



High Performance Buildings: Data Centers *Uninterruptible Power Supplies (UPS)*

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

This report was prepared as a result of work sponsored by the California Energy Commission (Energy Commission). It does not necessarily represent the views of the Energy Commission, its employees, or the State of California. The Energy Commission, the State of California, its employees, contractors, and subcontractors make no warranty, express or implied, and assume no legal liability for the information in this report; nor does any party represent that the use of this information will not infringe upon privately owned rights. This report has not been approved or disapproved by the Energy Commission, nor has the Energy Commission passed upon the accuracy or adequacy of this information in this report.

Table of Contents

1. Executive Summary	1
1. Current UPS Efficiencies	1
2. Proposed UPS Efficiency Levels and Labeling.....	1
3. UPS Energy Use and Potential Savings.....	2
4. Conclusions and Recommendations	4
2. Current UPS Market	6
2.1 Assessment of UPS Topologies.....	6
2.1.1 Overview of Offline or Standby UPS Topology.....	6
2.1.2 Overview of Inline or Line-Interactive UPS Topology	7
2.1.3 Overview of Delta-Conversion Topology	9
2.1.4 Overview of Double-Conversion or Online UPS Topology.....	10
2.1.5 Overview of the U.S. UPS Market (Data Center/IT Sector)	11
2.2 UPS Test Protocol	14
2.2.1 Factory/Laboratory UPS Efficiency Test Method	15
2.2.2 Field UPS Efficiency Test Method	15
3. UPS Efficiency Data	16
3.1 UPS Efficiencies	16
3.1.1 UPS Efficiency and Percent Load	16
3.1.2 UPS Efficiency and Load Characteristics.....	17
3.1.3 Results of UPS Efficiency Testing: Field Measurements	18
3.1.4 Results of UPS Efficiency Testing: Factory/Laboratory Measurements .	19
3.1.5 Investigation of Special High Efficiency Modes in Select UPSs.....	21
3.2 Redundancy Evaluation.....	26
3.2.1 Capacity or “N” Configuration	26
3.2.2 Isolated Redundant Configuration	27
3.2.3 “N+1” Parallel, Single-Bus Configuration.....	27
3.2.4 “2N” and “2(N+1)” Dual-Bus Configurations	28
4. Proposed Efficiency Specifications	30
4.1 UPS Operation	30
4.1.1 Total Cost of Ownership.....	30
4.2 Proposed Efficiency Label.....	31
4.2.1 Labeling UPS Efficiency	31
4.2.2 Proposed UPS Efficiency Standard for the U.S. Market.....	33
5. DC Architecture Scoping Study	40
5.1 Market Trends in IT Load Power Configuration.....	40
5.2 Efficiency of AC- and DC-based Systems	41
5.3 Demonstration Plans for DC Architecture	42
6. Conclusions and Recommendations	43

1. Executive Summary

As part of the California Energy Commission’s PIER (Public Interest Energy Research) initiative on efficient data centers, LBNL tasked Ecos Consulting and EPRI Solutions with investigating uninterruptible power supplies (UPSs) used in data centers and other high-tech facilities to guard against interruptions in services, primarily from electrical outages. The goals for this project included a complete analysis of distribution of UPS design efficiencies, which included the development of a UPS test protocol, and the broad circulation of efficiency findings to the industry through the CEC PIER program, Lawrence Berkeley National Laboratories (LBNL), and other industry and efficiency venues, such as www.efficientpowersupplies.org. Using the efficiency findings and market data, we also estimated the overall energy consumption of UPSs 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 proposed efficiency levels and a labeling scheme for various types of UPS units as a way to encourage the use of more efficient UPSs.

1. Current UPS Efficiencies

Our results indicate that there can be a wide variation in efficiency even between various UPS configurations. The table below summarizes the characteristic efficiency of a number of UPS topologies at various load conditions and shows the average efficiency for all of the UPSs measured. The table below also shows that regardless of the configuration or type, UPSs tend to be more efficient at full rather than part load.

Table ES1: Characteristic Efficiency of UPS Topologies

UPS Topology	Efficiency at 25% Load	Efficiency at 50% Load	Efficiency at 75% Load	Efficiency at 100% Load
<i>Delta-Conversion</i>	93% - 94%	96% - 97%	97%	97%
<i>Double-conversion</i>	81% - 93%	85% - 94%	86% - 95%	86% - 95%
<i>Line-Interactive</i>	NA	97% - 98%	98%	98%
<i>Standby</i>	NA	NA	NA	NA
AVERAGE OF ALL UNITS	86%	89%	90%	90%

Note: A straight average was used in the table above – we lacked the data for a weighted average calculation.

Due to the scope of this research effort, it was impractical to determine what specific elements of a model’s design contributed to its measured efficiency; however, our results indicate that some broader design decisions, such as the general UPS topology, can be indicative of efficiency. Research on the efficiency of both flywheel and delta-conversion UPSs should continue as manufacturers continue to improve their designs and diversify product lines.

2. Proposed UPS Efficiency Levels and Labeling

In the early stages of this project, we uncovered a proposed energy efficiency and power quality labeling scheme for single-phase, “small” UPSs being developed by the Swiss Federal Office of Energy. The Swiss’ proposed label, documents the measured power losses of the UPS in different modes of operation and provides an estimate of the annual energy consumed by energy losses in the UPS design. The label also reports the UPS’s ability to handle different types of power disturbances and the power quality (power factor and total harmonic distortion) that the UPS presents to the grid. It was designed to match the style of existing EU labels for other electric appliances.

In the interests of harmonizing with international efforts to report and promote UPS energy efficiency in a standardized format, our team created a modified version of the proposed Swiss label and a draft report summarizing the criteria that would be used to evaluate UPS efficiency in the U.S. market under the label. Our proposed label, shown below in Figure ES1 includes the following changes over

the Swiss label (note that both standards still do not address UPS attributes other than efficiency, as this is an area best handled with input and leadership from manufacturers):

- The label reports differing levels of energy conversion efficiency rather than energy losses. It was our opinion that reporting efficiency was more appealing to a potential UPS purchaser.
- We have eliminated a field in the label that reports energy losses incurred by operating the UPS for 2000 hours with no load.
- To help direct the focus of the label solely on energy efficiency and to simplify its design, we decided to eliminate the various tables that report on filtering of power disturbances and power quality.




Figure ES1

3. UPS Energy Use and Potential Savings

Using efficiency data and our estimates of UPS stock, we constructed a first-order estimate of UPS energy use for the US. This energy use and savings estimate is limited to the data center/IT sector due to the scope of our research, in which we examined the efficiency of data centers only.

Based on our estimates, the data center/IT sector currently consumes roughly 7.1 TWh (7.1 billion kWh) of electricity per year. We estimate that the State of California alone, which contains roughly 15% of the nation’s data center floor space,¹ consumes about 1 billion kWh of electricity and spends about \$100 million per year in electric bills due (mostly) to power conversion losses in data center UPSs.² Our national energy use figures are noticeably higher than past estimates made by Arthur D. Little for the U.S. Department of Energy, which estimated UPS energy consumption in the U.S. “IT/telecom” sector at 5.8 TWh per year.³ Both Ecos’ and Arthur D. Little’s energy use estimates are shown in Table ES2.

Using a “high efficiency” and “typical” cases for UPS energy consumption, we have estimated the energy savings potential for the U.S. The introduction of a voluntary UPS efficiency specification that would be administered under a program such as ENERGY STAR or other institutions may result in a

¹ R. Juarez, M. Alic, K. Chetan and B. Johnson. *SpaceDex III – Hosting Space: Not All Hosting Space is Created Equal – Smart, Complex Space Takes Data Center Stage*. Boston: Robertson Stephens Inc. 2001.

² Because California contains 15% of the nation’s data center floor space, its UPSs consume 15% of the energy consumed by all UPSs in the US.

³ K. Roth, F. Goldstein and J. Kleinman. *Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings, Volume I: Energy Consumption Baseline*. Cambridge, MA: Arthur D. Little, 2002. p. 84

percentage of the UPS market meeting efficiency levels, resulting in partial realization of the savings potential.

Table ES 2: Comparison of UPS Energy Use Estimates

	DOE/ADL 2002 Estimate of UPS Annual Energy Consumption in U.S. IT/telecom Sector	Ecos/EPRI 2004 Estimate of UPS Annual Energy Consumption in U.S. IT/data center Sector
<i>Estimated National UPS Stock</i>	151,761	250,343
<i>Nationwide Annual Energy Consumption (TWh)</i>	5.8	7.1
<i>Estimated California UPS Stock</i>	n/a	37,551
<i>California Annual Energy Consumption (TWh)</i>	n/a	1.1

If, for example, 25% of the UPSs installed in the U.S. were able to comply with the proposed UPS efficiency levels, the nation would save about 700 million kWh of electricity per year, amounting to about \$70 million in cost savings. We have summarized our energy savings scenarios for mandatory and voluntary specifications in Table ES3, along with similar energy savings estimates that apply only to the State of California, which we assume contains about 15% of the nation's UPSs.⁴

Table ES3: Energy Savings Scenarios for UPS Efficiency

	Mandatory Specification – 100% of Installed Units Comply	Voluntary Specification – 25% of Installed Units Comply
Nationwide Annual Energy Savings (TWh)	2.8	0.71
Nationwide Annual Utility Cost Savings (million USD)	\$280	\$71
California Annual Energy Savings (TWh)	0.42	0.11
California Annual Utility Cost Savings (million USD)	\$42	\$11

The energy savings figures shown here could be significantly larger when examining the *entire* UPS market and not just the data center/IT sector. For UPSs above 5 kVA in power output, the data center/IT sector only accounts for roughly two thirds of the installed units. An additional one third of UPSs installed in small offices, industrial facilities, hospitals and the like have not been included in our energy use and savings estimates. Our analysis also excludes the below 5 kVA UPS market segment, more typical of small office applications, residential installations, etc. The combined sales of UPSs below 5 kVA amounted to over 8 million units in 2004, close to 200 times the unit sales volume

⁴ Based on estimates that California contains 15% of the nation's data center floor space. R. Juarez et al. 2001.

of UPSs greater than 5 kVA in power output.⁵ Although these are lower power devices than the larger UPSs examined in this report and would thus have smaller per-unit power losses, they might represent an attractive energy savings opportunity simply due to the large number of units in operation that should be studied in the future.

4. Conclusions and Recommendations

This project served well as an initial investigation into this particular area of high-tech buildings and data centers, and has provided a good overview of UPS utilization and applications. Due to the broad scope of our work, we were able to identify a number of important issues. However, we were not able to carry out in-depth investigations into any one issue. Below are a number of observations and recommendations based on the findings of this study:

Focus on Large UPS Units: Standby UPSs are most common in low-power applications (0 – 2 kVA) such as individual workstation backup power, small office server backup power, etc. The lack of complete power conditioning and load isolation generally make standby UPSs unsuitable to critical loads such as industrial facilities or data centers. We estimate that there are about 225,000 UPSs in operation in the U.S. data center/IT sector today, about 15% of which reside in the State of California. Half of the national UPS stock are units in the 5.1 to 20 kVA range of apparent power output.

UPS Loading and Sizing Can Significantly Affect Energy Use: UPSs in mission critical environments typically operate at points between 30% and 50% of nominal and in situations where the data center is not fully populated this can be even lower resulting in very large inefficiencies. There are a number of reasons why data center UPSs operate at part loads, one of which is that UPSs are commonly sized to meet the maximum utilization of space in a data center, even though maximum utilization rarely occurs.⁶ This results in over sizing of the UPS for the actual amount of installed mission-critical infrastructure.

Topology and Efficiency Considerations: Due to the scope of this research effort, it was impractical to determine what specific elements of a model's design contributed to its measured efficiency; however, our results indicate that some broader design decisions, such as the general UPS topology, can be indicative of efficiency. Research on the efficiency of both line-interactive and delta-conversion UPSs should continue as manufacturers continue to improve their designs and diversify product lines. Nevertheless, the high efficiency of some UPS topologies in our current data set suggest that they deserve serious consideration by facility managers looking to lower operating costs in mission critical facilities that utilize large UPSs (greater than 50 kVA in capacity) and policy makers wishing to identify the most efficient technologies.

Reliability Considerations: Desire for high efficiency should naturally be balanced with concerns over load isolation and reliability. Although the delta-conversion UPSs that we tested performed better in regards to overall efficiency compared to double-conversion units, some manufacturers have argued that delta-conversion UPSs do not provide the same load isolation as a "true" double-conversion UPS. This report does not attempt to answer the question as to which topology is the overall best choice for mission critical applications. We simply report the observed efficiencies of the various topologies, and facility managers will ultimately need to judge whether a particular UPS can handle load disturbances and provide an acceptable level of load isolation. Additional research and discussion within the industry can help to provide more clarity on this issue.

Considerations in Creating a Standard for UPS: Because of the wide variation in performance among UPSs and the equally wide range of available sizes, it is difficult to create a one-size-fits-all UPS specification. UPSs can pay a penalty in overall efficiency by

⁵ Frost and Sullivan. 2004.

⁶ Madsen J. "Continuous UPS Availability: How Important is it to Your Company?" *Energy User News*. August 11, 2000. <http://www.energyusernews.com/CDA/ArticleInformation/features/BNP__Features__Item/0,2584,14489,00.html>

providing better load isolation, filtering, and other types of improved performance, and therefore it would be unfair to compare UPSs with a high-quality output to those with a lower-quality output. Furthermore, smaller power conversion devices typically cannot achieve the same levels of efficiency as larger devices, and so it would be equally unfair to compare the efficiency of a 100 kVA, facility-level UPS to that of a 1 kVA, workstation-level UPS. As a result, any proposed UPS efficiency specification should take the size and performance of the UPS into account.

Finally, it should be noted that additional energy savings can be realized from reduced facility cooling needs because less waste heat is generated from the conversion and storage equipment. LBNL benchmarking suggests that the secondary savings from reduced cooling needs as a result of more efficient facilities is on the same order of magnitude as the direct, primary savings from efficient conversions. There is also additional savings from reduced capital investment in equipment if a more efficient system is selected, as well as reduced real estate investment from smaller equipment footprints. These additional savings can be realized up front rather than through reduced operational costs. Thus, additional investigations into the areas discussed above, as well as the feasibility of DC powering architectures for data centers will provide further insight and can help in increasing the overall efficiencies of future data centers.

2. Current UPS Market

2.1 Assessment of UPS Topologies

Our first task was to survey the various types of UPS design topologies available on the market today – particularly those that are commonly used in data centers – and assess the efficiency of various designs through a data collection effort. A general description of the topologies follows in this section, whereas our analysis of the efficiency data gathered from the field and through manufacturer participation appears in Section 1.3. Some of the major questions that we hoped to answer in our overview of UPS topologies included:

- What specific factors separate one design topology from another?
- Which topologies are likely to be relevant to data center applications?
- How many UPSs of the various topologies are installed in data center/IT facilities across the U.S.?

2.1.1 Overview of Offline or Standby UPS Topology

The offline or standby UPS is the simplest and oldest UPS topology. The operating principle of a standby UPS is, as the name suggests, standing by in case of a disruption in power to a critical load. The UPS only delivers power to its connected load when power from the utility does not meet specifications. The vast majority of the time when utility power is deemed acceptable, the UPS allows the connected load to draw power directly from AC mains, using a small amount of power to trickle charge the UPS batteries.

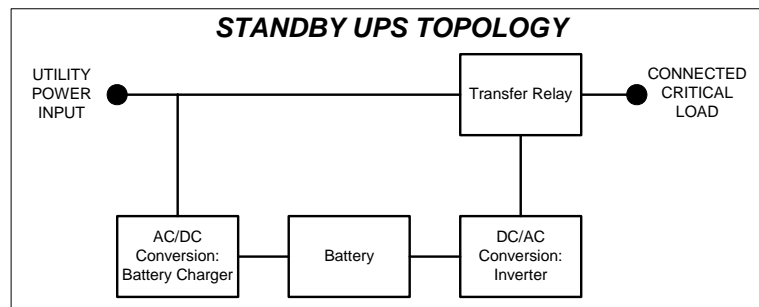


Figure 1

Figure 1 illustrates the basic subsystems of a standby UPS. During “normal mode,” when utility power is available and within acceptable quality tolerances, the UPS supplies the connected load directly with utility power. A small amount of AC power is used to trickle charge the UPS’s battery. For example, small standby UPSs designed to power individual workstations (rated less than 1 kVA) typically have fixed losses associated with this battery charging between 5 and 10 watts. An illustration of normal mode for standby UPSs is shown in Figure 2.

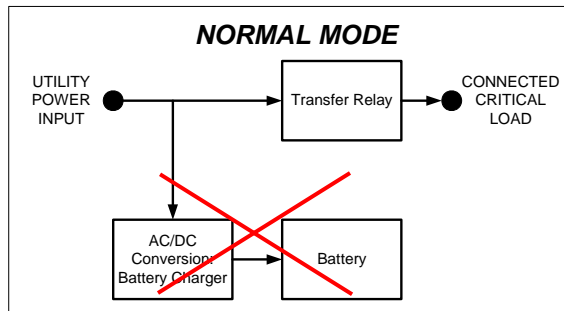


Figure 2

When there is an outage or the utility power goes beyond predefined quality specifications, the transfer relay isolates the load from utility power, and the UPS provides AC electricity through the battery and inverter in “stored energy mode.” This is illustrated in Figure 3.

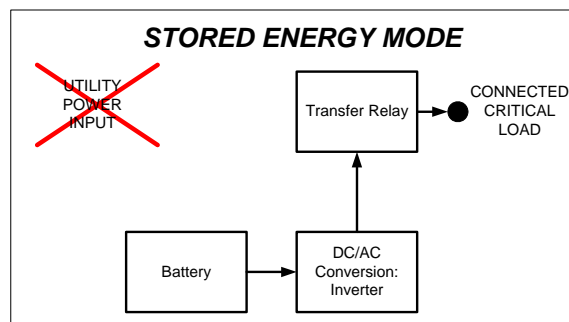


Figure 3

For most office and individual workstation backup applications, the UPS delivers unconditioned power to the connected load during normal mode. In some higher quality designs, utility power might be slightly conditioned upstream of the load so that it meets tighter voltage specifications.

Standby UPSs are most common in low-power applications (0 – 2 kVA) such as individual workstation backup power, small office server backup power, etc. Even though some higher quality designs might be capable of filtering utility power or protecting loads from power surges, the lack of complete power conditioning and load isolation generally make standby UPSs unsuitable for critical loads such as industrial facilities or data centers.

2.1.2 Overview of Inline or Line-Interactive UPS Topology

So-called line-interactive UPSs differ from standby UPSs because they are able to provide improved conditioning to utility power by interacting with the incoming electricity. They achieve this by placing inverter/battery charging circuitry or transformers in parallel with the AC utility signal. This design allows a line-interactive UPS to compensate for over- or under- voltages in the incoming utility power, but often allows commode mode noise and perturbations in AC frequency to pass through to the load. The line-interactive UPS has characteristics of both the standby and the on-line topologies. Operation is in a “conditioned” normal AC line mode (see Figure 9). If input voltage gets too high or too low, the line-interactive UPS may be able to correct for this condition without going to battery. This voltage regulation is accomplished through the use of either a tap-switching transformer, a ferroresonant transformer, or inverter magnetic components. Some line-interactive units are capable of supporting the load with no break in output when transferring to and from battery power. Other models are really a hybrid of the standby type, featuring voltage regulation but having a short break in output while transferring to and from battery.

The line-interactive topology is illustrated in Figure 4 below.

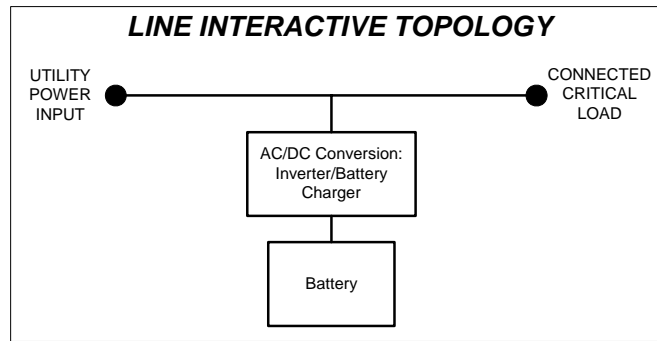


Figure 4

As with a standby UPS, a line-interactive UPS allows utility power to flow to the connected load in normal mode, but because of the parallel configuration of the inverter/battery charger in this topology, the UPS can condition incoming utility power for over- or under-voltages by interacting with the AC main. Small perturbations in frequency and common mode noise usually cannot be filtered out. The battery is continuously trickle charged during normal operation through the inverter/battery charger similar to the standby UPS topology. The normal mode of operation for a line-interactive UPS is illustrated in Figure 5 below.

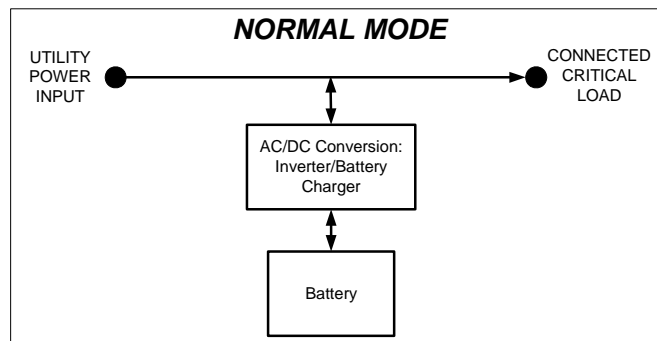


Figure 5

When utility power is unavailable or reaches unacceptable limits, a line-interactive UPS, like a standby UPS, will enter stored energy mode. The UPS disconnects the load from utility power and reroutes this load with a static switch to backup AC power, provided by the battery through the inverter. This is illustrated in Figure 6 below.

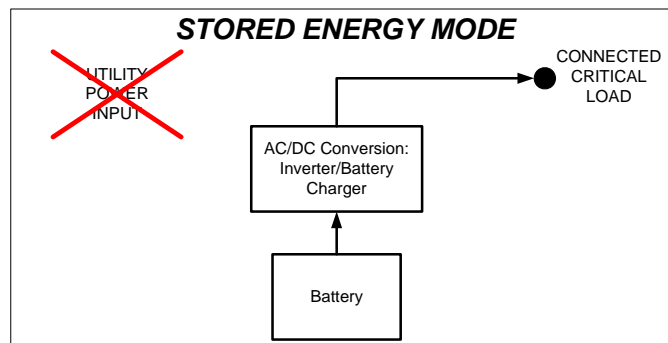
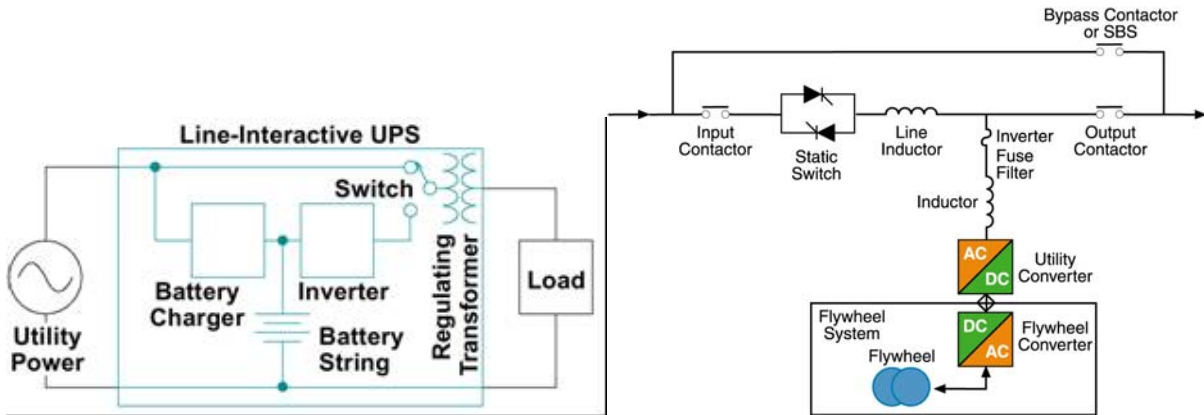


Figure 6

During normal mode, utility power is conditioned and delivered to the critical load through the rectifier and the inverter. The details of this conversion may vary from model to model, but the end goal is the same as with the double-conversion topology: keep the critical load isolated from unconditioned utility

power at all times. We can show two line interactive topologies in order to show these different details of conversion, as below.



Source: EPRI Solutions, based on APC design

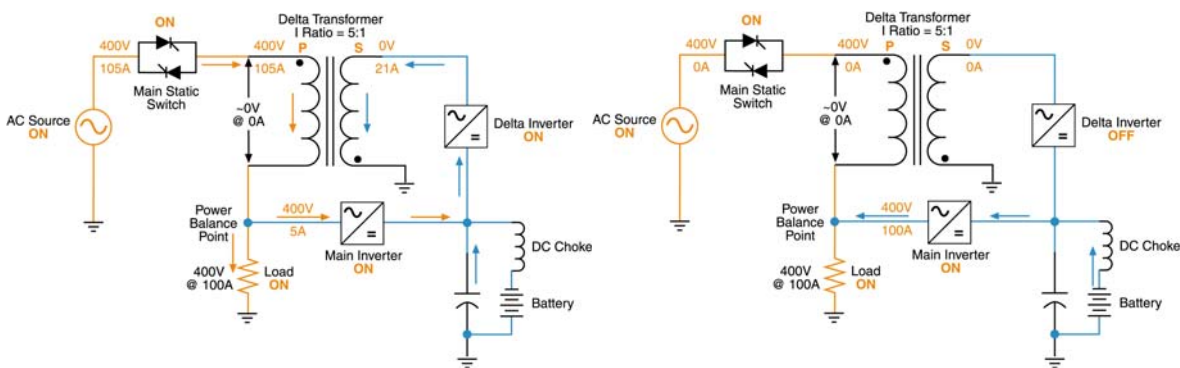
Figure 7

In Figure 7 above, (left) the regulating transformer is a ferroresonant type that maintains its output voltage when its input is *sagged* – for interruptions, its output will collapse but will still manage to shorten the duration of the interruption seen by the load. One of the flywheel-based designs uses what would be classified as a line-interactive topology but accomplishes this in a slightly different way. This design (right) controls reactive current flow through the line inductor to buck or boost the output voltage as necessary. The static switch opens to isolate the load from the line, but there is no transition seen by the load.

2.1.3 Overview of Delta-Conversion Topology

The delta-conversion topology is a proprietary topology developed by a single manufacturer, sold in sizes large enough to accommodate multiple racks or even an entire data center room. It is a unique design compared to line-interactive and double-conversion topologies. Delta-conversion UPSs use a special transformer configuration to interface between the load utility power, with a “delta” inverter in the transformer secondary to regulate input current and power. With this configuration, the UPS can regulate the magnitude, wave shape, and power factor of the current supplied at the UPS input, while still controlling the voltage very accurately at the load. This results in effective load isolation and very high power factor at the UPS input.

The basic modes of operation for delta-conversion UPSs are shown below (Figure 8).



Source: EPRI Solutions, based on APC design

Figure 8

Under normal conditions, (left) the Delta inverter and the Main inverter are both on, so that the delta inverter controls the transformer secondary current, including its magnitude, phase and wave shape. Based on the turns ratio, the primary current, (which feeds the load), is controlled and maintained. In the event the utility is removed, (right) the delta inverter turns off, forcing the primary current to zero, and the battery discharges through the main inverter to support the load. In this way, the transition time experienced by standby and line interactive topologies is avoided, and performance is on par with double conversion units. In addition, the efficiency is better than double conversion, since the delta inverter running in normal operation is sized for secondary current, much smaller than primary current. The main inverter in this condition is only conducting a small amount of leakage current required by the secondary circuit.

2.1.4 Overview of Double-Conversion or Online UPS Topology

Online or double-conversion UPSs are the most commonly used battery-based UPSs in data center environments because they are capable of completely isolating sensitive IT loads from unconditioned utility power. They receive their name, predictably, because they convert unconditioned utility power two times under normal operating conditions: first from AC to DC electricity and then back again from DC electricity into a highly conditioned AC signal. Double-conversion UPSs always provide the load with a high quality, conditioned AC signal, even during normal operation when utility power is available (Figure 9). For this reason, double-conversion UPSs are more common in high-availability, high-power mission critical applications such as industrial facilities and data centers. Typical output power ratings used in a data center setting for these types of UPSs range from 10 to over 1,000 kVA.

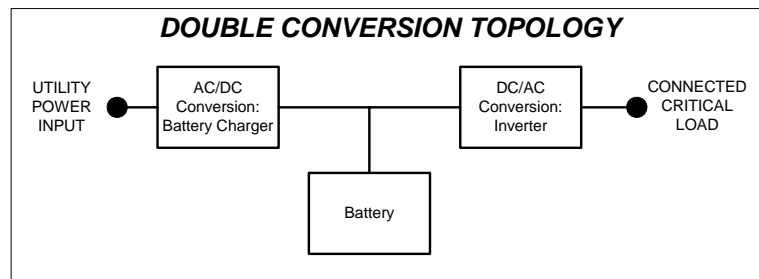


Figure 9

During normal mode, utility AC power is converted into DC power and is then converted back to a conditioned AC output for critical loads. During this phase of operation, the battery receives a continuous trickle charge to keep it at maximum charge. A flow chart of this mode of operation is illustrated in Figure 10 below.

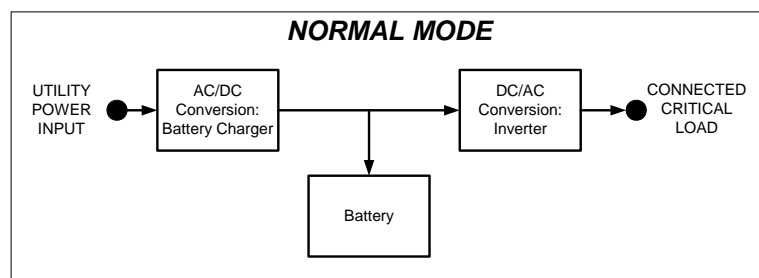


Figure 10

When utility power goes beyond acceptable limits, the double-conversion UPS enters stored energy mode, which is identical to stored energy mode in the line-interactive topology. The inverter draws DC power from the battery and continues to put out a conditioned AC signal to critical loads. In other words, the output of a double-conversion UPS is *always* conditioned. This often makes double-conversion UPSs the topology of choice for highly sensitive loads in mission critical applications, although delta conversion line-interactive UPSs are making inroads.

2.1.5 Overview of the U.S. UPS Market (Data Center/IT Sector)

We recognized early in the process of determining the size of the UPS market and the installed stock of UPSs in the data center/IT sector that new market research would be required to provide an accurate estimate. The tech boom of the late 1990s translated into rapid development of data center facilities around the country and along with it, increased purchases of UPSs to help protect mission-critical IT loads; however, this growth did not sustain itself. The “bursting” of the “tech bubble” in 2001 combined with market declines associated with the September 2001 terrorist attacks on the United States had a dramatic effect on the data center/IT sector for UPSs, which experienced two consecutive years of negative growth.⁷ Many earlier estimates of UPS stocks based their calculations on the rapid tech boom growth curve of the late 90’s and were unable to take into account the effects of the tech downturn in 2001. Figure 11 illustrates just how serious the downturn was for the UPS market in the U.S. Sales are not expected to recover to pre-2001 levels until after 2006.⁸

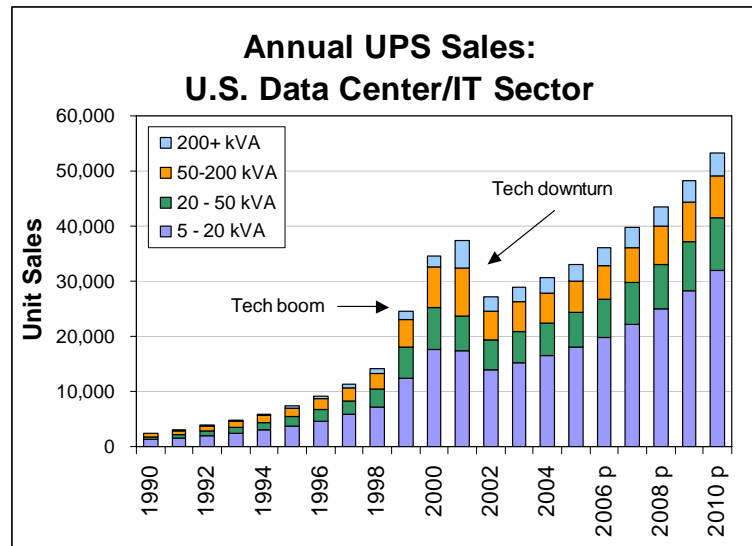


Figure 11

To provide a more accurate estimate of UPS stocks, it was necessary to take the tech downturn of 2001 into account by obtaining up-to-date UPS sales figures. We have chosen to utilize a Frost and Sullivan report published in 2004 entitled *World UPS Market* to gain access to more current market information. In general, the report describes a two-year period of negative growth in the global UPS market followed by a rebound in sales and renewed growth in 2003. The report cites a few trends that contributed to this growth, including an upswing in IT spending/development and lessons learned from the costly 2003 northeast blackout about the reliability of the aging U.S. power grid.

We were able to dissect the report to determine what percent of UPS sales occurred in the U.S. data center/IT sector and, from these sales figures, built a model to estimate the stock of installed data center/IT UPSs in the nation. The report provided detailed sales information for several different ranges of UPS sizes, measured in kVA, and broke down percent of sales based on geographic region (allowed us to characterize North American market) and end-user market (allowed us to characterize the percent of sales dedicated to data center/IT applications).

⁷ Frost and Sullivan. *World UPS Markets*. 2004. p. 1-6

⁸ Data source: Arthur D. Little for sales data from 1990 – 2000, Frost and Sullivan for sales data from 2001 – 2010.

Several basic assumptions were required in order to estimate unit sales in the U.S. data center/IT sector. First, we assumed that the U.S. market comprises 90% of all North American UPS sales, in keeping with past estimates made for the U.S. Department of Energy by Arthur D. Little on the size of the U.S. UPS market.⁹ Secondly, we excluded UPSs below 5 kVA in our analysis because these units are small enough that they would not likely be used to support data center equipment. Thirdly, we assume that the percent of sales by end-user market presented in the Frost and Sullivan report apply to all geographic regions. In other words, we assume that data center/IT spending is relatively similar in all regions of the world and particularly that the global data center/IT market share presented by Frost and Sullivan is representative of purchasing behavior in the United States. Finally, in order to arrive at an estimate of the stock of installed UPSs in the data center/IT sector, we assumed that the average lifetime for a UPS is 10 years. This number is based on estimates provided in both the Frost and Sullivan 2004 market report as well in the aforementioned Arthur D. Little report on office equipment energy use.

The sales data presented in the Frost and Sullivan report only begins in 2001, and because we have assumed an average UPS lifetime of 10 years, it was necessary to obtain earlier sales data provided in Arthur D. Little's 2002 analysis to be able to sum the cumulative sales over 10 years. Our stock estimate, thus, begins in the year 1997 and is projected out to the year 2010, even though our supplemental sales data begins in 1988.

We estimate that there are about 225,000 UPSs in operation in the U.S. data center/IT sector today, about 15% of which reside in the State of California. Half of the national UPS stock are units in the 5.1 to 20 kVA range of apparent power output. We assume that many of these units would not be used to power entire data center facilities but rather could be used to support individual server racks and the like. UPSs in the last two categories – 50.1 to 200 kVA and 200+ kVA – would likely be used to power entire data center facilities. Table 1 provides an approximate outline of the current distribution of UPSs grouped by their rated output power along with typical price ranges for each category.¹⁰ Although we expect the stock numbers of UPSs to continue to rise, we do not expect the percent distribution of these categories of UPSs to change very much in the near future. Figure 12 plots the expected growth in the stock from the years 1997 through 2010. The dashed line indicates the projected growth in stock that would have occurred had the tech boom of the late 90's continued past the turn of the century.

Table 1: Summary of UPS Stock, U.S. Data Center/IT Sector, 2004

	Stock of Installed Units	% of Stock	Typical Price Range (USD)
5.1 to 20 kVA	114,431	51%	\$4,000 - \$16,500
20.1 to 50 kVA	46,159	20%	\$8,000 - \$37,000
50.1 to 200 kVA	46,639	20%	\$19,000 - \$98,000
200 + kVA	19,080	9%	\$30,000 - \$207,000

⁹ K. Roth, F. Goldstein and J. Kleinman. *Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings, Volume I: Energy Consumption Baseline*. Cambridge, MA: Arthur D. Little, 2002. p. 84

¹⁰ Source of prices: Frost and Sullivan.

Estimated UPS Stock U.S. Data Center/IT Sector

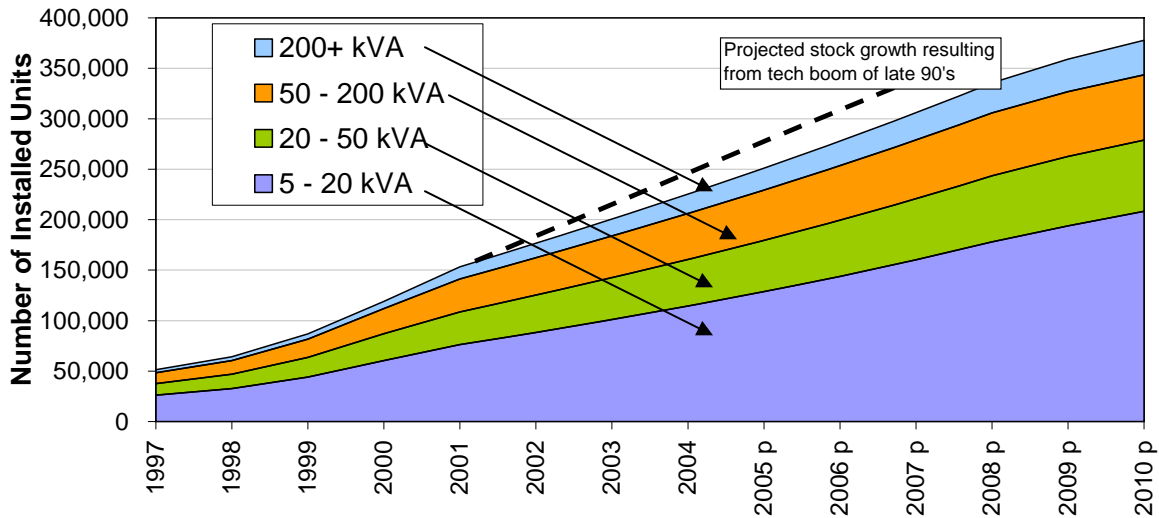


Figure 12

Using the estimated stock data presented in Figure 13 and a few assumptions about the composition of the UPS market, we have further estimated the basic composition of the UPS stock by design topology. Based on information from the 2004 Frost and Sullivan market report, the 2002 Arthur D. Little report on office equipment energy use, and common practices in data centers, we assume that there are essentially no standby UPSs in use in the data center/IT sector. This topology simply does not provide the type of availability that the market requires for protecting mission-critical loads.

Finally, it was necessary to make some assumptions about the split in sales between the double-conversion/online and offline/line-interactive topologies. For UPSs in the 5.1 to 20 kVA category, we assume that 81.3% of the stock is double-conversion and 18.7% is line-interactive, based on past sales data.¹¹ For UPS sizes above 20 kVA, some past literature has assumed that all sales are in double-conversion topologies; however, APC has been making inroads with its delta-conversion in this power range, and our analysis acknowledges this fact by breaking out delta-conversion UPSs as part of the UPS stock above 20 kVA. Since APC is the only manufacturer producing delta-conversion UPSs in sizes above 20 kVA (the company manufactures delta-conversion units up to 1000 kVA in size) and since the company sells exclusively delta-conversion units at these sizes, we assume that the percent of delta-conversion UPSs in the stock matches APC's market share for the 20.1 – 50 kVA, 50.1 – 200 kVA, and 200+ kVA market segments.¹²

Table 2 outlines our assumptions about the relative distributions of various UPS topologies in the U.S. data center/IT sector. The estimates are presented in terms of percent of stock and number of installed units for several size categories. Figure 14 graphically illustrates how the stock breaks down by topology. For UPSs larger than 5 kVA used in data centers, the market is essentially dominated by one UPS topology, the double-conversion/online topology. The data labels in the two right-hand columns indicate the estimated number of delta-conversion units installed in the 50 to 200 kVA and 200+ kVA ranges.

¹¹ Taylor J. and Hutchinson J. "Uninterruptible Power Supplies, Parts 1, 2, and 3." *Electronic Buyers News Power Supplement*. September 2000.

¹² Frost and Sullivan. pp. 6-19, 7-18, and 8-20.

Table 2: Composition of U.S. Data Center/IT Sector UPS Stock by Topology

	Line-interactive (% stock) (# units)	Delta-conversion (% stock) (# units)	Double-conversion (% stock) (# units)	Flywheel (% stock) (# units)	TOTAL UPSs in Stock (# units)
5.1 - 20 kVA	18.7%	0%	81.3%	0%	114,431
	21,399	0	93,033	0	
20.1 - 50 kVA	0%	11.1%	88.9%	0%	46,159
	0	5,124	41,035	0	
50.1 – 200 kVA	0%	2.0%	97.5%	< 0.5%	45,639
	0	913	44,498	< 228	
200+ kVA	0%	0.6%	99.0%	< 0.5%	19,080
	0	114	18,889	< 95	

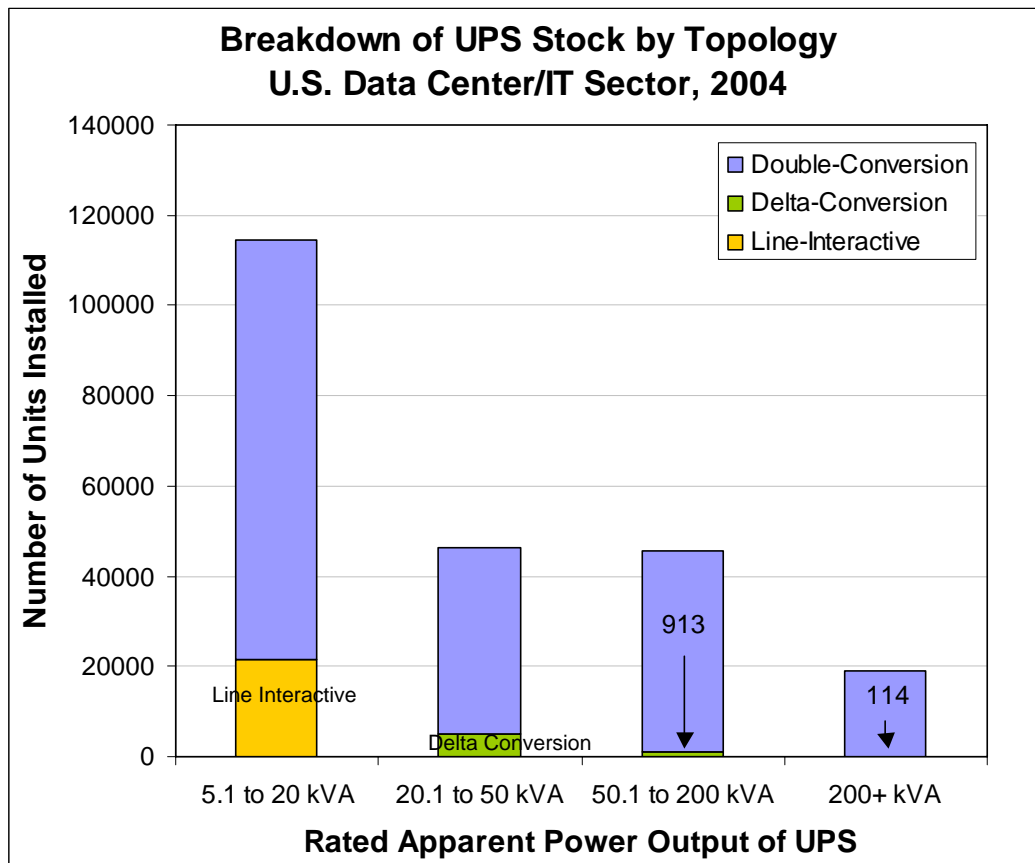


Figure 13

2.2 UPS Test Protocol

One of our most important tasks was to develop clearly written UPS efficiency test protocols that were to be used either in a factory/laboratory setting or a field setting to gather credible, comparable data on UPS efficiency, tested at resistive and non-linear loads. The International Electrotechnical Committee (IEC) has created a detailed and comprehensive international standard (IEC 62040-3) to

guide manufacturers in testing their UPS systems,¹³ which devotes 2 pages to UPS efficiency testing. The standard discusses the definition of efficiency, discusses approved linear and non-linear loads, and provides a suggested format for recording the data, but unfortunately only recommends testing for efficiency at 50%, 75% and 100% of the UPS's rated active and apparent power outputs. Past research indicates that UPSs in a data center environment typically operate at or below 50% of their rated active power output, where the efficiency of the system can decline significantly, and therefore it was of particular interest to our team to develop a UPS test method that was more specific about measuring efficiency at part loads.

2.2.1 Factory/Laboratory UPS Efficiency Test Method

In order to gather data to benchmark data center UPS efficiency, encourage manufacturers to measure UPS efficiency at part loads, and to raise awareness of the importance of properly characterizing UPS efficiency, a test procedure was developed that exclusively guides manufacturer testing of UPSs that are designed for data center and mission-critical applications, entitled *Guidelines for Manufacturers' Testing to Benchmark Data Center UPS Efficiency*.¹⁴ The test method was circulated widely to UPS manufacturers in mid-2004 to spur on an effort to obtain more accurate UPS efficiency data. Five of these manufacturers have agreed to participate in our project by sharing the results of UPS efficiency tested guided by our test procedure. These results are described in Section 1.3 below.

The test method is largely based on the IEC 62040-3 standard and references the document regarding the appropriate load characteristics for resistive and non-linear loads, the definition of UPS efficiency, and the types of UPS testing typically performed by manufacturers. The test method developed expands upon IEC 62040-3 by recommending efficiency testing at no load (stand-by), 10%, and 20% load conditions, which has allowed us to better characterize the shape of UPS efficiency curves at part loads. The manufacturer test method also provides several pages of useful tables and forms to aid manufacturers in reporting detailed efficiency and power quality information about their models.

2.2.2 Field UPS Efficiency Test Method

In the absence of any formal test procedure for measuring the power consumption and efficiency of field-installed UPSs, we developed a field test method as a compliment to our manufacturer guidelines, entitled *Guidelines for Field Data Collection to Benchmark Data Center UPS Efficiency*.¹⁵ Because the design implementation of UPSs in the field can vary and because there are a number of different ways that power/efficiency data could be obtained from an installed unit, the method needed to provide explicit guidance for many different types of configurations. The document provides guidelines for measuring field UPS efficiency through a number of means, including:

- Visual measurements from a front-panel display
- Data logging using UPS monitoring software
- Data logging using permanently connected power meters
- Data logging using a building energy management system
- Direct measurement with power meters

This field test method has been used as a guideline for testing UPS loading and efficiency at several data centers and other mission-critical facilities in the U.S. The results of this testing are presented in Section 1.3 below.

¹³ IEC 62040-3, *Testing Procedures for UPS Systems*. International Electrotechnical Committee. April 30, 2004.

¹⁴ *Guidelines for Manufacturers' Testing to Benchmark Data Center UPS Efficiency*. June 7, 2004.
<http://hightech.lbl.gov/Documents/UPS/UPS_Efficiency_Measurement.pdf>

¹⁵ *Guidelines for Field Data Collection to Benchmark Data Center UPS Efficiency*. June 7, 2004.
<http://hightech.lbl.gov/Documents/UPS/UPS_Field_Data.pdf>

3. UPS Efficiency Data

3.1 UPS Efficiencies

EPRI Solutions, Ecos Consulting, and Lawrence-Berkeley National Laboratories have gathered measurements on the efficiency of a variety of commercially available UPSs under controlled, factory conditions as well as field conditions using the test methods described above in Section 1.2.

The efficiency of a UPS, as defined by the International Electrotechnical Committee, is “the ratio of (active) output power to (active) input power under defined operating conditions,”¹⁶ where defined operating conditions refer to a specific percent load and load type (linear/resistive versus non-linear). We are primarily concerned with UPS efficiency while the unit is in its normal mode of operation, when utility power is within accepted tolerances and the UPS can use this power to feed the critical load. UPSs spend the vast majority of their time in this mode of operation, so this is where efficiency can have the most impact on energy use and total cost of ownership.

The efficiency of a UPS can be impacted by a number of factors that can be controlled during testing, including the percent of load being supplied by the unit and the power quality characteristics of the load (resistive vs. non-linear loads). We describe how these factors can affect UPS efficiency below, followed by an analysis of the data that manufacturers have provided.

3.1.1 UPS Efficiency and Percent Load

UPS efficiency varies with the amount of active power being supplied to a load and tends to decrease steadily at part loads. UPS efficiency curves are usually relatively flat between loads of 50% and 100%, but decrease steadily below 50% and precipitously below 20%. The efficiency of a UPS can vary by 7 to 12 percentage points depending on whether it is partly or fully loaded. The lower efficiency at part loads can result significant conversion losses compared to the same UPS operating at close to full load. Figure 14 illustrates a typical load curve of a UPS in our current data set. Note that there is no “typical” operating range for any UPS.

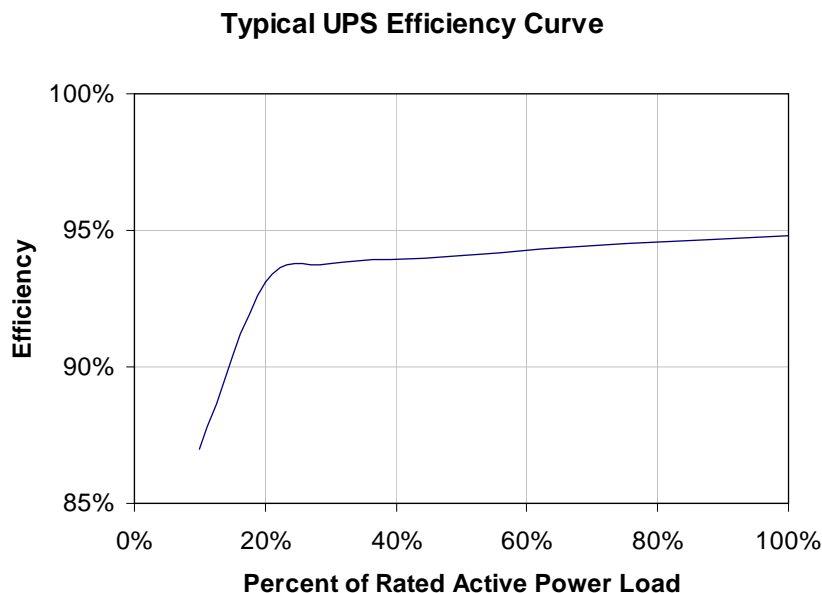


Figure 14

¹⁶ IEC 62040-3... p. 52

UPSs in mission critical environments typically operate at points between 30% and 50% of nominal and in situations where the data center is not fully populated this can be even lower resulting in very large inefficiencies. Therefore knowing the efficiency of a UPS at loads below 50% is extremely important in estimating real world energy use. As the figure above shows, efficiency tends to decline steadily below about 50% load, so understanding the behavior of the efficiency curve at these low loads is crucial to making accurate estimates of energy use and operational costs. Unfortunately, manufacturers usually only report UPS efficiency at 100% load, with occasional reporting at 50% and 75% load. This is one of the main reasons why the attempt to characterize UPS efficiency across a wide range of loading was so important.

There are a number of reasons why data center UPSs operate at part loads, one of which is that UPSs are commonly sized to meet the maximum utilization of space in a data center, even though maximum utilization rarely occurs.¹⁷ This results in over sizing of the UPS for the actual amount of installed mission-critical infrastructure.

UPSs in mission-critical facilities operate in redundant configurations, which also act to reduce the theoretical percent loading of the unit to below 50%. In a redundant UPS configuration, the load of the facility is shared between two or more UPSs so that, if one of the units fails, the other(s) will still be available to supply the facility's critical load. Any one UPS in a redundant configuration must be capable of supplying the *entire load* required by the facility. Because the load is shared equally between two or more of these full-sized UPSs, the largest load that a data center UPS could theoretically experience under most conditions is 50% of its rated active power load. A more thorough discussion of redundant UPS configurations and their effects on overall efficiency is provided in Section 1.4.

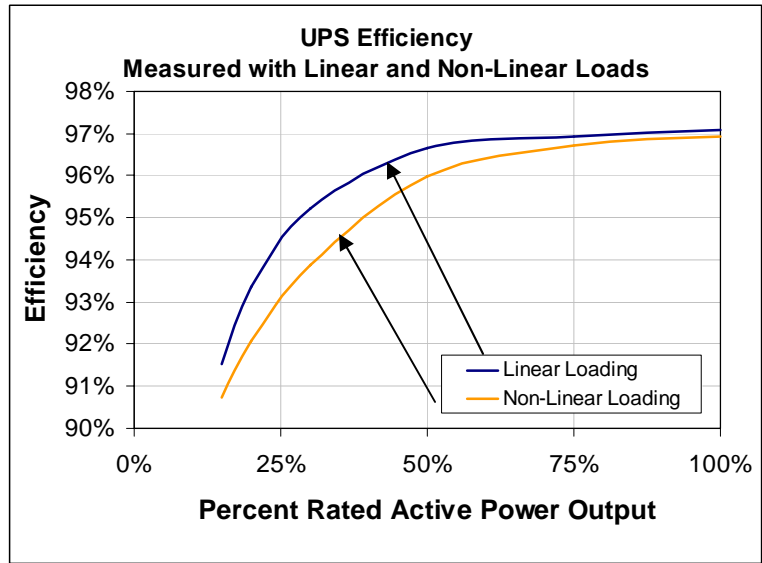
3.1.2 UPS Efficiency and Load Characteristics

The characteristics of a UPS load can have a noticeable effect on the measured efficiency of the unit. Manufacturers often report UPS efficiency tested with resistive or linear loads; however, UPSs in many mission critical installations usually have to power highly non-linear loads, like the switch mode power supplies (SMPS) used in desktop and low-end servers. These power supplies often have poor power quality characteristics, including lower power factor (0.7) and high total harmonic distortion (in excess of 100%). Low power factor loads like a SMPS cause the UPS to deliver more peak current, changing the way in which the system is loaded. A non-linear load like a SMPS can also create harmonic currents in distribution wiring, which can dramatically increase the amount of current in distribution wiring if the UPS does not include special filters.

These different characteristics of non-linear devices change the way the UPS is loaded and can have an impact on efficiency. Figure 15 shows how the measured efficiency of a given UPS in our data set differed by 1% to 2% between a linear and non-linear load. The additional current demands of the non-linear load, drives the efficiency down by a small but measurable amount. This is important because, as mentioned earlier, UPSs in mission critical environments very often support non-linear loads. Testing with non-linear loads will be most indicative of operational efficiency.

¹⁷ Madsen J. "Continuous UPS Availability: How Important is it to Your Company?" *Energy User News*. August 11, 2000. <http://www.energyusernews.com/CDA/ArticleInformation/features/BNP__Features__Item/0,2584,14489,00.html>

Figure 15



3.1.3 Results of UPS Efficiency Testing: Field Measurements

Results of recent data center UPS efficiency testing support the claim that part load efficiency – efficiency measured at less than 50% of a UPS’s rated active power load – is of greatest importance in estimating the operational costs and total cost of ownership of a UPS. EPRI Solutions, Ecos Consulting, and LBNL benchmarked the efficiency of UPSs operating under real world conditions in data centers. Ecos and EPRI Solutions built upon research conducted in 2001 in which LBNL benchmarked the performance of 14 data center facilities along with their UPS equipment. We added the UPSs from two new UPSs, both supporting IT loads in the EPRI Solutions data center.

Over 80% of the UPS systems measured were operating at below 50% of their rated load. Many of these units were partially loaded because they were sharing their load in a redundancy configuration, which we denote in the chart; however, four of the partially loaded units were not operating redundantly. They were simply oversized and underutilized. Figure 16 displays the results of our testing, plotting both the percent of the UPS’s nominal load utilized in the field as well as efficiency of the UPS under real world conditions. Table 3 summarizes the average load factor and operational efficiency for UPSs tested in the field.

Table 3: Average Loading and Efficiency of UPSs in the Field

Average Load Factor of UPSs Tested in the Field	Average Efficiency of UPSs Tested in the Field
37.8%	85.2%

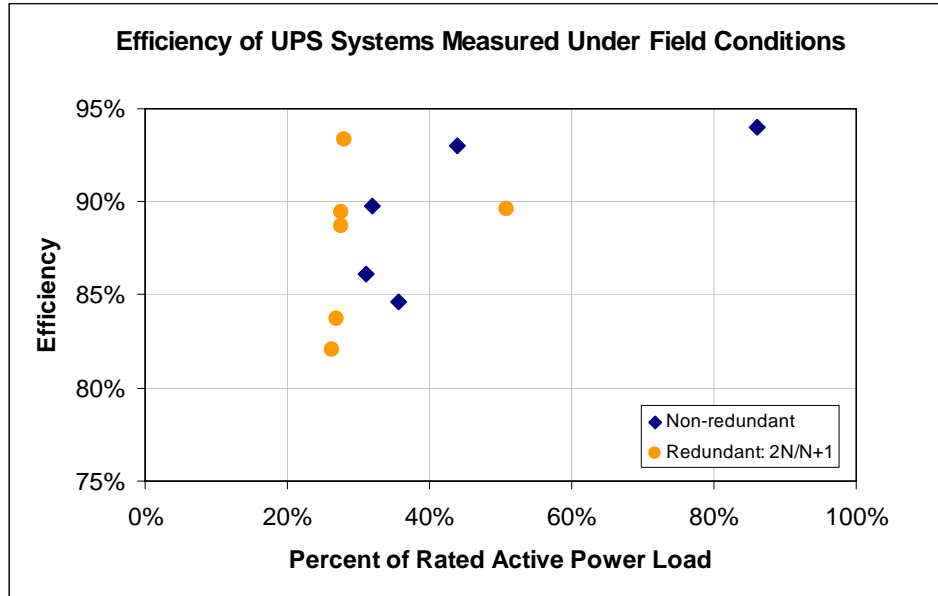


Figure 16

3.1.4 Results of UPS Efficiency Testing: Factory/Laboratory Measurements

Our results show that UPS efficiency varies significantly between different models. At low load conditions (20% of rated active power load), our findings indicate that UPSs are anywhere from 81% to 93% efficient. At full load conditions, our measurements show that UPSs are anywhere from 86% to almost 98% efficient. Although these figures span a range of less than 15 percentage points, upgrading from one of the lower efficiency UPSs to a model at the high end of the range could cut UPS operating costs by over a half. The range in UPS efficiency also suggests that there may be an opportunity to encourage higher efficiency designs through product labeling and standards, a technique which has been applied in the past to other high-power devices but which has not yet been adopted for UPSs.

Figure 17 summarizes UPS efficiency test results that have been provided by manufacturers. The results presented in this chart were conducted under factory conditions with linear loads and were guided by the manufacturer test guidelines mentioned earlier in the report. A second chart, Figure 18, presents the results of testing conducted using non-linear loads.¹⁸ The efficiency curves are identified by their general topology, including flywheel, double-conversion, and delta-conversion.¹⁹

¹⁸ Note: some manufacturers reported efficiency measurements conducted under both linear and non-linear loads; however, a large number of the UPSs were only tested with resistive loads, hence the larger number of lines in Figure 17.

¹⁹ Note: due to the scope of this project, only UPSs common in data center applications were tested. Standby or offline UPSs are uncommon in high-power, mission critical applications and, thus, were not tested.




Figure 17

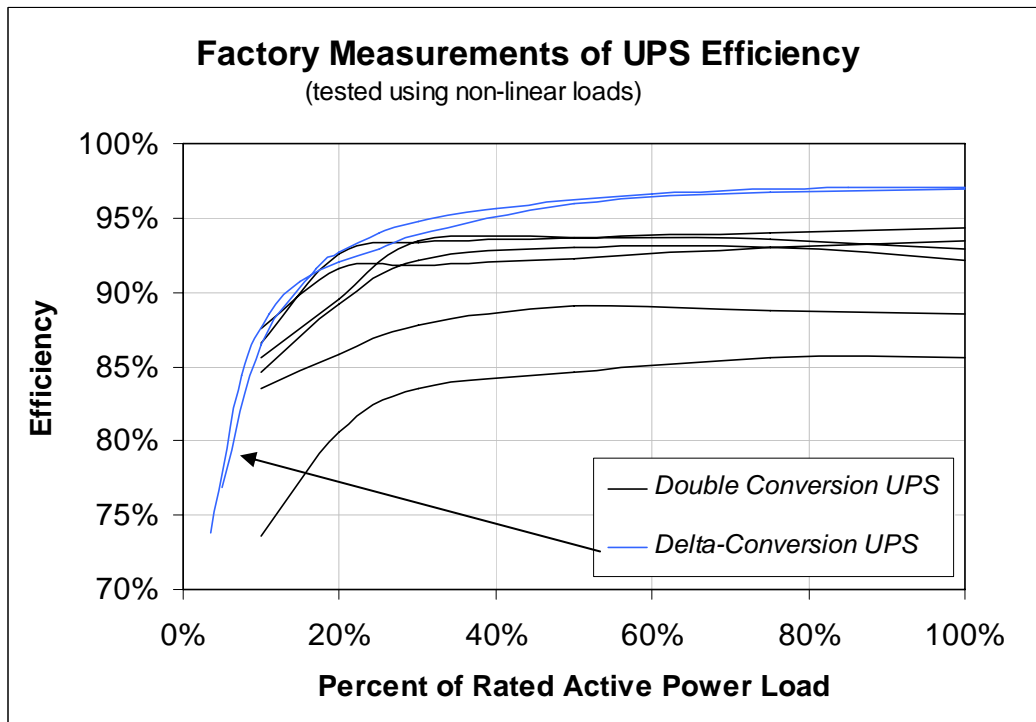


Figure 18

Our results indicate that there can be a wide variation in efficiency even between UPSs of the same general topology design, such as the double-conversion UPSs plotted in Figure 18 above, which vary from 86% to almost 95% efficiency at full load. The table below summarizes the characteristic efficiency of a number of UPS topologies at various percent load conditions and shows the average efficiency for all of the UPSs measured.

Table 4: Characteristic Efficiency of UPS Topologies

UPS Topology	Efficiency at 25% Load	Efficiency at 50% Load	Efficiency at 75% Load	Efficiency at 100% Load
<i>Delta-Conversion</i>	93% - 94%	96% - 97%	97%	97%
<i>Double-conversion</i>	81% - 93%	85% - 94%	86% - 95%	86% - 95%
<i>Line-Interactive</i>	unknown ²⁰	97%	98%	98%
<i>Standby</i>				
AVERAGE OF ALL UNITS	86%	89%	90%	90%

Due to the scope of this research effort, it was impractical to determine what specific elements of a model's design contributed to its measured efficiency; however, our results indicate that some broader design decisions, such as the general UPS topology, can be indicative of efficiency. In particular, flywheel and delta-conversion UPSs exhibited best-in-class efficiency compared to the double-conversion UPSs that we measured. The three line-interactive UPS designs and two delta-conversion designs had consistently higher efficiency than the double-conversion models for loads above 30% of nominal and similar performance in the 10% to 20% range compared to the highest efficiency double-conversion models. Research on the efficiency of both line-interactive and delta-conversion UPSs should continue as manufacturers continue to improve their designs and diversify product lines. Nevertheless, the high efficiency of these UPS topologies in our current data set suggest that they deserve serious consideration by facility managers looking to lower operating costs in mission critical facilities that utilize large UPSs (greater than 50 kVA in capacity) and policy makers wishing to identify the most efficient technologies.

Desire for high efficiency should naturally be balanced with concerns over load isolation and reliability. Although the delta-conversion UPSs that we tested performed better in regards to overall efficiency compared to double-conversion units, some manufacturers have argued that delta-conversion UPSs do not provide the same load isolation as a "true" double-conversion UPS. This report does not attempt to answer the question as to which topology is the overall best choice for mission critical applications. We simply report the observed efficiencies of the various topologies, and facility managers will ultimately need to judge whether a particular UPS can handle load disturbances and provide an acceptable level of load isolation.

3.1.5 Investigation of Special High Efficiency Modes in Select UPSs

A handful of UPS manufacturers now advertise double-conversion units with programmable high efficiency modes, sometimes referred to as "power-saver" or "eco" modes. We identified the following units in Table 5 as having some means of enabling higher efficiency through a user control. Although most of the units occupy the sub-50 kVA range of output and would not be capable of powering entire data centers, the Chloride 90-NET models are available in sizes ranging up through 800 kVA, proving that such technology could already be applied to entire facilities.

²⁰ Data on flywheel efficiency was only presented down to about 30% load.

Table 5: Summary of UPS Models with High Efficiency Modes

UPS Manufacturer	Model	Power Range (kVA)	Mfg.-Reported Double-Conversion Efficiency	Mfg.-Reported Energy-Saver Mode Efficiency
MFR#1	9170+ Online	3 – 18	88%	97%
MFR#2	90-NET	60 – 800	92% - 93%	97%
MFR#2	70-NET	10 - 60	91% - 92%	97%
MFR#3	Galaxy 3000	10 – 30	85% - 86%	unknown

EPRI Solutions was able to test a 3 kVA model of the Powerware 9170+ Online UPS. The unit is, by design, a double-conversion UPS but offers a “power-saver” high efficiency mode that the manufacturer claims can boost efficiency from the 88% range to 97%. EPRI Solutions conducted controlled laboratory efficiency measurements on the 9170+ to validate the manufacturer’s claims. Although the high efficiency mode only allows for the reported 97% efficiency at full load, our tests found that the feature did enable 5% efficiency improvements across the range of percent loads, as illustrated in Figures 19 (resistive load testing) and 20 (non-linear load testing).

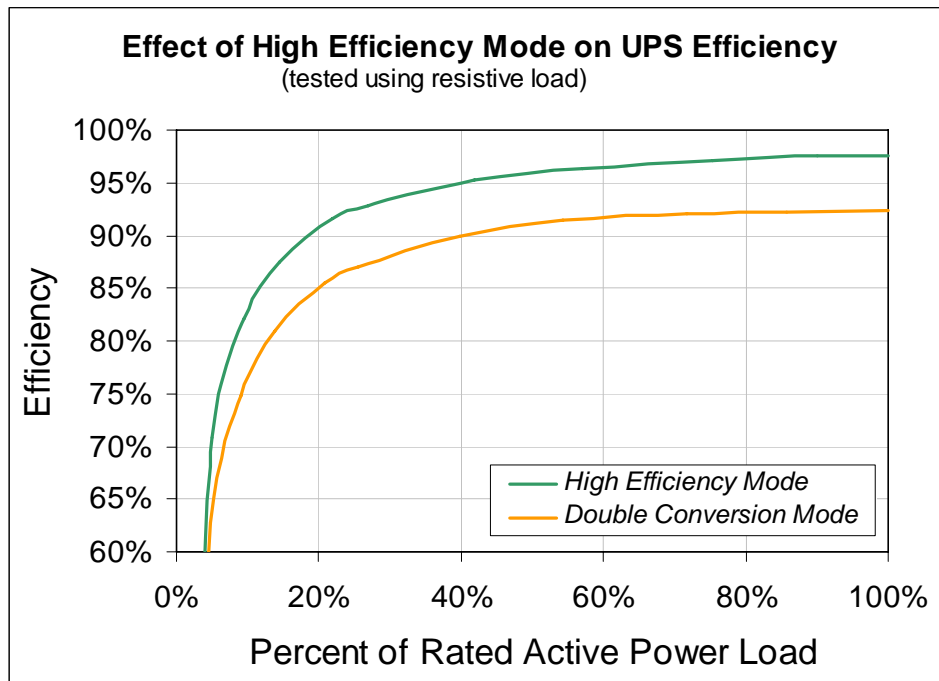


Figure 19




Figure 20

One of the potential drawbacks of enabling high efficiency modes in double-conversion UPSs is that the UPS no longer operates in double-conversion mode, potentially leaving mission critical loads more exposed to power abnormalities like voltage sags and line noise. It would seem that there is an inherent trade-off between high efficiency and complete load isolation, but is this trade-off significant enough that data center managers should choose not to enable high efficiency modes for fear of exposing their IT equipment to more power irregularities?

These concerns were investigated by conducting additional tests on the Powerware 9170+ Online UPS, exposing the unit to power abnormalities in both double-conversion and high efficiency modes. The tests sought to determine the response of the UPS output when the UPS input (i.e. utility power) experienced an 80% and 30% sag in voltage for various durations of time. In particular, these experiments were designed to answer questions such as:

- Is the output of the UPS seamless regardless of voltage sags at the input?
- Does the UPS transfer over to battery when it sees power abnormalities at the input?
- If the output of the UPS is interrupted, how long does the interruption last, and would this interruption pose a threat to mission-critical infrastructure like servers?

In general, the performance of the 9170+ in its normal double-conversion mode – what Powerware calls “Auto” mode – was flawless for both 80% and 30% voltage sags. The AC output maintained a constant frequency and voltage regardless of disturbances to the input whether the UPS transferred over to battery backup power. Figure 21 illustrates one case in which the UPS responds to a 1-cycle (16.7 ms) 30% voltage sag. The unit transferred to battery backup power in this case, but maintained a seamless AC output. The behavior shown in Figure 22 is generally indicative of the unit’s performance in 80% voltage sag tests as well, with the main difference being that, in the 80% sag tests, the UPS maintained a seamless AC output without transferring to battery backup power.

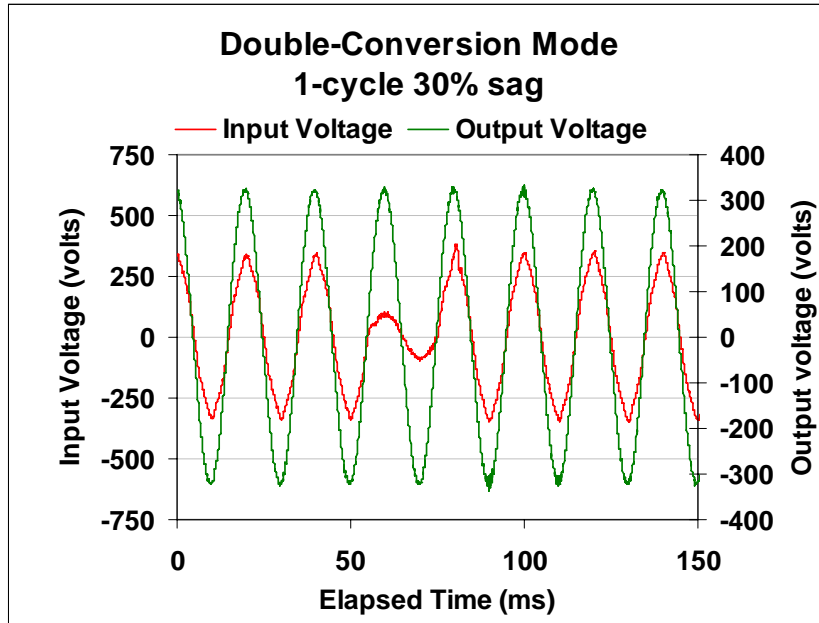


Figure 21

Unfortunately, the 9170+ does not appear to have the same “flawless” response to power abnormalities when its high efficiency mode is enabled. Take the case of the 1-cycle 30% voltage sag test. In double-conversion or “Auto” mode, the 9170+ was able to ride out this disturbance with no noticeable change to its output. As illustrated in Figure 22, the unit handles the disturbance much differently when high efficiency is enabled by the user. The output voltage first undergoes a 30% drop for one half of a cycle, followed by a 3/4-cycle disturbance while the UPS transfers over to its batteries for power. In total, this adds up to 20.9 ms of unconditioned power at the output of the UPS for a voltage disturbance that only lasts one AC cycle or 16.7 ms. We experienced the same type of behavior when conducting 80% voltage sag tests on the 9170+, except that the duration of the voltage sag on the UPS output was longer (2.5 AC cycles) before the unit made the transition.

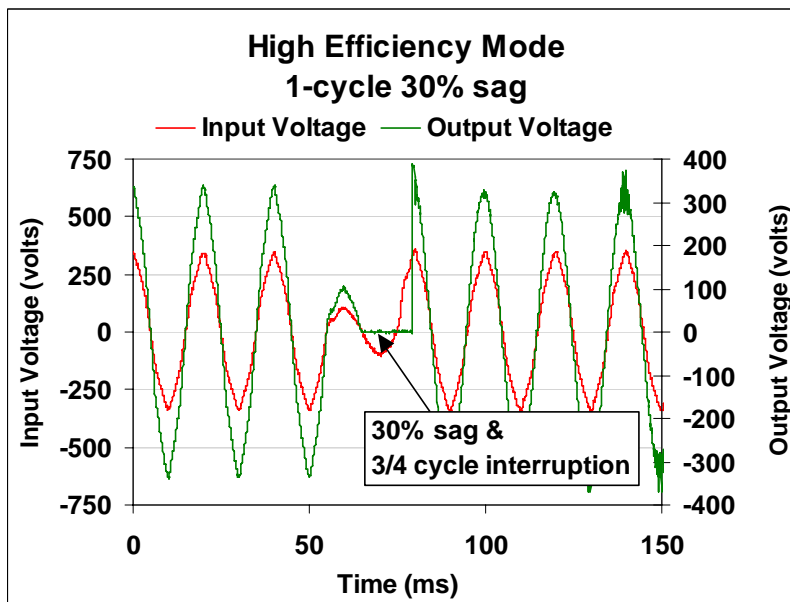


Figure 22

The tests not only indicated that the high efficiency mode would introduce delays in transferring to backup power, exposing connected loads to unconditioned power for brief periods of time, but they also suggested that the unit would have to transfer to backup power more frequently because it was unable to “ride out” small voltage disturbances. Recall that the 9170+, operating in double-conversion mode, was able to maintain a seamless output during the 80% voltage sag tests for disturbances up to 30 AC cycles or half a second in duration without transferring to battery power. When the same tests were performed with the unit set to its high efficiency mode, the unit switched to backup power for disturbances of 10 AC cycles or longer.

Tables 6 and 7 summarize all of the voltage sag testing performed on the 9170+ Online UPS operated in double-conversion and high efficiency modes, respectively. We have included a column that indicates whether a server power supply designed to Server System Infrastructure (SSI) specifications would be able to ride out the disturbance to the UPS output.²¹

Table 6: Summary of Voltage Sag Testing, Double-Conversion Mode

% Input Voltage Sag	Sag Duration (AC cycles & ms)	UPS Action	Length of Disturbance to Output (AC cycles & ms)	Disturbance Can be Handled by SSI PS
80%	1 cycle, 16.7 ms	No battery transfer	0 cycles, 0 ms	Yes
	10 cycles, 166.7 ms	No battery transfer	0 cycles, 0 ms	Yes
	30 cycles, 500 ms	No battery transfer	0 cycles, 0 ms	Yes
30%	1 cycle, 16.7 ms	Battery transfer	0 cycles, 0 ms	Yes
	10 cycles, 166.7 ms	Battery transfer	0 cycles, 0 ms	Yes
	30 cycles, 500 ms	Battery transfer	0 cycles, 0 ms	Yes

Table 7: Summary of Voltage Sag Testing, High Efficiency Mode

% Input Voltage Sag	Sag Duration (AC cycles & ms)	UPS Action	Length of Disturbance to Output (AC cycles & ms)	Disturbance Can be Handled by SSI PS
80%	1 cycle, 16.7 ms	No battery transfer	1 cycle, 16.7 ms	Yes
	10 cycles, 166.7 ms	Battery transfer	3.25 cycles, 54.2 ms	No
	30 cycles, 500 ms	Battery transfer	3.25 cycles, 54.2 ms	No
30%	1 cycle, 16.7 ms	Battery transfer	1.25 cycles, 20.9 ms	Yes
	10 cycles, 166.7 ms	Battery transfer	1.25 cycles, 20.9 ms	Yes
	30 cycles, 500 ms	Battery transfer	1.25 cycles, 20.9 ms	Yes

In general, the test results have led to the following conclusions about the 9170+ Online UPS. We stress that these conclusions should not necessarily apply to all of the UPS models that we have identified above that have a user-enabled high efficiency mode:

- The unit will tend to transfer to battery backup power more frequently when the high efficiency mode is enabled.
- Because the output is not fully conditioned, small or momentary abnormalities in power may pass through to the critical load when the unit is operating in high efficiency mode.

²¹ According to SSI, power supplies for servers must be able to ride out voltage sags less than 1 cycle in duration.

- The unit currently shows a fairly consistent 3/4 –cycle delay in transferring to battery backup power for large input power disturbances, combined with 1/2- to 2.5-cycle disturbances to the unit’s output before it recognizes the need to transfer to backup power.
- Most computers and servers tested in the past are able to ride through interruptions greater than a few cycles, so it is unlikely that the loads will be affected by these relatively slow transfers. We believe that it is a matter of educating the customer as to what is absolutely necessary to effectively protect their loads. In this way, change can be effected in the market that will allow higher efficiencies.
- It is unlikely that data center managers would be willing to operate a double-conversion UPS in high efficiency mode. This mode effectively turns a fully isolated double-conversion UPS into a less isolated line-interactive UPS. If customers are paying for the full performance of a double-conversion unit, they are likely to want all of the benefits that go along with it, including complete load isolation even though there is an energy penalty.²²
- Both standby and energy efficient modes for double conversion UPSs exhibit the “slow” transfer behavior. However, it should be noted that the conventional wisdom regarding the speed of these transfers is changing with today’s much faster switches and detection circuitry. It should be pointed out that the transfer is the sum of the detection time and the actual switching time. In many modern applications, these transfers are effected on the order of a few milliseconds. This all but eliminates the concern for shutdown of the critical computer loads. However, as stated before, it is still largely a matter of educating the customers.
- Another conclusion is that the data center operators should begin to consider some of the line-interactive topologies for their higher efficiencies, since their active voltage regulation offers the same performance as the double conversion units.

3.2 Redundancy Evaluation

Today’s facility managers expect previously unheard-of levels of availability from their UPS systems. Due to the expense of down time, world class facilities strive for UPS systems with six nines of availability (99.9999% available systems), meaning that mission-critical loads should have an available supply of power 99.9999% of the time. To put this in perspective, a UPS system with six nines of availability can only have 32 seconds of down time per year. Unfortunately, a single UPS system is typically capable of 99.92% availability, leaving a data center at risk of about 7 hours of down time every year.²³

Redundant UPS configurations help to secure higher levels of availability by preventing a facility’s power system from being interrupted by a single UPS failure. This is typically done by sharing the facility’s total load between two or three UPSs so that, even if one or two units simultaneously fail, there will still be another backup UPS to condition utility power and supply backup power in the case of power outages. Each individual UPS in a redundant configuration usually must be capable of supplying the entire critical load by itself to achieve the highest levels of availability. Because of redundancy configurations, a data center’s load is commonly shared equally between two or more UPSs that, individually, could power the entire facility by themselves. This load sharing reduces the percent load in individual UPSs and has the effect of decreasing their operational efficiency.

Below we provide an overview of some common types of UPS configurations along with an indication as to the level of availability that they provide and how the overall configuration affects efficiency.

3.2.1 Capacity or “N” Configuration

In the simplest UPS configuration, a single UPS supplies power to the entire data center facility. “N” refers to the predicted size of the load protected by the UPS. Since the configuration only requires

²² Personal communication with Tom Geist of EPRI Solutions. May 10, 2005.

²³ K. McCarthy. “Comparing UPS System Design Configurations.” American Power Conversion. 2004. < ftp://www.apcmedia.com/salestools/SADE-5TPL8X_R0_EN.pdf>

one UPS capable of supplying this load, the name for the configuration is simply “N”. This configuration is highly uncommon in mission critical facilities like data centers because there are no backup UPSs available to protect the load in case the main UPS fails or requires maintenance. As mentioned above, the estimated availability of a single UPS is 99.92%, meaning that the critical load will be unprotected from power abnormalities and failures for about 7 hours every year.²⁴

The main advantage of the “N” configuration is that it has no effect on the percent load of the UPS and, thus, no effect on the operational efficiency.

3.2.2 Isolated Redundant Configuration

The next level of complexity in the design of UPS systems is the isolated redundant configuration. This design is fairly similar to the “N” configuration except that it adds a second UPS to protect the critical load in the case that the main UPS fails in some way or needs to be bypassed for routine maintenance. Rather than sharing power between the two UPSs, the main UPS handles all of the critical load unless it fails or requires maintenance, in which case a static transfer switch (STS) quickly and automatically transfers the load over to the second UPS. The STS must be sufficiently fast in its switching (< 1AC cycle) so that the critical load receives fairly continuous power. Figure 23 shows a basic block diagram view of an isolated redundant configuration.

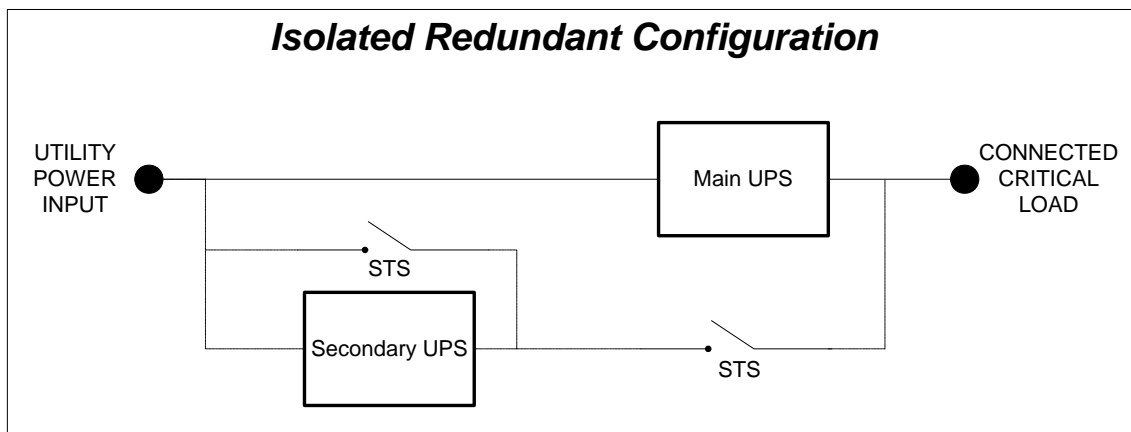


Figure 23

For the additional capital cost of the secondary UPS, the data center manager can make some small gains in overall availability. The availability of an isolated redundant configuration is estimated at 99.93%, which may not seem like much of an improvement over the “N” configuration, but which reduces the amount of unprotected load time to about 6 hours.²⁵

As with the “N” configuration, the percent loading of the system remains relatively unchanged with this configuration because the load is borne solely by the main UPS. Although the load and operational efficiency of the main UPS remain unchanged, the addition of the secondary UPS means that there will be some fixed energy losses required to keep batteries charged and keep the unit operating.

3.2.3 “N+1” Parallel, Single-Bus Configuration

In the “N+1” configuration, higher redundancy is achieved by sharing the data center load between two UPSs arranged in parallel. Both UPSs have an output rating of “N,” meaning that either one is capable of supplying the entire data center load. A system of STS's are not required for this system design because both UPSs are tied to the same inputs and outputs; however, the UPSs usually need to be the same make and model in order to maintain a synchronous output to the load. Figure 24 illustrates the “N+1” configuration in a block diagram.

²⁴ K. McCarthy. “Comparing UPS System...”

²⁵ K. McCarthy. “Comparing UPS System...”

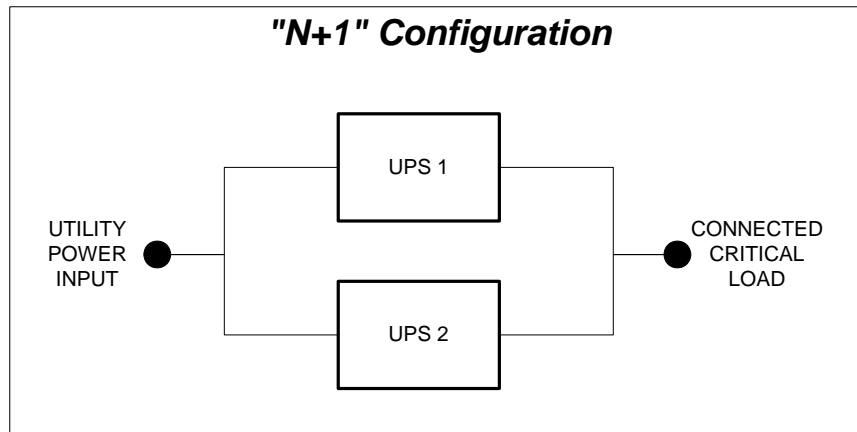


Figure 24

Industry experts estimate that “N+1” UPS configurations have an availability of 99.93%, similar to the isolated redundant configuration described above, without the need for some of the more complicated switch gear.²⁶

One of the major disadvantages is that, by sharing the load between two UPSs, the “N+1” configuration has lower operational efficiencies than the isolated redundant and “N” configurations. When a load is shared equally between two UPSs, each of which is rated to handle the full load by itself, then the largest percent load that could theoretically be achieved by either UPS is 50%. Because UPSs are typically oversized for the loads they support, the percent load and operational efficiency of units arranged in the “N+1” configuration drop further. Based on actual data center measurements (see Section 1.3.3), we estimate that typical operational load factors range from 30% to 50% for individual UPSs arranged in an “N+1” configuration, bringing operational efficiency down into the 80% to 90% range.

3.2.4 “2N” and “2(N+1)” Dual-Bus Configurations

Many servers today employ redundant power supplies, which essentially contain two separate, equally rated power supplies that share the server’s DC load. Should one “half” of the power supply fail, the other will assume the full load. This is analogous to the “N+1” UPS configuration described above in which two UPSs share a common load. Redundant power supplies do not, however, share their AC input through one cord. Rather, they contain two cords – one for each “half” of the unit – so that the unit can be powered by two separate AC buses.

This is where the “2N” UPS configuration comes in. This configuration allows servers with redundant power supplies to be powered by two completely separate AC sources by protecting the power on two independent utility lines. Rather than sharing the load placed on one utility line, the “2N” configuration places one UPS on each of two incoming utility lines, thus preventing single points of failure. All devices downstream of the UPSs have redundant power supplies, receiving half of their power from the first utility line and half from the second. Should both the utility power *and* the UPS on one of the lines fail, the redundant power supply will still have the other protected utility line as a backup. Similarly, if one of the UPSs requires servicing, critical loads can still be protected by the second UPS. Figure 25 shows a block diagram configuration of the “2N” configuration, which is the most common.

²⁶ K. McCarthy. “Comparing UPS System...”

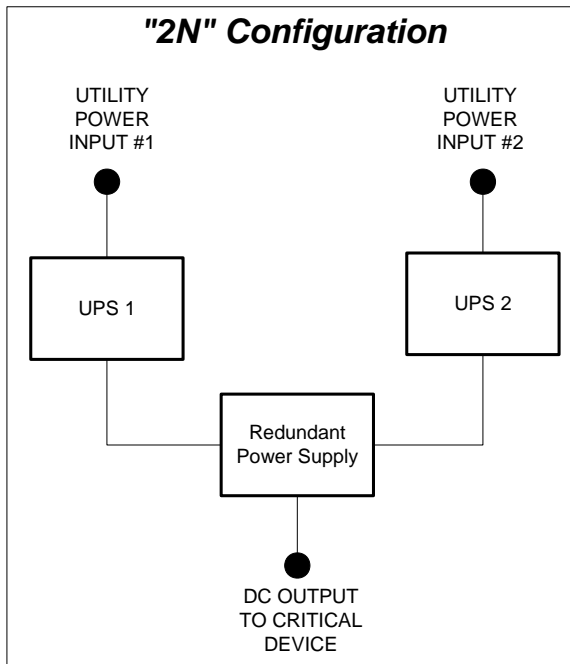


Figure 25

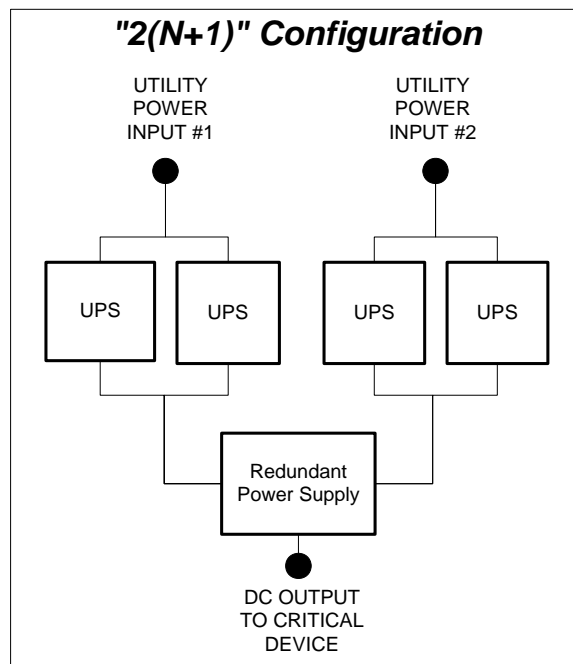


Figure 26

An easy way in which to improve the availability of this system is through the “2(N+1)” configuration. This approach takes the same, dual-bus approach used in “2N” configurations, but adds a full, “N+1” configuration on each utility line for a total of four UPSs. Figure 26 illustrates the approach in block diagram form.

Because the “2N” configuration uses redundant utility sources of power, it further improves on the availability offered by “N,” “N+1,” and isolated redundant configurations, making over four nines of availability (> 99.99% available) possible. With the added redundancy provided through a “2(N+1)” configuration, theoretical redundancy can exceed 99.9999%.²⁷

As with the “N+1” configuration, operational efficiency suffers with improved redundancy, because of load sharing and a reduction in percent load for each individual UPS. From an efficiency standpoint, the “2N” configuration is identical to an “N+1” setup, because the data center load is shared between two UPSs. This reduces theoretical maximum loading for each UPS in the “2N” setup to 50%. Field measurements of UPSs operating in “2N” configurations indicate that operational load factors range from 30% to 50% of the UPS’s rating. The maximum theoretical percent load is even further lowered in the “2(N+1)” configuration, where load is shared between four UPSs. The theoretical maximum load factor for each individual UPS is 25%, and operational load factors are likely in the range of 15% to 25%.²⁸ At these low percent loads, UPS efficiency drops off precipitously from the 90% range to below 80% (see figures 19 and 20, for example).

²⁷ K. McCarthy. “Comparing UPS System...”

²⁸ We have not measured “2(N+1)” UPS configurations under operating conditions in the past, so our estimate of load factor is extrapolated from measurements on “2N” and “N+1” configurations.

4. Proposed Efficiency Specifications

4.1 UPS Operation

4.1.1 Total Cost of Ownership

A calculator was developed for evaluating rack-level energy savings in data centers derived from higher efficiency power conversion. The calculator takes efficiency-related energy losses into account in several key power conversion devices: the data center UPS, individual server power supplies, and motherboard-level voltage regulator modules for processors. The calculator compares the energy use of a high efficiency case to a base case and reports energy savings through improved efficiency in each of the devices mentioned above. The user can modify various efficiency and power use assumptions to customize the calculations to their needs and can view annual dollar savings as well as net present value of savings over a user-defined period of time.

This calculator was further revised and developed into a more complete total cost of ownership (TCO) tool that focuses on energy savings for an entire data center facility rather than just one rack. We used the results of recent UPS and power supply efficiency testing to inform our assumptions for the base case and high efficiency case. We also built in options which allow the user to specify different UPS redundancy configurations, which as we have shown above, can have a dramatic effect on efficiency and, likewise, operational costs. The user now can select to evaluate the effects of efficiency in any combination of the above devices (UPSs, power supplies, and VRM's). The calculator now allows the user to define and adjust the following parameters through an interface like the one shown in Figure 27. This calculator is currently available on LBNL's website devoted to efficient high tech centers (www.hightech.lbl.gov):

- Number of racks in data center
- Number of servers per rack
- Cost of electricity
- Duration to calculate net present value of energy savings
- Discount rate to calculate net present value of energy savings
- Cooling system efficiency
- Power supply nominal power rating
- UPS nominal power rating
- Redundancy configuration of both the base and high efficiency cases

	VRM Options
Evaluate VRM efficiency?	No
	Power Supply Options
Evaluate PS efficiency?	No
PS Nameplate Rating	400 watts

Figure 27: example interface for TCO calculator

4.2 Proposed Efficiency Label

4.2.1 Labeling UPS Efficiency

In the early stages of this project a proposed energy efficiency and power quality labeling scheme for single-phase, “small” UPSs being developed by the Swiss Federal Office of Energy was uncovered. The Swiss’ proposed label, shown below, documents the measured power losses of the UPS in different modes of operation and provides an estimate of the annual energy consumed by energy losses in the UPS design. The label also reports the UPS’s ability to handle different types of power disturbances and the power quality (power factor and total harmonic distortion) that the UPS presents to the grid. It was designed to match the style of existing EU labels for other electric appliances and was intended to be used by manufacturers to voluntarily report efficiency and power quality information to potential UPS purchasers.²⁹ The Swiss Federal Office of Energy is currently investigating the possible mandatory use of the label and is also working with members of the European Union on developing a voluntary Code of Conduct (CoC) for UPS systems.

UPS-System		SFOE USV1A	
Manufacturer			
Model			
Nominal power kW ¹⁾ / kVA ²⁾		XXX / XXX	
Mode of operation			
Low losses			
Losses < 2 %	A		A
Losses < 4 %	B		
Losses < 6 %	C		
Losses < 8 %	D	D	
Losses < 10 %	E		
Losses < 12 %	F		
Losses >= 12 %	G		
High losses			
Energy losses kWh / year ³⁾		xx.x	xx.x
Energy losses kWh at 2'000 h standby		xx.x	xx.x
Filtering of net disturbances		$U_n = \text{~~~~~}$ ⁴⁾	
Outage		✓	> X ms
Voltage interruption		✓	> X ms
Over- and undervoltages		✓	> X ms
Voltage sags/brownouts		✓	> X ms
Harmonic voltages		✓	
Frequency variations		✓	> X ms
Fast transients		✓	< XXX % U_n
Energy loaded transients		✓	< XXX % U_n
Power factor and harmonic distortion		λ / THD ⁵⁾	
<small>No declaration for UPS-systems with a nominal power higher than 10 kVA</small>			
at nominal power in kW ¹⁾		x.xx / xx.x %	x.xx / xx.x %
at nominal power in kVA ²⁾		x.xx / xx.x %	x.xx / xx.x %
at asymmetric nonlinear load ²⁾		x.xx / xx.x %	x.xx / xx.x %
<small> ¹⁾ at ohmic load ²⁾ at non-linear load according to EN 60091 ³⁾ Energy losses at ohmic continuous load with 75 % of nominal power ⁴⁾ U_n = Nominal output voltage ⁵⁾ Filtering is sufficient, if the output voltage fulfills EN 50160 SFOE-Directions for UPS-Systems October 2002 </small>			

Figure 28

In the interests of harmonizing with international efforts to report and promote UPS energy efficiency in a standardized format, our team created a modified version of the proposed Swiss label and a draft report summarizing the criteria that would be used to evaluate UPS efficiency in the U.S. market

²⁹ Schnyder Engineers Ltd. *Label for UPS Systems*. The Swiss Federal Office of Energy. October 2002. pp. 1 - 3

under the label. Our proposed label, shown below in Figure 29, differs from the Swiss version in several distinct ways:

- We have modified the label so that it reports differing levels of energy conversion efficiency rather than energy losses. It was our opinion that reporting efficiency was more appealing to a potential UPS purchaser.
- We have eliminated a field in the label that reports energy losses incurred by operating the UPS for 2000 hours with no load. This information did not seem of much use to UPS purchasers, especially data center users who will likely keep their UPS at least partially loaded with servers year round.
- To help direct the focus of the label solely on energy efficiency and to simplify its design, we decided to eliminate the various tables that report on filtering of power disturbances and power quality. Manufacturers readily report this type of information in product literature should potential purchasers wish to compare different UPSs based on these criteria.

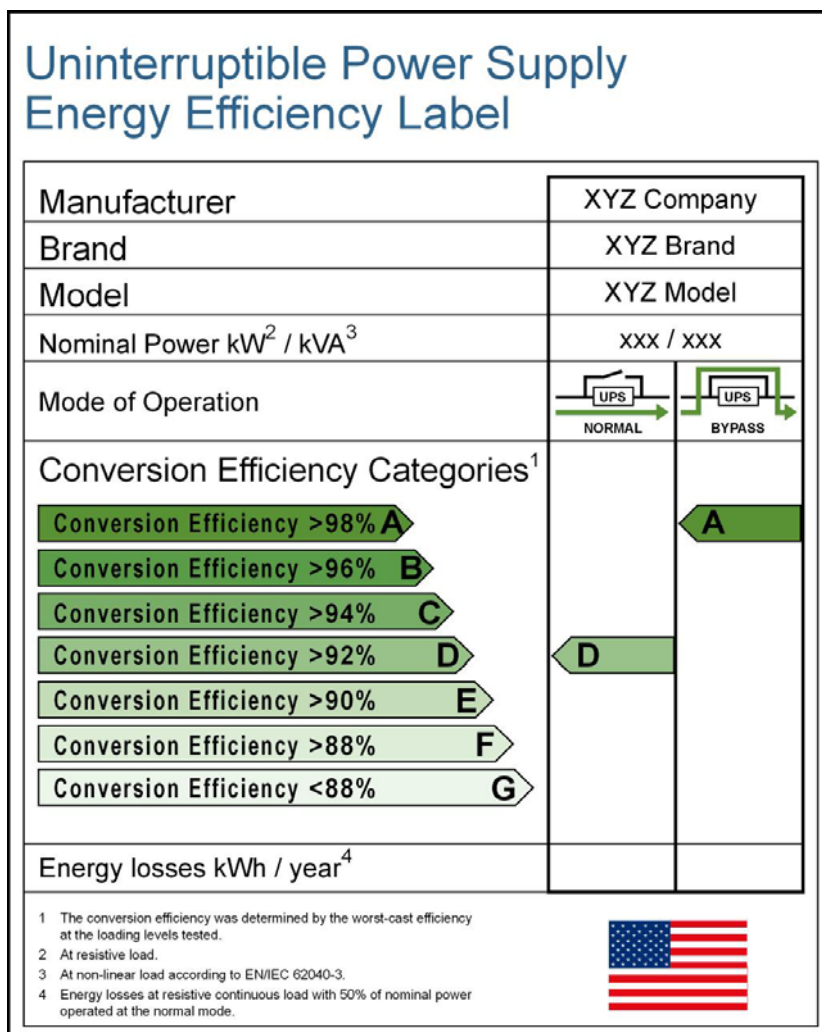


Figure 29

Manufacturer Comments on Proposed UPS Label

The draft label and report were distributed to numerous leading UPS manufacturers to seek comment on the label's design and to ensure that the documentation for the label would provide ample guidance for manufacturers to measure and report UPS efficiency in the proposed format. Several manufacturers responded with comments within a three-week period. We have distilled several common threads of discussion from the various comments received and summarize them below. The main concerns from manufacturers regarding the proposed label were:

Display of performance-related information: Most UPS manufacturers that submitted comments shared a concern that the proposed energy efficiency label does not provide enough information about the UPS's general design and performance. Manufacturers not only want potential purchasers to judge their models based on efficiency, but also performance issues relating to the quality of the UPS output, the ability of the unit to filter out power disturbances, the UPS input power factor, etc. In short, manufacturers want consumers to be able to judge for themselves the potential tradeoffs between efficiency and performance. They feel that consumers should be able to use the label to distinguish the highest efficiency UPS models of a given level of performance.

Bypass mode of operation: Several manufacturers recommended that the column in the UPS label used to show UPS efficiency in the bypass mode of operation be removed altogether. They argue (rightly) that the bypass mode is used to periodically take the UPS out of operation for servicing, which happens only several hours out of the year. During this period of time, the UPS cannot perform its design function: to provide backup power and power conditioning to critical loads. As a result, they argue that bypass mode is not indicative of typical UPS operation and should not be included in the label.

Loading conditions: Manufacturers had two major concerns about testing and loading conditions, namely that lower load levels below 50% of nominal load should be required in testing and that the displacement power factor of linear loads is specified in the labeling guidelines. Manufacturers felt that lower load levels should be required in testing mainly because, as our research has shown, the vast majority of UPSs operate at below 50% of their nominal load. Since the efficiency of a UPS at these lower percent loads can be significantly lower than at 50%, 75% and 100% load, manufacturers thought that an additional loading point in the 20% to 30% range should be required. Due to the effects that load power factor can have on UPS efficiency measurements (one manufacturer cites 2% lower efficiency with increased load power factor), one UPS manufacturer thought that the power factor of the UPS test load should be explicitly stated on the label.

The categorical letter scale: Several manufacturers were quick to point out the past failings of the letter scale used in the proposed UPS label. Recall that the original Swiss UPS label upon which our proposed label was modeled was developed to match European-style efficiency labels used in the EU to report refrigerator and air conditioner efficiency to consumers. These labels all use a categorical letter scale with "A" corresponding to the top efficiency level. Manufacturers noted that, as refrigerators and other appliances have become more efficient in Europe, the original letter scale was no longer adequate to separate out the most efficient models. Over time, additional categories such as "A+" and "A++" have been added to distinguish the top performers. Manufacturers are worried that, as UPS efficiency improves, the categorical scale used in our proposed label could become similarly obsolete. Some also argued that the categorical scale was not the correct approach for the U.S. market because it would be unfamiliar to U.S. consumers and provided too little information about UPS efficiency over a wide range of loads. One manufacturer suggested listing efficiency in a table at 50%, 75% and 100% load as an alternative, although this is not as visual of an approach.

4.2.2 Proposed UPS Efficiency Standard for the U.S. Market

In addition to examining a potential mandatory label for UPS systems sold in the U.S. market, Ecos Consulting and EPRI Solutions also developed and investigated the energy savings potential for a voluntary/mandatory energy efficiency specification. Voluntary labeling programs, such as the U.S. Environmental Protection Agency's ENERGY STAR® program, and mandatory efficiency programs,

such as the California Energy Commission's title standards, have been highly successful in encouraging the efficiency of other electric devices such as lighting, ceiling fans, and refrigerators; however, these programs have never focused attention on uninterruptible power supplies. Programs may have avoided larger UPSs used in mission-critical facilities in the past because these devices are not typically purchased by most consumers.

Below we examine the hurdles to creating a UPS efficiency specification, propose a sample specification based on our measured UPS efficiency data, and estimate the energy savings achievable through the specification.

Challenges to a UPS Efficiency Specification

Creating market transformation mechanisms for UPSs is not as straightforward a task as for other power conversion devices. In many simpler commodity-type power conversion devices, such as external power supplies, there might be a large variation in efficiency between the most and least efficient models, but very little variation in the basic performance and function of the devices. For example, server power supplies are now sold in common form factors designated by the Server System Infrastructure group (SSI) that helps to ensure similar performance characteristics, output voltages, adapter pin-outs, etc. Typical power ratings on server power supplies are also fairly similar and fall within a range from 200 to 600 watts DC.

However in UPSs, the size and performance characteristics of the devices can vary drastically depending on model, manufacturer, etc. UPSs come in a wide variety of sizes ranging from models designed to power individual workstations (< 1kVA units) to those designed to power entire facilities (> 100 kVA), meaning that the largest models can handle two orders of magnitude more power than the smallest models. The way that UPSs deal with power abnormalities and the extent to which UPSs can isolate a load from interruptions in utility power also vary drastically. Standby UPSs and many line-interactive UPSs can pass along voltage sags, line noise, or changes in AC frequency to connected loads, whereas most double-conversion UPSs can filter out all of these disturbances.

Because of the wide variation in performance among UPSs and the equally wide range of available sizes, it becomes increasingly difficult to create a one-size-fits-all UPS specification. UPSs can pay a penalty in overall efficiency by providing better load isolation, filtering, and other types of improved performance, and therefore it would be unfair to compare UPSs with a high-quality output to those with a lower-quality output. Furthermore, smaller power conversion devices typically cannot achieve the same levels of efficiency as larger devices, and so it would be equally unfair to compare the efficiency of a 100 kVA, facility-level UPS to that of a 1 kVA, workstation-level UPS. As a result, any proposed UPS efficiency specification should take the size and performance of the UPS into account.

Sample Efficiency Specification

Our data set of factory and laboratory UPS efficiency measurements were used to create a UPS efficiency specification that could be used by mandatory and voluntary programs. We have mimicked the approach taken by SSI with server power supplies by recommending efficiency levels at 20%, 50% and 100% of the UPS's nominal power rating. Based on the results of field testing of UPS efficiency, we believe that the 20%- and 50%-load efficiency levels are the most important for obtaining real world energy savings, because these points are closest to the range of percent loads at which most UPSs operate.

As described above, it would be difficult to create a "one-size-fits-all" UPS efficiency specification due to large variation in the size and features, both of which have a measurable effect on efficiency. Although we do not currently have enough information to devise an unbiased mechanism that accounts for differences in UPS performance, we were able to account for the UPS size issue by creating three slightly different specifications for UPSs in different size ranges: those with nominal power ratings less than or equal to 20 kVA, those greater than 20 kVA and less than or equal to 100 kVA, and those greater than 100 kVA. These three ranges of power ratings correspond to the regions

in Figure 30 outlined with dashed lines. The average efficiency measurements³⁰ of UPSs in a given region are fairly similar to each other (usually to within a few percent) and do not exhibit any large variations in efficiency as a function of the UPS's rated power output, at least according to our current data set. Our "bracketing" of the UPS specification into three different size groups is merely indicative of the approach that we recommend be taken in formal specification development. In the development of a more formal UPS efficiency specification, a more careful statistical study might need to be performed, coupled with industry stakeholder consultation to determine exactly where the boundaries between different size brackets should be.

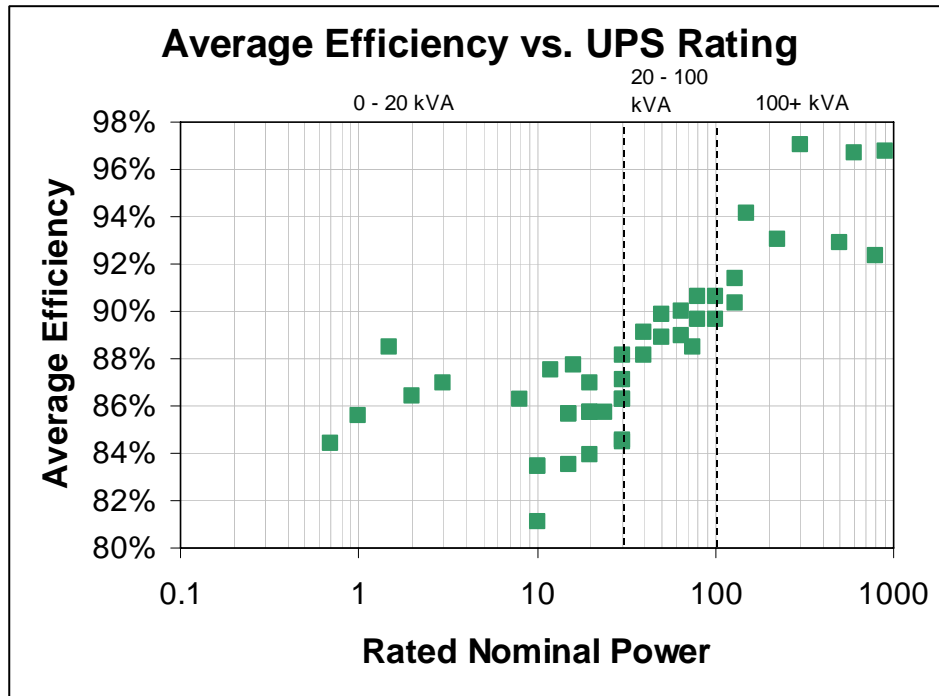


Figure 30

By examining the collected efficiency measurements in our data set within each of the UPS size groups described above, we developed recommended efficiency levels for our proposed specification, shown in Table 8.³¹ In order to pass the specification, a UPS must exceed the proposed efficiency levels at 20%, 50%, and 100% load. We designed the specification such that roughly the top 20% to 30% most efficient UPSs in our data set would pass the criterion, which ensures that our proposed efficiency specification is technologically achievable even today. Charts provided in Appendix A plot the various efficiency levels alongside our UPS efficiency measurements.

Table 8: Proposed UPS Efficiency Specification Levels

UPS Nominal Power Output	Efficiency at 20% Load	Efficiency at 50% Load	Efficiency at 100% Load
Below 20 kVA	80%	87%	88%
20 kVA to 100 kVA	83%	89%	90%

³⁰ Note: average efficiency is calculated by taking the mean efficiency at 25%, 50%, 75%, and 100% load. In several cases where manufacturers did not measure efficiency at these exact load points, it was necessary to use another load point within +/- 10% of 25%, 50%, 75% or 100%.

³¹ We used UPS efficiency measured with linear loads to determine specification levels, because the majority of bulk of our data set is comprised of these types of measurements. In a formal specification, non-linear loads should be used to measure efficiency (described in IEC-62040-3) because these loads more closely resemble the types of loads seen by most UPSs in operating environments.

Above 100 kVA	90%	95%	95%
---------------	-----	-----	-----

Estimated Energy Savings of Proposed Specification

We estimate that our recommended improvements to the efficiency of uninterruptible power supplies could save millions of dollars in utility costs every year by trimming power conversion energy losses. We have performed a first-order energy use and savings analysis to determine the effectiveness and potential energy savings of the proposed UPS efficiency standard and have built this analysis upon several key pieces of research that have emerged from work on UPS efficiency:

- Estimates of UPS stock presented in Section 1.1.5 for the US data center/IT sector
- Field measurements of average UPS load factors and operational efficiency, presented in Section 1.3.3.
- Factory and laboratory measurements of UPS efficiency across a wide range of loads, presented in Section 1.3.4.

Our general approach was to compare the energy losses or energy consumption of “typical” or “average” UPSs to those operating at our proposed efficiency levels within each of the three size/power categories described above. We first required several simplifying assumptions about UPS loading and operational efficiency. Based on our field measurements of data center UPSs, we assumed that most large UPSs servicing IT equipment and data centers operate at about 38% load and with an average efficiency of about 85%. Since UPSs smaller than 20 kVA were not tested in the field, we had to refer to our factory and laboratory measurements submitted by manufacturers to determine an average efficiency for this category of UPS – 83.6% at our typical percent load of about 38%.³² Due to the wide range of UPS sizes, we also needed to assume a typical size for UPSs in each of the three categories. We used the midpoint of the range to approximate average UPS size in the < 20 kVA and 20 – 100 kVA ranges, and we assumed a typical value of 200 kVA for UPSs larger than 100 kVA.³³ Finally, we assume that the amount of active output power (in kW) that a typical UPS can provide to loads is roughly 80% of the UPS rating in kVA (this is the apparent power rating and is always larger than the apparent power rating). This assumption is based on typical UPS nameplate ratings from manufacturer data sheets. Table 9 presents some of the key assumptions about “typical” or “average” UPSs.

Table 9: Assumptions for Typical UPS

Output Power Range (kVA)	Assumed Output Power (kVA)	Assumed Output Power (kW)	Load Factor During Operation	Efficiency at Operational Load Factor
< 20 kVA	10	8	38%	83.6%
20 – 100 kVA	70	56	38%	85.2%
> 100 kVA	200	160	38%	85.2%

In the case of high efficiency UPSs, we kept assumptions regarding output power and load factor the same but raised the UPS efficiency based on the levels in the proposed efficiency specification in the section above. Since the efficiency specification does not specifically require an efficiency level at average operational loads (38%), we linearly interpolated between the 20% and 50% efficiency levels to estimate the kind of operational efficiency that one could expect from a UPS that complies with the standard. Table 10 shows our complete assumptions for high efficiency UPSs.

³² This average efficiency had to be linearly interpolated based on efficiency measurements taken at other load points.

³³ This assumption is based on past sales data from Frost & Sullivan, Roth, and others that indicate that most UPS sales above 100 kVA cluster heavily in the 100 to 200 kVA range. 200 kVA seems to be an acceptable median value.

Table 10: Assumptions for High Efficiency UPS

Output Power Range (kVA)	Assumed Output Power (kVA)	Assumed Output Power (kW)	Load Factor During Operation	Efficiency at Operational Load Factor
< 20 kVA	10	8	38%	84.2%
20 – 100 kVA	70	56	38%	86.6%
> 100 kVA	200	160	38%	93.0%

With the given assumptions about the output power of the UPS, the typical percent load placed on the UPS, and the efficiency of the UPS at operational loads, we were able to estimate average power losses in different sizes of UPS for average and high efficiency UPSs. Table 11 shows the estimated power losses and estimated annual energy consumption (AEC) based on year-round operation.

Table 11: UPS Annual Energy Consumption Estimates

Average UPS		
Output Power Range (kVA)	Power Losses (kW)	Annual Energy Consumption (kWh)
< 20 kVA	0.59	5,197
20 – 100 kVA	3.68	32,211
> 100 kVA	10.51	92,032
High Efficiency UPS		
Output Power Range (kVA)	Power Losses (kW)	Annual Energy Consumption (kWh)
< 20 kVA	0.57	4,971
20 – 100 kVA	3.28	28,693
> 100 kVA	4.55	39,878

Using this information and our estimates of UPS stock, we constructed a first-order estimate of UPS energy use. This energy use and savings estimate is limited to the data center/IT sector due to the scope of our research, in which we measured the efficiency of data center UPSs both in the lab and in the field. Based on our estimates, the data center/IT sector currently consumes roughly 7.1 TWh (7.1 billion kWh) of electricity per year as a result of power conversion losses in UPSs – an amount equal to about 0.3% of combined commercial and industrial electricity consumption in the US.³⁴ Assuming an average electric utility charge of \$0.10/kWh, this amounts to over \$700 million spent every year by the data center/IT sector in the US to pay for electricity that is simply converted to heat inside the UPS.³⁵ We estimate that the State of California alone, which contains roughly 15% of the nation’s data center floor space,³⁶ consumes about 1 billion kWh of electricity and spends about \$100 million per year in electric bills due to power conversion losses in data center UPSs.³⁷ Our national

³⁴ Percentages based on “No. 916. Electric Energy Sales by Class of Service and State: 2001.” *Statistical Abstract of the United States: 2003*. U.S. Census Bureau. 2003. p. 593

³⁵ Note: this figure does not include electricity expenses for the devices connected to the UPS

³⁶ R. Juarez, M. Alic, K. Chetan and B. Johnson. *SpaceDex III – Hosting Space: Not All Hosting Space is Created Equal – Smart, Complex Space Takes Data Center Stage*. Boston: Robertson Stephens Inc. 2001.

³⁷ Because California contains 15% of the nation’s data center floor space, its UPSs consume 15% of the energy consumed by all UPSs in the US.

energy use figures are noticeably higher than past estimates made by Arthur D. Little for the U.S. Department of Energy, which estimated UPS energy consumption in the U.S. “IT/telecom” sector at 5.8 TWh per year.³⁸ Both Ecos’ and Arthur D. Little’s energy use estimates are shown in Table 12. Also included are Ecos Consulting’s estimate of UPS stock and energy use in California, even though the 2002 DOE/ADL report did not make such an estimate. There are several reasons for the estimated increase in UPS energy consumption:

- The UPS stock has grown significantly since the publication of the Arthur D. Little report, due to heavy sales at the tail end of the tech boom and continued, although weaker, sales of UPSs in the three years since. Overall, our stock estimate is 65% higher in terms of numbers of units.
- Research conducted by Ecos Consulting, EPRI Solutions, and LBNL indicate that UPSs typically operate at lower load factors and with lower efficiency than was assumed in the Arthur D. Little estimate, which assumed roughly 50% loading. We now know that UPSs more often operate in the 30% to 50% range of loads, where power conversion efficiency begins to decline.

Table 12: Comparison of UPS Energy Use Estimates

	DOE/ADL 2002 Estimate of UPS Annual Energy Consumption in U.S. IT/telecom Sector	Ecos/EPRI 2004 Estimate of UPS Annual Energy Consumption in U.S. IT/data center Sector
<i>Estimated National UPS Stock</i>	151,761	250,343
<i>Nationwide Annual Energy Consumption (TWh)</i>	5.8	7.1
<i>Estimated California UPS Stock</i>	n/a	37,551
<i>California Annual Energy Consumption (TWh)</i>	n/a	1.1

Based on a comparison of our “high efficiency” and “typical” cases for UPS energy consumption, we have estimated the energy savings potential for enacting our proposed UPS efficiency standard in the U.S. Again, due to the scope of our work, this estimate applies to the data center/IT sector and does not indicate savings that might be achievable in residential or industrial UPS installations. Although it might not be likely that U.S. federal authorities would implement a national, mandatory standard governing the efficiency of UPSs, we estimate that our proposed efficiency measures would have saved 2.8 TWh of electricity in 2004 if all UPSs in use in the U.S. were forced to comply. This would result in about \$280 million in savings from lower electric utility bills. Applying the proposed efficiency specification on a mandatory basis could cut nationwide energy use by UPSs by about 40%; however, the proposed savings would amount to only about 0.1% of combined industrial and commercial annual electricity consumption.

A potentially more feasible way to encourage energy savings in UPSs would be the introduction of a voluntary UPS efficiency specification that would be administered under a program such as ENERGY STAR. In this energy savings scenario, it is likely that only a fraction of the UPS market will meet the proposed efficiency levels. If, for example, 25% of the UPSs installed in the U.S. were able to comply with the proposed UPS efficiency levels, the nation would save about 700 million kWh of electricity per year, amounting to about \$70 million utility cost savings. We have summarized our energy savings scenarios for mandatory and voluntary specifications in Table 13, along with similar energy

³⁸ K. Roth, F. Goldstein and J. Kleinman. *Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings, Volume I: Energy Consumption Baseline*. Cambridge, MA: Arthur D. Little, 2002. p. 84

savings estimates that apply only to the State of California, which we assume contains about 15% of the nation's UPSs.³⁹

Table 13: Energy Savings Scenarios for Proposed UPS Efficiency Specification

	Mandatory Specification – 100% of Installed Units Comply	Voluntary Specification – 25% of Installed Units Comply
Nationwide Annual Energy Savings (TWh)	2.8	0.71
Nationwide Annual Utility Cost Savings (million USD)	\$280	\$71
California Annual Energy Savings (TWh)	0.42	0.11
California Annual Utility Cost Savings (million USD)	\$42	\$11

The energy savings figures shown here could be significantly larger when examining the *entire* UPS market and not just the data center/IT sector. For UPSs above 5 kVA in power output, the data center/IT sector only accounts for roughly two thirds of the installed units. An additional one third of UPSs installed in small offices, industrial facilities, hospitals and the like have not been included in our energy use and savings estimates. Our analysis also excludes the below 5 kVA UPS market segment, more typical of small office applications, residential installations, etc. The combined sales of UPSs below 5 kVA amounted to over 8 million units in 2004, close to 200 times the unit sales volume of UPSs greater than 5 kVA in power output.⁴⁰ Although these are lower power devices than the larger UPSs examined in this report and would thus have smaller per-unit power losses, they might represent an attractive energy savings opportunity simply due to the large number of units in operation that should be studied in the future.

It should be noted that additional energy savings can be realized from reduced facility cooling needs because less waste heat is generated from the conversion and storage equipment. LBNL benchmarking suggests that the secondary savings from reduced cooling needs as a result of more efficient facilities is on the same order of magnitude as the direct, primary savings from efficient conversions. Finally, there is also additional savings from reduced capital investment in equipment if a more efficient system is selected, as well as reduced real estate investment from smaller equipment footprints. These additional savings can be realized up front rather than through reduced operational costs.

³⁹ Based on estimates that California contains 15% of the nation's data center floor space. R. Juarez et al. 2001.

⁴⁰ Frost and Sullivan. 2004.

5. DC Architecture Scoping Study

5.1 Market Trends in IT Load Power Configuration

New generation data processing equipment used in data centers and general business services provides ever-increasing processing speed and bandwidth utilization benefits, at the same time, increasing growth in Internet-based communication and commerce activities continually stimulate the need for more performance. As a result of this self-reinforcing spiral, the amount and concentration of electrical energy being dedicated toward data processing and storage infrastructure is also on an exponential growth path. This growth creates new challenges regarding energy and resources to power these systems.

Data processing equipment is generally built around microprocessor cores, or individual server units in dense racks and clusters. Data centers use AC in a conventional distribution to deliver power to the data processing equipment (servers, storage disk arrays, etc.). This typical delivery scheme involves power conversion devices (power supplies) to convert AC to DC at the equipment component and processor levels. In the past these microprocessors-based server unit would typically utilize less than 100W of energy per unit. However, modern equipment quite often exceeds 200W per unit, depending on configurations, with future equipment forecast to go well in excess of this value. As these channels are paralleled together into clusters, the total energy being utilized presents unique challenges.

Inside the server boxes, power supplies can provide power factor correction as well as load isolation from the incoming power line. This conversion generally involves at least two stages. In addition, most modern microprocessors require very low voltages at fairly high currents, such as 1.1V at 100A. The precision of the voltage required is such that voltage regulation circuitry must be located directly next to the microprocessor. In order to effectively realize this circuitry, most processors require that an intermediate DC voltage, such as 12VDC, be delivered to the processor/local regulator combination. Thus, from the power supply, there can be up to six or more power conversion stages between facility power entry and the microprocessor. (480VAC to 208VAC, 208VAC to 400VDC, 400VDC to 208VAC, 208VAC to 400VDC, 400VDC to 12VDC, 12VDC to 1.1VDC)

Depending on processor and server loading, which can dynamically range from 0% to 100%, the efficiencies of power supplies can be lower at lower load levels and can significantly impact overall system efficiency. In many cases, redundant power supplies are used to deliver this power, either on standby or in load-sharing configurations inside servers, both of which reduce individual power supply loading.

The need for reliability also necessitates the use of UPS units to condition the AC power provided to the servers. The central AC UPS generally involves conversion from incoming raw AC power to DC, and then reconverting from DC back to AC. Moreover, since the UPS generally requires a maintenance bypass switch, facility-level entry power usually must be transformed from 480VAC down to 208VAC prior to the UPS. The central battery is then connected as an alternate input to the DC-AC converter so that if incoming raw AC power is interrupted, the system automatically switches over to battery power. Uninterruptible AC power is then passed through an AC power distribution grid and fanned out to individual data processing apparatus. The servers receive the AC output of the UPS (converted from DC) and then convert again to DC with its power supply.

Typically, uninterruptible power is viewed at the facility level, as opposed to the equipment level. This perspective provides an easy division between the facility power equipment and data processing equipment, with each focusing on a different part of the power delivery. However, this view also makes it difficult to ascertain and optimize overall operating efficiency and total cost of ownership, since losses from each of these conversion processes directly translate into heat, adding to both the server cooling load and the overall data center's cooling load.

5.2 Efficiency of AC- and DC-based Systems

Inside the server units, the limiting factors on power conversion efficiency are voltage rating of semiconductors, along with their corresponding conduction losses. In addition, power converter package size limitations, and the economics of power supply design and manufacturing can also play a part. All of these factors combine to ultimately limit available efficiency improvements for the power conversion processes. Power conversion efficiency for best-in-class computer grade AC-DC power supplies currently do not provide efficiency higher than about 80% (See related report on “best-in-class” testing results for server power supplies). For servers, on-board non-isolated converters that provide final processor power conversion can range as high as 98% efficiency.

The other area of complexity within large data processing installations is uninterrupted AC power distribution. Facility-level AC UPS systems generally have a centralized circuit breaker panel with power monitoring and bypass provisions for UPS and battery servicing. Often, these units are located at some distance away from the equipment being powered; leading to possible confusion regarding which equipment is fed by what breaker and introducing the possibility of inadvertent information technology equipment shut down due to operator confusion. In some cases the nature of the UPS being utilized is very sensitive to load current harmonics, resulting in a great level of care being required to assure data processing equipment and UPS compatibility.

The centralized battery plant utilized in facility-level UPS systems can be a compromise between what is required for successful system realization and component limitations. The actual battery run time required for acceptable system operation is often just a few minutes. The time represents the delay needed to switch to alternate utility power feeds, or to bring an auxiliary source of power on line (such as a motor generator). However, when batteries for centralized AC UPS systems are sized, the nature of the voltages required or battery type chosen can often result in hold up times well in excess of what is required, resulting in wasted energy storage and delivery capabilities. Additionally, the lower efficiency presented by the off-line power supplies utilized by data center equipment produces an extra load on batteries that only goes into producing heat instead of power conversion.

Using the typical efficiencies of the different conversion steps listed below – from facility entry AC to the processor:

- Step Down Transformer (99.5%)
- AC UPS (85%)
- Computer Power Supply (70%)
- Point of Load Converter (98%)

An estimate of the overall net power efficiency of about 58% can be estimated as follows:

Step Down Transformer (99.5%) x AC UPS (85%) x Computer Power Supply (70%) x Point of Load Converter (98%) = 58%

The above result indicates that for every watt of power utilized to process data, another 0.4W to 0.5W is required to support power conversion. In addition, for air-conditioned facility cooling, a conservative estimate suggests that another watt of power will be required for each watt utilized to cool the power conversion equipment. While additional cooling power can seem insignificant at the individual microprocessor level, when overall data processing activities reach power usage levels on the order of 200kW, or more (for example, a large Internet 4 hub installation), almost 100 kW of power would be wasted in the overall power conversion process, not including additional cooling loads.

5.3 Demonstration Plans for DC Architecture

5.3.1 Demonstration Objectives

The demonstration has the following goals:

1. Show that DC-powered server(s) and/or server rack exists in the same form factor or can be built and operated from existing components with minimal effort.
2. Show that DC-powered server(s) and/or server rack can provide the same level of functionality and computing performance when compared to similarly configured and operating server (and/or server rack) containing AC power supplies, as measured with industry standard measurement devices and software benchmarking tools.
3. Measure and document any efficiency gains from the elimination of multiple conversion steps in the delivery of DC power to the server hardware.
4. Identify areas for follow up investigations.

5.3.2 Approach

The proposed demonstration will compare two power distribution schemes to servers and/or server racks that perform the same data processing functions. The two systems are shown in Figures 31 and 32 below.

- System #1 is comprised of a double conversion UPS (local to the rack) powering a server rack with AC.
- System #2 utilizes a single rectifier (local to the rack) to provide DC to the server rack containing DC powered server(s).

Although many data centers use a centralized UPS, and a data center DC distribution design might be based on the plant telecom system which will provide similar functionality, for the purposes of this demonstration, a local system will be used in order to better control all the variables to be monitored.




Figure 31: System #1, Utilizing A UPS And AC Powered Servers.

Power analyzers on the input of each system will provide power consumption data as the servers go through a standard benchmarking that uses a controllable and repeatable data processing routine.




Figure 32: System #2, Utilizing Single Rectifier and DC Powered Servers.

The power consumption by System #1 will be compared to the power used by System #2, to determine if there is any difference in the power consumption by the two systems while performing identical tasks. The results of this demonstration will allow comparison of power consumption by the two approaches, and to determine whether increases in efficiency are possible.

6. Conclusions and Recommendations

This project served well as an initial investigation into this particular area of high-tech buildings and data centers, and has provided an overview of UPS utilization and applications. Due to the broad scope of our work, we were able to identify a number of important issues. However, we were not able to carry out in-depth investigations into any one issue. Below are a number of observations and recommendations based on the findings of this study:

Focus on Large UPS Units: Standby UPSs are most common in low-power applications (0 – 2 kVA) such as individual workstation backup power, small office server backup power, etc. The lack of complete power conditioning and load isolation generally make standby UPSs unsuitable to critical loads such as industrial facilities or data centers. We estimate that there are about 225,000 UPSs in operation in the U.S. data center/IT sector today, about 15% of which reside in the State of California. Half of the national UPS stock consists of units in the 5.1 to 20 kVA range of apparent power output.

UPS Loading and Sizing Can Significantly Affect Energy Use: UPSs in mission critical environments typically operate at points between 30% and 50% of nominal and in situations where the data center is not fully populated this can be even lower resulting in very large inefficiencies. There are a number of reasons why data center UPSs operate at part loads, one of which is that UPSs are commonly sized to meet the maximum utilization of space in a data center, even though maximum utilization rarely occurs.⁴¹ This results in over sizing of the UPS for the actual amount of installed mission-critical infrastructure.

Topology and Efficiency Considerations: Due to the scope of this research effort, it was impractical to determine what specific elements of a model's design contributed to its measured efficiency; however, our results indicate that some broader design decisions, such as the general UPS topology, can be indicative of efficiency. Research on the efficiency of both line-interactive and delta-conversion UPSs should continue as manufacturers continue to improve their designs and diversify product lines. Nevertheless, the high efficiency of some UPS topologies in our current data set suggest that they deserve serious consideration by facility managers looking to lower operating costs in mission critical facilities that utilize large UPSs (greater than 50 kVA in capacity) and policy makers wishing to identify the most efficient technologies.

Reliability Considerations: Desire for high efficiency should naturally be balanced with concerns over load isolation and reliability. Although the delta-conversion UPSs that we tested performed better in regards to overall efficiency compared to double-conversion units, some manufacturers have argued that delta-conversion UPSs do not provide the same load isolation as a “true” double-conversion UPS. This report does not attempt to answer the question as to which topology is the overall best choice for mission critical applications. We simply report the observed efficiencies of the various topologies, and facility managers will ultimately need to judge whether a particular UPS can handle load disturbances and provide

⁴¹ Madsen J. “Continuous UPS Availability: How Important is it to Your Company?” *Energy User News*. August 11, 2000. <http://www.energyusernews.com/CDA/ArticleInformation/features/BNP__Features__Item/0,2584,14489,00.html>

an acceptable level of load isolation. Additional research and discussion within the industry can help to provide more clarity on this issue.

Considerations in Creating a Standard for UPS: Because of the wide variation in performance among UPSs and the equally wide range of available sizes, it is difficult to create a one-size-fits-all UPS specification. UPSs can pay a penalty in overall efficiency by providing better load isolation, filtering, and other types of improved performance, and therefore it would be unfair to compare UPSs with a high-quality output to those with a lower-quality output. Furthermore, smaller power conversion devices typically cannot achieve the same levels of efficiency as larger devices, and so it would be equally unfair to compare the efficiency of a 100 kVA, facility-level UPS to that of a 1 kVA, workstation-level UPS. As a result, any proposed UPS efficiency specification should take the size and performance of the UPS into account.

Finally, it should be noted that additional energy savings can be realized from reduced facility cooling needs because less waste heat is generated from the conversion and storage equipment. LBNL benchmarking suggests that the secondary savings from reduced cooling needs as a result of more efficient facilities is on the same order of magnitude as the direct, primary savings from efficient conversions. There are also additional savings from reduced capital investment in equipment if a more efficient system is selected, as well as reduced real estate investment from smaller equipment footprints. These additional savings can be realized up front rather than through reduced operational costs. Thus, additional investigations into the areas discussed above, as well as the feasibility of DC powering architectures for data centers will provide further insight and can help in increasing the overall efficiencies of future data centers.