

Table of Contents

Week of 090202.....	1
WLCG Baseline Versions.....	1
WLCG Service Incident Reports.....	1
GGUS Team / Alarm Tickets during last week.....	1
Daily WLCG Operations Call details.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	3
Wednesday.....	4
Thursday.....	5
Friday.....	6

Week of 090202

WLCG Baseline Versions

WLCG Service Incident Reports

- This section lists WLCG Service Incident Reports from the previous weeks (new or updated only).
- Status can be received, due or open

- Reports due from FZK for LFC/FTS DB b/e problem last w/e & CMS "lost files" (see MB report)

GGUS Team / Alarm Tickets during last week

- Tickets: see https://gus.fzk.de/ws/ticket_search.php and select "VO" - this will give all tickets, including team & alarm - or use the following link:
 - ◆ In the GGUS Escalation reports every Monday

Daily WLCG Operations Call details

To join the call, at 15.00 CE(S)T Monday to Friday inclusive (in CERN 513 R-068) do one of the following:

1. Dial +41227676000 (Main) and enter access code 0119168, or
2. To have the system call you, click here

General Information

- Web archive: <https://mmm.cern.ch/public/archive-list/w/wlcg-operations>
- CERN IT status board: <http://it-support-servicestatus.web.cern.ch/it-support-servicestatus/>
- elogs: <https://prod-grid-logger.cern.ch/elog/CCRC'08+Logbook/>

See the weekly joint operations meeting minutes

Additional Material:

Monday:

Attendance: local(Jamie, Maria, Jean-Philippe, Harry, Markus, Jan, Uli, Gavin, Andrea, Roberto, Simone, Julia);remote(Michael, Michel, Gonzalo, Gareth, JT, Brian).

Experiments round table:

- ATLAS (Simone) - main ongoing activity is MC production. Many jobs running everywhere and a lot of data movement. Over w/e exceeded 1M file transfers per day. Production transfers still run with old version of site services - 10M file test was with new version. Critical to upgrade - will be done by critical clouds, RAL foreseen for this pm. Then DE, IT, DL clouds(?) During w/e various problems at sites - main issue as downtime of LFC@SARA - came back a few hrs ago. Consequence of moving Oracle b/e to RAC? JT - got caught with unexpected situation. LFC for ATLAS moved to RAC i/f. Experience on the cluster had been with LFC R/O. Tuning not appropriate for "mostly writes" - put back on old server. Going to stay with this h/w for ATLAS LFC Until new 3D h/w comes in - ATLAS will have its own m/cs on this new h/w. Write level particularly high this w/e? Simone - big

network intervention tomorrow and day after? JT - yes, not clear why "so big". Router intervention - OPN down. Failover to GEANT? Discuss tomorrow in T1 ops meeting. Simone - leave site services for SARA draining overnight and upgrade tomorrow. JT - a hopefully transparent network intervention at NIKHEF (20" downtime for new firmware - 20" to back out). Planning: one of main activities of next week and one after are pre-staging tests before next reprocessing round. As those run before Xmas - ATLAS site services to bringonline and use SRMLs to check that files are indeed online. Graeme reported that RAL CASTOR SRM has bug that prevents bringing files online. Can this be confirmed? Brian - issue with difference between bringonline command: srmls or checkbringonline. Discuss with CASTOR people at RAL and get back. CASTOR issue generically or specific to RAL setup? Jan - please take to hep project castor srm mailing list with Simone in cc? ASGC & CNAF are also on this list... Brian - clarify main reason RAL is hardest hit is because RAL has largest number of associated T2s - FTS channels etc. Not any other problem? Simone, yes, correct. Harry - what's the size of staging test? Simone: bringonline 20TB worth of data per Tier1 - measure speed to bring online and whether you can rotate buffer as expected - keeping files online for 8 hours and rotate. Test of tape system, ATLAS services, etc.

- CMS (Andrea) - data ops people still working that very small files will be created during reconstruction.
- LHCb (Roberto) - famous dummy MC prod continues. During w/e some issues "grid-wide" due to application used for mc generation. Will restart with newest version of MC simulation chain. FEST: over the weekend we reconstructed all FEST data that could be reconstructed. This week there is no FEST activity running continually. We may perform an afternoon of nominal running but this is to be confirmed with the online people. So far - as reported last week - some problems, mainly Lyon and CNAF. CNAF prob fixed - Isf scheduler 4 castor - now ok. GGUS ticket has no explanation. 2nd issue - main - at Lyon where status of transferred files were "unavailable" and not "online". Saw before Xmas in staging activity. Problem with config of write pool at Lyon? Team GGUS ticket with fairly high priority - still waiting for a solution. Miguel - does data volume change or file size? A: filesize stays the same so more data. 1/6 data is redistributed to T1s, i.e. 1 file every 10'. T1s will also get 20x more data - still "peanuts".
- ALICE (Patricia) - the experiment has ramped up until 1600 concurrent jobs during the weekend. This has helped a lot to the testing of the last WMS submission module which has shown a good behavior. This morning the Russian federation has been fully migrated to the last submission module (it was still pending).

Sites / Services round table:

- RAL (Gareth) - couple of hours break in OPN link Sat evening. Failed over to normal network and came back ok. Have an "at risk" tomorrow for Oracle patches behind CASTOR.
- PIC (Gonzalo) - starting migration WNs to 64b tomorrow. Should be transparent.. 3 CEs that act as gateways to qs. Will deploy new qs. One CE will be reconfigured tomorrow to publish new qs. ~1 week of time when WNs will be reinstalled in bunches. In 1 week will decommission 32b qs. Warning - no hard-coding of q names! New names in IS.

AOB:

- USAG (Maria) - meeting last Thursday with excellent representation, ALICE and CMS reminded that they should use GGUS routinely. A faq page is in preparation explaining differences between TEAM and normal tickets that now go directly to sites.
- SAM + GridView (Nick) - outage Wed morning due to h/w changes. Offline for a few hours.

Tuesday:

Attendance: local(Maria, Jamie, Uli, Julia, Simone, Harry, Miguel, Roberto, Olof, Jan);remote(Gonzalo, Jeremy, Michael, Gareth, Daniele).

Experiments round table:

- ATLAS (Simone) - migration of site services for UK and NDGF clouds took place -took a bit longer than expected, a few probs, completed at 21:00 (e.g. probs recreating db schema). Since then restarted at quite high rate. FTS job slots in RAL filled in a way that is fast enough - improvement over before. ATLAS ready to migrate other boxes at CERN to new version of SS. Today plan is today all serving clouds except SARA & CNAF - backlog for those clouds. Reinstall scratches DB and hence qs lost. These VO boxes need ~12h to drain - probably done tomorrow. This completes upgrade for SS run at CERN. At end of day will contact US people for BNL then T2s in US. Miguel will circulate email with rel. notes & upgrade procedures. Gareth - had some FTS probs overnight. Not directly linked to ATLAS changes. Since sometime last week FTS running on older h/w in prep for move to new computer building. Will bring up new FTS on new h/w in new building and flip across with minimal interruption - old FTS h/w will not move to new building. Somewhat independent of ATLAS changes but ATLAS can now load us even more....
- CMS (Daniele) - on Tuesdays a bit longer following Monday pm OPS. Replay and Cruzet repack in prog. LSF problems seen recently due to misconfig now understood and fixed. MC prod continues - issues fixed rel. WMS. Some new issues affecting MC, e.g. Pisa T2: prod agent related?? Savannah tracked. Transfers - 400TB last week over full phedex top. Reproc up to date. T1 sites getting jobs running. Still digesting transfer reqs. Issue of small files size at IN2P3 now addressed. T1 site issues: IN2P3 filesize issue; all jobs of skim d/s got stuck on WNs at IN2P3 - run state for several days (normally hours). Errors = file opened; unexpected failure from dcache errno 33. CMS site contacts noticed that transfers HPSS to dcache not transferred after Fri down. IN2P3: FTS server having probs transferring out to Florida T2 (at least). FZK: many CMS files lost due to dcache bug - files being deleted while being copied to tape. ASGC: back running CMS jobs since while. PIC, RAL, CNAF, FNAL running stably. ~12 savannah tickets to CMS contacts at T2 sites. No GGUS tickets as CMS-related so far.
- LHCb (Roberto) - tomorrow in pm full nominal rate test for FEST09 from 13:30 - 17:30 - write at 1.8kHz freq - about 25kb event size. Online -> Castor 45-50MB/s. Each t1 gets 1/6 for recons - run ~100 jobs per site. Issues: gridftp server issue & problem of copying out using lcg_utils. Prob due to lcg_utils version still in cert - most likely will grab from AA into dirac dm. (Mainly affects user activity). Lyon: wrong status report (unavail and not online) waiting dcache patch expected Thursday. This brings IN2P3 out of tomorrow's test... Finally: CNAF shared area now fixed! 😊 Run happily 😊 MC activity on T2 site of batch farm - some WNs still point to old shared area but being fixed. Luca - yes, you are right!
- ALICE (Patricia) - currently running 158 concurrent jobs, which are coming from the last production bunch but as Alice already announced it is a non-massive production. It is also important to mention that during the last Alice TF meeting, FIO members were invited (Ulrich and Daniel) to explain the last CPU usage calculations which have been applied to LSF at CERN in order to be used by Alice. It is a very interesting development that will be used also to the experiment level.

Sites / Services round table:

- NL-T1 (JT) -
 - ◆ network router firmware at Nikhef upgraded successfully. Connectivity was lost for a short period of time (order of a minute).

Tuesday:

- ◆ new router at SARA : installation and configuration proceeding ahead of schedule as of 14.00.
- ◆ ATLAS jobs at Nikhef using zero CPU time. Ganga problem (not sure what kind). User has been notified and working with ganga developer to diagnose.

AOB:

Wednesday

Attendance: local(Jamie, MariaG, MariaDZ, Gavin, Alessandro, Uli, Harry, Nick, Steve);remote(Michael, Daniele, JT, Jeremy).

Experiments round table:

- CMS (Daniele) - wrt yesterday 1 site issue: CNAF went out of upgrade of FTS server to 2.1 (glite update 38 on sl4 64bit.) successful - downtime over, expt should notify any issues. Sched downtime of PIC Tuesday 10th-February (announcement was sent today): affecting several CMS+ activities. WMS at CERN only: prompt help (please!) on problem seen only at CERN, i.e. not CNAF. Is it a CMS or WMS? Outstanding since some days: Vincenzo Miccio - last reply to thread some days old. Some follow up please? Some lines in twiki. Nick: GGUS ticket to go with WMS issue at CERN? A: no, Vincenzo contacted wms.support, reply from Ewan. more details
- ATLAS (Alessandro) - yesterday site services put in draining mode for ASGC, CERN, FZK, PIC and IN2P3. Today upgrading on 3 boxes - still ongoing, hopefully will finish before 21:00 (as happened yesterday) FTS in RAL: GGUS ticket open, 5-10% inefficiency. Problem still occurs this morning so maybe still after move back. Jeremy - ticket is being followed up.
- LHCb (Roberto) - LHCb data quality team gave green light after analysing express stream from pit to proceed with full recons at T1s. Data flowing, jobs will come soon. CERN issue with diskserver - news from Jan that rolled back to external gridftp configuration. Jan - few qs to Andrew yesterday - response please! This was done due to probs with lcg_cp fixed in code Dec. Like to revert to internal gridftp asap. Roberto - testing of this activity on AA - several bugs reported.

Sites / Services round table:

- RAL (Gareth) - In preparation for the move to a new computer building the FTS agent was moved to different hardware last week. This aim of this was to reduce the need for any interruption to the service during the move. However, performance issues have arisen and the FTS was moved back to the former hardware this morning. The dates for the move to the new building are still uncertain.
- NL-T1 (JT) - progress on sched downtime in SARA. Was going well, network part fine, problem with SRM. postgres db was very slow. 15' ago had fixed problem with postgres - should be out of sched. down. soon (now out=15:20). NIKHEF: at risk at same time - some network router upgrades. Were apparently successful. Problem targeted by firmware worked. Some ongoing network problems to NIKHEF trying to fix asap. Still within scheduled "at risk". Until 18:00 UTC.
- dCache issue (From Onno Zweers, one of the dCache honchos at SARA) - We saw a lot of IOWait, but we were unable to pinpoint which process caused this other than Postgres. There were no SRM connections from outside. A simple vacuum of Postgres took two hours until we killed it. We did not reboot the machine immediately because first we were trying to find out what was wrong. After reboot, the same vacuum finished normally in much less than half an hour. Before rebooting, we restarted Postgres but that didn't solve it. After reboot, we needed to fix also our Postgres backup script which had a bug.

- DB (MariaG) - ASGC has been reinstated using transportable tablespaces and PIC as source. Now both running in separate setup until tomorrow morning. Waiting for green light from ASGC to rejoin to normal setup.
- MariaDZ GOCDDB / GGUS - NIKHEF & SARA - please check names your sites have in GOCDDB and names in Twiki TierOneContactDetails for Tier1s. Please use same names and same address.
- FZK lost files - CMS data ops looking for copies of files and files that cannot be recovered.
- PIC (RS) - started upgrade to 64bit O/S. Issue with libraries for 32bit compatibility & those installed natively. Problem with packaging of dcache client - both libs in same path.
- gLite 3.0 (Nick) - point from weekly grid operations meeting to retire all glite 3.0 m/w by end April. Discussion needs to start with expts + sites to see if feasible. Sites don't have to upgrade < this date but would be date of end of support. Testing > this date would not include glite 3.0 services. occ.support@cernNOSPAMPLEASE.ch in case of comments. JT - RB = gone? A:Y

AOB:

Thursday

Attendance: local(Jamie, Jean-Philippe, Harry, Alessandro, Andrea, Maria, Roberto, Simone, Uli, Gavin, Nick);remote(Michael, Gareth, Luca).

Experiments round table:

- ATLAS (Alessandro) - upgrade to VO voxes of ASGC, FZK, PIC & TRIKUMF. Still draining SARA & CNAF. Should be finished draining tonight so upgrade tomorrow. Simone - some pre-pre-staging tests. Graeme will report at today's ATLAS meeting.
- CMS (Andrea) - today nothing to report (R.A.S.)
- LHCb (Roberto) - fest activity. Various runs yesterday 15:00 - 22:00. Site by site number of jobs failed / succeeded. All sites v. v. good except Lyon - no jobs due to outstanding issue of wrong status reported by dcache. Awaiting feedback from Lionel (opened GGUS ticket...) CERN: between yesterday and today 40 jobs failed trying to stage files (out of 500...) 1FTS issue at CNAF - were not able to query FTS - migration to SL4 ongoing. Announcement? Luca- also upgraded to FTS 2.1. Roberto - announced? Luca - sure, in gocdb. Problem that quattor config procedure didn't work. Problem also reported by CMS. Only OPS vo enabled at first. CNAF: issue transferring data to CASTOR@CNAF - fixed as run jobs. Luca - glitch in SRM. As soon as I saw team ticket - restarted e-p and up again. Still investigating reason. No clear ideas at the moment. Not the first time - at same time also planning to upgrade this instance.

Sites / Services round table:

- FZK (Andreas) - preliminary "SIR" regarding incident w/ FTS/LFC Jan 24-26:

At GridKa/DE-KIT the FTS/LFC Oracle RAC database backend was down from January 24 to 26 (Sat, app
On Sat, our on-call team immediately received Nagios alerts. From approx. 9:30 on Sat, our DBA wo
Oracle backup archive logs had been filling up the disks.

By trying to add an additional disk, ASM (the Oracle storage manager, i.e. the file system) got b
Due to the fact, that the LFC data was on the affected RAC system and it was unclear if the last
At approx. 16:30 on Sat. she opened an Oracle Service Request. After info/files exchange with an
another supporter (in our CET zone) came back to us on Mon, approx. 11:00. With his aid, the prob

Remarks:

Thursday

* It is unclear to me, why it took more than a day until we got an Oracle supporter in our time zone. As soon as I get to talk to the DBA I will try to clarify on which side mistakes happened.

* My personal opinion: even though the disk to be added to the ASM was not prepared correctly, the

- RAL (Gareth) - brief big ID problem on DB at end of afternoon. Unsched out on FTS this morning. Yesterday reported moving FTS f/e back to h/w used up to ~1 week before. Backed out... As soon as we put service back had disk fault in raid array. Now sorted - hopefully a bit more stable for a while now...
- FTS (Gavin) - which version of FTS? FTS 2.1 is current glite 3.1 prod. release. 2.2 is still in dev - don't wait - upgrade asap to FTS 2.1!
- GGUS (MariaDZ) - GGUS FAQ on direct routing of tickets to sites.
- DBs (MariaG) - try an upgrade of all cluster interconnect private switches to latest firmware prior to the March 14 upgrade of switches in CERN CC. This should allow to decrease # that will be upgraded March 14 but also to see that bonding config is redundant. Will try to do in a transparent way for RAC 5 & 6 last week of Feb.

AOB:

Friday

Attendance: local(Jamie, Gavin, Jean-Philippe, Simone, Uli, Nick, Sophie, Ewan);remote(Gonzalo, Michael, Daniele, Gareth, JT).

Experiments round table:

- CMS(Daniele) - follow-up: small files at IN2P3, as of late morning rerunning of involved skims done at 90% level. subscribed to in2p3 and caf@CERN. Then to t2s. Will cover 80% of caf analysis - should fully address in2p3 experts. GGUS placed on issues with WMS@CERN; support v prompt, discussions in progress, alot understood in last hours, some update uploaded later & close ticket. Report also Monday. Otherwise nothing special... Some T1 issues: ASGC dataops got permission denied error whilst creating subfolder stageout of merge area - Jason had look, problem with namespace, fixed, now ok, v good reaction but still to see if jobs run ok. Some news soon.. CNAF: production FTS server 2.1 - some major problems during downtime due to lack of support to CMS VO there - only dteam enabled - promptly fixed hence no visible impact. Ewan - ticket in today about extra subVO - channel for officially getting such roles?
- ATLAS (Simone) - received several emails concerning FTS. Reminder of Gav's clarification yesterday. Jamie sent reminder with snippet of minutes: point is that FTS 2.2 is not around the corner ☹️ Don't wait as it won't come in next month. Delegation patch will be introduced as patch in FTS 2.1 and not previous ones. As soon as available ATLAS expects sites will install it. "Manual patch" is manpower. Site decision but please consider moving to FTS 2.1 so that can apply patch as soon as it is available. Waiting for FTS 2.2 not realistic... FZK running with 20-80% failure for file transfer / stagein-out. This comes from pnfs overload, due to amount of files and problems with DB load. Would be cured with ~1/2 day operation on pnfs database. FZK would like to wait ~1 month when they will split SRMs per vo. This would mean fairly degraded mode for 1 month. Rod will push for this 'healing' done asap and not in one month. Would like to reinforce - FZK please cure problem now even if extra work - problem especially for MC PROD. All FTS have their own monitor which is functional. CNAF opened FTM to world. NDGF installed a couple of days ok - all sites now have FTM which is good. Daniele - FTM in general. Do we maintain a list of FTM URLs for all Tier1s? Nick - not really! Simone - we have a list! Nick - Savannah ticket requesting that these are published.

Asked that all Tier1s update the list at last weekly ops.

- LHCb (Roberto) - We had some jobs stalled at NIKHEF; Jeff isolated the problem to be most likely due to the network intervention on Wednesday and then dCache not able to recover the broken connection.

Sites / Services round table:

- NL-T1 (JT) - ALICE SAM tests not running.
- WMS & SAM tests failing (Sophie) - a fix coming so that reported as WMS failure. Being looked at by developers...

AOB:

This topic: LCG > WLCGDailyMeetingsWeek090202

Topic revision: r15 - 2009-02-06 - JamieShiers



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback