



CERN
Tape Archive

Large Scale Tape Storage: CTA

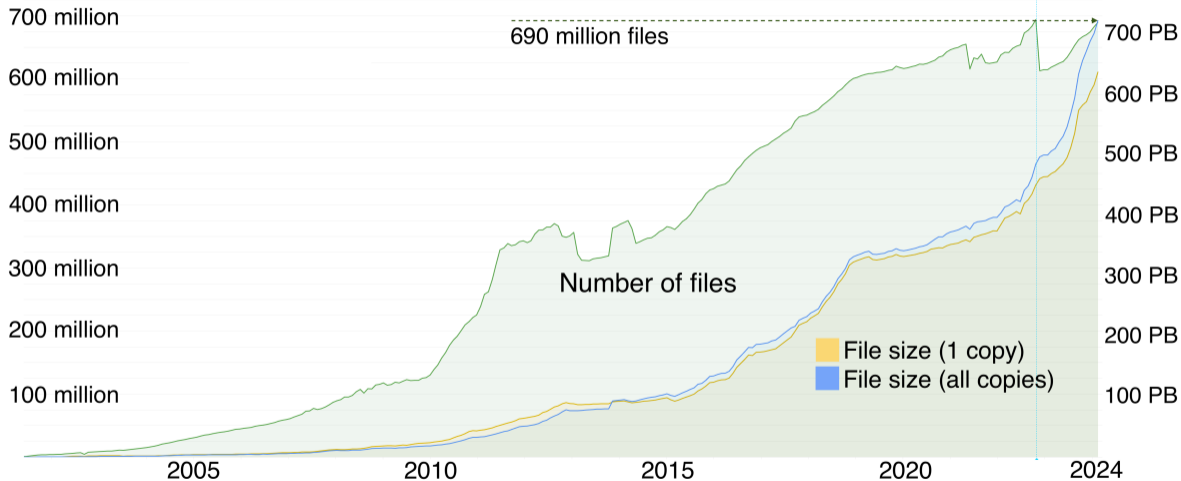
Joao Afonso

Tape Archive and Backup Section
Storage and Data Management Group
IT Department, CERN

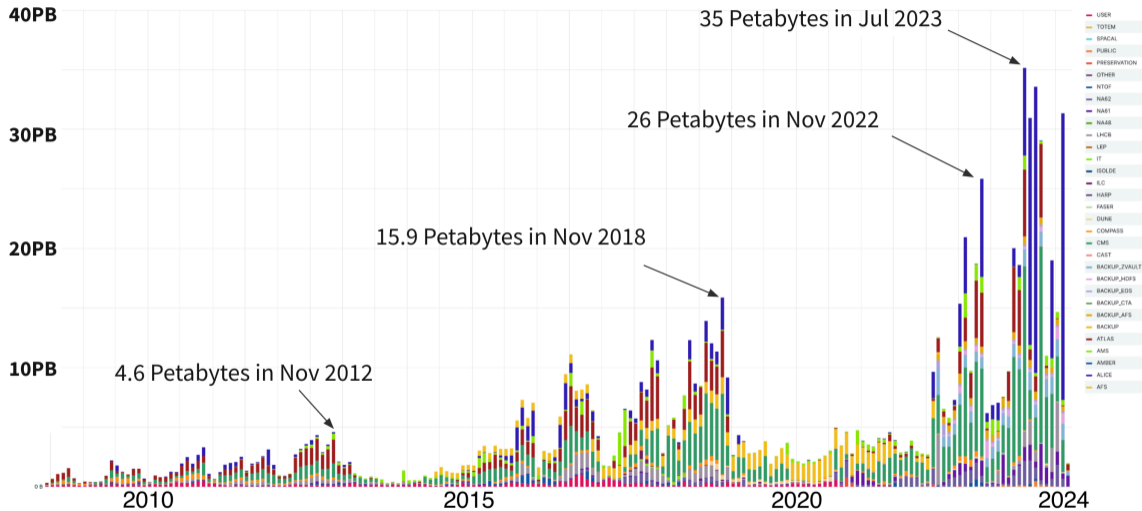
CERN/LHC Data Archival

- All LHC physics data—720 PB—is stored on tape in the CERN Data Centre, Geneva (Tier-0)
- Additional copies of the data are distributed across 13 national computing centres (Tier-1)
- Tier-1s also provide tape archival storage


Physics Data on Tape



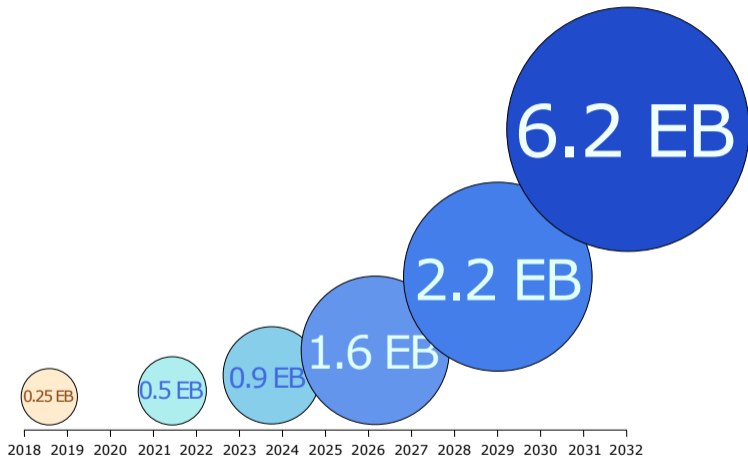
Data rates written to Tape



CERN Tier-0 Data Rates (2022–2026)



Predicted Tape Archival Storage Needs




Why use tape? Isn't it obsolete?



CERN Data Storage in the 1970s



140 MB 9-track tape (1974)



10 MB disk platter from CDC
7638 Disk Storage Subsystem
(1974)

Tape and disk have evolved in step

IBM TS1170 Tape Drive
50 TB capacity

WD DC HC670 Hard Drive
26 TB capacity




Comparing tape and disk

	Tape	Disk
Data transfer rate	400 Mb/s	200 Mb/s
Positioning type	Fast Sequential Access	Fast Random Access
Average positioning time	30 seconds (610 m @ 12 m/s)	5 milliseconds
Latency to first byte	A few minutes	5–10 milliseconds

Advantages of Tape : Reliability and Data Security

- Two heads are better than one :
read after write verification
- No data loss if a drive fails
- Immutable files
- Air-gap security
- Long media lifetime (30+ years)



Advantages of Tape : Energy Efficiency

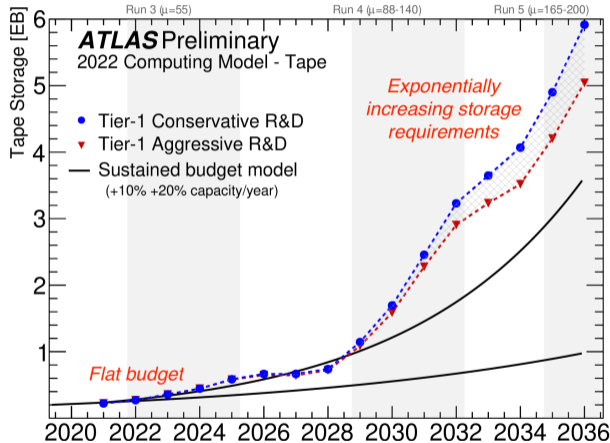
Hard disks are always on. They constantly consume power and generate heat.

- Expensive to run
- CERN Tier-0 Data Centre is at the limit of how much power and cooling it can deliver (3.5 MW)
- Disk storage requires power and cooling

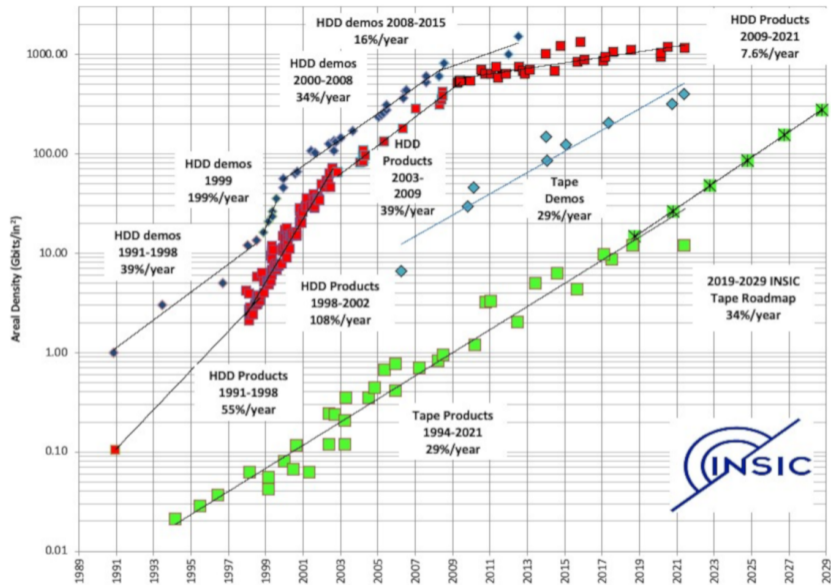
Tape cartridges don't consume any power when they are not mounted in a drive.

- Tape capacity can be increased without requiring additional power


Advantages of Tape : Cost!



- Storage needs are increasing exponentially
- Budget is not increasing to match needs
- Tape storage is 3–5× cheaper than disk storage

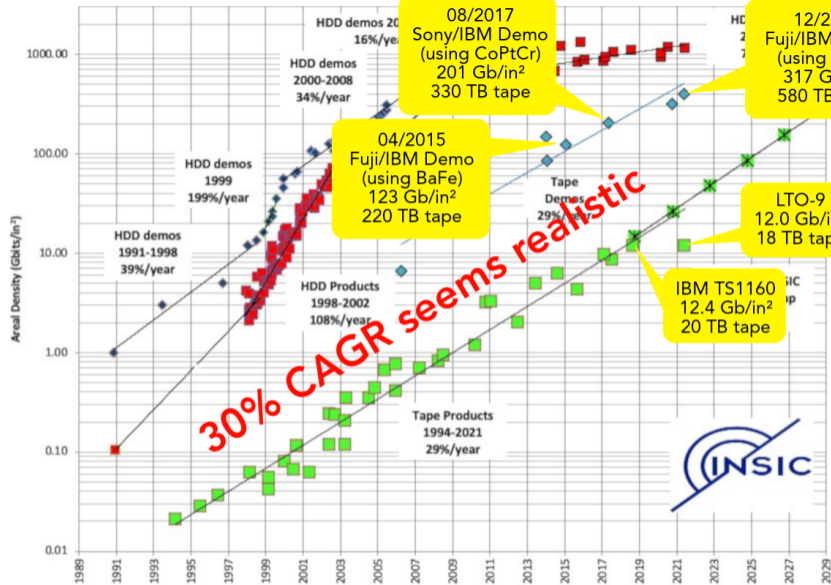


Magnetic Media : Evolution of Areal Density




Magnetic Media : Evolution of Areal Density





Magnetic Media : Evolution of Areal Density





Next Monday on CTA workshop...

The Future of Tape
by Dr. Mark Lantz (IBM)

9h15, IT Amphitheatre, CERN

Magnetic Media : Evolution of Areal Density



The Archival Storage Solution from CERN IT Storage Group



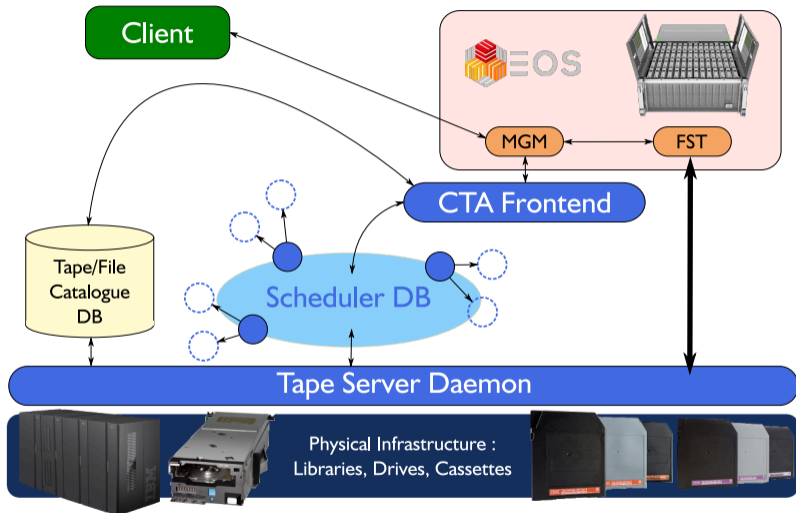
EOS+CTA : “Best of Both Worlds”

- CTA is the tape back-end to EOS (CERN’s high-performance disk system)
- Interface, file operations and disk pool management provided by EOS
- Scheduling and tape operations provided by CTA


EOS+CTA : "Best of Both Worlds"

Function	Provided by
File Metadata Operations	EOS (MGM/XRootD)
Namespace	EOS (QuarkDB)
Disk Buffer for tape files	EOS (FST)
Tape File Metadata Ops	CTA (Frontend)
Archive/Recall Requests	CTA (Scheduler DB)
Tape File Catalogue	CTA (Catalogue DB)
Tape Operations (libraries, drives, cassettes)	CTA (Tape Server)

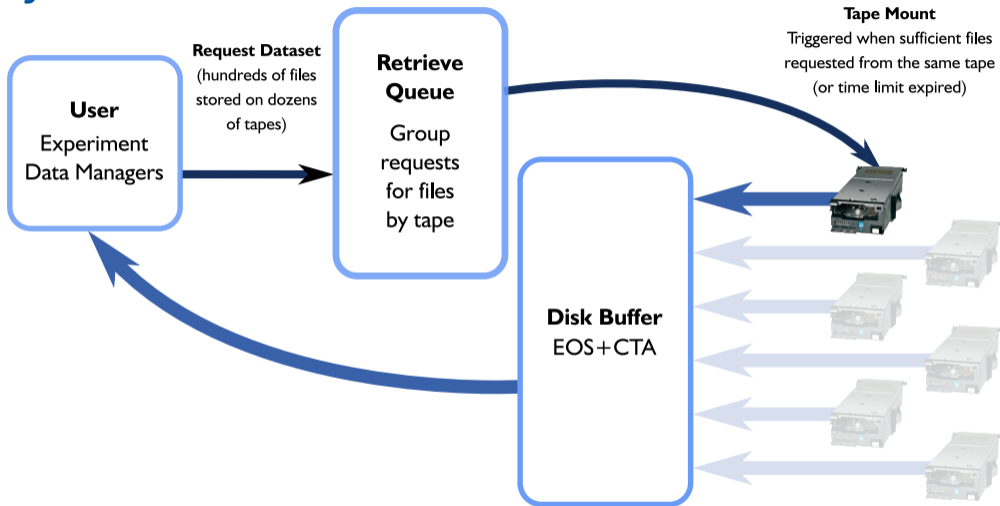
EOS+CTA Architecture



Physics Data Flow : Archival



Physics Data Flow : Retrieval




EOS for analysis vs. data archival/retrieval

“Big EOS” (analysis)

- Tens of PB of storage for physics jobs and retrieval to Tier-1s.
- File replicas have a long lifetime.
- Spinning disks.

“Little EOS” (archival/retrieval)

- Small buffer for copying files to/from tape.
- File replicas have a very short lifetime. Deleted as soon as tape copy exists (archival) or copied to “Big EOS” (retrieval).
- SSDs: reduce contention and give the best price/performance ratio.

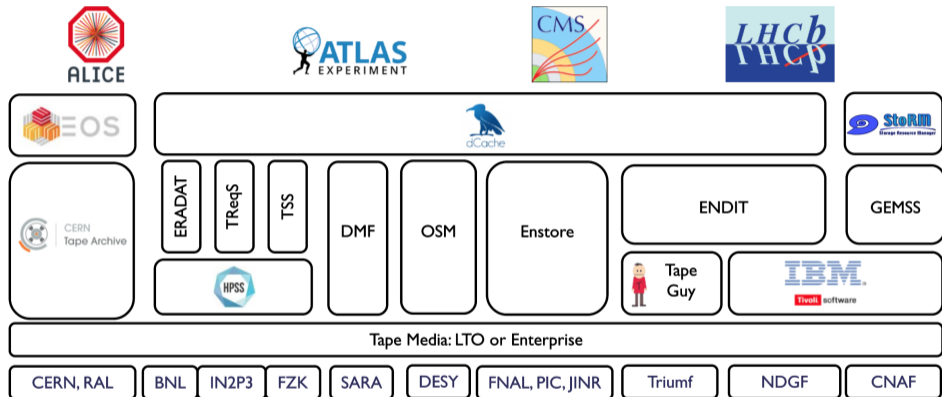


HDD icon: <https://commons.wikimedia.org/wiki/File:Hard-drive.svg>
SSD icon: <https://commons.wikimedia.org/wiki/File:Ssd.svg>
Tape icon: https://commons.wikimedia.org/wiki/File:Tape_cinta_cassette_backup.svg

Complementary roles of Disk and Tape

	Disk	Tape
Use Case	Online : Data Analysis	Offline : Long-Term Archival
Drives	220 000 disks (85% HDD/15% SSD)	200 tape drives (in 5 libraries)
Media	—	55 000 tape cartridges
Capacity	Nominal capacity ≈ 870 PB Effective capacity is lower due to redundancy	Capacity ≈ 920 PB but can be easily extended Currently 720 PB on tape (690 million files)

In the last years, the situation has evolved



In the last years, the situation has evolved

- The tape software landscape is consolidating
 - Changing license model/costs for commercial solutions
 - Some free software solutions at end-of-life (CASTOR, Enstore, ...)
- The CERN Tape Archive is an attractive option
 - Free and open source
 - Modern software integrated with the latest WLCG standards and protocols
 - Supported by CERN; long-term roadmap
 - Includes operational management tools

Beyond CERN : CTA at other sites

Triumf



Rutherford
Appleton Lab
(RAL)



DESY



Institute of High
Energy Physics
(IHEP)



Fermi National
Accelerator Lab
(FNAL)



Port d'Informació
Científica
(PIC)



Australia's Academic
and Research Network
(AARNet)



Beyond CERN : CTA at other sites

- EOS+CTA is Free and Open Source Software (GPLv3)
<https://cta.web.cern.ch>
- Active and growing CTA Community
<https://cta-community.web.cern.ch/>

Summary

- Tape is the best currently-available technology for archival storage, in terms of reliability, stability over long periods of time and cost
- Power and cooling constraints make it difficult to add more disk storage capacity; tape storage capacity can easily be increased
- Cost benefits over disk storage are significant and look set to increase over the next decade
- Storage needs are growing but budgets are flat
 - The CERN physics archive is 720 PB and will soon grow to 1 EB
 - Data retrievals already exceed 1 EB/year
 - The storage demands of HL-LHC will mean more data on tape and new tape workflows
- CERN is investing in tape as its primary archival storage medium for LHC Run-3 and Run-4

Tech Week Storage 2024 - CTA events

- Meet the team: CERN Tape Archive (CTA)
 - 14 March, 14:00 - 17:00, room 600/R-001
- CTA Workshop 2024
 - 18-19 March, CERN

Exit through the gift shop



