

## Disk-based backup and restore: reaching the next level

*Alexandre Delcayre looks at the advantages of using a Virtual Tape Library solution for data protection and disaster recovery.*

Escalating data growth, shrinking backup windows, over-burdened IT staff and overtaxed budgets... the problems facing storage administrators are all too familiar. Despite familiarity with the problems, for many the solution is as elusive as ever. Here, Alexandre Delcayre, outlines and compares conventional disk-based backup and restore methods with new disk-based techniques, including the virtual tape library.

If you're reading this, you're almost certainly familiar with the challenges associated with data storage; you may well have your own checklist for what your chosen state-of-the-art technology needs to offer to meet those challenges. Among others, you're probably looking for high-performance and reliable backup through the elimination of media error and robotic failure, rapid restore, and the old favourites of reduced operating cost (total cost of ownership, or TCO) and maximised return on investment (ROI).

To understand what the 'state-of-the-art' of those technologies is, it's helpful to recap how we got here. In the beginning, things were very simple: the application server was directly connected to the tape drive, and would backup its own disk(s) to a dedicated tape. In fact, this first model worked very well - until the number of application servers reached, say, double-digits.

The second, and current, approach to backup and restore was originally established on the premise that each physical tape drive can support ONE backup or restore session at any given time. And, as the amount of data to be backed up increased, coupled with business operations on a 24x7 basis becoming the norm, multiple tape drives/libraries needed to be purchased, deployed and maintained to accommodate concurrent backup and restore sessions.

So, with the high cost of connecting numerous application servers to tape libraries, IT managers instead designated backup/media servers to front-end the tape libraries. Special backup 'agents' are then installed on each application server to push data to the backup server for staging and streaming to the tape drive. This two-step approach is designed to address both the connectivity issue as well as differences in performance between disk and tape.

In a typical Fibre Channel (FC) SAN environment, the tape library is accessed by multiple backup servers. Since each tape drive can only be used by one backup server at a time, the backup servers need to co-ordinate the use of tape drives in the library to avoid conflicts. For example, Veritas backup products require the Shared Storage Option to co-ordinate access to tape libraries and tape drives in a FC LAN.

Now, in the current backup model, the backup server relies on the remote backup agent to push the data from the application server over an IP-based LAN, which limits the speed of any single backup to the speed of the LAN. Since the speed of the LAN is often less than the optimal speed of the tape drive, backup applications use techniques such as interleaving or multiplexing to allow multiple backup agents to stream to the tape drive.

While this offers the advantage of allowing the tape drive to operate at its best speed, it requires that backups be co-ordinated across multiple application servers to provide the necessary data. It also has the unfortunate side-effect of actually increasing the time needed to restore data, since the backup application now needs to stream through multiple backup sessions to restore the data for the desired application server. On top of all this complexity, the output of the backup agents adds to the workload of the LAN.

Given that many of the application servers are on a FC or iSCSI SAN, the easiest and most effective way to eliminate this bottleneck is to provide a SAN-based fast path to the backup server to access the application server's disk(s) at the speed of the SAN (up to 2 GB/s). For example, the HyperTrac Backup option of FalconStor's IPStor Enterprise Edition software delivers a SAN-based fast path by using database-aware snapshot technology to capture a consistent image of the application host's disk, provision the snapshot to the backup server over a FC or iSCSI SAN, and – crucially – enable the entire backup operation to take place without impacting the application server's performance.

Given that the transfer rate of a FC disk array exceeds the speed of a tape drive by approximately 3:1, backup software vendors are beginning to offer disk-to-disk (D2D) backup options that use high-speed disk arrays to cache the backup data stream. This eliminates the tape drive as the bottleneck on the data path (if the backup software can stream the data faster than the maximum speed of the tape drive) and accelerates the restore process if the tape session is already on disk.

Typically, though, IT managers must configure individual backup servers to perform D2D-based backup and assign the iSCSI or FC disk storage needed to stage or cache the backup data stream. The third party backup vendor's D2D option allows the IT manager to use the high-speed disk to cache/stage the backup stream, provided the backup software's throughput is greater than the speed of the tape drive.

### **The Virtual Tape Library**

While such D2D backup offers improvement on the current models widely employed, it still requires reconfiguration of storage and provisioning of disk storage for caching. But the concept of the Virtual Tape Library [VTL] uses disk to simulate the tape drive/library, thereby empowering the IT manager to create and attach virtual tape drives and libraries to application or backup servers on a FC or iSCSI SAN as if they were physical tape drives or libraries, without incurring any additional material costs.

Such emulation of industry standard tape drives and libraries allows existing backup products to interact with the virtual tape drives and libraries in a seamless manner, without changing the existing configuration or policy. And, in that each backup server can perform tape backup/restore to its dedicated tape drive or library, the VTL model also eliminates the need to share a tape drive among multiple backup servers on the SAN.

The benefits of a Virtual Tape Library solution are:

- \* Elimination of the need to purchase, install and maintain tape drives, generating immediate fiscal savings;
- \* Elimination of media and robotic errors that prevent successful backups, increasing the success rate of daily backups;
- \* The ability to leverage the high transfer rates of state-of-the-art disk arrays to accelerate backup, and restore the throughput of existing backup servers;
- \* The facilitation of tape vaulting by replicating disk-based virtual tape cartridges to the remote vault, over an IP network;
- \* The facilitation of tape vaulting by enabling the backup server to copy tapes to disk-based virtual tape cartridges in the vault;
- \* Enabling third party backup software and industry standard NAS appliances such as Network Appliance filers to perform backup and restore using the NDMP protocol over an IP network;
- \* Allowing IT managers to deploy virtual tape library technology at branch offices, enabling data to be backed-up on site, then transferred to the main data centre over IP.

With its higher speeds and greater flexibility, together with falling disk prices, disk-based backup is, almost inevitably, an attractive proposition. So, while today's tape-based approach

to backup has worked well, the introduction of Virtual Tape Library solutions will change the way we backup our data. VTL doesn't replace tape, but it does change its place in the backup schema.

*Alexandre Delcayre is EMEA technical director at FalconStor Software.*

*FalconStor are exhibiting at Storage Expo the UK's largest event dedicated to data storage. Now in its 4th year, the show features a comprehensive free education programme, and over 90 exhibitors at the National Hall, Olympia, London from 13 - 14 October 2004. [www.storage-expo.com](http://www.storage-expo.com)*