

# Implementing Energy Efficient Data Centers

By Neil Rasmussen

White Paper #114

**APC**<sup>®</sup>  
Legendary Reliability<sup>®</sup>

## Executive Summary

Electricity usage costs have become an increasing fraction of the total cost of ownership (TCO) for data centers. It is possible to dramatically reduce the electrical consumption of typical data centers through appropriate design of the network-critical physical infrastructure and through the design of the IT architecture. This paper explains how to quantify the electricity savings and provides examples of methods that can greatly reduce electrical power consumption.

# Introduction

Electrical power usage is not a typical design criterion for data centers, nor is it effectively managed as an expense. This is true despite the fact that the electrical power costs over the life of a data center may exceed the costs of the electrical power system including the UPS, and also may exceed the cost of the IT equipment. The reasons for this situation are as follows:

- The billed electrical costs come after the charges are incurred and are not clearly linked to any particular decisions or operating practices. Therefore they are viewed as inevitable.
- Tools for modeling the electrical costs of data centers are not widely available and are not commonly used during data center design.
- The billed electrical costs are often not within the responsibility or budget of the data center operating group.
- The electrical bill for the data center may be included within a larger electrical bill and may not be available separately.
- Decision makers are not provided sufficient information during planning and purchasing decisions regarding the electrical cost consequences.

This paper will show that all of the above can and should be corrected, because substantial financial savings are possible for typical users. The greatest advantage can be gained in the design of new facilities, but some savings are possible for existing and evolving facilities as well. **Simple no-cost decisions made in the design of a new data center can result in savings of 20- 50% of the electrical bill, and with systematic effort up to 90% of the electrical bill can be avoided.**

## What is the cost of electrical power consumption?


A typical value for the cost of electrical power is \$0.12 per kW hr. Given this cost, the annual electrical cost per kW of IT load is approximately \$1,000. Over the 10 year life of a typical data center this translates to approximately \$10,000 per kW of load.

As a general rule, approximately half of the energy used in a data center goes to the IT loads. The other half goes to the network-critical physical infrastructure (NCPI) equipment including power equipment. This means that **for each kW of IT load the 10 year electricity cost is approximately \$20,000**. For example, a 200 kW data center would have a 10 year electricity cost of \$4,000,000. This is a material cost for any organization and all IT professionals should understand where this expense is going and that it is avoidable.

## Where does the energy go?

Approximately half or less of the energy used in a data center goes to the IT loads. The other half goes to the network-critical physical infrastructure (NCPI) equipment including power equipment, cooling equipment, and lighting. **Figure 1** shows the electrical power flow in a typical high availability data center. Note that all the energy consumed by the data center ends up as waste heat which is rejected outdoors into the atmosphere. **Figure 1** is based on a typical data center with 2N power and N+1 cooling equipment, operating at approximately 30% of rated capacity.

*Figure 1 – Power flow in a typical data center*



The above data center is said to be 30% efficient, based on the fraction of the input power that actually goes to the IT load. For a more detailed understanding of where the power goes and how the different types of equipment contribute to the load, consult APC White Paper #113, “Electrical Efficiency Modeling for Data Centers.”

## Efficiency is a poor metric

Many discussions regarding electrical power consumption use the term “efficiency”. While the underlying meaning of terms like “improving efficiency” are well understood, the technical use of the term “efficiency” for quantitative assessment of data centers leads to confusion. Discussions are much clearer when electrical power consumption (kW) is the metric used, rather than efficiency metrics. For example, if two different devices in a data center are 50% and 80% efficient, it is not clear how to combine their efficiencies into a single number that relates to costs. In fact, the electrical costs would actually be dependent on the quantity of power flowing through each device. Furthermore, some devices like computers or lighting are zero

percent efficient which is a confusing concept and conveys no quantitative information regarding their electrical use.

In contrast, using electrical power consumption for a metric is simple and unambiguous. The total electrical consumption is simply the sum of the consumption of all devices in the data center. If one device uses \$10 of electricity per month and another uses \$20, it is a simple matter to sum these values. Therefore, in this paper electrical power consumption will be the quantitative term used instead of the more common but ambiguous term “efficiency”. A complete discussion of the modeling of data center energy consumption is provided in APC White paper #113.

## The value of a Watt

Electrical power is sold in units of energy called kilowatt-hours (kW-hr), which is the amount of energy delivered in one hour at a power level of 1000 Watts (1 kW). The distinction between **power** and **energy** is very important for the economic analysis. **Power capacity** costs are those associated with the systems that deliver energy and increase with the design power level of the system. Examples of costs driven by power capacity are UPS costs, generator costs, air conditioner costs, and power distribution equipment costs. **Energy** costs are those associated with the electrical utility bill.

A key principle to understand is that **reducing energy consumption can reduce the power capacity related costs as well as the energy costs**. That is, an implementation that saves electricity in many cases can also save on the NCPI infrastructure costs, which are primarily driven by the load power demand. A companion principle that is key to understand is that there is a difference between reducing energy consumption temporarily and reducing energy consumption permanently. Temporary savings like load shedding or server power management reduce electricity costs but do not necessarily reduce the power rating of the NCPI systems and the related NCPI infrastructure costs. Permanent or structural changes like high efficiency servers or high efficiency UPS systems reduce both the electricity costs and the infrastructure costs. These principles are illustrated in **Table 1** along with example savings values.

*Table 1 – The economic benefits of saving a kW of electrical consumption in a typical high availability data center, comparing temporary and structural consumption avoidance*

	Temporary Consumption Avoidance	Structural Consumption Avoidance	Comment
Method of savings	Power Management Load shedding Economizer	High Efficiency Servers High Efficiency UPS Right-sizing	
1yr electrical savings	\$960	\$960	Assuming \$0.12 per kW hr
10yr electrical savings (IT)	\$9,600	\$9,600	Typical design life of data center
10yr electrical savings (NCPI)	\$960	\$13,760	Structural avoidance allows reduction in capacity-related electrical consumption
NCPI CapEx Savings	\$0	\$13,300	Structural avoidance allows reduction in equipment capacity
NCPI OpEx Savings	\$0	\$6,600	Reduction in equipment reduces operating expenses such as maintenance
Total 10yr Savings per kW	\$10,560	\$43,260	

In the example above, the data center is 2N redundant and operating at a typical 30% load. Note that for a non redundant data center the savings would be greatly reduced, to about half of the savings shown. Note also that in a typical situation not all of the installed power and cooling capacity requirement can be avoided by a structural reduction, so the savings may be further reduced. However, in general a reasonable estimate is that structural consumption avoidance is worth twice as much as temporary avoidance.

**The key point to understand here is that there are two kinds of energy consumption reductions: Those that avoid energy consumption, but do not reduce power capacity requirements, and those that also allow the reduction of installed power capacity. We will refer to those reductions in consumption that avoid energy use without reducing installed power capacity as “temporary consumption avoidance” and those that allow the reduction of installed power capacity to be “structural consumption avoidance”. Furthermore, for data centers, a general rule is that structural consumption avoidance is worth approximately twice as much as temporary consumption avoidance.**

# Energy consumption reduction in IT equipment

Clearly the primary driver of power consumption is the power draw of the IT equipment. IT equipment power consumption directly contributes to the electrical bill, and it indirectly contributes by requiring various power and cooling equipment that also consume comparable amounts of electricity. Therefore, all IT personnel should be concerned with controlling the power consumption of IT equipment.

The methods for control of IT power consumption have historically been very weak. For example, IT equipment vendors have not provided adequate information to allow users to make decisions based on power usage. Users typically do not understand that they have IT choices that can affect power consumption. However, the situation is improving and users can take both operational and planning actions that will systematically reduce power consumption.

The reduction of power consumption of IT systems consists of several approaches:

- Operational actions: retiring systems, operating existing systems in an efficient manner, and migrating to more energy efficient platforms
- Planning actions: virtualization, and standardization

Each of these will be discussed in turn.

## Operational: Retiring IT systems

Most data centers have old technology platforms that remain operational for archival or research purposes. In fact, most data centers actually have application servers which are operating but have no users. It is useful to inventory these systems and create a retirement plan. In many cases, systems can be taken off line and powered down, even if they are not physically retired.

A related opportunity exists where multiple old technology platforms can have their applications consolidated onto new servers, essentially reducing the total server count. This type of consolidation does not require virtualization, which is discussed later.

A power consumption reduction of up to 20% is possible in typical cases. Even if the floor space is not recovered, **the power capacity recovered can be very valuable** as users deploy higher density IT equipment.

## Operational: Operating existing systems in an efficient manner

Today, most new servers have power management features. That is, they are able to reduce power consumption at times of reduced computational load. This was not true a few years ago, when the power consumption of virtually all IT equipment was constant and independent of computational load. Users should be aware of this change in IT technology, and be aware of the status of the power management features on

their IT systems. Where possible, power management should be enabled on all devices with such capabilities. Note that many equipment manufacturers supply equipment with these features disabled by default. This may require upgrading applications to ensure that they take maximum advantage of the power management features. Power management features reduce the total electrical usage but do not reduce the power capacity requirement.

### Operational: Migration to energy efficient computing platforms

Migration to more electrically-efficient platforms is another effective strategy for reducing power consumption. Most data centers have so-called “low density servers” that are 3-5 years old. Typically these servers draw the same or less power per server than today’s blade servers and are physically much larger per server. Migration to modern blade servers from legacy servers on a server-by-server basis typically does NOT reduce the total power consumption and may even raise it. However, such migration will permit much higher packing densities for servers. Blades do not create more heat than equivalent 1U servers, but they do create heat in a smaller area which gives rise to heat removal problems that create the perception that blades create excess heat.

When a new server deployment is planned, the use of blade servers as opposed to alternative server form factors will generally give a 20% reduction in power consumption. This is because blade servers generally have higher-efficiency power supplies and share some overhead functions such as fans. It is important to understand that selecting the blade form factor reduces power consumption relative to other server form factors for newly deployed equipment, but blades do not necessarily consume less power than older servers.

This discussion suggests that a server-by-server migration from existing server technologies does not necessarily cause a significant reduction in power consumption. To determine the potential to save power by migrating to blades on a server-by-server basis, the power consumption of the existing server should be compared with the power consumption of any proposed blade server. Furthermore, the performance of both servers should be compared so as to arrive at a performance per watt metric. Today, major OEMs such as Dell, HP, and IBM provide user configuration tools that accurately report actual power consumption for various blade server configurations. To determine power consumption values for legacy servers, the only realistic way is to measure example servers using a Wattmeter. By comparing the values obtained in this manner, the power savings due to a large scale server migration can be estimated. Nevertheless, the following migration strategies are generally the most effective:

- Use a 2-way server or a single processor dual core server to replace 2 or more old servers
- Use a blade based on a low-voltage or mid-voltage processor to replace an old server
- For servers with dedicated disk drives, use lower power enterprise class 2.5” drives instead of 3.5”
- Use a single dual core processor server to replace a dual processor server
- Use a 2-way dual core server in place of a 4-way server



This discussion suggests that migration is not typically the most effective tool for power consumption reduction. The major way that new server technologies can help reduce power consumption is when consolidation of applications on servers is used to reduce the total server count, or when servers are virtualized.

### Planning: Virtualization

Virtualization of servers results in a dramatic reduction of IT power requirements. Virtualization almost always dramatically reduces the number of installed servers. The elimination of a server is a structural consumption avoidance of approximately 200-400W, depending on technology. Therefore the electricity consumption avoided is approximately \$380 per year per server eliminated, and the total 10 year TCO cost saved by this structural avoidance is approximately \$7,680 per server eliminated. **This savings is substantially greater than the cost of the server itself.**

### Planning: Standardization

Standardization on energy efficient servers is a very effective approach, even if virtualization is not used. Today, blade servers are the most electrically efficient form of server. However, the available blade types for a given blade server system can vary dramatically in performance and power consumption. Often it is difficult to predict in advance the performance requirement for a server-based application, so users often specify the highest available performance at a substantial power consumption penalty.

When servers are virtualized, the strategy of using the highest performance server is generally the best approach to minimize overall power consumption. However, when servers are deployed application-by-application it can make sense to match the server performance to the application requirements to save energy.

For users that standardize on a blade server system and deploy servers per application, the option exists to standardize on two blades, a high performance / high power blade and a lower performance / lower power blade. The range of power consumption can be greater than two-to-one. A logical strategy is to deploy applications on the lower performance blade by default and only move to a higher performance blade if the need is demonstrated. This is facilitated by the ease of provisioning of blade servers. In this way, structural IT load consumption avoidance of 10% or more is possible for a typical business data center.

## Energy consumption reduction in NCPI equipment

The reduction of energy consumption of NCPI equipment is accomplished using the following techniques: right-sizing the NCPI system to the load, using efficient NCPI devices, and designing an energy-efficient system. Users may have some awareness of electrical efficiency of NCPI devices during the purchasing process, but the fact is that the data that is provided by manufacturers is typically not sufficient to determine

actual energy consumption differences, and furthermore right-sizing and system design each have a much higher impact on the electrical consumption than the selection of the NCPI devices.

## Right-sizing

Of all of the techniques available to users, right-sizing the NCPI system to the load has the most impact on NCPI electrical consumption. Most users do not understand that there are fixed losses in the power and cooling systems that are present whether the IT load is present or not, and that these losses are proportional to the overall power rating of the system. These fixed losses are the dominant form of NCPI electrical consumption in typical installations. In installations that have light IT loads, the fixed losses of the NCPI equipment commonly exceed the IT load. Whenever the NCPI system is oversized, the fixed losses become a larger percentage of the total electrical bill. For a typical system that is loaded at 30% of rating, the electrical cost per kW of IT load is approximately \$2,300 per kW per year. If the system were right-sized to the load, the electrical cost per kW of IT load falls to approximately \$1,440 per kW per year which is a 38% savings in electrical costs as shown in **Table 4**.

*Table 4 – The economic benefits of right-sizing a data center showing the 10 year cost per kW*

	Baseline Case	Right-Sized	Comment
IT Electricity	\$9,600	\$9,600	Assuming \$0.12 per kW hr
NCPI Proportional loss	\$960	\$960	
NCPI Fixed loss	\$12,800	\$3,840	Structural avoidance allows reduction in capacity-related electrical consumption
NCPI Capital Cost	\$13,330	\$4,000	Structural avoidance allows reduction in capacity equipment
NCPI Operating Cost	\$6,667	\$2,000	Reduction in equipment reduces operating expenses such as maintenance
Total NCPI Electrical Cost	\$13,760	\$4,800	Total of Fixed and Proportional losses
Total Electrical Cost (NCPI + IT)	\$23,360	\$14,400	
Total 10yr TCO	\$43,360	\$20,400	Including NCPI power and cooling capacity and power consumption expenses

Note that in addition to the electrical savings, right-sizing results in a savings of \$1,400 per kW per year of IT load in the NCPI capital and operating costs, which is almost as great as the electricity savings. These are the potential savings for a specific example; actual savings will vary and will be less for non-redundant systems

Right-sizing has the potential to eliminate up to 50% of the electrical bill in real-world installations. The compelling economic advantage of right-sizing is a key reason why the industry is moving toward modular, scalable NCPI solutions.

## Energy-efficient system design

Many users assume that the electrical consumption of a system is controlled by the efficiency of the individual components, and therefore that the main approach to reducing power consumption is to focus on the efficiency of the individual devices. This assumption is seriously flawed. The system design has an enormous effect on the electrical consumption of data centers, and two data centers comprised of the same devices may have considerably different electrical bills. For this reason, the system design is even more important than the selection of power and cooling devices in determining the efficiency of a data center.

Here are examples of system design issues that commonly reduce the efficiency of data centers to a much lower value than would be expected by summing the losses of the individual parts:

- Power distribution units and / or transformers operating well below their full load capacities.
- Air conditioners running with low output temperatures, continuously dehumidifying the air which must then be continuously re-humidified using a humidifier.
- Air conditioners that are actually heating while others in the same room are cooling.
- Air conditioners forced to consume excessive power to drive air against high pressures over long distances.
- Air conditioners operating with much lower return air temperature than the IT equipment exhaust temperature, which causes them to operate at reduced efficiency and capacity.
- Cooling pumps which have their flow rate adjusted by throttling valves which dramatically reduces the pump efficiency.

Note that this list is mainly comprised of design problems related to air conditioning. In fact most poor design practices that waste electrical power are related to air conditioning, because the power systems architectures are more standardized and therefore less prone to design-related errors.

The short list of problems above routinely cause data centers to draw **twice as much** NCPI power as necessary. Furthermore, all of these problems are avoidable at little or no expense by simple design decisions. There are two ways to avoid these problems:

- 1) Ensure the design has been fully engineered and tested to avoid the above problems, including complex Computational Fluid Dynamics modeling and comprehensive commissioning testing, or
- 2) Obtain a complete NCPI system based on a standardized design, comprised of modules that have been pre-engineered, pre-tested, and specified to avoid the above problems.

Due to extreme costs and variability of the first approach, the second of the above alternatives will become the standard way that data centers are specified and acquired in the future.

### Using efficient NCPI devices

Although the selection of NCPI devices such as power and cooling equipment has less effect on the overall system electrical consumption than does IT architecture, NCPI right-sizing or NCPI system design, device selection is nevertheless an important element in designing a power-efficient data center.

There is a substantial variation in the electrical losses between NCPI devices of the same type operated under the same conditions. For example, in a December, 2005 paper by the U.S. Electric Power Research Institute, it was found that different UPS systems operated at 30% of load rating varied in losses from 4% to 22%, which is a 500% variation. It is important to note that this variation cannot be ascertained from the specification sheets for these products. This paper and other white papers from APC clearly demonstrate that the electrical losses in real applications can only be correctly predicted if the appropriate models are used and that typical manufacturer's data is inadequate to make quantitative predictions of the electrical consumption of data centers. An example of the appropriate way to compare the electrical consumption of two NCPI devices is provided in APC White Paper #108, "Making Large UPS Systems More Efficient".

## Practical overall energy consumption reductions

This paper has demonstrated the magnitude of the electrical consumption problem and suggested various strategies to reduce consumption. Combining the approaches, it is possible to summarize the potential savings if a data center were optimized to reduce electrical consumption, as compared to a typical design.

**Table 5** summarizes 10 effective strategies that can be used to reduce electrical power consumption along with a range of savings when compared to typical data centers. These strategies are effective for new data centers, and some can be deployed immediately or over time in existing data centers.

*Table 5 – Practical strategies for reducing electrical power consumption for data centers, indicating range of achievable electrical savings*

	Savings	Guidance	Limitations
Right-size NCPI	10 – 30%	Using a modular, scalable power and cooling architecture Savings are greater for redundant systems	For new designs and some expansions Difficult to retrofit
Virtualize Servers	10 – 40%	Not technically a physical infrastructure solution but has radical impact Involves consolidation of applications onto fewer servers, typically blade servers Also frees up power and cooling capacity for expansion	Requires major IT process changes To achieve savings in an existing facility some power and cooling devices may need to be turned off
More Efficient Air Conditioner Architecture	7 – 15%	Row-oriented cooling has higher efficiency for high density (APC White Paper #130) Shorter air paths require less fan power CRAC supply and return temperatures are higher, increasing efficiency, capacity, and preventing dehumidification thereby greatly reducing humidification costs	For new designs Benefits are limited to high density designs
Economizer Modes of Air Conditioners	4 – 15%	Many air conditioners offer economizer options This can offer substantial energy savings, depending on geographic location Some data centers have air conditioners with economizer modes, <i>but economizer operation is disabled</i>	For new designs Difficult to retrofit
More Efficient Floor Layout	5 – 12%	Floor layout has a large effect on the efficiency of the air conditioning system Involves hot-aisle / cold-aisle arrangement with suitable air conditioner locations (APC White Paper #122)	For new designs and expansions Difficult to retrofit
More Efficient Power Equipment	4 – 10%	New best-in-class UPS systems have 70% less losses than legacy UPS at typical loads Light load efficiency is the key parameter, NOT the full load efficiency Don't forget that UPS losses must be cooled, doubling their costs	For new designs or retrofits

	Savings	Guidance	Limitations
Coordinate Air Conditioners	0 – 10%	<p>Many data centers have multiple air conditioners that actually fight each other</p> <p>One may actually heat while another cools</p> <p>One may dehumidify while another humidifies</p> <p>The result is gross waste</p> <p>May require a professional assessment to diagnose</p>	For any data center with multiple air conditioners
Locate Vented Floor Tiles Correctly	1 – 6%	<p>Many vented tiles are located incorrectly in the average data center or the wrong number are installed</p> <p>Correct locations are NOT intuitively obvious</p> <p>A professional assessment can ensure an optimal result</p> <p>Side benefit – reduced hot spots</p>	<p>Only for data centers using a raised floor</p> <p>Easy, but requires expert guidance to achieve best result</p>
Install Energy Efficient Lighting	1 – 3%	<p>Turn off some or all lights based on time of day or motion</p> <p>Use more efficient lighting technology</p> <p>Don't forget that lighting power also must be cooled, doubling the cost</p> <p>Benefit is larger on low density or partly filled data centers</p>	Most data centers can benefit
Install Blanking Panels	1 – 2%	<p>Decrease server inlet temperature</p> <p>Also saves on energy by increasing the CRAC return air temperature</p> <p>Cheap and easy with new snap-in blanking panels such as those by APC</p>	For any data center, old or new

**Table 5** above provides a summary of some of the most powerful and practical tools to reduce electrical power consumption of data centers. The energy consumption reduction values were estimated using energy calculations based on APC White Paper #113 (referenced earlier in this paper), applied to a range of data center designs. In addition to the items on this list are other sophisticated IT architectural strategies that are mentioned earlier in this paper.

Some of the savings described above can be built into the equipment provided by vendors, but the majority are related to the system design and installation. Pre-engineered standardized system designs are available from some vendors that have been optimized and verified for high efficiency; consult your supplier. For users with existing facilities that seek to reduce the electrical consumption, customers can attempt to implement the guidance provided above, or some suppliers such as APC offer a Data Center Energy Efficiency Assessment Service that uses specialized tools and methods specifically designed for data centers.

## Conclusion

The cost of electricity for data centers is a substantial operating cost that can and should be managed. A data center designed for reduced power consumption also saves other costs such as capital and operating costs associated with power and cooling systems, as well as saving space.

Electrical consumption of existing data centers can be reduced through various low cost methods but primarily via migration to more energy efficient computing platforms. For new data centers, there are additional options both in the IT architecture and in the NCPI architecture that can gain much greater savings.

The electrical power consumption is typically shared evenly between the IT loads and NCPI devices. Any rational approach to reduction in electrical usage must treat the combined IT / NCPI design as a system in order to maximize the benefit.

Some equipment providers offer complete standardized data center designs specifically engineered for efficiency, and energy efficiency audit services are available for users desiring to reduce power consumption in existing data centers.

The cost savings opportunities have been shown to be very large **yet the investment required to achieve them is small or even zero in some cases**, when compared with legacy approaches to data center design.

### About the Author:

**Neil Rasmussen** is a founder and the Chief Technical Officer of American Power Conversion. At APC, Neil directs the world's largest R&D budget devoted to Power, Cooling, and Rack infrastructure for critical networks, with principal product development centers in Massachusetts, Missouri, Denmark, Rhode Island, Taiwan, and Ireland. Neil is currently leading the effort at APC to develop modular, scalable data center infrastructure solutions and is the principal architect of APC's InfraStruXure system.

Prior to founding APC in 1981, Neil received his Bachelors and Masters degrees from MIT in electrical engineering where he did his thesis on the analysis of a 200MW power supply for a Tokamak Fusion reactor. From 1979 to 1981 he worked at MIT Lincoln Laboratories on flywheel energy storage systems and solar electric power systems.