

Tape or Disk: Why Not Both?

Corrected Perspectives on Tape and Disk



January 2011

Contents

Abstract	3
Introduction	3
Slow Data: Disk's Advantage	3
Defining Slow Data	4
One Physical Stream per Drive: Tape	4
Multiple Streams: Disk	5
Slow Data to Disk, then Rapidly to Tape	5
Unrealistic Expectations of Disk	6
Reliability	6
Storage Media Costs: Tape is Cheaper than Disk, After All	7
Updating Perceptions of Tape	8
Tape is Reliable	8
Tape's Strong Performance and Capacity	9
Ease of Use	9
Corrected Perspective: When to Use Each	11
Conclusion	11

Abstract

Why did a perception arise in the storage industry that it's disk versus tape, not disk *and* tape? Articles about the issue suggest that a storage professional is required to pick a side: disk OR tape. Both disk and tape are very useful in data protection strategies, so long as their implementation takes advantage of their strengths.

Introduction

Both disk and tape were created for purposes aside from protecting data; both have moved into the data protection arena; and both look to be there to stay. The arguments that some disk-only vendors use to push consumers away from tape and to promote only disk use are built on a shaky foundation made up of:

- unrealistic expectations of disk (easy to use, fast, never fails)
- an idea (heavily promoted by disk-only vendors) that disk and tape cost the same
- outdated perceptions of tape (unreliable, slow, hard to use)

These claims rest on outdated evidence. Such claims both dilute the potent advantage and use of disk in secondary storage - very rapid access to data - and discount the long-term advantages of tape - cost-effective, easy to use storage that ensures unmatched longevity. This paper teases out the facts obscured by these biases. It supports an understanding of the value of each kind of media so that each can be used to its best advantage.

Slow Data: Disk's Advantage

Disk has a reputation for speed, which is well-earned—for specific uses. There are a host of factors that affect the speed of data transcription and retrieval. The following speed advantages go to disk:

- When backing up “slow data,” defined below, disk has an advantage due to its ability to handle multiple streams of data (slow or not) simultaneously. Disk is also more tolerant of a variance in the flow of the data stream.
- When accessing only a few files from a server, the random access method of data retrieval is much, much faster than the time it takes to wind the tape to reach the correct file.

Note that disk is slower than tape when it comes to streaming quantities of data anywhere, including to and from tape. To restore a lot of data fast, for example in the case of disaster recovery, tape serves more efficiently in most cases.

Defining Slow Data

Not all data is backed up equally, at least in terms of speed. Often, networks are set up with many small data sources. These include granular data from applications such as Microsoft mail, and unstructured files on an employee’s computer.

Moving live/active data to a backup/archive device involves several stages. The data needs to be gathered by the backup application and moved to a server that is host to the backup target. Any of multiple variables in the path of this process can introduce bottlenecks and prevent a steady flow by introducing latencies and interruptions to the data stream.

As an example, assume that a total of seven data streams are being backed up—some of the streams are made up of structured data that is handed off to a backup application by the host application; the rest are “unstructured” data – hundreds or thousands of files to be backed up.

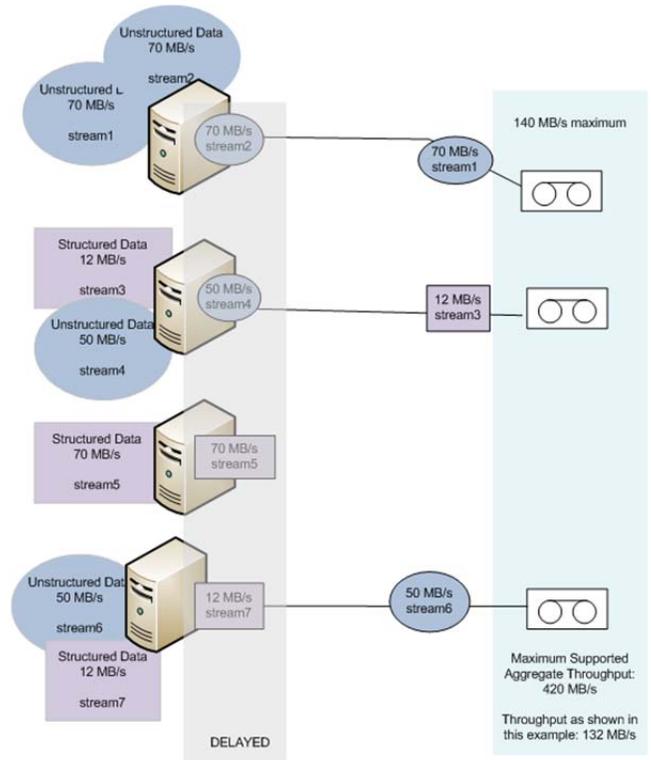
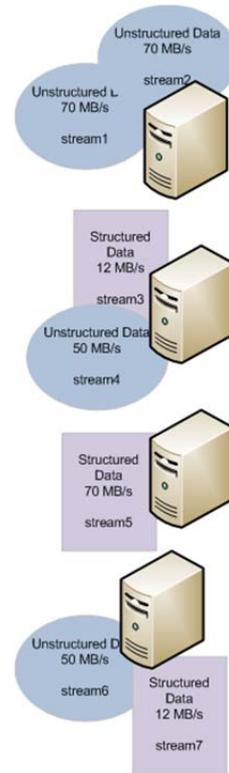
One Physical Stream per Drive: Tape

The *potential* transfer rate that data can be backed up is defined by performance of the target. If the initial target is tape, then the aggregate performance is the rate per tape drive times the number of drives. On a system with three LTO-5 drives, the total maximum throughput (each drive moving data at a maximum throughput of 140 MB/s) is 420 MB/s. However, as shown in this example, most systems can’t hand data to the drives at this rate.

Part of the reason data isn’t transferred rapidly is that each tape drive writes only one stream of data at a time.

Even when backup software interleaves (mingles) multiple data streams into a single stream, the interleaving itself slows data transfer. Interleaving data also penalizes restore performance significantly. Data restored from a tape with interleaved data takes a lot longer before of the constant re-positioning of the tape to locate the data. This is not tape’s strength.

To summarize, if each drive is being supplied data at a much slower rate (in this example, the rates per job are 70 MB/s, 12 MB/s, 50 MB/s), then the aggregate



Tape or Disk: Why Not Both?

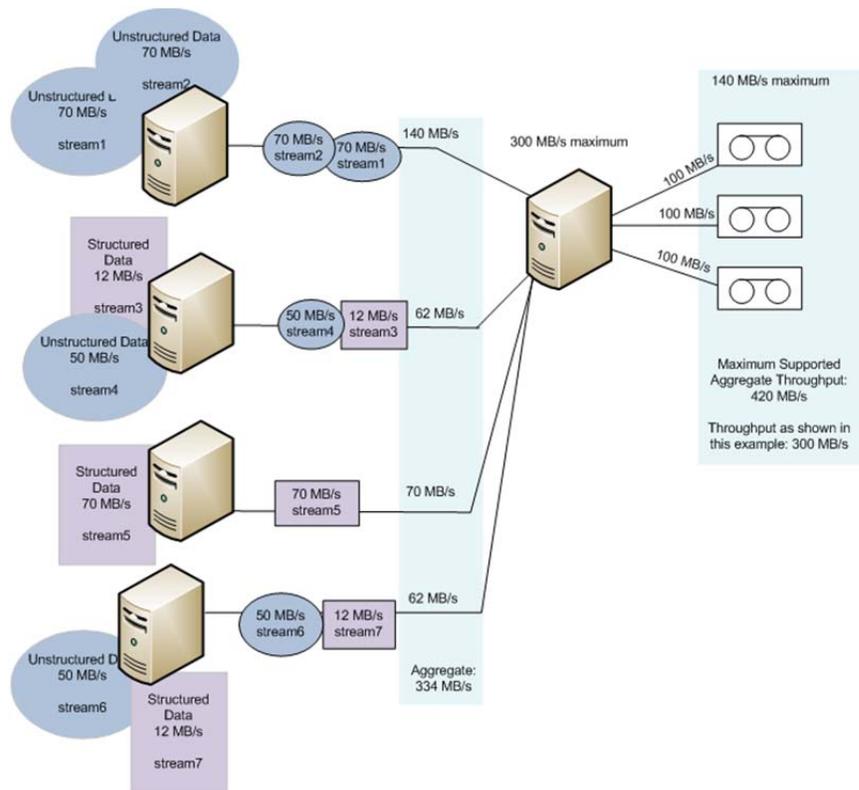
transfer rate of the three jobs totals just 132 MB/s throughput, well below the 420 MB/s the drives can potentially support.

Multiple Streams: Disk

This is where disk has an advantage. Disk has more flexibility in managing multiple streams of data that flow at different rates. If disk is presented as a file system, multiple files can be written at once independent of each other. If disk is presented as a VTL multiple virtual tape drives can be independently written as well. In this example, the aggregate speed at which data can be moved is 334 MB/s. Assuming the disk can accept and write data at 300 MB/s, the disk is well-used and the data at least in this example is moving nearly as fast as possible given network and backup constraints.

Slow Data to Disk, then Rapidly to Tape

Ultimately, slow data can be mitigated by using disk as an intermediate step. Using one disk server instead of multiple tape drives is like replacing multiple one-lane roads with a multi-lane highway. Just as each lane of a highway can support a car that is traveling at a different speed than cars in neighboring lanes, disk accepts multiple data streams regardless of the stream rate (up until it hits the disk's maximum capacity, at least). Data can simultaneously be written to disk in parallel.



In this example, disk handles all jobs concurrently, with only a little delay introduced once the disk throughput is maximized. Data moves to disk at or near 300 MB/s, faster than it can to tape, even though three tape drives support a potential throughput of 420 MB/s. Once the data is stored on disk, the formerly slow data can now be transferred rapidly to tape, since all the restrictions in the backup data path have been removed – that is, the data is logically closer to the tape target. Further, assuming that this is implemented without data interleaving by the backup application, data from tape can then also be speedily restored.

Some disk-only vendors insist the tape part of this set-up is unnecessary. Assume, for the sake of this example, that the storage administrator understands that disaster recovery best practices demand the use of tape for long-term, off-site, secure data protection. This point—that using tape for long-term retention and disaster recovery is critically important to organizational continuity— is reviewed in much greater detail later in this paper.

To determine if your site has slow data, examine the backup reporting to figure out how fast data is moving along during backup. If you have a lot of slow data, consider adding disk (such as Spectra's nTier line of disk) to your arsenal of tools for data protection. Tape still has a significant role in marshalling large quantities of data safely—but the tape drive performance can be increased through the use of disk.

Unrealistic Expectations of Disk

Widely held perceptions, created in part by heavy advertising and marketing paid for by disk-only vendors, portray disk as seamless, faster than tape in all cases, extremely reliable, and as cheap or cheaper than tape. Over the last ten years, as disk has moved into the secondary storage arena, its complexities as a backup medium are becoming better understood.

For example, if disk for longer term storage were as seamless to use as has been heavily marketed by disk-only vendors, then surely administrative burdens would be significantly reduced through the addition of disk as secondary storage. This hasn't happened. It turns out that that a single administrator typically is responsible for about 30 TB¹ of data stored on non-mainframe disk, while in a tape-based environment a single administrator typically manages up to more than 1 PB of data. This logically works out to about 33 administrators to manage data stored on disk compared to one administrator to handle that same amount of data stored on tape. While this is not necessarily a linear translation depending on the site, management tools, and data under management, the scale shows that data stored to tape is in fact easier to manage than data stored on disk.

In fact, disk-only hasn't been working out well for data centers. 68% of disk-only data centers plan to add tape back into the storage mix, according to a study² completed by the Fleishman-Hillard analyst group.

Reliability

Disk's reputation for reliability has also suffered in the past decade. Research at Carnegie Mellon³ showed that, "for drives less than five years old, field replacement rates were larger than [...] datasheet

¹ Moore, Fred. "Tape Technology Leaps Forward in the Third Era," *White Paper for Oracle*, June 2010.

² Fleishman-Hillard. "Over Two Thirds of Disk-Only Users look to add Tape Back into Storage Infrastructure," *White Paper for LTO Consortium*, March 12, 2008.

MTTF ... by a factor of 2–10. For five to eight year old drives, field replacement rates were a factor of 30 higher.” This means that data on disks is more vulnerable than data on tape, given that contemporary tape technologies, such as LTO, can store data securely for 30 years and more.

Data reliability of disk has improved 5x⁴ but still isn’t as reliable as tape. According to Harry Newman, a consultant on the high-performance computing and storage industries, “LTO tape has a much better bit error rate than enterprise disk, 1 bit in every 10⁻¹⁷ bits. **That is 40,960 times greater than enterprise disk, [emphasis added]** which is 1 sector (512 bytes) in 10⁻¹⁶ bits of data transferred⁵.”

Data protected on disk-only is further vulnerable to the network environment in a way that tape isn’t, simply because disk is online—hot. Disgruntled employees, viruses and malware, catastrophic disk failure such as array controller crashes, and natural disasters all can bring down a disk and make the data inaccessible. That includes mirrored data (data copied to an alternate set of disk) and RAID’d disk.

Storage Media Costs: Tape is Cheaper than Disk, After All

An idea heavily promoted by disk-only vendors is that disk and tape cost the same. Simply put, they don’t. Tape is significantly cheaper both in terms of acquisition costs and longer term cost of ownership. Virtually everyone agrees that tape is greener than disk if only because tape doesn’t use energy while it’s stored off-site in a vault. In an article focused on disk in CIO-magazine, staff writer Zach Church said “...disk storage generally costs about 3.5 times that of tape storage.⁶” The cost advantage to tape may be larger even than that, as shown in the following analysis⁷ by an analyst from the Clipper Group (note the different scales for each analysis):

³ Schroeder, Bianca and Garth A. Gibson, “Disk failures in the real world,” *Proceedings, FAST ’07: 5th USENIX Conference on File and USENIX Association Storage Technologies*, Computer Science Department, Carnegie Mellon University, 2007.

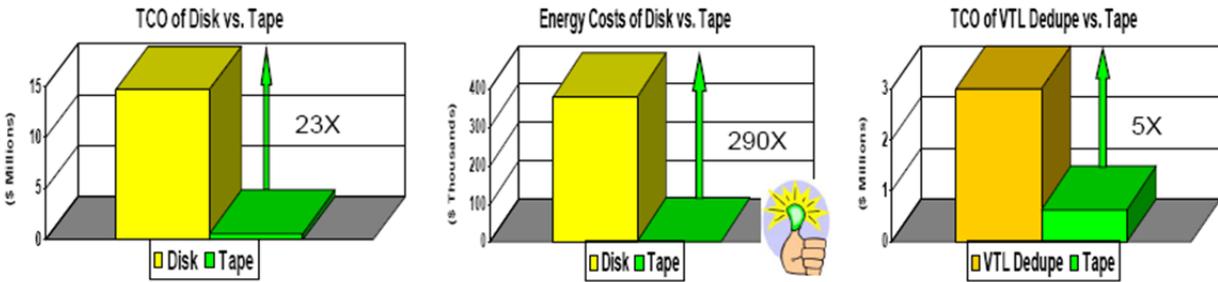
⁴ Moore, *ibid.*

⁵ Newman, Henry, “Why Enterprise Tape Can’t Get No Respect,” *Enterprise Storage Forum*, June 17, 2010 <http://www.enterprisestorageforum.com/continuity/features/article.php/3888366>.

⁶ Church, Zach. “Data deduplication making disk backup more realistic for midmarket.” SearchCIO-Midmarket.com, 23 July 2008.

⁷ Jelitto, Jens, *ibid.*

Tape or Disk: Why Not Both?



"Tape continues to provide the fiscal responsibility and functional value that enterprises require in the twenty-first century." The Clipper Group

The numbers vary wildly (that is, one article⁸ states that disk is 3.5 times the cost of tape, the second that disks costs up to 290 times the cost of tape), but all support a claim that tape is significantly less expensive than disk.

Updating Perceptions of Tape

Ten years ago, tape *was* technically weak. The reliability problems were significant, the backup software that wrote data to tape was less sophisticated, and tape was slow. None of that is true any longer. The advent of an open tape standard, Linear Tape Open, introduced a high-performing, highly reliable tape technology.

Tape is Reliable

Contemporary tape technologies are very reliable. Over the last decade tape reliability in terms of raw data integrity has increased⁹ 700%. Tape now automatically performs a data check on writing any data (the read-after-write feature); an enclosed tape path protects media from contamination; tape tension control through hub-lock technology allows for smooth writing and accurate transmission; writing guided by servo bands on magnetic tape is precise and fast; and many smaller advances have made tape in the 21st Century accurate and reliable.

Because of tape's very low bit error rate, tape is the choice of information experts at the Library of Congress¹⁰ involved with long-term data preservation. Tape's reliability is good for the long-term, as well. Even given the years when tape technology was not as robust as it is now, once data made it to tape, the data could be retrieved years later. Tape cartridges have a life of up to 30 years, and even

⁸ Moore, Richard L. "Disk and Tape Storage Cost Models." San Diego SuperComputer Center. San Diego, CA. Available through http://chronopolis.sdsc.edu/assets/docs/dt_cost.pdf.

⁹ Beech, Debbie. "The Evolving Role of Disk and Tape in the Data Center," *White Paper, Sylvatica Group*, 2009. Available through www.lto.org.

¹⁰ Kerr, Beth. Library Technician, Collections Care Section: Library of Congress. Email Correspondence. 1 Dec 2010. preserve@loc.gov

longer; tapes from the 1960s have been restored—after five decades of storage, the data remains accessible.

Tape's Strong Performance and Capacity

An updated perspective shows that tape performance and capacity now exceed most forms of disk. When backing up data, the bottleneck is the network servers, not the tape drive. At 140 MB/s maximum native throughput, LTO-5 tape and other high-capacity tape formats now are engineered to accommodate *slower* disk/networks, through features that include speed-matching levels and large buffers, so that tape can ratchet down to match the speed of data as it comes in. That way, data streams to the tape and tape stop/start action is minimized. LTO-1, released in 2000, supported throughput of 15 MB/s—with LTO-5, the improvement to 140 MB/s is a factor of 9.

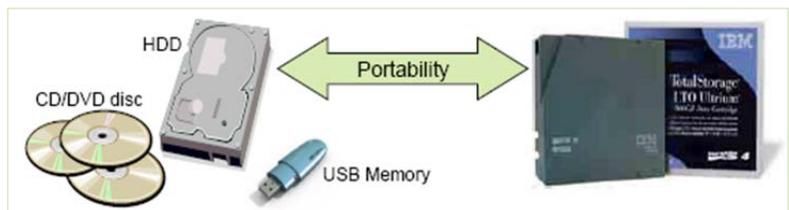
In fact, tape is so fast that it may be worth considering inserting disk between the initial data backup and moving data to tape. As described under slow data in an earlier section, the best use of tape and its single-stream relationship to incoming data is to set up an environment where data is aggregated on disk so that the data can clip along to the tape drives—fewer tape drives and better performance resulting in a shorter backup window.

Tape capacity has also exceeded that of most disk formats. Over the last ten years, LTO tape capacities grew by a factor of 15, from LTO-1's 100 GB to LTO-5's 1.5 TB (native format).¹¹

Ease of Use

To define how easy or hard tape is to use, first consider that tape itself is simply a medium. Its accessibility depends on the software interface and the library automating the tape. (Of course, the same is true of disk.) Advances in tape libraries like Spectra's also simplify tape management and integrity through the library's sophisticated software. Interfaces designed for efficiency and ease of use allow tape's speed and reliability to be used to their best. For example, with a Spectra library, administrators can track the health of tapes and of drives, minimizing interruption in the data protection process.

Data protection software has increased in sophistication and in ease of use over the last ten years. Backup software applications have become smarter and easier to use, and some now manage data regardless of where it resides—disk or tape or both. Filetek and FalconStor have some applications, for example, that let you set up an automated strategy that moves data across



¹¹ "LTO Ultrium Generations." November 28, 2010. <http://www.ultrium.com/technology/generations.html>.

tiers—from primary disk, to SATA disk, then to tape. With this flexibility, frequently used data resides on disk, permitting rapid data access. As the data ages or is used less frequently, the software moves the data to tape, providing long-lasting, reliable data storage, even if power goes out.

However, tape's ease of use is evolving so that it will soon approximate that of disk through a new data structuring system. Data stored on tape can now be presented in the same way data on disk is presented, through the use of a file-system front-end. This data access method is codified as the Linear Tape File System (LTFS) Format¹², announced in April 2010. Even better, the LTO consortium (the group that defines LTO standards) has made LTFS an open source file system.¹³

The LTFS standard builds on the LTO-5 partition feature, providing a file system for data stored on tape. The file system information, also referred to as metadata, is stored in one partition, the data in the other partition. Using a typical directory structure as a front end on tape makes it very simple to get to data stored on tape. By making data on tape as easy to find as data stored on disk, the media on which the data is stored becomes less relevant. Users no longer need to use backup applications to identify the data to be restored, restore it, save it to disk, and then access it. They can now click on a file. Essentially, tape becomes a portable media that provides easy access to data, as shown in the preceding illustration¹⁴. The opportunity to access data with only a brief delay, around 75 seconds, lets administrators have all data accessible, all the time. It removes the considerable barrier of involving IT staff and the time they need to open the backup application and pull the data.

Having all data on disk is, as one analyst¹⁵ put it, like having the car always running so you can jump into it and go instead of starting it up each time. With this front end, tape now provides a sensible alternative to having everything instantly available, including data that is virtually never accessed. Now you can put that data on tape and still access it with only a brief delay. This makes it possible to keep all of an organization's data accessible, independent of the media it resides on.

Tape libraries, such as Spectra tape libraries, further simplify the use of tape. Along with robotics to handle tape movement into and out of drives to write and read data, Spectra T-Series libraries provide a slew of features that assure data integrity, track tape drive and media health, and provide simplicity of use through a single graphical interface. It permits users to manage the library from anywhere through a web browser and through the library operator panel. By integrating features into a single, graphical, icon-driven interface, no additional hardware or software is required to set up, maintain, service, and

¹² "Linear Tape File System (LTFS) Format Specification," version 1.0. April 12, 2010. http://www.trustlot.com/LTFS_Format_To%Print.pdf.

¹³ Ferelli, Mark. "LTO-5 and LTFS: Shaking the Pillars of Heaven." Editorial, *Computer Technology Review*, August 17, 2010.

¹⁴ Jelitto, Jens, Mark Lantz et al. "Magnetic tape storage advances and the growth of archival data," *Proceedings of the First International Workshop on Standards and Technologies in Multimedia Archives and Records (STAR)*, Lausanne, 2010. http://mmspl.epfl.ch/webdav/site/mmspl/shared/star2010/ppt/star2010_jelitto.pdf

¹⁵ McAdams, Diane. Clipper Report

learn. Given the savings in energy use, purchase price, higher reliability, and long-term retention capabilities, including tape seems the best future for long-term storage.

Corrected Perspective: When to Use Each

By updating the perspectives on tape and disk, administrators are freer to use disk and tape to their best advantage. Disk gets to data fast. Tape stores data securely, for a long interval, and cost effectively. New technologies make it easy to access any data in a tape library.

Use disk for data the organization frequently uses. For cases as disparate as e-discovery, rapid data retrieval for scientific research, and online structured data, disk is hands-down the best choice. It's fast. You can cut costs using deduplication technology, and migrate data efficiently to less expensive storage using the right storage application. Have the data you need at your fingertips.

Use tape for less used data. You can still access the data through active archive applications and file system front end applications like LTFS, described earlier in this paper. Use tape for data that has long retention periods, and always for disaster recovery and organizational continuity. It is the most secure and long-lasting storage media available, period.

Disk and tape both have a big role in protecting data. Once the misperceptions are dispelled, the answer is obvious: use both disk and tape.

Conclusion

Technology is hard enough to keep up with—and it's even harder to do so when administrators are bombarded with industry clamor about the latest and supposedly best technology that can obscure simple facts. This white paper dispels misperceptions of both tape and of disk, and helps users identify the real strengths of each type of media. Once that is clear, it's easier to decide when to use disk and when to use tape, and also that it's wisest to use both.

Tape or Disk: Why Not Both?



www.SpectraLogic.com

Spectra Logic Corporation
6285 Lookout Road
Boulder Colorado 80301 USA
800.833.1132
303.449.6400

For more information, please visit www.SpectraLogic.com