

ALICE FTS data transfers status report

6 Sep 2006

1. Plans and basic principles of data transfer

Part of the ALICE Physics Data Challenge 06 (PDC'06) plans is to integrate the LCG FTS service into the ALICE File Transfer Daemon (FTD) and to test the operation and stability of the combined system. The FTS integration is performed through a FTD-FTS wrapper incorporated in the ALICE Data Management system.

The ALICE computing model foresees three types of data transfers:

- T0-T1: migration of raw data produced at T0 and 1st pass ESDs also produced at T0
- T1-T2, T2-T1: transfers of ESD and MonteCarlo data and AODs for custodial storage respectively
- T1-T1: replication of ESDs and AODs

The main task for PDC'06 is to test the T0-T1 transfers at up to 300 MB/sec tape-to-tape rate, which is the Pb+Pb data replication speed from the CERN T0 during the LCG winter shutdown. The individual T1 share of the 300 MB/sec aggregate rate is calculated according to the mass storage resources pledged for ALICE in the LCG MoU:

- CNAF: 20%
- CCIN2P3: 20%
- GridKA: 20%
- SARA: 10%
- RAL: 10%

It is not necessary to store permanently the transferred data at the remote T1 sites, therefore a tape recycling procedure should be foreseen at the T1s. It is mandatory that the sites provide for ALICE FTS Endpoints and SRM-enabled storage.

As a test of FTS service stability, ALICE expects to have seven days of sustained transfer rates to all T1s.

Table 1 contains the services provided by all T1 sites in terms of SRM endpoints.

SITE	SRM (Tape permanent)	SRM (Disk permanent)	SRM (Tape recycle)	SRM (Disk recycle)
SARA	X	X	X	X
CCIN2P3	X	X		
RAL	X			
CNAF	X		X	
GridKA			X	X

Table 1: Services provided by the T1 sites in terms of SRM endpoints

1. Current status of the exercise

Multiple successful transfers have been performed to all T1 sites involved in the exercise, however the target rate of 300MB/sec sustained for a week has not been met. The exercise

allowed to expose a number of critical areas and as such was a very important step toward achieving full nominal rates under realistic conditions. Major problems were the stability of the FTS service and the associated SRM-enabled storage at the T1 centres. We also had problems with the VO-box (critical for the FTD operation) at CERN and at SARA. Brief summary of the issues per participating centre are given in the table below. For support and problem reporting, ALICE uses the GGUS ticketing system. We had mixed experiences with the service, with some of the problems solved immediately, while others pending for many days, even before being assigned to an expert for solution. As a fall-back solution in critical periods, ALICE also uses direct contact with the FTS experts. We found that this is still the most effective method to solve the outstanding problems in a timely manner.

CNAF	Unstable overall with different sources of errors, most frequent are related to inaccessible storage (CASTOR2). Max rate achieved: 28.4 MB/s.
RAL	Joined the exercise late (site-wide issues with disk storage). Difficult to debug problems, transfers stay in waiting status without clear reason or fail. Resources for ALICE are not sufficient for the duration of the exercise – 1.8 TB of disk without garbage collector. Still in a setup phase.
CCIN2P3	Generally very stable. Problems with srm_get. Max rate achieved: 121.4 MB/s.
SARA	Problems with the LFC catalogue (backend ORACLE), associated to null comments inserted by ALICE. VO-box instabilities have adverse effect on transfers. Max rate achieved 47.6 MB/s.
GridKA	Unstable overall with a variety of error: SRM connection refused and transfer timeouts. Max rate achieved 164.3 MB/s.
CERN	Hardware problems with the VO-box, affecting the transfers to all centres. VO-box replaced.

Table 2: FTS transfers summary of problems per T1

For the 2nd Phase of the exercise (T1-T2, T2-T1 transfers), ALICE has prepared a list of the participating T2 sites and their relations with the hosting T1s. For that phase, ALICE requires LFC local catalogues deployed at all T2 sites. At this moment a certain number of T2 sites are entering the production without this service. Although this will not affect the production phase of the PDC'06, the sites have been made aware that they have to provide this service for the 2nd phase of the FTS exercise.