Advisory White Paper

Power Play:

Factoring Power and Cooling Costs Into Tiered Storage Decisions

Dogenes Analytical Laboratories, Inc.

Synopsis

Power shortages, rising costs and the specter of restrictive regulations will drive forward-thinking IT managers to factor energy into their data center deployment decisions. Storage is quickly becoming the largest consumer of energy in the data center, creating a need for optimization and conservation strategies. Existing technologies can provide the answer.

The worsening power problem

Cheap, plentiful, there-when-you-need-it power can no longer be taken for granted. The combination of an aging national power grid and the increasing demand of a technology-hungry economy has driven consumption to the point that supply and demand are nearly equal in many regions. As a result, little buffer exists to meet peak power demands. "Brown outs" or rolling black-outs have hit the most populous industrial centers of the Northeast and Pacific Coast. The problem is especially acute in summer as residential and commercial interests compete for electricity to power air-conditioning. Such A/C units, of course, are vital to data center operations.

In the early days of computing, most data was stored offline (e.g., magnetic tape) because the cost of disk on a per MB basis was prohibitive. As disk storage has become cheaper, the balance has shifted such that nearly all data is kept online with only archive data stored on tape. The trend is continuing as IT organizations transition traditional backup operations to disk (e.g., disk-to-disk, virtual tape, continuous data protection). While disk-based backup reduces the recovery time objective (RTO), these devices contribute to and exacerbate power consumption. Even with de-duplication technology, these devices can increase disk power requirements 25% or more.

Government to the rescue?

During the 2006 California brownouts, the operating reserve of electricity dipped to just 5%, prompting Gov. Arnold Schwarzenegger to order state agencies to reduce consumption by 25%. Under the government's alert, certain large businesses voluntarily reduced electricity consumption.[1] Tens of thousands of businesses and residences in northern and Southern California suffered power disruptions.

In 2006, Congress mandated that the Environmental Protection Agency (EPA) examine data center power consumption. In this study, the EPA determined that IT energy use can be reduced by up to 80%, with a target of 45%, through "best practice" operations. The report describes these best practices as "moderate server and storage consolidation" and the adoption of energy-efficient devices.[2]

Among the EPA's recommendations was the metering of data centers (starting with federal facilities) and developing incentives through electric utilities to entice data center consumers to adopt energy-efficient practices. The report did not include specific recommendations for

[1] Randal C. Archibald, "As Heat Soars in California, Power Supply is Strained." New York Times, 25 July 2006, http://www.nytimes.com/2006/07/25/us/25power.html?_r=1&oref=slogin.

[2] "Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431," Environmental Protection Agency, Energy Star Program, 2 August 2007, pp. 8-9.

incentives, but such incentives could range from voluntary (i.e., reduce usage during peak periods) to promotional (i.e., price discounts for energy-efficient practices) to regressive (i.e., increasing rate structure as energy use increases). In any event, these solutions are all driven by the market.

Although the EPA did not recommend specific laws, it did highlight a need for an Energy Star designation for IT products. Because Energy Star involves standards, it can be assumed that Congress will take some role in establishing these standards or other regulations. We anticipate consumption regulations or mandatory use restriction regulations will impact data centers by 2011/2012.

The direct impact of tiered storage on energy consumption

One of the more interesting findings of the EPA study was that energy consumption can be reduced 45% to 80% using existing technologies. Although the report is not sufficiently specific, we suspect that the findings are based on the difference between the energy consumption of low capacity, high rotational speed drives versus high capacity, low rotational speed drives. While this assertion is true "by the numbers," it ignores business-oriented data access requirements. Obviously, not all data can be migrated to high capacity drives. Nevertheless, it is instructive to see the huge difference of energy consumption between different drive types. Table 1 below is a sample of some common Seagate drives and is based on manufacturer specifications. The costs noted assume $.12 per KwH.

KEY

Older disk devices can cost as much as 800% more to operate than newer technologies on a per GB basis.

Drive	Type	Capacity (GB)	Rotational Spd. (rpm)	Watts/hr. (typical)	Annual cost	Annual $/GB
Cheetah 10K.7	FC	146	10,000	15.5	$16.29	$.110
Cheetah 15K.5	FC	300	15,000	18.8	$19.71	$.066
Barracuda ES	SATA	500	7,200	13.0	$13.67	$.027
Barracuda ES.2	SATA	1000	7,200	11.6	$12.19	$.012
Barracuda ES.2	SAS	1000	7,200	13.0	$13.67	$.014

Table 1 – Sample costs of different drive specifications.

Seagate also claims that its latest, sixth generation, Cheetah 15K drives are optimized to reduce energy consumption by 61%, but has not published watt specifications.

The real point of this cost exercise is to demonstrate that older drive technologies may cost nearly 800% more per GB in electricity to retain. Thus, repurposing older arrays for "online archive" may be a false economy. Our research indicates that less than 10% of applications need to retain archive data for immediate access. The solution, however, is not moving archive data to new arrays. Any energy cost savings would be negated by the acquisition cost of the array.

The most cost- and energy-efficient array is the one you never use

The total amount of archive storage on disk in the data center is not the most important metric. Rather, the issue is what mix of older versus newer disk drives are deployed in the data center. To illustrate the impact, we will use an example of a data center with 100 TB of data evenly distributed across the five drive types shown in Table 1. We will assume that 20% of the data is archive data and stored on the oldest drive (Cheetah 10K). The purpose is not to create a real-world example but

rather compare costs where capacity is constant. The result, rounded off to whole numbers, is shown in the Table 2, below.

Drive	Amount of data	Capacity (GB)	Number of drives	Total KwH/yr	Annual cost	Percent of total cost
Cheetah 10K.7	20 TB	146	137	18,602	$2,232	48%
Cheetah 15K.5	20 TB	300	67	11,034	$1,324	29%
Barracuda ES	20 TB	500	40	4,555	$547	12%
Barracuda ES.2	20 TB	1000	20	2,032	$244	5%
Barracuda ES.2	20 TB	1000	20	2,278	$273	6%

Table 2 – Comparison of drive-type costs to support a given capacity

We can obviously conclude from this chart that deciding to repurpose obsolete arrays into archive devices is not likely to be a cost-effective solution; one such array could nearly double the power cost of storing data on disk.

For those applications that do demand online archive, IT organizations should consider a new generation of disk array in which the disk drives power down to a standby mode when not in use. Sometimes dubbed "MAID" (massive array of idle disk), these devices are specifically designed to deliver nearline data access times, while minimizing power and cooling requirements. The acquisition cost of these devices is further mitigated because we estimate that their useful life can be extended to five years versus the three years normally assigned to disk arrays.

Magnetic tape: Still the best choice for archive requirements

Although the claim that disk is cheaper than tape has been perpetuated for several years, our research indicates that this claim cannot be supported with objective data. The problem with these claims arises from acquisition-cost comparisons based on arbitrary capacity assignments. (See Diogenes Labs Advisory White Paper, *"It's 'and' not 'or': Understanding the changing roles of disk and tape,"* November 2006.) In our example above, we are concerned with the energy costs alone. Using LTO-4 media as an example, we can store 20 TB of data on 13 tapes (compressed). At $105 per cartridge, this is a one-time cost of $1,365; the media itself has no energy cost. Obviously, the media must be created and managed by a tape drive and tape library. The operating cost of an LTO-4 device is $252 per year at 100% duty cycles; most drives operate at 60%-80% duty cycles.

TIP

Seek to eliminate older arrays before replacing them.

Although the EPA correctly points out that existing technology can significantly reduce energy consumption, this advice must be applied appropriately. Replacing older disk technology with the latest devices would certainly reduce energy consumption, but finding ways to eliminate the older arrays entirely is the optimum solution. Thus, IT organizations should consider both storage consolidation for data that must be online, as well as offloading archive data to tape.

User Action: Tiered storage not only reduces product-acquisition costs, but saves money in terms of direct power consumption, as well. IT managers should consider storing archive data on tape, when it has a recovery time objective greater than 24 hours, to cut power consumption costs by 50% compared with storing the data on obsolete disk.

About Diogenes Analytical Laboratories, Inc.

Diogenes Analytical Laboratories, Inc. is an independent organization dedicated to helping information technology buyers reduce the inherent risk and uncertainty associated with technology purchases. Our goal is to create an informed IT consumer and provide the complete information needed to make smart purchase decisions. The opinions expressed in this report are those of Diogenes Analytical Laboratories, Inc.

Diogenes does not endorse products or vendors. Nothing in this document should be construed as an endorsement or recommendation of any technology or product for a particular purpose. From time to time vendors pay a fee to Diogenes for the right to reprint certain materials.