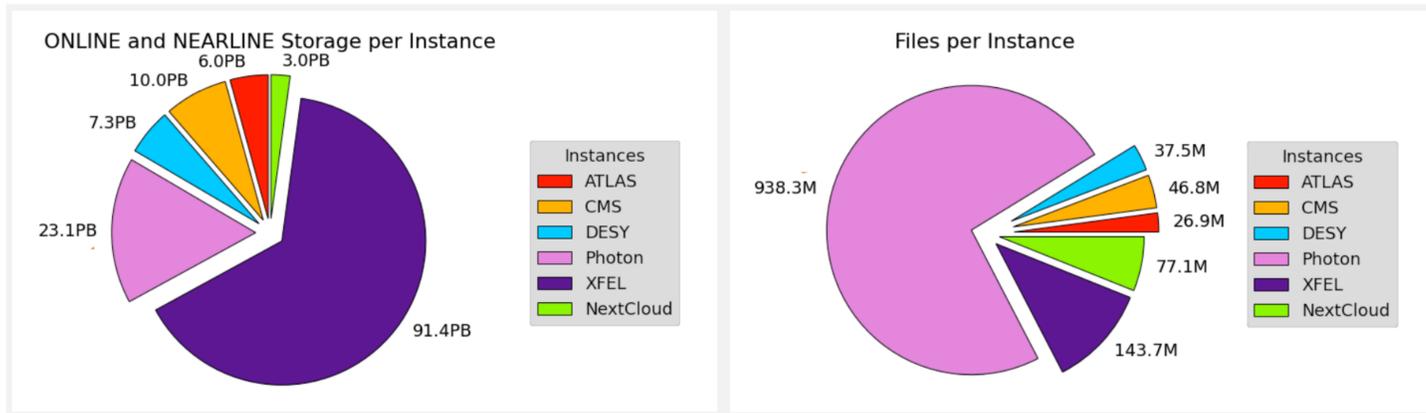


dCache Integration with CERN Tape Archive (CTA)

Mwai Karimi on behalf of the dCache Collaboration
Annual MT Meeting, 26.9-27.9/2022, DESY, HH

DESY as Storage



Tape Tech in place

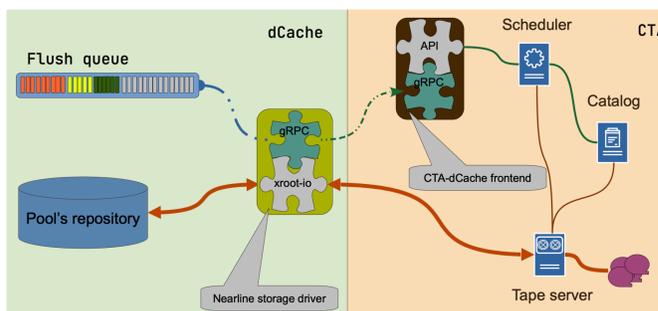
- 2x SL8500 - EOL
 - 26x LTO-8 Drives
- 2x IBM TS4500
 - 20x Jaguar Drives
 - 18x LTO-9 Drives
- TSM (IBM Spectrum project)
- dCache - Interface to HSM
 - Scientific data
 - AFS/MAIL backup
- OSM (Open StorageManager)
 - Since 1994
 - Proprietary software

dCache <=> CTA Integration

- dCache**
- Nearline driver to add (dCache ≥ 7.2.2)
 - Can run in parallel with other HSMs
 - Pre-scheduling on pools should be disabled/reduced
 - File path, uid, gid not preserved

- CTA**
- Additional *cta-dcache* frontend service (packaged as own rpm)
 - Limited to dCache required minimal functionality
 - *cta-frontend* still needed for admin commands

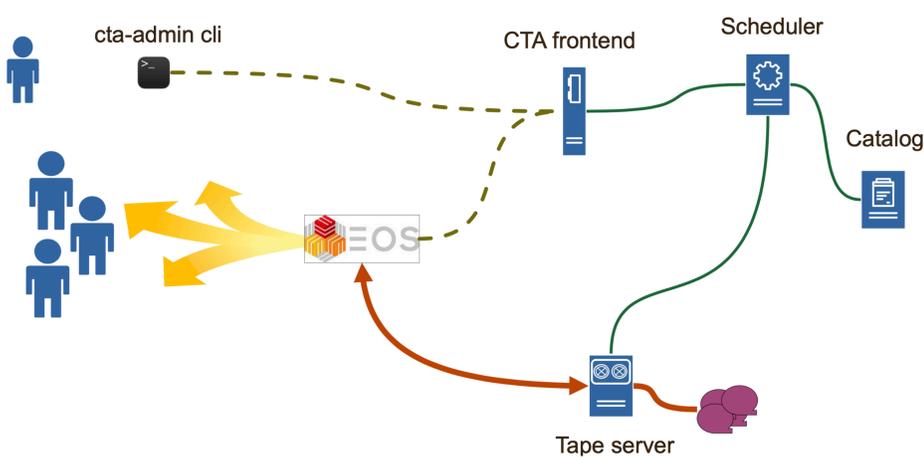
Nearline Storage Driver



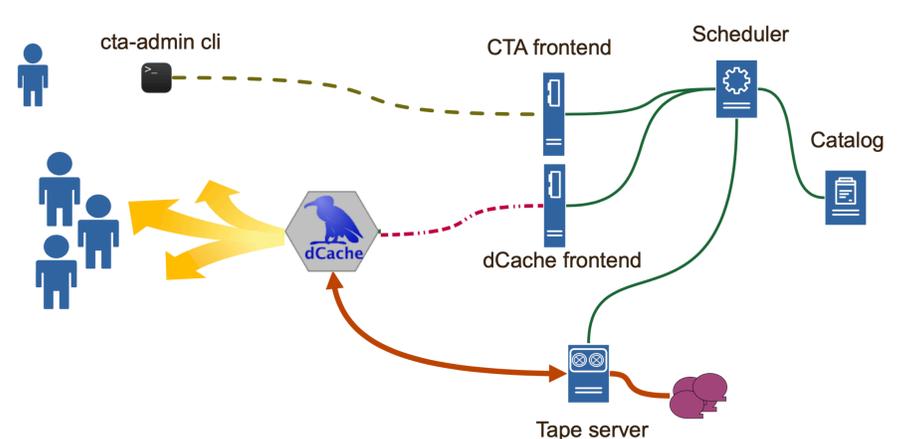
Tape Software Requirements

- Maximise tape efficiency
 - Integrate into DESY ecosystem
 - Integration with dCache tape interface
- Stable operation for the next decade
- Open source
- Wide user and tech community

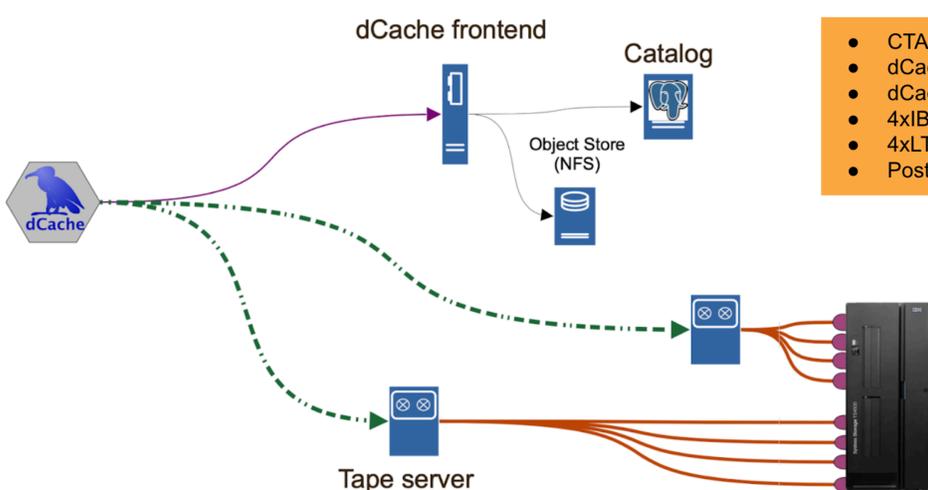
CTA Design



CTA Design with dCache



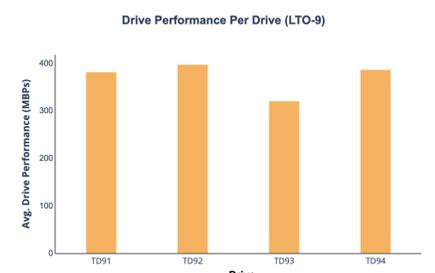
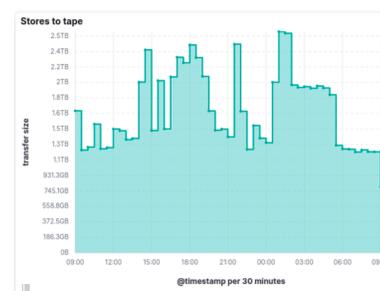
Test Deployment



- CTA 4.7.9
- dCache-cta 0.7.0
- dCache 7.2.8, 8.0.0
- 4xIBM T1160
- 4xLTO-9
- PostgreSQL 14

Tested Use Cases

- Multiple drives on same host
- Multiple tape copies per file
- Multiple Storage classes
- Concurrent stores and recalls



Work in Progress

- Out of the box dCache support
- Puppet config
- Catalog only migration
- Multiple tape format support

