

## Ops Meeting 050213

Pete Gronbech (minutes), John Green, Sam Skipsy, John Bland, Daniela Bauer, Steve Jones, Mark Norman, Elena Korolkova, Mark Mitchell, Mohammad Kashif, Jeremy Coles(Chair), Matt Doidge, John Hill, Chris Walker, Gareth Roy, Robert Frank, Gareth Smith, Mark Slater, Raja Nandakumar, Daniel Traynor, Duncan Rand, Andrew Washbrook, Rob Harper, Emyr James, Duncan Rand, Chris Brew, Ewan MacMahon, Alessandra Forti, Wahid Bhimji

CMS, DB: Nothing exciting this week all UK sites back in green. RALPP had had problems but ok now, Bristol is good. Brunel are now using xrootd so no longer have to use DPM hacks. What are Atlas using at Brunel? Almost certainly rfcv as Atlas not really ready to automatically use xrootd.

LHCb RN: Nothing to report everything quiet, no major problems from UK, some minor issues across the grid but nothing major.

Now Durham has upgraded they may be available to be in the mask. RN asked JC to let him know when they would be ready for testing.

Does LHCb use the GOCDDB when submitting to the ce, no it is manual. SAM jobs will continue until they are physically banned when pilots have failed. They then do check the GOCDDB and if it is down they block it otherwise they open a ticket.

Sheffield CE was in downtime but incoming jobs came in and caused problems. The WN acted as a black hole.

LHCb does have an RSS module to check the GOCDDB but it is not yet reliable enough to use automatically.

SAM jobs will still come if the site is banned, but not production jobs.

KM if a site is in downtime SAM jobs always continue coming. If they want to stop getting all jobs you can run a command on the cream ce to stop getting jobs (set offline) or set queues in draining. SJ: There is a need for a document to describe this.

If you are publishing in the bdii then you will attract jobs. If you have two ce's driving the same WNs then you need to disable the ce in downtime to avoid problems.

**ACTION** SJ: Will check for documentation and work with JC & CW to document if required.

Atlas:

AF: This week QMUL had problems with jobs that got stuck in transfer state. They did not arrive in FTS, eventually the jobs failed with timeout. There has been much exchange with Storm SRM people. CW is happy to upgrade StoRM if required. He is currently running the EMI 1 version. He has not upgraded to EMI 2 as often upgrades are problematic. Problems are probably load related, and a race condition comes up, as the problem is not seen always. Ueda has spotted DDM being slow sometimes.

Seems like it is returning some control characters to FTS which goes to the Atlas s/w and is not handled properly. Only happening at QMUL so not attracting much support.

A command has been supplied to check the s/w packages installed at a site:

```
export ATLAS_LOCAL_ROOT_BASE=/cvmfs/atlas.cern.ch/repo/ATLASLocalRootBase
```

```
source $ATLAS_LOCAL_ROOT_BASE/user/atlasLocalSetup.sh
```

```
diagnostics
```

```
checkOS
```

At Glasgow they did not change the list of packages they installed, but there must have been different dependencies.

JC: Is there a CVMFS useful commands page? AF: No but we should have one.

Oxford and RALPP had storage load related issues for Atlas and have tickets. Oxford appears to be the RFIIO high load issue.

Using GridPP VOMS server for ngs VO, all sites should remove the ngs voms server from all VO configurations.

Bristol, Durham, Glasgow, Oxford, qmul, rhul, Sheffield had been seen trying to contact the old NGS voms server. Most have enabled the new one but may still have references to the old one. (This is not critical but should be tidied up)

EMI Upgrades , well done everyone. One or two sites still behind. Eg Edinburgh some WNs not yet using the new EMI tar ball. Should be completed this week.

Please update the ticket with the plan if you have not already.

Some other countries having problems with the tar ball. Two issues, python path, not working with pacman. Other one (we never thought to test) if you use the SRM commands (No one in UK does) it is extremely verbose. The JAVA setup is too verbose.

These were never tested in the UK, in fact the SRM commands were broken in the old glite tar ball and no one ever noticed. It is Canadian sites that are complaining. If you are using CVMFS you don't need pacman.

IC is upgraded now.

There is a workaround for CMS sites for the python path issue. They have to unset the path, it is documented.

Monitoring page links have been updated.

WLCG ops coord info this week, standard agenda.

Glue validator is going to be a Nagios test.

Tier 1 report

Bdii now running new s/w. Three machines behind alias and all upgraded now. Looks good so far (no longer having to automatically restart demons)

SL6 WN test queue. CW: should he encourage small VOs to test against it. GS will find out if all VOs are supported.

AFS client planning to be removed on WNs. Assuming other sites do not install AFS client, it is certain

Storage and data management:

Discussions last week over the possibility of a workshop. At the moment not planning on doing one at the GEI Forum. Could have a BoF session with those that are there. Still favour setting one up later.

No Accounting news.

Need to review documentation more regularly.

<https://www.gridpp.ac.uk/php/KeyDocs.php?sort=area>

If there are any docs that should be retired then they should let SJ know.

Monitoring: Will focus on probes and sharing useful tools, email David Crooks and JC or the list if you have something to offer.

AB has updated the status page.

JG Warning Red Hat released a large number of Java critical updates. Once they are in Scientific Linux we should upgrade asap.

NGS VOMS will be switched off this week, they will setup a web page redirecting people to the new GridPP one. RF will email TB-support once he knows the time.

Tickets:

[http://www.gridpp.ac.uk/wiki/Operations\\_Bulletin\\_Latest](http://www.gridpp.ac.uk/wiki/Operations_Bulletin_Latest)

T1 had a packet loss perfsonar ticket.

FTS problems being investigated.

Neiss have been enabled at RAL WMS.

CMS tickets do not appear to have CMS set as involved VO.

Glasgow tickets in progress.

A lot of progress at Durham. Old ce's and batch system decommissioned, so there is nothing old at Durham. This eliminated most of the problems.

The Atlas disk migration tickets are being used for tracking rather than a problem as such. Sites are good at dealing with real problem tickets and they often do not get discussed at the ops meeting.

Imperial and RHUL MTU discovery networking problems are being tracked.

CW originally checked and has rechecked since. Glasgow believe it is not a bug and Clydenet are unlikely to fix it.

ECDF pool nodes and some WNs to be updated asap.

[https://www.gridpp.ac.uk/wiki/UK\\_EMI2\\_Deployment](https://www.gridpp.ac.uk/wiki/UK_EMI2_Deployment)

RHUL has not completed its pool node upgrades and Govind is away for the month.

Any sites running EMI1 will need to upgrade to EMI in April.

- We need a process to identify and follow-up discussions that have an impact on documentation (and from other discussions dissemination too).
- It is thought that some tasks such as monitoring need a new focus on sharing enabling solutions (such as scripts for monitoring and configuration management solutions).
- The priority of SL6 and IPv6 is now such that we will increase our engagement in testing efforts.
- There is concern about divergence in small VO needs and middleware capabilities - Ganga and a DIRAC instance may help bridge the developing job submission gap, but effort to provide/support them is a concern.
- In light of encouragement to partner more with industry, it would be useful to clarify with JANET the circumstances under which the academic network can be used where there are commercial implications.
- It is felt that we need a more joined-up (informing) effort in regards to new technologies. [Sites would benefit if testing on behalf of GridPP rather than a single individual or on a site-by-site basis].
- Several new areas are opening up that have the potential to change our daily operations. These include integration of clouds and increased use of federated storage.

AOB there will be a joint PMB ops meeting on the Tuesday morning of the GridPP30 meeting.

Cloud computing and research councils impact agenda.

Program being developed.

Please register soon for GridPP30:

<https://www.gridpp.ac.uk/gridpp30/>

Neasan is always in search of good news items. We need to collate ideas!

Chat Window:

[11:03:45] Duncan Rand [https://cms-site-readiness.web.cern.ch/cms-site-readiness/SiteReadiness/PLOTS/T2\\_R+Wcorr\\_perc\\_last15days.png](https://cms-site-readiness.web.cern.ch/cms-site-readiness/SiteReadiness/PLOTS/T2_R+Wcorr_perc_last15days.png)

[11:07:33] Elena Korolkova Does lhcb use gocb when submitting jobs to ce?

[11:10:40] Jeremy Coles We are just coming to ATLAS

[11:14:03] Mark Slater I believe that's the case

[11:15:41] Elena Korolkova we have two ce's for the same wn's

[11:17:47] Elena Korolkova yes

[11:18:29] Elena Korolkova it was a particular case when ce's was set up and not tuned

[11:18:34] Elena Korolkova yes

[11:22:29] Elena Korolkova It's also because QM is a storm site

[11:22:47] Alessandra Forti export

```
ATLAS_LOCAL_ROOT_BASE=/cvmfs/atlas.cern.ch/repo/ATLASLocalRootBase
```

```
source $ATLAS_LOCAL_ROOT_BASE/user/atlasLocalSetup.sh
```

diagnostics

checkOS

[11:23:09] Emyr James yep...i did that last night. much easier than working with big long lists of packages :-D

[11:23:10] Steve Jones Last week, I got to the meeting (vi SeetheVo) just in time to say goodbye!

[11:24:46] Ewan Mac Mahon What would be useful would be for the VOs to maintain the HepOSLibsSLx package

[11:24:57] Ewan Mac Mahon Do it once, and do it right.

[11:25:04] Elena Korolkova Alessandra, could you send the command to the list

[11:25:08] Elena Korolkova please

[11:25:16] Ewan Mac Mahon Not once, more-or-less correctly, at every single site on the grid.

[11:25:16] Elena Korolkova thanks

[11:26:45] Elena Korolkova The changes can destroy transfers to/from other sites

[11:28:56] Jeremy Coles [https://ggus.eu/ws/ticket\\_info.php?ticket=91029](https://ggus.eu/ws/ticket_info.php?ticket=91029)

[11:31:14] Ewan Mac Mahon Ours was the well(ish) known RFIO DPM load problem

[11:31:32] Ewan Mac Mahon RALPPD is dcache though, which is interestingly different.

[11:31:36] Sam Skipsey What release are you on, Ewan? 1.8.6 is supposed to improve that.

[11:31:47] Sam Skipsey (We're not on it either, and we still see the load issues.)

[11:31:48] Ewan Mac Mahon And no, it wasn't set read only - the load was basically all rfio read load.

[11:32:05] Rob Harper nothing to add -- chris is afk

[11:32:10] Wahid Bhimji how does 1.8.6 help with the load issues sam - I don't believe I heard that -

[11:32:14] Ewan Mac Mahon It wasn't getting much write activity anyway (and it was the would-be source for the failing FTS transfer)

[11:32:29] Sam Skipsey It's the first release that includes some changes to limit load, I think.

[11:32:35] Elena Korolkova If i put the overloaded disk server in RDONLY I'll get other one overloaded. I've checked this

[11:32:43] Sam Skipsey (and, yes, we basically only see read load issues with DPM)

[11:32:47] Wahid Bhimji with rfio ? I thought they won't going to do that

[11:32:58] Ewan Mac Mahon We're all still 1.8.5 atm.

[11:33:32] Wahid Bhimji or is in it in DMLite the load limiting - but there is no rfio plugin for dmlite so it will only work with...

[11:33:41] Duncan Rand when people say read only do they mean that or zero filesystem weight ?

[11:33:51] Sam Skipsey I mean Read Only.

[11:34:02] Ewan Mac Mahon I (at least) am entirely happy with the idea of ATLAS either moving to xrdcp based file-staging or xrootd direct access though.

[11:34:08] Robert Frank not yet

[11:34:24] Sam Skipsey Sure, although our experimental testing of that here didn't seem to show much improvement (with file-staging).

[11:34:29] Duncan Rand the advantage of the latter is that you keep the space available

[11:34:38] Sam Skipsey Erm.

[11:34:44] Sam Skipsey Oh, on the publishing, yes.

[11:35:00] Duncan Rand yes it's less disruptive

[11:35:05] Wahid Bhimji (ewan you are welcome to try swtich to xrootd - perhaps we should role it out to a few prod sites

[11:35:17] Wahid Bhimji @sam - we only tested direct access actually not filestage

[11:35:22] Jeremy Coles [https://ggus.eu/pages/help/help\\_findchildren.php?ticket\\_id=90356](https://ggus.eu/pages/help/help_findchildren.php?ticket_id=90356)

[11:35:27] Elena Korolkova I'm working on that. I need to check whether re-yaim disk server stopped contacting ngs server

[11:35:29] Sam Skipsey Well, we should test filestaging then!

[11:35:33] Ewan Mac Mahon AIUI it's an ATLAS internal config. I'm cool with them trying it out though.

[11:35:38] Sam Skipsey It is.

[11:35:42] Sam Skipsey Wahid can twiddle bits.

[11:35:50] Ewan Mac Mahon I'm not planning on turning rfio off or anything, mind you.

[11:35:59] Ewan Mac Mahon Not for a week or two, anyway.

[11:36:04] Duncan Rand I thought glasgow was going to try xrootd staging as an alternative to rfc

[11:36:09] Wahid Bhimji OK - @ewan yes it's an atlas change (which I can do) - but I think it's a site choice

[11:36:40] Sam Skipsey Duncan: well, we tried xrootd direct access, and that... didn't help, so we looked at some other approaches.

[11:37:19] Wahid Bhimji @sam - we can test it tomorrow - I will change `analy_glasgow_xrootd` - but it may be screwed up due to paul / peter test pilot testing -

[11:37:37] Sam Skipsey True. Still worth a look, though, Wahid.

[11:37:46] Wahid Bhimji I will try it tomorrow

[11:38:06] Wahid Bhimji I put a reminder date of end Feb for the ticket

[11:38:15] Wahid Bhimji so it shouldn't be necessary to update it until then

[11:38:47] Ewan Mac Mahon Feel free to try any combination of xrootd direct, xroot filestage, and webdav filestaging that you feel like. Please just change it back again if it all goes horribly wrong.

[11:39:23] Ewan Mac Mahon It'd be nice to have all the shiny new stuff we've set up actually get some use.

[11:40:07] Sam Skipsey (sorry, I tell a lie, it was 1.8.7 that had some improvements for DPM load things)

[11:45:24] Ewan Mac Mahon Pete - are you typing while unmuted?

[11:47:03] Alessandra Forti someone who isn't muted is typing loudly

[11:47:33] Pete Gronbech sorry

[11:47:45] Jeremy Coles I was waiting for a chance to comment...

[11:49:37] Jeremy Coles <https://www.gridpp.ac.uk/php/KeyDocs.php?sort=area>

[11:56:22] Gareth Smith Update on SL6 Test queue at RAL: This is currently being tested by Atlas. Once we have checked the systems out with Atlas and are happy that they are running OK we will open up to other VOs.

[11:58:54] Wahid Bhimji I have to go at 12 - sorry to arrive late and leave early. Gareth - would be interested to know any issues you encounter in atlas sl6 testing at RAL if they are documented anywhere

[11:59:00] Wahid Bhimji bye

[11:59:02] Wahid Bhimji left

[12:00:53] Gareth Smith I'll see what we can do about reporting the SL6 issues we find. (Some may just be configuration issues at our end).

[12:01:17] Rob Harper RALPP Atlas disk ticket will be getting an update very shortly!

[12:02:15] Ewan Mac Mahon We're still intending to look again at the BNL networking issue, it's just round tuit dependent.

[12:02:56] Ewan Mac Mahon AIUI it's not improved, but I'm not quite ready to drop it yet.

[12:03:06] Ewan Mac Mahon I think the ticket status is a fair reflection of reality.

[12:03:56] Ewan Mac Mahon It hasn't but it will be soon.

[12:04:03] Ewan Mac Mahon The transfer succeeded.

[12:04:49] Alessandra Forti Ewan: are you closing it?

[12:05:23] Gareth Smith left

[12:05:26] John Green left

[12:06:23] Ewan Mac Mahon I've just closed it; I think that makes sense.

[12:07:30] Mark Slater Indeed

[12:07:38] Ewan Mac Mahon

[12:13:20] Alessandra Forti they should be reviewed now that they have updated the site

[12:16:17] Christopher Walker If Clydenet formally say they don't care about the MTU issue, and that it isn't a problem - can you formally record that in the ticket and mark it unsolved on that basis. I'm not marvelously happy about you just closing it because they are taking a long time getting around to doing something about it.

[12:19:21] Mark Mitchell We will pursue this with Clydenet, I can add in their formal response

[12:23:17] Duncan Rand feb



[12:23:17] Raja Nandakumar Apologies - leaving now.

[12:23:20] Raja Nandakumar left

[12:23:33] Christopher Walker Thanks Mark.

[12:23:51] Ewan Mac Mahon Pool nodes still count.

[12:23:58] Daniela Bauer I second what Matt says.

[12:24:13] Daniela Bauer They don't raise any alarms

[12:24:28] Ewan Mac Mahon They're harder to see and monitor, but that just makes them a hidden risk, not less of a risk.

[12:24:57] Ewan Mac Mahon We certainly can cheat and just keep schtum, but let's not be under any illusions that that's anything other than cheating.

[12:25:02] Daniela Bauer well, they will get updated, but there's no point going through a big upheaval for the sake of 6 weeks

[12:25:13] Matt Doidge That's why I was being the Devil's Advocate when I made the suggestion

[12:26:19] Daniela Bauer I guess I shouldn't have put that in teh chat window - no written evidence

[12:26:46] Ewan Mac Mahon I'm inclined to agree, with the caveat that they haven't lagged because of exeptional reasons (e.g. the tarball), they just didn't update on time.

[12:27:41] Ewan Mac Mahon There should really be some recognition that that is not OK.

[12:43:46] Jeremy Coles <http://pprc.qmul.ac.uk/~lloyd/gridpp/hs06.html>

[12:49:57] Jeremy Coles <http://accounting.egi.eu/egi.php>