

Disk-to-Disk Backup Solutions

By Ben Ginster

COMPANIES LARGE AND SMALL USE MORE DISK STORAGE EACH DAY. More important than this, their primary businesses depend more on electronic data than ever. Healthcare providers need larger amounts of disk space to hold their medical images. Banks face ever-growing requirements for electronic data retention. Even small businesses have begun to rely on programs with more advanced storage requirements like database-driven applications. And whether a multinational conglomerate or a local law firm, businesses cannot afford to lose this access to data for even an hour. The bottom line is that companies need to backup more data than ever before, faster than ever before.

For years, when people mentioned the word backup, tape media immediately sprang to mind. But as companies have realized the importance of speed—of reducing the time it takes both to backup data and to restore from backups, a new backup paradigm has emerged—one called disk-to-disk, or D2D.

It's a trend that both small and large sized companies are noticing and one that has already prompted many to implement, according to forward-looking statistics that show that disk-to-disk backup usage is on the rise. In late 2004, Storage Magazine's 2004 Purchasing Intentions Survey found that 53% of respondents planned to increase spending on disk-to-disk backup methods in 2005 while only 31% of respondents planned to increase their use of tape backup, down from 56% in 2003.

There are three primary drivers for this rising demand. First—speed. Disk offers users faster backups and more importantly, faster restores. Second—cost. The relative cost of hard drive media for the amount of storage it provides (\$ per MB) has significantly dropped and removed any cost advantage for tape media. The third driver for disk-to-disk backup adoption is reliability—tape is simply no longer cutting it when it comes to backing up and restoring data effectively. Because of the linear nature of tape, tape backup has inherent difficulties with backing up and restoring large amounts of data. Other potential risks surrounding tape include robotic library parts failure, backup and restore failures due to tape being exposed and brittle media. Too often, using tape leaves a single stressful question in the back of your mind: Is my backup going to be successful tonight?

With data increasing and backup windows decreasing, disk-to-disk has become the de facto solution for these problems. But even after doing a bit of research, users can be easily confused as to what exactly disk-to-disk backup is and how to implement it.

Type “disk-to-disk backup” or “D2D backup” into any search engine and a slew of companies that offer a variety of technologies claiming to be D2D will be listed. Some companies offer backup over the WAN to a colocation center, or hosted disk solution. Others offer

“Virtual Tape,” which in essence, makes the disk array appear as a tape library. There are NAS, SAN, RAID, MAID, and many other options. The choices among flavors of disk-to-disk backup to implement can seem as dizzying as the selection at an ice cream store.

In an effort to clarify the confusing group of options for disk-to-disk backup solutions, below is a basic overview. D2D comes

down to two types of offerings—disk-to-disk solutions that still require tape and those that don't.

DISK-TO-DISK SOLUTIONS THAT REQUIRE TAPE—D2D2T

When looking at a D2D backup solution, read the fine print to see what you will actually be doing with your data. Many on the market today tout the benefits of backing up to disk, pointing out the obvious flaws with tape backup as the justification for using their systems. But a closer look at their packaging often reveals a word that may confuse you—“tape.”

Many of these “disk-to-disk-to-tape” backup solutions use one form of disk or another to quickly backup data to a disk based target to cache or store your data. But once the data has been written to this disk, sub-system tape is still required if you want to take data offsite to adhere to best business practices or comply with regulatory measures. Although these systems add an extra step in the process, D2D2T is for now the more common D2D method in use.

These systems can indeed quickly backup and restore data that is nearline and connected to the network. However, when companies want to restore historical data, they still face a time-consuming process. Historical restores must be run from offsite tapes, which



involves copying the restored data onto the network attached device, and then requiring it to be copied back to the server in order to access that data. This three-stage process is complicated and inefficient, and

ties your company to a single format that may be increasingly hard to procure and even rise in cost over time. And in times of disaster, when quick restores are most important, the media will only be readable by

the proprietary system, which may itself be in crisis. Standard parallel and SATA hard drives can, on the other hand, be run from nearly any standard computer in the world.

Finally, look for solutions that enable you to create a sound, offsite rotation of disks. An example is a “Grandfather-Father-Son”

The alternatives that are currently available for backing up your data to disk are excellent if you know what to look for. In all, when considering solutions, the main thing is to choose the one that will most swiftly and thoroughly protect your data, and enable you to get that data back most quickly in a crunch.

leaves the door open for some of the earlier problems discussed when getting data off of tape media—in the end, sometimes the tape drives just won’t read the tapes.

DISK-TO-DISK SOLUTIONS THAT DON’T REQUIRE TAPE—D2D

Pure D2D backup uses disk as both the portable and offsite media. With the reliability, cost and capacity of hard disks being more favorable than ever before, this method is making a comeback and a few innovative companies are leading the way with truly tape-free disk-to-disk backup solutions.

The benefits of using disk as a replacement to tape are many. Disk media uses a random access approach to data storage rather than a linear one, which makes backup and restore much faster. Scalability and disk capacity are other advantages of disk over tape. Parallel and SATA disk capacities have increased from 300GB to 500GB over the last 12 months and are estimated to grow to 850GB sometime in 2006. Disk systems that use hard drives in place of tape take advantage of this trend, and will enable companies to increase their backup capacity as the need arises, simply by adding or swapping in larger disks.


Also, hard drives are inherently more reliable than tape drives because disk media is not exposed to the environment and is not flexible media. Unlike disks in a PC or RAID array, ejectable disks are used expressly for backup, adding to the usable life of these disks. Some manufacturers even use caddies and protective cases to further protect disks, making them infinitely more durable than tapes.

WHAT CRITERIA SHOULD I USE WHEN SELECTING AN EJECTABLE DISK BACKUP SYSTEM?

First in mind when selecting an ejectable disk backup system should be ease of use, flexibility and ease of integration into existing network and backup schemes. These systems should not require new backup software. Most of today’s backup software applications support backup to disk and have advanced in usability such that backing up to removable disk is as easy as backing up to tape. Some even have additional options such as spanning disks and storing multiple days’ worth of backup in separate folders on one large disk. The best systems will enable the user to use any capacity hard drive from any manufacturer and allow them to be interchangeable within the same system.

As with any technology, avoid removable disk systems that utilize a proprietary format, in this case, for ejectable disks. Proprietary media

rotation that enables you to store daily, weekly and monthly full backups offsite.

The alternatives that are currently available for backing up your data to disk are excellent if you know what to look for. In all, when considering solutions, the main thing is to choose the one that will most swiftly and thoroughly protect your data, and enable you to get that data back most quickly in a crunch. 

NaSPA member Ben Ginster is a Channel Manager for IdealStor.

