ESG Enterprise Strategy Group

IT Knowledge • Business Results

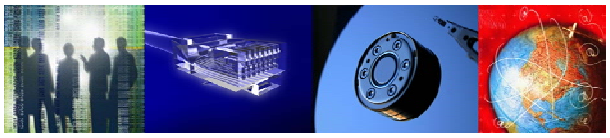

Power, Cooling, Space Efficient Storage

By Tony Asaro
Senior Analyst

July 2007

Table of Contents

Storage Power, Cooling, and Space Efficiency

Approximately a year ago the Director of IT for a company that has approximately 80,000 employees told ESG that his biggest issue was power, cooling and floor space. We conducted a series of interviews with different end users and depending on who you spoke to the answers were extreme. ESG has found that power, cooling and floor space issues are binary to end users: it is either a problem or it is not. However, we feel that over time the majority of end users will be faced with this issue. Regions with denser populations are more likely candidates for power and cooling to be a real issue. The availability of electricity is decreasing while demand is increasing, resulting in higher cost per Kilowatt.

From a technology perspective, increased server, storage and connectivity density provides significant 'power budget' challenges to data center managers. HVAC systems draw even more electricity as they work harder to remove increased amounts of excess heat. If the power and cooling systems in the data center are near capacity, expensive upgrades may become unavoidable.

Rising energy prices and drastic increases in the density and power requirements of server and storage solutions over the past few years have driven IT managers to consider strategies for reducing power consumption in the data center. According to the Department of Energy, the average cost of power in the United States has increased 25% since 2000 to 8.74 cents per kilowatt. As shown in Figure One, the cost of power varies significantly within the United States and is 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, limited power availability and floor space.

Figure One: 2007 US Average Commercial Retail Price of Power

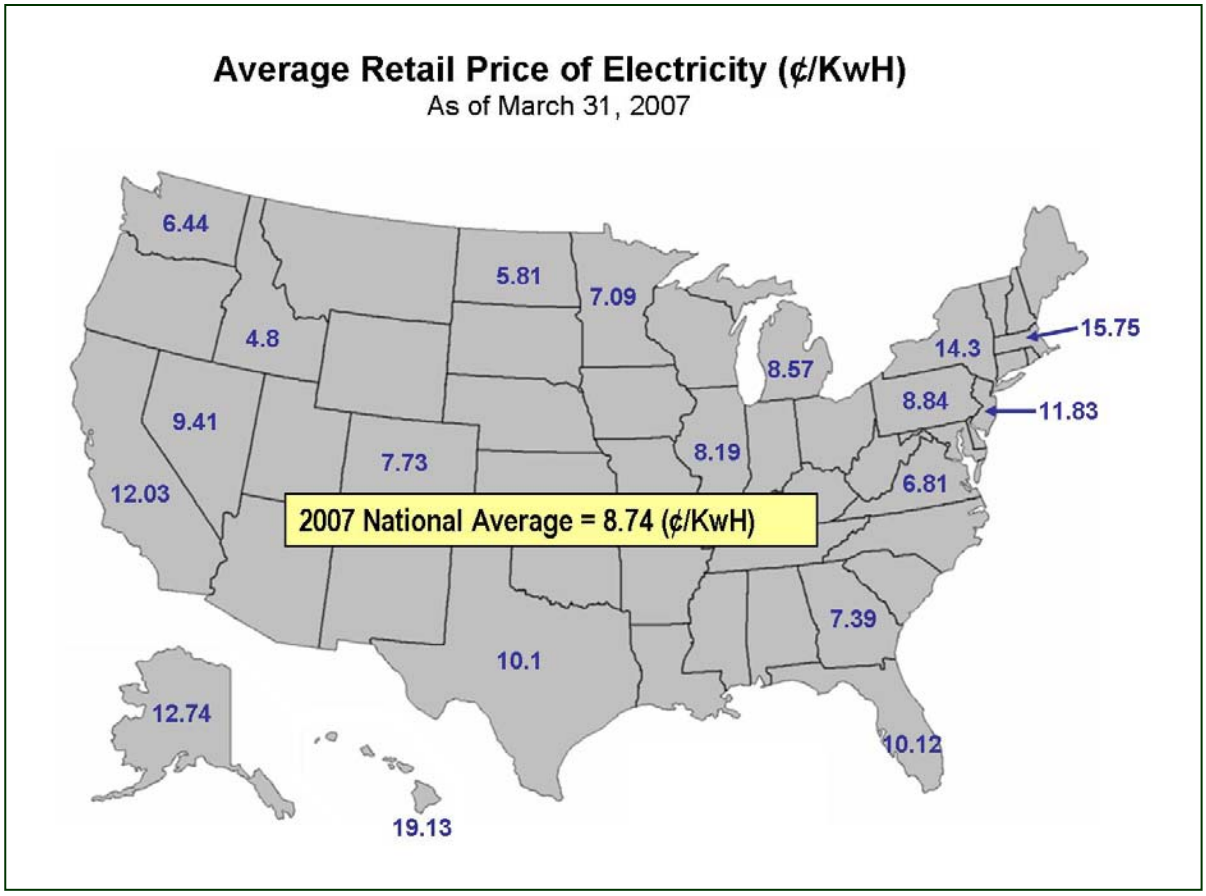

The high cost of energy is only one of the factors contributing to the overall rising costs of power and cooling in the data center. The amount of power and cooling needed for each square foot of data center floor space is increasing dramatically as racks are filled with high compute-density equipment.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

There are a number of approaches that can be used to address power and cooling problems in the data center. This report focuses on what can be done in storage driven by innovative technology that provide power, cooling and space efficiency (PCSE).

Storage Innovation Applied to PSCE

Network storage systems will increasingly consume power and cooling resources. There will be greater consolidation of server technology, which will lead to more storage consolidations. Additionally, storage growth will continue to escalate and the amount of capacity will be compounded by increasingly longer retention periods. The more capacity that is required the more power and cooling is consumed and the greater amount of physical floor space is required.

There are various technologies that can be used to reduce power, cooling and space consumption provided by advanced storage systems. These capabilities can have a significant impact and should be considered especially if power, cooling and floor space are major issues for your data center.

> "Data center electricity consumption increased 97% from 2000 to 2005 and now accounts for between 1-2% of all U.S. electricity consumption."
>
> –Source: *Estimating Total Power Consumption by Servers in the U.S. and the World*, Jonathan G. Koomey, Ph.D., February, 2007.

Fundamentally, the goal is to evaluate storage systems that enable you to minimize your capacity requirements and that allow you to utilize the systems resources to their fullest. ESG has identified following technologies and capabilities are ones to consider:

Thin Provisioning

Thin provisioning provides the ability to provision volumes while only consuming capacity based on the amount of data you have. Thin provisioning essentially eliminates allocated but unused capacity and stranded storage.

For example, let's say you have 200 GB of actual data and create a 2 TB volume for it to be stored on based on your growth projections. Al 2 TB of capacity is allocated to that volume even though only 200 GB of actual data is stored on it. That means that 1.8 TB of that volume consists of empty blocks that are consumed but unused. Thin provisioning solves this problem. You still provision a 2 TB volume but only 200 GB of it is consumed and the other 1.8 TB is available to be used for other applications.

There is a great deal of value provided by thin provisioning including:

- Since there is no stranded storage capacity, less storage is required over the life of the storage system
- Additional storage systems will not be required based on having stranded storage
- More applications/servers per storage system
- Greater levels of consolidation
- The time and resource required to perform storage provisioning tasks is reduced
- Thin provisioning-aware" local and remote volume copy services consume less capacity
- Reduce storage system software licenses based on capacity
- Power and cooling cost and consumption are reduced since end users can use less capacity

In the context of PCSE thin provisioning enables you to use less capacity, i.e. disk drives, to address the same or more application and business requirements. ESG has found that end users have between 30% and 50% allocated but unused capacity. That is a significant amount of waste that is consuming power, requires cooling and takes up space. Thin provisioning eliminates this stranded storage and has an immediate impact to the PCSE problem.

Dynamic Volume Management

Dynamic Volume Management (DVM) is different from thin provisioning but is often confused with it. DVM and thin provisioning can actually be complementary technologies. DVM capability is provided by the host operating system and the storage system must be able to support it. DVM is the ability to grow the volume size while online and without performing a re-boot of the system. The OS will recognize that the volume has increased.

DVM can also support a reclamation capability. Let's say you moved or deleted data off of a volume or created a big volume and after a few months you realize that the application just isn't going to need that much capacity. You can shrink the volume and put it back into the capacity pool. The ability to reclaim is just as important as expansion if not more so.

DVM has a similar value proposition to thin provisioning. You can more smartly allocate storage capacity and minimize waste.

Snapshots

One of the most popular and useful features provided by storage systems is snapshot technology. A snapshot is a logical copy of primary data stored on the storage system. A snapshot essentially creates a frozen image of the primary data at a specific point in time. The snapshot data is read-only and cannot be edited, deleted or corrupted.

> "A 50,000 square-foot facility might have cost $20 million a few years ago, but building for the projected future heat load could cost as much as $250 million today."
>
> –Source: *Power and Cooling in the Data Center* – AMD White Paper 2005

Snapshots can typically be created almost instantly. This enables systems administrators to rapidly protect data. Additionally, making snapshot copies does not consume storage system resources. This is important since a storage system's primary function is to service read and write requests from users and applications. The storage system shouldn't have to sacrifice its hardware resources to make copies of data.

A Snapshot is space efficient, meaning that it does not require a large amount of capacity to protect primary data. This is very different than full volume copies, which make an exact replica of primary data. Making a full volume copy of 5 TB of primary data requires an additional 5 TB of capacity. Making a snapshot copy of 5 TB of primary data requires no additional capacity at its initial creation. In fact, making 100 snapshots of 5 TB of primary data also requires no additional capacity. The snapshot will only consume capacity when *new* data is created. As a result, snapshots require much less capacity than full volume copies, providing a lower cost method for protecting primary data.

Using logical snapshots versus full volume copies can significantly reduce capacity requirements. Snapshot technology is a well-known storage system feature and is used broadly across the industry.

Writeable Snapshots

Every storage system that offers snapshot capability also provides a mechanism to restore snapshot images and make them read/write so that users can edit them again. When you mount a snapshot copy another physical copy of that data is created thus consuming an equal amount of capacity as the original primary copy. This process is different than a writable snapshot.

A writable snapshot is used to provide another working copy of primary data, which creates a second set of that primary data. However, this is a logical copy sharing the same source data as the original primary data.

Writable snapshots have the same advantages as read-only snapshots including near instantaneous creation and space efficiency, as described above. Like with read-only snapshots, at creation a writable snapshot requires no additional capacity and will only require disk space when new data is generated.

Writable snapshots enable companies to create another logical set of primary data copied from original data to be used for production purposes. Writable snapshots are used for application testing, data mining, and even data protection.

Consider the capacity savings of using writable snapshots compared to full volume copies of data. ESG is aware of companies that make six or more copies of primary data. If you are creating 6 replicas of 10 TB of primary data it results in an additional 60 TB of capacity used for copies. However, with writeable snapshots you can create 6, 10 or 100 copies and initially consume no additional capacity[1].

Writeable snapshots can have a major impact on PCSE depending on the number of copies you make. Only a small number of storage systems support writeable snapshot technology today and it is not a well understood technology. However, ESG believes that over time more solutions will support it. Certainly, in the context of PCSE it provides a major benefit.

Data Compression

While data compression is widely adopted in various aspects of IT, it has not been embraced at the network storage layer. ESG believes that there are two reasons for this. First, data compression is CPU intensive and needs dedicated or semi-dedicated resources. Interjecting any process into the I/O path will cause some level of latency. The goal is to minimize this to the point where it doesn't impact application performance. Perhaps the bigger issue is that data compression has not been a major requirement. However, with the requirement for more capacity efficient storage the awareness is being raised. Data compression is valuable and is tried and true standards-based open technology that should be considered when evaluating solutions for PCSE.

Data De-Duplication

Data de-duplication is the process of examining data to identify and then eliminate any duplicates. Data de-duplication solutions optimize physical storage capacity and bandwidth by using less of each. Reducing capacity results in fewer disks needed to store the same effective amount of data. And less bandwidth is required to move and copy the same amount of effective data across the WAN. Data de-duplication significantly reduces the amount of capacity required to store data – 5:1, 10:1; 20:1 and beyond.

Consider the value of reducing the amount of disk capacity to store your data. You can store 20 TB of data onto 1 TB of disk. That is landscape changing technology. ESG is aware of end users that are achieving these types of data de-duplication ratios. Today data de-duplication is found mostly in disk-to-disk (D2D) backup and active digital archival storage. However, it is moving up to primary storage as well.

Data de-duplication has the potential to change the landscape of storage. If you consider the impact that 10 to 1; 20 to 1; or more reduction ratios can have on your capacity requirements then you will understand the far reaching implications. As it relates directly to PCSE, data de-duplication is perhaps one of the most compelling technologies because of its ability to substantially reduce your capacity requirements while effectively storing the same amount of data.

Internal Storage Virtualization – Wide Stripe Groups

Some storage systems have the ability to stripe data across large numbers of drives providing greater levels of performance. This has an impact on PCSE since disk drives are not sitting idly but are being

[1] There will be some capacity consume for the snapshot metadata.

fully utilized. It is about reducing any wasted resources and having them work to perform read/write operations.

Storage System Consolidation – Multi-protocol Support

Storage consolidation is a well worn initiative in the data center with the goal of reducing the number of storage systems with the latest and greatest. ESG has seen bigger IT organizations consolidate six storage systems into one, which has a direct impact on PCSE. We are also witnessing a twist to storage consolidation that includes multi-protocol support. End users are consolidating iSCSI, FC, CIFS and NFS into single storage systems, which drives storage consolidation even further.

Intelligent Internal Tiering

ESG has found that 60% to 80% of all data is dormant 90 days after its creation. That means that no one accesses the majority of stored data 90 days or more after its initial creation. For each 10 TB of mission-critical data, 6 to 8 TB is not being used, and for each 100 TB of mission-critical data, 60 to 80 TB is not being used. This means a significant amount of the data that resides on your most expensive storage and is protected by your most stringent polices and processes are not being used by anyone.

Internal tiered storage is typically associated with different drive types, RAID levels and protection policies. Internal tiers are valuable, especially for those companies and organizations that don't want or need multiple storage systems. PCSE is impacted by intelligent internal tiering by moving data off of high performance drives such as 146 GB FC disk onto high density drives such as 500 GB SATA drives. You can store more data onto the 500 GB drives. The cost per KwH is lower for the denser drives and since they are the same physical size you get more GB per square inch.

End users must also consider tape as a part of a tiered storage strategy. ESG has been a big proponent of D2D backup storage systems, there is certainly a continued role for tape. Indeed, with PCSE becoming a priority issue, tape must be factored into the equation as a way to further improve efficiencies. Since tape is a removable media it can store data without any sustained power requirements and can be stored offsite thus consuming no data center floor-space. End users need to consider the right policies that determine when data is moved to tape (and removed from disk) and stored offsite.

MAID

MAID ensures that only disk drives in active use spin while the rest remain inactive. This provides the obvious benefit of drastically reducing not only the power a MAID array consumes, but less spinning drives results in less heat generated, thereby lowering cooling requirements.

MAID platforms are being positioned and used for both data protection and archival. A number of storage vendors are beginning to see the value of MAID and have begun implementing it into their platforms. MAID is a compelling technology with the potential to exist on all types of disk arrays including primary storage.

Storage System Architectures

Storage vendors are looking for new ways to reduce power and cooling consumption at a hardware and design level. They are relying more heavily on faster processors and cache memory to offset the use of disk drives for performance. Additionally, there are system processors that are designed to be more power efficient. Probably the biggest bang for the buck is power supply efficiency. Like the macro trend in data centers, the efficiency of power distribution and cooling systems can have a dramatic effect on the power efficiency of storage systems.

ESG's View

Data center power, cooling and floor space issues are becoming top of the mind with many end users at the executive level, IT professionals and data center facilities managers. IT managers have reported that they cannot deploy new applications because power demands cannot be met and that whenever they want to turn something on they have to turn something else off. Some CFOs are realizing that their large data centers are consuming more power than a small town. It is often the case that containing the cost of power and cooling is a strategic initiative that can reduce operating costs and positively impact bottom line growth. There are also "green" initiatives that are part of a bigger program that includes telecommuting, paper recycling, car pooling, etc. that also extends to the data center.

There are two major aspects of storage that lead to PSCE waste: 1) Buying more storage than you need; and 2) Underutilizing what you have. When analyzing any storage system it is important to understand how they address both these problems. ESG published a report in 2006 that found that the majority of end users storage system's capacity was underutilized by 30% to 50%. Consider the waste that this results in terms of PCSE. If you have 100 TB of storage that means that 30 TB to 50 TB are not being used but are still consuming power and taking up space. This is a gross misuse of resources in an environment that can't afford it.

Innovation applied to this problem is essential and technologies such as thin provisioning, dynamic volume expansion, snapshots, writeable snapshots, data compression, data de-duplication, virtualization, multi-protocol support, MAID, and advancements in storage system architectures can make a material impact on these issues. These technologies can streamline your capacity requirements and ensure that you fully utilize what you implement.

ESG recommends that you ask your vendors whether they support these technologies today or plan to in the near future. Discuss with them what they have found from the utilization of PCSE promoting technologies and to provide you with any quantification that they've done on consumption savings. Power, cooling and space efficiency is an ongoing concern that must be factored in as a requisite concern going forward.