



The CERN Tape Archive (CTA) : Archival Storage for Scientific Computing


Michael Davis

CERN, IT Department, Storage Group




CTA is the tape archival storage
back-end to EOS

CERN Tier-0 Data Rates (2022–2025)




Data archived to tape storage each month since 2008

26 Petabytes in November 2022




Physics Data on Tape

Total Data on Tape in CASTOR/CTA



Predicted Tape Archival Storage Needs



Comparing tape and disk technologies


IBM TS1160 Tape Drive
20 TB capacity

WD DC HC 530 Hard Drive
14 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



Magnetic Media : Evolution of Areal Density






Magnetic Media : Evolution of Areal Density



Magnetic Media : Evolution of Areal Density

The Outlook for Tape Technology




New Advanced Materials

- Very fine magnetic particles
- Smooth surfaces with low friction
- 3D stacking of magnetic particles

Spinning disk technologies are pushing the limits of storage density. Tapes have plenty of room to improve capacity.

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

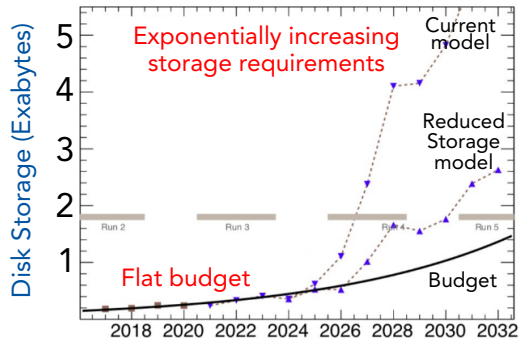
Hard disks are always on. They constantly consume power and generate heat.

- 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 compared to budget (ATLAS)

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

High-Luminosity LHC (2029–32)




AUTOMATIC REVERSE CASSETTE CHANGER

Hi-Lumi LHC : Data Carousel

Store data for online analysis on tape

Data storage challenge of HL-LHC :

- 'Opportunistic storage' basically doesn't exist
- Format size reduction and data compression are both long-term goals, require significant efforts from the software and distributed computing teams
- Tape storage is 3~5 times cheaper than disk storage, increasing tape usage is a natural way to cut into the gap of storage shortage for HL-LHC



'Data Carousel' R&D → to study the feasibility to use tape as the input to various I/O intensive workflows.

The Archival Storage Solution from CERN IT Storage Group



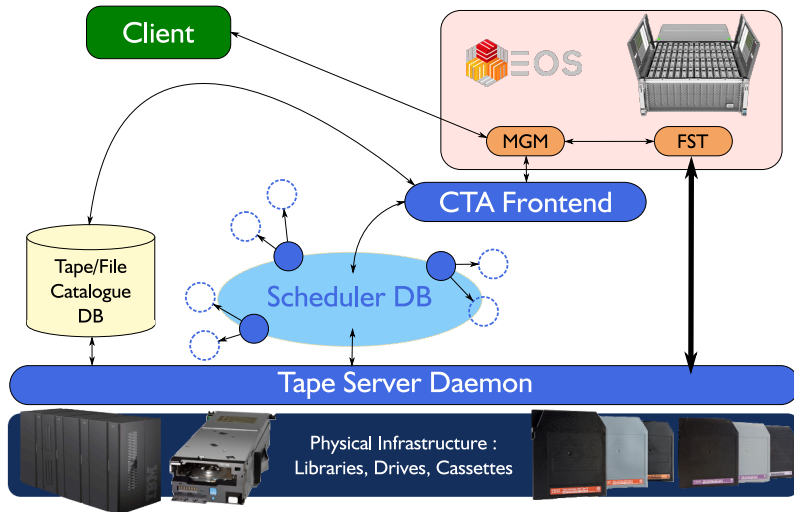
EOS+CTA : “Best of Both Worlds”

- CTA is the tape back-end to EOS
- 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 Staging	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




EOS for Analysis vs. EOS for Staging

"Big EOS" (Analysis)

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

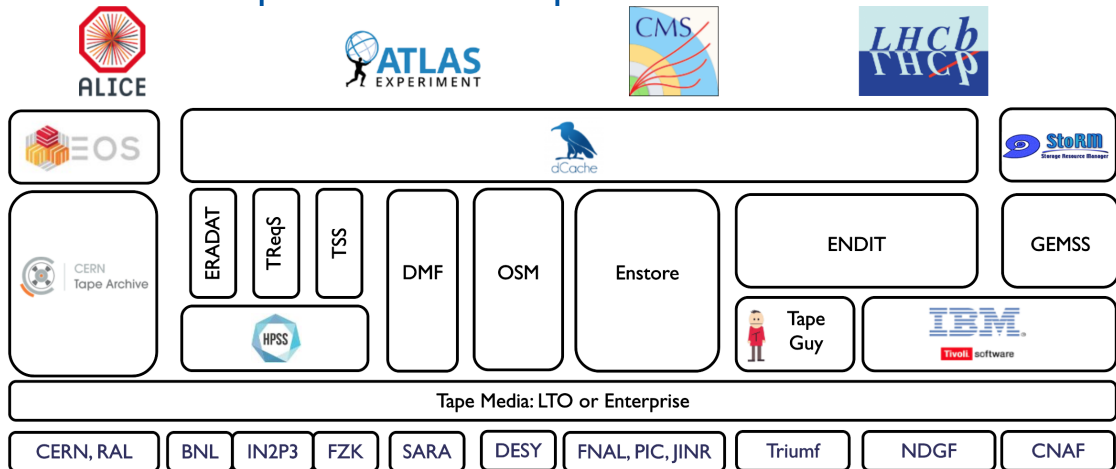
"Little EOS" (Staging)

- 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

WLCG Tape Landscape in 2021



Beyond CERN : CTA at other sites

- Wide diversity of tape software deployed at Tier-1 sites, consolidation likely in coming years
- Increasing license fees for commercial alternatives/ risk of lock-in
- CTA is Free and Open Source Software (GPLv3)
<https://cta.web.cern.ch>
- Active and growing CTA Community
<https://cta-community.web.cern.ch/>

CTA Day at the EOS Workshop :

Wednesday 26 April

- Site reports
- Technical presentations
 - Operations
 - Software development
- Hands-on session
- CTA Roadmap

