TANDBERG DATA

WHITE PAPER

SLR100 Technology Advancements and Features

by: Willy A. Johansen
Corporate SLR Product Manager
Tandberg Data ASA

SLR100 - 150% THE TRANSFER RATE AND 2 TIMES THE CAPACITY OF PREVIOUS SLR MODELS

Tandberg SLR100 deliver high capacity, high transfer rates and data integrity for "trouble-free backup and restore"

1. Why Backup

Lost Data Cost U.S. Economy $11.8 Billion in 1998; Pepperdine Professor Estimates Average Incident of Lost, Stolen or Damaged Data Runs More Than $2,500 per Affected PC. (With reference to "The Cost of Lost Data" report from Stac Software Inc.)

The report projects costs are likely to increase as the need for distributed computing increases. Extra costs are likely to be incurred if more than one PC on a network is affected and without protection such as virus software and backup systems, large-scale problems over the network could cripple an organization. Smith also states that the costs may even be much higher, as some factors are difficult or impossible to quantify, such as potential sales losses, lost opportunity, damage to company reputation and complete loss of the knowledge itself.

According to the report, hardware failure, human error, software corruption and viruses are the leading causes of data loss. Theft, which is prevalent among laptops, was responsible for five percent of lost information while the remaining three percent of incidents was the result of complete hardware destruction from events such as floods, brownouts and lightning.

"There are two possible outcomes for lost data," said researcher Smith. "Either it is recoverable with the help of technical personnel, or it is completely lost and needs to be re-entered. Both scenarios involve significant hours of work and considerable sums of money."

2. The Most Important Features for a Backup Solution?

The three most important functions for a backup device are:

1. **Trouble-free restore**
   - Restore of data is normally not an every-day-operation, but for most companies restore of lost data will be needed at some time. When a restore is needed, the most important function to have is a trouble-free restore. Trouble-free restore is only secured by having a reliable tape drive solution, reliable software, and a good backup routine.
2. **High performance**
   - Backup of the data within a shortest possible time window. The time window wanted depends on how long the host ideally can be occupied with backup. The time needed to spend on a backup operation depends on the total system performance including the speed of the hard disk, size of the files, speed on the network, performance of the operating system and the backup application software, and the performance of the tape drive for the different types of data.
   - Data to be backed up for most companies will be a mix of data that can be compressed from no compression at all (1:1) to several times of compression (x:1). The tape drive should also be able to have a linear increase of transfer

rate as the compression ratio increases. Compression ratios of e.g. 3:1 - as for text files and graphical files - should accommodate a three times higher transfer rate.

3. **Scalability and Compatibility**
   - As trouble free restore is very important for a backup system, being able to stay with a known, reliable technology over years becomes more important. Reliability means that the user can feel assured that a restore operation can be performed without any problems. This requires that the tape technology chosen can offer a growth path in capacity, transfer rate, and automation products that covers the growing requirements for the customers needs for data protection in the years to come.
   - For many customers it is also important to be able to read older media, and also for some customers to be able to write to older media formats for data transfer and software distribution.

3. Importance of Trouble-free Backup

Data Protection is an insurance against data loss. Using highly reliable linear tape drive technology and well planned routines for backup will minimize the work and hassle when the need for restoring of data occurs, and it will occur for every company. A welll planned backup routinecan be one of several e.g.:

Routine	Description	Advantage
Full backup every day	The whole disk is backed up	One or one set of backup media contains all data at the time of full backup. Maximum of data loss will be one day of data input
Incremental	Full backup once a week, and every day backup of only new data	One or one set of backup media contains all data at the time of full backup. One or a smaller set of media per day contains the new data for each day. The daily incremental backup takes shorter time but restore of data can mean searching for the media with the required data.
Differential	Full backup once a week, and every day backup of new data since the full backup	One or one set of backup media contains all data at the time of full backup. One or a smaller set of media per day contains the new data since the last full backup. The daily differential backup takes shorter time than the full backup. Restore of data means using the full backup media and the latest differential backup media.

Like for an insurance it is when an incident occurs one see if the insurance was worth paying for. Being able to recover from the incident or being able to restore the data after data loss is the whole and single reason for the insurance and the backup.

For a backup device it is essential that data can be restored at any time. A mechanical device like a hard disk drive, and a tape drive will break at one time. Disk drives are much more in use than a tape drive, and disk drives are thus more likely to break if everything else is similar in regards to mechanical parts.

Even a RAID 5 system is not fully fault tolerant (see section 1. Why Backup). More mechanical parts, and the more moving mechanical parts, the more likely it is that a failure will occur earlier. Especially this is true if the mechanical parts have an influence on writing to and reading data from a media. There are several techniques on how to write data to, and read data from a magnetic tape media. The linear technique the one using less mechanical and moving mechanical parts that can influence reading and writing, and thus linear is regarded as the most reliable technique. The linear technique is used on most higher capacity tape drives for backup like DLT, LTO, Magstar and SLR.

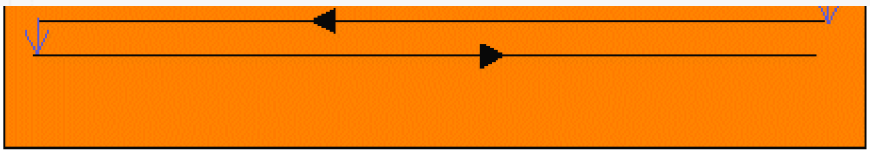

Linear technique means writing data on a longitude track in the whole length of the tape media. Only at the end of the media the write/read head is moved to another position and to a new track.

The Tandberg SLR (Scalable Linear Recording) product is the tape drive – within linear technology -that has the fewest moving mechanical parts – the head, and the capstan motor. The read/write head is moved to another position/track at the end or beginning of the media whilst the capstan motor drives the tape media. Also for the SLR technology, the media is always inside the cartridge, and all media wear components are inside that cartridge. This means again higher quality as most of the wear parts are replaced when the media is changed.

4. Importance of High Performance

For more and more companies, the time window for backup is shrinking while the amount of data is increasing. For some server platforms, the time window for backup is down to one hour or even less. To cope with such requirements one needs not only a fast tape drive, but also a tape drive that can adopt to the speed of the data coming to the tape drive. Also the data path from the hard disk to be backed up to the tape drive is essential. A 100Mbit-network connection line only can support about 4 to 5 Mbytes/sec of data transfer in total for all devices connected. This means only 14 to 18 gigabytes per hour shared between all devices communicating on the network connection. Thus a direct connection of the backup tape drive to the host where the hard disk is installed, is the only solution. For such a connection using a Wide Ultra 2 SCSI interface, the possible burst transfer rate of the bus is 80 Mbytes/sec – or 288 gigabytes per hour. Then the speed of the backup is limited to the maximum speed of the tape drive, and how well the tape drive can buffer, compress, and process different types of data.

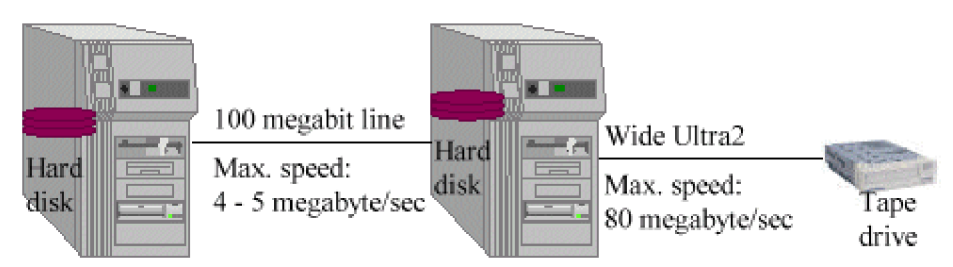

Data transmitted from remote server – speed limited to the speed of the line vs. directly attached to the server – speed limited to maximum sustained transfer rate of the tape drive.

Some types of data can not be compressed, and others can be compressed more than ten times. In an office environment there is a mix of these types of data. Dependant on type of company, and what kind of data the company normally uses, an average compression ratio can be much more than 2:1 as the tape drive industry normally is using. This means that a tape drive with hardware data compression should ideally be able to increase the transfer rate with the same factor as the data is compressed. If then data is compressed two times, the actual transfer rate should double, and if data is compressed six times, the actual transfer rate, and capacity should be six times the native transfer rate.

The actual speed of the backup is however limited not only to how fast a tape drive can write data on the tape media, but also the how fast the electronics and firmware inside the tape drive can process the data before the data is written to the tape. A slow electronics or firmware processing inside the tape drive will limit the actual transfer rate to a certain maximum speed even if the data is compressed several times.

5. Tandberg SLR100

The Tandberg SLR 100 announced in November 1999 is the newest tape drive product in the SLR series. The capacity is 50 MB native, and 100 MB assuming 2:1 compression whilst the transfer rate is 5 MBs in native mode and 10 MBs with 2:1 compression. The product will be available to the market from December 22, 1999.

To accommodate the requirement for trouble-free backup and high performance, SLR100

To accommodate the requirement for trouble-free backup and high performance, SLR100 has several advanced technical features ensuring "trouble free backup" with high reliability, transfer rate and capacity. The most important technical advancements are detailed below.

6. SLR100 Features for Trouble-free Backup

I. Improved Rewrite Algorithm

The SLR tape drives has embedded read-while-write functionality. During the write operation the data is read back (see Four-channel Head section below), and checked for any error. If the read signal during the read-while-write operation for SLR100 is outside factory set limits, the write current, and/or read signal amplification is adjusted so the read-while-write signal again is within the correct limits.

For every block of data, six bytes of CRC (Cyclic Redundancy Check) data is calculated and written just after the data bytes. During the read-while-write operation, the CRC data is again calculated and checked with the written CRC. If any error, the data is re written to a different track. A block can be rewritten on a different track to end of tape before the tape drive reports a hard write error. In theory all data can be written on four, three, two or on one single track if the media or the other tracks are not functioning.

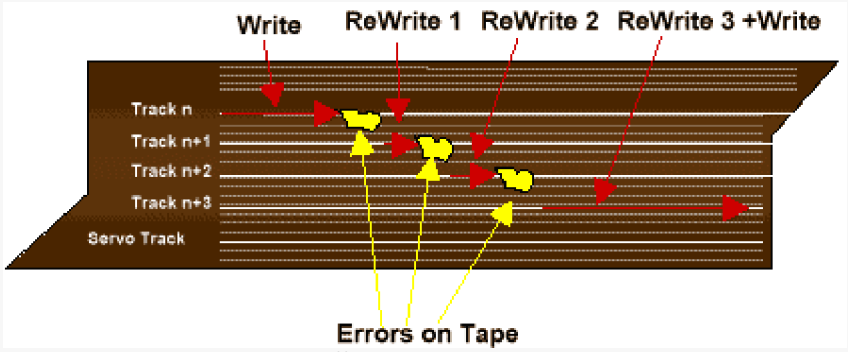

Blocks are rewritten on a different track to improve reliability during write

II. Improved ECC Level 10 (RS 64, 54)

The Tandberg SLR100 writes 192 tracks of data on the 0.315-inch (8mm) media. In addition there are 24 prewritten servo tracks which are used to assure data is always written on the exact correct position on the tape media.

Data are written on the tape in blocks of 512 bytes. Each data block consists of:

- Random Seed – 3 bytes
  Seed for derandomisation of blocks used by the VR 2 technology
- Control Bytes - 8 bytes
  Type of block (data, ECC, etc), block number, etc,
- Data bytes - 512 bytes
  Actual data or other information like ECC(Error Correction Code) data
- CRC (Cyclic Redundancy Check) Bytes – 4 bytes
  Calculated number based on the actual data, and used for error detection during read-while-write and read operations.

SLR100 features an improved ECC named ECC Level 10 or RS 64, 54. Level 10 means there are 10 ECC blocks while RS 64, 54 are more detailed and mean Read_Solomon, and 64 blocks whereof 54 are data blocks in an ECC system (ECC Set). This ECC system can correct data in up to 10 concurrent incorrect blocks when reading the data back.

Before written on the tape, the data is sectioned into frames consisting of 108 data blocks and 20 ECC blocks. In a frame there are two ECC sets – even and odd. Data block 0, 2, 4, … 106 belongs to the even ECC set, while data blocks 1, 3, 5, … . 107 belong to the odd ECC set. So in a frame there are calculated 10 ECC blocks for the even data blocks and 10 ECC blocks for the odd data blocks – or 20 ECC blocks per frame of total 128 blocks.

frame of total 128 blocks.

A SUPER FRAME consists of four frames (labeled A, B, C and D) are divided among the four simultaneously written tracks on the tape. This will give a longitudinal distance of 8 blocks (6 mm) between blocks in the same ECC set or interleave. The idea is to spread the blocks of the same ECC set as fare away from each other as possible. The blocks of a super frame are distributed across the four tracks as shown in the figure below.

Track 0	**A0 (e)**	D1(o)	C2 (e)	B3 (o)	D4 (e)	C5 (o)	B6 (e)	A7 (o)	**A8 (e)**	D9(o)	C10 (e)	...	A127(o)
Track 1	B0 (e)	A1 (o)	D2 (e)	C3 (o)	**A4 (e)**	D5 (o)	C6 (e)	B7 (o)	B8 (e)	A9 (o)	D10 (e)	...	B127(o)
Track 2	C0 (e)	B1 (o)	**A2 (e)**	D3 (o)	B4 (e)	A5 (o)	D6 (e)	C7 (o)	C8 (e)	B9 (o)	**A10 (e)**	...	C127(o)
Track 3	D0 (e)	C1 (o)	B2 (e)	A3 (o)	C4 (e)	B5 (o)	**A6 (e)**	D7 (o)	D8 (e)	C9 (o)	B10 (e)	...	D127(o)

Channel Block Distribution. (e) - Even ECC set, (o) Odd ECC set.

The ECC data is used to recreate data lost during a read operation. Up to 20 completely missing blocks in a frame can be recreated. Since the blocks are spread over the media with a distance of eight blocks, there needs to be up to 56 mm or 2.2 inches of missing data on a track on the tape media. This is very unlikely to happen, and this will only happen if the media is heavily worn or damaged. Media with such defects as this will be be ejected during the write operation as all incorrect blocks are rewritten. More "normal" errors on the media are small dots of missing magnetic material spread all over the media, scratches along the media, derbies or dust on the head.

For a SUPER FRAME, up to 80 blocks can be recreated given there are not more that 10 concurrent missing blocks in the same ECC set on the same track. This gives the Tandberg SLR100 an ability to recreate almost 20% of all data on the whole media.

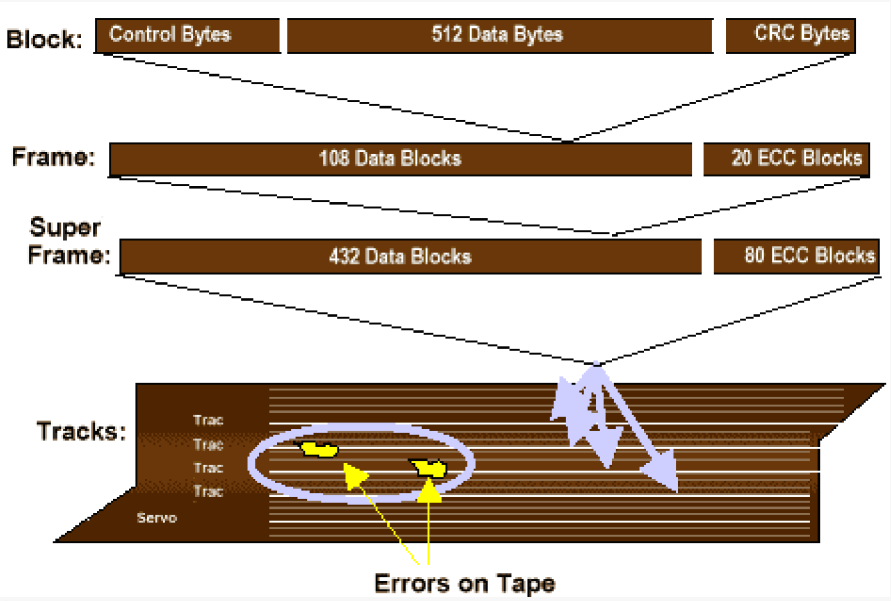

More than 2 inches (5.6 cm) of missing data on a track and almost 20% of all data can be recreated to assure reliability during restore

Data and ECC blocks are spread over four tracks to assure highest possible ability to recreate any lost data during restore.

III. Adaptive Data Channel

During a write or read operation, the read-while-write or read signal is checked against factory set limits. These limits are in the area of signal level, phase shift, and timing. If the read signal is outside any of these limits, the channel characteristics (amplification, write current, signal phase, other) are readjusted.

This functionality will assure correct written data even for worn media and/or worn head. As an additional assurance for reading correct data, the read channel

characteristics are also corrected if any error should occur during read - after the tape media was written.

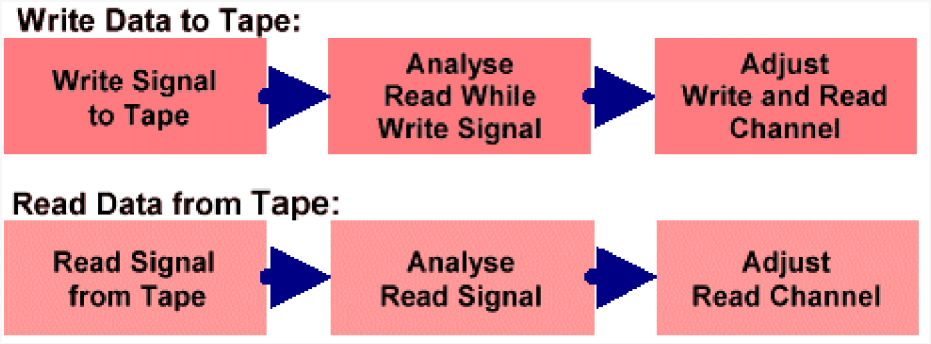

The read signal is always at the correct level to assure reliability during restore

The read signal is always verified both during read-while-write and read operations, and if needed, the channel characteristics are readjusted to assure correct writing, and trouble-free restore

7. SLR100 Features for High Performance

I. SLR100`s In-line Data Compression

Tandberg Data is using a specially designed circuit/ASIC to allow compression of data "on the fly" before the data is entered into the buffer. This allows the whole buffer to be used for intermediate storage of the data before written to the tape media. Using this method called In-line compression; the SLR100 can receive data sufficient for writing in 1.6 sec. instead of .8 sec if a Look-aside method is used.

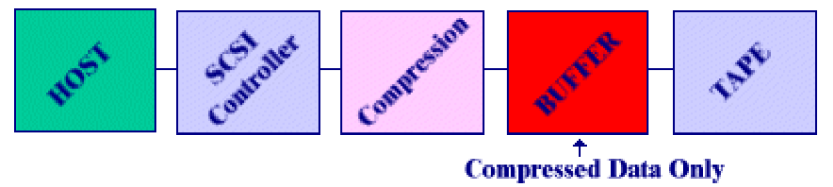

In-line compression: The data is compressed in a separate data-compression-chip before entered in the buffer. Thus 100% of the buffer can be used for intermittent storage of data before written to the media, which gives higher transfer rate.

Other technologies uses half of the buffer for the "raw" uncompressed data transfer the data to a compression chip, compress the data, and enter the compressed data to the other half of the buffer. This means both reduction of the buffer capacity, and need of additional operations for the tape drive's processor.

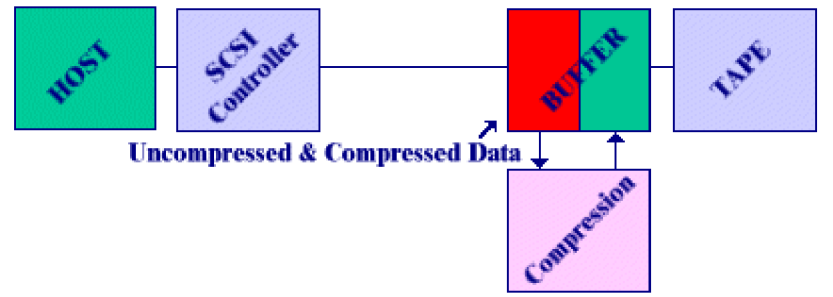

Other technologies' Look-aside Compression: The buffer is divided into two sections: Compressed and non-compressed that reduces the effective buffer size significantly

The *in-line compression* method also gives another advantage – better ability to increase capacity and transfer rate for highly compressible data -, as no additional operation is needed to compress the data. Benchmarks shows that the SLR100 can increase the transfer rate – and capacity – up to more than six times the native transfer rate and capacity when writing highly compressible data.

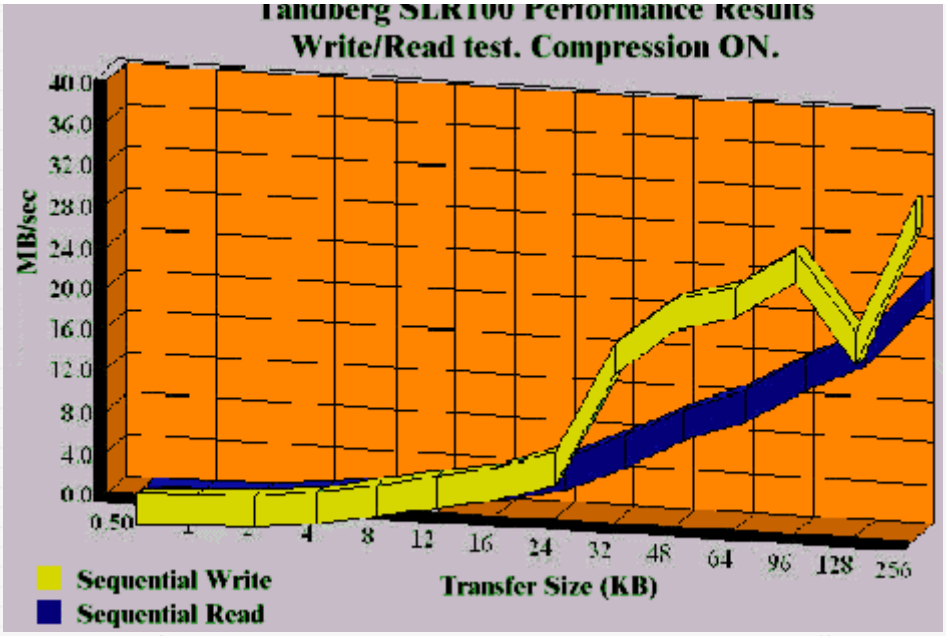

Benchmark performed by Tandberg Data writing highly compressible data using different block sizes (Transfer Size)

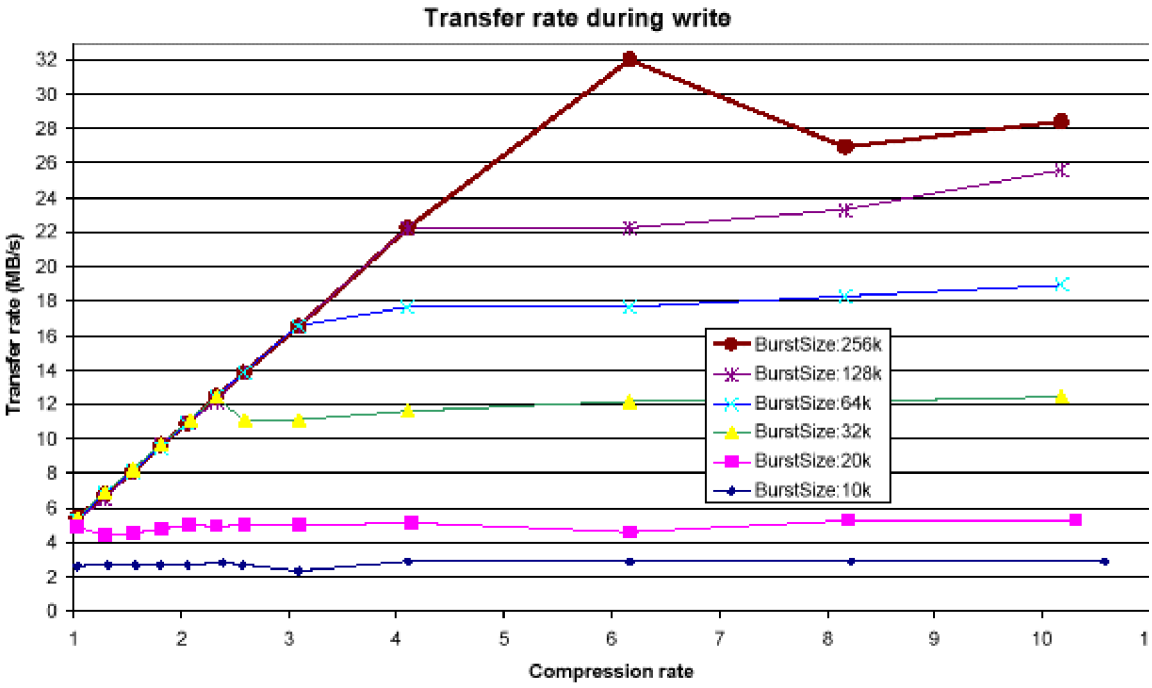

Test performed by Tandberg Data writing data with different compression ration, and buffer sizes

Tandberg SLR100 features in-line compression method which gives SLR100 the ability adapt to variable transfer rate on the host bus, and to increase transfer rate and capacity with increased data compression.

II. Large 8 MB Buffer

The SLR100 features an 8-MB buffer. With a burst transfer rate of 40 megabytes per sec. data feed to the tape drive from the host, an empty buffer fills up in 0.2 sec. With a sustained transfer rate – the speed of writing data to the tape media - of 5 megabyte per sec., the buffer is emptied – data written to the tape – in 1.6 sec. This means that the host and the SCSI bus can perform other operations in 1.4 sec. while the tape drive is busy writing data.

Large 8 MB buffer to store data before written to the tape media gives higher sustained transfer rate - optimized to the host's bus speed

III. SLR100 "Auto-sense" Transfer Rate

If the buffer is not filled or filling again after the 1.6 sec, the tape speed is automatically adjusted to accommodate 4 or 3 megabytes per sec. The tape speed for 5 MBs is 87.5 ips (inch per sec.), while the tape speed for 4 MBs is 70 ips, and 52.5 ips for 3 MBs.

This feature allows the SLR100 to adjust to the optimum speed on the data transfer from the host computer and to prevent underrun to allow maximum performance. Underrun means that the tape is stopped, and is rewinded to where the last block of data was written. The tape will start moving forward?when the buffer receives data, and the write operation is continued. This normally slows down the write operation. Another method to avoid underrun is to keep the tape media running, and keep on writing "filler", dummy, or the last block – also called forced streaming. This method however reduces the total capacity on the media.

SLR100 optimizes to the host's bus speed, which gives higher transfer rates and maximum performance.

IV. Wide Ultra2 interface

Most hard disks now in the market are using the SCSI3 Wide Ultra2 interface, which allows a burst transfer rate up to 80 megabytes per sec. (MBs). The next generation of the SCSI road map – Ultra3 allows up to 160 MBs. Ultra3 interface is just coming to the market on hard disk drives.

Also for the SCSI interface, backward compatibility is important. The Ultra3 protocol also allows devices with older interfaces to be connected to the same bus. Ultra2 devices can be connected to an Ultra3 interface bus, and an Ultra2 bus supports the older Wide Ultra and Ultra with only 20MBs transfer rate.

The Ultra2, and Ultra3 interfaces also utilizes the Low Voltage Differential (LVD) Signaling feature. LVD allows up to 12 meter of cable between the host interface adapter and the last device, whilst the Single Ended (SE) interface used for older SCSI protocols only allows 1.5 meter of cable. Using an interface that allows a long cable becomes more and more important as rack mounting is more and more used. The host computer or server can then be installed in one rack, while other devices like a tape drive or an automation product can be connected on the same bus in another rack several meters away.

The SCSI backwards compatibility has some limitations – e.g. if an Ultra or Wide Ultra device which only allows the SE interface - is connected to a Ultra2 bus, the whole bus is limited to the slowest SCSI protocol, and shortest cable length. If this device is connected to the same bus as hard disk drives or other devices, the Ultra/Wide Ultra device will slow down the performance of the whole system. Also the flexibility of installation will be limited of the shorter cable.

Eight devices can each send/receive 20 MB of data per second

Four devices can each send/receive 20 MB of data per second

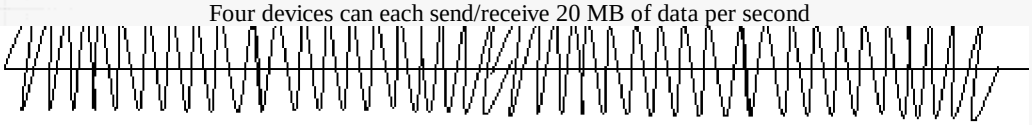

Only one device can send/receive 20 MB of data per second

SLR100 burst transfer rate of 40 MBs, supporting both LVD and SE interfaces gives improved overall system performance.

8. SLR100 Features to Assure Compatibility and Scalability

I. Four-channel TFMR Head

The Tandberg SLR100 features a thin-film, inductive write/magneto-resistive read head with five dual channels, and an additional dual channel for the SLR24, SLR32, SLR50 formats. The magnetic elements are all small elements that for the write operation can magnetize the particles on the media with a low write current.

Of the five dual channels for the SLR100 format, one is used for reading the servo track while the other four tracks are used for parallel writing and reading. Each of the dual channels accommodates two combined read/write elements. When the tape media is moved in forward direction the first element acts as a write element, and the other as a read element for the read-while-write operation.

When the tape media is moving in the reverse direction, the other element, which was the read element in forward direction now, acts as the write element. The read-while-write operation is then performed by the last element, which was the write element in the forward direction. With this head technology adding more channels for higher future higher capacity and higher transfer rate SLR tape drives is relatively easy.

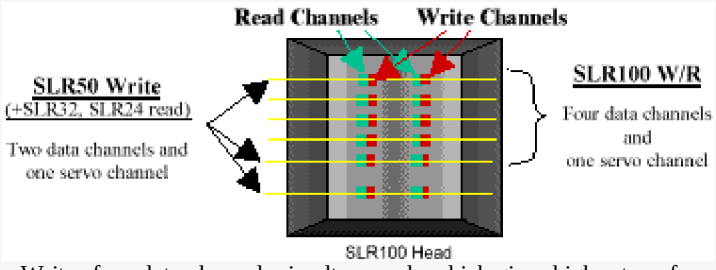

Writes four data channels simultaneously which gives higher transfer rate and improved reliability
(see the Improved ECC section for more information).

II. Advanced Data Coding - VR2™

In previous generations of linear tape drives, a peak detection algorithm was used to read back the data. To be able to secure the distinction between neighbor bits is was necessary to put in "dummy" bits after every second real bit. 1.7 RLL - Peak Detection Read Channel is commonly used in linear recording. For detection of each peak independently, every 2 data bits are recorded as 3 bit-cells on the tape.

Conventional 1.7 RLL coding- max 67% coding efficiency

With the PRML technique the bits can overlap and still the algorithm will be able to retrieve the original content. In addition more bits can be coded into each signal transition on the tape. This enables a higher logical bit density and transfer rate although the physical density remains the same.

Tandberg Data is using the VR2 (Variable Rate Randomizer) algorithm patented by Overland Data in the SLR100 products and product derivatives. This is the first linear tape drive in the market utilizing PRML (and VR 2) encoding.

tape drive in the market utilizing PRML (and VR 2) encoding.

VR2 technology improves on these inefficiencies - approaching 99% coding efficiency

The VR2 writing technology is more area efficient, and will thus help assure future higher capacity and higher performance SLR tape drives.

9. SLR Compatibility and Scalability Chart

Tape Format Compatibility	Tape Drives							
	SLR100	SLR50	SLR32	SLR24	SLR5	SLR4	SLR3	SLR2
SLR100	RW							
SLR50	RW	RW						
SLR32	R [1]	RW [1]	RW [1]					
SLR24	R	RW	RW	RW				
SLR5		R	RW	R	RW			
SLR4		R	RW	R	RW	RW		
SLR3					RW	RW	RW	
SLR2					R	RW	RW	RW
SLR1					R	RW	RW	RW

Note 1) : Includes 13/26 GB Media

The SLR product line is scalable from 525MB and 0.7BG/hr to 8TB and 346GB/hr, and backward compatible to at least one backward write, and two backward reads.