

Commentary

December 12, 2005

Tape: An Ongoing Bulwark for Data Protection

Face up to it and plan accordingly. Tape is here to stay for data protection. Disk will have an expanded role in data protection, but tape will continue to use its strengths to complement disk. Disk is good for front-line data protection defense for high availability and low short-term data loss risk. Tape supplies long-term data preservation that is a mandatory part of any data protection policy for an enterprise. To meet those tape needs, the LTO Program offers a well-accepted tape technology with a roadmap for the future. This open tape technology will continue to serve the interests of IT organizations for open systems tape for the foreseeable future.

The Challenge to Tape

Tape — by its sequential nature and the fact that it is not the active production copy — is of necessity a lower availability data protection solution than a four or five nines availability solution. And many organizations, for good and sufficient reason, have that level of availability burned into their service level agreements (SLAs).

And that requirement of availability (as well as low short-term data loss risk) is leading to the greater use of disk-based data protection both as a first line defense (such as mirroring and continuous data protection) and a second line defense (such as disk-based backup including virtual tape libraries). The greater attention paid to and the expanded role of disk has led some pundits (and vendors) to proclaim that disk will replace tape.

Nothing could be further from the truth.

Setting the Case for Tape

On the surface, the inevitability of tape is not intuitively obvious. Let's take a zero-based approach to understand why if tape had not been invented, it would have been necessary to invent it.

The discussion starts with data preservation — the second of the key data protection principles. One measure of data preservation is recovery point objective (RPO) — the amount of data that is exposed to permanent loss. SNIA suggests 1 minute in a five nines environment and 10 minutes in a four nines environment.

However, the RPO represents only the first fallback design point to which the data protection infrastructure can respond in the event of a failure. Additional failures may expose all data to the risk of permanent data loss. For example, the RPO for a single failure in a

RAID 5 array is zero. However, a second failure in the array before the first failed drive is rebuilt represents the potential loss of all the data (which is why there is backup).

And complete (or nearly so) data preservation, which is tape's strength, is mandatory (i.e., long-term business survival depends upon it), whereas availability is only vital (the risk of revenue and market valuation loss has to be balanced with feasibility and costs to achieve).

The argument for tape will revolve around why disk-based data protection cannot achieve all the necessary goals for data preservation technically, economically, and process-wise.

Why High Availability All the Time Is Unrealistic

High availability is possible under typical data protection scenarios. It may not be feasible in edge case scenarios, such as a regional disaster. Data protection also has to shield against uncommon situations where the probability of a potential problem may be low, but the potential loss would be huge.

Take the best case in a disaster situation where an enterprise is able to afford a fully-equipped remote spare data center and make either a hot or warm failover after the predicted lengthy loss of its original production data center. High availability has been preserved, but the enterprise is at risk because it is down a data center. That risk is not only the loss of availability, but also the risk of the permanent loss of the information that is vital to the business's long-term survival.

Reconstruction, Not Just Restoration

And no organization can afford to keep one more than absolutely necessary data centers just in case. What it needs is a "virtual" data center that is affordable and tape media provides that capability.

When all else fails, tape media provides the data building blocks upon which working applications in an IT infrastructure can be reconstructed. Tape media represent a "virtual" data center that can be translated into a real data center as needed. The location of the real data center can be determined after a crisis has occurred so that it can be put in a non-affected location.

The media properties that are essential for this functionality are removability and portability. That means that the media can be taken from its creation point and taken physically to where it is needed for either a virtual or real data center.

Tape buys both time to reconstruct a real data center and peaceful sleep at night knowing that tape represents an emergency "spare tire" as a "virtual" data center.

Process and Technology Diversification

No rational investment manager would be caught without a balanced portfolio as forecasts for any one type of investment vehicle are highly uncertain. Even if by chance an organization could afford to have enough degrees of protection through disk alone, it would not be a wise idea.

For example, all the disk mirrors imaginable would not prevent against the propagation of logical corruption swiftly to each and every disk mirror. Continuous data protection systems can defend against logical data protection problems, but are subject to the same physical

failure as any other fixed disk system. Tape provides protection against both physical and logical data protection problems in extreme situations where, for whatever reason, disk systems are unable to meet the need.

When one uses a common process and technology, such as disk, a data protection system can be exposed to a systematic or a random problem, such as a human error that is not detected until too late, but still has a far-reaching impact. For example, replication from a full to an empty array can backfire if the replication is accidentally done the wrong way (as has happened). And that error doesn't have to be only manual. Automation is the computerization of human processes and errors can creep into automated processes (c.f., software) as well.

Tape offers a different management process and technology that offers a fallback position in case of emergency. Yes, the backup/restore software management process with tape may not always be a popular one, but it is a safety measure. No, absolute certainty is never possible, but having two separate and distinct processes adds a margin of safety.

I/O Isolation — an Extra Measure of Safety

At least one data protection copy should be safe from inappropriate write I/Os, say from an accidental "delete all" command from a super-user, who thought that the command was being applied locally and not globally. Tape media outside of tape drives in slots of a tape automation system or stored offsite comply with this dictum.

Disk-based Backup Is Not a Panacea

But what about disk-based backup, such as the use of a VTL? Say two copies of data are kept at both the production site and a spare site. Isn't that enough?

The answer is not the number of copies, but the fact that there is no offsite storage of media to act as the third leg in a data protection stool to act as the virtual data center.

Moreover, while there is some process diversification, there is no technology diversification (everything being disk) or real I/O isolation (as the disk is technically online).

Removable Disk Is Emulated Tape

What about removable disk? Don't confuse fixed disk, which serves as production data, with removable disk, which serves as a data protection copy. Removable disk is "tape" in a different form factor.

The question of scale comes into play. Tape solutions scale well (and "carbon copies" of a particular backup can be made). Removable disk would have to assume the characteristics of tape; including compression and sequential access, and automation using the same management processes (including media management). Enterprises would ask themselves why not use tape, which also provided technology diversification and is already automation-friendly. The answer might be cost, but for scalable solutions tape still has a cost advantage over both removable and fixed disk.

Doing Tape Well

To recap, the purpose of tape is to ensure the long-term data preservation of all of an enterprise's data even in the face of an extreme challenge. For dealing with

extremes, fixed disk in fixed locations is geographically-challenged in addition to not providing the requisite variety in both process and technology. Yes, disk works well the vast majority of the time, but tape is there as the last bastion of defense.

Doing Tape Right

But tape cannot be done halfheartedly. Tape management disciplines, such as media management to keep track of what information is available on both onsite and offsite pieces of media, has to be in place. At a minimum two copies have to be kept — one onsite and one offsite. And even starting from a zero-based approach, an enterprise may very well come up with the tape rotation strategy that is currently in place (except for perhaps one copy that is held in a disk-based backup).

Enterprises can turn to a number of well-seasoned tape technologies for their tape requirements. The technology towards which many enterprises are turning comes from storage vendors that support the LTO (Linear Tape Open) Program-developed specifications. Those specifications cover both tape drives and the tape media itself for what is called LTO Ultrium technology.

Four companies supply tape drives that meet LTO Ultrium specifications as follows:

- Quantum
- HP
- IBM
- Tandberg

All four companies have a long history of success in the manufacturing of tape drives.

In addition, IT buyers have a choice of buying LTO Ultrium tape media from a number of established suppliers as well including:

- FujiFilm
- Imation
- Maxell
- Sony
- TDK

The objective of any open technology is interchangeability and compatibility of products. With LTO, any certified LTO tape cartridge can work with any certified drive of the same — at least — LTO tape generation.

Although IT buyers typically would buy tape drives from one supplier and tape media from one or two tape cartridge suppliers, those buyers have the peace of mind in knowing that they could switch if they had to for whatever reason. They know that competition among the suppliers of both the drives and the media will continue to drive quality, availability, price, and the development of the next generation of product. LTO technology offers compatibility between generations. It is designed to read back two generations and write back one helping users to ease implementation and protect investments.

LTO Ultrium 3 — Popeye after Eating His Spinach

LTO Ultrium is now on its third generation. LTO Ultrium 3 has a native capacity of 400 GB per tape cartridge and up to 80 GB per second native performance. Compression at an assumed ratio of 2 to 1 is typically used with tape, so IT buyers can plan on actually getting double the native numbers. What that translates into — as an illustration — is that the entire 19 million books in the U.S.

Library of Congress could be contained on about 12 Ultrium 3 cartridges and that the movie *Finding Nemo* could be stored in about half a minute.

LTO Ultrium — Working to Make Tape Even More Reliable

Tape drive manufacturers continue to strive to make tape drive technology even more reliable than it has been. One of the key reliability development efforts in LTO tape technology involves servo tracking mechanisms that improve the precision tracking of the head and media in order to help ensure accurate reads and writes of data.

Another is read-after-write verification, where data is immediately read after writing. That not only helps ensure that the data was written correctly in the first place, but also should increase a system administrator's confidence that the data can be restored if and when necessary.

To further increase reliability, LTO technology uses advanced microcode to detect and correct errors. This capability is important for reducing the number of possible failures when trying to restore information from tape.

LTO Ultrium Is Up to Speed

An issue with linear tape technology prior to LTO Ultrium Generation 2 was back hitching. Back hitching, also called "shoe shining," is the stop and go action of a tape drive when the tape drive cannot match the speed at which it is receiving data to its tape speed. Stopping tape that is moving fast dead in its tracks and then restarting again and again is not good for performance, nor for the reliability of a tape drive and its associated tape cartridges.

Starting with LTO Ultrium 2, all generations of LTO technology have (or will have when available) speed matching, where the drive can sense how fast it is receiving data and take the proper action to avoid back hitching. Speed matching not only improves performance, but also reliability. Another feature added with LTO Gen 3 is the expansion of the buffer from 64MB to 128MB. The larger buffer helps to keep data streaming to the tape.

WORM — Getting an Extra Use from Tape

Using tape technology for storing compliance data is a good idea for a number of reasons including data protection and low archival costs, and especially where tape technology is already in place. LTO technology offers a WORM (write-once, read-many) formatted cartridge. An LTO WORM cartridge helps to address compliance regulations for being non-rewritable. Data can be appended to a partially written WORM tape, but parts of the tape that have been already written to are unalterable. A WORM LTO cartridge is two-tone to enable easy visual identification to distinguish it from a standard non-WORM LTO cartridge.

LTO Ultrium — For All Generations

The LTO Program has already demonstrated the robustness of its technology, not only in the first and second generations, but now in the third generation of the technology.

The LTO roadmap now extends into the next three generations 4, 5 and 6. Enterprises can be comfortable knowing that LTO Ultrium is a "living, breathing" technology that will continue to evolve as rapid data growth continues.

Commentary

Conclusions

Tape will continue to be a best practice for enterprises as part of their data protection infrastructure portfolio. Mobility (i.e., the combination of removability from a tape system and portability, which enables physical transportability over long distances) is not an option, but rather a necessity for the virtual data center. In addition, process and technology diversification as well as I/O isolation provide those extra safeguards that may prove all the difference if called upon.

So enterprises should maintain their tape management disciplines and not worry about how to replace tape, but rather how to put together a balanced data protection portfolio where disk and tape complement each other in the appropriate mix. And the LTO Ultrium technology, with shipments of over 1 million tape drives and 30 million tape cartridges, points the way with demonstrated great success in the real world.

The bottom line is that tape is an insurance policy that covers an enterprise against a risk that it cannot afford to take.

David Hill

Analyst Name: David Hill
Topic Area: Tape

Mesabi Group LLC
26 Country Lane
Westwood, MA 02090
www.valleyviewventures.com

Mesabi Group LLC is an affiliate of Valley View Ventures that aims to provide thought leadership and sound advice to both vendors and users of information technology.

Phone: (781) 326-0038
email the author: davidhill@mesabigroup.com

The information contained in this publication has been obtained from sources Mesabi Group LLC believes to be reliable, but is not warranted by Mesabi Group LLC. Commentary opinions reflect the analyst's judgment at the time and are subject to change without notice. Unless otherwise noted, the entire contents of this publication are copyrighted by Mesabi Group LLC, and may not be reproduced, stored in a retrieval system, or transmitted in any form or by any means without prior written consent by Mesabi Group LLC