

Is Tape Backup Dead? We Expose the Common Myths and Rumors.

Trust NovaStor DataCenter for Fast, Reliable Data Security

Introduction

People have been proclaiming the death of magnetic tape for decades. There is a recurring tendency in some IT circles to declare backup to tape as obsolete or near death as 'the cloud' increasingly makes its way into the data centers of the world, as a backup destination. But even in these times of private cloud services, SAN snapshots, SSD's, and cheap public cloud storage we shouldn't be too quick to lower the coffin. The rumors circulated by cloud service providers and disk vendors might have you believe that tape is slow, unreliable, or simply outdated. Perhaps you even have a personal experience, where a poorly implemented system was inherited or you got stuck using archaic backup software. Whatever the reasons, if you have preformed opinions, we ask that you temporarily suspend them as we tackle the most common misconceptions surrounding tape as a backup medium. Tape does not fit all situations and may not be the ultimate storage solution, but it may still have a few tricks up its sleeve when it comes to the long term retention and archival storage of your critical data.

Today we tackle some of the most popular, and highly circulated myths about tape.



#1 Tape Is Dead

Whoops, who said that? A global survey conducted by Kroll Ontrack in 2017 showed that while cloud storage use had doubled over the previous year, magnetic tape has also experienced an impressive boost in popularity. It is predicted that with explosive data growth and the decreasing cost of high-performance storage, SSD and flash technologies will come to serve as the primary storage location for software applications while hard disk drives and tape storage move into archiving and long-term storage roles.

Among our own NovaStor's DataCenter customers, approximately 50% are still using tape devices like tape libraries in some form. For example; To back up data for retention purposes in a remote, off-site location (bank vault or similar).

Thanks to the rapid rise of ransomware, cyber-crime and other types of unauthorized surveillance, thousands of companies are revising their privacy and retention policies, and re-evaluating data storage locations. Can it be argued that having all your data held in a public cloud for long term data retention is safer and more secure than a tape locked away in a vault where only you have the key?

#2 Tape Is Slow

Just take a look at the raw numbers. With regards to speed you could stream data at 140MB/s when LTO-5 was released in 2010, and looking at the current generation 8 years later, LTO-8 can now stream data to tape at 360MB/s. Think about the disk array or NAS you would need in order to be able to handle 300MB/s+ backup streams going to it as a backup target, you better have some SSD caching going on that target to handle that type of speed at the very least and minimally 10Gbe connectivity to the nodes you are backing up to that target. We know what you are thinking, "Well if I have 10Gbe connectivity I can push a lot more than 300MB/s across that link, and now my tape drive won't keep up". The great thing about tape is if you need more speed, all you need to do is add another tape drive to your tape library (assuming you have space for another tape drive).

Another speed related misconception is how tape is "only good when sending large files to tape", and how "sending large numbers of small files causes tape back-hitching which causes slow speeds and drive/tape wear". Historically this would be correct as 'back in the day' backup



software was not designed to handle multiple streams of small files in a way to stream to tape at an optimized speed, instead the backup software that was being used would just do a single stream to tape file by file and at a very low speed. Modern enterprise backup software such as NovaStor Datacenter, Commvault, Netbackup, and Arcserve all have options to handle multiple streams of small files in order to keep tape backups streaming at the speed they need to be.

#3 Not Enough Capacity

We all know that our backup sets keep growing. As a sysadmin it seems that regulations such as HIPAA, PCI, GDPR and others are asking the impossible from you with regards to data longevity and retention requirements. This leads many to believe that their only choice for keeping up with growing data capacity, is to ship everything to some cloud storage or purchase additional local and remote storage.

Administrators may not realize that the current generation of LTO tapes store 12TB of uncompressed data per tape, and even the less expensive previous generation can accommodate 6TB uncompressed to tape. To put this in perspective, for well under \$200 with currently recently released LTO-8 tape cartridge you have 12TB of uncompressed space that you can write to at 350MB/s. Also, understand that 12TB does not include any compression benefits you would achieve, While achieving the 30TB of compressed data that manufacturers advertise is quite unlikely, getting 20TB of data storage with compression while backing up regular business data is actually quite achievable.

#4 Won't Scale

There are, in our experience, people with two very different visions of what a tape library looks like and how it can function in an enterprise backup environment. One type of person has experience with tape libraries that may take up a full rack of space. Their configuration may include six or more tape drives in a unit with hundreds of tape slots. This type of person knows that tape can scale to accommodate a very large and diverse network. And there are people whose only experience with tape is a single standalone tape drive or perhaps an old clunky tape library that had one tape drive with 8 tape slots in it. It was likely running an outdated version of backup



software that was new when Windows 2000 came out. They had nothing but trouble with tape and have sworn off tape as something archaic, problem bound, and frankly not worth the hassle.

These are the people that we talk to that are surprised to learn that using the right software a tape drive can handle backups from multiple machines at the same time, and be used as the final medium that is taken off-site for their disk-to-disk-to-tape strategy for multiple servers. The tape libraries of today have dramatically improved over the technology available 10+ years ago. With a long hardware refresh cycle, some companies may still be using that 10+ year old library so it's hard to fault them for this perception.

#5 Cost Prohibitive

This point may be the biggest misconception when talking about the tape medium for backups, price and total cost of ownership. When correctly structured, tape is relatively inexpensive when compared with other storage mediums.

To illustrate this let's look at a simple example for ease of understanding;

A fictional company, Acme Corp., has decided to keep 7 years of data in a simple GFS (Grandfather-Father-Son) type configuration. They have 5TB of data they need to backup and it will not grow, and daily changes are 500GB. Their backup plan consists of the following:

- Keeping 7 yearly backups (Full)
- Keeping 4 quarterly backups (Full)
- Keeping the last 3 monthly backups (Full)
- Keeping the last 4 weekly backups (Full)
- Keeping last 4 weeks worth of daily backups M-F (Incremental)

In summary, 7 yearly, 4 quarterly, 3 monthly, and 4 weekly backups translates to keeping 18 full backups at 5TB each, so 90TB worth of data. Then 4 weeks worth of daily backups at 500GB each which is another 10TB of data. That means they require 100TB of storage space in order to fulfill their backup plan.

Now let's compare costs between a disk-based and tape based solution. Please note that prices change, and that the infrastructure required to support the solution is not included.



For a disk based solution at current prices, the 8TB hard drives are at the best price point. Each 8TB hard drive (whether Western Digital or Seagate) currently cost about \$250 USD for each drive, and we would need 16 of them in order to achieve a 15 drive RAID6 for about 100TB of usable space and also having a hot spare. Then for an iSCSI target to put the 16 drives in, lets assume we're going with a QNAP device that can handle 16 drives for \$4500. The total disk based solution would be about \$8500, which honestly for 100TB of usable space is not bad at current prices. Though we must consider that the speed of this type of setup would probably rarely break 100MB/s.

For a tape based solution we would go with a LTO-7 generation tape library. At the time of the writing of this document, LTO-8 has recently been released, but libraries and drives are not quite readily available to be purchased. LTO-7 tapes also give us 6TB of native capacity, and to be fair in this scenario we will say one tape per full backup (though the daily backups can span more than just one tape). For the library we will go with a Quantum i40 that has a single SAS LTO-7 tape drive in it and 25 slots at a price of \$4800. We need a minimum of 20 tapes, but to give some excess room we will go with 24 tapes to fill up the library which already comes with a cleaning tape. A 20 pack of LTO-7 tapes cost \$1450 and then 4 more at \$75 each. That would be a total price for the tape solution of \$6550. This solution would have a total uncompressed storage capacity of 144TB, and has the capacity to be written to at 300MB/s. This offers us almost \$2000 in savings that could be used to purchase additional disk storage for the server that is attached to the tape library, giving us a disk backup location for a disk-to-disk-to-tape type backup configuration. Not bad!

#6 Unreliable

LTO tape is designed for 15 to 30 years of archival storage and has almost two decades of history behind the specification and architecture of the medium. What hard drive, disk array, or cloud solution can say the same? Then there is the relative longevity of the medium itself when compared to spinning media or even solid-state drives. While spinning hard drives, can hold data on their platters for decades, other parts like the hard drive head or spindle are [a much bigger question](#) over decades of storage. Consider what might happen if someone were to accidentally drop or knock a HDD while it is writing data. When talking about using SSDs for long term data retention you must consider the possibility of drives being unpowered



for extended periods, and consider how this might affect the data. Magnetic tape is specifically designed for long term retention and archiving, so as long as it is stored correctly and the lifespan of the tape is taken into account, your data is safe and sound.

Summary

NovaStor works side-by-side with System Administrators to implement modern, efficient backup solutions that utilize the best technologies for the client's unique backup scenario. It is important to fully analyze the speed, capacity, scalability, and reliability, and budgetary requirements when designing backup solutions and not get persuaded into using the "latest and greatest" options without carefully considering tried and trusted technologies. While magnetic tape storage is not a fit for all situations, it has come a long way in recent years and maintains an undeniable ability to persist in an industry of constant change.

Any Questions?

If you'd like to learn more about the speed and scalability of NovaStor DataCenter, or begin testing in your environment, please contact our Solutions Team directly at:

Tel.: +1 (805) 579-5445

Monday-Friday 9AM-5PM Pacific Time

Email: DCInfo@novastor.com



About NovaStor

NovaStor® (<http://www.novastor.com>) represents "Backup for the Rest of Us" by empowering overwhelmed and underfunded IT administrators with all-inclusive, fast, highly scalable, budget sensitive data backup solutions for both physical and virtual environments. NovaStor's disruptive approach redefines service by including personalized local, expert level professional services as part of every solution - helping ease the enormous expectation being placed on maintaining a working, compliant backup under even the strictest of budgets. NovaStor's global references range in size from small highly regulated verticals to F500 environments with thousands of servers and PetaBytes of data.



- 📍 NovaStor Corporation
29209 Canwood Street
Agoura Hills, California 91301
- ☎ Tel.: (805) 579-6700
Fax: (805) 579-6710
- ✉ Mail: sales@novastor.com
- 🏠 www.novastor.com



Copyright © 2018 NovaStor Corporation. All rights reserved. "NovaStor" and the NovaStor logo are registered trademarks of NovaStor Corp. Windows is a registered trademark of Microsoft Corporation. Other names may be trademarks or registered trademarks of other holders of rights. Technical changes, deviations from the illustrations excepted.



NovaStor Corporation
29209 Canwood Street
Agoura Hills, CA 91301 USA
Tel +1 805 579 6700
Fax +1 805 579 6710

NovaStor Software AG
Poststrasse 18
CH-6301 Zug
Tel +41 41 712 31 55
Fax +41 41 712 31 56

NovaStor GmbH
Neumann-Reichardt-Str. 27-33
D-22041 Hamburg
Tel +49 40 638 09 0
Fax +49 40 638 09 29



Microsoft Partner
Gold Application Development

www.novastor.com