

Dteam meeting 09-06-2009

Present

Jeremy Coles  
Pete Gronbech  
James Thorn  
Alessandra Forti  
James Cullen  
Raja Nandakumar  
Brian Davies  
Dug McNab  
Daniela Bauer  
Sam Skipsey  
Matt Hodges  
Gareth Smith  
Graeme Stewart  
Kashif Mohammad  
Mingchao Ma

**11:00 Experiment problems/issues - STEP09 focus (45')**

Review of weekly issues by experiment/VO

- LHCB

Raja :

Started Step09 last Friday. Progressing very well. 1 billion monte-carlo simulation. Currently running 6-7 thousand Jobs. Transfer of data started yesterday, lhcb requested tier1 to delete tape cache. Will start testing of staging files from tape today and will also start bulk job from today afternoon for step09

Problem at Cambridge, All jobs failed with maradonna error while in Condor-G queue.

Faced a Small Problem at RAL T1 due to castor configuration, but now fixed.

Graeme:

In terms of CPU allocation Glasgow allocates 42% to Atlas and 42% to lhcb, but atlas is running 100's of jobs and lhcb only 18 ?

Raja :

Lhcb does not send jobs to a particular site but it goes to a site if it looks attractive to lhcb.

Graeme :

Is lhcb oversupplied with cpu. As there is no jobs in waiting at Glasgow.

Raja:

Currently running 6000 jobs and quite few are waiting, will check about it.

Pete : Oxford is also not seeing many lhcb jobs. Oxford is

full of atlas, no pilot job coming from lhcb.

CMS

David:

Step09 is low key affair for cms, pre-staging works fine at cms.

We can do better now because data was badly distributed earlier.

Users are complaining that there analysis jobs is slowing  
There was problem initially as Jobs were going to RAL without lazy download, it was corrected.

1200 user analysis job can bring down d-Cache .

Jeremy : Can we know the site distribution of cms and atlas jobs to check the contention between cms and atlas.

Graeme : only at ralpp there is contention between cms and atlas. Atlas would like cms to be running fully at the t1, so they can see what effect it would have on atlas work

Jeremy: is it possible to start cms analysis job at this week

David: I have to ask but not optimistic as people are busy in other things.

- ATLAS

Graeme : its extremely useful , lots of part of system worked and learned a lot from parts which didn't worked.

Data distribution has gone very smoothly to T1's. FZK problem caused a backlog

We have 5 raw data sets which have to transfer from tier 0 to tier1. 25% to 30% of raw data ended up at RAL. Transfer rate of 500 MB/s was recorded through out UK. It's very good and it will keep network people happy.

Cosmic raw from last autumn, converted to ESD, and they get written back to tape, RAL is one of the best sites for reprocessing, 2-3 times nominal rate

One problem at RAL, in spring reprocessing campaign, started sending to 3gb queues, to ensure that it does not go over 2gb limit. Have kept on doing this, but because others are requesting 2gb, scheduler not seeing enough to start atlas jobs. Need to discuss a solution with RAL, even though they have less job slots. Overall Tier 1 doing extremely good

MATT : allow more 3gb jobs to run, alice queues a lot of jobs but atlas has less jobs/

Graeme : we keep small no of jobs queued .if you suddenly free a lot slots then it would be a problem.

300 jobs waiting and 32 jobs running. We will send a mixture of 3 gb and 2gb jobs, if 3gb is picked by 2gb, it will die.

Brian : 500 MB/s over weekend, highest was 1400 MB/s at morning.

Graeme : running production also, problem at QMUL , imperial is fine. RHUL is also fine, oxford is not running much and Birmingham is running nothing. problem at RALPP

Pete : full of atlas jobs

Jeremy : what is the problem at Birmingham

Pete: Not sure, I will check.

Graeme : user analysis jobs were very successful , More than

half million user analysis job in 6 days, analysis jobs are very stressful for tier 2 SE, Lots of site fail in data analysis

Theirs is problem user analysis runs, production system uses different stage for it. it uses 3 access method,

1 Panda uses whatever you have in production, so generally at DPM use rf-cp. But oxford lcg-cp is being used but rfcop would work better.

2 Athena background process stages file n+1 uses lcg-cp a bit like panda puts a load on your srm.

3 Third type is classic lan rfio access

We saw at Glasgow,

Dropped both analysis to 100 jobs each, then started raising the bar slowly, started raising panda up to 500. Then dropped it back to 100 again, then did the same with the wms jobs.

We dropped wms analysis job and found that first time load storage below maximum.

Rfio seems to flatten the storage, but jobs run for 12-24hrs. Tried changing rfio read ahead buffer, changed load pattern on servers but did n't help. Still trying to find out the correct settings.

.

QM lustre setup went down over the weekend, caused by a user trying to write a 20TB file to the lustre, crashed kernels. If you have multiple users on storage and if one behave badly then it will be problematic for other users also.

Pete: implemented channel bonding and changed to rfio and it did help.

Cpu efficiency for lhcb is 40% and for wms analysis job it is 4% and job is aborted after proxy expire.

Alexendra : can block the max no of jobs for different group of jobs.

Graeme: if you kill a lot of wms jobs then may be blacklisted.

Liverpool is blacklisted at the moment.

Tier 2 is slow in data distribution. Three reasons for it.

1 Too few fts slots.

2 too many user analysis jobs

3 if you have some fraction of the transfers that go slow, (eg 80 fast 20% slow), you end up with 80% of transfer slots blocked with the slow transfers. So overall throughput is about 20% Lancaster and RHUL? Seen this.

Best part to see blocked throughput

[http://atladcops.cern.ch:8000/drmon/ftmon\\_tier2s.html](http://atladcops.cern.ch:8000/drmon/ftmon_tier2s.html)

Alexendra : only one pool account is used, we could try to use a random no of pools.

Graeme : dpm uses round robin, so it will take as much data from old as much from new. DPM drain was broken before 1.7 Glasgow did not get round to upgrading.

Pete : we are waiting to dpm to drain so we can rearrange data

- Other

- Experiment blacklisted sites: check on sites in this state

- Site performance

-- <http://pprc.qmul.ac.uk/~lloyd/gridpp/ukgrid.html>

Site observations

\*\*\*\*\*

Tier-1 summary

Tier-2s

- Server headnodes (recommendations) loads vs throughput  
Previously we have problem with dpm headnode but now it is pool headnodes

- Tuning issues for DPM

- LAN bottlenecks

- Failing jobs

- Job efficiencies

Sites should check job efficiency

- Fairshare issues

We have large amount of alice work it may be effecting usage form other vo,s

- UK job distributions

- Tools and cross-checks

- Tools available to site admin's to update

Pete: we would like to know about channel bonding form sites with 10gbps links.

Allesendra : we are planning to test it.

Gareth : Machine room change : wms will move tomorrow and we will carry on for next week and we will be unavailable for next week. We have a network issue also so we may be unavailable on 23rd also.

Outstanding issue of migration to tape.

Main concern is that everybody knows the move is happening, see the blog.

<http://www.gridpp.rl.ac.uk/blog/>

## 11: ROC update (25')

45 ROC update

\*\*\*\*\*

From the EGEE ops meeting: Nothing to report.

From the site reports: RAL sBDII problems.

Fro the PMB: A request to verify site contact information in the GOCDB.  
From the PMB: Clarify status of top-level BDII deployment  
Manchester has one top level BDII apart from one at Glasgow  
London : its not on high priority.  
We should configure bdii client so if one top level fails it can automatically switches to other top BDII. Can we check it before moving of BDII  
Discussion: Can we simply swing the DNS alias to an alternate tBDII (see what happens).

WLCG update

\*\*\*\*\*

There is a pre-GDB on storage support at CERN today. The agenda is here:

<http://indico.cern.ch/conferenceDisplay.py?confId=51915>.

On Wednesday there is a GDB.

<http://indico.cern.ch/conferenceDisplay.py?confId=45476>. The agenda consists of a number of regular update items plus a short review of STEP09 so far and a summary of HEPiX. Duncan will be the T2 rep.

Ticket status

\*\*\*\*\*

[https://gus.fzk.de/download/escalationreports/roc/html/20090608\\_EscalationReport\\_ROCs.html](https://gus.fzk.de/download/escalationreports/roc/html/20090608_EscalationReport_ROCs.html)

40624 - DN publishing. In progress.

48353 - ATLAS-ECDF. Issue with file access.

## 12:10 AOB (05')

- ROD review meeting on Thursday at 14:00.

Regional operator on duty next week, plan meeting on Thursday and may be Friday.

Availability figures for Manchester were poor because the data was taken from gstat and that had not been changed.

GIIS.

[11:11:56] Raja Nandakumar Site ranking from LHCB is at <http://santinel.home.cern.ch/santinel/cgi-bin/lcg-voview>

[11:12:17] Raja Nandakumar Glasgow : ERT = 74492 min

[11:12:39] Raja Nandakumar Oxford : ERT = 901598 min

[11:40:24] Pete Gronbech /etc/shift.conf

[11:40:47] Pete Gronbech RFIO IOBUFSIZE 134217728

[11:41:47] Sam Skipsey Pete: how much memory do your disk servers have? With that setting, (some) of our disk servers were exhausting their memory during the last week.

[11:42:14] Pete Gronbech 16GB

[ [11:43:02] Sam Skipsey Our older servers, which were the dodgy ones, have only 8GB.

[11:45:43] Graeme Stewart [http://atladcops.cern.ch:8000/drmon/ftmon\\_tier2s.html](http://atladcops.cern.ch:8000/drmon/ftmon_tier2s.html)

[11:49:53] Sam Skipsey DPM drain was mentioned - Brian even brought it up as UK feedback about DPM.

[11:50:16] Alessandra Forti shouldn't we give this feedback to michel?

[11:50:25] Alessandra Forti he is schedule to talk about dpm

[11:50:33] Alessandra Forti if I'm not mistaken

[11:51:13] Pete Gronbech already given his talk, this am,

[11:51:28] Alessandra Forti Ah

[11:52:03] Sam Skipsey And hence the feedback I mentioned was given after Michel's talk.

[11:53:36] Graeme Stewart Matt: the problem at RAL is this

[11:53:37] Graeme Stewart 018 (11032632.049.000) 06/08 11:34:49 Globus job submission failed! Reason: 22 the job manager failed to create an internal script argument file ...

[11:53:49] Graeme Stewart i have run out of inodes on the CE again!

[11:59:07] Matt Hodges Derek left some notes on this. OK for me to just clean this up?

[11:59:15] Graeme Stewart Yes, please do

[11:59:32] Graeme Stewart I usually run a `find . -mtime +1`

[11:59:42] Graeme Stewart there are a few running jobs

[12:05:47] Gareth Smith Tier 1 blog (including STEP 09 comments) at: <http://www.gridpp.rl.ac.uk/blog/>

[12:06:32] Matt Hodges Graeme: clean up done. I asked Derek about future mitigation.