DISK-TO-DISK BACKUP: WHY, HOW, WHERE

Andrew Brewerton, technical director, BakBone Software

Backup technology is changing, and it is changing fast. Not so long ago, backing up meant copying your primary data from hard disk to tape - initially to the spools of half-inch tape beloved of film directors, and more recently to various types of tape cassettes and cartridges.

Now though, more and more organisations are using hard disks for their backups as well as their primary data, a process that has become known as disk-to-disk backup, or D2D for short.

There are a whole host of reasons for this shift. In particular, the cost of hard disks has fallen dramatically while their capacity has soared, and disk arrays have much better read/write performance than tape drives - this is particularly valuable if an application must be paused or taken offline to be backed-up.

In addition, tape is quite simply a pain to work with, especially if a cartridge must be retrieved, loaded, and scanned in its entirety, just to recover one file. Tapes can be lost or stolen, too. While we put up with all that in the past because we had to, that is no longer the case.

Sure, a tape cartridge on a shelf - albeit in a climate-controlled storeroom - is still the cheapest and least energy-consuming way to store data for the long term, but that is increasingly the role of an archive not a backup. So while tape is unlikely to vanish altogether, its role in backup is declining fast.

Disk can be incorporated into the backup process in many ways, from virtual tape libraries (VTLs) through snapshots to continuous data protection (CDP), and each method may suit some applications or user requirements better than others. An organisation may even use more than one D2D backup scheme in parallel, in order to address different recovery needs.

Disk is also used in many other forms of data protection, such as data replication and mirroring, although it is important to understand that these are not backups. They protect against hardware failure or disasters, but they cannot protect against data loss or corruption as they offer no rollback capability.

Why backup?

It is an inescapable fact that computers and their disk storage go wrong - not often, but they do. Even when they do not, they can be stolen, destroyed by fires or other disasters, damaged by power failures, or suffer one of a host of other events that cause data to be lost, such as a user deleting a file by mistake or saving the wrong version. When that happens, you need a backup.

Replication and mirroring do fill some very important data protection needs, because if one whole system is lost, the other should be able to take over. However, a replica or mirror is not a backup, because if data on one half of a pair is corrupted or deleted, it will also be deleted or corrupted on

the other half.

Backup is all about saving data at a specific time and in a consistent state. It allows you to go back to a specific recovery point, for example to recover an earlier version of a corrupted or deleted file or database, to restore a crashed system to its previous working state, or to take a copy off-site for safekeeping.

A backup traditionally has been stored on non-volatile and removable media, such as tape, due to that need to move copies off-site. Increasingly though, this role is being filled by non-removable media at a different site - either your own secondary site, or a facility managed by a service provider - with data networks replacing lorries as the transport mechanism.

When it comes to restoring data, disk's big advantage over tape is that it is random-access rather than sequential access. That means that if you only need one file or a few files back, it will be faster and easier to find and recover from disk.

What backup and recovery methods you use will depend on two factors - the recovery point objective (RPO), i.e. how much data the organisation can afford to lose or re-create, and the recovery time objective (RTO), which is how long you have to recover the data before its absence causes business continuity problems.

For instance, if the RPO is 24 hours, daily backups to tape could be acceptable, and any data created or changed since the failure must be manually recovered. An RTO of 24 hours similarly means the organisation can manage without the system for a day.

If the RPO and RTO were seconds rather than hours, the backup technology would not only have to track data changes as they happened, but it would also need to restore data almost immediately. Only disk-based continuous data protection (CDP) schemes could do that.

Ways to use disk

Most current disk-based backup technologies fall into one of four basic groups, and can be implemented either as an appliance, or as software which writes to a dedicated partition on a NAS system or other storage array:

* **Virtual tape library (VTL)**: One of the first backup applications for disk was to emulate a tape drive. This technique has been used in mainframe tape libraries for many years, with the emulated tape acting as a kind of cache - the backup application writes a tape volume to disk, and this is then copied or cloned to real tape in the background.

Using a VTL means there is no need to change your software or processes - they just run a lot faster. However, it is still largely oriented towards system recovery, and the restore options are pretty much the same as from real tape. Generally, the virtual tapes can still be cloned to real tapes in the background for longer-term storage; this process is known as D2D2T, or disk-to-disk-to-tape.

Simpler VTLs take a portion of the file space, create files sequentially and treat it as tape, so your save-set is the same as real tape. That can waste space though, as it allocates the full tape capacity

on disk even if the tape volume is not full

More advanced VTLs get around this problem by layering on storage virtualisation technologies. In particular this means thin provisioning, which allocates a logical volume of the desired capacity but does not physically write to disk unless there is actual data to write, and it has the ability to take capacity from anywhere, e.g. from a Storage Area Network, from local disk, and even from Network Attached Storage.

* **Disk-to-disk (D2D):** Typically this involves backing up to a dedicated disk-based appliance or a low-cost SATA array, but this time the disk is acting as disk, not as tape. Most backup applications now support this. It makes access to individual files easier, although system backups may be slower than streaming to a VTL.

An advantage of not emulating tape is that you are no longer bound by its limitations. D2D systems work as random-access storage, not sequential, which allows the device to send and receive multiple concurrent streams, for example, or to recover individual files without having to scan the entire backup volume.

D2D can also be as simple as using a removable disk cartridge instead of tape. The advantage here is backup and recovery speed, while the disk cartridge can be stored or moved offsite just as a tape cartridge would be.

* **Snapshot:** This takes a point-in-time copy of your data at scheduled intervals, and is pretty much instant. However, unless it is differential (which is analogous to an incremental backup) or includes some form of compression, data reduction or de-duplication technology, each snapshot will require the same amount of disk storage as the original.

Differential snapshot technologies are good for roll-backs and file recovery, but may be dependent on the original copy, so are less useful for disaster recovery.

Many NAS (network attached storage) vendors offer tools which can snapshot data from a NAS server or application server on one site to a NAS server at a recovery location.

However, in recent years snapshot technology has become less dependent on the hardware - it used to be mainly an internal function of a disk array or NAS server, but more and more software now offers snapshot capabilities.

* **Continuous data protection (CDP):** Sometimes called real-time data protection, this captures and replicates file-level changes as they happen, allowing you to wind the clock back on a file or system to almost any previous point in time.

The changes are stored at byte or block level with metadata that notes which blocks changed and when, so there is often no need to reconstruct the file for recovery - the CDP system simply gives you back the version that existed at your chosen time. Any changes made since then will need to be recovered some other way, for example via journaling within the application.

CDP is only viable on disk, not tape, because it relies on having random access to its stored data. Depending on how the CDP process functions, one potential drawback is that the more granular you

make your CDP system, the more it impacts performance of the system and application. So technologies that do not rely solely on snapshot technology offer an advantage.

In addition, it can be necessary to roll forward or backward to find the version you want. One option here is to use CDP to track and store changes at very granular level, then convert the backed-up data to point-in-time snapshots for easier recovery.

Beyond data protection, a well designed CDP solution can bring other advantages, such as a lower impact on the application and server. It also moves less data over the network than file-based protection schemes, as it sends only the changed bytes.

Coherency and recovery

In order to be useful, a backup has to be coherent - a copy of something that is in the middle of being updated cannot reliably be restored. With traditional backup methods, applications would be taken offline for backup, usually overnight, but newer backup methods such as snapshots and CDP are designed to work at any time.

Snapshots provide a relatively coarse temporal granularity, so are more likely to produce a complete and coherent backup. However, they will miss any updates made since the last snapshot. The fine-grained approach of CDP is less likely to lose data, but it may be harder to bring the system back to a coherent state.

How you achieve a coherent backup will depend on the application or data. For instance, with unstructured file systems you need to find a known-good file version - typically the last closed or saved version. For files that can stay open a long time, you need to initiate a file system flush and create a pointer to that in the metadata.

To recover data, you would then find the right point in the CDP backup, wait for the data to copy back to the application server and then reactivate the application. However, that means that the more data you have, and the slower your network is, the longer recovery will take.

Fortunately, technologies are emerging to speed up this process. These provide the application with an outline of the restored data that is enough to let it start up, even though all the data has not yet truly been restored; a software agent running alongside the application then watches for data requests and reprioritises the restoration process accordingly - in effect it streams the data back as it is called for.

Schemes such as this can have applications up and running in less than 10 minutes, as the quickly recovered shell-file is just a few megabytes. Of course it does still take time to fully restore the application, but it does allow users to start using it again immediately.

One other issue that may affect the choice of snapshots or CDP is the level of interdependency within the application and its files. If there is too much interdependency, it will be more difficult to find a consistent recovery point. A potential solution is to choose software that is application-aware and can apply granular recovery intelligently, because it knows the dependencies involved.

Power and efficiency issues

One thing that must be said in tape's favour is that its power consumption for offline data storage is very low - potentially as low as the cost of the air-conditioning for the shelf space to keep the cartridges on. Removable disk cartridges can match that of course, but only for traditional backup processes with their attendant delays.

To use newer backup processes such as snapshots and CDP requires the disk storage to be online. D2D hardware developers have therefore come up with schemes such as MAID (massive array of idle disk), which reduces power consumption by putting hard disks into a low-power state when they are not being accessed.

MAID-type systems from the likes of Copan, Hitachi Data Systems and Nexsan, and related technologies such as Adaptec's IPM (intelligent power management) RAID controllers, therefore allow banks of disk drives to operate in different power states at varying times.

For instance, they can automate drives to go into standby mode or even spin down completely during idle periods. If a drive is accessed while powered down, the controller will spin it back up; alternatively the administrator can define peak IT activity periods when drives will never be spun down. The controller also monitors drives that have been powered down for a while, to make sure they still work OK.

Conversely, when drives do need to be accessed these storage arrays implement staggered spin-up techniques. This is to avoid overloading an array's power supply by trying to power up all its drives at the same time.

It is claimed that these power management techniques can be configured to reduce a drive's power consumption by up to 70 percent, without sacrificing performance. Higher reductions are possible, but may come at the cost of added latency and/or lower throughput.

Duplication duplication

There is more to using disks for backup than merely speed. A big advantage of disk over tape is that disk storage is random-access, whereas tape can only be read sequentially. That makes it feasible to reprocess the data on disk once it has been backed up, and as well as snapshots and CDP, that has enabled another key innovation in backup: de-duplication.

This is a compression or data reduction technique which takes a whole data set or stream, looks for repeated elements, and then stores or sends only the unique data. Obviously, some data sets contain more duplication than others - for example, virtual servers created from templates will be almost identical. It is not unusual for users to report compression ratios of 10:1 or more, while figures of 50:1 have been reported in some cases.

In the past, de-duplication has typically been built into storage systems or hardware appliances, and has therefore been hardware-dependent. That is changing now though, with the emergence of backup software that includes de-duplication features and is hardware-independent.

The technology is also being used for backups between data centres, or between branch offices and headquarters, as it reduces the amount of data that must be sent over a WAN connection.

D2D in branch offices and remote offices

There are many challenges involved in backing-up branch offices and remote offices. Who changes the tapes and takes them off-site, for instance? Plus, local data volumes are growing and more sites now run applications locally, not just file-and-print, so what do you do when the backup window becomes too small?

One possibility is to backup or replicate to headquarters, preferably using CDP or de-duplication technology to reduce the load on the WAN by sending only the changed data blocks. The drawback with anything online or consolidated is how long it takes to restore a failed system, however. Even if you have the skills on hand and a fast connection, it can take an enormous time to restore just a few hundred gigabytes of data.

D2D is the obvious next step - it can be installed as a VTL, so it functions the same way as tape but faster, but it also gives you a local copy of your files for recovery purposes. That local copy will probably answer 90 to 95 percent of recovery needs.

Add asynchronous replication to headquarters, and you can store one generation of backups locally with more consolidated at the data centre. Layer de-duplication on top, and there is less data to backup from the branch office and therefore less bandwidth consumed.

Consolidating backups at the data centre can bring other benefits too, in particular it enables information to be searched and archived more readily. It also takes the backup load off the branch offices as their backups are simply for staging and fast local recovery, so they no longer need to be retained.

Should the entire branch or remote office be lost, there are techniques to speed up the process of restoring a whole server or storage system. An example is the use of external USB hard drives, sent by courier and used to 'seed' the recovered system.

Even faster though are data-streaming technologies. This virtualises the recovery process, presenting the application with an image of its data and streaming the underlying data back as it is called for.