SPECTRA

DEEP STORAGE EXPERTS

May, 2015

OPTIONS IN TAPE TECHNOLOGY: TS1150 TECHNOLOGY and LTO

Abstract

There are two major tape technologies families available today for large scale, long duration data storage – midrange and enterprise. Both technologies offer advantages beneficial to users, so how should you choose which tape technology to deploy? The answer often comes down to the scale and economics of your storage requirements. The following paper briefly outlines the primary factors that you should consider when choosing whether to deploy midrange tape like LTO or enterprise tape like TS1150 Technology.

Overview

The analyst firm IDC has been tracking tape technology over several decades. During this time it has categorized tape technologies as entry level, midrange and enterprise. From 2000 onward, midrange tape generally meant LTO technology which has been the predominant tape technology in the market by volume shipped. By definition, midrange tape has consisted of tape technology with lower performance, capacity, and durability characteristics than enterprise tape, but also a lower initial price point.

The definition of enterprise tape technology has been reserved by IDC for proprietary tape drives like the TS11x0 and T10000 families from IBM and Oracle respectively. These drives have historically been considered premium drives by virtue of their operational characteristics and initial price point relative to LTO.

Within the universe of tape customers, users generally find themselves in one of two categories. Those who absolutely need the biggest, fastest, most reliable tape technology available and those who view it as optional. The former often include customers ranking in the top 20% to 25% by storage capacity with 2.5 PB or more of data, having demanding operational requirements, and the willingness to acquire premium technology. The latter typically have data footprints below 2.5 PB with less stringent operational requirements and who may be more price conscious at the point of purchase. Table 1 provides a general summary of user characteristics that may require Enterprise Tape (TS1150) technology versus those who are just as well served by LTO tape technology.

Table 1

Use TS1150 when	Use LTO when
Required	**Required**
Large Scale (> 2.5 PB*)	Multiple vendors
Highest performance	Lowest power consumption
Highest reliability	Read back 2 generations
Investment protection (legacy media reformat)	Public roadmap
Advance technology roadmap	Initial price valued above features / performance

*Release of LTO-7 in 4Q15 may increase capacity break point to 5 PB.

TS1150 Technology or LTO: When to choose.

Enterprise technology tape drives have historically been the leader in tape drive capacity, performance, and reliability. LTO technology drives, frequently called midrange, generally feature lower capacity, performance, and reliability metrics as well as a lower purchase price.

Choosing the right drive technology becomes a function of matching the features and their resulting benefit levels with your most critical operational requirements. Doing so can help optimizing your price-to-value choice when selecting a particular drive technology.

For customers with large environments and 24 x 7 operational requirements, enterprise drives may be the most cost effective solution over all. Customers with less critical duty cycle requirements can optimize their technology selection with LTO drives carrying a lower entry price point and comparably lower performance metrics.

Capacity and Performance

Data often grows unconstrained. Unfortunately, floor space needed to accommodate that growth is not always available. As a result, maximizing storage density within the space available becomes critical. If your organization is faced with unlimited data growth and limited floor space, then deploying the largest capacity i.e., most dense, tape technology available may be high on your priority list. TS1150 technology currently provides nearly 67% greater density than LTO-7 which will not begin shipping until late in 2014 (Table 2).

High levels of data growth often *mean* performance becomes important in order to address required data availability requirements. If large-scale data growth results in the need for greater performance for your organization, then the current generation TS1150's ability to deliver 17% greater throughput performance than the next generation of LTO should be strongly considered (Table 2).

Table 2 –Capacity and Performance

Feature	TS1150 Technology	LTO-7*
Capacity (Native)	10.0 TB	6.0 TB
Transfer Rate (Native)	360 MB/s	300 MB/s
R/W Compatibility	R/W TS1140 Reformat TS1140 To higher capacity / performance	R/W LTO-5 & 6 R only LTO-5

*LTO-7 projected to begin shipping in 4Q15. Specifications are estimated based upon best available information.

Reliability

Large scale data and high rates of data growth also require a great deal of technology reliability to ensure the data is available when needed and that it is accurately delivered. Enterprise tape technology like TS1150 is designed for high duty cycle environments. This means it is built to tolerate very high levels of usage e.g., tape loads and unloads which can stress the drive. The current generation of TS1150 technology tape drive is three times more durable than the next generation of LTO technology is projected to be. This means fewer fails, less down time, and lower replacement cost while ensuring a higher rate of data availability (Table 3).

Writing data to storage typically includes the possibility of write errors that may not become apparent until the data is recalled at which point in time it may be too late for correction. The more data being written, the greater the number of errors a user might expect to incur during the write process. For users with very large data repositories, this can mean large numbers of write errors unless the technology writes data with an extraordinarily high level of accuracy.

TS1150 technology tape drives write data at a corrected error rate that's three orders of magnitude better than the best LTO technology available and several orders of magnitude greater than that of disk. To put this into perspective, the best LTO technology drives may write up to 111 PB of data before a bit is written in error without correction. In contrast, the TS1150 technology drive may write up to 111,000 PB of data before a bit is written in error and not corrected. As a result, if your organization is writing large quantities of data to tape, the greater the Bit Error Rate (BER) the fewer the errors you will have to address over time (Table 3).

Table 3 –Reliability

Feature	TS1150 Technology	LTO-7*	SAS Enterprise Disk
Load / Unload Cycles	300,000	100,000	N/A
Bit Error Rate (Uncorrected)	$1x10^{-20}$	$1x10^{-17}$	$1x10^{-16}$
Bit Error Rate (Undetected)	$1.6x10^{-33}$	$1.6x10^{-33}$	N/A

Sources: IBM tape drive specifications documents and engineering briefing. Instrumental Inc., White Paper, "Tape: Comparison of LTO and Enterprise", April 19, 2013. Uncorrected BER are bits written error not corrected after application of error correction codes. Undetected BER could not be established for disk.

Investment Protection

Technology re-use and economic efficiency are important factors in any long-term or large scale storage strategy. LTO technology provides read / write capability with prior generations of LTO tape. It also enjoys a declining cost curve over time as successive generations of media become available. Consequently, the economics for LTO users, particularly those who are concerned with recurrent media pricing, are very attractive. Chart 1 demonstrates the rate at which LTO media pricing historically declines over time.

Chart 1 –LTO Media Pricing

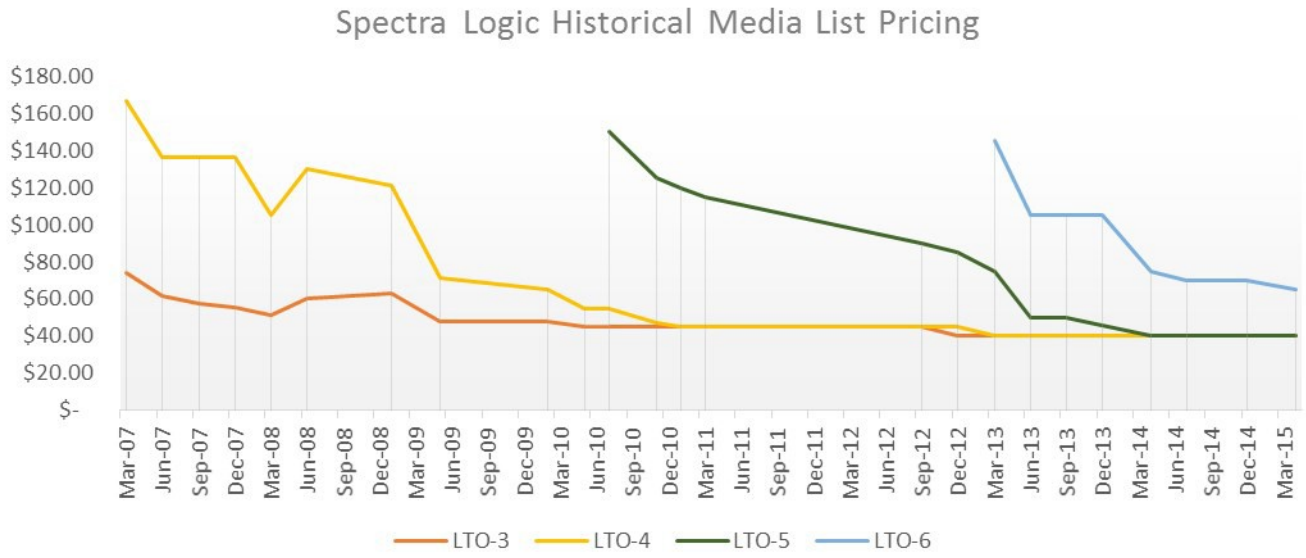

Curves based upon historical Spectra List Prices.

LTO-7 is expected to be released in 4Q15. Although actual media pricing is not yet available, current price estimates for LTO-7 media range from $180 to $220 per tape based upon the general cost ($ / GB) trend experienced with each successive generation of LTO media that is released. Assuming LTO-7 media declines in cost at a similar rate going forward as its predecessors have historically, users may expect to see 6 TB LTO-7 media cartridges dropping into the $120 range within a year and possible as low as $90 to $100 per cartridge within 18-24 months of release.

While LTO tape technology delivers favorable media price declines to its purchasers over time, enterprise tape technologies such as TS1150 typically don't see as great a decline in purchase price. However, TS1150 technology delivers what may be an even bigger benefit – media reformatting – and do so quickly. For example, a TS1150 tape drive can ingest a TS1140 tape cartridge and, from beginning of tape (BOT), reformat it to deliver not only greater capacity, but better performance as well. LTO does not offer this capability.

Charts 1 & 2 below display the added capacity and performance available with prior generation TS1140 tape media when reformatted in a TS1150 technology tape drive. The capacity increase of 75% represents growth in storage available that is in excess of data growth rates for many customers. In other words, by purchasing new tape drive technology, the legacy tape media may provide enough extra storage space to preclude the purchase of additional media otherwise needed to handle the data storage increase. The result is unparalleled investment protection as well as a reduction in the total cost of ownership over time because fewer media purchases may be required.

Assuming users purchased TS1140 media at $.04 - $05 per GB for a 4 TB cartridge, they would have paid $160 - $200 per tape. By reformatting the TS1140 cartridge to hold up to 7 TB of uncompressed data, the effective cost per gigabyte for that same TS1140 tape drops to $.02 - $.03 per gigabyte while also delivering 20% more performance. Furthermore, this kind of price decline can be controlled by a customer rather than waiting on the historical economic trends in price erosion normally found with LTO.

Charts 1&2 –Investment Protection

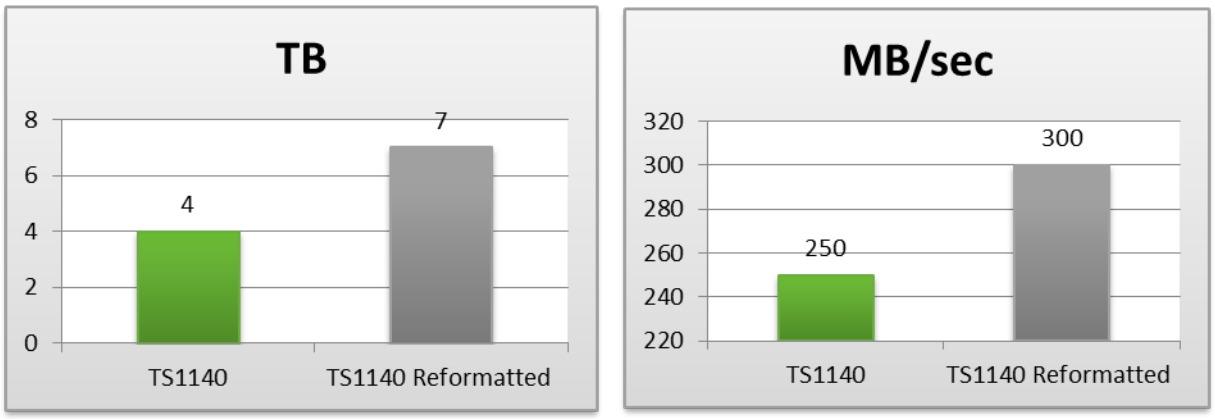

Source: IBM TS1150 OEM Specifications.

Cost Comparison

For many customers, the ultimate decision factor about the storage technology they choose to purchase is the cost of the solution. Generating apples-to-apples comparisons between storage offerings normally involves building similar configurations in terms of capacity and performance, then using a common metric such as cost per gigabyte ($ / GB) for the actual comparison.

Relative to LTO, the cost per gigabyte advantage of TS1150 becomes apparent as configuration size increases from the 1 PB level. Chart 3 below compares $ / GB based upon similar performance and capacity configurations for LTO and TS drives and media while using representative prices for Spectra library solutions.

Chart 3—Storage Cost per Gigabyte

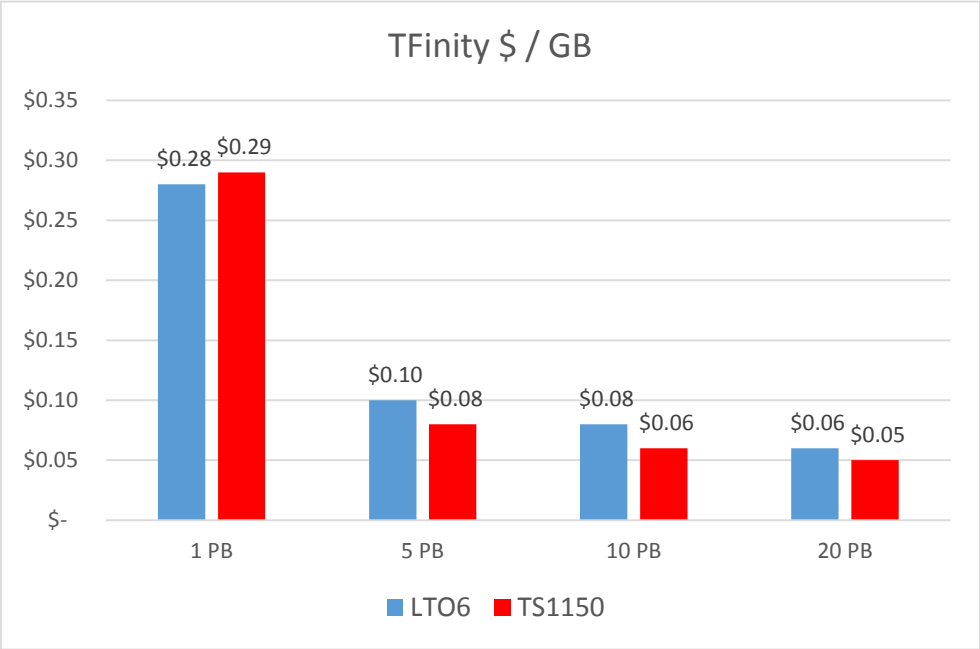

Pricing uses LTO-6 within the library configurations rather than LTO-7 since LTO-7 price points are not yet available.

The increased density and performance of TS1150 tape media and drives result in customers purchasing fewer tape drives and fewer tape cartridges than they would with LTO in order to achieve similar performance and capacity characteristics. Because of the smaller number of drives and tape media, users can purchase fewer library frames while licensing a smaller number of library slots. The aggregate reduction in drives, media, slots, and library frames helps drive down the total solution cost of TS1150 technology relative to LTO on a $ / GB basis.

Summary

All tape technology is not created equal. IDC is correct in defining and tracking tape technology as either midrange or enterprise given the feature, function, cost differentials, and target markets of the two categories. For users with 1 PB or more of data, TS1150 technology tape makes logical, economic sense. For users with less than 1 PB of data, but high data growth rates, TS1150 technology may be a fit depending on their planning horizons. And for users with much smaller data footprints and lower growth rates, LTO is a great fit.