

# Achieving Energy Efficiency with EMC CLARiiON CX3 Storage Systems

*Best Practices Planning*

---

**Abstract**

The energy efficiency of external storage systems depends on many interlinked variables. When designing a strategy to reduce system power usage, you need to consider performance, cost, energy, and cooling. This white paper documents best practices for configuring and utilizing EMC® CLARiiON® energy-efficient technologies to achieve optimal system performance while considering power, cooling, and cost.

October 2007

---

---

Copyright © 2007 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 H4009

---

## Table of Contents

<b>Executive summary .....</b>	<b>4</b>
<b>Introduction .....</b>	<b>4</b>
Audience .....	4
<b>Current and future energy consumption considerations .....</b>	<b>4</b>
<b>The Fundamentals .....</b>	<b>5</b>
Terminology .....	5
Btu .....	5
Watts .....	6
Volt Amperes (VA).....	6
Power/cooling projections .....	6
<b>CLARiiON solutions.....</b>	<b>7</b>
Multi-tiering.....	7
Disk drive classes.....	7
High-capacity 750 GB SATA drives .....	8
Drive types and energy consumption .....	8
RAID types .....	9
RAID 0: Speed (widely used, but not recommended).....	9
RAID 1: Disk mirroring (widely used) .....	9
RAID 3: Speed and parity protection.....	10
RAID 5: Speed and parity protection (widely used) .....	10
RAID 6: Double parity protection.....	10
RAID 1/0: Disk mirroring.....	10
Power consumption .....	10
Storage consolidation and utilization .....	11
System utilization .....	12
MetaLUNs .....	12
Virtual LUN technology .....	14
Example of cost savings with virtual LUN migration .....	14
Navisphere Quality of Service Manager .....	15
Prioritizing applications during peak energy hours.....	15
Performance and power.....	17
Application use cases .....	17
<b>EMC solutions .....</b>	<b>18</b>
EMC virtualization technologies.....	18
VMware server virtualization (server consolidation).....	19
Storage virtualization, EMC Rainfinity, and Invista .....	20
Avamar (data de-duplication).....	20
EMC Power Calculator.....	20
EMC data center energy efficiency services.....	20
The Green Grid .....	21
<b>Conclusion .....</b>	<b>21</b>
<b>References .....</b>	<b>21</b>

---

## Executive summary

Many data centers are running out of power and cooling capacity. In some cases, there is not enough physical room left to power up new systems or add more resources. Facility managers cannot keep up with the cooling and airflow required for hotter-running server/storage equipment. Furthermore, the cost of energy has been on a steep incline. There is a significant need to increase energy efficiency in data centers, and IT professionals continue to look for ways to reduce the cost of cooling and powering their data centers.

## Introduction

Energy considerations impact business operations and customer buying decisions. This white paper describes the best practices for improving power usage and cooling requirements by reviewing the following key topics:

- General energy and cooling concepts that relate to storage and data center power efficiency.
- How to best use CLARiiON<sup>®</sup>'s intelligent embedded technologies to reduce energy consumption. These technologies include:
  - Virtual LUN technology for online nondisruptive balancing
  - MetaLUN technology for efficient provisioning
  - Navisphere<sup>®</sup> Quality of Service Manager (NQM) for monitoring and managing resources
- IT solutions for tiered storage, archiving, and single instancing that consume less energy
- Storage and server virtualization technologies that reduce the infrastructure's consumption of energy and cooling resources
- Tools and services to help you design and implement deployment strategies that achieve significant savings

## Audience

The intended audience for this white paper is EMC customers and field personnel who need to understand the impact of power consumption and heat dissipation in data centers.

## Current and future energy consumption considerations

Over the past few years, rising energy prices and drastic increases in density and power requirements for server and storage solutions have driven IT managers to consider strategies for reducing power consumption in their data centers. According to the Department of Energy, the average cost of power in the United States increased 25 percent since 2000 to 8.74 cents per kilowatt hour. As shown in Figure 1, energy costs vary significantly within the United States, and are highest for regions along the densely populated East and West coasts, Alaska, and Hawaii. Other parts of the world, including Europe and Asia, have even greater challenges with the high cost of energy, and limited power availability and floor space.

---

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

As of March 31, 2007

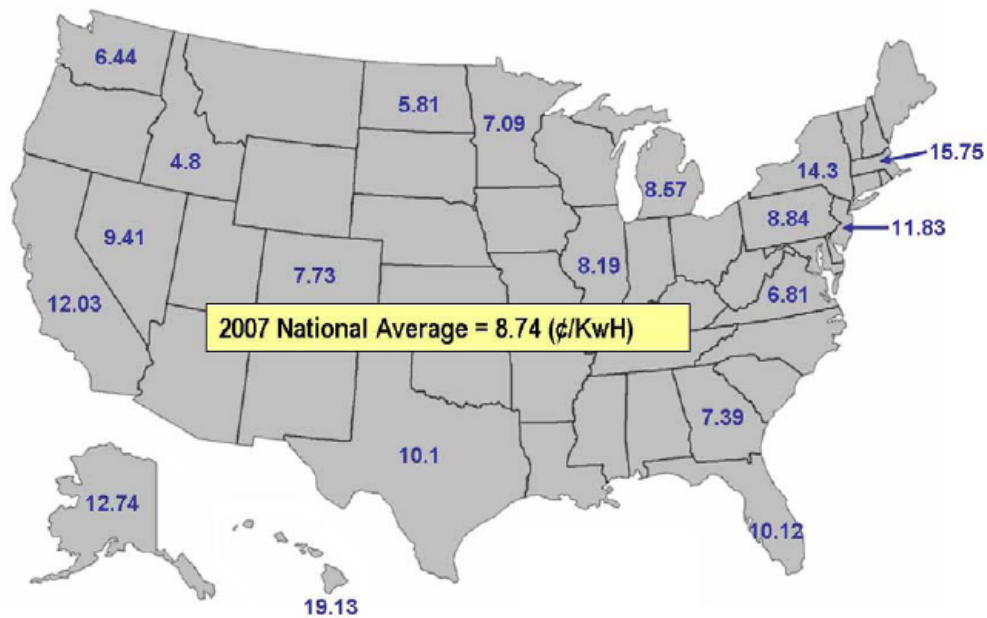


Figure 1. Average price of electricity<sup>1</sup>

The high cost of energy is only one of the factors contributing to the rising costs of power and cooling in data centers. The amount of power and cooling needed for each square foot of data center floor space is increasing dramatically as racks are filled with high compute-density equipment. While IT managers historically have been delighted that processor speeds double every 12-18 months (in accordance with Moore's Law), the new concerns are the power and heat problems that future acceleration will bring.

There are a number of approaches that can be used to address power and cooling problems in the data center. This white paper focuses on what can be done in a CLARiiON storage system environment driven by innovative technology that provides power, cooling, and space efficiency.

## The Fundamentals

### Terminology

#### Btu

The **British thermal unit** (Btu) is a unit of energy used by power, steam generation, heating, and air conditioning industries. Although it is commonly used in these industries, in scientific circles it has been replaced by the SI unit of energy, the joule (J).

In North America, the term *Btu* is used to describe the heat value (energy content) of fuels, and 1 cubic foot of natural gas  $\approx$  1000 Btu (to within a few percent).

---

<sup>1</sup> Asaro, Tony, *Power, Cooling, Space Efficient Storage*, ESG Enterprise Strategy Group, 2007

---

## Watts

The **watt** (or **W**) is the SI derived unit of power, equal to 1 joule per second. A human climbing a flight of stairs is doing work at the rate of about 200 watts. A first-class athlete can work at up to approximately 500 watts for 30 minutes. An automobile engine produces mechanical energy at a rate of 25,000 watts (approximately 30 horsepower) while cruising. A typical household incandescent light bulb uses electrical energy at a rate of 40 to 100 watts.

## Volt Amperes (VA)

A **volt Ampere**, in electrical terms, means the amount of apparent power in an alternating current circuit equal to a current of 1 ampere at an electromotive force of 1 volt. It is equivalent to watts for non-reactive circuits.

- 10 kVA = 10,000 watts capability (where the SI prefix k equals kilo)
- 10 MVA = 10,000,000 watts capability (where M equals mega)

While the volt Ampere and the watt are dimensionally equivalent, one may find products rated in both VAs and watts with different numbers. This is common practice on uninterruptible power supplies (UPSs). The VA rating is the apparent power that a UPS is capable of producing, while the watt rating is the real power (or true power) it is capable of producing, as opposed to reactive power. Reactive power arises due to the effects of capacitance and inductance of components in the load that is powered by the AC circuit. In a purely resistive load (incandescent lights, for example), the apparent power is equal to the true power and the amount of VAs and watts used would be equivalent. However, in more complex loads, such as computers (which UPSs are intended to power) the apparent power used (VAs) will be larger than the true power used (watts). The ratio of these two quantities is called the power factor.

## Power/cooling projections

Power and cooling have increasingly become a major consideration when designing new storage solutions. Due to the rapid increase of data storage requirements, as well as server and storage systems becoming more powerful, power plants have started facing major challenges in meeting overall power requirements.

The growth of a new generation of more powerful servers continues to draw more power and generate more heat. New classes of servers, such as blade servers, are designed to provide higher processing speed and capabilities in the same footprint as their predecessors. The net result is that as this trend continues, data centers are struggling to meet required power and cooling.

Energy use is on the rise, along with the amount of cooling and airflow required to maintain an acceptable data center temperature. As seen in Figure 2, the cost of powering and cooling data centers is trending to match that of acquiring new servers. Indeed, servers are the major contributor to increasing power and cooling demand.

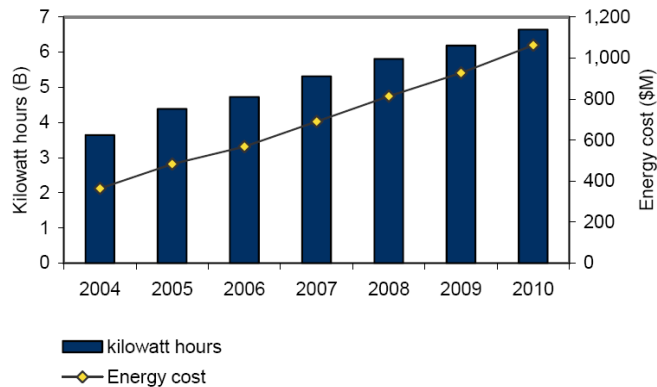
In another study of only storage-based solutions for the data center, IDC found that energy usage and costs are following a similar trend<sup>2</sup>. Clearly, both server and storage systems are driving more energy use and require more cooling to counter the generated heat.

In the following sections, we consider how advanced, energy-saving features in CLARiiON and EMC products help reduce storage-related power consumption. These sections provide a conceptual description of the technologies for the features and the benefits they provide. They also provide use cases for proof points and increased clarity.

---

<sup>2</sup>[Borovick, Gray, Reinsel, and Roch, \*Storage Power Management Is Heating Up: Challenges and Solutions\*, IDC, July 2006, IDC #202573](#)

Worldwide Power Consumption of New Disk Drive Shipments into Enterprise Solutions, 2004–2010



**Figure 2. Worldwide power consumption of new disk drive shipments into enterprise solutions, 2004 – 2010<sup>3</sup>**

## CLARiiON solutions

CLARiiON’s advanced storage technologies are designed to help improve overall system performance and optimize storage energy use. In the following sections, we discuss some of the foundations of energy usage and various CLARiiON best practices for improving energy efficiency.

### ***Multi-tiering***

CLARiiON has the ability to make effective use of multiple tiers of disk drive classes within the same system. CLARiiON’s UltraScale™ architecture allows you to use a wide variety of disk drive technologies within the same array. For example, the UltraScale disk-array enclosure (DAE) allows concurrent use of SATA, Fibre Channel, and low-cost Fibre Channel drives with interface speeds of 2 Gb/s and/or 4 Gb/s. Furthermore, the flexibility and configuration options provided by UltraScale technology and virtual LUNs (discussed later) enable you to reap the benefits of these technologies simultaneously.

### **Disk drive classes**

The CLARiiON 15k rpm 73 GB, 146 GB, and 300 GB 4 Gb/s hard disk drives are designed to increase performance through improvements in disk operations such as rotational latency and seek rates—the factors that most directly affect access times. The 146 GB and 300 GB 15k rpm hard disk drives deliver 3.5 ms average read seek times and 2.0 ms average rotational latency times—both the fastest available in today’s hard-disk drive market.

The CLARiiON 146 GB and 300 GB 10k rpm 2 Gb/s disk drives increase capacity per storage system through improvements in disk features such as linear/area densities and tracks per inch—the factors that most directly affect the overall disk-drive capacity per spindle. The 300 GB disk drive delivers two times the capacity per spindle than previous-generation 146 GB drives, with equivalent or better performance on direct spindle-to-spindle performance.

<sup>3</sup> [Borovick, Gray, Reinsel, and Roch, \*Storage Power Management Is Heating Up: Challenges and Solutions\*, IDC, July 2006, IDC #202573](#)

---

## High-capacity 750 GB SATA drives

The CLARiiON 750 GB 7.2k rpm SATA II disk drives dramatically increase capacity per storage system. The 750 GB disk drive delivers a 33 percent increase in capacity per spindle than past-generation 500 GB drives, with equivalent performance in large block, sequential I/O environments. Due to their increased density, when compared to the 500 GB drives, the 750 GB drives use the same power yet deliver higher capacity, resulting in a 33 percent power improvement.

It is common practice in the CLARiiON storage systems to mix drive types according to end-user requirements. This is where 750 GB drives should be factored into the data capacity/performance mix. With the current capacity points of the CLARiiON hard disk drives at 73 GB (10k and 15k rpm), 146 GB (10k and 15k rpm), 300 GB (10k and 15k rpm,) and 750 GB (7200 rpm), you can apply these different capacity and performance drives to suit the various applications within your operating environment.

Within the CLARiiON storage array, being able to move data dynamically from 15k rpm disk drives to 7.2k rpm disk drives results in energy savings as these storage tiers have a different power profile, as shown in Table 1.

**Table 1. Power usage between active and idle drives**

Rotation speed	Active power per disk drive
15k rpm	19.36 watts
7.2k rpm	13 watts

You should consider several factors when planning performance, capacity, and energy requirements for any given environment. The performance, pricing, power consumption, and capacity requirements should be completely understood before selecting the proper disk drive to use in your configurations. Select the disk drives that meet the storage capacity, price point, power consumption, and performance requirements for your application. For more information, please see the [EMC CLARiiON Best Practices for Fibre Channel Storage: CLARiiON Release 26 Firmware Update](#) white paper on [EMC Powerlink](#).

## Drive types and energy consumption

Table 2 lists typical power consumption and related energy costs for different disk drive technologies in a standard CLARiiON environment. Please note that these figures compare only the drive technology choice and not their relative energy per capacity.

**Table 2. Annual energy costs for drive technologies**

Drive type	Number of drives	Line current	Power consumption	Annual energy costs	Heat dissipation
7.2k rpm SATA II	15 drive	1.4 A	.29kVA	\$1,111.00	900 Btu/hr
10k rpm	15 drive	1.5 A	.31kVA	\$1,215.00	1000 Btu/hr
15k rpm	15 drive	1.6 A	.32kVA	\$1,256.00	1100 Btu/hr

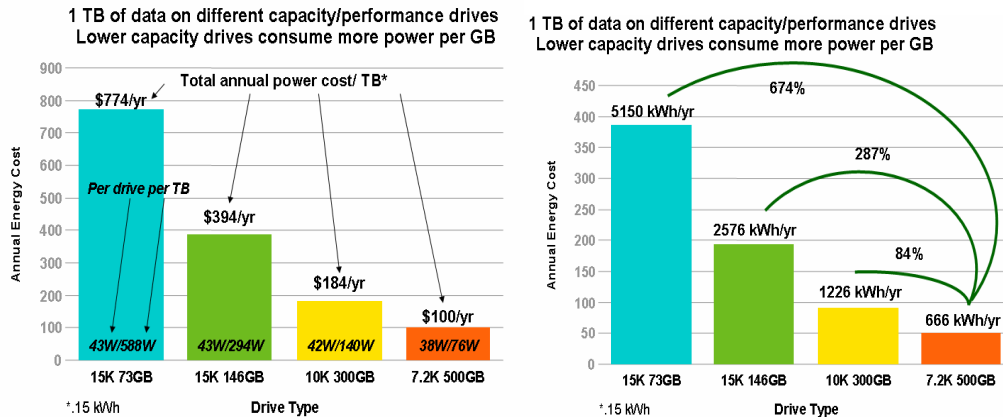
Enclosure power consumption calculations based on an average local utility rate of 0.1537 \$ /kW-hr

Examples of various disk drive usage are listed below:

- Tier 1a (log files): 73 GB 15k
- Tier 1b (hot data, tables): 146/300 GB 15k
- Tier 2a (warm data, tables): 300 GB 10k

- Tier 2b (clones): 300 GB 10k
- Tier 3 (B2D, archive): 500/750 GB SATA

Using one type of disk for all types of loads on storage systems is not a recommended practice, although some vendors may suggest doing this. Each storage tier has different information and power requirements throughout its lifecycle; thus, the different tiers of storage/disk drive classes should be a major consideration when you design the configuration and layout for storing and managing data in your system.



**Figure 3. Comparison of power consumption for different drive types**

## RAID types

Another major consideration for designing storage solutions is choosing the right RAID type, as well as the right storage tier, for your application. CLARiiON is designed to handle more than one RAID type in the same array and storage group. This capability allows much more flexibility to designate the right storage environment for your application and power requirements. The following sections highlight various aspects of performance and power for different RAID types.

### RAID 0: Speed (widely used, but not recommended)

RAID level 0 is disk striping only, which interleaves data across multiple disks for performance. Widely used for test and development when costs are a major consideration, RAID 0 has no safeguards against failure. Since there is no mirroring and replication of data on another drive, RAID 0 uses less power when compared to other RAID types for the same capacity.

### RAID 1: Disk mirroring (widely used)

RAID 1 uses disk mirroring, which provides 100 percent duplication of data. This RAID type offers the highest reliability but doubles storage and power costs.

RAID 1 is widely used in business applications that require maximum high availability. RAID 1 is a good choice when a dedicated file system is required, and the storage needs are small enough to make a RAID 1/0 LUN too costly. Often, RDBMS transaction logs are implemented on a RAID 1 LUN.

However, RAID 1 is not striped, which means you will not gain the performance advantages of a striped RAID level. Furthermore, RAID 1 is not suitable for writing very large I/O sizes. If your application/OS/HBA are exchanging requests larger than 128 KB in size, you will suffer some latency at the disk level. (A striped RAID level breaks these larger requests up as it stripes across the disks.)

---

## RAID 3: Speed and parity protection

In RAID 3, data is striped across three or more drives. This RAID type is used to achieve the highest data transfer rates, because all drives operate in parallel.

RAID 3 is a specialty solution. There are only two RAID group sizes valid for CLARiiON RAID 3 — five-disk and nine-disk. The target problem for RAID 3 is large, sequential data. For best performance, the stripes should match the I/O size or be an even multiple or fraction of the I/O size.

## RAID 5: Speed and parity protection (widely used)

In RAID 5, data is striped across three or more drives for performance, and parity is written for data protection. The parity bits from two drives are stored on a third drive and are interspersed with user data. RAID 5 is widely used in data storage systems and typically is the *most power efficient* when factoring in drive count versus capacity.

RAID 5 should be considered for data mining, medium-performance media serving, and RDBMS implementations in which the DBA is effectively using read-ahead and write-behind.

## RAID 6: Double parity protection

RAID 6 has the highest reliability because it can recover from a failure of two disks. Similar to RAID 5, RAID 6 performs two different parity computations or the same computation on overlapping subsets of the data. RAID 6 also provides a relatively similar capacity and good power efficiency as RAID 5.

RAID 6 is generally described as data striped across a number of drives within a RAID group, with two independent parity fields maintained for redundancy. Thus, RAID 6 is also referred to as “N+2” or “double parity.” Key points regarding the general concept of RAID 6 are:

- Data is still available even if two drives within the RAID group cannot be read.
- RAID 6 can tolerate a hard media error during a single drive rebuild.
- RAID 6 can tolerate a second drive failure during a single drive rebuild.
- RAID 6 can tolerate hard media errors from two drives on the same logical block address.
- RAID 6 may have some performance impact compared to either RAID 5 or RAID 1/0, though not necessarily significant.

## RAID 1/0: Disk mirroring

RAID 1/0 is RAID 1 + 0. The drives are mirrored for fault tolerance and the mirrors are striped for performance (RAID 0). Similar to RAID 1, RAID 1/0 almost doubles storage and power consumption when compared to a RAID 3 or 5 configuration.

RAID 1/0 outperforms RAID 5 in workloads that use very small, random, and write-intensive I/O. In these cases, the need to write a parity segment hurts RAID 5. Examples of random small I/O workloads are:

- Microsoft Exchange and similar products
- Real-time data/brokerage records
- RDBMS data tables containing small records that are updated frequently (account balances)

## **Power consumption**

Proper planning is crucial when designing a storage system for your environment. You need to understand the requirements for each application in your environment, including:

- Capacity
- Performance
- Reliability

- Price point
- Power consumption

Analyzing these requirements will help you select the best RAID group(s) for your system.

Table 3 compares the power consumption and related energy costs for different RAID groups in a standard CLARiiON environment.

**Table 3. RAID type vs. power consumption**

RAID type	Number of drives	Line current	Power consumption	Annual energy costs	Heat dissipation
RAID 0	4 drive	0.7A	.14kVA	\$532.00	400 Btu/hr
RAID 3	5 drive	0.7 A	.15kVA	\$597.00	500 Btu/hr
RAID 5	5 drive	0.7 A	.15kVA	\$597.00	500 Btu/hr
RAID 6	6 drive	0.8 A	.17kVA	\$663.00	600 Btu/hr
RAID 1/0	8 drive	1.0 A	.20kVA	\$795.00	700 Btu/hr
Disk drive usage is 300 GB, 15k rpm drives, RAID group capacity is 1073.48 GB					

As shown in Table 3 and Table 4 , various configurations or tiers within the CLARiiON storage can have a significant impact on the power and cooling characteristics of the CLARiiON storage system. Depending on the drive types and RAID configurations of the storage system, aging data, for example, can be moved via virtual LUN (covered later) to a Tier 3 based CLARiiON technology system (RAID 5 with SATA II disk drives) and effectively reduce the power and cooling consumption of the storage system. As shown in Table 4, significant cost savings can be realized by utilizing multi-tiers of disk storage and migrating aging data or less performance-dependent data to SATA II disk drive technology. Also, note the overall increase in storage capacity per system when SATA II disk drives are used in Tier 3 configurations.

**Table 4. Annualized cost savings when comparing a 90-drive, multi-tiered CX3-80 system**

Typical 90 drive, CX3-80 system	300 GB, 15k rpm Fibre drives	750 GB, 7200 rpm SATA II drives
6 enclosures, 90 drives Installed	90 - 300 GB Fibre Channel	15 - 300 GB Fibre/75 - 750 GB SATA II
Line current	11.2 A	10.3 A
Power consumed	2.35 kVA	2.17 kVA
Btus consumed	7,700 Btu/hr	7,100 Btu/hr
Total system capacity	~ 18.29 TB	~ 41.27 TB
Annualized energy costs at 0.1537 \$/kW-hr	\$9,168.00	\$8,441.00

## **Storage consolidation and utilization**

One of the best ways to improve energy efficiency for data storage is to consolidate storage by moving your data from a DAS to a SAN environment. Recent user analysis<sup>4</sup> showed storage utilization rates average between 44 percent for DAS to 70 percent for SAN topologies. In addition to more efficient utilization of storage, all of your applications share storage that is centrally powered, which further reduces energy costs. Storage consolidation also provides single management, lower maintenance costs, and increased availability.

<sup>4</sup> [Barrett, Alex, Survey Says: Fibre Channel SANs have the best utilization rates, Storage magazine, May 2007](#)




Figure 4. Storage capacity and utilization<sup>5</sup>

## System utilization

Due to their superior architectural design, CLARiiON systems are capable of providing much higher utilization rates without sacrificing system performance. As Figure 5 shows, CLARiiON supports 74 percent more Microsoft Exchange users with a much higher capacity utilization than the FAS3050c.

Low system utilization rates translate into wasted power consumption. In most storage systems, utilization rates have a reverse relationship to overall system performance. When planning an overall storage solution, both the system utilization and required application performance must be considered.

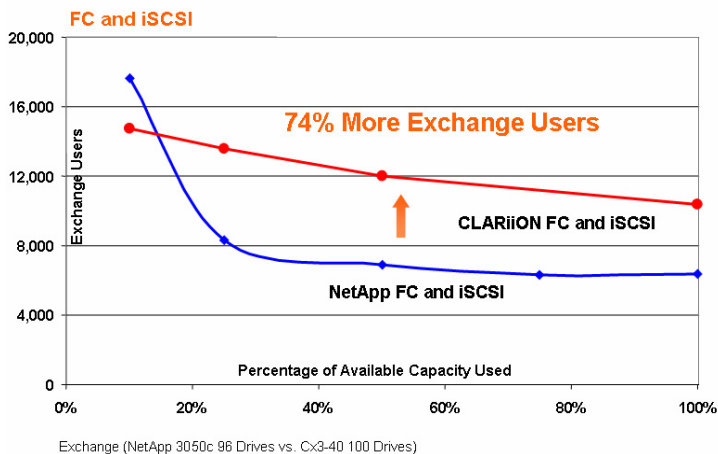


Figure 5. System utilization comparison

## MetaLUNs

MetaLUN technology allows users to expand the capacity of any given RAID group and individual LUN to increase storage volume and accommodate growing storage needs. This technology is designed to provide efficient provisioning of storage resources and accommodate real-time LUN expansion without impacting application performance and availability.

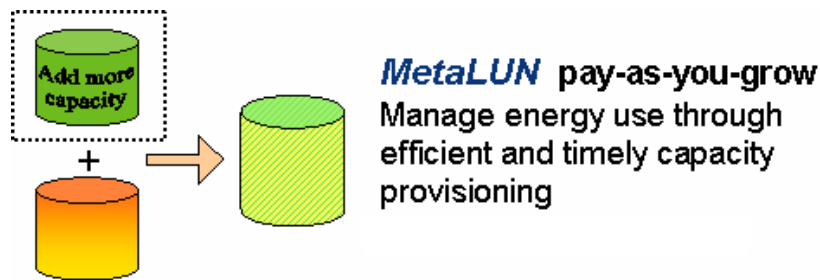
<sup>5</sup> [Alex Barrett, Survey Says: Fibre Channel SANs have the best utilization rates, Storage magazine, May 2007](#)

With metaLUN technology, you can expand a LUN whenever you need more storage capacity. For this reason, when you install your system you do not need to create extra LUN capacity to support future growth. As a result, you have more control over the number of drives in your system. MetaLUNs provide a substantial savings in terms of not having to power up drives that may not be needed until a much later date.

The effective use of metaLUN technology enables better provisioning of data growth on the array in a way that power consumption can be managed in a *pay-as-you-grow* fashion. Increasing a given LUN with the addition of more drives allows the customer or end user to grow their storage as required by their applications.

Additionally, metaLUN technology improves the overall system utilization by allowing data to be expanded to disks that are underutilized. This is a unique software functionality that meets the growing demand for storage while improving the system use rate. As stated earlier, proper system utilization is the best way to manage overall power utilization.

There are two ways to implement LUNs: concatenation and striping. These two options are designed for different application and capacity criteria. For more information on metaLUN technology implementation and planning, please see the white paper [EMC CLARiiON MetaLUNs: Concepts, Operations and Management](#) on [EMC.com](#) and [EMC Powerlink](#). MetaLUN expansion technology is available on CLARiiON release 16 or later systems, and is built into the FLARE<sup>®</sup> operating environment. There is no need for additional software licenses or any installation steps to take advantage of this technology.



**Figure 6. Benefits of metaLUN expansion technology**

**Example:**

Without the use of metaLUN technology, the user has to provision all the capacity initially whether it's needed or not. As shown in the example in Table 5, the total power used for this configuration (32 TB) is considerably higher than that of a 16 TB LUN made up of the same disk drives.

With metaLUN technology, the user is able to expand their storage requirements as required, thereby saving the operating cost of utilizing a 32 TB LUN until it becomes necessary. Again, as shown in the example in Table 5, the total power used for this provisioning technology is considerably higher than that of a 16 TB LUN made up of the same disk drives.

**Table 5. 16 TB LUN compared with a 32 TB LUN utilizing the same type of drives**

Drive type	16 drive RAID 5	Total capacity	Line current	Power consumption	Annual energy costs	Heat dissipation
300 GB - Fibre Channel	64 drives	16.102 TB	6.9 A	1.43kVA	\$5,557.00	4600 Btu/hr
300 GB - Fibre Channel	128 drives	32.204 TB	13.4 A	2.80kVA	\$10,845.00	9100 Btu/hr
Annualized energy costs at 0.1537 \$/kW-hr						

For more information on the metaLUN technology, see the white paper [EMC CLARiiON MetaLUNs: Concepts, Operations, and Management](#).

## Virtual LUN technology

EMC's unique virtual LUN technology is a built-in feature on CLARiiON arrays that allows users to seamlessly migrate data between LUNs in the array without requiring any application downtime. Virtual LUNs allow a user to migrate data to other LUNs (with better performance or other characteristics) within the array, providing a greater level of control over the system. The characteristics of a LUN or drive type that can be changed are:

- **RAID geometry**
  - RAID type
  - Number of disk spindles
  - Alignment offset
  - Stripe element size
  - LUN size
- **Drive type**
  - Fibre Channel → SATA
  - SATA → Fibre Channel

With the advent of additional drive capacities and technologies, virtual LUN is an easy to use, seamless migration tool that transparently migrates data between drives with different energy profiles. This advanced technology does not take down the applications that access the LUN, and is able to move data across different LUNs of different types.

This advanced technology is completely transparent to the hosts. Virtual LUNs are available on all CLARiiON release 16 or later systems, and are built into the FLARE operating environment. There is no need for additional software licenses or any installation steps to take advantage of virtual LUNs



Figure 7. Benefits of virtual LUN technology

### Example of cost savings with virtual LUN migration

As shown in Table 6, when using a 16-drive RAID 5 storage unit residing on 15k rpm, 300 GB FC drives, annual energy costs are about \$1,590 per year.

Also shown in Table 6, when migrating aging or archival data to a seven-drive RAID 5 storage unit on 7.2k rpm, 750 GB drives, total annual energy costs drop to a total of \$661, resulting in considerable energy savings over the life of the storage system.

Table 6. 15k rpm disk drive RAID groups vs. 7.2k rpm disk drive RAID groups

Drive type	RAID 5	Total capacity	Line current	Power consumption	Annual energy costs	Heat dissipation
750 GB - SATA II	7 drives	4.127 TB	0.8 A	.17kVA	\$661.00	600 Btu/hr
300 GB - Fibre Channel	16 drives	4.025 TB	2.0 A	.41kVA	\$1,590.00	1300 Btu/hr
Annualized energy costs at 0.1537 \$/kW-hr						

For more information on virtual LUN technology implementation and planning, please see the white paper [EMC Virtual LUN Technology - A Detailed Review](#).

## Navisphere Quality of Service Manager

Navisphere Quality of Service Manager (NQM) measures, monitors, and controls application performance on the CLARiiON storage system. The monitoring feature is an excellent first step to improving the performance of high-priority applications, because it gives you a more logical view of system performance, both for the entire storage system and for specific applications. This can be a powerful method of evaluating the storage system to determine the current service levels and to provide guidance on what service levels are possible, given the specific environment.

NQM controls application performance by giving you the ability to set performance targets for high-priority applications or performance limits for low priority applications. These performance targets and limits are one of three key performance characteristics – response time, bandwidth, and throughput. For example, an email application running on a CLARiiON storage system can be given a specific response time performance target to achieve (on the storage system itself). Similarly, a non-critical application can have its bandwidth limited to a certain level so that the storage system has resources available for more important requests.

In addition, NQM has a built-in scheduling feature allowing these performance targets to be adjusted dynamically based on business needs. The scheduler allows users to give priority to certain applications during specified “windows.” For example, a database application can be given a performance target during the day, during its peak hours of operation. Overnight, it is critical that this important information be backed up in a timely fashion, so NQM creates another performance target for the backup application during that time period. Both of these can be scheduled to run automatically.

## Prioritizing applications during peak energy hours

Delivery of energy to industrial and commercial environments costs more during peak hours. Managing application access and reducing load on the storage system energy consumption can reduce energy use. As shown in Table 7, active disks consume more power than inactive disks.

CLARiiON NQM’s ability to schedule highly disk-active tasks, such as backup to disk, at off-peak hours results in significant overall energy savings. Furthermore, NQM directly links policy associated with the application to the individual LUN. By enforcing explicit policy criteria for applications, NQM pushes off the active disk use to night time to save energy costs. The example in Table 7 is a practical illustration of using NQM scheduling technology to better utilize the storage array, and the cost benefits of managing active disk applications at off-peak hours.

**Table 7. Benefits of NQM scheduling**

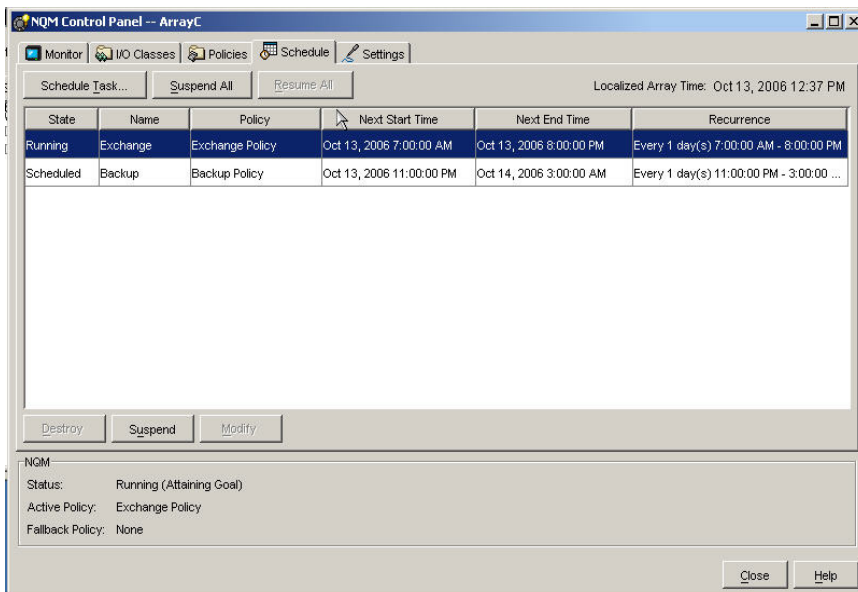
	Energy use	Yearly cost	Annualized savings
<b>Without NQM</b> (applications running all the time)	4 TB of DSS application active 24 hours 4 TB of BU application active 24 hours	<b>\$22,222.46</b>	
<b>NQM scheduling active</b> (application split between peak and off-peak)	4 TB of DSS application active and 4 TB of BU application idle 12 hours peak 4 TB of BU application active and 4TB application idle 12 hours off-peak	<b>\$20,738.14</b>	<b>\$ 1,484.32</b>

Users are also constantly being asked to adjust to workload changes in their environment and adapt to multiple service level requirements for different applications during different times. This can be a challenge to maintain and the man-hours devoted to the task can be costly. NQM's scheduling functionality allows you to create different service level goals throughout the day, so you can set performance targets for certain applications during their peak usage.

**Table 8. Example of peak vs. off peak rates for commercial delivery service<sup>6</sup>**

Customer Charge:	\$70.72/month
Distribution Demand Charge:	\$3.80/kW
Distribution Energy Charge:	
<i>Peak Hours</i>	<i>1.249¢/kWh</i>
<i>Off-Peak Hours</i>	<i>0.017¢/kWh</i>
Transmission Charge:	1.032¢/kWh
Transition Demand Charge:	\$0.75/kW
Transition Energy Charge:	0.140¢/kWh
Demand Side Management Charge:	0.250¢/kWh
Renewables Charge:	0.050¢/kWh

Let us consider the following example: A company has an Exchange server and a backup application residing on the same CLARiiON system. They want to ensure that the Exchange environment meets its throughput target during the day (7 A.M. - 8 P.M.). At the same time, they want the backup application to complete its nightly job during the current backup window (11 P.M. - 3 A.M.). To accomplish this, the customer creates two separate policies, one for Exchange and one for the backup application, and schedules each policy to run at the appropriate time, as shown in Figure 8. This methodology takes advantage of significantly lower energy costs during off-peak hours as shown in Table 8.

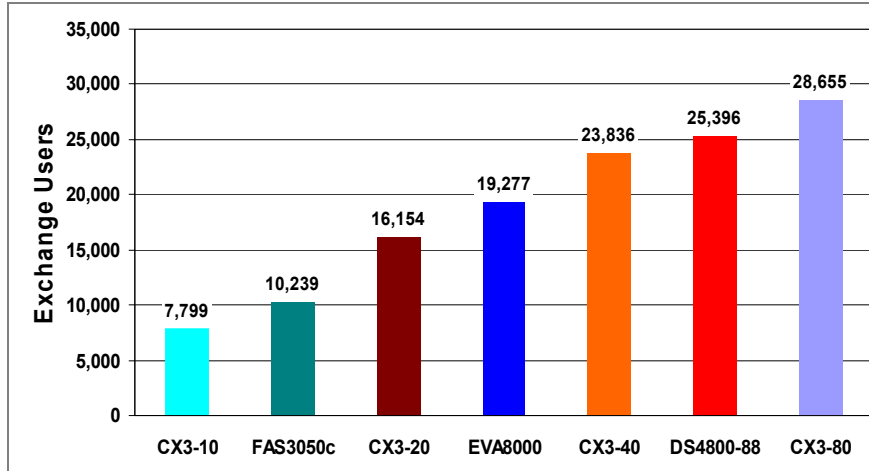


**Figure 8. NQM scheduling policy management**

<sup>6</sup> US National Grid, commercial, and industrial customer rates for delivery of services

## Performance and power

The CLARiiON CX3 UltraScale platform architecture delivers industry-leading levels of scalability and performance. CX3 uses a highly efficient storage processor design and higher-speed interconnect technology, creating the most powerful storage system in its class, as shown in Figure 9.



**Figure 9. CX3 series versus competition: Number of supported Exchange users**

Due to its superior performance and scalability, CLARiiON CX3 series is capable of handling a larger workload in a given timeframe. This overall system efficiency directly translates to better utilization of available disks. In general, compared to other storage systems, CLARiiON can use fewer disks to meet application requirements. Using fewer disks has a direct correlation to the energy usage of the array. As outlined earlier, most of the power draw in midrange storage systems is for spinning disks. CLARiiON eliminates the need for larger drive count to achieve application performance requirements. For application-specific comparisons of CLARiiON CX3 to competitive systems, please see the white paper [CLARiiON CX3 Application Performance: A Comparison of Midrange Storage Platforms](#).

**Table 9. Application performance results for the CX3 series and competitive storage systems<sup>7</sup>**

Application workload	CX3-10c 60 disks 2 GB cache	FAS3050c 262 disks 8 GB cache	CX3-20 120 disks 4 Gb cache	EVA8000 240 disks 8 GB cache	CX3-40 240 disks 8 GB cache	DS4800-88 224 disks 16 GB cache	CX3-80 480 disks 16 GB cache
OLTP (transactions/min)	19,148	12,396	34,097	44,810	59,630	63,559	72,372
Exchange (# users)	7,799	10,239	16,154	19,277	23,836	25,396	28,655
Backup-to-disk (MB/s)	393	229	471	480	881	985	1,121

### Application use cases

A key benefit of a higher performing storage system is its ability to consolidate multiple storage tiers in one place. CLARiiON is designed to handle multiple workloads and application requirements. This makes it very easy to match energy consumption to the specific application and storage tier. In doing so, a substantial power savings is achieved with CLARiiON storage systems.

<sup>7</sup> [CLARiiON CX3 Application Performance: A Comparison of Midrange Storage Platforms](#) white paper

---

In the following example, CLARiiON uses a combination of consolidation, drive technologies, and data migration tools to migrate from a less efficient system to a CX3-based array. As shown in Figure 10, by making use of key features in CLARiiON, the power usage per hour is reduced almost by a factor of two.




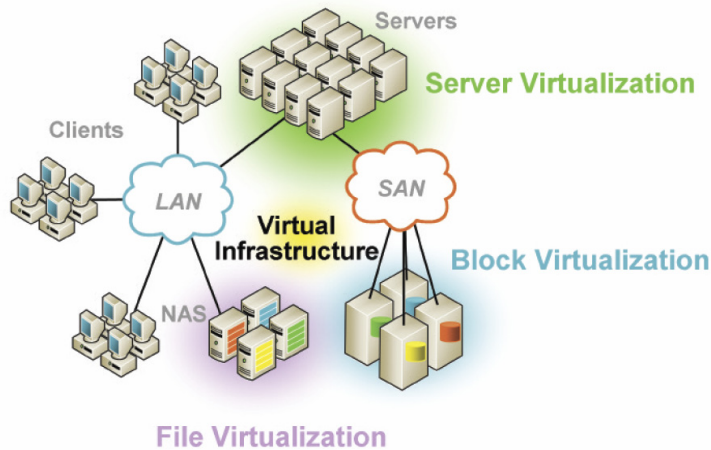
Figure 10. Energy costs before and after consolidating tiers

## EMC solutions

### *EMC virtualization technologies*

EMC virtualization solutions help optimize your organization's computing assets (including applications, databases, servers, and storage) into a simple, cost effective, flexible-tiered architecture that saves energy and aligns with your business requirements. Solutions include:

- **Server virtualization:** EMC leverages VMware, the leader in virtual-infrastructure technology, to transform a customer's server infrastructure, enabling applications and databases to be supported and maintained independent of the server platform. With VMware, EMC addressed the critical storage consolidation and data migration issues that may result as part of this process.
- **File virtualization:** EMC leverages the Rainfinity® Global File Virtualization software platform to optimize IP-based storage, NAS, and file server. It builds upon industry-standard namespace services to provide location transparency and full read/write access to file-based data that is being relocated dynamically, which dramatically improves utilization and speeds migrations.



**Figure 11. File virtualization**

- **Block storage (SAN) virtualization:** EMC leverages Invista<sup>®</sup>, a solution that enables the virtualization of storage in networked storage environments, thus reducing downtime associated with data movement across storage tiers.

EMC virtualization solutions provide increased energy efficiency in your data centers by:

- Lowering data center and infrastructure costs to reduce server and storage power consumption
- Optimizing capacity and performance across the infrastructure by increasing and balancing utilization
- Retiring old equipment, eliminating the need for additional equipment, and reclaiming data center floor space

### VMware server virtualization (server consolidation)

IT organizations can reduce hardware and operating costs by as much as 50 percent by implementing a VMware server consolidation solution. In addition to the significant reductions in server hardware expenses made possible by improving server utilization, they achieve savings in costs for administration, power, cooling, and data center infrastructure. These savings have provided customers a rapid return on their VMware software investment. This software:

- **Boosts utilization.** VMware infrastructure aggregates x86 server resources into pools that can reliably support CPU utilization exceeding 80 percent with the continuous virtual machine load balancing provided by VMware DRS.
- **Significantly improves manageability.** Server consolidation and containment with VMware software reduces data center complexity by reducing the number of servers that IT organizations need to manage. With VMware management tools, organizations can also simplify and centralize the monitoring and management of large virtual infrastructure environments.
- **Simplifies server provisioning.** IT departments can reduce the time it takes to provision new servers by 50-70 percent. Virtual machines are as easy to copy as software files and are hardware independent, so deploying new workloads takes minutes instead of days.
- **Increases IT efficiency.** A VMware solution streamlines and eliminates common administrative tasks such as enabling IT to manage a growing server environment with existing resources.
- **Increases responsiveness.** By streamlining common tasks and continuously balancing workloads, VMware software makes it possible for IT to respond more rapidly to requests for new servers and maintain service levels as resource demands fluctuate.

---

Following is a link to a report generation tool that highlights cost savings you can achieve using VMware products: <http://www.vmware.com/go/calculator>.

## Storage virtualization, EMC Rainfinity, and Invista

Storage virtualization creates an abstraction layer between physical disk drives and the logical file or block data access form. By separating the physical location of data from the logical access, storage can be dynamically optimized based on specific application requirements. Storage virtualization allows single access to data from various underlying storage systems, thus improving management. Balancing the usage across various storage platforms enables improved consolidation of systems and ultimately reduces the number of deployed arrays. Virtualization technologies reduce the need for additional equipment, enable seamless load balancing, and improve storage tiering.

### ***Avamar (data de-duplication)***

EMC Avamar<sup>®</sup> software solves traditional backup and storage challenges by reducing the size of data at the source, before it is transferred across the network. The global data de-duplication technology stores just a single copy of redundant sub-file data segments across sites and servers. This effectively reduces the required storage and backup volume by up to 300 times, enabling companies to utilize fewer resources for backup and recovery. Avamar's Redundant Array of Independent Nodes (RAIN) technology provides high availability and its grid architecture delivers online scalability.

Avamar technology reduces the increased need for storage by removing redundant data. Utilizing this technology is a proven method to reduce overall power and cooling requirements for data centers.

### ***EMC Power Calculator***

Use the EMC Power Calculator to estimate energy consumption and cooling requirements for specific storage configurations. This is a stand-alone tool that now combines calculators for:

- Symmetrix<sup>®</sup> DMX-3
- Symmetrix DMX-3 950
- Celerra<sup>®</sup> NS40
- Celerra NS80
- Celerra NSX
- CLARiiON CX3 (all configurations)
- EMC Disk Library

---

**Note:** This tool is available only to EMC staff and partners. Complete user instructions and access to the Power Calculator can be found on Powerlink. This tool replaces the separate Symmetrix and CLARiiON power calculators.

---

The power calculator lets EMC staff and partners provide customers with accurate energy and air conditioning consumption estimates for a specific configuration before or after installation. It includes energy consumed in typical, maximum, and recharge loads, as well as British thermal unit (Btu) output measurements. This is *far more precise* than the data sheet information that is typically used by EMC competitors.

### ***EMC data center energy efficiency services***

EMC Energy Efficiency Services takes a holistic approach to energy efficiency by conducting an analysis of the data center and IT infrastructure to develop the appropriate energy efficiency strategies. This service entails development of a comprehensive understanding of data center power and cooling based on customer requirements and needs. It identifies efficiency opportunities utilizing the areas covered in this paper, such as consolidation and virtualization of servers and storage.

---

This comprehensive consulting service assesses infrastructure energy efficiency and provides an optimized design to improve energy use and cooling/airflow.

Assessments look at two areas:

- The IT infrastructure:
  - Server, storage, and network energy consumption
  - Server and storage capacity/utilization
  - Server and storage configurations
- The data center:
  - Facilities' conditions and systems upgrade options
  - Data center tiering

The assessment services produce detailed recommendations and energy efficiency options, including the trade-offs and projected benefits of:

- Server/storage consolidation and virtualization
- Configuration alternatives and storage tiering options
- Upgrades to data center facilities and systems

EMC's comprehensive energy efficiency consulting services are based on substantial experience in developing and delivering a complete information infrastructure. In addition to providing tools and extensive knowledge in managing information services and systems, this EMC offering is based on explicit proof points and best practices. We highly recommended these services to support the growing demands on your data center.

## **The Green Grid**

As part of our commitment to improve data center power consumption and efficiency, EMC/VMware are board members of the Green Grid consortium. The Green Grid is a consortium of information technology companies and professionals seeking to lower the overall consumption of power in data centers around the globe. The organization is chartered to develop meaningful platform-neutral standards, measurement methods, processes, and new technologies to improve energy efficient performance of global data centers.

The specific efforts of the Green Grid are based on defining meaningful, user-centric models; metric, developing standards; and measurement methods to improve performance against the metrics and to promote the adoption of energy efficient standards and processes.

EMC continues to develop and market technologies designed to improve power efficiency and cooling for today's data centers.

## **Conclusion**

CLARiiON's advanced storage systems provide an extensive feature set designed for optimizing power consumption and cooling requirements. Configuration options that support energy-efficient disk technologies, combined with intelligent management options (such as virtual LUN and metaLUN storage optimization technologies), achieve both high performance and low energy consumption. Combined with other EMC technologies, such as VMware virtualization, infrastructure management, or Avamar data-duplication, CLARiiON delivers an outstanding choice for power and cost performance.

## **References**

- [Asaro, Tony, \*Power, Cooling, Space Efficient Storage\*, ESG Enterprise Strategy Group, 2007](#)
- [White Paper: \*Power Efficiency and Storage Arrays - Technology Concepts and Business Considerations\*](#)

- 
- [White Paper: EMC CLARiiON MetaLUNs, Concepts, Operations and Management](#)
  - [White Paper: EMC Virtual LUN Technology - A Detailed Review](#)
  - [US National Grid, commercial and industrial customer rates for delivery of services](#)
  - [Service Overview: EMC Energy Efficiency Services](#)
  - [White Paper: CLARiiON CX3 Application Performance: A Comparison of Midrange Storage Platforms](#)
  - [White Paper: EMC CLARiiON MetaLUNs - Concepts, Operations, and Management](#)
  - [The Green Grid Consortium: http://www.thegreengrid.org](http://www.thegreengrid.org)
  - [Borovick, Gray, Reinsel, and Roch, Storage Power Management Is Heating Up: Challenges and Solutions, IDC, July 2006, IDC #202573](#)
  - [Barret, Alex, Survey Says: Fibre Channel SANs have the best utilization rates, Storage Magazine May 2007](#)