



Disk-to-Disk- Technology



Leveraging multi-tier
architecture improves
backup performance
and expedites restores

to-Tape

Backup, archiving, retrieval and restoring of data are important parts of every company's IT process. Employees mistakenly delete files and need them back, or they need retrieved copies of older information they thought they were done with. Servers holding databases or files may crash, or their disks may fail. And copies of data need to be archived offsite, in order to provide disaster recovery.

Tape has historically been used for backups, especially for long-term archival backup. This is due to its compactness, lack of electromechanical parts and comparatively low price per megabyte, gigabyte, and now, terabyte — typically about half that of hard drives.

Hard drives, however, have always been useful for backup, retrieval and restore, thanks to their intrinsic speed. And because, unlike tapes, it's possible to get data from any part of the hard drive nearly instantly, while tapes need to be linearly read. And while still more expensive per gigabyte than tape, the gap is closing.

For companies with ongoing backup, archiving, retrieval and restoral needs — which is all companies — the right strategy often means the difference between productivity, quick resumption of business and rapid response versus lost time, missed deadlines, unhappy customers and even monetary fines. With this in mind, should you be backing up from disk-to-disk (D2D) or from disk-to-tape (D2T)? The answer, increasingly, is both: disk-to-disk-to-tape (D2D2T). ▶

Disk-To-Disk-To-Tape: Using Each Where It's Best

A D2D2T backup solution combines both disk and tape as backup media, for solutions that offer IT the best mix of speed, flexibility, convenience and cost-effectiveness.

“Disk-to-disk and disk-to-tape are different,” says McClain Buggle, manager of vendor alliances, Software Alliances, Exabyte Corporation. “D2D is backup, D2T is backup and archiving. Tape is the only one that can handle the archiving side while disks have speed and performance.”

“Tapes offer price, performance and archiving,” adds Kelly Beavers, vice president product marketing, Exabyte. “To archive, and remove for long-term storage, tape wins.”

But D2D2T isn't just about media cost; it's about using each technology where it makes the most sense. “Combining the two pieces lets you optimize,” says Exabyte's Buggle.

A disk-to-disk-to-tape strategy starts with a disk-based backup device. It can be a separate device on your network. It can also be capacity on an existing network attached storage (NAS) system, plus a tape drive or library, either directly attached to the disk storage, or connected over the network, via iSCSI — plus backup/archiving software.

For bulk file transfer over a network, such as backups, copying to tape is inefficient, points out Mark Greenlaw, senior director, NAS Product Marketing, EMC. “The tape drive's speed is pretty high, but the challenge is to keep it streaming, which means no interruptions. To sustain that high speed means it's writing steadily.”

Any blip, whether in the network or because you're switching from one server to another, means the tape has to stop and restart. So any backup process will typically take several times the actual time to save that volume of information to tape.

Disk in the Middle

Once data from remote disks has been collected to the backup disk, it's available for quick online restoral and retrieval. It can also provide maximum-performance archiving to tape, transferring a stream of data to the tape. If the tape drive or library feeds directly off the backup disk device, there is no performance impact to the network.

If you're doing incremental nightly backups, tape presents another challenge, Greenlaw points out: “There's a fair amount of time being spent deciding if a file needs to be backed up. Consequently, there's a lot of time where the tapes aren't being streamed fully; they start and stop a lot. So the backup windows on these incremental backups tend to be longer than people would expect.”

It therefore can make sense to do the initial backup and aggregating to disk, says Greenlaw. “By streaming to disk, the data rates are similar, but the overall time is much better. It's these people who most want some disk environment. They may do the full backups to tape, but only do incrementals to disk. And 99 percent of restores are for an accidental delete. From disk, that's milliseconds, versus minutes for tape. And you can keep researching and reusing a disk with backups over and over, without worrying about a tape breaking.”

Disk's time-to-find speed compared to tape also means much faster retrievals, adds Greenlaw. “Tapes read on a stream. If the data is on the end of the tape, you have to wind to the end first. For a single-file restore, this might be a big delay. So it's hard to predict exact time to save or restore, where tapes are involved.” By having backup data online, on disk, IT can find it much more quickly, since disks are easier and quicker to search and don't need to be mounted (unless you're using removable disk cartridges, of course).

While, depending on how much backup data you have, you probably won't have everything kept on disk, you'll probably have at least the most recent backup, and other likely data such as e-mail. Tapes, being offline, also provide an additional layer of security for your data, notes Exabyte's Beavers. “Your backups are safe from predatory attacks by computer viruses or other intruders, because they're offline.”

D2D2T Gaining in Popularity

Disk-to-disk-to-tape is a hot topic, according to Randy Kerns, senior partner, at storage consultancy Evaluator Group. “Primary drivers are to speed up the backup process and make it more reliable.”

For companies reviewing their backup strategy, Kerns advises, “Look at the recovery time objectives and the recovery point objectives.” There's an economic value to getting it back faster versus days to get a tape back. The recovery point objective is about recovering data to a specific point that may make sense, for example, a backup ten minutes ago. You'll want that one; if the previous backup is three days ago, it may not be of value.”

For D2D2T, Buy, Don't Build

While it's possible to create your own D2D2T solution, you don't have to. Bundled solutions are available, such as the Exabyte/Snap Appliances/BakBone Software bundle or, for small offices, Quantum's DAT-based solutions.

“The three pieces all work together so you don't have to do trial-and-error troubleshooting,” says Exabyte's Buggle. “And all three companies support it. We understand how it works. We put the effort into testing the package at all three companies. Buyers get that compatibility right off the bat.”

And pricing for a bundled D2D2T solution may be better than buying the components separately.”

“A bundled D2D2T package can be particularly valuable for small- to medium-sized businesses,” notes Robert Amatruda, research manager, tape and removable storage, at tech research firm IDC. “In many SMBs, you find the same needs as small enterprises, but you don’t necessarily find the budgets or sophistication or time.”

D2D2T solutions are also easily scalable. “Our solution with Adaptec is very scalable,” states Exabyte’s Beavers. “You can add terabytes a disk at a time, and can go from LTO-2 [Linear Tape Open-2] to LTO-3 tape to double the capacity of your library, or add more drives.”

Virtual Tape Libraries

While saving initially to disk can add efficiency, many organizations have already implemented tape-oriented solutions. So does interposing a “layer” of disks require new development and more management?

It doesn’t have to. “VTL (virtual tape library) software lets D2D backup disks present themselves as tape libraries, so no changes are needed to backup software,” says EMC’s Greenlaw. “For example, EMC’s CLARiiON disk library is a virtual tape environment in which the disk drives appear to the server to be backup tapes.”

Where D2D2T Makes Sense for Your Business

Is D2D2T right for you? “It goes into need for backup,” says Beavers. “If you’re a 24-hour global shop, transaction-based, where a user can’t afford a single minute of downtime, or a restore/backup process of eight hours, disk becomes a mandatory item for you; tape becomes mandatory for offloading/offsite. And/or if data is distributed, to collect it, you need to be going to disk.

“Another good reason would be your frequencies of restores. If you have an environment with reasonable activity, such as a few events per day, D2D2T will add value. If those events are big — several hundred gigabytes, or even a terabyte — tape may still have the right performance. But only when you’re looking for a few tens of megabytes, or less, D2D will be better.”

Niche D2D2T Solutions for Archiving, Retrieving E-mail

D2D2T (disk-to-disk-to-tape) can also play a role in addressing specific concerns in a company’s data retention process — namely e-mail retention for regulatory compliance.

Already, financial institutions face regulatory obligations to retain e-mail relating to electronic transactions for typically up to seven years — and be able to produce these files on request from the courts or government agencies within days. Failure to produce these records can result in significant fines; even failure to produce them in the specified time frame can result in large fines.

Medical and pharmaceutical firms, government agencies and educational institutions are facing a growing number of similar requirements. And compliance with laws like Sarbanes-Oxley mean even more businesses need to be retaining e-mail in case of future demands to produce it. (And, of course, employees often need to refer to e-mail they deleted, accidentally or deliberately, weeks or months after the fact.)

Traditional e-mail server backup addresses the problem of full server restores, but isn’t designed for quick finding and retrieval of specific messages. Also, regulatory compliance includes requirements that the archive be time- and date-stamped and unalterable.

For small- to medium-sized businesses (SMBs) e-mail archiving solutions have been expensive and require IT time to buy, integrate and manage.

One recent solution, with a price-point affordable even by SMBs, is the Sony AIT Intradyn ComplianceVault, an all-in-one e-mail archiving solution. This package combines Intradyn’s ComplianceVault e-mail archiving/retrieval appliance, which comes in 1U rack-mount models, currently with 250, 500 or 800GB of disk space, plus any of Sony’s Advanced Intelligent Tape (AIT) StorStation 1U 400GB or 800GB AIT-2 or 1.6TB AIT-3 library tape autoloaders.

Once connected to the network — which can be a matter of ten or more minutes — the ComplianceVault begins capturing all user e-mail continuously. Intended as an archival, rather than backup medium, Sony’s WORM (write once, read many) tapes let data be appended, but existing data can’t be overwritten — suitable for providing a legal audit trail, says Brett Schechter, senior manager marketing, Sony Components Solutions Division. Depending on the configuration, package prices range between \$10,000 and \$15,000, significantly less expensive than other alternatives.