
Subject: Re: Meeting on Thursday, November 1, 2007 - minutes

Posted by [Kilian Schwarz](#) on Thu, 01 Nov 2007 13:27:26 GMT

[View Forum Message](#) <> [Reply to Message](#)

1) shifts to write protocols of the meetings

It has been agreed that the task will be rotated according to the alphabetical order of the family name. The start will be done by Kilian SCHWARZ.

Next meetings should start at 9:30 am because Peter Malzacher has to leave at 10:45.

order of people to write protocol:

Schwarz Kilian,

Zynovyev Misha,

Andronic Anton,

Ivanov Marian,

Kreshuk Anna,

Malzacher Peter,

Masciocchi Silvia,

Miskowiec Dariusz,

Penso Victor,

Preuss Carsten,

(people who know in advance that they cannot be there or cannot write minutes must find a replacement).

2) status new XRD/PROOF cluster

- new PROOF test cluster is up and working

- temporary setup: redirector lxgrid2 and 4 WNs (lxb284, lxb285, lxb286, lxb287). (only 2 machines can be used by outside users)

- functionality (also with large files and AliRoot) has been tested by Ana and Marian.

- Tree->Draw does not work with current head version of AliRoot. It did work with v4-06-Release. The problem appears locally and in PROOF and is no PROOF related problem.

AliEn SE:

- plain xrdcp to new cluster works

Plan:

- start functioning AliEn SE on new cluster

- move more and more of the new ALICE machines to the cluster

- When farm management is understood, GSIAF will be shut down and all machines will be moved to new cluster. 1.5 days will be needed for the switch. An announcement will be done beforehand.

- start new cluster including all machines as GSIAF

- Marian wants to test memory checker on test PROOF cluster.

3.) AliXRDPROOFtoolkit

- currently a standalone macro in \$ALICE_ROOT/TPC/macros, usage directly from ROOT. The functionality is described in .cxx file.

- this tool should be used ASAP. The presentation can be sent to Federico. Spelling mistakes have to be corrected beforehand.

- for most implemented functions data servers are accessed via lsrn, which would not work on file servers. For file servers a solution still has to be found. Eventually a neutral account with limited rights could be created, or proxy or alien-token might be used for authentication.

- includes: make chains, create file lists, check file inconsistency, copy directories, find file

functionality

- todo: user interface, more functionality, documentation, convince users
- motivation: users need to produce their own data on xrootd. Not everybody in ALICE, though should have the right to stage datasets.
- deleting files should happen via web page where files to be deleted are published. Every week, e.g. Sunday night a cronjob should run which actually removes the published files.
- order on filesystems should be maintained via predefined directory structure, but everybody can do anything under any name.

4.) next PROOF tutorial.

- last November week is PROOF 2007 @ CERN.
- GSI PROOF tutorial shall be held on the week of November 19. Silvia reserved the IT Schulungsraum for Friday, November 23. The tutorial shall be from 10 am to 5 pm with lunch break in between. Participants: around 10 persons.

5.) GRID users jobs at GSI

- Ana's analysis task: compilation problems
 - info sent to CERN experts for clarification, so far no answer
 - so far not enough time for detailed investigation
 - will be done with high priority, is on TODO list.
- Silvia experienced 100% job failure because of NFS time outs
- Anna's monitoring will be put to central place

TO BE INVESTIGATED:

- what do "priorities" in AliEn mean ? / share of aliproduct/user jobs
- from INSERTED to ASSIGNED up to 200 minutes. Before 50 minutes. Why from ASSIGNED to RUNNING 25 minutes ? SAVING to SAVED 20 minutes ?

6.) status of file transfers

- Kilian knows how to do AliEn transfers from CERN but to be able to mirror TPC test data he needs to be daquser.
 - lots of files transferred successfully
 - Nov-1: 2394 transfers, 99.71% efficiency
 - Oct-31: 5698 transfers, 97.20% efficiency
 - Oct-30: 5588 transfers, 99.45% efficiency
 - Oct-29: 2539 transfers, 95.52% efficiency
 - Oct-28: 187 failed transfers, 0% efficiency
 - Oct-27: 210 failed transfers, 0.94% efficiency
 - Oct-26: 189 failed transfers, 5.97% efficiency
- reason: migration problem at CERN. After the upgrade of CASTOR2 and changes of pools it was not possible to get the correct configuration.
- Nov-1: GSI vbox crashed because CERN initiated too many parallel transfers.
 - maximum transfer rate: 0.5 MB/s (ML)
 - average incoming traffic during last week: 35 KB/s (ML)
 - network group measure at the same time incoming data rate of 70 Mb/s, all routed through the vbox (bottleneck !!!)
 - file list of transferred PDC07 data still to be created

7.) first speed test results

- see presentation.
- method: 1-8 AliRoot jobs analysed data in parallel by reading either from local disk on lxb281 or from filesystem lxf51 via xrootd: 1 mio pp events have been read and processed using

AliRoot v4-05-Rev-03.

- result: local data on local disk: 1 mio events in 43 minutes. (after 500 kevents processing goes faster)
 - 8 parallel jobs on local disk: 7:12 hours. Already 2 jobs need more time !!!
 - question: why does train of tasks need more time than single task ?
 - main result: I/O speed roughly the same from local disk and from fileserver
 - maximum speed about 500 to 600 events / s corresponding to 15-20 MB/s, all jobs together. For clarification: 1 job manages to analyse 500 to 600 events /s and 8 jobs, too, but all 8 together. The rate per individual job goes down accordingly
 - remark: old D-Grid machines have no raids, new ALICE 16 GB batch farm machines have raid.
 - remark: raid does not seem to scale at all. Eventually files too small and no hardware caching. Only 1 PROOF worker/machine ?
- Future AliRoot versions may solve this problem in storing several ESDs per file. Test should be repeated with larger files, e.g. AliESDfiends.
- should we move to new data ? 2 mio events with new analysis manager ?

8.) copy of AliEn database to GSI

DONE.

stored in v4-07-Rev-01(32bit)/Full/

/Ideal/

/Residual/

/PartialGeom/Full

/Ideal

/Residual

9.) installation of AliRoot v4-07-Rev-01

32bit: DONE

64bit: DONE (currently problem with large binary output)

- remark: in December run zero supression is needed
 - remark 2 (PBM): with the current GSI budget the planned 1 Gb bandwidth for GSI might be problematic.
-