

CERN Tape Archive: a distributed, reliable and scalable scheduling system

Wednesday, 19 May 2021 11:29 (13 minutes)

The CERN Tape Archive (CTA) provides a tape backend to disk systems and, in conjunction with EOS, is managing the data of the LHC experiments at CERN.

Magnetic tape storage offer the lowest cost per unit volume today, followed by hard disks and flash. In addition, current tape drives deliver a solid bandwidth (typically 360MB/s per device), but at the cost of high latencies, both for mounting a tape in the drive and for positioning when accessing non-adjacent files. As a consequence, the transfer scheduler should queue transfer requests before the volume warranting a tape mount is reached. In spite of these transfer latencies, user-interactive operations should have a low latency.

The scheduling system for CTA was built from the experience gained with CASTOR. Its implementation ensures reliability and predictable performance, while simplifying development and deployment. As CTA is expected to be used for a long time, lock-in to vendors or technologies was minimized.

Finally quality assurance systems were put in place to validate reliability and performance while allowing fast and safe development turnaround.

Primary authors: CANO, Eric (CERN); BAHYL, Vladimir (CERN); CAFFY, Cedric (CERN); CANCIO, German (CERN); DAVIS, Michael (CERN); KEEBLE, Oliver (CERN); Mr KOTLYAR, Viktor (IHEP, Protvino, Russia); LEDUC, Julien (CERN); MURRAY, Steven (CERN)

Presenter: CANO, Eric (CERN)

Session Classification: Storage

Track Classification: Distributed Computing, Data Management and Facilities