

White Paper

Tape: Comparison of LTO and Enterprise

Instrumental, Inc.

<http://www.instrumental.com>

April 19, 2013



Contents

Overview	3
Reliability of Disk and Tape	3
LTO and Enterprise Tape Performance	4
Acquisition Cost	4
Other Features	5
WORM	5
Backwards Compatibility	5
Technology Refresh	6
Summary	6
Definitions	6
Disclaimers	7

Table of Figures

Figure 1 LTO Past and Future Public Roadmap	6
Figure 2 Enterprise Past and Future Public Roadmap	6

Table of Tables

Table 1 Disk and Tape Hard Error Rates	3
Table 2 Failure Impact of Hard Error Rates	3
Table 3 LTO and TS1140 Performance	4
Table 4 Tape drive and media cost comparison	4

Notice

Product names mentioned in this document may be trademarks and/or registered trademarks of their respective companies and organizations and are the property of these companies and organizations.

Tape: Comparison of LTO and Enterprise

Overview

Different technologies have different characteristics for reliability, performance, cost and usability. For storage this is true for SSD, disk and tape. This paper explains the advantages tape has over disk for reliability for archival data and explains the differences between enterprise tape and LTO tape for archival solutions. Large archives require a good understanding of not just the short term technology and cost issues, but a benefit from a long-term analysis of technology trends, and present and future value costs.

As part of this process this paper looks at some of the technology issues when making the determination of what technology provides the best value for larger archives. There are two different media choices from IBM in the market today. They are:

1. TS1140 enterprise tape
2. Linear Tape Open (LTO)

Each of these tapes has different characteristics and benefits; this paper attempts to present the facts for these two IBM tape types. It should be noted that IBM enterprise tape has been around for almost 30 years while LTO has been a standard for almost 15 years. Technology refresh and media reuse are a significant issue for the long-term cost of any archive. As an example, the support for the hardware interface to the tape drive such as fibre channel has about a 7 year life before the technology reaches end of life (EOL). In a critical archive you do not want any technology reaching EOL, and migration to new technology must be completed before EOL. Though the tape media might be rated to last 30 years there are a large number of issues in other parts of the system and operational environment that reduce the 30 year number to less than a decade.

Reliability of Disk and Tape

The hard error rate, which defines how many bits of data can be read before a read fails, for enterprise tape storage is at least 2 orders of magnitude better than that for enterprise disk storage as shown in the table below.

Table 1 Disk and Tape Hard Error Rates

Device	Hard error rate in bits	Equivalent in bytes	PiB equivalent
SATA consumer	10E14	1.25E+13	0.01
SATA Enterprise	10E15	1.25E+14	0.11
Enterprise SAS/FC	10E16	1.25E+15	1.11
LTO and some Enterprise SAS SSDs	10E17	1.25E+16	11.10
Enterprise TS11xx Tape	10E20	1.25E+19	11102.23

The table below shows what these hard error rates translate to in terms of failure rates with the devices running at their rated performance.

Table 2 Failure Impact of Hard Error Rates

Device Type	1 Device	10 Devices	50 Devices	100 Devices	200 Devices
	Hours to reach hard error rate at sustained data rate				
Consumer SATA	51	5	1	1	0
Enterprise SATA	301	30	6	3	2
Enterprise SAS/FC 3.5 inch	2,759	276	55	28	14
Enterprise SAS/FC 2.5 inch	1,965	197	39	20	10
LTO-5	23,653	2,365	473	237	118
LTO-6	20,696	2,070	414	207	103
Enterprise TS1140 Tape	13,797,371	1,379,737	275,947	137,974	68,987

Tape: Comparison of LTO and Enterprise

Based on the industry defined hard error rates, options with a tape storage component will have a higher level of reliability than disk or SSD. It should also be noted that a single bit in error on a disk drive translates to a full sector error (512 bytes or 4096 bytes for some newer devices) while with tape a bit in error results in only 1 bit of data being lost.

Besides the hard error rate, disk drives have an approximately five year life span based on numerous studies going back to 2007. Tape media on the other hand has a much longer life, and most media is rated to 30 years. Now, of course, there are many issues with trying to read media beyond 10 years, including drive interfaces, other hardware and software formats. But, tape has far more longevity than disk drives, and disk drives suffer from the same interface issues.

LTO and Enterprise Tape Performance

There are important performance differences between LTO-5, LTO-6 and enterprise TS1140 drive tapes.

Table 3 LTO and TS1140 Performance

	Number of Drive Speeds	Interface(s)	Native TB	Performance in MiB/sec uncompressed
LTO-5	13	FC-8 SAS-6 Gb	1.5	124
LTO-6	13	FC-8 SAS-6 Gb	2.5	153
TS1140	13	FC-8 SAS-6 Gb	4	238

What is important to note is that all IBM drives support multiple speeds to allow the drive to match the incoming data rate. This is important for the performance of the drive. If the data rate is not matched, then the drive will have to slow down, starting and stopping the tape and causing poor performance.

Clearly, the TS1140 drive has much higher performance and much greater native density than LTO-6. As many archival formats, such as video and audio, are pre-compressed, the native capacity is the most useful comparison for these formations.

Acquisition Cost

The following baseline cost numbers and floor space square feet for various configurations are based on standard GSA costs. By looking at the costs for tape drives, robot, slots, and media, it provides a good comparison of media types. Floor space is also compared, but the cost is not taken into account for the total pricing. The following table compares LTO-5, LTO-6, TS1140 and next generation TS11xx.

Table 4 Tape drive and media cost comparison

	GiB I/O	Cart Slots	Frame Type			Drives	Sq. Ft.	Total Cost	Cost/TB	Media Acq.
			L	D	HD					
5 PB										
LTO-5	5	3334	1	2	2	33	52	\$668,440.00	\$133.69	\$150,030.00
LTO-6	5	2000	1	2	1	31	41	\$593,650.00	\$118.73	\$126,000.00
TS1140	5	1250	1	1	1	20	31	\$737,547.50	\$147.51	\$249,669.00
Future media use based on historical data	5	770	1	1	0	16	21	\$533,567.50	\$106.71	\$153,808.00
Future based on historical data	5	556	1	1	0	14	21	\$470,470.00	\$94.09	\$130,660.00
10 PB										
LTO-5	5	6667	1	2	5	33	83	\$903,025.00	\$90.30	\$300,015.00
LTO-6	5	4000	1	1	3	31	52	\$747,850.00	\$74.79	\$252,000.00
TS1140	5	2500	1	1	2	20	41	\$1,015,435.00	\$101.54	\$499,375.00
Future media use based on historical data	5	1538	1	1	1	16	31	\$715,175.50	\$71.52	\$307,216.00
Future based on historical data	5	1111	1	1	0	14	21	\$600,895.00	\$60.09	\$261,085.00

Tape: Comparison of LTO and Enterprise

			Frame Type								
10 PB											
LTO-5	10	6667	1	6	3	67	103	\$1,342,945.00	\$134.29	\$300,015.00	
LTO-6	10	4000	1	6	1	62	83	\$1,182,130.00	\$118.21	\$252,000.00	
TS1140	10	2500	1	3	2	40	62	\$1,471,335.00	\$147.13	\$499,375.00	
Future media use based on historical data	10	1538	1	1	1	33	31	\$1,054,750.50	\$105.48	\$307,216.00	
Future based on historical data	10	1111	1	2	0	28	31	\$908,745.00	\$90.87	\$261,085.00	
15 PB											
LTO-5	10	10000	1	6	5	67	124	\$1,549,330.00	\$103.29	\$450,000.00	
LTO-6	10	6000	1	6	3	62	103	\$1,364,530.00	\$90.97	\$378,000.00	
TS1140	10	3750	1	3	3	40	72	\$1,749,222.50	\$116.61	\$749,063.00	
Future media use based on historical data	10	2294	1	2	3	33	62	\$1,290,361.50	\$86.02	\$458,227.00	
Future based on historical data	10	1667	1	2	1	28	41	\$1,067,605.00	\$71.17	\$391,745.00	
15 PB											
LTO-5	15	10000	1	8	5	100	145	\$1,977,970.00	\$131.86	\$450,000.00	
LTO-6	15	6000	1	7	2	94	103	\$1,725,490.00	\$115.03	\$378,000.00	
TS1140	15	3750	1	4	1	60	62	\$2,120,522.50	\$141.37	\$749,063.00	
Future media use based on historical data	15	2294	1	4	1	50	62	\$1,629,936.50	\$108.66	\$458,227.00	
Future based on historical data	15	1667	1	3	0	42	41	\$1,347,255.00	\$89.82	\$391,745.00	
20 PB											
LTO-5	15	13334	2	8	7	100	176	\$2,217,770.00	\$110.89	\$600,030.00	
LTO-6	15	8000	1	7	4	94	124	\$1,907,890.00	\$95.39	\$504,000.00	
TS1140	15	5000	1	4	3	60	83	\$2,426,610.00	\$121.33	\$998,750.00	
Future media use based on historical data	15	3077	1	4	2	50	72	\$1,814,540.75	\$90.73	\$614,631.00	
Future based on historical data	15	2222	1	3	0	42	41	\$1,477,680.00	\$73.88	\$522,170.00	

The cost of using LTO media, drives and robots is currently slightly cheaper than using enterprise media today, but it is expected the future enterprise density will significantly reduce the costs. The advantages today of using enterprise media include significantly reduced floor space, low number of tape drives given the higher performance, and significantly better media reliability.

Other Features

Both the LTO formats and TS1140 support other features not found on disk storage in most cases. These features provide tape with a significant advantage over disk for archival data.

- WORM (Write Once Read Many)
- Backwards compatibility

WORM

Both LTO and TS1140 have WORM support. This allows initial write to be guaranteed non-permutable. WORM ensures that critical data is not overwritten or destroyed.

Backwards Compatibility

Both LTO and TS1140 drives allow for backwards compatibility. LTO-6 drives allow LTO-5 cartridges to be read and written and allow LTO-4 cartridges to be read. TS1140 on the other hand can read back 4 generations and write previous generations' media at a higher density than the native density. Therefore TS1140 provides more flexibility than LTO. The following roadmaps summarize the direction of each media type.



Figure 1 LTO Past and Future Public Roadmap

LTO Generation	LTO-3	LTO-4	LTO-5	LTO-6*	LTO-7	LTO-8
Max Native Capacity	400 GB (L3)	800 GB (L4)	1.5TB (L5)	2.5TB (L6)	Up to 6.4 TB (L7)	Up to 12.8 TB (L8)
Other Native Capacities	200 GB L2 100 GB L1 R/O	400 GB L3 200 GB L2 R/O	800 GB L4 400 GB L3 R/O	1.5 TB L5 800 GB L4 R/O	2.5 TB L6 1.5 TB L5 R/O	6.4 TB L7 2.5 TB L6 R/O
Native Data Rate	80 MB/s	120 MB/s	140 MB/s	160 MB/s	Up to 315 MB/s	Up to 472 MB/s

* Data Compression engine enhancement from 2:1 to 2.5:1

Figure 2 Enterprise Public Roadmap

TS1100 Generations	Gen - 1 3592 J1A	Gen-2 TS1120	Gen-3 TS1130	Gen-4 TS1140
Max Native Capacity	300 GB (JA)	700 GB (JB)	1.0 TB (JB)	4.0 TB (JC)
Other Native Capacities with Media Reuse		500 GB JA	640 GB JA	1.6 TB JB 640 GB JA R/O
Native Data Rate	40 MB	100 MB/s	160MB/s	250 MB/s

As can be seen from the public roadmaps the density growth for enterprise tape is much greater than with LTO and the ability to read and write previous generations of media provides greater flexibility.

Technology Refresh

As technology is changing at a rapid rate migration must be expected for any technology, as both hardware and software tend to lack support for the data path after about 5-7 years. Simple things such as tape drive hardware interfaces and tape drivers in software tend to have a limited support window. Though tape media might last 25 or 30 years, it will become very difficult to read after 10 years given the lack of hardware and software support. Initial planning for technology refresh should be done as part of the architectural design process for the number of tape drives and bandwidth requirement to meet the Service Level Agreement (SLA) for the users.

Summary

In large archives it is clear that enterprise tape has a significant advantage over LTO in areas of reliability, density per square foot of floor space today and, given future enterprise densities and reuse, overall cost and the ability to read old data without migration to new media. LTO is a commodity product that provides inter-vendor operability, but it comes with the disadvantage of significant innovation in areas such as tape density and reliability. For large archives, IBM enterprise tape is a better choice with long-term cost advantages over LTO.

Definitions and Disclaimers

Definitions

- Bit Error Rate (Calculated) – Storage device. The average bits processed between bit error events affecting availability of data at a storage device level, excluding externally generated failures. This projection does not reflect any specific storage device or subsystem configuration, but is an average across all installed subsystems and configurations of a given type. Bit Error Rate is expressed as a value of bits processed per bit error.
- Load/Unload life – Storage device. The average load and unload cycles, (i.e., a complete load followed by an unload) between hardware impacts availability at a storage device or subsystem level, excluding externally generated failures. This projection does not reflect any specific storage subsystem configuration, but is

an average across all installed subsystems and configurations of a given type. TS 1140 supports 300,000 Load/Unload cycles

- Mean Time Between Failure (MTBF) – Storage device. The average time between hardware impacts availability at a storage subsystem level, excluding scheduled outages. This projection does not reflect any specific storage subsystem configuration, but is an average across all installed subsystems and configurations of a given type. MTBF values are expressed in hours, months or years. TS1140 MTBF is 237,000 hours

Disclaimers

This information is provided for planning and information purposes only, and does not represent an express or implied warranty, guarantee or contractual commitment. THE SUBMITTAL OF THIS DATA CREATES NO WARRANTY, EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR APPLICATION.

The reliability metrics set forth are based on the following statements, definitions and assumptions:

- Availability estimates are averages derived from measured field performance data and/or current engineering projection data. The following are excluded events from availability:
 - o Vandalism, abuse, fire, flood damage or acts of God, or any other factor beyond IBM's reasonable control.
 - o Scheduled outages of any nature, which include but are not limited to, remedial or preventative maintenance or engineering changes.
 - o Operator errors, operational procedures or procedural changes that cause system outages.
 - o Application or operating system errors outside of the storage subsystem which prevent access to data.
 - o Corrupt data introduced into the solution.
 - o Loss of or interruption to electrical power (IBM highly recommends use of an uninterruptible power system (UPS) to minimize power-related outages).
 - o Network related error or application hang conditions
 - o Installation of any new products, features or changes not defined as part of the proposed configuration.
 - o Any outages or degraded performance that occurs due to delay or postponement by the customer in obtaining support or required services needed to preserve integrity and availability of the proposed solution.
- Field performance data reflects an average across installed subsystem configurations and features in typical application environments. The data assumes industry standard operating system recovery procedures and that environmental / operational characteristics and maintenance is performed according to IBM maintenance standards, procedures and schedules.
- Engineering projection data assumes the use of the industry standard operating system error recovery procedures, assumes maintenance is provided according to IBM maintenance standards, procedures and schedules and assumes the units are operated in accordance with IBM environmental and operational specifications.