

DDM report

FZK cloud monthly meeting

Cédric Serfon


Ludwig-Maximilians Universität, München

January 12th 2010

Data distribution last month

- Overall transfer efficiency in the last 30 days in the cloud : 86 % (91 % last month). This is a bad efficiency (always stay over 90 % in last months). Due to many staging problem (but the Dashboard only report staging errors not success) + T3s issues.
- Volume transfered in term of files (\sim 3.13M files transfered vs 3.13M last month). 264 MB/s (867 MB/s last month) (summing up all transfers in the cloud) over 30 days.


Data distribution last month



Number of bytes exported to the sites of the cloud

- Low activity during Christmas break.
- No T0 export

Data distribution last month



Number of bytes exported to the T2s of the cloud

- Most of the transfers done via PD2P.

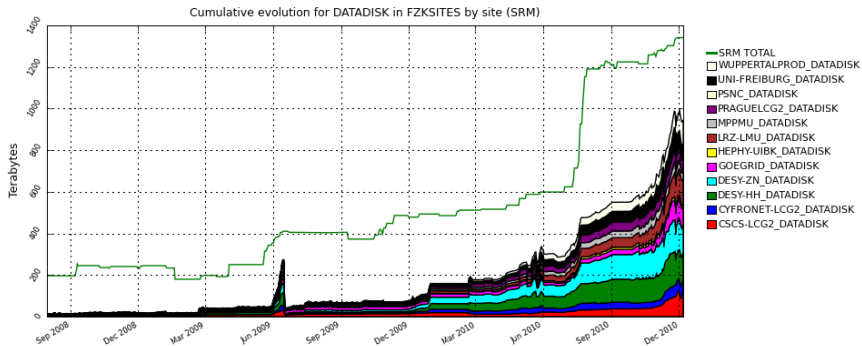
Space Tokens consolidation

- Decision to merge MCDISK and DATADISK announced during SW week approved by CREM.
- This consolidation is starting now in the DE cloud.
 - The aim is to have it done beg. Feb.
 - Actions of moving dataset from MCDISK to DATADISK will be done by DDM ops.
 - Sites are only required to shrink the MCDISK space token and reallocate space to DATADISK.
 - When all the datasets will be removed from MCDISK, the space token will be decommissioned.
- No timescale yet for consolidation at Tier1s.

Going to PD2P

- This was an other proposition during SW week that has been approved.
- The idea is that there is no more data preplacement. PD2P decide of what datasets are shipped to T2s. Cleaning agent (aka "Victor") decide when the datasets need to be deleted according to their popularity.
- Not decided yet when we stop preplacement, but IMHO the PD2P algorithm should be revised before going to full PD2P. For the time being, it triggers the replication of all mc*, data* datasets (except RAW,RDO,HITS and group datasets) as soon as the dataset has been touched once by a user job. Lead to too many transfers and full sites (see plot next slide).

Going to PD2P



Going to PD2P

- With this new model, we cannot decide anymore to have all data in the cloud in ATLAS managed area, i.e. DATADISK.
- We can still get all what we want in LOCALGROUPDISK which is under responsibility of site and cloud via DATRI.

Situation on LOCALGROUPDISK

DE Sites

- Top 3 on all DE sites :

Size (bytes) DN

52790884449296 /C=DE/O=GermanGrid/OU=UniGoettingen/CN=***

46851323776573 /O=GermanGrid/OU=DESY/CN=***

28617259185989 /C=DE/O=GermanGrid/OU=UniWuppertal/CN=***

- 61 (65 in last report) users over 1 TB. 11 (prev. 10) power users (over 10 TB). Total space used : 429 TB (382 TB prev.).