

## DTEAM Meeting 10<sup>th</sup> May 2011

### List of Attendees:

Andrew McNab  
Andrew Washbrook  
Ben Waugh  
Brian Davies  
Catalin Condurache  
Chris Brew  
Daniela Bauer  
David Crooks  
Derek Ross  
Duncan Rand  
Elena Korolkova  
Ewan Mac Mahon  
Gareth Smith  
Govind Songara  
Jeremy Coles  
John Bland  
Mark Mitchell  
Mark Slater  
Mingchao Ma  
Peter Gronbech  
Raja Nandakumar  
Rob Fay  
Rob Harper  
Sam Skipsey  
Santanu Das

### Experiment and Problem issues

LHCb : Raja has discussed with in lhcb that jobs at T2 should not run longer than 24 hours. He confirmed that sites can happily limit LHCb to 24 hour queues only. LHCb uses only queues which are published in information system. There are minor issues with different sites which are being discussed through ggus tickets. LHCb page may have been unavailable temporarily due to VOMS server update

CMS : Nothing to report

Atlas : Report Provided by Elena

Atlas report for the week 4-10 May 2011

=====

RAL:

Disk server gdss206, part of the ATLASDATADISK space token: two faulty hard drives on the machine were replaced. This machine was put back in production on Saturday(7th May) afternoon.

UKI-LT2-QMUL:

DT (power supply problem) Friday- Saturday morning. Was back in production.  
srm errors on Monday ([https://gus.fzk.de/ws/ticket\\_info.php?ticket=70384](https://gus.fzk.de/ws/ticket_info.php?ticket=70384)). Was blacklisted and set offline. The problem was solved and QM was back in all atlas activities last night. Was set in brokeroff in analysis by mistake. I set it back online. HC team is asked to investigate the problem.

UKI-LT2-RHUL:

was in DT Fiday-Monday morning (network problem). Back in all atlas activities.

new queue for new WNs was set up and under test.

UKI-LT2-UCL:

was in DT. Still offline in production and analysis. Test jobs run > 13h (instead of 30 Min) and failing.

UKI-SCOTGRID-GLASGOW:

Lost files after power cut. [https://gus.fzk.de/ws/ticket\\_info.php?ticket=70363](https://gus.fzk.de/ws/ticket_info.php?ticket=70363). Files declared as lost.

checksum error [https://gus.fzk.de/ws/ticket\\_info.php?ticket=70358](https://gus.fzk.de/ws/ticket_info.php?ticket=70358)

Sam thinks it's a network or firewall problem.

Some sites are available and some are missing. It is unusual error with switches

UKI-SOUTHGRID-BHAM-HEP

Problem with SQL server.

[https://gus.fzk.de/ws/ticket\\_info.php?ticket=70359](https://gus.fzk.de/ws/ticket_info.php?ticket=70359). Was set offline and blacklisted. Was in DT. Back from DT last night. Trasfers are still failing. UKI-SOUTHGRID-BHAM-HEP excluded from DDM again this morning.

-----  
From: atlas-ssb-notifications-noreply@cern.ch

Subject: [ATLAS SSB Notification] Cloud UK: Daily Résumé (Tue May 10, 2011)

Date: 10 May 2011 08:30:34 GMT+01:00

To: ATLAS UK Cloud Support <atlas-support-cloud-UK@cern.ch>

Cc: atlas-adc-ssb-devs@listbox.cern.ch

Cloud UK info:

UKI-LT2-UCL-HEP

brokeroff in analysis activity since Apr 26 09:30

Space token UKI-LT2-UCL-HEP\_HOTDISK under 20% (Free:0.196 Total:1.074) since May 4 09:30

UKI-NORTHGRID-MAN-HEP

ggus 69336 State:on hold Date:2011-04-04 Info:rftp failures at MAN-HEP

UKI-SCOTGRID-GLASGOW

ggus 70358 State:in progress Date:2011-05-07 Info:checksum error in RAL-LCG2\_DATADISK

UKI-SOUTHGRID-BHAM-HEP

ggus 70359 State:in progress Date:2011-05-08 Info:repeated transfer failure from UKI-SOUTHGRID-BHAM-HEP\_PRODDISK due to SRM error

offline in analysis activity since May 8 09:30

UKI-SOUTHGRID-OX-HEP

Space token UKI-SOUTHGRID-OX-HEP\_DATADISK under 20% (Free:8.676 Total:58.0) since Apr 29 18:24

UKI-SOUTHGRID-RALPP

Space token UKI-SOUTHGRID-RALPP\_HOTDISK under 20% (Free:0.15 Total:0.999) since Apr 29 18:24

UK cloud

savannah 120817 Date:2011-05-08 11:55 Info:"UKI-SOUTHGRID-BHAM-HEP set offline"

savannah 120564 Date:2011-04-26 00:45 Info:"UKI-LT2-UCL-HEP is set offline. The site in downtime"

**CE Testing before production:** Santanu raised the issue that whether it is possible to test a CE for atlas before being in production. Elena said that it is OK but you should use the same CE for production when it start passing test which was not the case with Cambridge

Ewan mentioned that atlas is completely ignoring information provided by information system. Like Oxford had a lcg-CE long time back and when the same hostname was used for CREAMCE it took enormous amount of time and many mails for enabling Atlas jobs for that CREAMCE because panda server was still considering it lcg-CE.

**ROD Team status:**

Nothing special to report from ROD team . Kashif is going to update gridppnagios and after this update it will start testing glexec . It will be tested for those sites only who have added glexec service in GOCDB

**Tier 1 state :** Successfully patched oracle backend database for LFC and FTS services. Some update are available for creamce and WMS and it is planning to update it. Moved CMS to different queue with long memory limit. Saw some load issue with CMS castor server but later it turnout to be cms configuration problem.

Jeremy asked about whether it is sufficient if T1 provide written report. The general feeling was that it would be better if some one from tier1 be present in meeting so it will keep a link between tier1 and T2 and T2 sites would also get advance information about impending updates or downtime so T2 sites can plan their downtime accordingly.

**Operation Security :**

Checking results at <https://pakiti.egi.eu> No critical update required

This week's GDB (<http://indico.cern.ch/conferenceDisplay.py?confId=106644>) : No T2 reprehensive from UK but everyone is encouraged to attend it through EVO even for one session.

**Availability Figures :**

Availability figures for April has been released. EGI policy have changed so now if a site availability is less than 75% for three consecutive months then sites would be banned. It is unlikely to affect any UK sites as availability figures are well above the mark for most of the time.

Tickets : Escalated tickets

[https://gus.fzk.de/download/escalationreports/roc/html/20110509\\_EscalationReport\\_ROCs.html](https://gus.fzk.de/download/escalationreports/roc/html/20110509_EscalationReport_ROCs.html)

57746 : Its going for long time so if anyone interested he can have a look.

GGUS open tickets: <http://tinyurl.com/6z6uq5v>

Sites can look at this url to see if they have any tickets open against them.

Chris mentioned about ticket which was allotted to Bristol as well QMUL. Problem solved at QMUL but ticket still open because he is not sure about Bristol.

**Security Challenge :** SSC5 will be at EGI level and every NGI will choose 4 sites for security challenge. It would be tested by Atlas.SSC5 will be across whole WLCG and sites would be measured through different metric like

response time, knowledge of procedures and competence in tackling the problem etc. It is important that UK site should come with good metrics result.

When SSC5 will be completed then same framework can be use for security challenge SS4 within NGI and that can be more wide spread.

For SSC5 sites were chosen randomly and anonymously by T2 manager.

## GlEXEC & Argus status

Snapshot taken from [http://www.gridpp.ac.uk/wiki/Site\\_status\\_and\\_plans](http://www.gridpp.ac.uk/wiki/Site_status_and_plans) for yesterday's PMB:

Deployed: 5 sites (26%)  
In progress: 6 sites (32%)  
To start: 5 sites (26%)  
No plans: 3 sites (16%)  
A tarball install is wanted/needed by 6 sites (for primary or shared clusters).

The individual site positions are:

1. Brunel: glEXEC to be deployed in April.
2. Imperial: testing a hacked tarball glEXEC install with SGE. Waiting for EMI-1 ARGUS release.
3. QMUL: Not yet deployed. Need version compatible with tarball WN install.
4. RHUL: Installing ARGUS now. Plan to install glEXEC by 20th May.
5. UCL: Waiting for tarball installation (will deploy on HEP and Legion).
6. Lancaster: Had earlier test installation. Tarball install wanted.
7. Liverpool: ARGUS running. GlEXEC on test node. Waiting for request to fully rollout.
8. Manchester: ARGUS installed. Working on glEXEC.
9. Sheffield: Will deploy in June.
10. Durham: glEXEC not needed for supported VOs. Could be deployed.
11. ECDF: No firm objections but have no current plan to deploy. SGE.
12. Glasgow: SCAS and glEXEC deployed since November. ARGUS planned for May. Will put in production upon request.
13. Birmingham: Deployed and in testing on local cluster. Tarball needed for shared cluster.
14. Bristol: Waiting for better resourced sites to deploy first.
15. Cambridge: Reviewing compatibility with Condor batch system.
16. EFDA-JET: Not a major analysis site. Low priority.
17. Oxford: Deployed across WNs with ARGUS server back end.
18. RALPP: Installed and working.
19. RAL Tier-1: Installed and working.

- What issues have you been having?

- Feedback from those sites that have deployed plus pointers to wiki entries that will help other sites.

[11:02:02] Daniela Bauer Duncan has just disappeared, no idea where to.

[11:02:16] Sam Skipsey Probably a cunning plan to avoid minutes...

[11:02:27] Skype Bridge joined

[11:02:29] Jeremy Coles Kashif is taking minutes today.

[11:02:53] Daniela Bauer @Sam: Indeed

[11:03:13] Daniela Bauer I think this should go in the minutes

[11:03:49] Jeremy Coles The chat does go into the minutes!

[11:03:54] Queen Mary, U London London, U.K. joined

[11:15:31] Ewan Mac Mahon I do mostly try to read T1 broadcasts, but .....

[11:23:55] Ewan Mac Mahon I think the problem is that there's a difference between 'notified site' and 'assigned'.

[11:27:27] Andrew Washbrook <http://lhcbweb.pic.es/DIRAC/LHCb-Production/visitor/jobs/SiteSummary/display> gives a server error

[11:27:33] Andrew Washbrook (for LHCB)  
[11:28:18] Raja Nandakumar VOMS is under maintenance  
[11:28:42] Andrew Washbrook ok thanks  
[11:28:46] Daniela Bauer We only have one queue and I don't want to start going to the 1 queue per VO system.  
[11:28:55] Elena Korolkova I heard that atlas voms server is back  
[11:29:04] Daniela Bauer but of course sge in principle supports it.  
[11:29:13] Queen Mary, U London London, U.K. I missed the notice, but have had other things to worry about...  
[11:29:17] Raja Nandakumar But this is not one queue per vo  
[11:29:41] Raja Nandakumar This is the different lengths of jobs you want in your system.  
[11:30:12] Jeremy Coles ATLAS report now uploaded.

Ewan Mac Mahon I do mostly try to read T1 broadcasts, but .....

[11:17:50] Rob Fay joined  
[11:23:55] Ewan Mac Mahon I think the problem is that there's a difference between 'notified site' and 'assigned'.  
[11:27:27] Andrew Washbrook <http://lhcbweb.pic.es/DIRAC/LHCB-Production/visitor/jobs/SiteSummary/display> gives a server error  
[11:27:33] Andrew Washbrook (for LHCB)  
[11:28:18] Raja Nandakumar VOMS is under maintenance  
[11:28:42] Andrew Washbrook ok thanks  
[11:28:46] Daniela Bauer We only have one queue and I don't want to start going to the 1 queue per VO system.  
[11:28:55] Elena Korolkova I heard that atlas voms server is back  
[11:29:04] Daniela Bauer but of course sge in principle supports it.  
[11:29:13] Queen Mary, U London London, U.K. I missed the notice, but have had other things to worry about...  
[11:29:17] Raja Nandakumar But this is not one queue per vo  
[11:29:41] Raja Nandakumar This is the different lengths of jobs you want in your system.  
[11:30:12] Jeremy Coles ATLAS report now uploaded.  
[11:39:01] Daniela Bauer @Raja: When we had different length queues we always had trouble with users not specifying their minimum CPU and they ended up in the wrong queue, so we gave up on that.  
[11:40:02] Daniela Bauer In fact as Valdimir noticed lhcb pilot jobs end up in our 1 h queue, which is meant for SAM tests et al  
[11:41:51] Elena Korolkova exactly  
[11:41:54] Raja Nandakumar Daniela, The description of the different VOMS groups and roles is in the VO-card  
[11:42:06] Ewan Mac Mahon FWIW Oxford has four queues, a one hour test queue that's only open to dteam and ops, and 12 24 and 48 hour queues that are open to everyone else.  
[11:42:38] Ewan Mac Mahon We don't seem to get complaints about people winding up in the 12 hour queue when they don't want to.  
[11:42:54] Daniela Bauer @Raja, I know, but I don't want to have a special queue for LHCB and that's what it boils down to.  
[11:43:32] Daniela Bauer Of course if MHCb wants to run the risk and run jobs 48 h are welcome to, but of course more can go wrong in 48 h than 24.  
[11:43:33] Raja Nandakumar Daniela - but if you want jobs of different lengths, what other way do you have?  
[11:44:12] Daniela Bauer The user determines how long their job is, not me. All LHCB needs to do is stop pulling jobs after 24 h.  
[11:44:26] Raja Nandakumar The actual jobs do not run for 48 hours of course. Only pilots do.  
[11:44:38] Daniela Bauer Most jobs here run 12 h, but if someone needs more time they are welcome to have a go.  
[11:45:05] Sam Skipsey Raja - the pilots \*are\* the jobs, as far as the batch system is concerned.  
[11:45:10] Raja Nandakumar Just curious - if most jobs run only for 12 hours, why give 48 hour queues?  
[11:45:23] Raja Nandakumar Why not just a single queue of 24 hours for everyone?  
[11:45:50] Daniela Bauer Because we support 25 VOs and I don't have the time to look after every one of them.  
[11:46:46] Raja Nandakumar Sam - fair enough. It was just to point out that the actual jobs are shorter and less likely to have problems. Pilot jobs can have higher rates of failure than actual jobs.  
[11:47:24] Daniela Bauer It's not a big deal, it's just that 48 h jobs live a much more dangerous life than 24 h ones (especially if we have to shut down the cluster on short notice)  
[11:47:43] Daniela Bauer but in the long run it'll probably all come out in the wash....  
[11:48:19] Raja Nandakumar Daniela - I still do not understand : if you do not want jobs longer than 24 hours from any VO, why at all are you running 48 hour queues? No matter how many VO-s you have.  
[11:49:40] Daniela Bauer It's for your benefit, not mine.  
[11:49:53] Raja Nandakumar Just to point out - Imperial has been quite good for LHCB  
[11:50:13] Daniela Bauer Sometimes people want to run long jobs even if 90% of their jobs are short) and it saves them (and me) from having to reconfigure anything.  
[11:50:25] Ewan Mac Mahon IIRC the original reason for wanting to limit LHCB to 24 hour queues was for the site - LHCB are quite happy with longer ones.  
[11:50:47] Ewan Mac Mahon The problem was that some sites weren't happy with LHCB blocking all their job slots for two days.  
[11:50:58] Raja Nandakumar Thanks. We are happy with 48 hour queues  
[11:51:01] Ewan Mac Mahon If you don't have a problem with that, then you don't need to change anything.  
[11:51:07] Daniela Bauer We have CMS take care of this.  
[11:52:12] Ewan Mac Mahon (Also I don't think I do have a problem with this at Oxford, so I'm probably not going to change anything)  
[11:52:29] Daniela Bauer Until the next security update when I kill 300 LHCB jobs because I can't wait 2 days to drain the queues  
[11:54:17] Raja Nandakumar Uh - oh. You will then have to listen to my lecture that you are welcome to drain the site rather than kill running jobs ...  
[11:54:40] Daniela Bauer @Raja, indeed I do remember that lecture  
[11:56:41] Govind Songara We suffering high load on batch server due to NFS of /home are of pooluser, how does other site deals with it.  
[11:57:43] Govind Songara I found on TB support thread someone copy it each worker node locally  
[11:58:36] Daniela Bauer Govind, what jobs are these ? If you have shared home dirs the jobs shouldn't run in the home dirs.  
[12:01:06] Mohammad kashif <https://samnag023.cern.ch/nagios/cgi-bin/status.cgi?hostgroup=United+Kingdom&style=detail>  
[12:01:21] Mingchao Ma back in a minute

[12:02:36] Govind Songara Daniela, i think cms jobs had lot of things copy to home shared area  
[12:04:30] Mingchao Ma back  
[12:05:13] Mohammad kashif <http://www.gridpp.ac.uk/wiki/Oxford>  
[12:05:57] Mark Mitchell [http://www.scotgrid.ac.uk/wiki/index.php/Glasgow\\_GLite\\_gLExec\\_installation\\_and\\_configuration](http://www.scotgrid.ac.uk/wiki/index.php/Glasgow_GLite_gLExec_installation_and_configuration)  
[12:06:49] Daniela Bauer @Govind: You are runnign prd jobs or user jobs ? (I assume its prod ?)  
[12:07:06] Govind Songara prd  
[12:07:50] Daniela Bauer Do you have an account name for me so I can have a look ?  
[12:08:00] Daniela Bauer Our prd homes are pretty empty  
[12:08:03] Jeremy Coles [http://gvdev.cern.ch/GRIDVIEW/downloads/Reports/201104/wlwg/WLCG\\_Tier2\\_Apr2011.pdf](http://gvdev.cern.ch/GRIDVIEW/downloads/Reports/201104/wlwg/WLCG_Tier2_Apr2011.pdf)  
[12:08:07] Govind Songara Ok ..i will email you  
[12:10:33] Ben Waugh left  
[12:10:35] Mark Slater left  
[12:10:35] Derek Ross left  
[12:10:35] David Crooks left  
[12:10:35] Catalin Condurache left  
[12:10:36] Brian Davies left  
[12:10:37] Ewan Mac Mahon left  
[12:10:38] John Bland left  
[12:10:38] Sam Skipsey left  
[12:10:39] Gareth Smith left  
[12:10:39] Mingchao Ma left  
[12:10:39] Raja Nandakumar left  
[12:10:40] Andrew McNab left  
[12:10:40] Rob Fay left  
[12:10:40] Mark Mitchell left  
[12:10:40] Rob Harper left  
[12:10:40] Elena Korolkova left  
[12:10:42] Andrew Washbrook left  
[12:10:42] Chris Brew left  
[12:10:43] Santanu Das left  
[12:10:47] Govind Songara left  
[12:10:48] Queen Mary, U London London, U.K. left