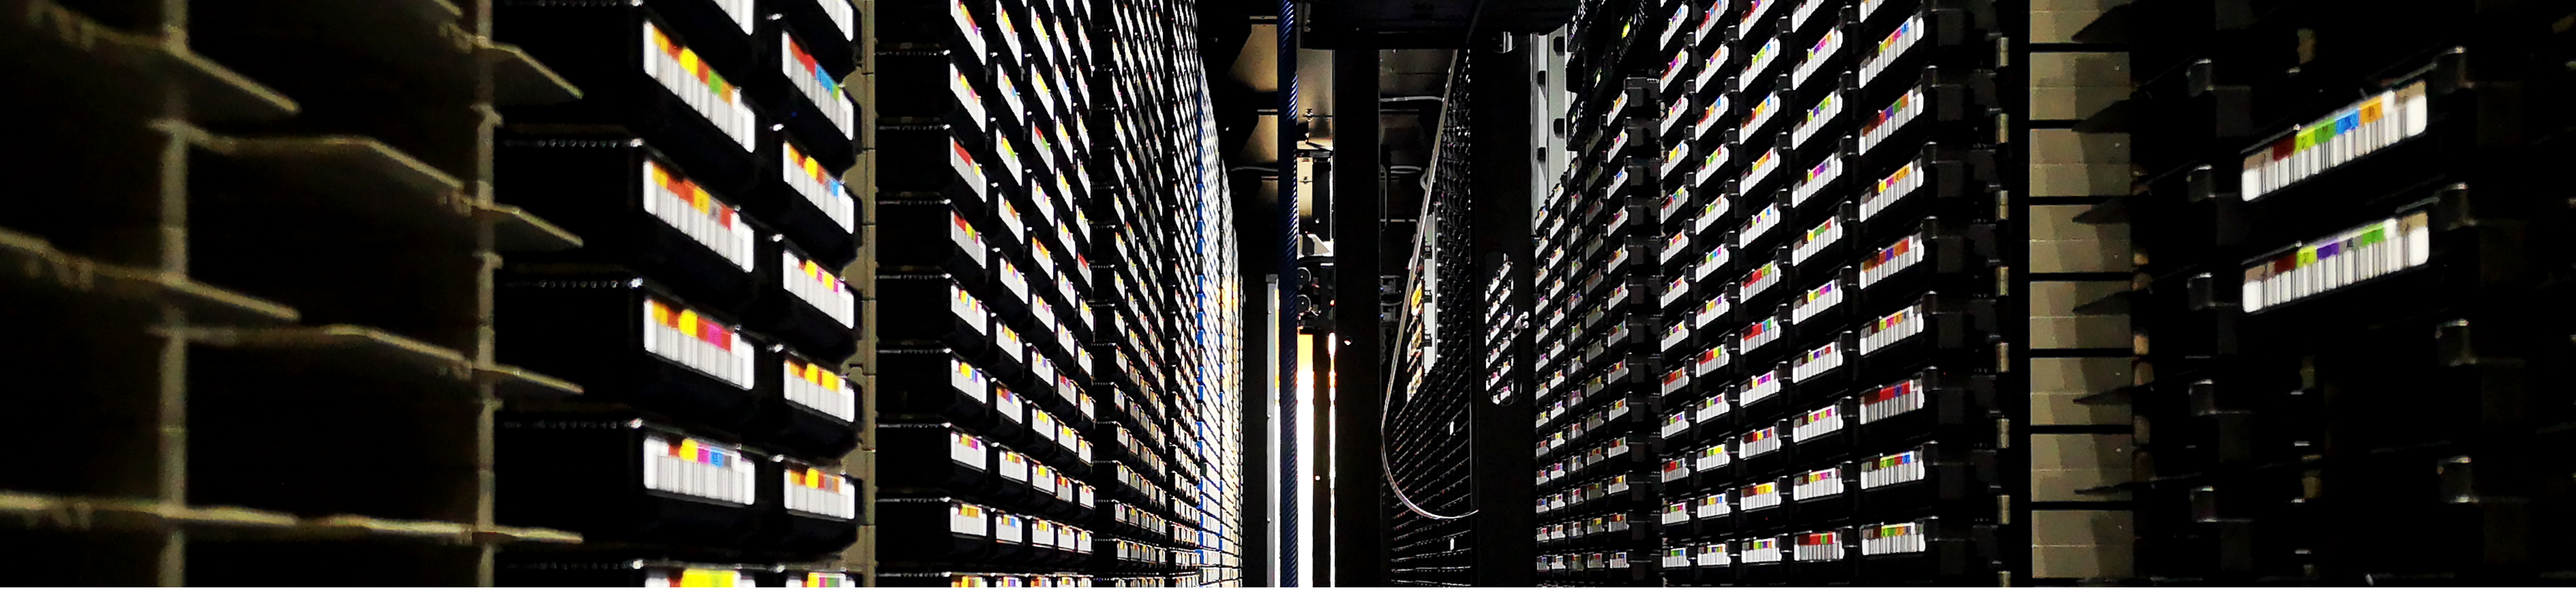

dCache-CTA Integration: One Year in Production at DESY

J.Chodak, T. Mkrtchyan, R. Leuken, P. Suchowski, M. Karimi, C. Voß

Role of Tape Storage at DESY

Tape storage is an integral component of the data management strategy, especially as it serves as the 'Tier-0' storage solution for PETRA-III and EuXFEL experiments. The vast amounts of data generated by these facilities place tremendous pressure on data management infrastructure. This strain is particularly felt in archival storage systems, where resources such as tape libraries often face challenges with accessibility and sharing. To maximise the efficiency of these resources, a close, seamless integration between hardware and software components is essential.

The Cern Tape Archive (CTA) is the archival storage system for the custodial copy of all physics data at CERN and is implemented as the tape back-end to EOS. Despite this, CTA offers flexibility for other frontend implementations making it adaptable for a wide range of environments and use cases beyond CERN. At DESY, CTA runs as the tape back-end to dCache

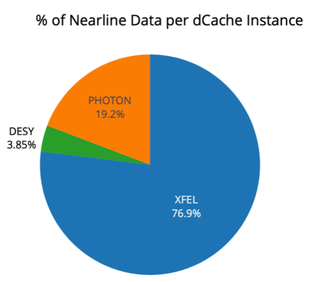

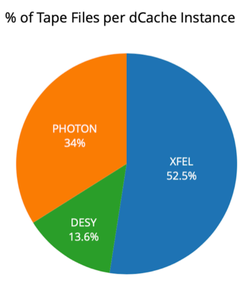

DESY as Storage Provider

DESY plays a pivotal role in data storage for various High-Energy Physics experiments, offering both nearline and online storage through dCache. In addition to serving as a Tier-0 storage site for EuXFEL and PETRA-III, DESY supports further collaborations as a Tier-1 site for Belle-II and Tier-2 site for ATLAS and CMS experiments

DATA Policies
- DESY - Single copy on tape
- PHOTON - Two copies on different media
- XFEL - Two copies, disk+tape

Migration to CTA

Our journey toward integrating CTA began in 2021 with the development of the dCache driver. By the spring of 2022, we had entered an intensive testing phase. Once confident in the system's stability, we initiated the metadata migration on our smallest instance in January 2023. This culminated in a successful full production launch in April 2023. Before CTA, the primary software used to archive data was The Open Storage Manager, which has been in production since 1994.

Tape Infrastructure

Tape Drives
- 24 x LTO9, 26 x TS1160

Library
- 2 x TS4500

Tape Servers
- 15

Deployment

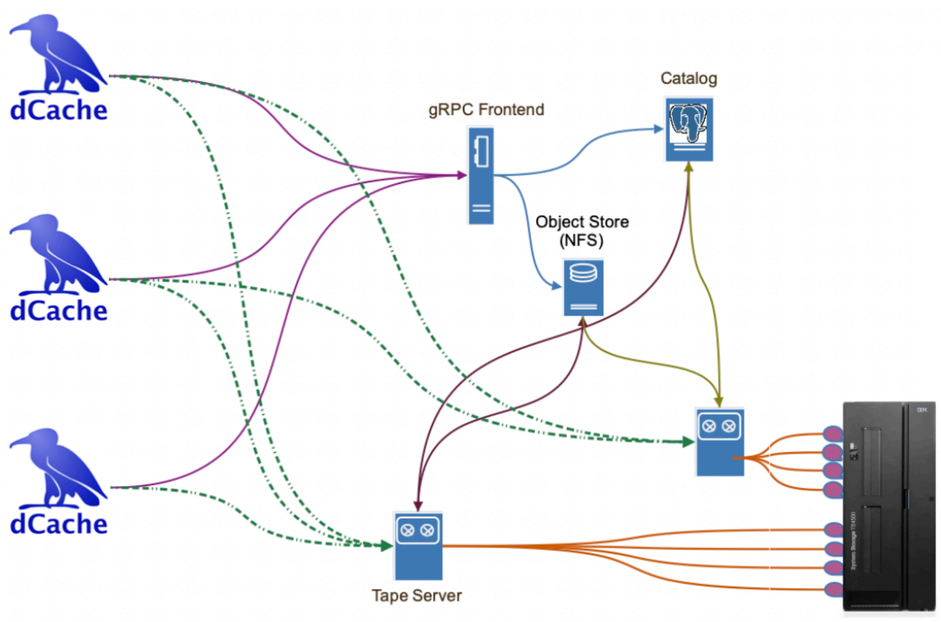

1 Frontend/gRPC
2- 4 drives/host
1 x ObjectStore (NFS)
PostgreSQL (CTA Catalogue)

The dCache Frontend is responsible for submitting requests. We run multiple drive per host setup and use a shared NFS mount serving as the Object Store and PostgreSQL for the tape catalogue.

Performance

With CTA with seen a significant improvement in optimising resource utilisation while maintaining excellent throughput and reliability. CTA's robust architecture has been instrumental in achieving our data management goals, consistently delivering high-performance results, supporting high ingestion rates and large-scale retrievals while maintaining stability

Since moving to production last year, over 50 PB of data has been successfully flushed to tape via CTA. Additionally, more than a petabyte of data has been repacked, and the seamless integration of CTA's operational tools into our environment has significantly streamlined workflows

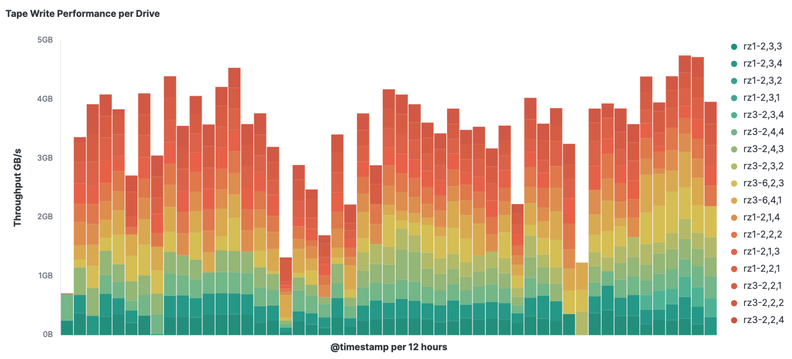

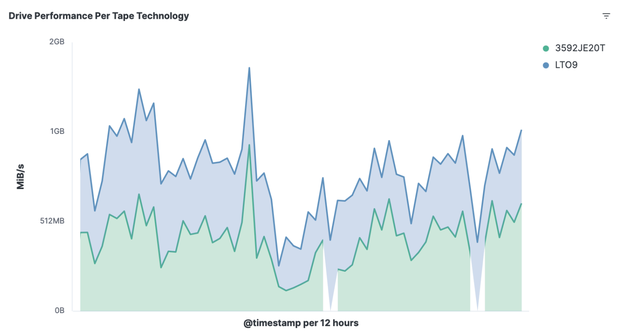

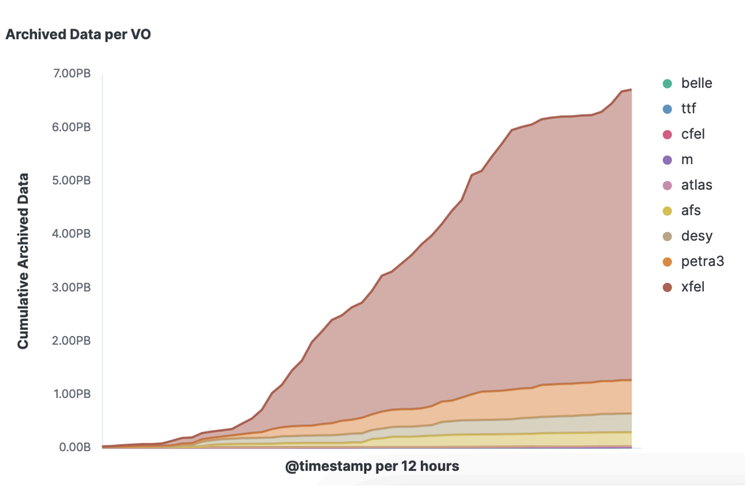

Next Steps

- Extend monitoring
- Fine tune the dCache-Frontend gRPC and dCache nearline driver
- Develop more operational tools to improve the workflow

More Info:

https://github.com/dCache

https://eoscta.docs.cern.ch/latest/

Deutsches Elektronen-Synchrotron DESY
Ein Forschungszentrum der Helmholtz-Gemeinschaft