

# STORAGE SWITZERLAND REPORT

## BRINGING GREEN IT TO DATA STORAGE DOING MORE WITH LESS POWER



George Crump, Senior Analyst

Going into 2010, reducing power consumption is at the top of every CIO's project list and, the storage manager will be asked to play a role. Despite the fact that storage manufacturers are claiming to be power-efficient, most of that power efficiency comes as the result of 'fuzzy math'. Bringing Green IT to data storage is going to require more than just creative calculations. Real reduction in the amount of power is going to require tools to help the storage manager.

Using tools like Storage Profiler from [Tek-Tools](#) will allow storage managers to identify data that has not been accessed for a period of time and is eligible for migration to less expensive and more power-efficient drive technology. Creating a secondary tier of storage or an archive not only saves the company costs by lowering the amount of primary storage that needs to be purchased, it also increases power efficiency by storing more data for the same amount of watts consumed. This is the fuzzy math.

In reality, however, the most power-efficient drive is the one that is powered off, no fuzzy math required. The downside of a powered off drive is, of course, that it is difficult to retrieve data off of the drive. MAID (massive array of idle disks) storage is the alternative. Systems of this type, like those from Nexsan, will go through a

progression of power saving steps as the period of inactivity increases. With each step, power savings increase but so does the time required for the drives to return to full speed and fulfill a request for data.

Placing data on these drives requires an understanding of how the data is being accessed. Too much activity means the drives are spinning up and down all the time, potentially wasting more power. Too little activity means that the storage manager was not aggressive enough in moving data off of the primary store and as a result, is not gaining maximum cost savings and power savings from the MAID investment.

A simple use case is to divide the MAID storage into three groups, each with its own power efficiency setting. Using storage analysis tools like those available from Tek-Tools classify the current data set along four lines. The first three will go into the groups on the MAID storage and the final one will be left on primary storage. Data that has not been accessed in the last year or so should be placed on the first MAID group. The MAID setting for this group should be at a maximum level of power savings. As stated earlier, this does mean that if some of the data does need to be recalled from this storage there will be a slower response time, but not one that the typical user is going to notice.

The first group is the classic archive use case. But MAID storage can be used for much more than archive. For MAID to save power it just needs to be idle more often than standard storage systems. A storage analysis tool plays an even more important role in determining which data fits into the final two use cases - "seldom accessed" and "daily accessed".

Seldom accessed data is data that has not been accessed in the last four weeks, but is not ready to be classified as "monthly or quarterly records" (quarterly report templates being good examples). This data can be positioned in the second group with a more modest power saving setting and less of an effect on performance.

The final use case is data that is accessed throughout the day, but does not demand the highest performance and is likely to be idle most of the evening. The best example of this data is user home directories. These can be positioned on the third group within the MAID array with the minimum power savings setting. This will still provide some power efficiency, especially in the evening, but not adversely affect performance.

The final data area that remains is the most active set of data, traditionally, not a good candidate for a power savings discussion. However, with a storage analysis tool the primary data set can also be considered.

The first step is to identify data, typically database data, that is to be written to RAID volumes with high drive counts. Often this is being done to give better storage I/O performance to the database, not for capacity. The storage analysis tool will show that most of the capacity is unused. These database applications could be moved to SSD drives that are far more power-efficient and will, of course, provide better performance.

With these actions completed, the storage environment becomes significantly more power-efficient and has likely led to significant cost reductions as well. The result of all of these moves should lead to significant amounts of free capacity on primary storage. This excess capacity should lead to the delay in the purchase of additional storage that would also need to be racked, powered and cooled.

If that is not the case then the storage administrator ends up with more free space on primary storage. That free space still has to stay racked, powered and cooled. Just because there is now 50% more free capacity does not mean that 50% of the drives and drive shelves can be powered off. This is because most RAID groups in an array are a combination of drives from different shelves. This is a solid best practice and protects data loss from a single shelf power failure bringing down an entire data volume. It does however present a challenge in powering down the excess capacity that deduplication, compression and archiving will certainly deliver.

There is one more option that a storage analysis tool allows; rearranging RAID layouts to use less shelves and then powering off the unused shelves. In fairness, this is a drastic measure, and no small undertaking. But, if there is extreme pressure to save power it may be justified. For example, many data centers can no longer get additional power to the building and more aggressive measures may be justified.

When it comes to increasing storage power efficiencies, a storage analysis software application like Tek-Tools Storage Profiler is a critical tool for users to maximize the capabilities of MAID storage, archive and deduplication.

### **About Storage Switzerland**

**Storage Switzerland, is an analyst firm focused on the virtualization and storage marketplaces. For more information please visit our web site: <http://www.storage-switzerland.com>.**