



## TECHNICAL BRIEF

# Are You Feeling Lucky?

## Three Key Questions to Consider Before Migrating to Disk-based Backup

Without proper due diligence, adding disk-based backup components to your environment can be like spinning a roulette wheel — the results can vary greatly. A common misconception I often hear is “Once I get rid of tape in my environment, all my backup management and performance problems will go away.” Tape can work well for backup, and has for decades...yet many have experienced growing pains while traveling the long road to maturity of tape devices in complex SAN environments. They can be difficult to configure, it is hard to keep stable, and doubly hard to maintain consistent data streaming rates in a shared drive environment.

Some people think that adding some random-access, low-cost disk will cure all their IT ills. Maybe, but maybe not. While there are many advantages to disk-based backup when properly sized, configured and managed, you might want to consider the following before embarking on this endeavor:

**Can I migrate to a completely disk-based environment or will I need to address some backup situations with tape?** It is often impractical and cost prohibitive to remove tape completely from most environments. Unless the bandwidth exists to replicate all backup data from the primary disk target to an offsite disk target, making a tape copy for offsite storage is usually a necessary evil. Therefore, test your current bandwidth to a remote location by transferring a large file with a simple tool like ftp. Also, don't underestimate the amount of data being transferred during a nightly backup cycle. Most disk backup target replication scenarios offer some kind of compression or delta-based transfers to reduce the size of the payload transferred. However, for testing purposes use the un-enhanced file transfer performance as a baseline and then calculate theoretical rates based on the selected product specifications.

**If both disk and tape will be required how will they interact?** If, for instance, the tape component is solely used for offsite copies, how will the data be migrated from the disk target to the tape target? Most backup applications support disk backup connectivity through SAN, NAS, iSCSI and DAS. Consider the maximum throughput of the disk target and the connection pipe when designing your disk solution. Try to avoid multiple network hops during the disk to tape duplication process to avoid latency delays.

Remember to leverage the strengths of combined disk and tape solutions. Newer high performance tape drives such as LTO5 handle large contiguous data streams very well. It is hard to beat a database backup streaming to a new high-end tape drive. Disk-based backups perform well in multi-client, small stream situations. In addition, the elimination of tape mounting and positioning make disk backups an excellent choice in environments where file level restores are handled on a regular basis.

**What class of service do I need to achieve?** With all the specialized disk backup appliances available today, from virtual tape libraries to single instances to de-duplication, it can be difficult to determine the class of service a solution will provide. The old rules of thumb used to calculate performance can be clouded in the virtual layer between the backup application and the actual device layer. The results achieved from performance tuning a physical SAN attached tape drive are usually easy to measure because of the one-to-one association. This may not be the case when working with a virtual tape drive assigned to a backup host. The one-to-one relationship between the backup host HBA and virtual tape device is masked when working in the virtual library configuration. It can be extremely difficult to identify disk spindle contention in this type of environment. Similarly, the ability to assign many virtual tape drives to many hosts via the front end fibre ports of this type of appliance are both its strength and weakness. It makes it easy to assign drives, but may make it impossible to understand contention and conduct any type of consistent performance tuning of the environment if the backend disk is virtualized too much.

So are you feeling lucky? With a little bit of planning, research and a clear understanding of the results you are trying to achieve, you can avoid the pitfalls that can exist in what at first glance seems like a sure bet.

#### ABOUT DAYMARK SOLUTIONS, INC.

**Daymark Solutions, Inc. is an experienced technology solutions provider focused on architecting, providing and implementing effective server, storage, network and security solutions for businesses in the Northeastern United States. Our mission is to help you use technology effectively to solve the business challenges you face today, while making sure your systems are agile enough to adapt to future requirements. For more information, call 781-359-3000.**

*Daymark Solutions – 23 Third Ave – Burlington, MA 01803*

*[www.daymarksi.com](http://www.daymarksi.com)*