



CERN
Tape Archive

3rd CTA Workshop (CTA 2024) Challenges and Roadmap

Dr. Michael Davis for the CTA Team

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

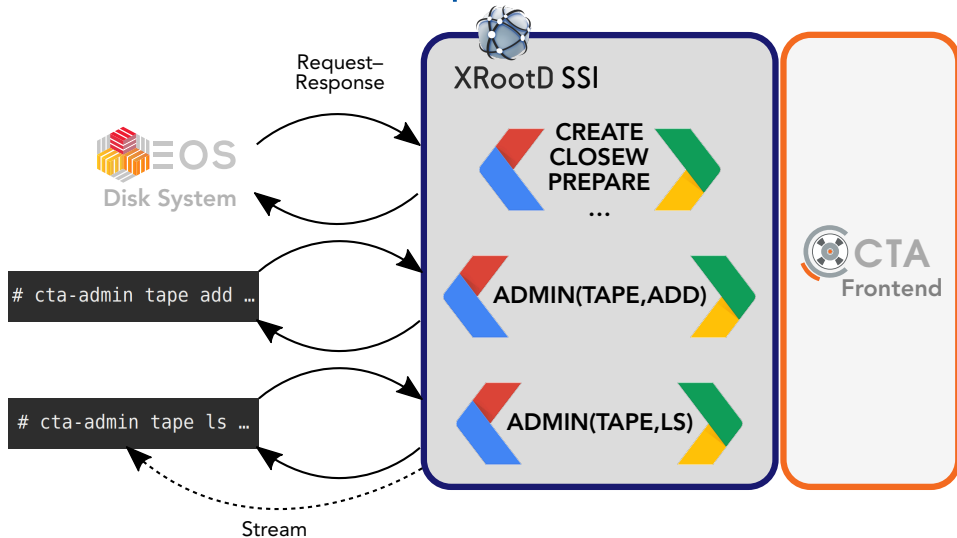
CTA Service : Plans for 2024

- Migrate from XRootD 4 to XRootD 5 **DONE**
- Migrate from CERN CentOS 7 to Alma Linux 9
 - Complete before CC7 end-of-life (30 June 2024)
- Continue to meet the challenges of Run-3 production workload
- Begin to collect Archive Metadata from ATLAS and CMS
 - Statistical analysis of dataset fragmentation over tapes

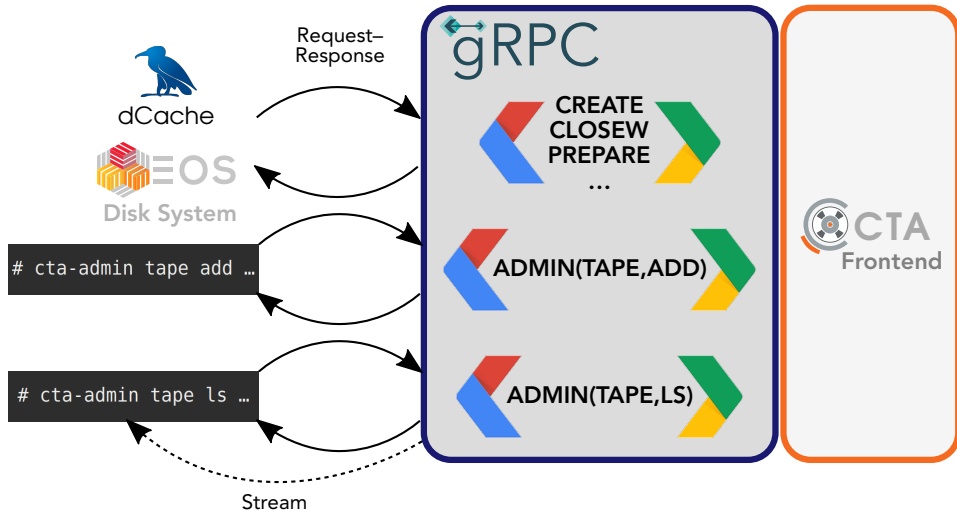
CTA Software : Plans for 2024

- Improved metrics,
monitor internal subsystem timings
 - Investigate OpenTelemetry/Prometheus
- Remove dependency on Oracle libraries
for sites running PostgreSQL
 - Implement plugin to dynamically load
Oracle or PostgreSQL libraries

CTA Frontend Transport Protocol



CTA Frontend Transport Protocol



CTA Frontend SSI → gRPC

- We don't want to support two Frontend implementations
- We are defining a migration path from SSI to gRPC
 - The protocol buffer definition must be the same
 - The request message processing code must be common
- See The CERN Tape Archive Beyond CERN—an Open Source Data Archival System for HEP, CHEP 2023
Forthcoming, accepted for publication in conference proceedings

CTA Frontend SSI → gRPC

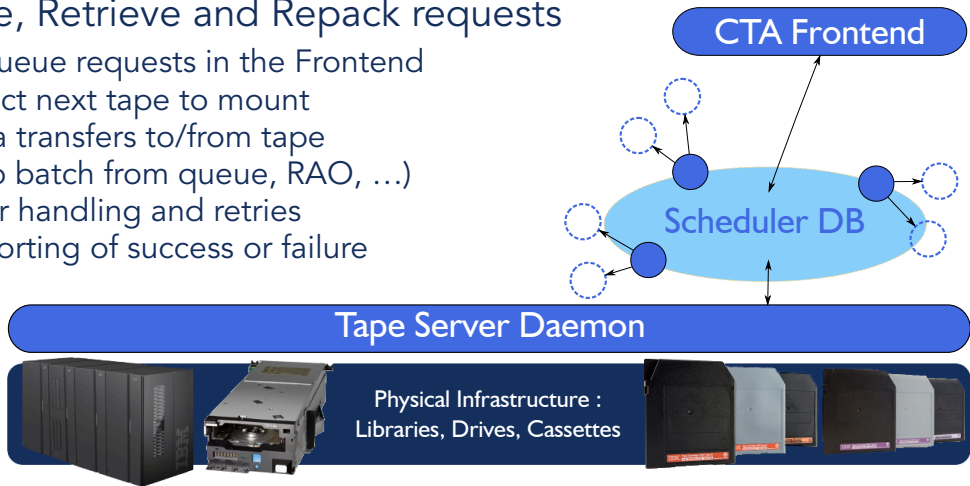
- Refactor SSI Frontend code **DONE**
 - Separate request message transport protocol layer from message processing functions
- Refactor gRPC Frontend code
 - Use common protocol buffer definitions and common message processing code
- Complete implementation of `cta-admin` commands
- Ensure authentication works as expected
 - gRPC token, Kerberos
- Add system tests in CI

CTA Software : Scheduler Database

- The CTA Scheduler controls the workflow and lifecycle of Archive, Retrieve and Repack requests
 - Enqueue requests in the Frontend
 - Select next tape to mount
 - Data transfers to/from tape (pop batch from queue, RAO, ...)
 - Error handling and retries
 - Reporting of success or failure
- The transient data on which the Scheduler works is stored in the Scheduler Database

CTA Software : Scheduler Database

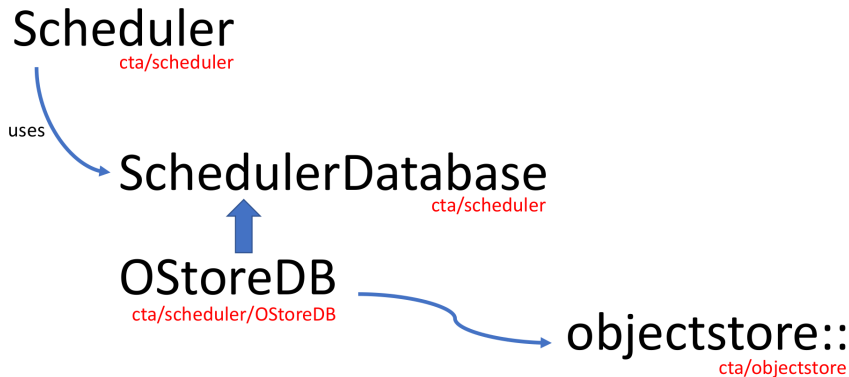
- The CTA Scheduler controls the workflow and lifecycle of Archive, Retrieve and Repack requests
 - Enqueue requests in the Frontend
 - Select next tape to mount
 - Data transfers to/from tape (pop batch from queue, RAO, ...)
 - Error handling and retries
 - Reporting of success or failure



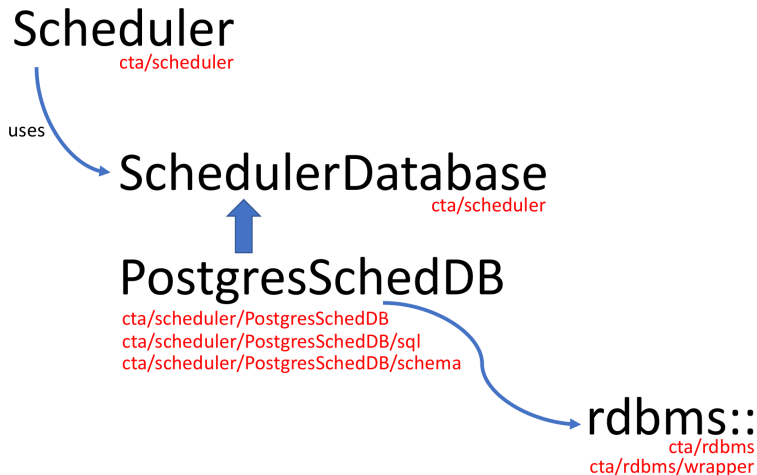
Scheduler Database : Objectstore

- Efficient and works well for FIFO queuing operations (archive/retrieve)
- Requires workarounds for non-FIFO operations (delete, priority queues)
- Limitations of the objectstore
 - Operational issues: difficult to change schema, opaque, difficult to debug scheduling problems, difficult to clean up if something goes wrong
 - Additional software dependency
 - Complexity, difficult to maintain, additional technology for new team members to learn
 - Puts constraints on how the CTA Scheduler code can be modified

Replacing the SchedulerDB component



Replacing the SchedulerDB component



PostgreSQL SchedulerDB Status

- New PostgresSchedDB class to replace OStoreDB class
 - Archive methods were mostly implemented by 1Q 2023
- Development restarted 4Q 2023
 - PostgreSQL Scheduler DB code has been integrated into CI
 - The code compiles and runs (!)
- Archive request handling functionality in progress:
 - Updated SchedulerDB schema
 - CTA Frontend can queue archive requests
 - CTA Tape Server functional implementation in progress
 - Reporting of the CTA Archive workflow **TO DO**

PostgreSQL SchedulerDB Plans

- Finish and test CTA Archive workflow implementation
- Implement functionality and reporting for the other workflows
 - Retrieve
 - Repack
- Goal is to be ready for a functional test by end of 2024
- Targetting repack as initial production use case

CTA Roadmap : Summary

- Important development tasks for 2024
 - Continue evolution of Repack functionality
 - Get ready to consume Archive Metadata from ATLAS/CMS
 - Refactor and complete gRPC Frontend
 - New SchedulerDB back-end
- CTA Website
 - Source Code, Documentation, Presentations and Publications
- CTA Community on Discourse



home.cern