

Extension of local dCache instance capacity using national e-infrastructure

Thursday, March 23, 2023 11:00 AM (30 minutes)

The Czech WLCG Tier-2 center for LHC experiments ATLAS and ALICE provides computing and storage services also for several other Virtual Organizations from high energy and astroparticle physics. The center deployed Disk Pool Manager (DPM) for almost all (only ALICE VO uses xrootd servers) supported VOs as a solution for storage until recently. The local capacity was extended by a separate instance of dCache server which was operated by CESNET Data Storage unit in a remote location. The exact location has changed during the project, the distance was between 100 to 300 km. This storage extension was based on HSM and was mapped as a separate ATLAS space token where higher latencies were expected. The intended usage was for a non-automatic backup of the LOCALGROUP disk used by ATLAS users from the Czech Republic. Since the usage was relatively low and the system had only one group of users from the ATLAS VO, the effort required for maintenance and frequent updates was not effective.

The DPM project announced the end of support, and we migrated the main Storage Element in CZ Tier-2 to dCache. This brought a possibility of unified solution for an SE. The dCache system at CESNET was stopped and we started to test a new solution with only one endpoint for all users. CESNET Data Unit also changed the underlying technology for data storage - they moved from HSM to CEPH. We mounted one file system as RADOS block device (RBD) on test dCache server and measured properties of the system to compare with storage based on local disk servers. This solution differs from a solution used in the Nordugrid Tier-1 center, where distributed dCache servers use caching on local ARC Computing Elements. Tests included long term stability of network throughput, duration of transfers of files with sizes from 10 MB to 100 GB and variation of transfer time for cases when several simultaneous transfers are executed. The network tests were first executed on an older diskless server and later on a new dedicated test server with surprisingly different results. We used the same tools also to measure differences in transfer performance between local disk servers which are of different age and connected by different speed. Since the results of tests were satisfactory, we will use the external storage first as a dedicated space token for ATLAS and later as a part of a space token located also on local disk servers. We may also use the solution for other Virtual Organizations if the external available space is increased by a sufficient volume.

Primary author: CHUDOBA, Jiri (Institute of Physics of the CAS, Prague)

Co-authors: MIKULA, Alexandr (FZU); SVATOŠ, Michal (FZU); VOKÁČ, Petr (FZU); CHUDOBA, Michal (Faculty of Mathematics and Physics, Charles University)

Presenter: CHUDOBA, Jiri (Institute of Physics of the CAS, Prague)

Session Classification: Data Management & Big Data

Track Classification: Track 6: Data Management & Big Data