



Migrate Your Existing Tapes & Tape Drives to Spectra's Stack Modern Tape Library

May 2019




Contents

Hardware Migration	3
Migration of Existing LTO Tape Drives.....	3
Working with various levels of LTO tape drives in a tape library	4
Migration Paths – Data Sets	5
Which data tapes need to be migrated and which do not?.....	5
Migration Paths – Library, Drive, and Media.....	6
Soft Migration	6
Hard Migration.....	8
Thoughts on Future Migration.....	10
Summary.....	12
Appendix	13

Copyright ©2019 Spectra Logic Corporation. All rights reserved worldwide. Spectra and Spectra Logic are registered trademarks of Spectra Logic. All other trademarks and registered trademarks are the property of their respective owners. All features and specifications listed in this white paper are subject to change at any time without notice. Reproduction without permission is prohibited.





Is the data or digital content you're writing, creating, or capturing going to last longer than the media it's written to? The majority of every organization that works with such information or content will answer, "Yes."

That means the information will have to be moved from the original technology it's stored on to a newer technology – welcome to digital migration. The life of disk in production is roughly 3 to 4 years. LTO tape cartridges have a shelf-life of 30 years, but that also requires a tape drive capable of reading the tape 30 years later. As data and digital content become ever more important in our modern society, so does the need for a modernized approach to data and content migration – the subject of this paper.

Hardware Migration

Migration of Existing LTO Tape Drives

Typically, migration conversations revolve solely around the migration of data. Using the Spectra® Stack Tape Library, the conversation begins around the movement of the actual LTO tape drives. There are a few manufacturers who allow users to move LTO drives from their existing libraries into a new library as long as the manufacturer in question provided both the older library and the newer library the drives are being moved from and to. Spectra is unique in allowing LTO drives to be moved from a library *not* originally manufactured by Spectra into a new Spectra Stack tape library.

The ability to re-purpose your existing tape drives offers tremendous investment protection for users, and as will be shown below, offers greater versatility and ease in migration. Spectra Stack accepts LTO-5 through LTO-8 tape drives. Most existing tape drives to be moved into a Spectra Stack library will fall into one of two categories: Straight Move or Sled Move.

Straight Move – Straight move allows the user to move LTO tape drives directly from their existing tape library into the Spectra Stack library. This can be done by the user (with phone support from Spectra if desired), or this can be done onsite by a Spectra Professional Services representative. It's a simple process requiring no new or additional hardware. See appendix 1 of this paper for a listing of known drives supported with Straight Move.

Sled Move – Any tape drive used in library automation requires a "sled" to mount into the library. The sled performs power regulation, control electronics and adds airflow via a fan(s). See figure 1. Not all tape drive sleds are compatible with the Spectra Stack library. If users are working with LTO tape drives which don't have a compatible sled, this can easily be remedied in most cases. Spectra offers a sled kit for most LTO drives that don't have a Spectra Stack-supported sled. See appendix 1 of this paper for a listing of known drives supported with Sled Move.



Figure 1: An LTO half-height drive sled and LTO half-height drive are shown above.

Again, changing the drive sled is a simple process which can be done by the user or done onsite by a Spectra Professional Services representative.

There are some models of LTO tape drives for which there are no compatible sled units available. In these instances, Spectra offers discounted refurbished and factory-warranted LTO tape drives to replace your existing LTO drives.

Your Spectra sales representative can provide a full assessment of your existing hardware and create a customized migration plan for your existing hardware.


Working with various levels of LTO tape drives in a tape library

Historically, LTO tape drives were able to read back 2 generations and write back one generation. LTO-7 is able to *read* LTO-5 and LTO-6 tapes. It can also *write* to LTO-6 tapes. LTO-8 has incorporated a new type of head technology referred to as Tunneling Magnetoresistance (TMR) technology. This helps assure a strong future roadmap for LTO. The LTO Consortium now shows a roadmap extending through LTO-12. Unfortunately, this change in head technology means that LTO-8 is only backward compatible one generation for *both* reads and writes. The Metal Particle (MP) tape used through LTO-6 is too rough to be used on the new TMR heads. LTO-7 switched to Barium Ferrite (BeFe) tape which is smoother and is now being used from LTO-7 onward. As of the writing of this paper, the LTO Consortium has not made a final decision on whether LTO-9 will read back 1 or 2 generations.

LTO-8 drives are also capable of initializing a new and unused LTO-7 media to create “LTO-7 Type M” media. This initialization process takes the LTO-7 media from 6TB per tape to 9TB per tape. That’s a 33% increase in capacity but *no* increase in pricing. This allows users one of the most affordable storage mediums on the market today.

Users with LTO-5 or LTO-6 tape drives might be hesitant to upgrade to LTO-8 for fear of losing read access to their current tapes. For this reason, many users are unable to move to new tape library technology. By allowing users to move their older LTO tape drives to the new Spectra Stack tape library, partitions may be configured for the older tape drives and tape cartridges meaning they can immediately move their drives and media to a new library and decommission the older library without migrating any of their data to new media. Given this option, users are now able to take advantage of the newest tape library technology as well as the latest LTO tape technology and not lose access to their existing tape archive or be forced into an immediate “hard” migration which will be discussed below.

Spectra Stack allows users to configure as many as 20 partitions in a single library. Each partition appears as a virtual, stand-alone library to the user and software application. The



minimum configuration for a partition is 1 tape drive and 5 tape cartridges. Given the large number of partitions supported, it would be possible to have many different versions of LTO tape drives, and associated tapes, being used in a single Spectra Stack library.

There are several elements to consider when migrating from an existing library to a new library, and this is where the conversation takes a step back to consider higher level questions in the migration process.

Migration Paths – Data Sets

Which data tapes need to be migrated and which do not?

Most organizations will have a very large number of tapes in their existing library as well as stored in archives – onsite, offsite, or both. Obviously it would be ideal to have all data that is stored on old media converted to the new media immediately. But in many cases this is not necessary, or for that matter, practical. It may be helpful to determine which data sets need to be migrated before determining the technology needed for the migration. Ideally users can categorize to some extent which media will be migrated and which will not. Media aging out in the near future for example should be allowed to age out and removed from the migration pool.

There are data sets which are highly likely to *never* be accessed again on one end of the spectrum and data sets which are *very likely* to be recalled on the other end. Below is a general categorization of parameters to consider when embarking upon migration and calculating both the time and resources which will be required.

1. Backup Tapes

Backup tapes refer to the tapes used to perform ongoing daily and weekly backups. The value of any given backup set diminishes quickly. Once a given set of full and incremental backups is followed by the next week's set of full and incrementals, the value of the first set is significantly reduced. Users typically go to the most recent backup copy for restoration. Although the words "backup" and "archive" are sometimes used interchangeably, in this discussion we do not consider a set of backups as "archive" data. Archive data will be discussed next.

- a. New backup operations will immediately be directed to the newest drives and media in the library.
- b. In most instances existing backup tapes are never migrated to the new library. In some instances, approximately 90 days' worth of backup sets may be migrated.
- c. Backup tapes will hereinafter be referred to as "no-touch" media.



2. Active Archives

We use the term “active archive” to refer to archived data that has a high likelihood of being accessed in the future and also has a significant retention life. In some vertical markets such as Media and Entertainment or High Performance Computing, data sets or digital content is retained forever.

- a. The tapes which hold this type of archived data are the highest priority for migration, either soft or hard (explained in the following section).
- b. Active Archive tapes will hereinafter be referred to as “high-touch” media.

3. Compliance Archives

Here we’re referring to data which has specifically been archived for compliance mandates – usually dictated by organizational, government or industry policies.

- a. They have a defined expiration date.
- b. User usually has a long window for finding and recalling.
- c. We are suggesting these data sets not be targeted for migration if lifespan is down to less than two additional years.
- d. If not migrated to new media, the user will require at least one drive, which is compatible with the media set retained, for the duration of the compliance retention period.
- e. Compliance archive tape will hereinafter be referred to as “low-touch” media.

Migration Paths – Library, Drive and Media

We will describe two scenarios for migration – *Soft* and *Hard*. It is also possible to have elements of both soft and hard migration when planning a tape migration project.

Soft Migration

Overview – A “soft migration” refers to moving existing LTO media and drives directly from an existing library into a Spectra Stack library. The existing data movement application or data manager will be directed to the new library, an inventory will be conducted and media sets will be available for restoration if needed. See figure 2 below:

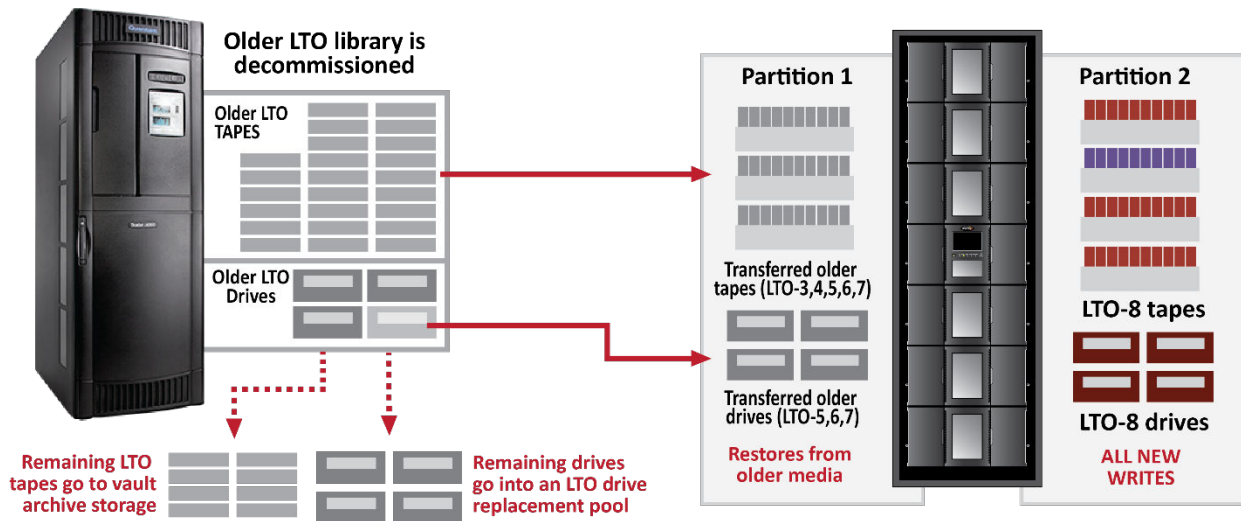


Figure 2: Here we see the process by which existing LTO drives and LTO media are moved from an existing tape library to a new Spectra Stack library which allows for immediate decommissioning of the older tape library.

If differing versions of LTO technology are not being mixed in the same library, partitions may be used for separating data on a departmental basis, operating system or application basis according to need. Migration would be complete.

As mentioned earlier, previous LTO drives have been able to “read back” 2 versions and “write back” 1 version of LTO tape. This changed with the introduction of LTO-8. If older version LTO drives and newer version LTO drives are being added at the same time, partitions would be created for each drive type and associated media. All *new* data writes will be directed to the partition(s) holding the new drives and media. Data retrieval from older media will be directed to the partition(s) holding that media and corresponding drive type. Over the course of time, older media types will be allowed to age-out as the data sets they contain come to end of retention.

There are several advantages to this approach. First and foremost, existing libraries could be decommissioned as soon as the installation is complete – significantly decreasing ongoing support, power and footprint costs.

In soft migration, after information is recalled from older tapes, it will then naturally be backed up or archived according to existing backup/archive schedules to the new media.

High-touch media, as defined earlier, would be moved in its existing form to the Spectra Stack library. Low-touch and no-touch media could be moved to the Spectra Stack or stored in vault or offsite storage. Some or all of the existing tape drives could be moved to the Spectra Stack. *Spectra is not able to provide support for customer’s existing tape drives if they were not originally purchased from Spectra.* Any existing, older LTO drives not installed in the Spectra Stack would be retained as replacement spares for problematic drives in the future. In this way, users could maintain their own pool of existing, older-model LTO drives as replacement drives.

Eventually that drive pool should be reduced to zero either through attrition or completion of all media migration required to the newest LTO media type. Using the older drives for read-only will certainly assist in lengthening the longevity of that drive pool.

If users don't have additional drives for a "support pool," Spectra does provide discounted pricing on refurbished drives that could be used as replacement drives. Support for the new or refurbished LTO tape drives purchased from Spectra and all other aspects of the new library system will be covered by Spectra's Price Lock Support which covers a wide range of offerings and SLAs for any user situation.

Hard Migration

Overview – Hard migration refers to restoring data from existing backup or archive tapes and then directing the data to the new tape library and tape media. Hard migration is most commonly used when an existing data mover application (backup, HSM or archive software) which writes in a proprietary format is being replaced as well. One might also refer to this as a "traditional" migration. See figure 3 below.




Figure 3: Hard migration from one an older tape library to a new tape library.

Under a traditional or hard migration, the old library and Spectra library would run in tandem. Hard migration of selected media will continue until completed at which point the older libraries will be decommissioned. This requires tremendous resources in the way of primary disk storage for staging and redirecting data, network bandwidth, and two tape libraries on the floor as well as all associated data center and support contract costs. Hard migration could go on for months or even years in large archive situations.

By moving older tape drives into the new Spectra Stack library, hard migration can occur in the background as time and resources allow without the need to maintain an additional library. Relying exclusively on hard migration is no longer required.

While hard migration of tapes is not required for a soft migration, it is still recommended to be performed in the background as time allows for high-touch or in some cases, low-touch media.

All new writes would be directed to the new LTO tape drive technology. Some backup/archive applications also support the feature of “tape stacking” where multiple, older LTO tape can be “stacked” onto newer LTO media without having to restore the data/content to primary storage first. This is an ideal feature for migration, but it’s not supported by all software packages. Figure 4 below shows all possible elements of migration revolving around a single Spectra Stack library after a soft migration has occurred.

Spectra MigrationPass Service

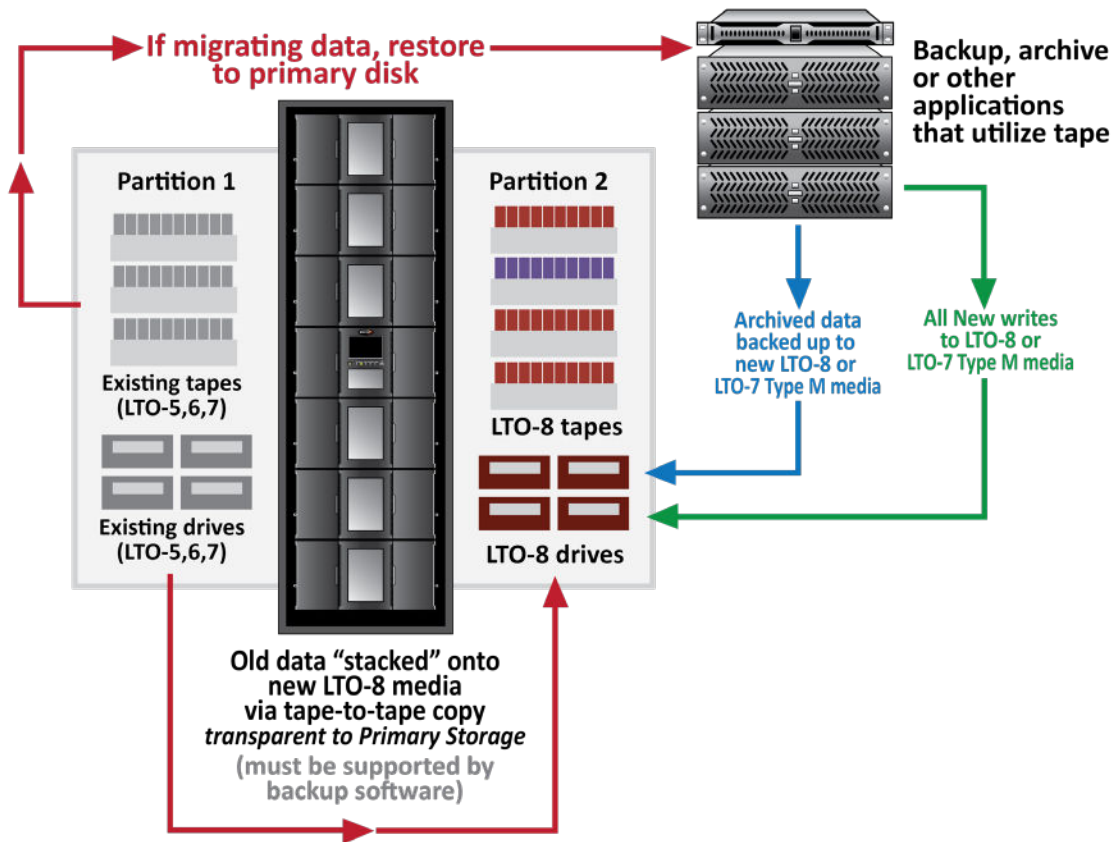


Figure 4: This hybrid approach to migration shows how various methods of migration can occur in a single library after the soft migration has completed.



Thoughts on Future Migration

This white paper started with the question, “Is the data you’re writing, creating or capturing going to last longer than the media it’s written to?” Chances are good that once the migration project is complete, it will only be a matter of time before the data needs to be migrated again. No migration project should be started without at least first considering what can be done to make future migrations easier.

There are two areas which are well worth investigating when considering any migration project: The format used to write the data or content and the file structure used to archive that data/content.

As mentioned above, many data mover software applications (backup, archive or HSM) use proprietary formats for writing data. These unique formats often add desirable features, but make it difficult to migrate existing archives from one application to a new application. As tape is the most common storage medium for long-term archives, it’s worth considering a tape solution which uses the Linear Tape File System (LTFS) format.

In addition to writing data to tape, LTFS also writes the metadata to each tape. That allows the tape to be “self-describing.” This works much in the same way a USB key would work. Insert the tape into a tape drive, and it can list all files on that tape even if the application that stored them isn’t present. This allows for better sharing of data and migration. The new data mover application would need to be able to read tapes in an LTFS format. Some do... some don’t.

In addition to data format, it’s also important to consider file structure. Traditionally, file systems have been used in creating, storing and archiving data or digital content. While file systems are very good at tracking, editing and sharing information, they aren’t as well suited for long-term archiving. File systems are limited in their ability to scale. Likewise, they incorporate the physical location of each file within the file system into the metadata. This hierarchical storage scheme becomes problematic when users wish to use tiered storage and move older, less frequently accessed data to a different medium for greater density and lower cost. In traditional backup/archive storage, the data mover which originally moved the data must retrieve the data to primary storage and then store it to the new location and update its tracking database.

When data is moved to the cloud, it’s typically stored using object storage. This structure allows for amazing scalability and no longer requires a set location for the data. When an application moves data into object storage, the object storage server or gateway provides a unique Object ID. This unique object ID is like a parking stub in valet parking. You pull up to an establishment with valet parking, give your car to the parking attendant, he gives you a parking stub, and you go into the restaurant or where ever you’re headed. Often your car is moved multiple times while you’re gone depending on how long you’re gone and how busy the establishment is. When you return, you present the parking stub and the parking attendant retrieves your car. You don’t need to know exactly where the car was parked or if it was moved from one space to another.

This is how the cloud works. You don't know the exact physical location of your data, you don't know what it's stored on, you don't know how often it's been moved around. All you need is the unique object ID, and the cloud provider retrieves your data. The service level you contracted for will determine the time to retrieve data. Some retrievals are virtually instantaneous; others may take days or more. While cloud is great for compute and many other services, it's tremendously expensive for long-term, large archives. Likewise, the bandwidth it takes to perform a large restore sometimes make the time and cost of retrieval unrealistic. Yet object storage is the ideal *method* for archiving information.

Spectra has combined object storage with the LTO format (when writing to tape) which allows users to create their own hybrid cloud on premise. This is accomplished via the Spectra BlackPearl Converged Storage Server which acts as a gateway to object storage on tape as well as nearline disk and spin-down disk. BlackPearl can even send a copy of the data to cloud if that's desired. While it's beyond the scope of this paper to cover the BlackPearl solution in full, figure 5 shows a diagram of possible storage targets using BlackPearl.




Figure 5: BlackPearl accepts data and then distributes it to one or more storage targets. By utilizing object storage, BlackPearl can make multiple copies of data or content, age out copies at different times, offer self-healing for data corruption and scale to billions of objects in a single name space.

The use of object storage and unique object IDs make it possible for BlackPearl to migrate tape from one version of LTO tape to a newer version without ever having to restore to primary storage. It can all be done via the BlackPearl and tape library. When writing data or digital content to tape, BlackPearl uses the LTO format. All tapes can be read with or without the BlackPearl solution in the future.

These are some of the ways in which new technology can help alleviate the past hardships of migration moving forward.



Summary

Migration offers the opportunity to re-evaluate size needs in a disk or tape library solution, assure the new system is “migration friendly” for future migrations, and ensure the partner you select to work with aligns with the business goals, ethics and direction of your own organization.

With 40 years in the storage industry, Spectra realizes that each migration project is unique. Business goals, existing infrastructure, technology preferences, timeframes and budgets all work together to create the opportunities and challenges that will be present in every event. Our professional services organization is always available to engage in conversation around the particulars of your project and bring decades of accumulated experience to the table.

Appendix 1

Tier 1 – Libraries supporting Straight Move Drives	Tier 2 – Libraries supporting Sled Move Drives <i>(Excludes drives listed under Tier-4)</i>	Tier 3 – Supported Drives	Tier 4 – Non Supported Drives
HPE StoreEver 1/8 G2	IBM TS3310	Half-height LTO-5 by HP (FC or SAS)	LTO-1 through LTO-4
HPE MSL2024	IBM TS4300	Half-height LTO-6 by HP (FC or SAS)	Full-height LTO-5 by HP (FC or SAS)
HPE MSL4048	IBM TS2900	Half-height LTO-7 by HP (FC or SAS)	Full-height LTO-6 by HP (FC or SAS)
HPE MSL8048	Quantum i500	Half-height LTO-8 by HP (FC or SAS)	Full-height, LTO-5 by IBM (SAS)
HPE MSL8096	Quantum i40	Half-height LTO-5 by IBM (FC or SAS)	Full-height, LTO-6 by IBM (SAS)
HPE MSL6480	Quantum i80	Half-height LTO-6 by IBM (FC or SAS)	
Overland NEOxl 8000	Quantum i3	Half-height LTO-7 by IBM (FC or SAS)	
Overland NEOxl 80	Quantum i6	Half-height LTO-8 by IBM (FC or SAS)	
Overland NEO T48	Quantum i6000	Full-height LTO-5 by IBM (FC)	
Overland NEO T24	Oracle SL150	Full-height LTO-6 by IBM (FC)	
Overland StorageLoader	Dell ML6000 family	Full-height LTO-7 by IBM (FC)	
Qualstar Q8		Full-height LTO-8 by IBM (FC)	
Qualstar Q24			
Qualstar Q48			
Qualstar Q80			
Fujitsu Eternus LT20			
Fujitsu Eternus LT40			
Fujitsu Eternus LT60			
Fujitsu Eternus LT260			
IBM TS3200			
IBM TS3100			

About Spectra Logic

Spectra Logic develops data storage solutions that solve the problem of short- and long- term digital preservation for business and technology professionals dealing with exponential data growth. Dedicated solely to storage innovation for 40 years, Spectra Logic's uncompromising product and customer focus is proven by the adoption of its solutions by industry leaders in multiple vertical markets globally. Spectra enables affordable, multi-decade data storage and access by creating new methods of managing information in all forms of storage—including archive, backup, cold storage, private cloud and public cloud. To learn more, visit www.SpectraLogic.com.