

# IBM's 3592 Storage Solution: A Taste of the Future

*Increased Capacity ... Record Speed ...  
Frequent Access to Data ... Intensive Use of the Drive*

Fujifilm Recording Media

Tel: +33 (0) 147 15 84 54

E-Mail: [elisabeth.gameiro@fujifilm.com](mailto:elisabeth.gameiro@fujifilm.com)



## TABLE OF CONTENTS

Introduction	Page 2
Increasing <b>the storage capacity</b> depends on securing the data integrity	Page 9
How can we <b>increase the storage capacity</b> of new tape generations?	Page 22
<b>BER (Bit-Error-Rate)</b> & the interpretation of BER	Page 32
The <b>write speed</b> will reach 400MB/s...	Page 35
Intensive use and <b>frequent access</b> to data	Page 59
How to <b>reduce the</b> storage system cost	Page 71

Every month, the equivalent of one exabyte on tape is delivered by tape vendors across Europe




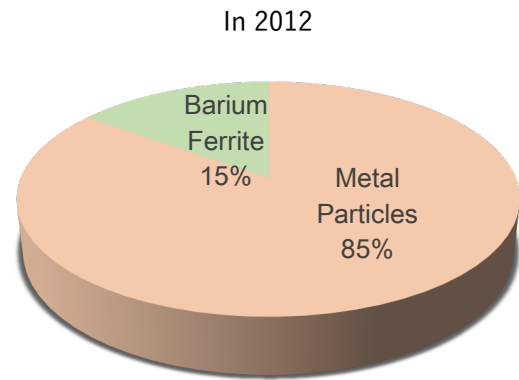
1 000 000



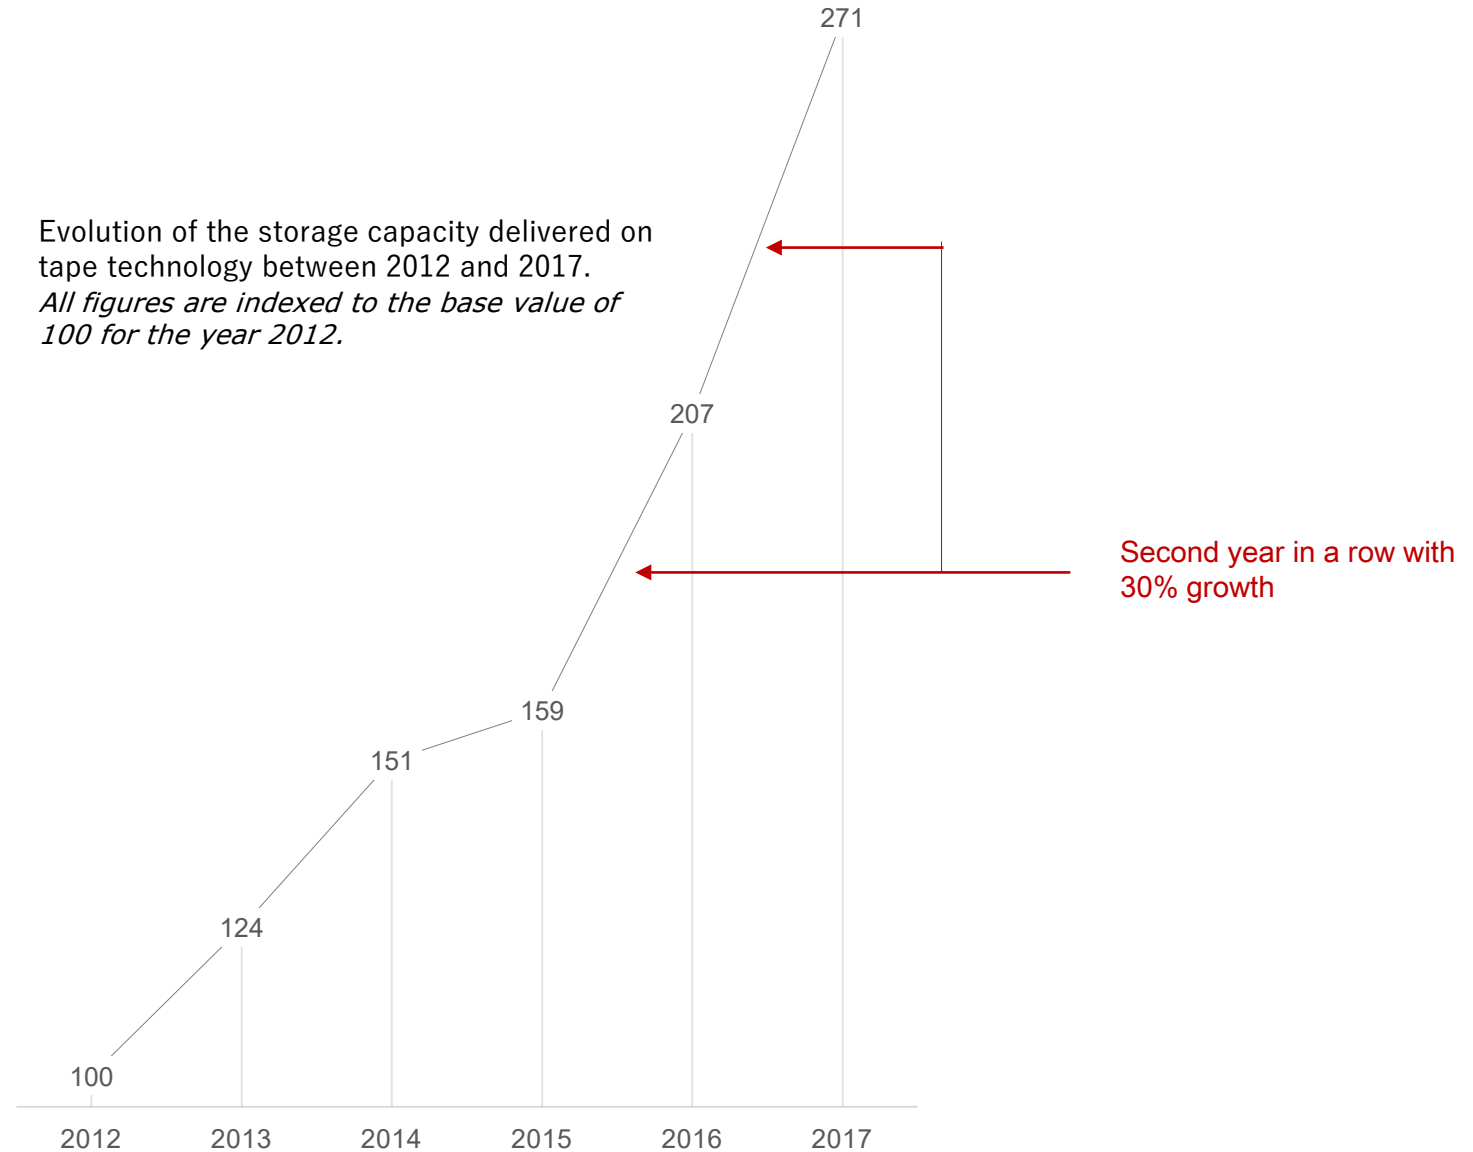
...that makes 1 Million hard disks  
of 1TB capacity

Which tape coating technology is the most requested?

Shown below are the global manufacturing shares of magnetic tapes per technology - all figures are in share % of storage capacity delivered.



Evolution of the storage capacity delivered on tape technology between 2012 and 2017. All figures are indexed to the base value of 100 for the year 2012.



# Which Companies Keep their Data on Tapes?

Your government keeps its data on tape

More than 97% of the 10,000 largest European companies keep their data on tapes

Your pictures on social networks

The school records of your children

The videos you watch on the web

The movies you are watching

Your taxes are recorded on tapes (unfortunately)

Your medical records

The scientific research area keeps its data on tapes

Your private e-mails

Your electricity supplier

Your bank details

TV Channels

Satellite images


The weather forecast

Railways


Car manufacturers

Flag carriers  
Airline transportation

Telecommunication mobile operators



# The 3592 ID Card




# The 3592 Tape Storage Solution in a Few Ideas ....

67% more storage capacity than LTO8

30% faster write speed than LTO8

3X faster write speed than newer hard drive generations

Allows for an intensive use of the drive - LTO does not allow this



10X higher data integrity level than LTO tape (less data loss)

The lifespan of the hardware is longer than for LTO8

A less frequent, easier and faster migration process

A data access system that is far superior to that of LTO

A roadmap established until the 2030s (SrFe technology)


IBM's new TMR-Terzetto write/read head

Cost of use - 3592 reduces the hard disk capacity needed to access data

The message we have received from users is that...


...the amount of new data created will multiply by 10 over the next 10 years

We need to rapidly increase the storage capacity of tape technology



Source: IDC's Data Age 2025 study, sponsored by Seagate, March 17

...87% of this data will have to be protected




- Does not require any data protection
- Requires data protection
- Data that is not yet protected

- Archive life
- Data Integrity
- Protection vs viruses and hackers

- Requires data protection
- Data that is already protected

Source: Storage Newsletter April 7, 2017

... this data is all types of sizes. Despite this, users need to be sure that they can access it.



- Optimize the backup time for small files
- Enable frequent access to the data





Increasing the  
Storage Capacity  
depends on  
Securing the Data  
Integrity

1





Barium Ferrite

# The Tape Market Segmentation




# TS 4500 = The Best Capacity / Floor Space Ratio in the IT Market



**4.23 m<sup>2</sup> of floor space occupancy can carry 11 PB of data (550 slots).**

Speed: this library can produce a transfer rate of up to 4.8GB/s (up to 12 TS1160 drives).


TS 4500 = The Best Capacity / Floor Space Ratio in the IT Market



6.29 m<sup>2</sup> of floor space occupancy can carry 31 PB of data (1 550 slots).

Speed: this library can produce a transfer rate of up to 11.2GB/s (up to 28 TS1160 drives).

TS 4500 = The Best Capacity / Floor Space Ratio in the IT Market



Speed: this library can produce a transfer rate of up to 17.6 GB/s (up to 44 TS1160 drives).


TS 4500 = The Best Capacity / Floor Space  
Ratio in the IT Market



The maximum capacity of a TS 4500 library:

- ✓ 351 PB of data
- ✓ 51.2 GB/s transfer rate

TS 4500 configuration: you can build your library by using three different types of frames



FRAME "L" TYPE – "L" stands for Library

- It is the Base Library or the Base Frame
- In position 1: it can hold up to 550 tape cartridges and 12 drives.
- In any other position: it can hold up to 660 tape cartridges and 16 drives.


FRAME "D" TYPE – "D" stands for Drives

- It is an additional Frame or Expansion Frame
- In position 1: it can hold up to 590 tape cartridges and 12 drives.
- In any other position: it can hold up to 740 tape cartridges and 16 drives.

FRAME "S" TYPE – "S" stands for Storage

- It is an additional Frame or Expansion Frame
- In position 1: it can hold up to 798 tape cartridges.
- In any other position: it can hold up to 1000 tape cartridges.
- However, it does not contain any drives.

# TS 4500 configuration: you can build your library by using three different types of frames (II)



### FLEXIBLE CAPACITY


- The TS4500 is scalable, and can expand to the right, or to the left, depending on the client's choice.
- Depending on the location of the frame, the maximum number of drives and cartridges that the library can hold may vary, as some positions become unusable (please see the table below).

		L25	D25	S25
Max capacity of the frame in position 1 (leftmost)	Nb of cartridges	550	590	798
	Nb of drives	Up to 12 drives		Only tape cartridges
Max capacity of the frame in position 2 (to the right of the basic cabinet)	Nb of cartridges	660	740	1 000
	Nb of drives	Up to 16 drives		Only tape cartridges

### DIFFERENT CAPACITIES PER FRAME

Why do different frames have **two different capacities** depending on their location within the Library?

Regardless of the frame model, L, D or S, the frame that will be placed in position 1 (leftmost) will have its capacity reduced, as part of the frame will not be reachable by the arm of the library. Please see on the drawing below, the portion of the library that becomes unusable when placed in position 1.






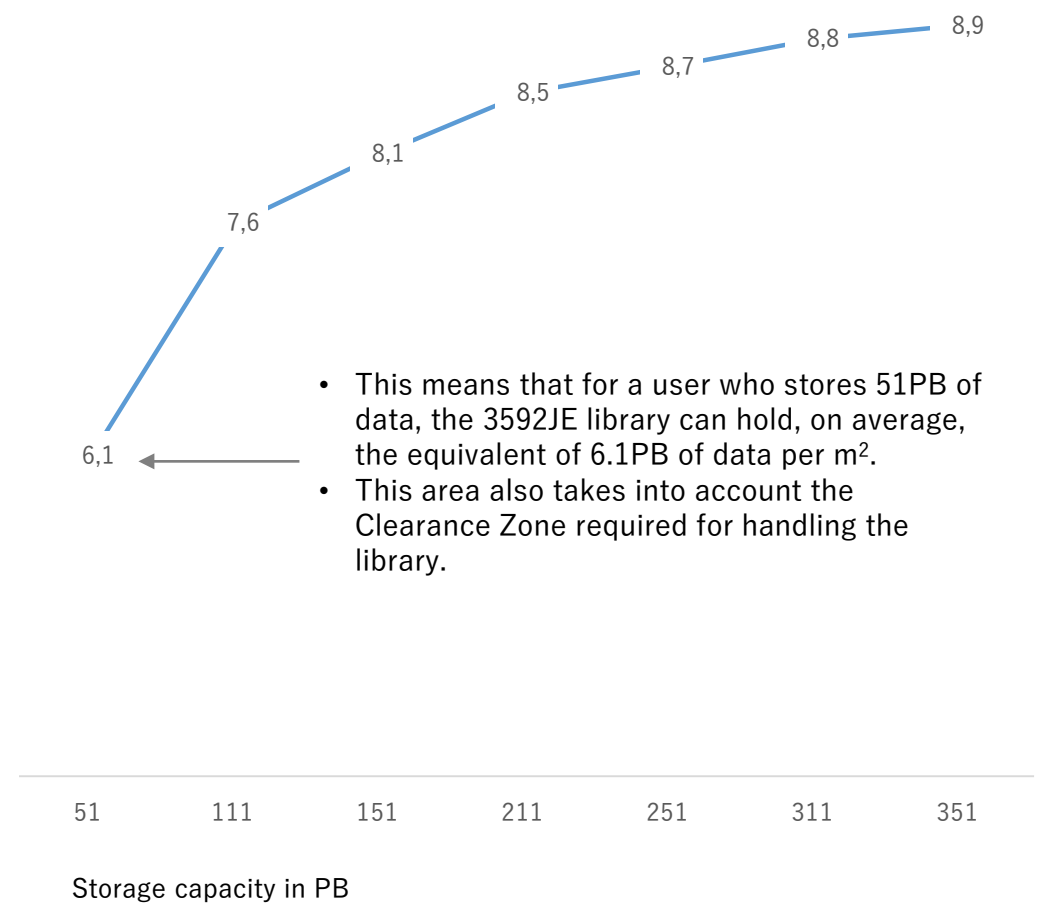
# TS 4500 = Floor Space Estimation per Storage Capacity



Floor Area - what square footage in m<sup>2</sup> is necessary to hold 5 to 30PB of data on a TS4500 library?



What storage capacity in PB can we get per m<sup>2</sup> on a TS4500 / 3592JE library?



- This means that for a user who stores 51PB of data, the 3592JE library can hold, on average, the equivalent of 6.1PB of data per m<sup>2</sup>.
- This area also takes into account the Clearance Zone required for handling the library.



You can generate your own estimate of the floor space that your TS4500 library would occupy.

TS4500	Nb slots max per frame (jaguar)	Cumulated Slots (jaguar)	Size in m <sup>2</sup> (w/o service clearance)	Size in m <sup>2</sup> (with service clearance)	Capacity with 3592JE
Base	550	550	0.95	4.23	11
1 exp	1 000	1 550	1.86	6.29	31
2 exp	1 000	2 550	2.78	8.36	51
3 exp	1 000	3 550	3.69	10.42	71
4 exp	1 000	4 550	4.61	12.49	91
5 exp	1 000	5 550	5.52	14.55	111
6 exp	1 000	6 550	6.44	16.62	131
7 exp	1 000	7 550	7.35	18.68	151
8 exp	1 000	8 550	8.26	20.75	171
9 exp	1 000	9 550	9.18	22.81	191
10 exp	1 000	10 550	10.09	24.88	211
11 exp	1 000	11 550	11.01	26.95	231
12 exp	1 000	12 550	11.92	29.01	251
13 exp	1 000	13 550	12.84	31.08	271
14 exp	1 000	14 550	13.75	33.14	291
15 exp	1 000	15 550	14.67	35.21	311
16 exp	1 000	16 550	15.58	37.27	331
17 exp	1 000	17 550	16.50	39.34	351

HOW TO READ THIS TABLE

- With a configuration combining:
  - \*One Base Frame "L" on position 1 (550 slots)
  - \*Two Expansion Frames "S" on other positions (1000 slots each)
 It is possible to set a Library that can offer up to 51PB storage capacity.
- In the table opposite, we can see that the area occupied by the library is 8.36m<sup>2</sup>.
- We can also notice that a TS4500 library can consist of up to 18 Frames, and can therefore contain 350PB of data.
- This type of configuration implies that the user does not need more than 12 drives. Otherwise, the user would have required a "D" type of Expansion Frame that could have complemented the user's drive requirements. An Expansion Frame type "D" can contain up to 740 slots.

# How to Store your Data in a Limited Amount of Space?





# How to Store your Data in a Limited Amount of Space (II)?

The 3592JE compared to older LTO tape generations.

A 3592JE tape cartridge can contain:

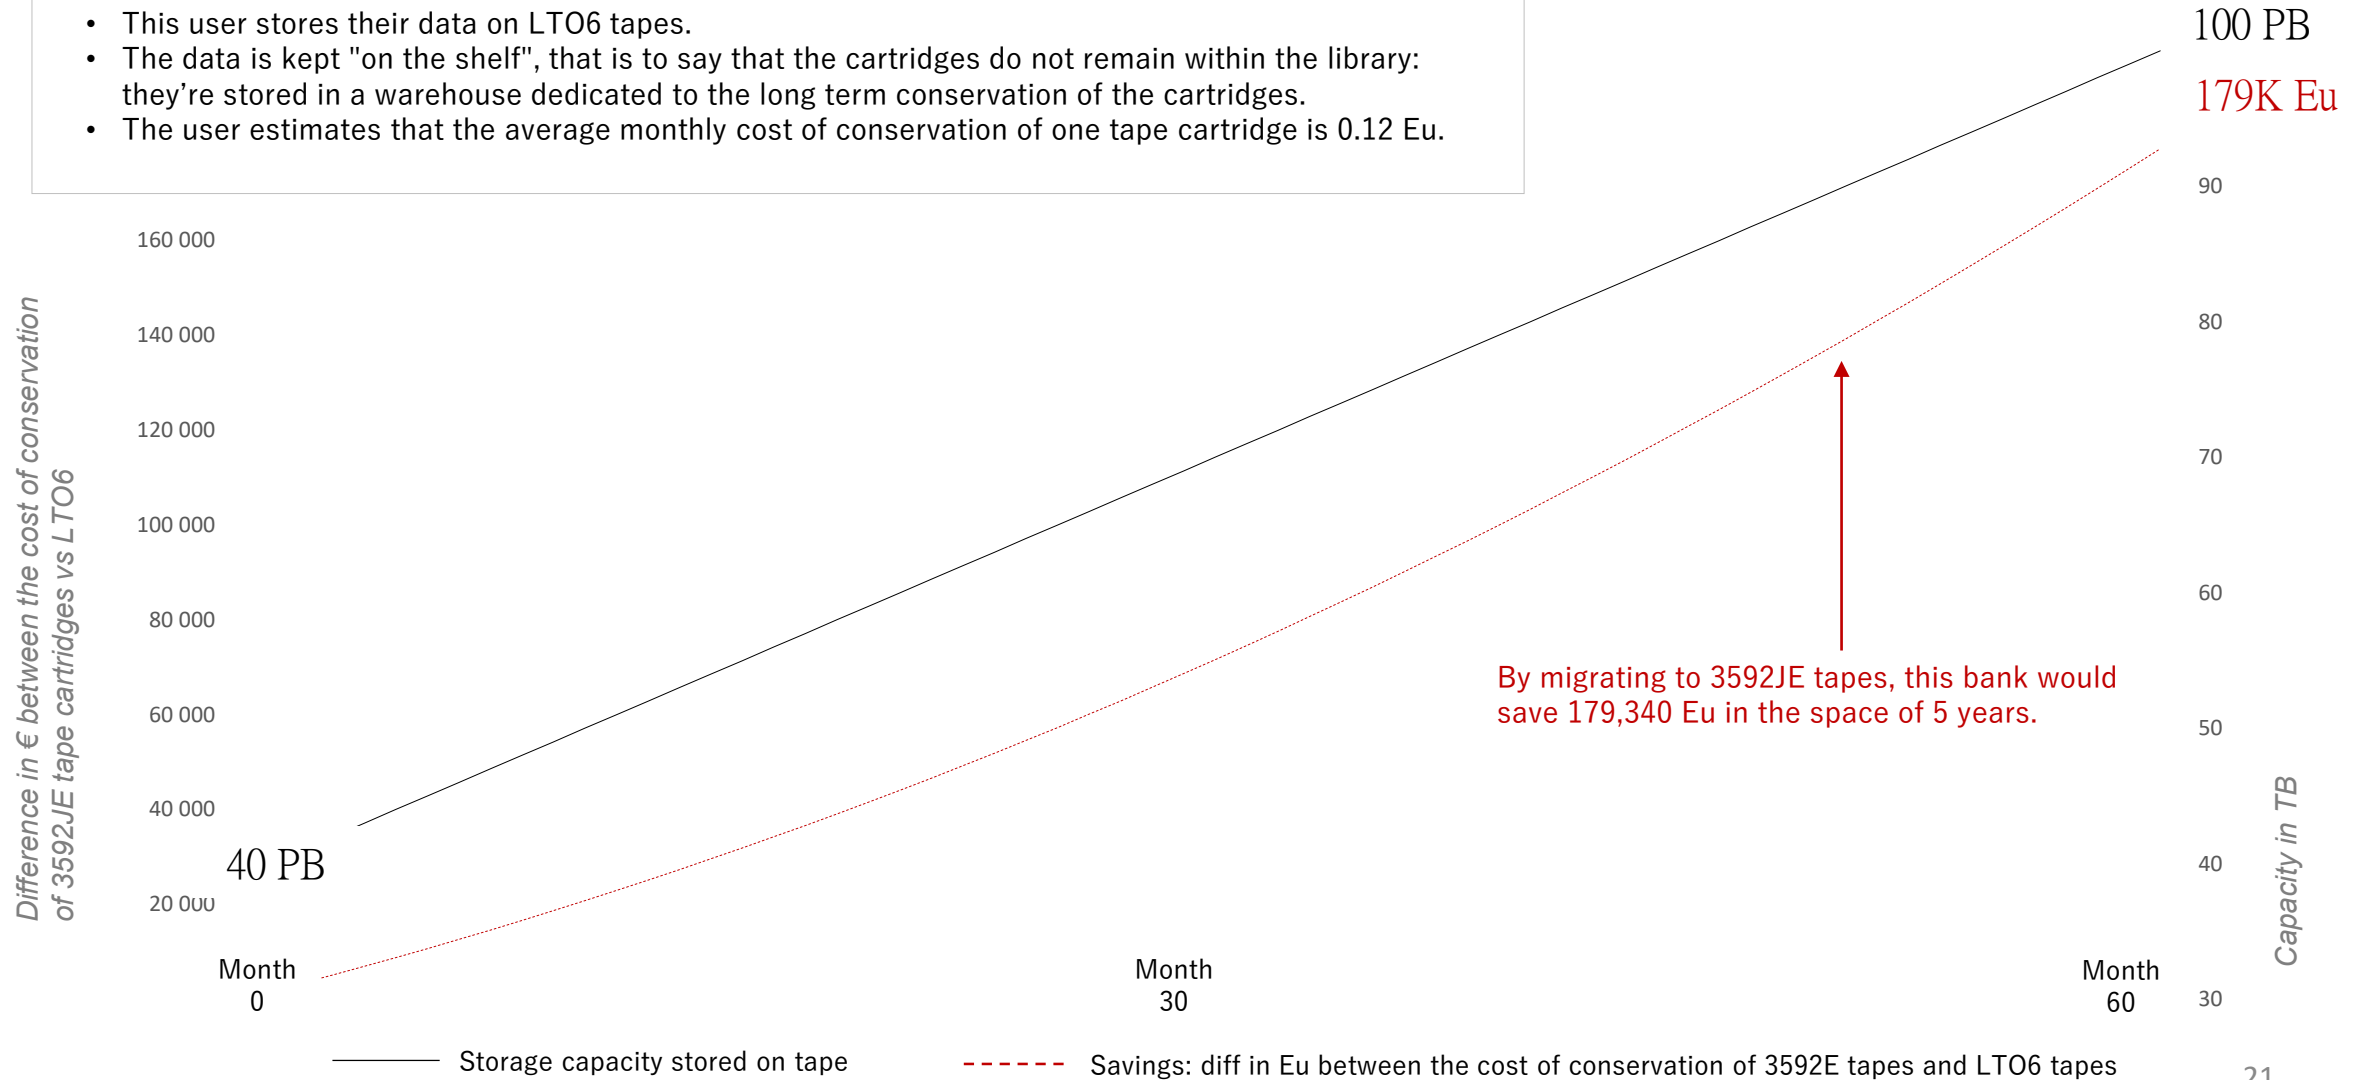
- 8 x LTO6 tapes
- 13 x LTO5 tapes
- 25 x LTO4 tapes



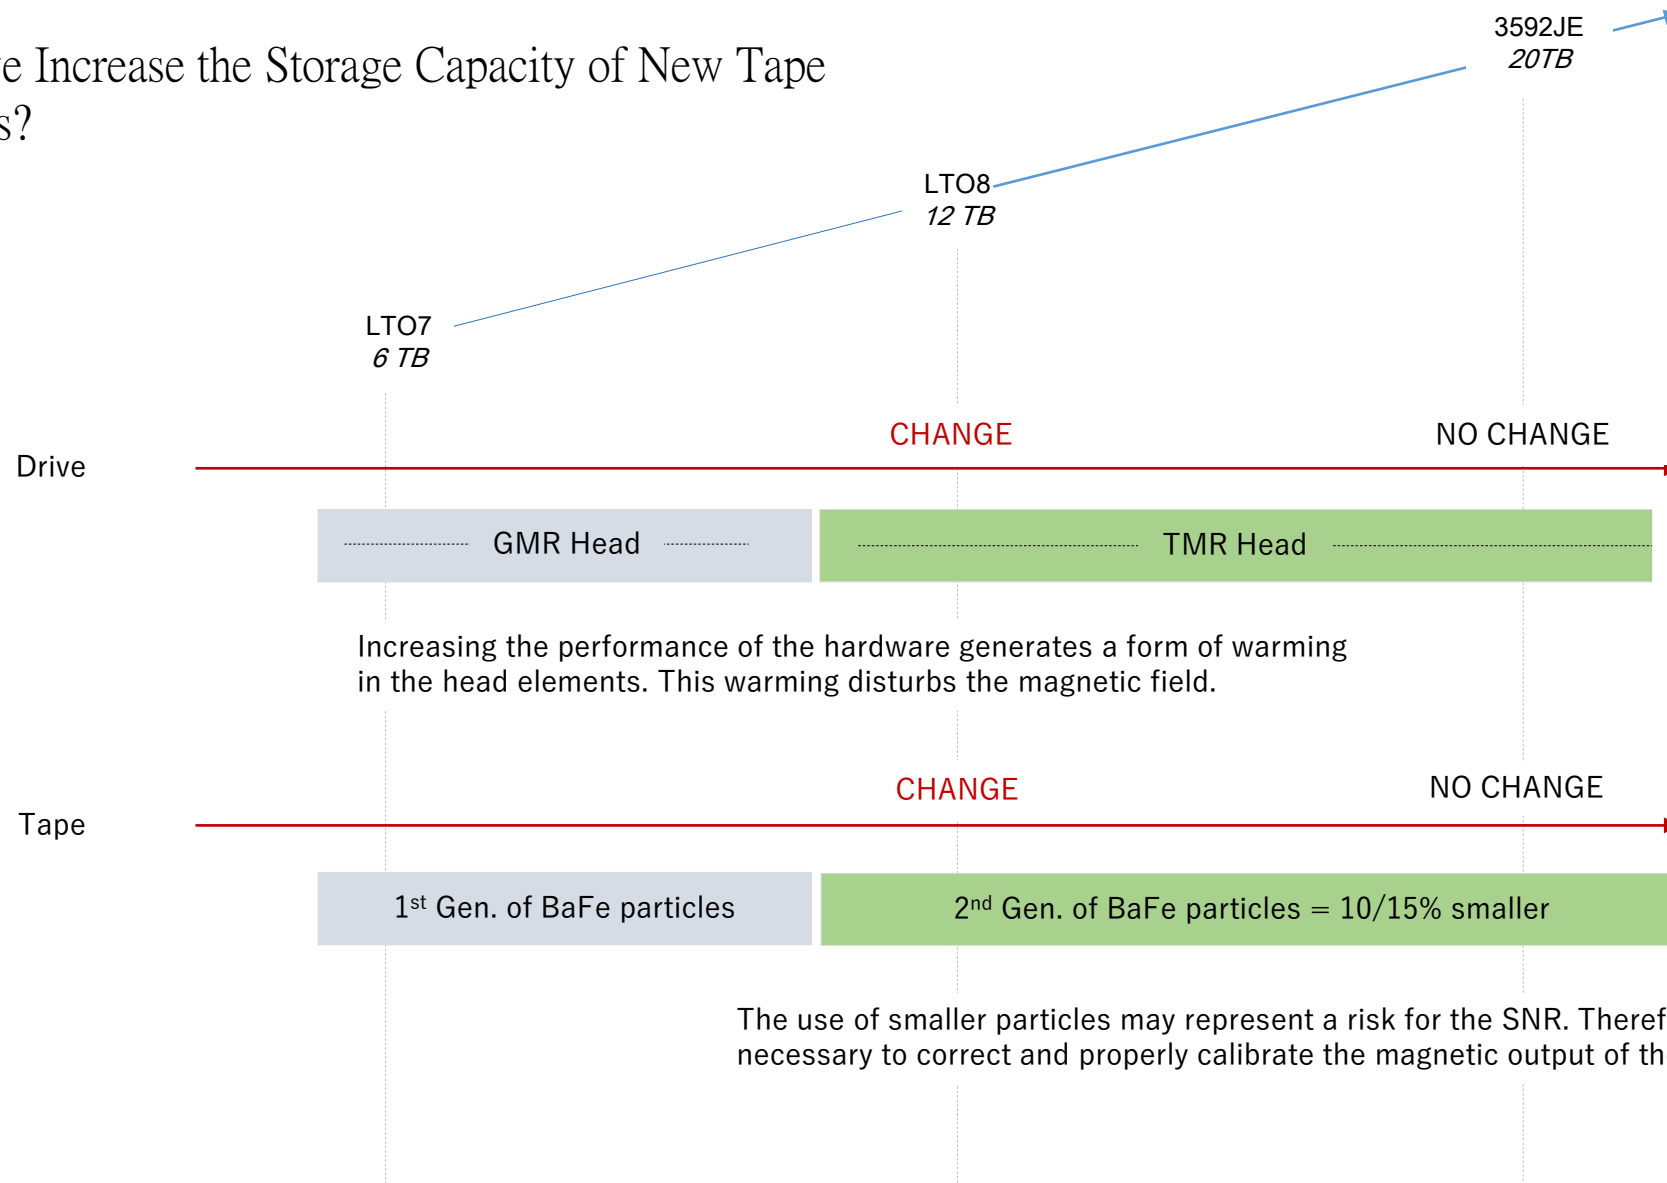
# Floor Space Reduction - Migrating from LTO6 to 3592JE = Substantial Savings



- A European bank will see its data grow from 40PB to 100PB within 5 years. The average monthly creation of new data: 1PB
- This user stores their data on LTO6 tapes.
- The data is kept "on the shelf", that is to say that the cartridges do not remain within the library: they're stored in a warehouse dedicated to the long term conservation of the cartridges.
- The user estimates that the average monthly cost of conservation of one tape cartridge is 0.12 Eu.





# How can we Increase the Storage Capacity of New Tape Generations?

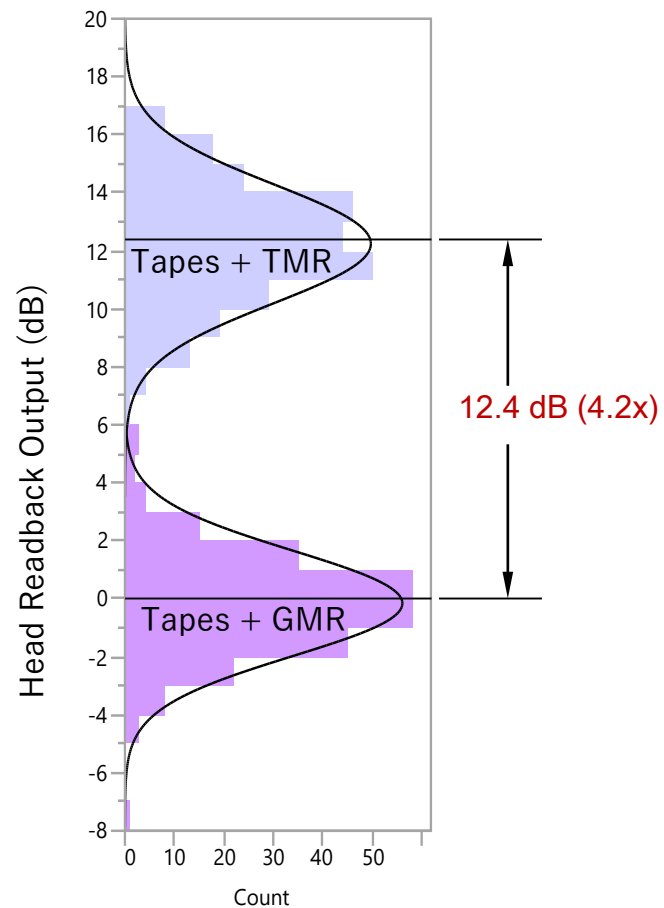


# The TMR Head Generates Spectacular Progress in the Field of Data Integrity (I)

Beyond a certain level of system performance, the current flow can hit the SNR.



## The TMR Head Generates Spectacular Progress in the Field of Data Integrity (II)



The TMR-Terzetto head captures 4X more signals than the GMR-Terzetto head.

- The higher magnetic output enables a better reading of the data.
- The TMR head improves the level of SNR.
- The higher performance level of the TMR, therefore, allows the manufacture of tapes of over 15TB of data.


This result is obtained by testing both heads under similar conditions:

- The same test platform
- The same size populations
- The same track width, bias
- The same tape type
- Etc...etc...




# SIGNAL-TO-NOISE RATIO

## WHAT CAN HURT THE SNR ...



MP Particles



Horizontal polarization




Time (long-term)



High speed


## WHAT CAN IMPROVE THE SNR ...




IBM's Terzetto head



Barium Ferrite tapes



Vertical polarization



Bring the tape close to the head



## CONSEQUENCES OF A LOW SNR



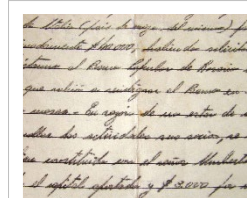
Capacity Loss



Low Read/Write speed




Lifespan of the Hardware



Read & Write errors

MP particles are larger in size: in absolute terms, their Magnetic Output is higher than that of BaFe particles

BaFe's SNR is superior to MP's since:  
 \* The clear noise is stronger  
 \* The noise pollution is weaker




MP Particles

BaFe Particles  
1st generation

The very nature of Nanocubic Barium Ferrite Technology allows a constant increase of the SNR

- Smaller particles combined with Nanocubic technology allow for thinner and smoother tape layers.
- This results in a significant reduction of the noise pollution and an improvement of the SNR.
- "The ideal would be to use only one row of particles on one layer of tape"

- Dr.Schmeinck-

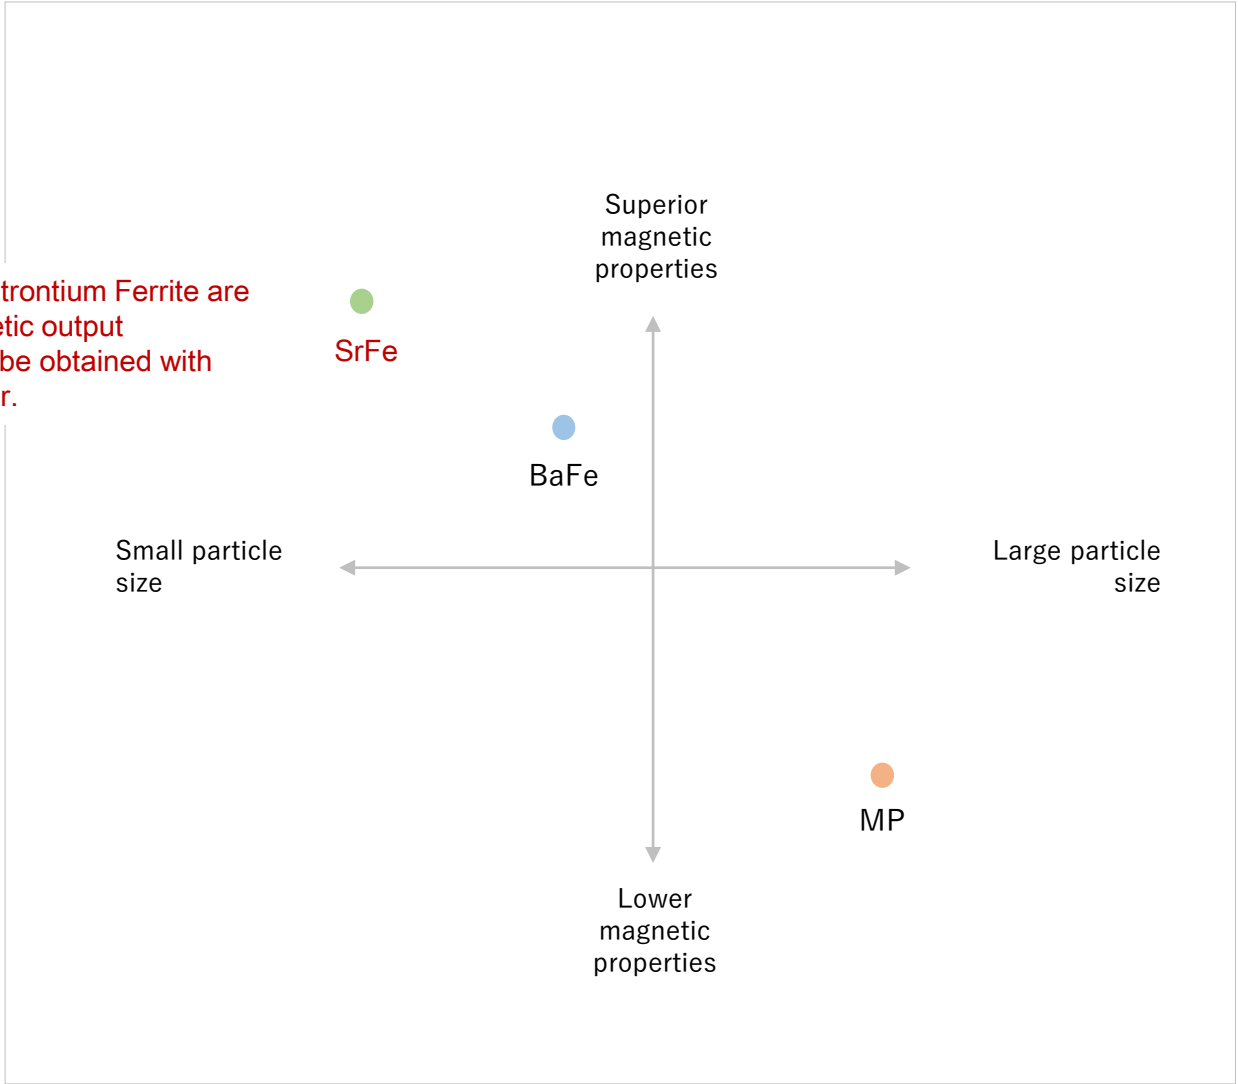



The reduction of noise pollution increases the total SNR, despite the smaller particle size.

BaFe Particles  
2nd generation


# Strontium Ferrite - Smaller Particles with Higher Magnetic Output

The magnetic properties of Strontium Ferrite are such that they allow a magnetic output equivalent to that of BaFe to be obtained with particles that are 40% smaller.




## Evolution of Particle Size

Barium Ferrite  
First generation




- Current BaFe particles (used for LTO7, for instance)
- Volume of a particle:  $1900 \text{ nm}^3$

Barium Ferrite  
220TB record



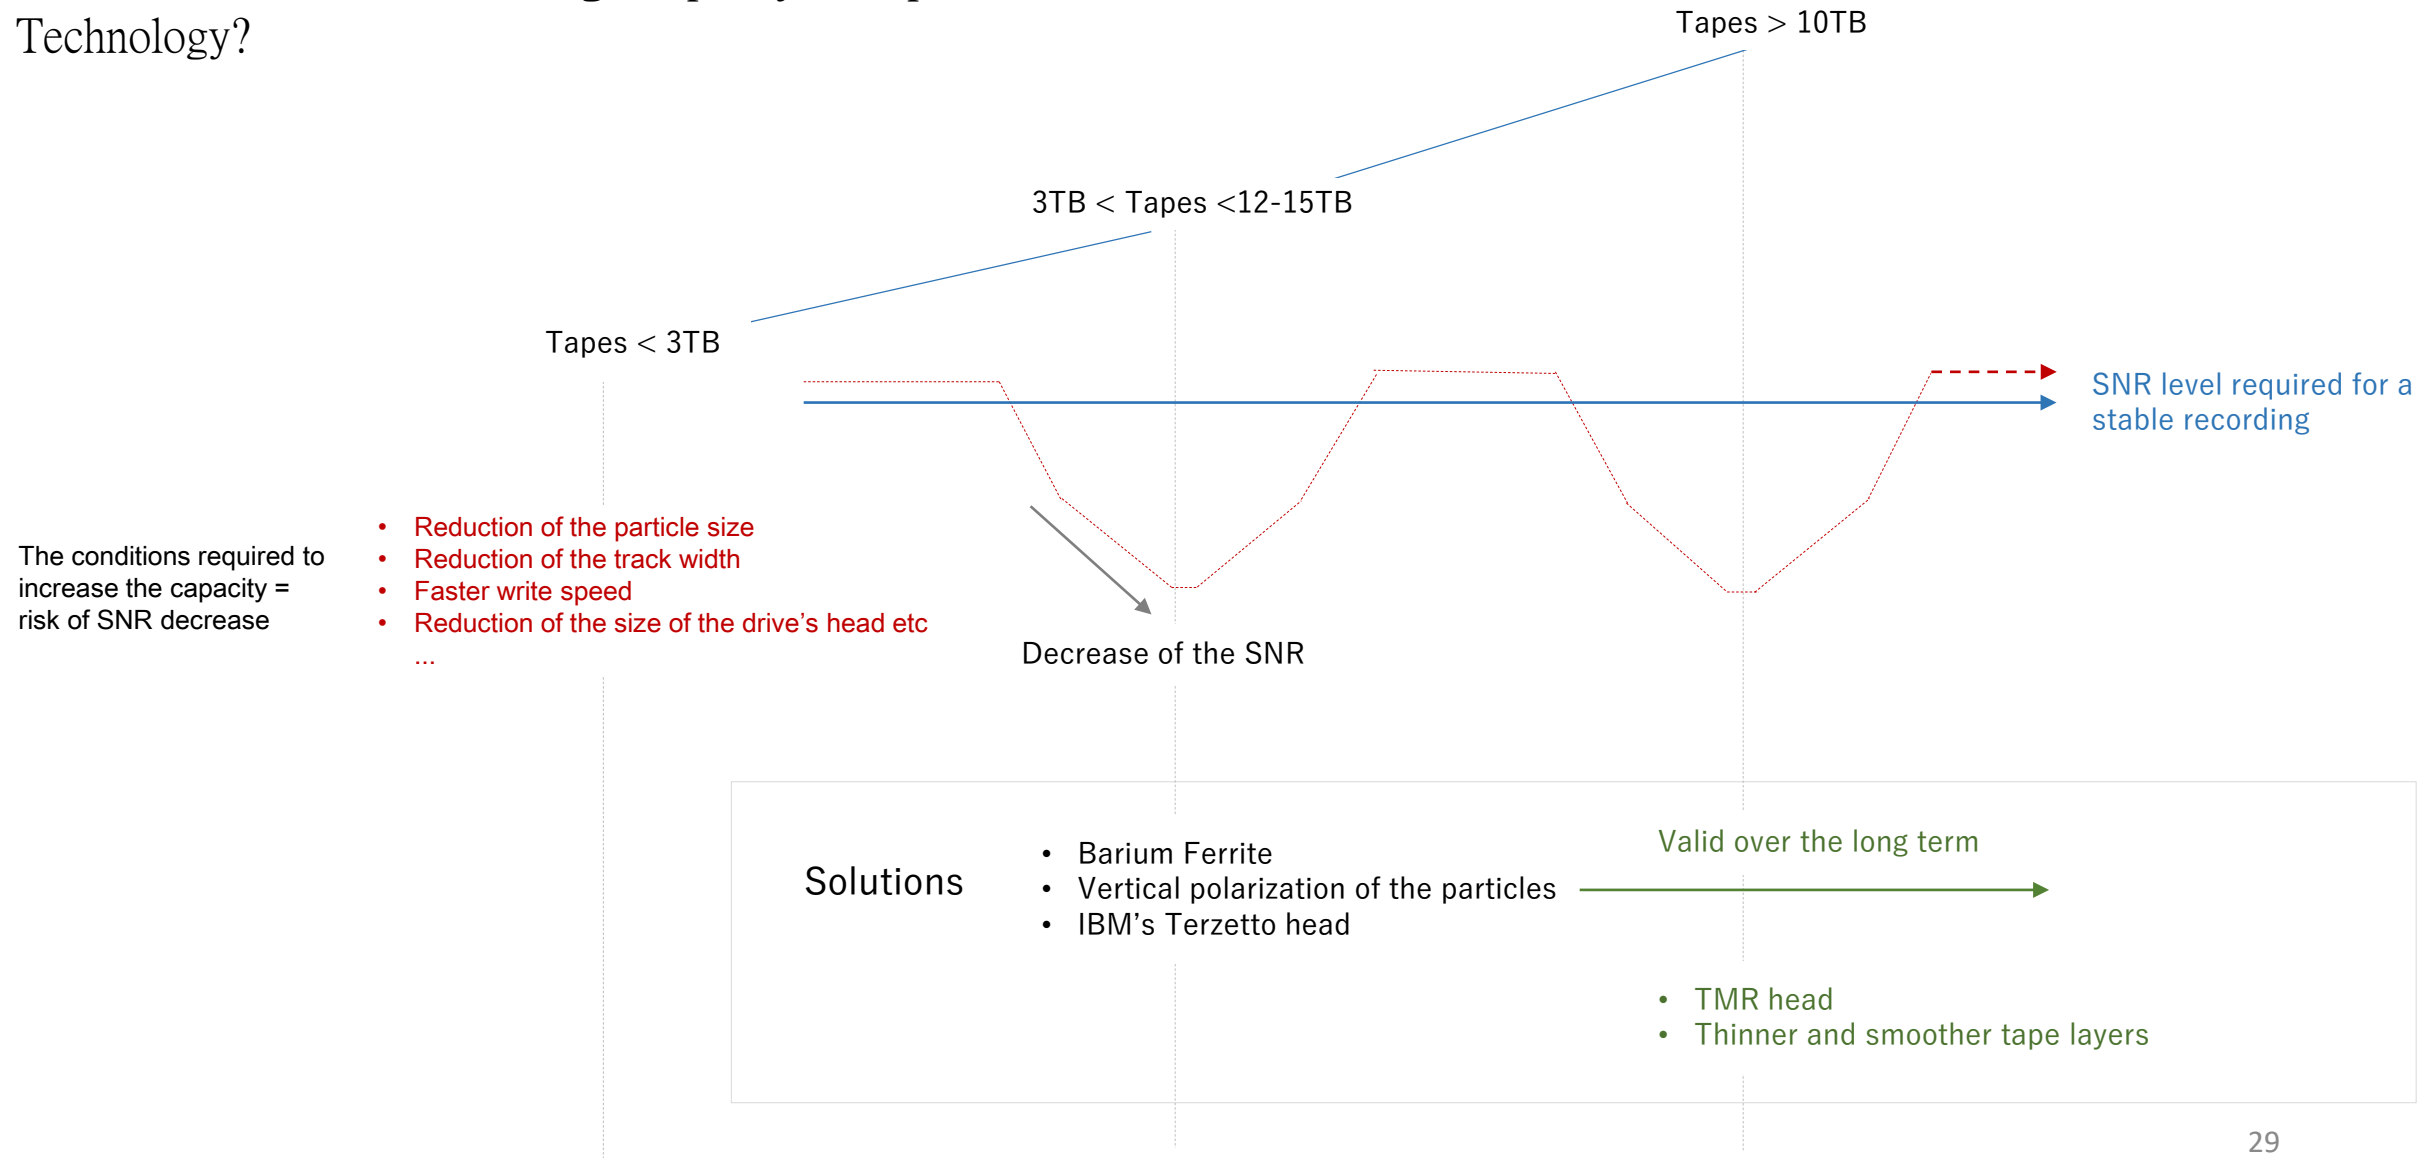
- BaFe particles used for a technical demonstration in 2015 (220TB for one single tape cartridge)
- Volume of a particle:  $1600 \text{ nm}^3$

Strontium Ferrite  
First generation




- Strontium Ferrite particles
- Volume of a particle:  $900 \text{ nm}^3$
- Capacity = 400TB for a single tape

# How can we Increase the Storage Capacity of Tape Technology?



# Our Vision of the Future, in 5 Phases ...




THE DAY BEFORE  
YESTERDAY

The period prior to  
Barium Ferrite: LTO4,  
LTO5, LTO6 until 2012.




YESTERDAY

Beginning of the Barium  
Ferrite era: 6 TB (LTO7)  
for SMEs, 15 TB (3592JD)  
for large companies.




TODAY

Breaking new records with  
the 3592JE: 20 TB and  
400MB/s, LTO8 offers 12 TB  
and 360 MB/s.



TOMORROW

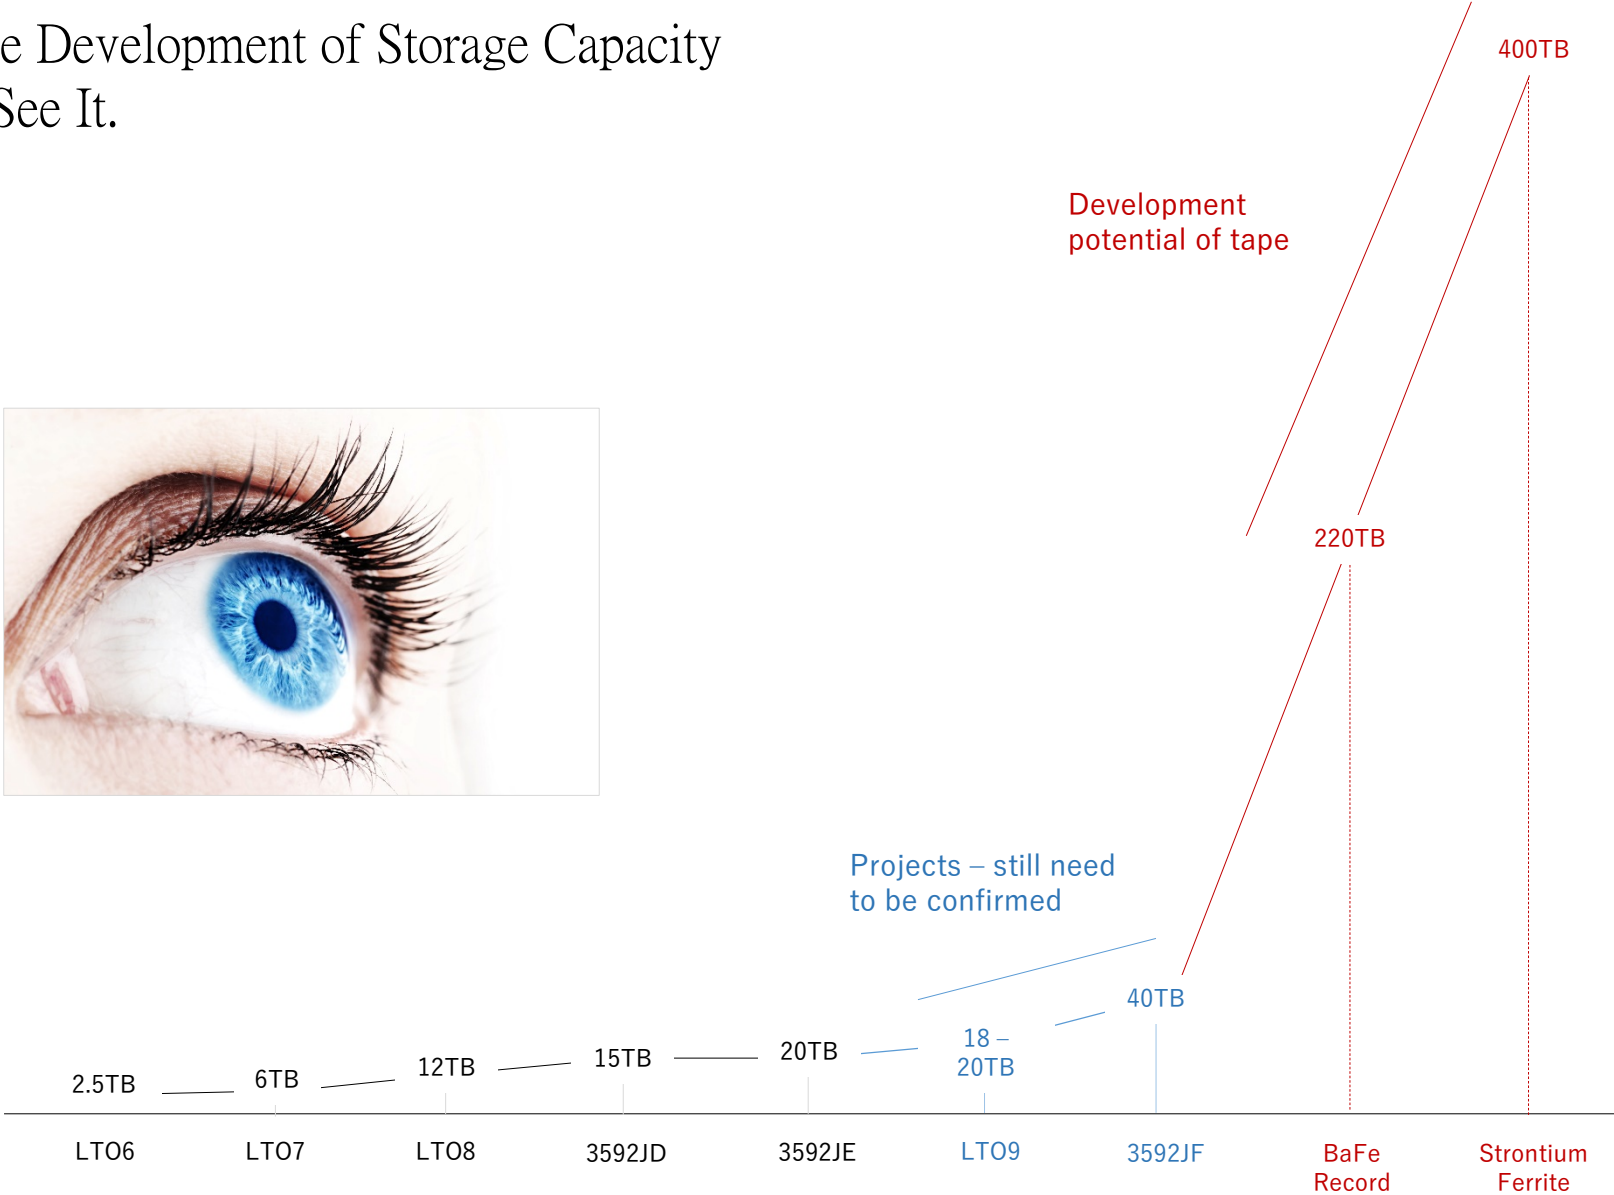
Development of tapes of more  
than 50 TB-60 TB, planned for  
2022-2023.




THE 2030s

Use of Strontium Ferrite  
technology: tapes with native  
capacities above 100 TB!

IBM + Fujifilm. The Development of Storage Capacity as our R&D Team See It.




# What are the BER (Bit Error Rates) of Tape Technology and Hard Disk?



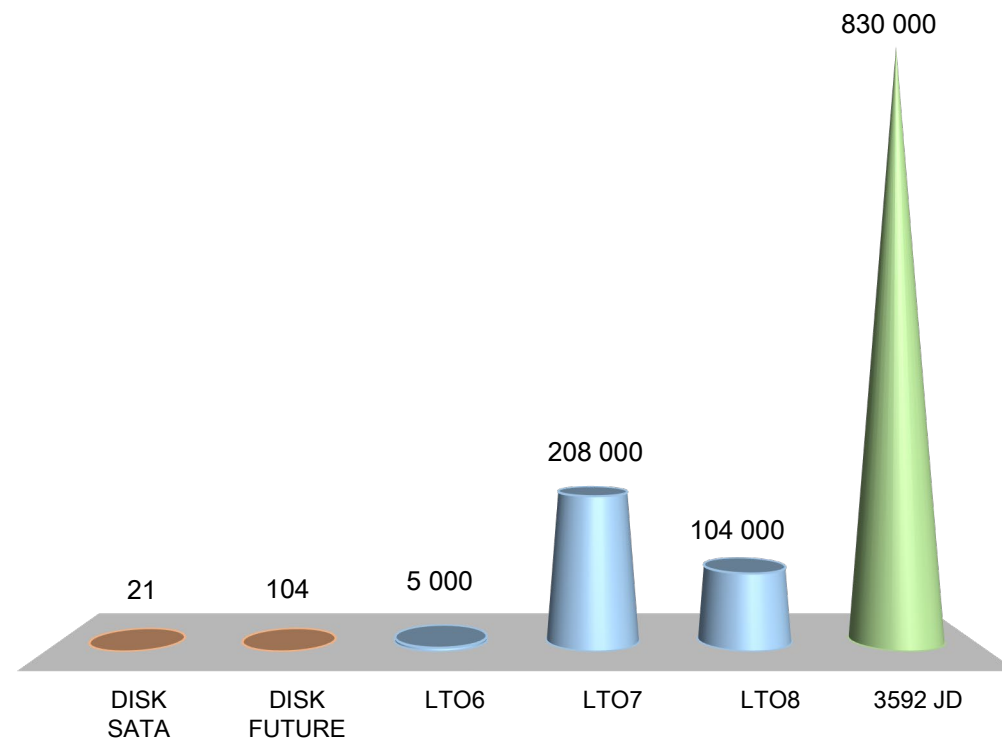


# How can we Translate the BER into Concrete Information?

What storage capacity can you save before risking a write error?

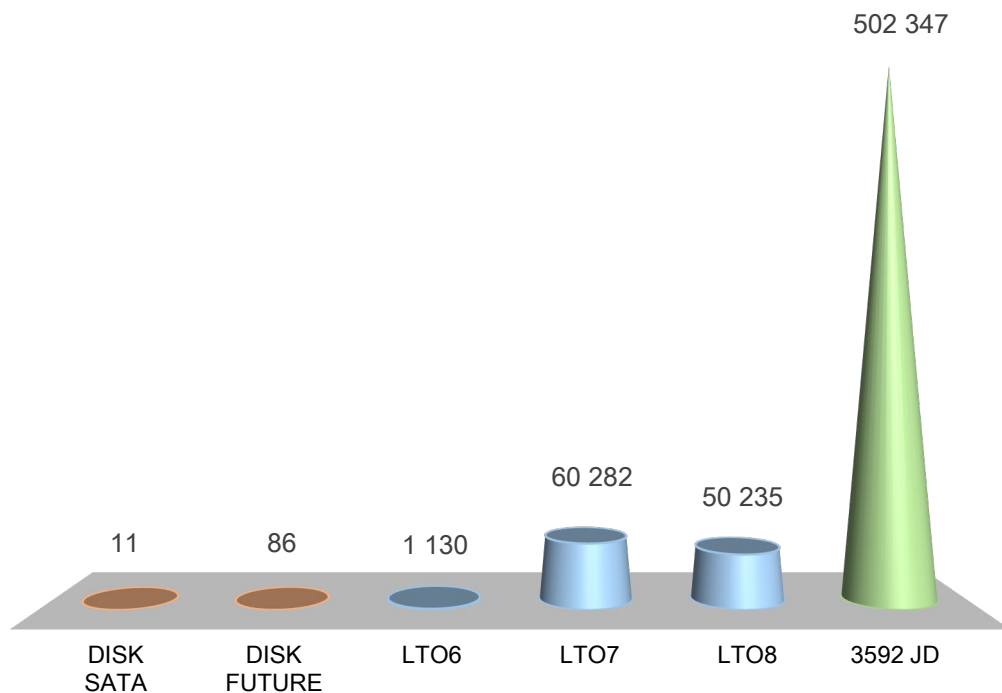


How many tape cartridges or 6TB Hard Disks can we use before risking a write error?



## How can we translate the BER into Concrete Information (II)?

The number of days of full use of the drive or disk before the risk of a write error.



### ATTENTION:

- The data shown in the table opposite are statistical test results.
- Only the empirical use of the system makes it possible to obtain the closest result to reality.
- However, these tests have the merit of highlighting the big difference in reliability between 3592 technology and other data storage solutions.

The Write Speed  
will Reach  
400MB/s ...



2


# What is the Real Backup Time on Tape Technology?

1 / When addressing the question of write speed, what users ask us is not necessarily to increase the official transfer rate.

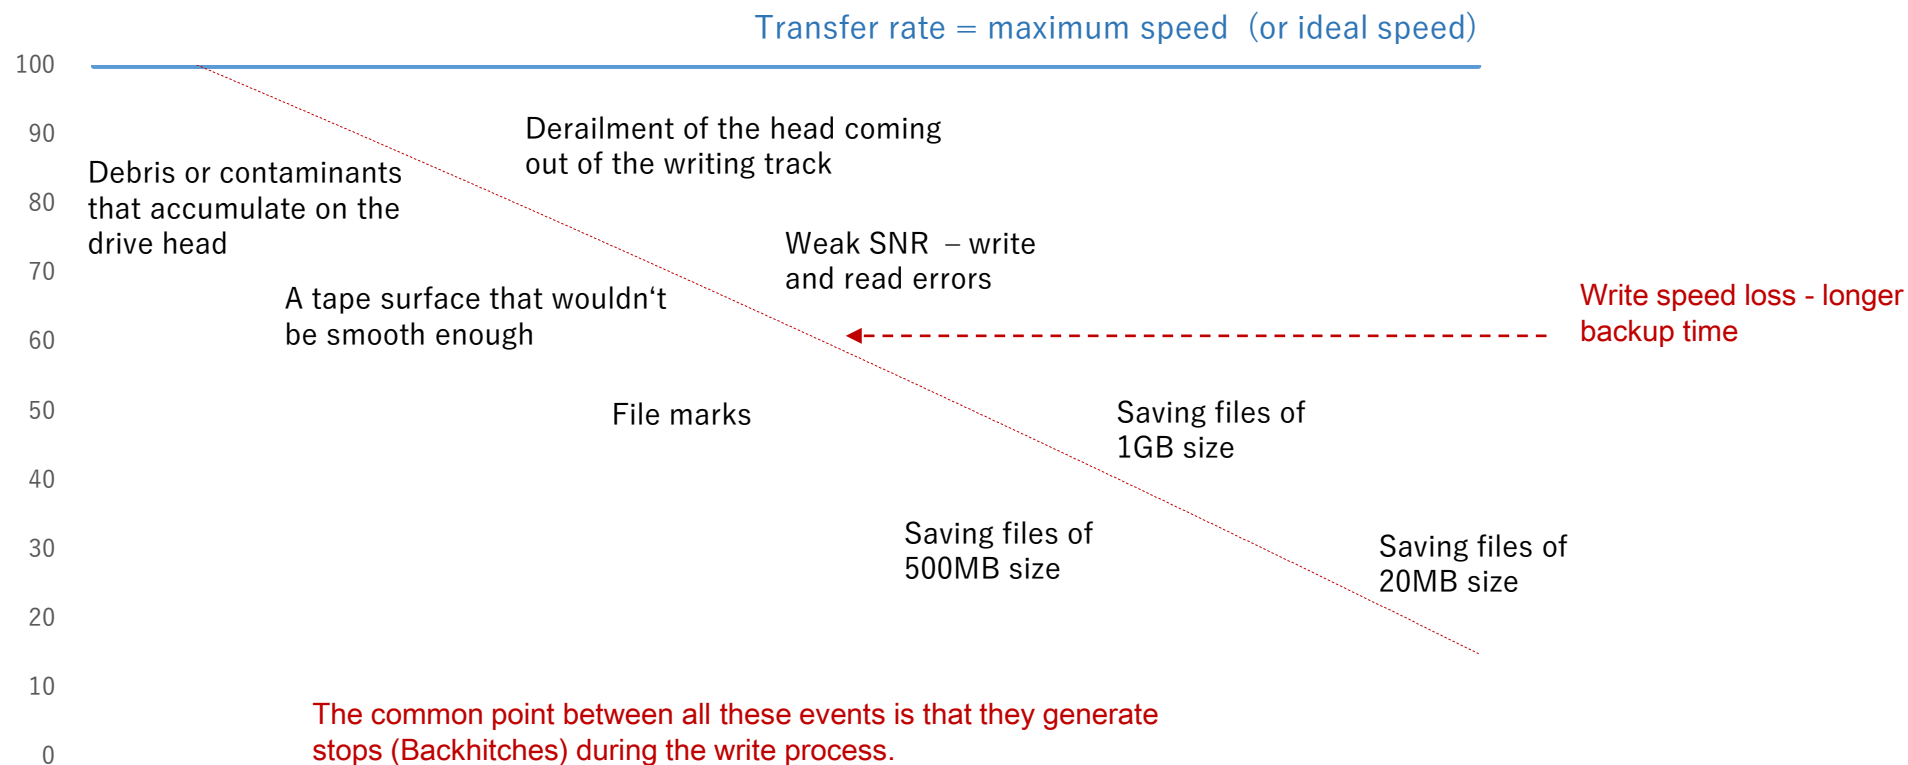
2/ During the writing process, a number of phenomena will hit the backup time. In the end, the actual backup time is often longer than that promised by the official transfer rate. Please see the test on the opposite graph.

3/ What users ask us, in priority, is to prevent the degradation of the operational write speed as much as possible, and to reduce the backup time.

✓ We know the official transfer rate. On the other hand, what is the real operational speed of recording data on a tape?



# What is the Real Backup Time on Tape Technology?



## General Principles on the Writing Speed

1. The official transfer rates of 3592JE, 3592JD and LTO8 are almost comparable:

\*3592JE: 400MB/s

\*3592JD: 360MB/s

\*LTO8: 360MB/s

2. However, when looking at the operational writing speeds of 3592, we see larger differences. It is indeed more realistic to estimate the backup time by considering the following theoretical speeds:

\*3592JE: 280MB/s

\*3592JD: 252MB/s

\*LTO8: 216MB/s

3. The basic principle is that, during the backup process, a number of events will mechanically slow down the backup time. The fastest system is the one that limits the loss of time during the backup process the most.

4. This illustrates that, despite a transfer rate that is only 11% higher than that of LTO8, the actual writing speed of 3592JE is:


- 30% faster than LTO8 for 1GB capacity files.
- More than 50% faster for files smaller than 200MB.



„The fastest system is the one that most limits the loss of time during the backup process “

# Official Transfer Rate vs Operational Writing Speed


When saving 1GB files



# Evolution of the Operational Writing Speed

When saving 1GB files

The priority of 3592 = reduce the speed degradation when saving small files.





# How to Optimize the Backup Time?

## Using the Buffer in the Writing Process.

1. During the backup process, the system must prepare each file for writing. It must prepare the description of the files: date of recording, title, location etc...



2. Constant stops in the writing process:

- The drive starts writing the file as soon as the description is ready.
- If the file is too small, the drive will save the data too quickly and will have to wait for the next file to be ready to be written.
- We call these stops in the writing process 'Backhitches': a Backhitch lasts 6-7 seconds on average, but can exceed 20s depending on the circumstances.
- Saving a large number of files will generate constant write stops and significantly extend the backup time.

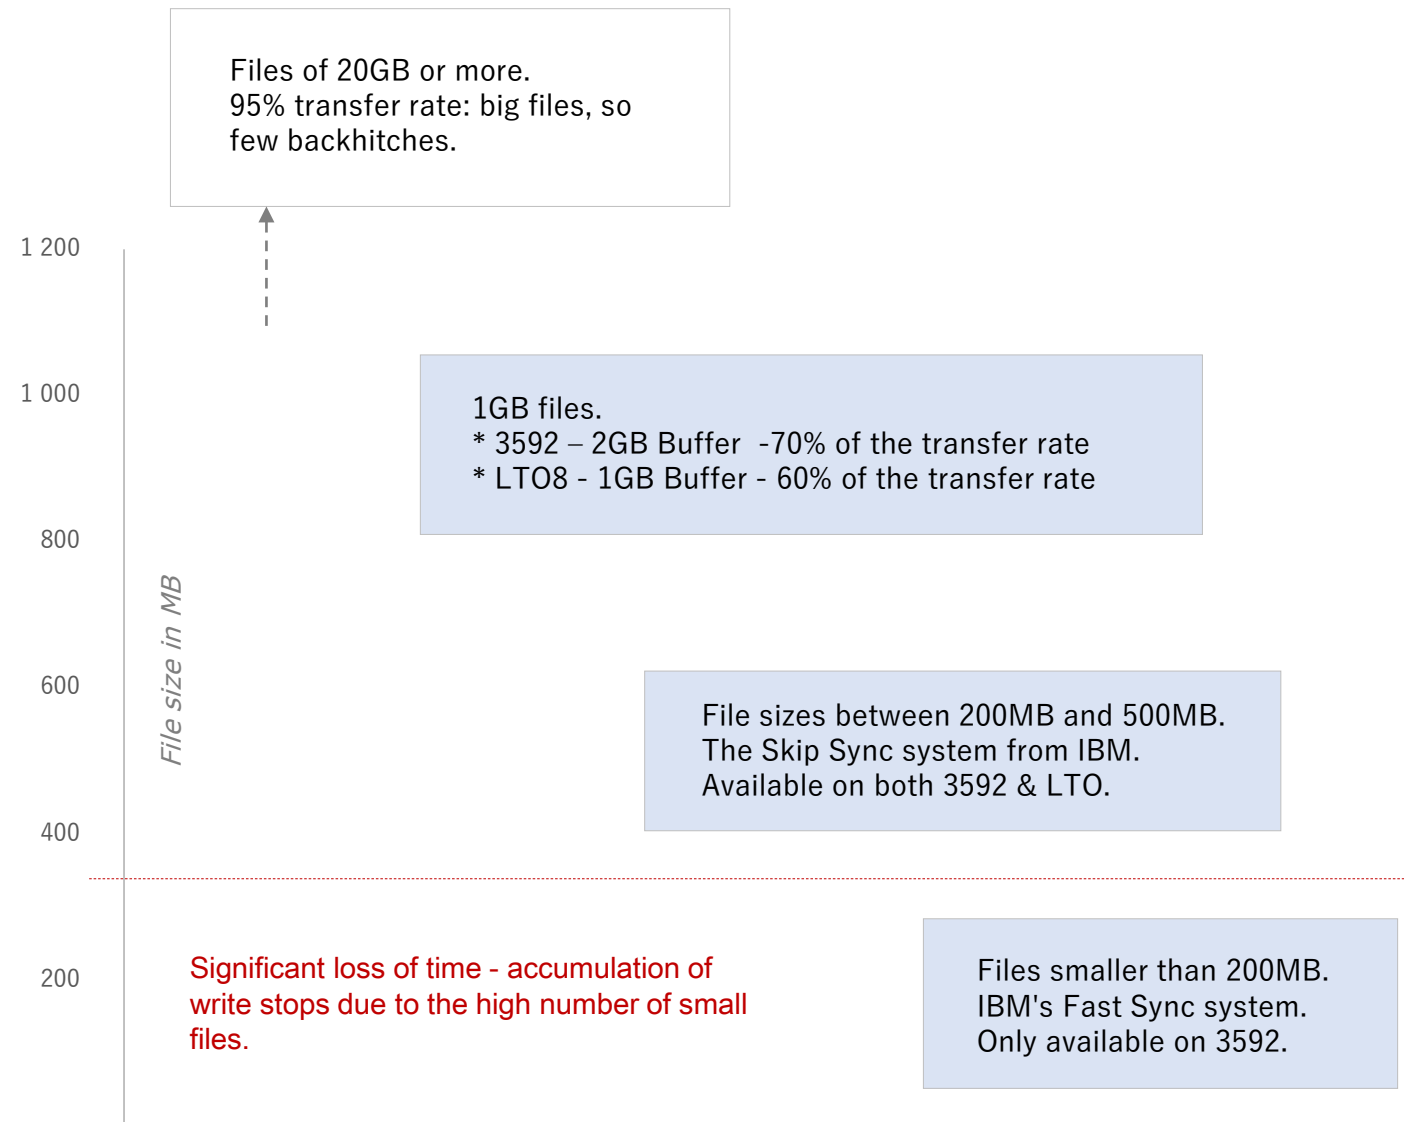
3. Using the Buffer helps to limit the loss of backup time:

- The purpose of a Buffer is to serve as a tank as the data is transferred to and from the tape media.
- The Buffer reserves files for writing: in the event of a write down, the drive will draw new files to write for the Buffer in order to avoid any "Backhitch" and, therefore, to limit the loss of time.
- CAUTION: The Buffer loses its effectiveness when you save files smaller than 500-600MB.

- ✓ The Buffer size of 3592JD or JE is 2GB, while that of LTO8 is 1GB.
- ✓ Concretely, this means that when calculating the backup time of small files (around 1GB), it is better to take the operational speeds as follows into account:
  - \* 70% of the official transfer rate of the 3592 and,
  - \* Only, 60% of the official transfer rate of the LTO8.

# How to Optimize the Backup Time?

1. The Buffer **gradually loses its effectiveness** as the size of the files goes below 500MB.
2. The risk is to notice **a drop in the operating speed below 10-15MB/s** when recording a large number of very small files.
3. The use of new features is, therefore, necessary when saving small files.
4. It is physically impossible to maintain the initial transfer rate in this type of scenario. However, **IBM has developed new features that maintain** a theoretical transfer rate of between 50MB/s and 100MB/s when saving files smaller than 80MB.
5. We have to compare this figure with the performances of the LTO tapes (15-20MB/s) and the hard disk (4-5 MB/s), for these kind of interventions.



## How to Optimize the Backup Time?

### IBM's Skip Sync system

*For writing files of sizes between 200MB and 500MB*

- The Skip Sync is a system developed by IBM.
- The principle of this is to consider that when the system is confronted to a write stop, the restarting of the drive generates an additional loss of a few seconds. **Therefore, an ideal situation is that the drive should never end in a "stopping" mode.**
- In this case, the drive acts like it is writing during the few seconds of stops. In other words, the head keeps moving without saving any new data, in order to avoid the process of restarting the writing process once new files are ready to be recorded.
- On this specific point, both 3592 and LTO technology carry this feature.

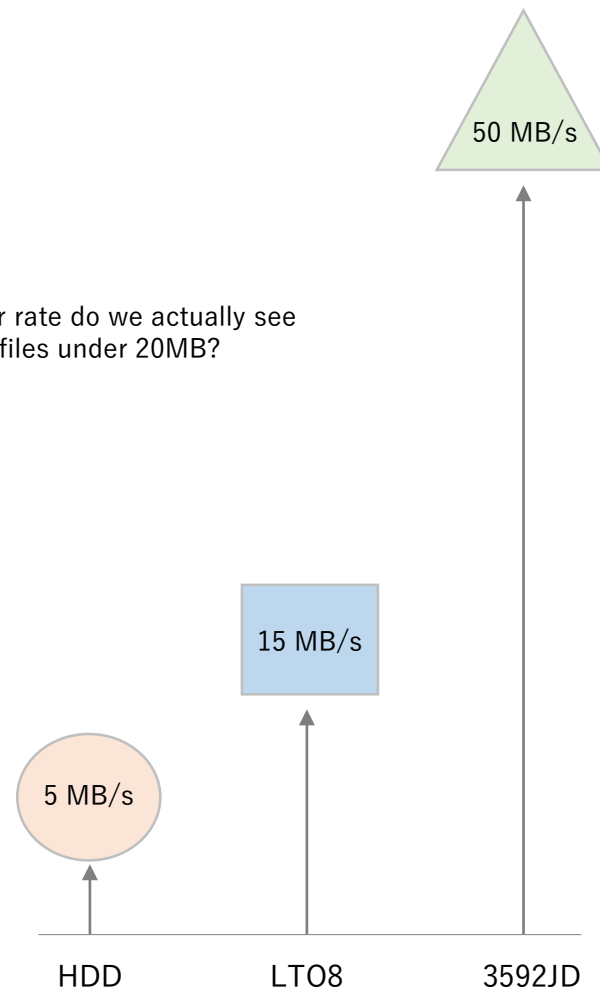
### IBM's Fast Sync system

*For writing files of sizes below 204MB*

- The Fast Sync is a system developed by IBM.
- It follows the same principle as Skip Sync, but with a different solution: the common point is that it must avoid stopping writing, even if the drive is in a waiting situation.
- In this case, the drive will write the data twice. Therefore, it re-writes the data that has already been recorded on another track or wrap of the tape.
- **Fast Sync is a feature that only exists on 3592.** Therefore, the loss of time generated by an LTO8 drive is even more important when saving small files.

# The Backup Time of Files Smaller than 20MB

What transfer rate do we actually see  
when saving files under 20MB?



ADVANTAGES OF IBM'S 3592


- ↑ The Official Transfer Rate:
  - \* 11% higher than LTO8
  - \* 2X higher than hard disk
- +
- ↑ The Buffer Size:
  - \* Twice as large as LTO8
  - \* 16X larger than that of HDD
- +
- ↑ IBM's Fast Sync system

## Transfer Rate + Buffer + Fast or Skip Sync, Oracle had measured the Effect of an Equivalent System on T10000D


Test comparing the write speeds of LTO6 and T10000D for small files:

- The LTO6 drive - transfer rate: 160MB/s - 1GB Buffer
- The T10000D drive - transfer rate: 252MB/s - 2GB Buffer

✓ Please note that T10000D includes a system called File Sync Accelerator, comparable to IBM Fast and Skip Sync systems.

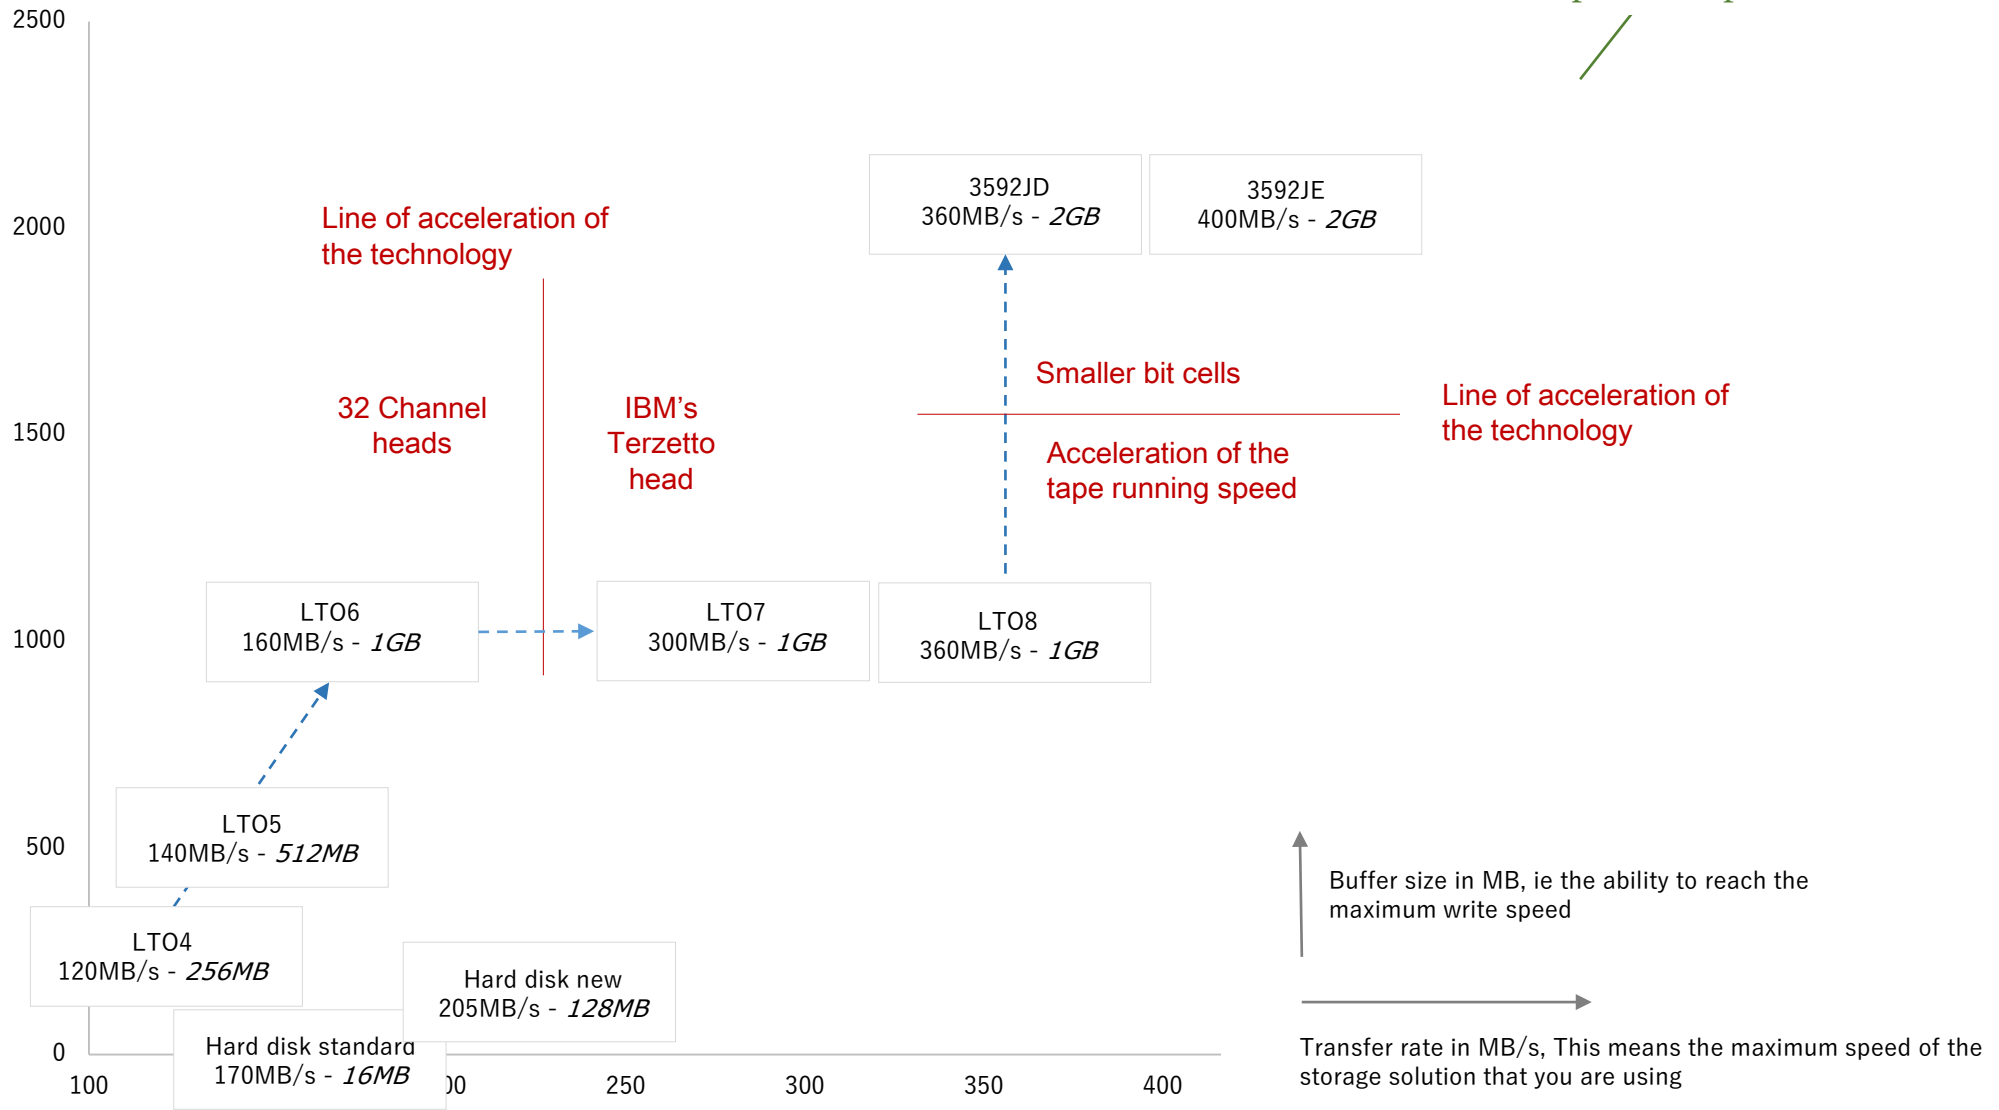
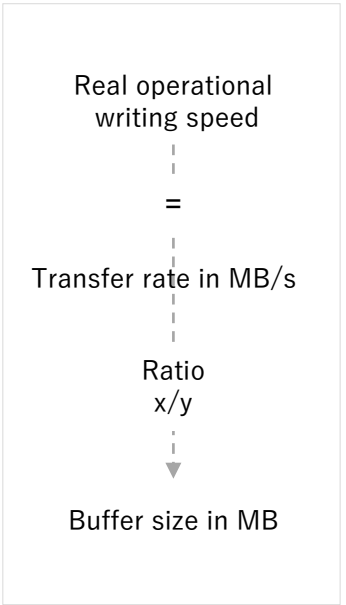


The test consisted of estimating the time needed to write 25GB of small 13MB files (256 KB blocks).




The Backup Time Reduction Mechanism - the Optimal Speed is the Result of an Equation that Combines the Transfer Rate and the Size of the Buffer.

Optimum speed





# Evolution of the Transfer Rate According to the Number of Channel Heads...



# Backup Time


How long does it take to save 20TB of data?  
(1GB files)






# Saving a Large Amount of Data


Backup time in hours on LTO8 and 3592JE/TS1160 drives



Backup time in days:  
 \*8 hours of use of the LTO8 drive per day  
 \*14h of use of the 3592JE / TS1160 per day



Number of drives needed to save data within 5 days



## Data Migration – 3 Key Ideas




### Less Data Migration

The 3592 technology reduces data migration cycles: the capacity growth of 3592 is greater than that of LTO technology.



### Easier Data Migrations

IBM provides more backward compatibility than LTO technology.




### Faster Data Migrations

The higher write speed, coupled with the fact that a 3592 drive can be used for a longer period of time over the day.



## Data Migration - Higher Capacity Reduces the Number of Migration Cycles

Evolution of the highest capacity available in LTO technology and 3592 - all figures in TB.

- In the early 2010s, LTO used to offer 1.5TB of capacity and 3592, 1TB.
- It took 4 years for 3592 to reach 10TB of capacity.
- It took 9 years for LTO to exceed 10TB of capacity.



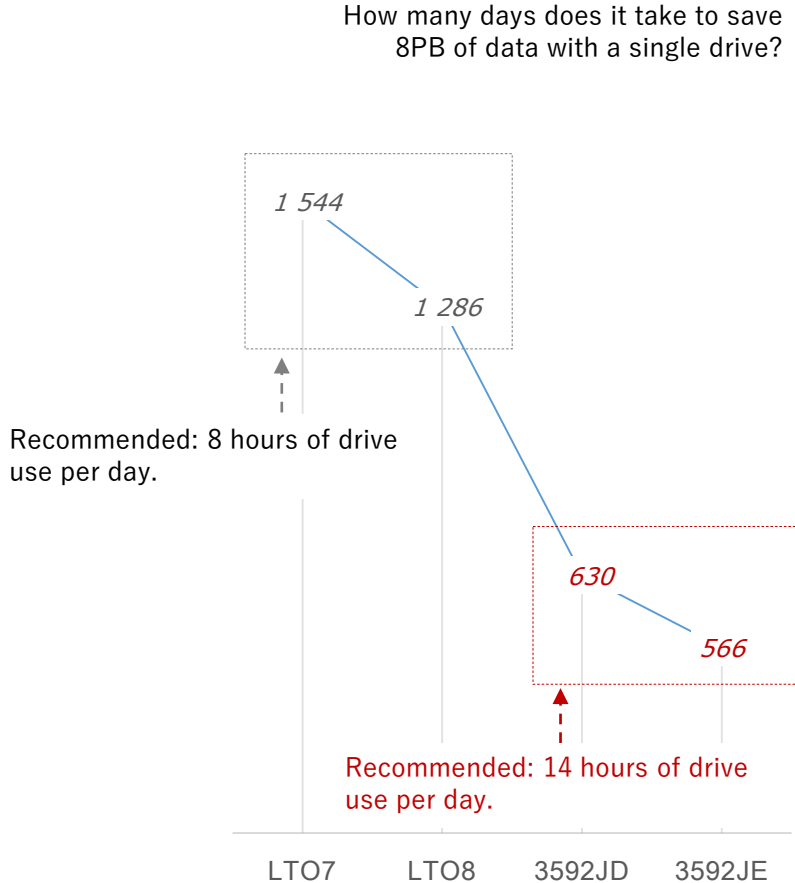


# Data Migration - Enhanced Backward Compatibility



Projects: all specifications and launch dates remain to be confirmed.

# Migration Time...


How long does it take to migrate 8PB of data



# Migration Time (II)...

How long does it take to migrate 8PB of data

How many drives does it take to save 8PB of data in under 60 days?



## Summary

30%

A 3592JE drive will write data at a speed that will be 30% faster than an LTO8 drive.

X 2.3

The duration of a data migration on LTO8 is 2.3 X longer than on 3592JE.

14h

A 3592 drive can be used more intensively and for longer than an LTO drive.

X 2

The duration of a data migration on LTO8 is 2X longer than on 3592JD.

# On the Importance of Using Drives of New Generations

## Data migration - data access

- The data migration time does not only depend on the write speed of the new generation drives.
- Firstly, you need to read the data saved on the old generation tapes.



### LTO6

- It takes almost 7000h to read 4PB of data on LTO6 with a single drive.
- This would almost make 870 days, with an 8 hours per day drive use.



### 3592JE

- It takes almost 4000h to write 4PB of data on an 3592JE tape with a single TS 1160 drive.
- This would almost make 284 days, with a 14 hours per day drive use.


## CONCLUSION:

- Even though the new TS1160 can write 4PB of data in 284 days, the actual migration time will be 870 days, which is the time it will take to "read" the data stored on the LTO6 tapes.
- ✓ Continuing to use slow drives will complicate and slow down the data migration process.

# Data Migration: The time it takes to read your data on older generations of tapes.


What storage capacity in TB can we read in 90 days?

- \* 8h/day of use of the LTO drives
- \* 14h/day of use of the 3592 drives



How many days does it take to read 4PB of data?


- \* 8h/day of use of the LTO drives
- \* 14h/day of use of the 3592 drives



Using new 3592JE tapes will allow you to anticipate your future data migrations: they will take 2X less time than if you use LTO8 drives and tapes.





## Increasing the Writing Speed Presents Multiple Challenges



You have to imagine that:


- The head must run at a speed of 5 m/s on a track of about 1072 m in length and about 2.5  $\mu\text{m}$  in width (TS 1150).
- This is equivalent to running a 100m race in 20 seconds, while staying aligned on a track so fine that we can not see it with the naked eye.
- The condition being that the tracking devices should not derail.
- The tightrope misalignment (for a 1072 m length + a 2.5  $\mu\text{m}$  width) should not exceed 0.06  $\mu\text{m}$ .



Please bear in mind that:

- We cannot address the issue of improving the transfer rates without considering that we also have to consistently produce higher capacity tapes.
- An increased storage capacity implies the use of tracks that are getting thinner for every new generation.
- So, as we have to increase the writing speed of the system, we need to take into account that the head will operate more quickly on a narrower track.
- Avoiding “getting off the road” is a major challenge: stable recording is one of the fundamental criteria in many aspects.

High performance tapes will need to stabilize the tape pass better than before





3

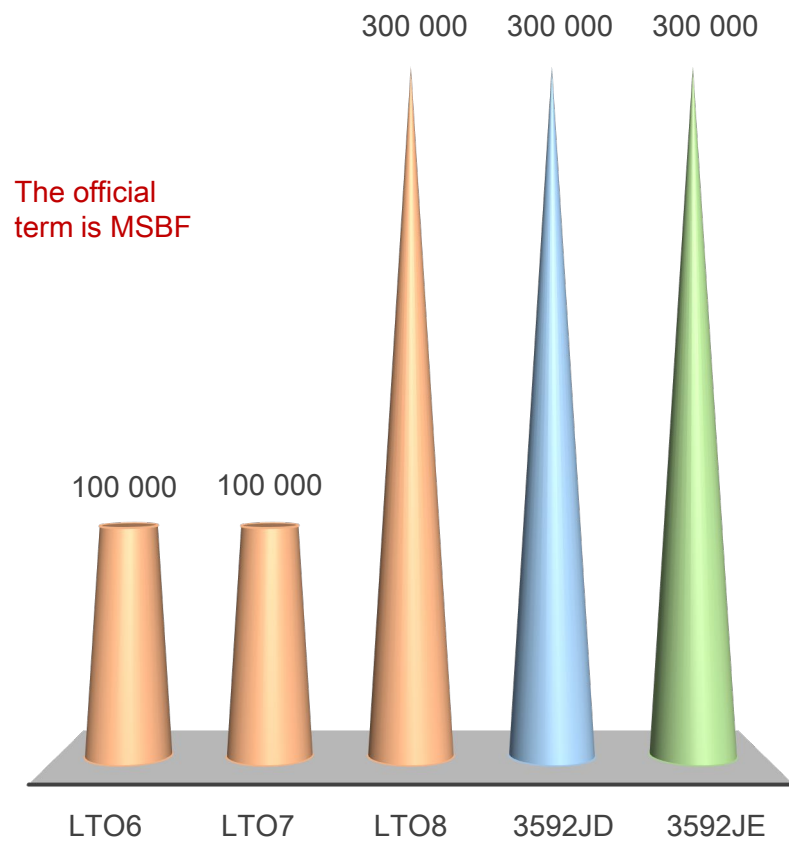
Intensive Use and  
Frequent Access to Data






# Access to Data: A World of Difference between 3592 and LTO

Maximum Number of Loads and Unloads of a Tape Cartridge:




How Many Times can we Access a File Saved on a Tape?




# More Robust Tape Cartridges

LTO tapes



Thickness of the shell: 1.65 mm

IBM 3592 tapes



Thickness of the shell: 3.15 mm  
+ strengthened shell (ribs)

The 3592 tape cartridge is built for an intensive use of the drive...


#### Robust Cartridge Design

- Thicker plastics (vs. LTO)
- Ribs to hold a reel
- Five screws to tighten shells
- Spec'd to withstand 1m drops from all 6 axes without data loss (not recommended!)

#### Dust-Proof Curved Door Design


- Effective dust-proof design for higher recording density
- Passed an open/close test more than 50K cycles

Ribs



Access to data: 3592JE compared to LTO8

50%




The sum of the 4 advantages of 3592JE on LTO8 makes the data access time on 3592JE 50% shorter.




# The Theoretical\* Unwind/ Rewind Time of 3592JD Compared to that of LTO8

Motion speed of the drive head


Waiting time in seconds when loading the tape




Head motion speed (in m/s)




Length of the usable portion of the tape (in m)



Unwinding time from the beginning to end of the tape (in seconds)



Load time + maximum rewind time (in s)

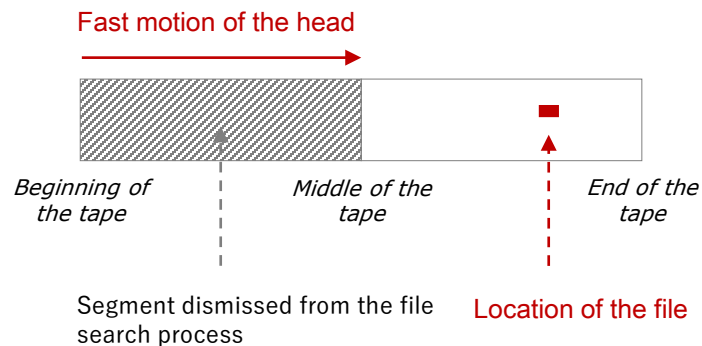


\*Doesn't take into account the advanced features developed by IBM.

# The HRTD Tape Directory of 3592 Speeds Up the Access to your Files


HRTD = High Resolution Tape Directory

## 1-The file search in the traditional mode




- During the unwinding process, the system can help the head to locate the file.
- It informs the head of whereabouts on the segment of the tape the file is positioned.
- When the head reaches the segment in question, it just remains for it to search for the file by capturing the magnetic signals emitted by said file.
- It is as if, in the search for an address, the system indicates the name of the street and that what is left for the head is to look for the number of the building.
- On the other hand, the system is relatively limited since it only works with three points of reference: the beginning, the middle and the end of the tape. As a result, the system can only divide the tape into two segments. Therefore, this is not a particularly effective help.

## 2-With the 3592 Tape Directory system



- The Tape Directory system on the IBM 3592 performs the same operation.
  - The advantage is that it is able to divide the tape into 128 segments.
  - In fact, it will determine the 127 segments that the head will eliminate from its search.
  - The area that the head needs to explore is reduced, which considerably speeds up the search for files. Instead of a long avenue, the head only has to explore a small alley.
  - More importantly, the head will be able to move at maximum speed on the "grey" areas that are removed from the search, while it must slow down on the white area where the file is located, in order to be able to capture the signals emitted by the particles.
- ✓ In some cases, we find that the search for files, and therefore, the overall unwinding time is reduced by 35%.

## IBM's RAO System Generates a Quick File Search Order when Accessing Multiple Files on a Single Tape Cartridge




### Recommended Access Ordering – an example:

- The recall of multiple files on the same tape cartridge is the trickiest operation on tape.
- The drive must unwind/rewind, find the first file, open it, then go back to the beginning of the tape and unwind again in order to locate the second file etc ...
- With the 3592 RAO system, the drive calculates and produces the list of file recall orders in its optimal version. In other words, the RAO calculates the fastest scenario or file recall path.
- The RAO performs its calculation by taking several criteria into account, such as the size of the file, its positioning (on which wrap, which databand etc. ..), as well as the attributes of the system (transfer rate).

# IBM's RAO System - Some Examples of the Reduction of the Access Time in the Scientific Environment

- 20% faster access to the data when it comes to opening a small number of files: 2 or 3 files.
- Access to data 30% faster when it comes to opening fifteen files of 10GB on average.
- Finally, for an almost constant access to the data, a major European research centre explained that the time of use of the drive is halved, thanks to the RAO system. Please see the graph opposite:





The 3592JE is a System whose Function is to Facilitate Frequent Access to Data



HDD

- Lifespan not exceeding 3-4 years
- More expensive, slower, and less secure
- Convenient for data access

3592 tapes

LTO tapes

- Archiving data over the long term
- A safer solution than hard disk
- Superior write speed

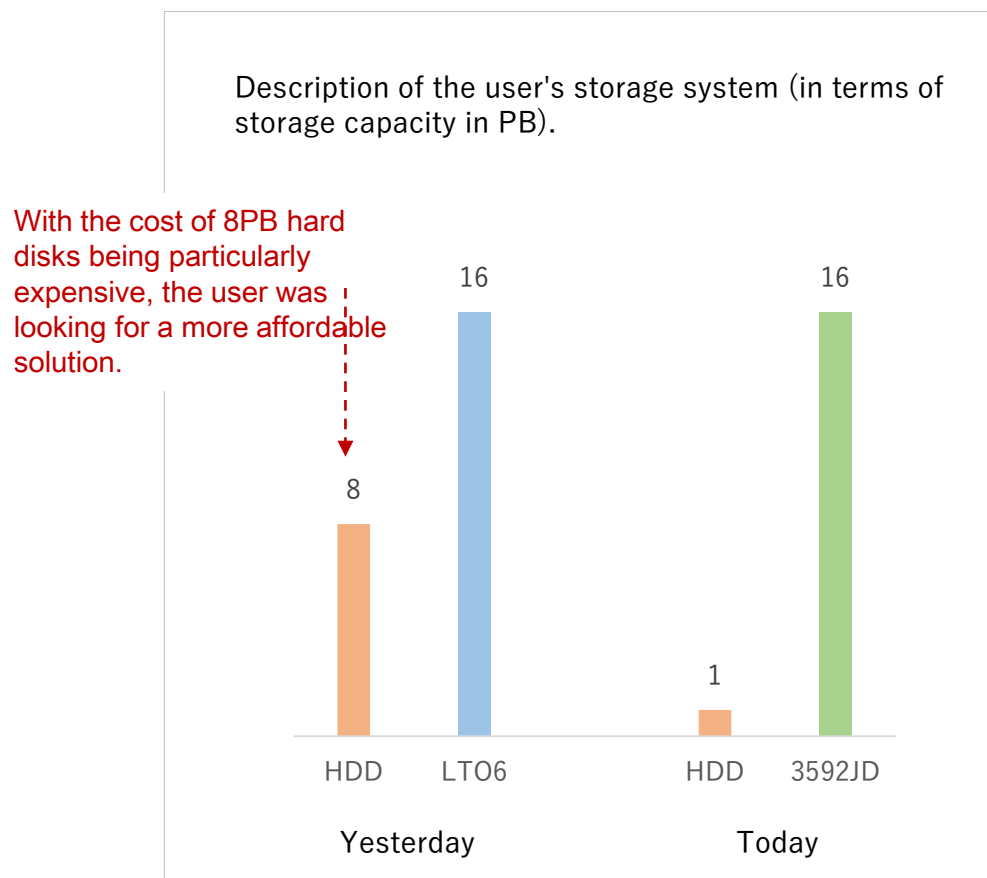
3592 combines the qualities of both systems.

-----> 3592JE is a 2-in-1 solution, the nature of which is to enable users to reduce their storage costs by reducing the hard disk capacity required for frequent access to data.

## The 3592JE is a System whose Function is to Facilitate the Frequent Access to Data


A European TV channel has reduced the overall cost of its storage solution by purchasing a 3592JD library.

Example of a TCO (total cost of ownership) reduction thanks to the acquisition of a 3592JD solution



1. The very nature of 3592 is that it performs the functions of LTO tapes and that of hard disk at the same time.
2. The cost of use of the HDD is between 5X and 7X more expensive than that of tape.
3. The price of a 3592JD solution is, on average, 25% more expensive than that of an LTO tape based solution.
4. The purchase of the 3592 storage solution has allowed the user to reduce its hard disk capacity from 8PB to 1PB.


✓ This allowed the user to reduce the overall cost of his storage system by 50%.



## An Example Description of a Storage System combining 3592JD and HDD

A company involved in the remote sensing area and in satellite imaging has moved from a 100% HDD solution to a HDD + 3592JD solution.

- The proliferation of regulations on the long-term retention of data has prompted the user to consider an option other than hard disk.
- This company stores 2.5PB of data and will create between 480TB and 500TB new data per year over the next 5 years.
- They explained to us that as they purchased 3592JD hardware, they do not require more than 0.5PB hard disk capacity (or "disk cache"):
  - \*250TB in order to keep 6 months of new data on HDD
  - \*and another 250TB of extra capacity for exceptional interventions (migration, vertical use etc ...).





# An Example Description of a Storage System Combining 3592JD and HDD (2)


An anecdote that is fairly representative of the advantage of 3592JD over hard disk: the user explained that they experienced the following operational speeds when recording their data;

- \* HDD - Files > 20GB: 190MB/s - 1GB files: 60MB/s
- \* 3592JD - Files > 20GB: 342MB/s - 1GB files: 252MB/s


✓ They also informed us that they keep a large number of small files - the operational speed for small files is, therefore, the one they mostly take into account.

Double comparison criterion:  
 Speed = 1 x TS1155 = 5 x HDD  
 Capacity = 2 x 3592JD = 15 x HDD

RAID – 3 copies in total



Backup time for 10TB of data on a single drive or hard disk (in hours).



\* The user wants to be able to save the equivalent of one week of new data within a day.

\* They're on TS1155 (15TB) and use 6TB hard disks.

\* Due to the robustness of the 3592JD, they estimate that they need 5 HDDs to generate the work of a single 3592 drive.

\* Opposite, you can see the result of a 3 week cycle of data Backup.

HDD = disastrous ecological footprint



## An Example Description of a Storage System Combining 3592JD and HDD (3)

The LTO8 option was rejected on three basic criteria:

1. Floor space reduction + fewer data migration cycles.
2. Superior performance (writing speed).
3. Frequent access to data + system robustness.


We can turn this company's TCO into an index, as follows:

- LTO8: 100
  - 3592JD: 135
  - HDD: 700
- } = at equal capacity, their cost on the 3592JD solution is 35% more expensive than that of the LTO8 solution.

Some complementary information:

- If they had opted for LTO8, they would have only reduced their disk capacity by half (1.25PB).
- Had they not needed frequent access to data, the cost of the HDD + LTO8 solution would have been 12% lower than that of the HDD + 3592JD. However, this user considered that 3592JD was still more profitable (floor space reduction, less write errors, lifespan of the hardware, robustness etc ...).

The 100% HDD solution is 150% more expensive than 3592JD



NOTES

**FUJIFILM**