

# Power Efficiency and Storage Arrays

## *Technology Concepts and Business Considerations*

---

### **Abstract**

Factors like cost and availability of energy for the data center are rapidly gaining in importance as both usage and costs increase. This white paper looks at electrical power, the ways that it's measured, and the design and configuration factors involved in developing efficient systems.

May 2008

---

---

Copyright © 2006, 2007, 2008 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number H2402.3

---

## Table of Contents

<b>Executive summary .....</b>	<b>4</b>
<b>Introduction .....</b>	<b>4</b>
Power and efficiency .....	5
Audience .....	5
<b>Electrical power fundamentals .....</b>	<b>5</b>
Terminology .....	5
Energy .....	6
Power .....	7
Efficiency .....	7
Current and power factor.....	7
<b>Facilities planning considerations .....</b>	<b>8</b>
Specifications and the electrician .....	8
Derating.....	8
Data center efficiency, keeping your cool .....	8
Calculating operating costs.....	9
<b>Configuring for efficiency .....</b>	<b>9</b>
Scalability .....	9
Drive performance versus larger capacity .....	10
Flash drives.....	11
Tiered storage.....	11
RAID level .....	11
The role of software .....	11
<b>Conclusion .....</b>	<b>12</b>
<b>References .....</b>	<b>12</b>

---

## Executive summary

As energy consumption and the costs of energy continue to increase, the necessity to deploy energy-efficient storage systems has been widely recognized. Technology suppliers must increase the efficiency of systems that they provide, and they need to offer informed choices to users where tradeoffs exist.

Users, on the other hand, must make choices about performance, availability, expandability, configurations, and workload according to their business requirements. Obtaining correct and complete information from the vendor is critically important in making the right selections and reducing overall costs.

In addition, there are the physical planning aspects to be considered, and determining energy efficiency is sometimes difficult as terminology and metrics may differ.

This white paper provides an overview of energy concepts and terminology, with recommendations on energy-efficient facilities planning and how to configure efficient storage systems.

## Introduction

Energy consumption is rapidly growing in importance in IT. The price of energy is rising sharply, with no relief in sight, and the amount of information to be stored continues to grow at a rate of 60 percent per year. Customers are demanding that the products that they buy are designed with power efficiency in mind, and that they have access to the information that they need to make informed choices.

Energy management can be complex. Consumption varies with workload, availability requirements, time of year, and other variables that make understanding the issues nontrivial. Power use produces heat, which requires energy to cool, and all this must frequently be managed with redundancy for high availability and disaster recovery.

### The energy capacity of data centers is a limiting factor

The limiting factor to growth for many customers today is the availability of energy in the data center. In many cases, there is no additional energy capacity available, or the limits of the air conditioning systems have been reached, so that nothing can be added without removing something else of similar power. Gartner, Inc. predicts<sup>1</sup> that by 2008 50 percent of data centers will have insufficient power and cooling capacity to meet the demands of high density equipment.

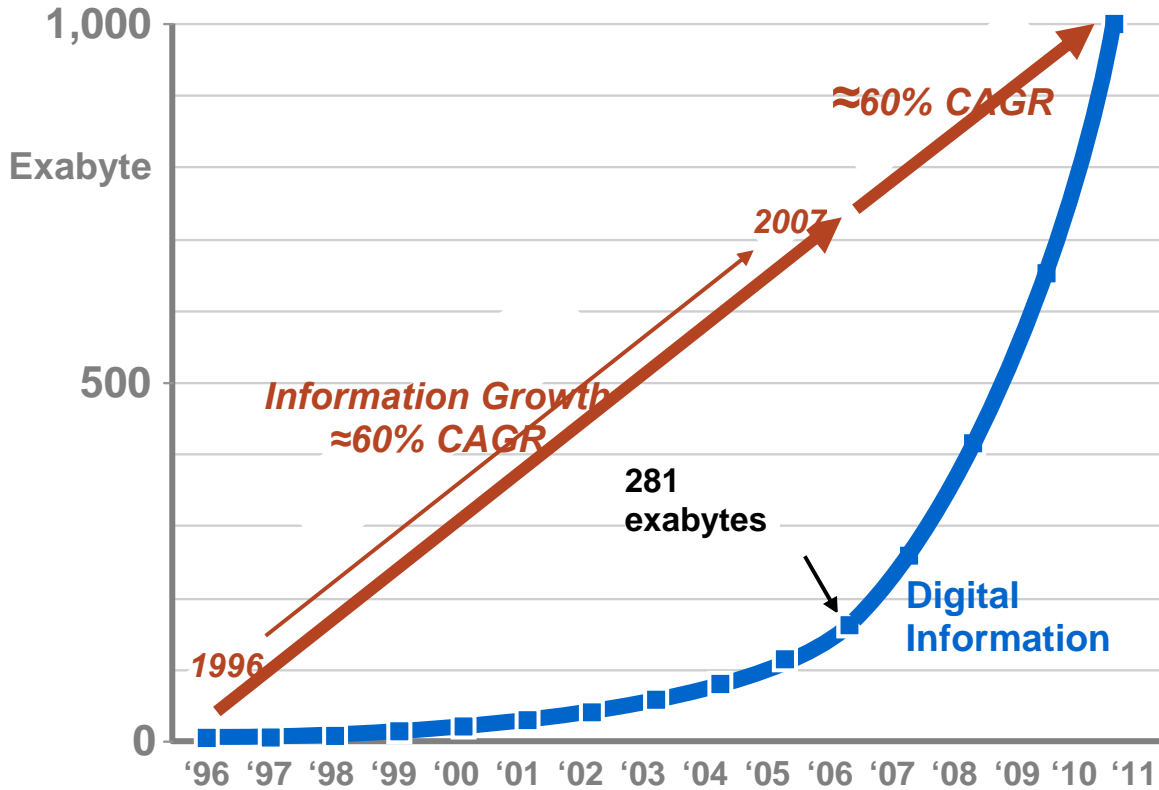
### Power usage and cost continue to grow

A recent IDC survey<sup>2</sup> indicated that the amount of data that business was expected to store was growing at nearly 60 percent per year. This means that a data center with 500 TB today will be faced with the challenge of storing around 3 PB in the next five years.

---

<sup>1</sup> Gartner, "Meeting the DC Power and Cooling Challenge," December 2006

<sup>2</sup> IDC, "End User Perspectives on Server Blade Adoption," May 2006



**Figure 1. The amount of data to be stored is growing exponentially**

Energy costs have increased about 56 percent since 2002, and no relief is in sight. Thus, something that was considered a minor nuisance a few years ago is now becoming a major expense. Figure 1 illustrates the current costs.

### Power and efficiency

To properly address these issues, it is becoming increasingly important to have the right information and to know how to interpret it. This white paper will introduce the important concepts necessary to ask the right questions and obtain the proper information. EMC offers a variety of resources and tools intended to make this easier, which are available from your local EMC office.

### Audience

This white paper is intended for a nontechnical audience wishing to understand how electrical power is measured and the importance of power engineering, and how to configure storage arrays for maximum power efficiency.

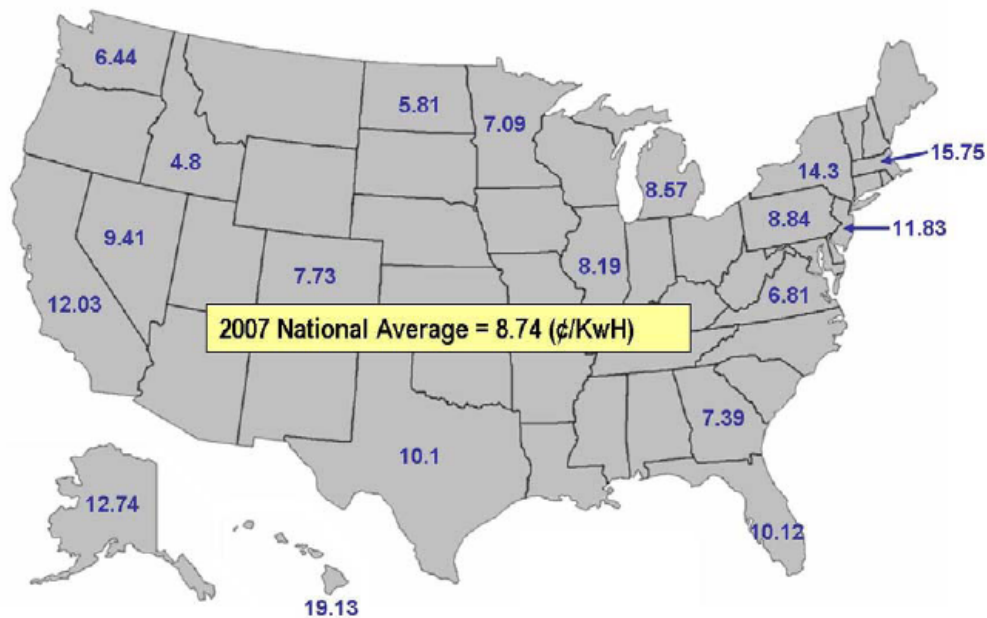
## Electrical power fundamentals

### Terminology

The subject of energy efficiency requires a thorough understanding of several terms and concepts. Some of the terms — *energy* being the most obvious — are used in everyday conversation in ways that are inappropriate for a somewhat technical discussion, and terms like *energy* and *power* are used interchangeably when they actually have quite different meanings.

## Average Retail Price of Electricity (¢/KWh)

As of March 31, 2007



**Figure 2. Average price of electricity** (source: The Enterprise Strategy Group, ESG Report, *Power, Cooling, Space Efficient Storage*, 4AA1-3700ENW)

## Energy

Energy is a fundamental entity. It's what you pay for when you get your electric bill, and can be measured in many ways. At its most basic, it's the ability to do work, where work can mean moving things, lifting things, heating things, or similar activities. It can be changed from one form to another, and each form has a different set of units to measure it, some of which can be confusing. Physicists are well acquainted with the law of conservation of energy, which holds that energy cannot be created or destroyed, and no counterexample to this has ever been found.

For our purposes, the two most useful forms of energy are electrical energy, and heat energy. Electrical energy is generally measured in watt-hours (Wh), or more commonly kilowatt hours (kWh). A kWh is what you pay for on your electric bill, generally between \$.05 and \$.15 per kWh.

The other useful measure of energy in this discussion is heat energy. The end result of using electrical energy in the data center is that it all ends up as heat. Electrical circuits, spinning drives, power supplies, UPSs, and spinning fans all turn electrical energy into heat energy. Heat energy is commonly measured in British Thermal Units, (BTUs), which is the amount of energy used to raise a pound of water one degree Fahrenheit. (A calorie, incidentally, is the metric version of a BTU, although we generally use it to measure the energy in food, and not data centers.)

Of course, there are many other forms and measurements that don't apply here. When you wind a clock, stretch a spring, push an object to put it in motion, or raise an object off the ground, you're changing the chemical energy in your food into some other form of energy.

### An energy imposter?

The recent popularity of so-called "energy" drinks make for an interesting contradiction. Some of them claim zero calories, which means that they have no energy at all! They may make you feel "peppy," but that's not at all the same, since the energy content is measured in calories, which is the metric version of BTUs.

---

## Power

In the simplest sense, power is the rate at which you use energy. (A physicist would say that it's the rate at which you do work, but it amounts to the same thing here.) Electrical power is measured in watts (W) or kilowatts (kW), and power stations are measured in megawatts (mW) or more. Thus, a 200W light bulb will consume energy at twice the rate of a 100W bulb. The amount of energy used is calculated by multiplying the wattage of the device by the time in service.

A 100W bulb that is operated for 10 hours will consume 100x10, or 1,000 Wh, or 1 kWh. At \$0.15/kWh, this will cost you 15 cents to run for 10 hours, 36 cents for 24 hours, and \$131.40 to leave it on for a year. (Don't forget to turn out the light when you're through reading this.)

In the case of heat, power is measured in BTU/hour (remember, power is that rate at which we use energy). When electricity is turned into heat (which is all the time in the data center), one watt produces 3.413 BTUs per hour. So, to summarize:

**Table 1. Common terms**

Useful units in the data center	Electrical	Heat
Energy	kWh	BTU
Power	kW	BTU/hr

## Efficiency

Efficiency, in this discussion, is best described as the relationship between the energy used and the quantity of some desired result. Sometimes this is straightforward. In a power supply, for example, you might be putting in 120V AC, and getting out 12V DC. If you compare the watts in and watts out, you can calculate the efficiency straight away. The difference is given off as heat, of course; we'll get to that a bit later. Depending on your energy source, and the desired result, the units of efficiency can be almost anything. A simple percentage might work some of the time, as we've observed for power supplies, and we see things like miles/gallon, BTUs per gallon of heating oil, BTUs per ton of coal, and hundreds of others.

In the data center, measuring efficiency can get a bit tricky, since it's difficult to specify and measure the desired output. In the case of servers, the logical unit would be some sort of performance metric, like MIPS/Watt, MHz/Watt or transactions/Watt, but there is no agreement on how this should be measured, and no readily accepted efficiency metric. (The EPA is looking at including IT equipment in the Energy Star program, but the inability to measure performance, a necessary ingredient in efficiency, is getting in the way.)

For storage, the problem is no easier. Most common is some sort of capacity metric, like GB/W (or GB/kW) of online storage, which can be raw capacity, usable capacity adjusted for RAID protection and spare drives, or other measures of usable capacity where various products may not perform well when fully used. Performance-based metrics like IOPS/Watt have also been proposed. Raw capacity is the most common today, as it is easily adjusted for other factors.

## Current and power factor

The power drawn by IT equipment is typically expressed in either watts (W) or volt-amperes (VA). Watts are the "real" power or useful work that you get from the utility. Volt-amperes are simply the product of the input voltage and current and is known as the "apparent" power. For common devices, where the load is resistive, such as a light bulb, real power and apparent power are equal. The situation gets a bit tricky with electronic power supplies, however, as the load is somewhat capacitive, and so some of the current whizzes by without doing useful work. In this case, simple multiplication (kVA) will yield a number that is higher than the wattage as measured. The difference, expressed as a percentage, is known as the power factor.

Power supplies for most high performance IT equipment today are power factor corrected, so the difference is only a few percentage points, and can almost be ignored. For older or simpler power supplies, such as

---

your home PC, the cheap power supplies may not be power factor corrected, and the power factor may be as low as 60 percent.

This difference is especially important in the data center, though, where power conditioners and UPSs are commonly used, since they must deliver current at a specified voltage and thus are rated on kVA. Watts, on the other hand, are what you pay the electric utility for, and also determine the amount of energy that is turned to heat. (Remember all of the electricity consumed by IT equipment in the data center ends up as heat.)

## **Facilities planning considerations**

### ***Specifications and the electrician***

The electrical requirements for IT equipment are normally specified in two ways: the input watts or volt-amperes, and the circuit breaker rating required to protect the equipment. EMC storage products are power factor corrected, which means that the volt-ampere and wattage specifications are virtually identical. The published documentation provides the customer with kVA ratings so that the wiring and power allocation for the customer Power Distribution Unit (PDU) can be easily determined.

One question that comes up quite often concerns the customer circuit breaker requirements and the allocation of power within the data center. If the IT equipment manufacturer specifies a 50A circuit breaker requirement for a system operating on three-phase line to line (Delta) installations and a 32A circuit breaker for systems operating from a three-phase line to neutral (WYE) installation, this does NOT mean that the IT equipment will consume 50A/32A of power. The consumer must provision power within the data center based on the 50A/32A circuit breaker rating called out in the documentation, so the amount of power allocated within the data center will be much greater than what's actually used. The problem is made more complex by the redundant AC input power connections that are required for virtually all vendors of storage array equipment. Customers should ask vendors for detailed energy usage for the specific configuration they are considering.

The local electrician also needs to verify that each bay of equipment is connected to two independent customer PDUs. These two AC connections are what provide redundant, high availability AC power to each of the system and storage bays.

### ***Derating***

As mentioned in the previous section, the input kVA rating for the storage array equipment may in some cases be much lower than the rating required for the data center circuit breaker and AC wiring infrastructure. In some cases this is done in order to allow the customer to easily upgrade their system with either additional logic (additional DA pairs) or larger capacity disk drives. There are also Safety Agency input connector derating requirements that vary by country and this can also lead to equipment current ratings that seem to be greater than expected based on the configuration of the equipment.

An example of connector derating would be the requirement that UL and CSA have in place for the United States and Canada. The requirement reads: "For equipment connected to a standard US or Canadian source of power, the attachment plug cannot be rated less than 125 percent of the rated current of the equipment at the nominal system voltage range as described by the configuration of the plug."

This requirement means that the equipment connector cannot be used at a current rating greater than 80 percent of its nameplate rating. If a connector is rated for 50A of current, the maximum current rating of the equipment cannot exceed 40A.

### ***Data center efficiency, keeping your cool***

So far, we've looked at ways to measure the energy used. But all this equipment, servers, power supplies, logic bays, spinning drives, BBUs, spinning fans, lights, and network equipment are putting out heat at the rate of 3.413 BTU/hr for every watt. Server energy densities keep getting higher, and the amount of power

---

in a single rack has gone from a few hundred watts to several thousand, and we can put 240 of the 15k rpm drives in a single bay, using over 5,000 watts. As a result, air circulation and cooling have become important considerations.

It turns out that in a very efficient data center, with proper design, clean space under the floor for air circulation, well-placed cold air outlets, hot and cold aisles segregated and other factors, for every watt of power used by the IT equipment, another watt of power is necessary to remove the heat produced. The most common way to measure this is Data Center Efficiency (DCE), which in this case would be 50 percent, since half the energy is used to do useful work, and the other half to cool it. Most data centers run at DCE figures of less than 50 percent, but there is no standard answer for this. Unless a customer is in the process of redesigning a data center, it's quite unlikely that he'll know what the DCE is for his own installation, so you'll have to make an estimate. If it was built in the last few years, you might start with 50 percent; if it's an older building that's been there for 20 years, try 40 percent and hope for the best. The lower the DCE is, the greater the benefit of efficient storage is.

## **Calculating operating costs**

There are three factors needed to calculate the annual cost of energy in the data center:

- **P:** price of electricity itself, given as the price / kWh
- **W:** energy consumption of the device, given in kW
- **DCE:** Data Center Efficiency factor

With this information, calculating the cost for a year (a year contains 8,760 hours) is quite simple. The energy cost for a year, EC, can be calculated as follows:

$$EC = (P \times W \times 8,760) / DCE$$

Here is an example for a 5 kW device. Assume that:

Electricity is \$.15/kWh

The DCE is 50% (half the power goes to IT equipment, half to air conditioning and power conversion)

$$EC = \$.15 \times 5\text{kW} \times 8,760 / .50$$

$$EC = \$13,140/\text{year}$$

## **Configuring for efficiency**

There are configuration considerations that can be made when selecting a storage array for power efficiency. Work with suppliers that are energy-conscious and that attempt to help customers reduce their total cost of ownership through utility costs. One industry consortium, 80 PLUS, is leading the way in promoting and certifying energy-efficient products. The 80 PLUS specifications, which EMC has adopted, stipulates that power supplies should be at least 80 percent efficient and have a minimum of 0.9 power factor. For further information, refer to their website at [www.80PLUS.org](http://www.80PLUS.org).

It is important for the storage provider to design power-efficient disk drives and other components to reduce the equipment power load. EMC uses specially designed disk drives that have a predetermined and consistent power budget. This allows the "next generation" disk drive to be upgraded in fielded legacy cabinets. The future disk drives will meet today's power and cooling budget. The disk drives are by definition, highly efficient already. Since the disk drives are the single most power-intensive subsystem in the array, this translates into ongoing power savings for the customer.

## **Scalability**

To add storage incrementally as demand grows, there are two options: Buy new systems and connect into a storage area network or buy additional storage and add it to the existing storage array. From a power efficiency standpoint, the latter is the better option. Adding storage incrementally can better match

increased storage demand, allowing fewer disk drives and logic infrastructure. This option also allows the existing logic infrastructure to be amortized over more drives, thereby reducing the watts per terabyte. Some storage providers require additional logic bays for every two or four storage bays. EMC<sup>®</sup> Symmetrix<sup>®</sup> DMX-3 and DMX-4 systems support up to 10 storage bays per logic or system bay, and CLARiiON<sup>®</sup> leads midrange systems by supporting 480 drives.

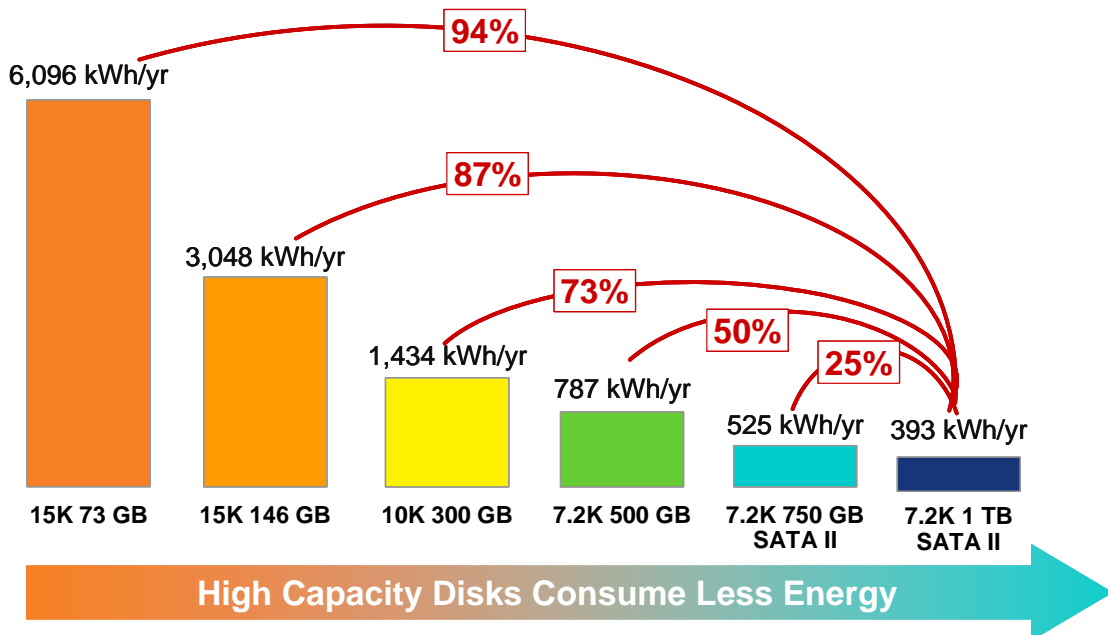
The number of host channel directors, disk adapters, and cache memory can also be increased as storage demand grows. It is best to design the system for the intended applications and environments, then scale using the smallest number of storage bays to reduce inefficiencies.

### Drive performance versus larger capacity

Since we're looking at the power requirement per TB of storage, the easiest way to reduce power is to use the largest capacity drives that your application and performance requirements will permit. Since the power used by the drive is about the same regardless of capacity, the larger the drive, the more efficient it will be in terms of watts/TB. With the availability of 500 GB and 1 TB drives, which are ideal for backup to disk, data warehousing, and similar applications, power savings can be significant.

## Today: Energy-Efficient Storage Design

### 1 TB SATA II Data on Different Capacity/Performance Drives



**Figure 3. Energy required to store 1 TB of data on different capacity drives**

Not all applications are suited for large capacity drives. Many require high I/O and fast response times. For these applications, such as Exchange and Oracle databases, high performance, low capacity drives such as the 15k 73 GB types are well suited. The advantage of low capacity is that overall performance depends on the number of spindles, so having more small drives will provide better performance than fewer large drives of the same speed, but power consumption will, of course, be higher. The goal, though, is to minimize the number of high-performance drives and use them for only the most data intensive applications. A larger capacity drive such as a 15k 146 GB could also be used, but the number of spindles will be halved, so careful analysis is required. It may be better to configure the system for the number of spindles required for the required performance rather than configure only for the required capacity.

The EMC Sales teams have tools that can calculate the power used by specific configurations, and others that can calculate the back-end requirements of the workload and determine the number of drives needed.

---

This can give you the information that you need to make informed choices about performance and energy use. A model of the proposed system with the consolidated workload can then be used to ensure the appropriate number of physical drives is configured.

Once the high-performance applications are satisfied, the remaining storage needs can be consolidated on larger capacity drives. This can include 10k 300 GB and 7.2k 500 GB and 750 GB drives. By consolidating the storage needs, fewer drives are required, which also reduces the supporting infrastructure that needs to be powered up.

## **Flash drives**

The recent availability of flash drives, as an alternative to rotating drives, offers a couple of interesting considerations. Most obviously, a flash drive uses about 38 percent less power than a high performance rotating drive of the same capacity. Of greater interest, though, are the performance-based savings: A flash drive can provide as much as 30 times the IOPS performance, so for applications that use a large number of spindles to get performance, the number of drives can be reduced by a factor of 30, resulting in a 98 percent reduction in energy use.

## **Tiered storage**

Another way to gain efficiency is to use fewer, large-scale storage systems. For this to work, the systems need to have the capability to configure mixed environments within a single system. For example, the smaller high performance drives can coexist with larger, slower drives, and data can be easily moved from one type of environment easily within the storage system. This is known as tiered storage, where multiple tiers can be combined into a single, large system. Careful configuration will provide the highest performance with the best predictability for critical applications and a more cost-effective, lower-performance tier for less critical applications. This type of configuration can be designed to most optimally meet your storage requirements and minimize the power needs.

## **RAID level**

This is well known, of course, but needs to be considered as a key factor in your calculations and not as an afterthought. RAID levels can differ quite a bit in their energy efficiency, and become part of the inevitable trade-off between performance and cost. Drives used for protection, whether mirrors or parity information, are not available for data, but use the same energy as any other drive. It also turns out that the higher the performance of a RAID group is, the higher the energy cost is. Table 2 is a quick summary of the various RAID types and their relative efficiency.

**Table 2. RAID types and efficiency**

<b>RAID level</b>	<b>Storage efficiency</b>
RAID 1	50%
RAID 5 (3+1)	75%
RAID 6 (6+2)	75%
RAID 5 (7+1)	87.5%
RAID 6 (14+2)	87.5%

## **The role of software**

This paper has looked carefully at the hardware used in storage. After all, this is what gets plugged in and draws power. It's important not to overlook the critical role that software plays in reducing energy use, though. The ability to transparently move data from one tier to another, to archive unused data, to de-

---

duplicate backup applications, virtual provisioning, and a number of others, can play a critical role in reducing your energy needs and saving money. While hardware draws the power, software can seriously reduce the amount needed, and should play a primary role in efforts to conserve.

## Conclusion

Controlling energy consumption in a data center environment is becoming more important with every passing day. There is a great deal that storage suppliers can do to make their products energy-efficient, and much that IT managers can do to develop efficient configurations.

To plan an energy-efficient data center, it is critical to review current power, energy, efficiency, and operating costs. EMC Sales teams have tools that can calculate the power used by specific configurations to provide the information that is needed to make informed choices about performance and energy use.

Working together, IT managers and storage suppliers can reduce energy costs significantly

## References

The following two documents provide additional sources of Symmetrix information. They are available on EMC.com and Powerlink<sup>®</sup>, EMC's customer- and partner-only extranet.

- *EMC Symmetrix DMX-3 and DMX-4 Electrical Power Estimation and Configuration Planning* white paper
- *EMC Symmetrix DMX-3 Best Practices Technical Note*