

# Technical Brief

**Infotrend**<sup>®</sup>

## Increasing the Efficiency of Your RAID Storage to Make it Greener

### **Abstract**

This document proposes some important factors users should take into consideration when they want to make their storage greener by making it more efficient.

### **Infotrend<sup>®</sup> Technology, Inc.**

Networked Storage Solution Provider

Revision 1.0

July, 2008

Infotrend Technology makes no representations or warranties with respect to the contents hereof and specifically disclaims, any implied warranties of merchantability or fitness for any particular purpose. Infotrend reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation to notify of such changes.

Infotrend, Infotrend logo, and EonStor are registered trademarks of Infotrend Technology, Inc.

## **Green Trend and the Storage Industry**

'Environmental friendliness' has been a common topic recently and now we are all aware that we should take better care of our planet. So, do we simply buy greener products? Is it really as easy as that?

In our personal and business lives we have a number of choices when it comes to choosing a product that will meet our needs and be environmentally friendly. From long-life light bulbs to partly eliminating the need for air travel by holding video conference calls, we can all help make a difference. The sector of the IT industry that has received most of the 'green' limelight is the storage industry. With organisations around the world generating extremely large volumes of data every day, the need to store, access and protect this information is paramount. So how can we ensure that data is suitably stored without unnecessarily harming the environment?

Product efficiency is the invisible, yet key, factor, as poor efficiency usually leads to higher power usage. Improve the efficiency of your storage infrastructure and suddenly your data centre will become greener and your TCO (Total Cost of Ownership) more appealing.

Some vendors have focused on data storage power consumption when the products are idle. However, much higher savings can be realised by reducing the power requirements of those products which are operational the majority of the time because it is these devices that will have a greater impact on energy bills.

So how can storage and data centre administrators get the best of both worlds? How can they make their infrastructure more efficient, more cost-effective and more environmentally friendly? Below are some considerations that will help put them on the right path.

## **Raise the Efficiency, Raise the Green Level**

No matter how effective a product's power-saving features are, if its overall efficiency is poor, it is not green. Every application has a unique requirement whether the priority is performance, capacity, cost, data availability (redundancy), connectivity, etc., or a unique balance of those features. The first priority should always be that the solution fits the purpose and then improve the efficiency. By doing so, you may find the storage products that you are currently using are greener than the ones marketed as 'green.'

## **Choosing the Suitable Storage**

Some RAID products are specifically built for data archiving, some are designed for 24x7, high-availability environments, and other multi-purpose subsystems can be tuned and used in different applications. These various types of storage products offer different levels of availability, security, performance and capacity. Because different types of disk drives are used in high-availability storage as opposed to archiving, the hardware, configurations and power requirements differ. A side-by-side comparison of their environmental impact will therefore lead to misleading results.

For example, demanding environments such as database applications should rely on disk drives with a high spindle speed (10,000 RPM or higher), and include more disk drives in the configuration so that more transactions per second can be performed. Reliability is also vital to support 24x7 data availability. In this way and with the help of the cache memory in the RAID controller, the RAID system can process vast numbers of small transactions coming in random order. On the other hand, disk drives used in data archiving need to be large capacity and have a lower unit cost. The performance requirement here is not high, thanks to the snapshot backup functionality; reliability is not critical for data archiving as in many cases the drives are left idle.

For these reasons, a RAID system designed or configured for data archiving cannot be used in a high-availability environment. Can you imagine online banking users having to wait 30 seconds for the drives to 'wake up' before being able to access their account details? In this case, all the disk drives must stay active in order to handle the high volume of random transactions, so the ability to have idle disks holds no appeal. On the other hand, putting a RAID system designed and configured for 24x7, high-performance operation into a data archiving environment will feature a capacity/price ratio that will make the company CFO jump.

## **Sharing Storage**

Efficiency is closely linked to maximisation of resources. In a typical office with 50 users, each desktop or laptop PC has a built-in disk drive that stores important data which the IT manager needs to back up along with the servers and any other shared data. Suppose each user was given a USB disk (or similar solution) to backup to; if each USB disk has 300GB capacity and uses 60W of power that corresponds to approximately 15TB of capacity and 3,000W of power consumption. Besides the management issue, some users may need more capacity, while others only have small amounts of data to backup.

By backing up to a central disk array via the office network, all the unutilised storage is consolidated to serve more users and make better use of available capacity. If 50 users share a 15TB, 500W disk array, individual power consumption is just 10W.

To make the most of centralised storage when choosing a disk array, the IT manager should favour high performance and capacity in order to serve more users and provide faster service. High performance helps increase efficiency which makes a product more environmentally friendly.

If there are two RAID systems and one performs four times faster than the other, the faster system can replace four slower systems. By consuming less power and generating less heat than four slower systems, the one faster system is much greener. Hence, disk array performance should be a major consideration when thinking about green storage.

### **Smart Disk Backup – how and when**

It's important to note that archiving, disk backup (write-once-read-occasionally) and snapshot backup are different in nature. As they have their own unique requirements and provide different types of services, in many situations the features do not overlap and should be used in conjunction with each other to customize the best storage solutions possible for your needs.

One way to ensure infrastructure efficiency is by using the "snapshot backup" function to minimize the backup window. Ideally, IT managers should use the snapshot function that is built-in and performed by the RAID system rather than the host software. Data transfer is accomplished without the host software intervening, thus avoiding unnecessary host CPU utilisation and occupation of the host access data bandwidth. As a result the disk backup takes just a few seconds, not hours. Moreover, snapshot backup decreases the frequency of full backups and archiving, thereby saving more energy because the devices not in use can be turned off.

Once data has been recorded onto off-line media, it stays there without using additional power. This makes an off-line storage product a greener choice for archiving purposes instead of using massive disk arrays with power-saving features.

When to do backups is another important factor concerning efficiency. Scheduled backups and other maintenance tasks should occur during non-peak work hours

utilising the scheduler function available on the disk arrays in order to avoid affecting the service performance. Analysing the applications' time and usage patterns and scheduling automated tasks further enable the arrays to be used around the clock in a more efficient manner.

## **Invest in the Right Technologies and Start Saving the Planet**

In summary, your storage can indeed be greener if you bear in mind the following factors:

- Understand your actual storage requirements and envisage the demands your organisation will place on them over the next three to five years.
- Ensure your storage infrastructure meets business requirements and priorities.
- Select the right storage product by considering the nature of the application, efficiency, performance, and technologies.
- Choose the storage vendor which provides a broad range of solutions.
- Choose a multi-purpose storage product to meet dynamic business requirements. A single-purpose storage product may not be suitable if it is not used as originally intended.

The extent to which a solution will conserve power usage is determined by the combination of business requirements, efficiency and performance.

Efficient storage is green storage.