

# Table of Contents

|   |          |
|---|----------|
| <b>Week of 09504.....</b>                       | <b>1</b> |
| WLCG Baseline Versions.....                     | 1        |
| WLCG Service Incident Reports.....              | 1        |
| GGUS Team / Alarm Tickets during last week..... | 1        |
| Weekly VO Summaries of Site Availability.....   | 1        |
| Daily WLCG Operations Call details.....         | 1        |
| General Information.....                        | 1        |
| Monday:.....                                    | 1        |
| Tuesday:.....                                   | 3        |
| Wednesday.....                                  | 3        |
| Thursday.....                                   | 4        |
| Friday.....                                     | 5        |

# Week of 09504

## WLCG Baseline Versions

## WLCG Service Incident Reports

- This section lists WLCG Service Incident Reports from the previous weeks (new or updated only).
- Service incident with the robotic library outage that occurred on April 25th at CCin2p3.
- Service incident with the cooling failure that occurred on May 3rd at CCin2p3.
- Service incident with tape backend outage that occurred on May 4th at SARA (NL-T1).

## GGUS Team / Alarm Tickets during last week

- Tickets: see [https://gus.fzk.de/ws/ticket\\_search.php](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
- Please observe the procedure to become a LHC Experiment VO TEAM member and other most recent GGUS FAQs.

## Weekly VO Summaries of Site Availability

- ALICE
- ATLAS
- CMS
- LHCb

## 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/>
- M/w PPS Coordination Work Log: PPSCoordinationWorkLog

See the weekly joint operations meeting minutes

Additional Material:

## Monday:

Attendance: local(Harry, Julia, Jean-Philippe, Dirk, Alessandro, Olof, Sophie, Nick, Ewan);remote(Gang, Angela, Daniele, Michael, Michel).

Experiments round table:

- ATLAS reports [↗](#)- Several issues during this long weekend: 1) The machine submitting functional tests and inserting monitoring data into SLS got stuck and was only fixed a few hours ago. Apparently an incorrect quattor template removed the arc (for privileged afs commands) command. ATLAS was however able to monitor the 4 sites receiving the current cosmics run data. 2) The prodsys dashboard stopped working. ATLAS experts were contacted but did not have enough instructions so it was only restarted this morning. To be followed up. 3) The AGLT2 (Michigan) muon calibration tier 2 site SRM endpoint has been very unstable and there is no agreement for out of hours support so ATLAS need to establish better procedures. 4) There are many FTS transfer timeouts on large (5GB) files and ATLAS want some expert advice. DB said CMS have seen such problems and have been looking to increase timeouts when only a few percent of a file remains to be transferred and data is still being moved.
- CMS reports - DB emphasised that the problem of large files at PIC (solved by deleting the datasets) needs a policy rethink by CMS. These files are input to reprocessing where the output files are larger than the input so this contributes to the problem. A ticket was submitted to afs support a week ago with no response yet- Olof will follow this up in FIO. CMS are soliciting input on the planning of middleware, sites and the other experiments for moving to SLC5 and plan to raise this at the next GDB (May 13). Nick reported that worker node middleware was available for slc5.
- ALICE -
- LHCb reports - The report includes an interesting presentation of the results of LHCb Analysis Tests at Tier1's of running a user application over large amounts of input data (up to 25 TB/day).

Sites / Services round table:

IN2P3 - message to WLCG from D.Boutigny, Director: CC-IN2P3 has been impacted with a major cooling problem during the week-end. Yesterday evening one of our chiller went down due to an overpressure alarm. This problem triggered a chain reaction in the cooling system and all the chillers went down. Of course the temperature went up very fast in the computer room, we even had temperature raising up to 50 degrees near the disk units before we had time to shut down everything. The cooling system has been restored during the night and we are now in the process to restart the services. We do not know yet how much broken hardware we will have to fix. I will keep you informed of the evolution of the situation.

ASGC: Have had CASTOR errors due to the well known BigID problem (understood but not fixed yet). Have completed file export to PIC, but a few files stucked in transferring to FNAL. Alessandro asked when they might be ready to restart running the ATLAS functional tests and the answer was the end of this week or the beginning of next but that DPM could be tested very soon.

FZK: Have changed several dcache parameters to try and improve stability. A CMS disk-only pool has developed a disk i/o error - being investigated.

CERN: 1) ATLAS CERN WMS are currently experiencing large backlogs due to resubmission of failing jobs from one DN. The cause of the failure is understood but there is not yet a clear idea on how to shift the backlogs. Investigations are continuing. 2) There was an outage of the myproxy service from 8.50 till 11.15 this morning. See the post-mortem at <https://twiki.cern.ch/twiki/bin/view/FIOgroup/PostMortem20090504> 3) CASTOR upgrades to 2.1.8 this week with ALICE tomorrow morning at 09.00 and ATLAS on Wednesday from 14.00. Each upgrade involves a 4 hour downtime.

AOB:

Monday:

## Tuesday:

Attendance: local(Julia, Nick, Sophie, Harry, Alessandro, Jean-Philippe);remote(with apologies the conference phone link did not work today).

Experiments round table:

- ATLAS reports [↗](#) - They would like to check the castor configuration at ASGC notably the setup of the space tokens. Alessandro will email Jason explaining in detail.
- CMS reports - CMS were running staging tests at CERN where staging stopped for 7 hours then resumed. To be followed up.
- ALICE -
- LHCb reports - 6000 jobs running for the physics MC production MC09. CNAF problem of yesterday was related with the propagation into StoRM of new bunch of local users (pilot users) followed by an ACL problem.

Sites / Services round table:

- ASGC: ATLAS data deletion on castor started yesterday evening, now around 40~50 TB has been removed.
- CERN: This mornings castoralice upgrade completed successfully. A new xrootd plugin is scheduled for Thursday. castoratlas upgrade is scheduled for 14.00 to 16.00 tomorrow.

AOB:

## Wednesday

Attendance: local(Harry, Alessandro, Jean-Philippe, Nick, Olof, MariadZ, Roberto);remote(Jeremy, Gang, Daniele, Ian).

Experiments round table:

- ATLAS reports [↗](#) - By email from Graeme Stewart yesterday: Panda database migration from the Oracle integration service (INTR) to the ATLAS offline production service (ATLR) will happen tomorrow (meaning Wednesday 6 May), which will involve downtime for the panda servers here at CERN. This will put the panda service "at risk" tomorrow afternoon for all clouds except FR and US.

The intervention is scheduled for 1300 UT (1500 CEST) and is anticipated to last for two hours. After the service is restored the switch of backend database will be transparent to clients. During the intervention the panda monitor will run - read access will continue, but write access for CERN clouds will fail and the data for CERN clouds will become a bit stale. Paul has extended the timeouts slightly on the pilot, up to one hour, however we do expect some jobs to fail after this as they cannot update their status in the panda server. Running jobs will be unaffected over the anticipated downtime period 2 hours. Shifters will see failed jobs appear as lost heartbeats 6 hours after the intervention.

After tomorrow's downtime the final act, the migration of the FR and US clouds, should take place on Monday 11th May

AdG ATLAS report: 1) We would like to know the ASGC castor pool names and how they map to space

Tuesday:

tokens. It was agreed he would mail the question to Gang. 2) There was a 1 hour downtime of the LFC at PIC after they installed a new 64-bit front-end. J-PB explained this was because both 32 and 64-bit clients were installed requiring /lib/lib\_uuid and /lib64/lib\_uuid respectively and the 32-bit version was initially missing. This problem does not happen in client versions 1.7.0 and higher. 3) There was a 2 day delay in SARA responding to a ticket following failure of SRM to export data. Apparently the delay was due to a public holiday. There was a reported 36 hour tape backend problem affecting SRM at SARA which may have been the cause - a SIR has been requested. 4) ATLAS are now training new shifters at the pit prior to restarting daytime shifts in June so there will be more elog entries and perhaps poor quality tickets while the training goes on.

Olof queried how ATLAS publicised this afternoon's castoratlas downtime as they were getting tickets from ATLAS shifters and users. AdG agreed shifters should check the IS service status board and will follow up.

- CMS reports - RS queried what were the file level investigations at CNAF leading to closing tickets. DB explained that there had been a small amount of disk data not migrated to tape when CNAF started their recent long scheduled downtime and that on restart these data had been lost. DB also explained that the reported 3 files lost at FZK were now ascribed to CMS prodagent file registration workflow problems.
- ALICE -
- LHCb reports - 1) MC09 production and merging is ongoing at Tier 1 sites. 2) There was an ACL problem with storm at CNAF that was quickly fixed but we still have problems accessing returned turl's this morning.

Sites / Services round table:

- ASGC: Cleaning of ATLAS data on ASGC castor has finished but not yet on DPM. Then some storage will be setup and ATLAS functional tests can be restarted.
- FZK (by email from X.Mol): as no one from our site (FZK) can join in the daily conference, we want

to give a statement regarding our status:

1) Stability of SRM service: The reason for this problem is identified: Several users have a huge number of active GET TURL requests. They probably use a client which does not release finished requests. We are in touch with dCache developers for finding an appropriate fast solution in order to reduce the impact on the production.

2) CMS specific SAM tests, that failed because of one disk-only pool throwing i/o-errors, are green again as this pool has been repaired.

\*CERN Castor: RS asked when castorlhcb will be upgraded and Olof explained that they wanted to run the upgraded castoralice and castoratlas for a week before deciding.

AOB:

## Thursday

Attendance: local(Harry, Sophie, Gavin, Jean-Philippe, Nick, Julia, Steve, MariadZ);remote(Gang, Angela, Jeremy, Gareth, Michel, Daniele, Roberto).

Experiments round table:

Wednesday

- ATLAS reports [↗](#) -
- CMS reports - Main issue is following up on the impact of the partial restoration of services at IN2P3 (see below). Also a large RAW file is failing repeatedly in transfer CERN to FNAL. File size is 50 GB so it is probably simply too big to be transferred before the timeout of 1 hour on the channel. Central teams are leaving up to FNAL to decide if they wish to transfer it manually or take some other actions.
- ALICE -
- LHCb reports - are getting close to the end of producing the first batch of MC09 minimum bias events. After quality tests generation of 10\*\*9 'no truth' events, expected to take some months, will start. Merging jobs are held back at CNAF due to ACL issues of pilot accounts in Storm.

Sites / Services round table:

IN2P3 - message to WLCG from D.Boutigny, Director: After the cooling incident we suffered last Sunday, CC-IN2P3 is now partially back into operation, but today, while people were powering on more CPU servers, we experienced another cooling system trip, fortunately this time it was possible to keep the system under control without tripping the whole computing center.

By beginning of June, a new 600 kW chiller will be put into operation. As it is impossible to speedup this installation we will probably have to live with reduced computing power during the coming month. We are studying the possibility to install and to power a temporary chiller but even if we have good hopes, it is not yet clear whether this is feasible or not within such a short time scale.

At the moment CC-IN2P3 is running at a little bit more than 1/3 of its CPU capacity. The full storage capacity that was available before the incident is back online.

CERN: Following the update of the CA rpms to version 1.29-1 on Wednesday 6th May 2009, several IT and Experiments services were affected by failing client based certificate authentication: ATLAS Central Catalog, Nagios services, SAM Portal, Gridview Portal. The problem is due to the increased number of certificate authorities (the length of the string). The rpms have been rolled back for a cleanup of the number of authorities though this is seen as a temporary solution. A post-mortem is available at <https://twiki.cern.ch/twiki/bin/view/FIOgroup/PostMortem20090507>

FZK: Short server outage. Access to gridmapdir lost. Would like ATLAS to note that the FORTRAN compiler available on their SL5 worker nodes is gfortran and not g77. In reply to a question from RS on SRM problems replied that they have identified 3 users whose applications are not releasing turl's and are trying to find out which clients they are using.

AOB: MariaDZ is looking for help on gstat from ASGC. Apparently the expert (Joanna) is currently at CERN working with Lawrence Field. Also they need information from CMS on GGUS/CMS ticket interworking (see <https://savannah.cern.ch/support/?106911#comment36> [↗](#)).

## Friday

Attendance: local(Harry, Jamie, Sophie, Dirk, Gavin, Nick, Jean-Philippe, Julia, Andrea, Simone);remote(Roberto, Gang, Michael, Jeff, Jeremy, Gareth, Brian).

Experiments round table:

- ATLAS reports [↗](#) - 1) Main activity is export of cosmics as a computing exercise (no physics interest). Exported according to MoU shares but not many data sets so variable quantities. Testing new model

Thursday

of exporting open datasets when they have enough files as will be done in STEP'09. All current subscriptions complete apart from to AGLT2 (Michigan) which is showing a 0 to 30% efficiency. Cosmics finish tonight but tail of exports will continue till Monday. 2) Also in preparation for STEP'09 ATLAS are replacing the fake data to be used for T1 to T1 transfers (does not go to tape). 3) Following the lack of knowledge of the castoratlas upgrade downtime the atlas-support mailing list was found to be out of date. This will be fixed starting with people who need to know. 4) We are querying why after the upgrade ATLASMCDISK has lost 4TB of free space and ATLASDATADISK has lost 20TB while the total disk space has stayed the same. Dirk replied the free space calculation algorithm had changed and he would send Simone a description (see under AOB below). 5) Under questions Jeff reported that since 16.00 yesterday an ATLAS batch application (only Graeme and Hurng jobs are running) is querying their site bdii at 100 Hz. Simone to follow up. 6) Brian would like to know how to get better access to FTS logs to better answer an ATLAS ticket. Gavin suggested they get in touch.

- CMS reports -
- ALICE -
- LHCb reports - 1) Problem with FZK worker nodes having a 4 GB file creation size limit when LHCb merge jobs want to create 5 GB files. 2) WMS at T1's are failing with different error messages. 3) Production UI at CERN must be upgraded to correct the problem with client hanging when SOAP message returned. Sophie reported she had just prepared a new version to be tested before release.

Sites / Services round table:

- ASGC:
  - ◆ Squid service at ASGC was migrated into new box and rma processes was extended with vendor. This caused the SAM frontier and squid tests down for about 11 hours.
  - ◆ 'Big ID' problem hurt ASGC frequently these days. Yesterday afternoon we restarted the SRM server and things got better, but this only last 17 hours, and today during the past 10 hours we received a lot of this 'big id' error in SAM SRM tests. Have informed our local admin to restart the SRM server again. Dirk reported that Shaun at RAL was trying to automate such restarts and he would pass on what he knew.
  - ◆ Simone reported he had now got a description of their castor pools and asked when functional tests could resume. Gang said they had not yet finished cleaning their dpm as CMS were busy using it but Simone said this could wait and they would like to start already testing the critical services of CASTOR and LFC.
- NIKHEF: Had degraded tape capacity yesterday with half their drives having a problem.
- RAL: Currently having 2 CE outages where one had to be modified and another extended. Network reconfiguration next Tuesday morning.
- CERN: 1) Changes to myproxy service next Monday that will lead to reducing the number of boxes. 2) Monday morning the CA upgrade will be done again with 5 unused (by WLCG) CA's removed. Managers should check their services.

AOB: Here is the explanation from Dirk of the 'missing' ATLAS CASTOR space under their point 4 above: It is indeed due to the change in the free space calculation, which was announced in the 2.1.8-6 release notes:

"- Modified the output of `stager\_qry -s` to report as FREE the available free space currently in production as opposed to the total free space. This means that space reports may fluctuate depending on hardware availability, but makes sure users don't rely on free space on unavailable hardware for their space accounting (e.g. through SRM)."

## WLCGDailyMeetingsWeek090504 < LCG < TWiki

This change was introduced to avoid the case that castor or srm report free space to users, but that space currently unavailable as the corresponding disk servers are offline eg because of service interventions.

-- JamieShiers - 28 Apr 2009

---

This topic: LCG > WLCGDailyMeetingsWeek090504

Topic revision: r18 - 2009-05-11 - HarryRenshall



Copyright &© 2008-2022 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](#)