

Power Efficiency and Storage Arrays

Technology Concepts and Business Considerations

Abstract

The cost and availability of energy for the data center is 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.

October 2006

Copyright © 2006 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 H2402.1

Table of Contents

Executive summary	4
Introduction	4
Audience	5
Electrical power fundamentals	5
Power, efficiency, and how power consumption is measured	5
Facilities planning considerations	6
Specifications and the electrician	6
Derating.....	6
The cost of removing heat	7
Configuring for efficiency	7
Scalability	7
Drive performance versus larger capacity	8
Tiered storage	8
Conclusion	9
References	9

Executive summary

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

Users, on the other hand, must make the choices about performance, availability, expandability, configurations, and workload according to their business requirements. The right selections can contribute significantly to reducing costs. Getting the correct and complete information from the vendor is critically important.

In addition, there are the physical planning aspects to be considered, which sometimes use different terminology and metrics than those used for efficient configurations.

This white paper examines each of these factors, with recommendations on what information is required, and on how to configure efficient storage systems.

Introduction

Energy consumption is rapidly growing in importance in IT. The price of energy is rising rapidly, with no relief in sight, and the amount of information to be stored continues to grow exponentially. Customers are demanding that the products that they buy tomorrow 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 produces heat, which requires energy to cool, and all this must frequently be managed for power failure as well. This white paper introduces the basic concepts and terminology so that you can make informed decisions about your storage and power requirements.

Power usage and cost continue to grow

A recent IDC survey¹ indicated that users expect power consumption to increase by 25 percent over the next two years. Since university research indicates that the amount of data to be stored is increasing by 70 percent per year, the challenge is to do more and store more while holding power increases to much less than that.

The importance of energy cost

That same survey² of CIOs indicated that power provisioning, the ability to obtain power for the data center, and power consumption were both among the top three concerns as shown in Figure 1, and had a weighting of over five on a scale of 0 to 6.

¹ IDC, End User Perspectives on Server Blade Adoption, Doc # TB20060525, May 2006

² IDC, End User Perspectives on Server Blade Adoption, Doc # TB20060525, May 2006

Power and floor space are top data center issues

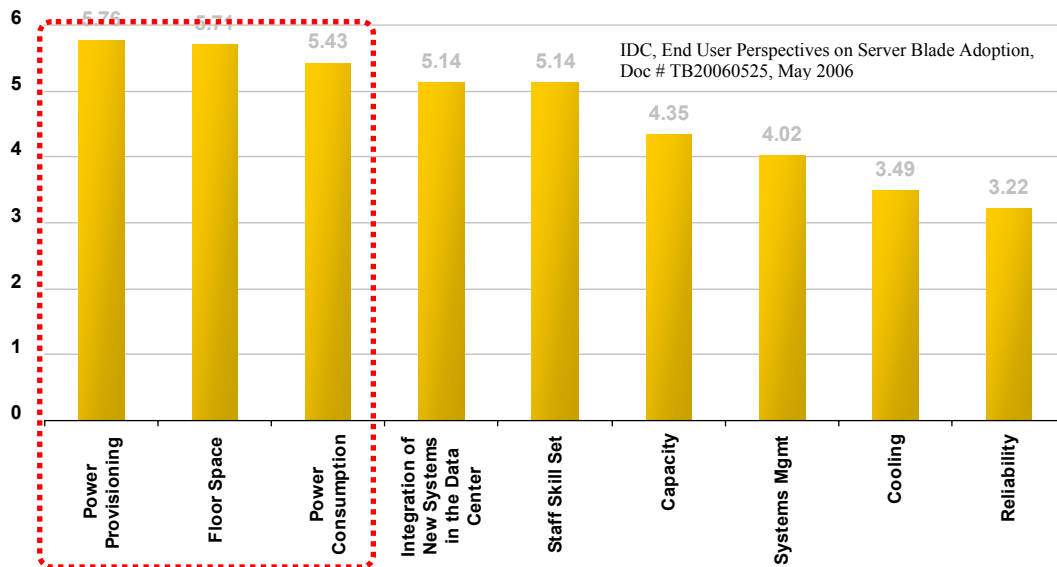


Figure 1. IDC's top data center issues

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

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

Electrical power fundamentals

Power, efficiency, and how power consumption is measured

If you manage a data center or own a home the general rule is that you pay an electric bill based on the number of kilowatt-hours (kWh) consumed on a monthly basis. Basically this is a simple calculation that takes the number of kilowatt-hours consumed multiplied by the \$/kW rate set by the local utility company. Power efficiency, in the simplest case, is typically the ratio of output power (watts) divided by the input power it takes to produce it. That's a good way of evaluating motors and power supplies, but it's not helpful in describing storage since the output can't be measured that way. Watts per terabyte of storage is a more useful way of comparing storage equipment.

Some manufacturers of storage arrays require the addition of a second system controller/logic bay once the number of disk drives exceeds a specific level, but today's storage systems must be designed with future scalability in mind. EMC has addressed this problem by providing an efficient and expandable design, reducing the need for duplicating the system bay logic as the total number of disk drives increases. EMC®

Symmetrix® DMX-3 and EMC CLARiiON® can be much more efficient due to a reduced need for additional logic bays, as described in the “Scalability” section.

Specialized EMC tools are used to calculate power consumption so that customers can plan facilities and budgets with increased precision.

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. The DMX-3 and CLARiiON power systems 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 section above, 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

Volt-amperes versus watts

The power drawn by IT equipment is typically expressed in either watts (W) or volt-amperes (VA). The 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 are known as the “apparent” power. The wattage rating determines the power that is needed from the local utility company to operate the equipment, and the VA rating is used to determine the size of the data center power infrastructure used to operate the equipment, that is, circuit breakers and the AC wiring supplying power to the system.

It’s also important to understand Power Factor when dealing with volt-amperes and watts. The Power Factor is the ratio between the real power (watts) and the apparent power (volt-amperes) and represents whether the IT equipment has a capacitive input such as in a switch-mode power supply, or an inductive input that is typical for a motor or a transformer. In practically all modern IT equipment the power supplies present a capacitive load on the AC line and also utilize power factor corrected switch-mode power supplies. This means that the power factor (or the ratio of the real to apparent power) of the IT equipment will be very close to 1. Simply stated this means that the wattage and volt-ampere ratings will be practically identical values. In the case of some PC power supplies and “adapter” type power supplies, power factor corrected switch-mode power supplies are not utilized and in these cases the power factor can be as low as 0.6. This translates into a wattage rating that is dramatically different than the volt-ampere rating. In cases like these it’s very important to know both ratings so that power in the data center can be distributed appropriately and the infrastructure can be sized correctly for the input current requirement.

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 as follows: "For equipment connected to a standard US or Canadian source of power, the attachment plug cannot be rated less than 125% of the rated current of the equipment at the nominal system voltage range as described by the configuration of the plug."

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

The cost of removing heat

Now it's clear that the lower the amount of input power a system requires, the lower the annual electric bill will be. The same can also be said for air conditioning costs to keep the IT equipment cool within the data center. The higher the level of efficiency of the power system, along with the use of a scalable storage system, will allow the customer to size the system according to their current requirements, which in turn leads to lower air conditioning costs. In general you can consider every watt that the storage system consumes to be released into the data center, which then needs to be removed by the air conditioning system. This adds to the annual cost of operating the equipment. Depending on the time of year and geographic location of the data center, the electricity cost to remove the heat is typically 40 percent to 60 percent of the input power to the cabinet. This also includes waste heat from the voltage transformers, UPS, and electric distribution system. Customers should choose storage systems that are designed to provide the lowest W/TB metric in order to save on their annual electric bill cost.

To assist in heat removal and more efficient air conditioning, it is best to place air return handlers near the source of exhaust heat and prevent warm air from mixing with colder air. This improves the efficiency of the air handlers and prevents preheating surrounding air and increasing operating temperatures.

EMC has a patent pending for the DMX-3 cabinet cooling design, which specifically addresses the issue of directing heated exhaust and simplifies the task of designing data centers to control its flow.

Configuring for efficiency

There are configuration considerations that can be made when selecting a storage array for power efficiency. First of all, it is a good idea to work with suppliers that are energy-conscious and 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.

It is important for the storage provider to design power-efficient disc drives and other components to reduce the equipment power load. EMC uses specially designed disc drives that have a predetermined and consistent power budget. This allows the "next generation" disc drive to be upgraded in fielded legacy cabinets. The future disc drives will meet today's power and cooling budget. The disc drives are by definition, highly efficient already. Since the disc 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, customers typically have two options: They can buy new systems and connect into a storage area network or they can 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 disc drives and logic

infrastructure. It 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 DMX-3 supports up to 10 storage bays per logic or system bay, and CLARiiON 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 drives, which are ideal for backup to disk, data warehousing, and similar applications, power savings can be significant.

Not all applications are suited for large capacity drives, though. Many require high I/O and fast response times. For these applications, such as Exchange and Oracle database, 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 back-end requirements of the workload and determine the number of drives needed. 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 drives. By consolidating the storage needs, fewer drives are required, which also reduces the supporting infrastructure that needs to be powered up.

RAID 1 full mirroring requires nearly double the drive overhead as other RAID schemes. Some critical applications may require full mirroring, but many do not, and sufficient protection can be achieved with other RAID levels, such as RAID 5. Using hot spares in the array can further increase availability without incurring a large power expense.

Tiered storage

Another way to gain efficiency is to use fewer, large-scale storage systems. For this to work, though, 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.

Conclusion

Controlling energy consumption in a data center environment is becoming more important with every passing day. As the need for processing power grows, along with the need for more storage, so does power consumption and the concentration of that consumption into smaller and smaller areas.

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. Together, it is possible to reduce energy costs significantly.

To do this, it is critical for the user to have the necessary information, both for the products themselves and the way that they will be used, to plan effectively.

References

The following documents can be found on the EMC website at <http://www.emc.com/techlib/>.

- *EMC Symmetrix DMX-3 Electrical Power Estimation and Configuration Planning*
- *Symmetrix DMX-3 Best Practices Planning*