cable, or small business/LAN			External AC power supplies.

Source: Dave O'Leary, Juniper Networks.

ADL/US DOE (2002) estimated the current stock of routers at 3.2 million by summing the shipments of routers, as reported in ITIC (2000) for the four years 1997 to 2000. We obtained more recent estimates of router shipments (segmented by router class) from Synergy Research Group (Table 23). This data shows that low-end, small and branch-office, routers account for over 80% of the current router stock assuming a four-year life as in ADL/US DOE (2002).

Table 23. Revised Router Stock

Router Class	2001	2002	2003	2004	Total Stock	% of Stock
Small	36,176	15,833	206,836	105,890	364,735	19%
Branch Office	382,928	325,912	273,732	237,688	1,220,260	62%
Mid-Range	81,178	78,131	70,103	58,774	288,186	15%
High-End Enterprise	10,453	11,689	5,874	1,503	29,519	2%
Service-Provider Edge	12,140	8,005	10,172	13,061	43,378	2%
Service-Provider Core	4,238	1,476	998	1,193	7,905	0%
Total	**527,113**	**441,046**	**567,715**	**418,109**	**1,953,983**	**100%**

Note: 2004 shipments are based on 3 quarters of actual data and one quarter of forecasted data.

Source: Synergy Research Group.

Next, we surveyed the manufacturer datasheets of over 65 different routers to determine the average power supply rating (in watts) and configuration for each router class (Table 24).

Table 24. Router Power Consumption

Router Class	Examples	Average PS Rating (Watts)	# of PS
Small	Allied Telesyn AR3XX Family Cisco 800, 1000, 1600 Tasman 1000 Family	15	1
Branch Office	3Com 3XXX Family Allied Telesyn AR4XX Family Cisco 1600, 1700, 2500, 2600 Tasman 4000	57	1 (or 2)
Mid-Range	3Com 5XXX Family Cisco 36XX, 4000	182	2 (or 1)
High-End Enterprise	3Com 6XXX Family Cisco 7500, 7206	513	2
Service-Provider Edge	Cisco 64XX, 7500, 7206, 7600, 10000, 7301, 7400 Juniper M5, M10, M20, ERX, M320 Laurel Networks ST200, ST Redback 400 Smart Edge, SMS 1800SL, SMS 1000SL Nortel MPE 9500	1,215	2 (or more)
Service-Provider Core	Alacatel 7770 OBX Avici TSR, SSR, QSR Caspian Aperio Cisco GSR, 12000 series, CRS-1 Juniper M40x, M160, T320, T640	3,660	2 (or more)

Based on this information, we determined the power consumption for routers of different classes (Table 25).

Table 25. AEC per Unit, Router

Router Class	Average PS Rating (Watts)	Hours / Yr	AEC, Per Unit (kWh)
Small	15	8,760	51
Branch Office	57	8,760	201
Mid-Range	182	8,760	1,275
High-End Enterprise	513	8,760	3,597
Service-Provider Edge	1,215	8,760	8,516
Service-Provider Core	3,660	8,760	25,653
Total			

5.1.2 Switches

LAN Switches

Roth (2002) estimated the stock of LAN switch ports in the United States to be 95 million.[57] This estimate is probably low, as recent shipments of switches are more than double those of just a few years ago.

In 2003, the worldwide LAN switch market grew 16.0% with total port shipments of 193.0 million versus 166.3 million in 2002. Despite this, market revenues decreased from $13.0 billion to $11.4 billion, due to falling average sales prices. In-Stat/MDR is predicting total port shipments of nearly 502.8 million by 2008.[58] Gigabit switches continue be one of the main drivers of this market, having surpassed Fast Ethernet in revenue during the first quarter of 2004. Cisco is by far the market leader in LAN switches, with over two-thirds of the market: HP, Nortel, 3Com, Extreme, and Foundry are also players. [59]

According to Dahlquist and Borovick (2000) [60] and other sources,[61] the United States accounts for about half of worldwide sales. Based on this information, and a two-year product life, we estimate the 2004 stock of U.S. LAN switches at approximately 213.4 million ports (Table 26).

Table 26. LAN Switch Stock – Worldwide and U.S. Market (Ports, millions)

Year	WW Switch Sales	U.S. Switch Sales	U.S. Stock
2001	137.2	68.6	
2002	166.3	83.2	151.7
2003	193.0	96.5	179.7
2004	233.7	116.9	**213.4**
2005	283.1	141.5	258.4
2006	342.8	171.4	312.9
2007	415.2	207.6	379.0
2008	502.8	251.4	459.0

Note: Worldwide sales data from In-Stat/MDR, assumes straight-line growth from 2003 to 2008. Assumes U.S. equals 50% of worldwide market. Stock calculated as two-year rolling average.

To determine the AEC of these switches, we surveyed over manufacturer datasheets for over 100 different products. Table 27 summarizes these results.[62]

Table 27. Power Consumption of a Switch

LAN Switch Type	Power Consumption per Port (Watts)	Hrs/Yr	AEC per Port (kWh/Yr)
Ethernet/Fast Ethernet	4.3	8,760	38
Gigabit	5.4	8,760	47

Note: Manufacturer datasheets did not typically differentiate between Ethernet and Fast Ethernet; therefore these categories were combined.

[57] ADL/US DOE, p. 71.
[58] http://www.instat.com/press.asp?Sku=IN0401449LN&ID=1037
[59] http://www.idc.com/getdoc.jsp?containerId=pr2004_08_19_165431
[60] ADL/US DOE, p. 71. Note: this estimate is probably low, as it is based only two years of sales data. Switches likely have a product lifetime similar to routers, 4 years.
[61] http://www.extremenetworks.com/aboutus/pressroom/releases/pr05_17_00.asp
[62] These numbers are similar to the findings of ADL/US DOE (2002), which estimated 35 kWh/yr per port.

WAN Switches

ADL/US DOE (2002) estimated a total of 50,000 WAN shelves installed in the United States.[63] However, this estimate is likely to be low as it appears to be based on only 1 year of sales data.

More recent market research from Dittberner Associates forecasts that the number of carrier WAN switch *and routers* shipped will increase from 24,841 units in 2001 to 82,157 in 2008, alongside a revenue increase from $3.76 billion to $18.2 billion.[64] However, other data sources describe a very different picture for this market indicating that "multiservice WAN switch sales rebounded nicely in 2004, after three consecutive years of declining sales."[65] The decline in sales was largely the result of service providers cutting capital expenditure budgets: about 90% of total multiservice switch sales are to service providers.[66]

2004 saw a turnaround in worldwide multiservice WAN switch sales. In-Stat/MDR reported sales of $2.3 billion, an increase of 9% over 2003 levels. Growth rates are expected to gradually decline over time as "switches face stiff competition from routers and service providers evolve to MPLS networks."[67] Assuming a 50% U.S. market share (as for LAN switches), Table 28 estimates the U.S. sales of WAN switches through 2008. Over the past three years, ATM/IP/MPLS-related equipment accounted for about 75-80% of WAN multiservice switch sales, while the remaining 20-25% was Frame Relay.[68]

Table 28. WAN Switch Sales –Worldwide and U.S. Market ($, billions)

Year	WW WAN Switch Sales	Projected Growth (%)	U.S. WAN Switch Sales
2001	$3.8		$1.9
2002	$2.3	-41%	$1.1
2003	$2.1	-7%	$1.1
2004	$2.3	9%	$1.2
2005	$2.5	7%	$1.2
2006	$2.6	6%	$1.3
2007	$2.8	6%	$1.4

Worldwide projected growth rates and sales from In-Stat/MDR.

Assuming an average sales price of $155,000[69] and a product life of four years, we estimate the current stock of WAN switches installed in the Unites States to be 34,000 (Table 29.)

Table 29. U.S WAN Switch Sales and Stock, Units

Year	U.S. WAN Switch Sales	U.S. WAN Switch Stock
2001	12,417	
2002	7,326	
2003	6,813	
2004	7,426	**33,983**
2005	7,946	29,512
2006	8,423	30,609
2007	8,928	32,724

Note: U.S. WAN switch estimated based on 50% market share.

[63] ADL/US DOE, p. 72.
[64] http://www.dittberner.com/news/newsrelease2.php
[65] http://www.instat.com/press.asp?ID=1286&sku=IN0501949WN
[66] http://www.instat.com/press.asp?ID=570&sku=IN030636WN
[67] http://www.instat.com/press.asp?ID=1286&sku=IN0501949WN
[68] In-Stat/MDR, Networking Quarterly - Multiservice WAN Switches, 2002-2004. Available at: http://www.instat.com/Catalog/ncatalogue.asp?ID=67&year=2004
[69] Based on Dittberner Associates 2001 data: $3.76 Billion and 24,841 units shipped.

Leading vendors in the multiservice WAN category are Alcatel, Ciena, Cisco, Ericsson, Lucent, Marconi, and Nortel.[70] Table 30 provides some example switching products, their power supply configurations, and approximate power consumption. As the table shows these products typically offer redundant power, with both DC and AC options available as required.

Table 30. Multi-Service WAN Switches

Manufacturer	Model	Watts	Power
Cisco	Cisco MGX 8950 Multiservice Switch	5000	Optional Ac
Cisco	Cisco MGX 8850 Series	1000	Optional Ac
Cisco	MGX8830	1200	Ac
Cisco	MGX8830	1050	Dc
Nortel	Multiservice Switch 6480		Ac or Dc, redundant
Nortel	Multiservice Switch 6440		Ac or Dc, redundant
Nortel	Multiservice Switch 6420		Ac
Nortel	Multiservice Switch 7420		Dc (Ac option), redundant
Nortel	Multiservice Switch 7440		Dc (Ac option), redundant
Nortel	Multiservice Switch 7460		Dc (Ac option), redundant
Nortel	Multiservice Switch 7480		Dc (Ac option), redundant
Nortel	Multiservice Switch 15000		Dc
Lucent	GX 550® Multiservice WAN Switch	1154	Dc
Lucent	GX 550® Multiservice WAN Switch	450	Ac
Lucent	CBX 3500	2880	Ac or DC Redundant chassis power distribution system
Lucent	CBX 500	1400	Redundant AC or Dc
Alcatel	OS/R	375	Ac, optional redundant
Alcatel	OS/R	650	Ac, optional redundant
Alcatel	Omni Switch 512	50	Internal Dc
Ciena	DN 7000	220	Dc
Ciena	DN 7050	350	Dc
Ciena	DN 7100	680	Dc
Ciena	DN 7200	2800	Dc

Source: Manufacturer datasheets.
Note: Nortel did not provide access to hardware specifications free of charge online.

Based on the manufacturer data collected, we determined the average power draw of a WAN switch to be 1,284 watts. Hence, assuming continuous operation the AEC of a WAN switch is approximately 11,250 kWh per year (Table 31).

Table 31. Power Consumption of a WAN Switch

	Power Consumption per Unit (Watts)	Hrs/Yr	AEC per Units (kWh/Yr)
WAN Switch	1,284	8,760	11,247

Power consumption per unit was determined by the average of 15 products which listed wattage ratings.

5.1.3 Hubs

A hub is common connection point for devices in a network. Hubs are commonly used to connect segments of a LAN.

[70] http://www.instat.com/abstract.asp?id=4&SKU=IN0501949WN

Silva (1998) estimated the number of hub ports installed in commercial office buildings in the United States to be 93.5 million.[71] However, over the last 5 years there has been a rapid migration from hubs to switches. In 1999, PC Magazine proclaimed, "The hub is dead….long live the switch!…No one should even think about installing a hub."[72]

Figure 13 shows a decline in worldwide shipments of Ethernet hubs from over 50 Million ports in 1998 to only 30 Million ports in 2003. Likewise, while shipments of Fast Ethernet hubs rose through 2001, these too have started to decline.

Figure 13. Worldwide LAN Hub Market, System Ports (000s)

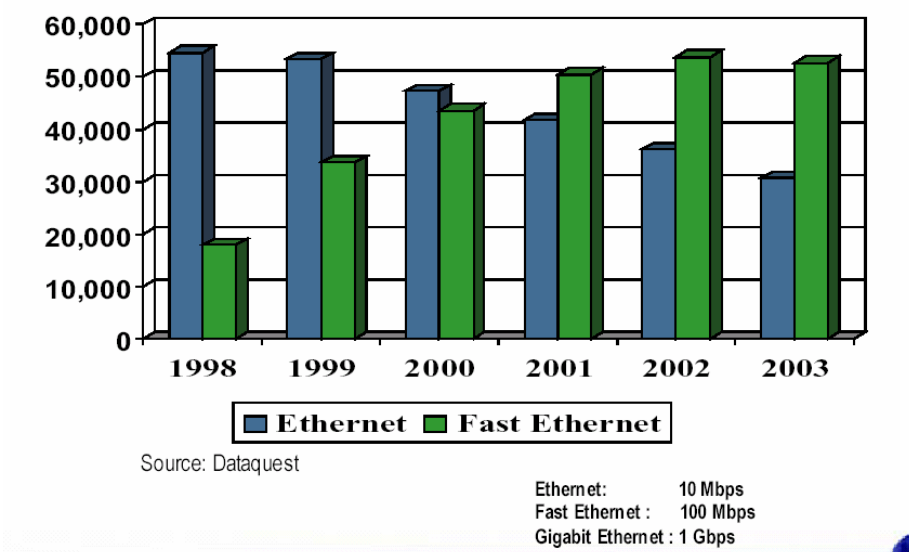

Source: http://www.xilinx.com/esp/consumer/home_networking/pdf_files/ethernet/complete.pdf

Using the above sales data and the same assumptions for hubs as for LAN switches (i.e., the United States represents 50% of the market and a product life of two years), we estimate the current stock of U.S. LAN hubs is approximately 77 million ports (Table 32).

Table 32. LAN Hub Stock – U.S. Market (Ports, millions)

LAN Hub Type	2001	2002	2003	2004e	2004 Stock	% of Stock
Ethernet	20	18	15	13	28	36%
Fast Ethernet	25	27	26	24	50	64%
Total	**45**	**45**	**41**	**36**	**77**	**100%**

Notes: Assume US market is 50%. Source: Ecos estimate based on 2003-2004 market data; 2004 data estimated based on trends in previous year sales.

In 2000, IDC projected that the worldwide *installed* base of 10 Mbps (Ethernet) hubs would drop from 184 million to 153 million ports (17%) and another 21% to 121 Million in 2001.[73] Assuming a 20% drop in subsequent years, by 2004[th] the worldwide installed base of Ethernet hubs would be projected to approximately 62 Million ports in 2004, and the U.S. portion about 30 million. This is very close to our estimate of 28 million ports in Table 32.

To determine the AEC of hubs, we surveyed over manufacturer datasheets for over 12 different products. Table 33 summarizes these results.[74]

[71] Roth, p. 73.
[72] http://www.onshore.com/downloads/partners/cisco/CIS_Mig_hub.pdf
[73] http://www.onshore.com/downloads/partners/cisco/CIS_Mig_hub.pdf
[74] These numbers are similar to the findings of Roth (2002), which estimated 35 kWh/yr per port.

Table 33. Power Consumption of a Hub

LAN Hub Type	Power Consumption per Port (Watts)	Hrs/Yr	AEC per Port (kWh/Yr)
Ethernet/Fast Ethernet	1.9	8760	17

Note: Manufacturer datasheets did not typically differentiate between Ethernet and Fast Ethernet; therefore these categories were combined.

5.1.4 External Data Storage

The demand for data storage is large.[75] IDC data (Table 34) indicates that in 2003 nearly 400,000 external storage units were shipped worldwide, of which about two-thirds was direct attached storage (DAS) and the remaining one-third was networked, including both network attached storage (NAS) and Storage Area Networks (SAN).[76]

In 2004, IDC reported that it was encouraging to see the "continued acceleration in the annual growth rate for external disk storage systems petabytes, which grew 63% in 2004. With revenue accelerating, but at a slower pace, average pricing remains competitive during the ongoing influx of higher-capacity drives and applications."[77]

EMC was the largest vendor of external storage in 2004, with a 21.1% revenue share. HP and IBM followed with 18.7% and 12.6% share, respectively. Hitachi and Dell are also dominant players.[78]

Table 34. Worldwide External Data Storage Shipments and Sales

	1999	2000	2001	2002	2003	2004est
Units						
DAS	503,608	509,667	425,255	298,264	270,379	301,492
Networked	74,215	138,455	140,902	141,148	128,599	143,397
Total External Disk Storage	577,823	648,122	566,157	439,412	398,978	444,889
Non-OEM Factory Revenue ($M)						
DAS	$13,773	$14,452	$9,357	$5,932	$5,504	$7,112
Networked	$4,368	$7,299	$7,838	$7,165	$8,087	$8,055
Total External Disk Storage	$18,141	$21,751	$17,195	$13,097	$13,591	$14,165

Note: Networked includes NAS and SAN.
Source: 2004 data estimated based 2004 IDC sales data and average prices. 1999-2003 data: IDC, 2004 Release as shown in *The Business for Storage Networks. Chapter 1, Industry Landscape: Storage Costs and Consumption*. Cisco Systems, Sept 3, 2004. Available at:
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns516/c1272/cdccont_0900aecd80257124.pdf

The United States disk storage systems market represents 39% of the worldwide market.[79] Assuming a product life of four years, we estimate the 2004 stock of External Disk Storage to be 721,280 units.

[75] The ADL/US DOE 2002 report did not look at data storage separately; it was combined with servers.
[76] Williams, Bill. The Business Case for Storage Networks. Cisco Press, 2004. Available on: http://www.ciscopress.com/content/images/1587201186/samplechapter/1587201186content.pdf
[77] http://www.idc.com/getdoc.jsp?containerId=pr2005_03_03_154203
[78] http://www.idc.com/getdoc.jsp?containerId=pr2005_03_03_154203
[79] Genereux, Scott. The United State Storage Market: A Perspective from Hitachi Data Systems Corporation.

Table 35. U.S. External Data Storage Stock

Year	Units Shipped			Estimated Stock
	DAS	Networked	Total	
1999	196,407	28,944	225,351	
2000	198,770	53,997	252,768	
2001	165,849	54,952	220,801	
2002	116,323	55,048	171,371	870,290
2003	105,448	50,154	155,601	800,541
2004	117,582	55,925	173,507	**721,280**

Notes: Assumes U.S. is 39% of worldwide market. Stock calculated as 4-yr rolling average.

There is a wide range of storage products. We briefly surveyed some of the storage products offered by Dell, EMC, HP, IBM, and once again we found that many manufacturers do not provide details about power supplies in their product datasheets. Table 36 lists some of the products that did report power supply information.

Table 36. Selected Storage Products

Manu	Model	Form Factor	Watts	Power Supply	Efficiency
Dell	PowerVault 22Xs	3U	600		
Dell	Powervault 745N	1U	280		
EMC	NetWin	1U	250	Auto-sensing 120/240V	
EMC	CLARiiON AX100 - single processor	2U	250	Single power supply	
EMC	CLARiiON AX100 - dual processor	2U	326	Redundant	
EMC	CLARiiON CX700	2U	510	Redundant	
EMC	NS700	8U or 4U	510		
EMC	NS500	6U or 3U	578		
EMC	CLARiiON CX500	2U	618	Redundant	
EMC	Clariion CX300	2U	618	Redundant	
HP	HP ProLiant DL100 Storage Server	1U	250	Single power supply	
HP	HP ProLiant ML110 Storage Server	Tower 5U	350	1 350W non-redundant power supply	
HP	HP StorageWorks DL380-SL Clustered Gateway	2U	575	CE Mark Compliant, Hot Plug Redundant power supply included	
HP	HP ProLiant DL380 G4 Storage Server	2U	575	Redundant Power Supplies; 2 x 575-Watt Hot Plug Power Supplies	73.1-77%
HP	HP ProLiant ML350 G4 Storage Server	Tower 5U	725	1 hot plug 725 Watt Power Supply (Base Model); 2 hot plug 725 Watt Power Supply (Int SCSI Storage)	
HP	HP ProLiant ML370 G4 Storage Server	Tower 5U	775	1 hot plug 775 Watt Power Supply (Base Model); 2 hot plug 775 Watt Power Supply (High Performance Model)	
HP	HP ProLiant DL580 G2 Storage Server	4U rack mount	800	2 x 800-Watt (low line or high line) Hot Plug Power Supplies	70-76%

Manu	Model	Form Factor	Watts	Power Supply	Efficiency
HP	HP ProLiant DL585 Storage Server	4U rack mount	800	2 x 800-Watt (low line or high line) Hot Plug Power Supplies	70-76%
IBM	IBM TotalStorage NAS Gateway 500	Rack-mount 4U	670	Two 670 Watt auto-ranging hot swappable power supplies	

Source: Manufacturer datasheets.

Based on the products listed above, we determine the average AEC of external disk storage units to be 4,591 kWh per year. Table 37 summarizes these results.

Table 37. Power Consumption of External Data Storage

LAN Hub Type	Power Consumption per Unit (Watts)	Hrs/Yr	AEC per Units (kWh/Yr)
External Disk Storage	529	8760	4,591

5.2 Existing Efficiency of Power Supplies in Other Devices

5.2.1 Manufacturer Data

We found that manufacturers of routers, hubs, switches, and storage devices provide very little, if any, information about the device's power supply in their product specifications. Manufacturer datasheets sometimes give the rated output and/or configuration of the power supply, but do not list any information about efficiency. Dave O'Leary of Juniper Networks confirms that this is standard industry practice.[80]

Aebischer (2003) notes that the power supplies contained in these other devices are similar to the ones found in servers, therefore one would expect to see similar results.[81] Aebischer also measured "the workload of several servers, routers and switches…" and found it "…lies typically between 50% and 30%."[82]

Juniper Networks provided us with some limited information about several of their products (Table 38).

[80] Personal communication with Dave O'Leary, Nov 29, 2004.
[81] Aebischer, Bernard, Energy- and Eco-Efficiency of Data Centres, 2003, p. 28. Available at: http://www.cepe.ch/
[82] Aebischer, Bernard, Energy- and Eco-Efficiency of Data Centres, 2003, p. 31. Available at: http://www.cepe.ch/

Table 38. Power Supply Efficiency in Juniper Network Devices

Model	Device Category	PS Rating (Manufacturer? Configuration?)	# of PS	Reported Efficiency (at Full Load)	Tested Efficiency ?
M5, M10	Internet Router	800W (AC); 700W (DC)	two load-sharing, AC or DC		
M7i	Internet Router	293W	one or two load sharing, AC or DC		
M20	Internet Router	750W isolated	two load-sharing, isolated, AC or DC		
M40e	Internet Router	2900W isolated (AC); 3000W nonisolated (DC)	two load-sharing, pass-through, AC or DC		
M160	Internet Router	2400W nonisolated; 3000W nonisolated (enhanced)	two load-sharing, pass-through power supplies, DC		
M320	Internet Router	1750W (AC) 2000W (DC)	four load-sharing AC power supplies, AC or DC		
T320	Internet Router	3200W	two load-sharing DC power supplies		
T640	Internet Routing Node?	3200W	two load-sharing DC power supplies		
TX	Matrix Platform	4560W	two load-sharing DC power supplies		

Source: http://www.juniper.net/techpubs/hardware/

In addition, we found one power supply manufacturer (Delta Electronics) that classifies its product offerings by application (Desktop, Server, Networking, and Other). Figure 14 provides a summary of the power supply ratings and efficiencies of all Delta Electronics *Networking* products. These products are further divided into five subcategories: Network, PC Peripheral, Storage, 1U Application, and Information Application. While there is a very wide distribution of power supply ratings (from about 5W to over 4500W, the majority of these products have efficiencies in the 60-75% range. Only a handful of the Network and Storage power supplies are 80% or more efficient; these appear to be the higher-rated models (700W and above). One of the units claims a remarkable 95% efficiency.

Figure 14. Efficiency of Delta Electronics Networking Power Supplies

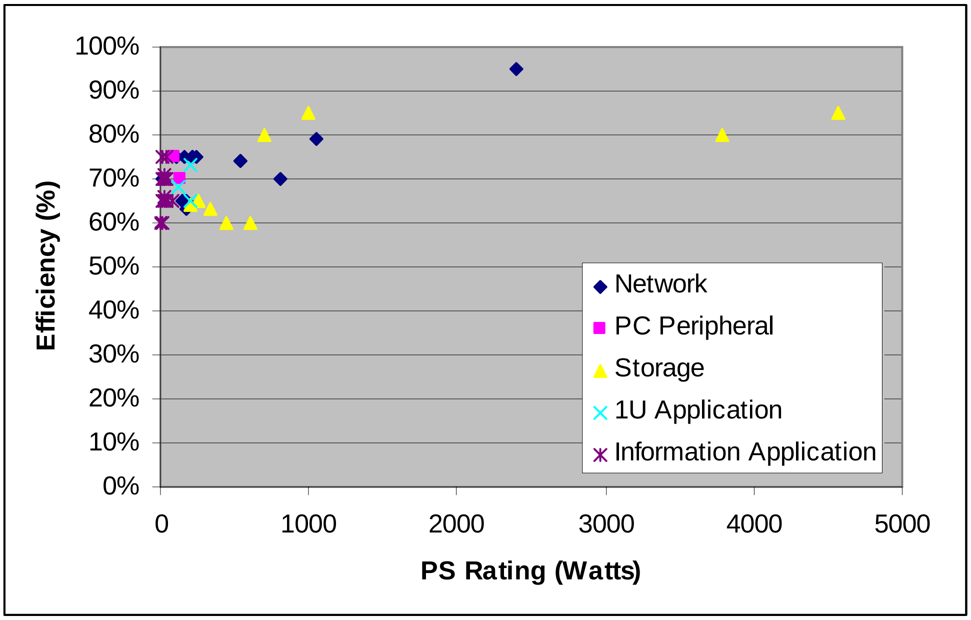

Source: Delta Electronics, product datasheets. Available at:
http://www.delta.com.tw/product/ps/sps/sps_main.asp

The average efficiencies for different subcategories of Delta Electronics Networking power supplies are provided in Table 39.

Table 39. Average Efficiency of Delta Electronics Networking Power Supplies

Networking Power Supply Subcategory	# of Power Supplies	Average Efficiency (%)
Network	17	73.2%
PC Peripheral	3	73.3%
Storage	10	70.7%
1U Application	3	68.7%
Information Application	18	67.3%
All	**51**	**70.4%**

Source: Delta Electronics, product datasheets. Available at:
http://www.delta.com.tw/product/ps/sps/sps_main.asp

5.3 Market Penetration of Other Devices and Energy Savings Potential

We estimated the energy savings potential of more efficient power supplies for “other” devices in three steps:

1. We developed a preliminary AEC estimate for each category of device from market research gathered on the number of devices and the power supplies, contained there in (Section 4.1).
2. We revised our AEC estimate to more accurately reflect the typical operation and efficiency of these devices (i.e., 30% loading and 70% efficient power supply) (Section 4.2).

3. We then estimated savings potential of more efficient power supplies for two cases (as for servers): a "recommended" 78% efficiency case and a high-efficiency 83% case.

5.3.1 AEC

The AEC for each category of device—router, switch, hub, and data storage—was calculated by multiplying the annual AEC per unit times the total stock. This methodology is similar to the one was used in ADL/US DOE 2002. Comparisons the ADL/US DOE 2002 estimate are provided here for reference.

Routers

The annual energy consumption (AEC) for routers is 1.31 TWh (Table 40). This estimate is similar to the ADL/US DOE 2002 estimate of 1.1 TWh. However, their number was derived very differently, assuming over 1 million more servers and flat average 40 watts per server (about 350 kWh per year).

Table 40. Annual Energy Consumption, Routers

Router Class	AEC per Unit (kWh)	2004 Stock	AEC (TWh)
Small	51	364,735	0.02
Branch Office	201	1,220,260	0.24
Mid-Range	1,275	288,186	0.37
High-End Enterprise	3,597	29,519	0.11
Service-Provider Edge	8,516	43,378	0.37
Service-Provider Core	25,653	7,905	0.20
Total		**1,953,983**	**1.31**

Switches

For LAN switches, we estimate AEC of 9.0 TWh. (Table 41). This nearly three times higher than the AEC estimated by ADL/US DOE (2002) for LAN switches (3.3 TWh). Our estimate of AEC per port is similar; so, the increase in overall AEC is due to the increase in the number of ports installed.

Table 41. AEC of LAN Switches.

LAN Switch Type	AEC per Port (kWh)	2004 Stock (Ports, Millions)	AEC (TWh)
Ethernet/Fast Ethernet	38	106.7	4.0
Gigabit	47	106.7	5.0
Total		**213.4**	**9.0**

Notes: Assumes 50% of market is Gigabit.

For multiservice WAN switches, we estimate 0.4 TWh AEC. (Table 42). Like LAN switches, this nearly three times the AEC estimated by ADL/US DOE (2002) for WAN switches (0.15 TWh). However, unlike LAN switches, for WAN switches the increase in AEC is largely due to a much higher estimate of energy consumption per device.

Table 42. AEC of WAN Switches.

	AEC per Units (kWh/Yr)	2004 Stock	AEC (TWh)
WAN Switch	11,247	33,983	0.4

Hubs

Our AEC estimate for LAN hubs is 1.3 TWh. (Table 43). This is slightly less than the AEC estimated by ADL/USDOE (2002) for hubs (1.6 TWh), which makes sense given the transition from hubs to switches.

Table 43. AEC of LAN Hubs

LAN Hub Type	AEC per Port (kWh/Yr)	2004 Stock (Ports, Millions)	AEC (TWh)
Ethernet	17	28	0.5
Fast Ethernet	17	50	0.8
Total		**77**	**1.3**

External Data Storage

Combining our stock and AEC per unit estimates yields an annual AEC of 3.3 TWh for U.S. external disk storage (Table 44).

Table 44. AEC of External Data Storage

	AEC per Unit (kWh)	2004 Stock	AEC (TWh)
External Data Storage	4,591	721,280	3.3

5.3.2 Adjusted AEC

The AEC estimates above use a similar methodology to the one employed by ADL/US DOE (2002). This methodology, however, does not account for existing power supply efficiency, nor does it account for the fact that devices normally operate at a fraction of their stated output rating.

We revised our initial AEC estimates to reflect an average of 30% loading and 70% power supply efficiency. Table 45 presents the "adjusted" AEC.

Table 45. AEC, adjusted for 70% PS Efficiency and 30% Loading

Segment	AEC (TWh)	Adjusted AEC (TWh)	% of Adjusted AEC
Routers	1.3	0.6	8%
Switches - LAN	9.0	3.9	59%
Switches - WAN	0.4	0.2	3%
Hubs	1.3	0.6	8%
Storage Devices	3.3	1.4	22%
Total	15.3	6.6	100%

Note: AEC is adjusted by a factor of .43 =.3/.7 to account for loading and PS efficiency.

5.3.3 Efficient Power Supply Energy Savings Potential

Together, these "other devices" consume a little less than half of the energy that servers do. LAN switches contribute the most to AEC simply due to their large numbers. Remember that, as with servers, not all of the energy consumed goes to power the device itself. Some energy is "lost" or consumed by the power supply itself; the amount of energy consumed depends on the efficiency of the power supply. In the typical existing case, where power supply efficiency is 70%, 30% or about 2.0 TWh is consumed by the power supply (i.e., energy losses)

As with servers, energy savings can be achieved by increasing the efficiency (or reducing the energy losses) of the power supply. We present two alternative cases to demonstrate the potential energy savings due to more efficient power supplies. The first is a "recommended case" where power supply efficiency is increased from 70% to 78%; the second is a "high-efficiency case" where power supply efficiency reaches 83%. Figure 15 shows that by increasing power supply efficiency by 8%, to 78% we are able to save 0.7 TWh by cutting power supply energy consumption by 34% to 1.3 TWh.

Similarly, under the high-efficiency scenario, savings of 1.1 TWh are realized by cutting power supply energy consumption 52% to 0.9 TWh.

Figure 15. Energy Savings Potential of PS in Other Devices

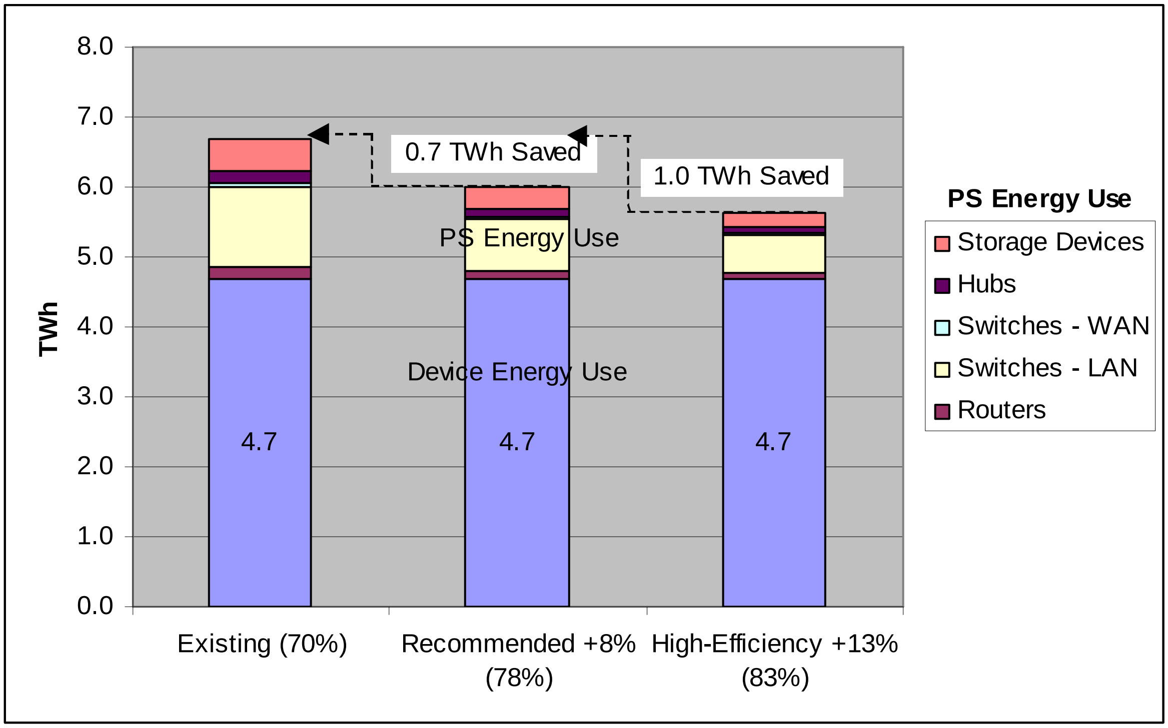

These high-level estimates show that there is significant energy savings potential from improving efficiency of power supplies in devices other than servers. While the savings potential is not as large as in servers in absolute terms, the extension of efforts to these “other’ devices would be relatively straightforward. Much of the groundwork for improving power supply efficiency and developing specifications for servers, desktops, etc. has already been done or is in process. This work could serve as a starting point for efforts targeting other devices.

The main challenge to encouraging efficiency in the power supplies of other IT equipment like routers and switches is that there is currently no industry body like SSI coordinating standards and efficiency improvements in these products. Power supply designs for this type of equipment are usually customized for a particular product and may vary from manufacturer to manufacturer, whereas in servers, only a few distinct form factors for power supplies exist. The diversity of power supply designs in equipment like routers and switches may complicate efforts to uniformly encourage efficiency improvements.

Fortunately, one does not necessarily have to address all of these product categories at once and can instead focus on the “low hanging fruits”. Our estimates help to provide some guidance as to which devices should be prioritized first. For example, LAN switches clearly represent the best energy savings opportunity because of the sheer numbers of devices in use.