



CERN
Tape Archive

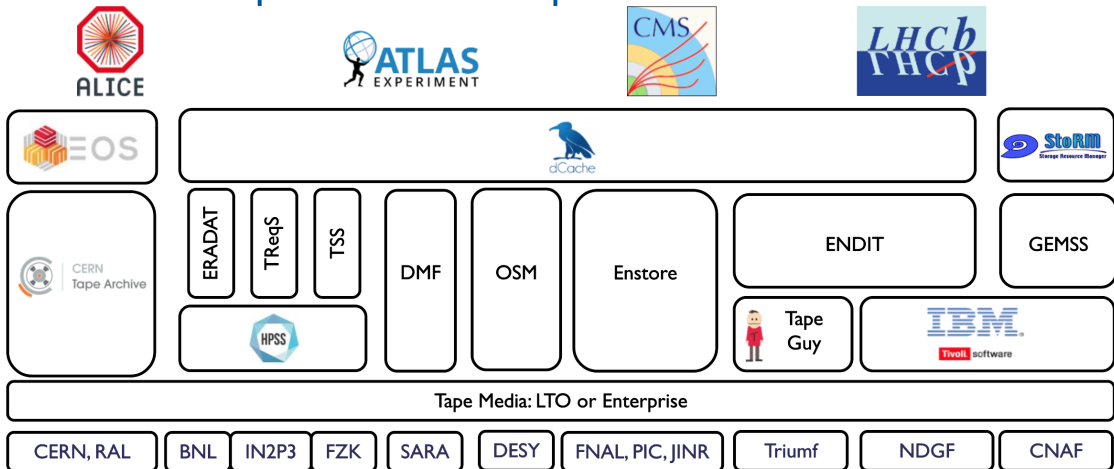
The CERN Tape Archive Beyond CERN

An Open Source Data Archival System for HEP

Dr. Michael Davis

On behalf of the CTA team at CERN
and contributors from DESY, Fermilab and RAL

WLCG Tape Landscape in 2021



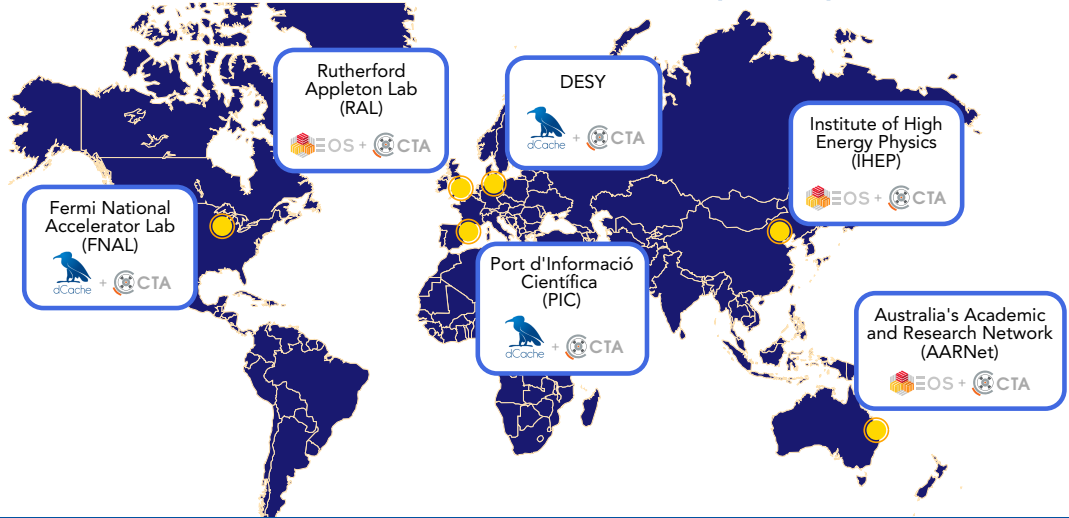
In 2023, the situation has evolved

- The tape software landscape is consolidating
 - Changing license model/costs for commercial solutions
 - Some free software solutions at end-of-life (CASTOR, Enstore, ...)

In 2023, 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

CTA Sites (7th EOS Workshop, Apr. 2023)



CTA Design Motivations

- Successor to CASTOR
 - 350 PB to migrate from CASTOR to CTA
 - Migration should be a pure metadata operation
- Tape back-end to EOS
 - CTA is a pure tape system
 - CASTOR disk to be replaced by EOS
 - Architected to be agnostic to disk system... almost

Issues for sites wanting to adopt CTA

- Migration path to CTA
 - Avoid having to rewrite existing archive
 - Migrate tape catalogue to CTA Catalogue
 - Provide disk namespace migration tools
- Allow disk cache/namespace other than EOS
 - Specifically: dCache
- Ease of distribution and deployment
 - Remove CERN-specific software dependencies
 - Alternative to Oracle for CTA Catalogue DB
 - Binary RPM distribution

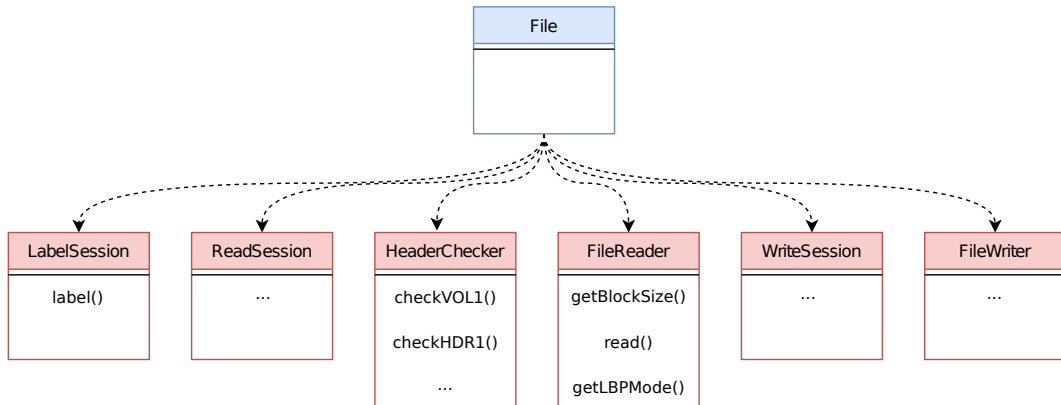
1. Migrating existing tape archive to CTA

- CERN use case: migrating from CASTOR to CTA
 - 350 PB to migrate, cannot afford extended downtime
- Design decisions
 - Key principle: migrate metadata not data
 - CTA tape format \Leftrightarrow CASTOR tape format
 - Migrate tape metadata \rightarrow CTA Catalogue
 - Migrate disk metadata \rightarrow EOS namespace
 - [Migration from CASTOR to the CERN Tape Archive](#)
(HEPiX Autumn Workshop 2021)
- Migration for other sites
 - RAL: used CERN migration tools to migrate from CASTOR to CTA
 - What about migration from other tape systems? Enstore, OSM, ...

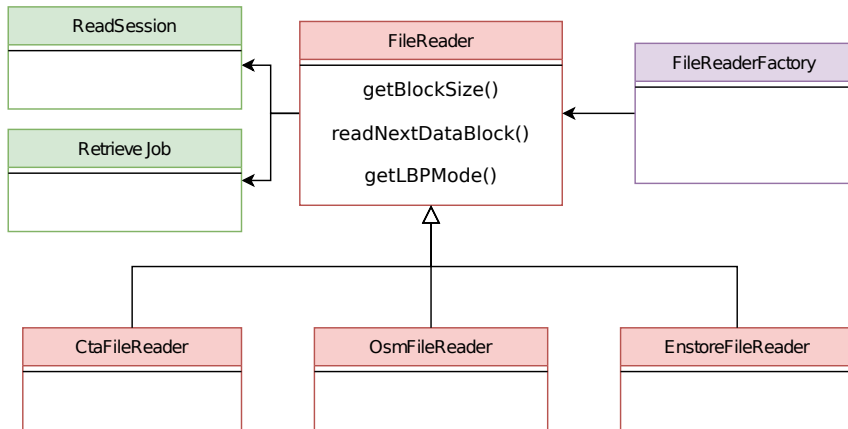
1. Migrating existing tape archive to CTA

- Solution: allow CTA to read (but not write) non-native formats
- Tape label format
 - Enstore label similar to CASTOR/CTA
 - OSM label format is different
- Tape file format
 - Enstore/OSM formats based on CPIO
 - Minor difference in representation of file size (octal/hex)

Read Non-native File Formats



Read Non-native File Formats



Read Non-native File Formats

- CTA can now read Enstore, OSM tapes
 - External tape readers:
Integration into CTA and OSM/Enstore cases
- Other file formats can be added as required
- Tools to convert metadata are left as an exercise for the reader
 - Disk File Metadata for Tape Files—
Migrating, Restoring, Replicating

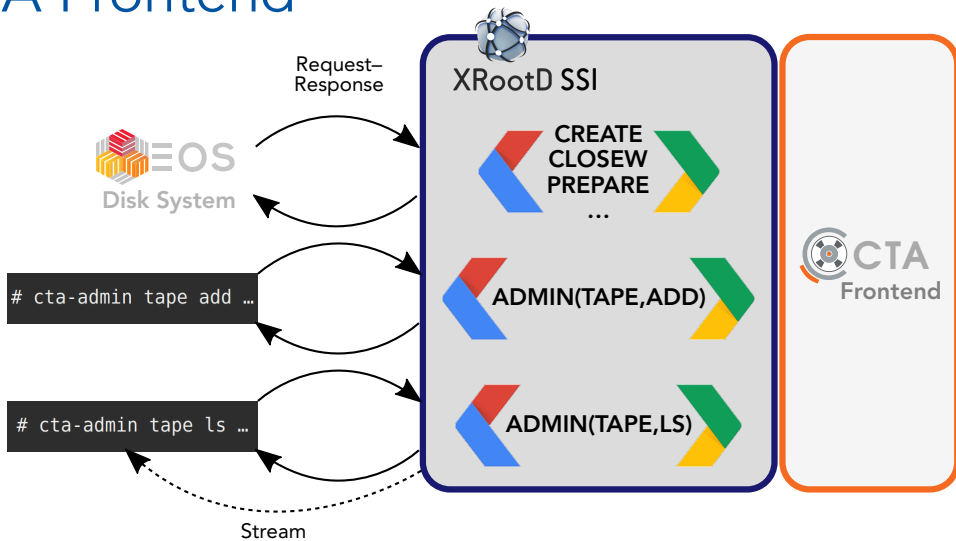
2. Allow non-EOS disk cache/namespace

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)

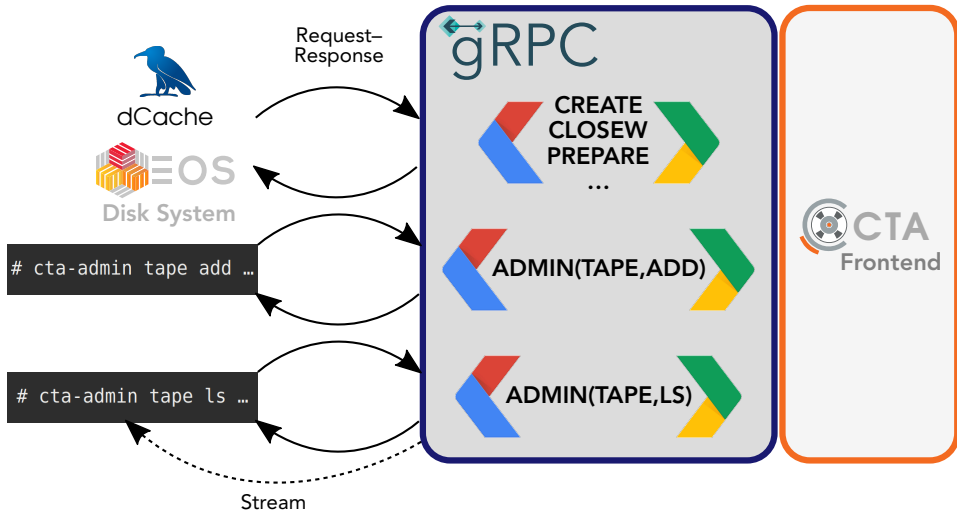
Disk System messages to CTA Frontend

- Client request messages to CTA Frontend are serialised in Google Protocol Buffers
- Transport protocol is XRootD Scalable Service Interface (SSI). This works well, but:
 - SSI not widely used; additional (non-standard) dependency
 - Specifically: SSI extensions not supported by dCache's XRootD Java client library
 - gRPC is the native transport for protobuf
- gRPC Frontend implementation/proof-of-concept contributed by dCache team

CTA Frontend



CTA Frontend



CTA Frontend Roadmap 2023

- **DONE:** Disk workflow events over gRPC
 - See poster contribution:
dCache integration with CERN Tape Archive (CTA)
- Refactor XRootD SSI/gRPC Frontend implementations to use common request message processing code

CTA Frontend Roadmap 2023

- Refactor CTA disk file ID code
 - EOS disk file ID is `uint64_t`
 - dCache disk file ID is a string
 - Refactor to use strings for disk file ID
- Finish implementation of gRPC “streaming”
admin commands: `cta-admin ... ls`
- Authentication (gRPC token, Kerberos)

3. CTA Catalogue Database

- CTA supports Oracle and PostgreSQL
 - Oracle is used in production at CERN
 - PostgreSQL is popular in external sites
 - Support for MySQL was dropped
- CTA Catalogue library depends on Oracle RPMs
 - WiP to refactor CTA Catalogue shared library code
- Database schema upgrade procedure and tools
 - Updating the database schema for services in production
 - Improve upgrade procedure between CTA Public releases

4. CTA Public Release

- “Public” binary RPM release of CTA
- No CERN-specific dependencies
- CTA Public release v4.8.7-1/5.8.7-1
 - [Announcement](#)
 - [GitLab repo](#)

5. CTA Operations Tools

- Public release of CTA operations tools
 - **ctautils** : collection of helpers and wrappers used across all the operations tools
 - **tapeadmin** : library for interacting with tape media and libraries
 - Repack automation **ATRESYS**
 - Documentation
- Operations Tools release v1.2
 - [Announcement](#)
 - [GitLab repo](#)

CTA Operations Tools

- Repack
 - ATRESYS—Automated Tape REpacking System, a tool for managing CTA repacks and tape lifecycle
- Monitoring
 - Monitoring your EOSCTA deployment—the general recipe
- Coming soon
 - RPM: alternative to `pip install`;
versionlocks operations tools to compatible CTA versions;
manages dependencies not handled by `pip pkg-mgmt`
 - EOSCTA namespace reconciliation scripts
 - ACL management tools

6. CTA Support



all categories ▸

Latest

Top

Categories

+ New Topic

☰ Topic

Replies

Views

Activity

Request for Assistance with CTA GitLab SSO Error

A

8

50

6d

last visit

Releasing the CTA operator tools

Site Administrators

M

S

7

128

8d

Public release 4/5.8.7-1 announcement

Releases

A

0

30

12d

CTA 5 and HTTP Tape API

General Discussion

T

J

3

49

29 Mar

EOSReporter fails with CTA+dCache

General Discussion

E

M

10

243

23 Mar

Summary : CTA Community Features

- Migrate your existing tape archive to CTA without re-writing the data
- Choice of EOS or dCache for namespace and disk cache
- Choice of Oracle or PostgreSQL for CTA Catalogue
- CTA “Public” binary RPMs
- CTA Operations Tools

Related CHEP Contributions

- Posters

- [The CERN Tape Archive Run-3 production experience](#)
- [dCache integration with CERN Tape Archive \(CTA\)](#)
- [ANTARES: the new tape archive service at RAL Tier-1](#)

- See also

- [CTA Website](#) : Presentations, Publications, Documentation, Source Code
- [CTA Community](#) on Discourse

