

Subject: Re: [LCG MB] Minutes of MB meeting - 26.09.2006
From: Dario Barberis <Dario.Barberis@cern.ch>
Date: Tue, 3 Oct 2006 14:23:45 +0200
To: Alberto Aimar <Alberto.Aimar@cern.ch>
CC: "worldwide-lcg-management-board \ (LCG Management Board)"
<worldwide-lcg-management-board@cern.ch>

Dear all,

just a couple of explanations after reading the minutes of last week's MB meeting, I had to leave early, for which I apologise, so I could not comment on some points that were discussed in the second half of the meeting.

1) The 10 TB of database resident data refers to the TAGs. It is the maximum size that would correspond to a nominal year with 2×10^9 events. We have still to test if it is practical at all to have this amount of data in a database, increasing linearly with time, even before thinking how to transfer the data to Tier-1s. There are both s/w and h/w implications that are non-trivial.

In any case, independently of the size, ALL database resident data MUST be replicated to ALL Tier-1s. If we decide to go for 10 TB of TAG data in a database, they will be at all Tier-1s. If this turns out to be impractical, we have to find another solution. We are working on this issue.

2) In the ATLAS Computing Model all "active" data for analysis are disk resident. The cache to tape deals only with the storing of RAW data and locally produced simulated or processed data.

Looking at the Tier-1 data flow diagram that I showed many times (the first time in Mumbai last February), we can work out the nominal data rate to/from tape for the first year of running, with concurrent data reprocessing for a Tier-1 covering a fraction F of ATLAS:

$F * (320 + 200 + 20) * 1.25 = F * 635$ MB/s to tape (RAW from T0, ESD and AOD from T1 reprocessing, 1.25 to take simulation into account)
If you assume a 2-day worth buffer, you need $F * 100$ TB of disk cache only for this flow.

$F * 320 * 1.25 = F * 400$ MB/s stagein for reprocessing (RAW and simulated RAW)
Here we assume that we use staged-in data will stay for, say, 10 days. Then we need $F * 350$ TB in addition.

The calculation can be further refined but these are the orders of magnitude.

Cheers,
Dario

On 2 Oct 2006, at 16:08, Alberto Aimar wrote:

Dear MB Members
the draft minutes of the latest Management Board meeting
are available at

<https://cern.ch/twiki/bin/view/LCG/MbMeetingsMinutes>

Due/urgent actions are in the action list,
both for several MB members. Please check them.

<https://cern.ch/twiki/bin/view/LCG/MbActionList>

Best regards.
Alberto Aimar.

--
CERN /Bldg:28-R-004 /Mail:alberto.aimar@cern.ch /Tel:163158 /Skype:alaimar

Dr Dario Barberis
CERN-PH Department
MailBox E25510
CH-1211 Genève 23 (Switzerland)

Tel.: +41.22.767.1302
Fax.: +41.22.767.8350
