CERN
Tape Archive

CTA Status and Roadmap

Dr. Michael Davis for the CTA Team

CERN, IT Department, Storage Group, Tape Archive and Backups Section

The archival storage solution from the CERN IT Storage Group

CASTOR to CTA Migration

- EOS replaces CASTOR disk
- CTA replaces CASTOR tape
- EOS+CTA offers the "Best of Both Worlds"
  - EOS—interface, file operations and namespace
  - CTA—highly performant tape operations based on CASTOR Tape Server

CASTOR to CTA Migration

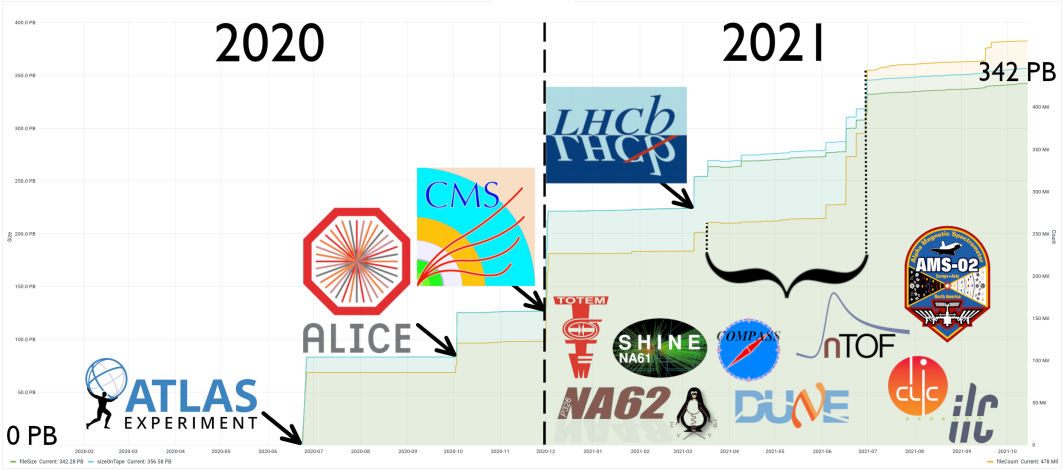

CTA Operational Schedule

- June 2020 : CTA in Production
- July 2021 : All experiments migrated to CTA
  SPS Restart : SMEs start data taking
- Now: LHC Data Challenges
- April 2022 : LHC Beam Commissioning
- Jul 2022 : Start of LHC Run–3, LHC physics data taking
- 2026 : End of Run–3, start of Long Shutdown 3 (LS3)

CTA Development Schedule
1H 2022 : Final preparations for Run–3

- Mount policies support a wider range of authentication
- Disk space reservation system for recalls
- Bug fixes (batch size for RAO recalls, …)
- Better abstraction of disk system/disk instance concepts
- Move Drive State from objectstore to DB
- Improved logging
- Improved handling of failures
- Review of retry logic

CTA Development Schedule
1H 2022 : CTA "Public" Release

- Code clean-up
- Remove unnecessary dependencies (MySQL, …)
- Update system dependencies (XRootD 5, EOS 5)
- Build OS remains as CC7/RHEL7
- Update build system to GCC8 (C++17)
- Build RPMs for distribution outside CERN

CTA Development Schedule

Upcoming Features (after start of Run–3)

- HTTP Stage and Transfer
- Repack
  - Repacking individual tapes works OK
  - Repack operations can interfere with normal archival/retrieval operations: prioritisation, resource management, monitoring
- Scheduler Database

Scheduler Database

- Current Scheduler Database implementation is the Objectstore. Underlying technology: Ceph RADOS.
- Efficient and works well for FIFO queuing operations (archive/retrieve)
- Requires workarounds for non-FIFO operations (delete, priority queues)

Scheduler Database
Additional Considerations

- Some operations are difficult in the Objectstore: schema changes, tracing problems, cleaning up
- Development and operations team members need to learn the Objectstore technology
- Maintenance burden of objectstore code (distributed transaction management system)
- Additional external dependency

Scheduler Database

- Plan : add a new Scheduler Database back-end implemented with a relational DB (PostgreSQL)
- Non-queuing operations (Drive Status) removed from Objectstore to the CTA Catalogue DB—**DONE**
- Queuing operations will be implemented with PostgreSQL backend
- Objectstore implementation will be used for Run–3 operations. No plans to change this.

CTA Operational Tools and Community Features

- Tape verification
- Restoring deleted files
- Reconciliation
- XRootD SSI vs gRPC
- CTA Tape Formats

CTA Status and Roadmap : Summary

- CTA is feature-complete and ready for Run–3
- CTA "public" binary release available later in 2022
- Important development tasks for 2022–2023 :
  - New SchedulerDB back-end
  - Repack priority and scheduling
- CTA Website : Source Code, Documentation, Presentations and Publications
- CTA Community (Discourse)


CERN