

Table of Contents

Week of 091102.....	1
WLCG Service Incidents, Interventions and Availability.....	1
GGUS section.....	1
Daily WLCG Operations Call details.....	1
Monday:.....	1
Tuesday:.....	2
Wednesday.....	4
Thursday.....	6
Friday.....	7

Week of 091102

WLCG Service Incidents, Interventions and Availability

VO Summaries of Site Availability				SIRs & Broadcasts	
ALICE	ATLAS	CMS	LHCb	WLCG Service Incident Reports	Broadcast archive

GGUS section

- GGUS Escalation reports every Monday (used for WLCG Service Report to MB)
 - ◆ Full search: https://gus.fzk.de/ws/ticket_search.php and select "VO" - this will give all tickets, including team & alarm
- LHCOPN Tickets in GGUS: see https://gus.fzk.de/pages/all_lhcopn.php and change your selection criteria. Future actions are also listed.
 - ◆ If network group participation is necessary, please invite them in time.
- OSG items selected from the GGUS escalation reports.
- Procedure to become a LHC Experiment VO TEAM member
- Other recent GGUS FAQs

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](#)
3. The scod rota for the next few weeks is at [ScodRota](#)

General Information			
CERN IT status board	M/W PPSCoordinationWorkLog	WLCG Baseline Versions	Weekly joint operations meeting minutes

Additional Material:

STEP09	ATLAS	ATLAS logbook	CMS WLCG Blogs
------------------------	-----------------------	-------------------------------	--------------------------------

Monday:

Attendance: local(Gang, Wei, David, Olof, JPB, Simone, Harry, Jan, Patricia, Luca, Alessandro, MariaD, Dirk);remote(Xavier/Kit, Michael/BNL, Daniele/CMS, Alessandro/CNAF, Gareth/RAL, Fabio/IN2P3, Jason/ASGC).

Experiments round table:

- ATLAS - (Simone) tests with FTS 2.2 since last Thu with checksums enabled. No problems over the weekend apart from small problems (5%) with CERN-RAL (possibly on one diskserver down 5%)
Jan: sure not related to FTS upgrade? Simone: doubt it. FTS people will check logs.
- CMS reports - (Daniele) T1s ok apart from ASGC: Ticket to CNAF for phedex monitoring - turned out to be not a cnaf issue, now closed. FNAL: low transfer quality - now closed. IN2P3: one squid server - now closed. RAL: 1k files still to be invalidated after castor db problems - in progress.
Transfer link issue between CNAF - CALTEC (need help from CALTEC). More details about T2

related issues as usual in the CMS twiki.

- ALICE - (Patricia) Since Fri afternoon 2 new VOboxes in production (one to wms and one to cream) worked fine. Will start quattorisation with FIO this week. Alice is ramping up again - some issue to discuss with FIO for myproxy - will report back.
- LHCb reports -

Sites / Services round table:

- Xavier/KIT - reminder downtime tomorrow for dcache LHCb/ALICE upgrade from 9-13
- Micheal/BNL: Diskarray problem on Fri, recovered by Fri noon. Issues after daylight savings time change: clock sync failed and FTS proxy delation stopped working. BNL time server refused to sync some clients (being followed up) - few seconds lack but affecting. The issue was fixed yesterday. Reminder: major outage tomorrow. ATLAS network upgrade from 8-24
- Alessandro/CNAF: ntr
- Gareth/RAL - some tape problems with castor migration queues - vendor engineer on site working on the problem
- Fabio/IN2P3: ntr
- Ron/NL-T1: currently one of two tape libraries down (robot hand replacement) should be up again soon. Dcache was upgraded without major problems, tape acls in place now. The network upgrade originally planned for tomorrow will be rescheduled due to security issues with switch config - working on a fix. Tomorrow: LFC/ATLAS and FTS will move to new Oracle RAC. Simone: LFC/FTS intervention all day? Ron: most of the day
- Jason: new RAC setup progression, some FC hardware problems on the weekend now fixed. In touch with DBA team at CERN. Hope to finish this night with cluster setup - database import more likely tomorrow. Will follow up in a dedicated call tomorrow morning 9:15.
- David: Network intervention announced to ASGC / TRIUMF : tomorrow 11:00 for 8h. (reduced bandwidth: 1.8Gb/s for ASGC, 1Gb/s for TRIUMF)
- Jan: Interventions on castor public tomorrow (goddb registered) and MyProxy - both expected to be transparent for WLCG.
- MariaD: Unscheduled GGUS downtime (which was broadcasted) due to OS misconfiguration after the release.

AOB:

- MariaD: How will xmas coverage for this year discussed? Will raise xmas coverage to appropriate meeting and report results back.

Tuesday:

Attendance: local(Gang, Nick, Jan, Ricardo, Luca, JPB, Wei, Julia, MariaD, Harry, Lola, Dirk);remote(David/NL, Michael/BNL, Gonzalo/PIC, Jason/ASGC, Alessandro/CNAF, Jeremy/GridPP, Gareth/RAL, Daniele/CMS, Angela/KIT).

Monday:

Experiments round table:

- ATLAS - (Simone) Activity stopped for SARA (DB intervention), also for BNL (network upgrade), ASGC (CASTOR recovery). Triumph on backup link and SRM problem at NDGF - several tier 1 sites unavailable. Luckily there is not much activity right now - otherwise this would be a problem.
- CMS reports - (Daniele) FNAL issue: transfer quality from T0 is bad, first thought to be due to proxy expiration problems, but still many errors observed at FNAL. Investigating...
- ALICE - (Lola) Alice is working with site managers on VObox migration.
- LHCb reports -

Sites / Services round table:

- David/NL: small power outage at one site (part of tape affected), power outage (2h) scheduled for tomorrow in Amsterdam, FTS move as scheduled today. Dcache 1.9.4-4 dcache upgrade went ok, site is fully on RH 5. Free space test started to fail for ops VO: don't understand why happened only now. Nick: nothing has changed - this test has always been critical. GGUS ticket https://gus.fzk.de/ws/ticket_info.php?ticket=52694: lcg-voms.cern.ch server usage causes experiment jobs to fail. LCG-voms certs from CERN are not on VO ID card and hence not accepted at NIKHEF. Maybe VO should update ID card (currently containing only voms.cern.ch). Reason could be a package dropped between glite 3.1 and 3.2. MariaD will follow up offline updating the ticket. MariaD spoke to Steve Traylen. Voms host cert is no more required in glite 3.2. More in the ticket.
- Michael/BNL: services down for network upgrade - in progress.
- Gonzalo/PIC: Tue next week scheduled intervention for part of storage services: partial recabling of dcache servers (10-20% of disk-only file will fail on reading). Handling will depend on experiment: CMS prefers to handle the errors due to (temporarily) missing files, LHCb/ATLAS prefer storage downtime (ATLAS will stop panda submission).
- Jason/ASGC: some progress in clusterware and ASM setup done. will follow work plan agreed this morning, Luca: both cluster installed - DB should up today, DB import more likely tomorrow, then back to ASGC for CASTOR restart (hopefully also tomorrow).
- Alessandro/CNAF - ntr
- Jeremy/GridPP - ntr
- Gareth/RAL: CASTOR problem with tape migration still being investigated
- Angela/KIT: ticket update today for LHCb finished. ATLAS copies in SRM failing: jobs seem to delete before they copy. Investigating..
- Jan/CERN: SRM-PUBLIC update to 2.8-2 and CASTORPUBLIC hardware move finished. Issue with CMS CAF disk space: CMS adding data while deployment team was removing disk servers for upgrade - now fixed.

AOB:

- Ricardo/CERN: Still issues with LHCb jobs creating high query load on the batch system: will limit number of jobs on originating queues to protect system.

Wednesday

Attendance: local(Gang, Jamie, Maria, Gavin, Jean-Philippe, Patricia, Roberto Maria, Antonio, Jan);remote(Angela, Michael, Gonzalo, Stephane, John, Onno, Fabio, Jason, Brian, Alessandro, Daniele).

Experiments round table:

- ATLAS - Had one issue yesterday evening - triggered ticket from Simone. Export of data from CERN to T1s. No answer yet... Looks like problem solved around 23:00. Would like post-mortem as to how / why problem was solved. 2nd issue: current problem to export from RAL to T1s and UK T2s. Team ticket - observe problems to PIC and NDGF and UK T2s. Answered by RAL that they are looking at UK T2s. John - aware of issue and looking into it. Think its a firewall issue. Brian - seemed strange: aware of issues T2->RAL but not RAL->T2s. Transfers to/from PIC. Stephane: now RAL to PIC and NDGF. John - all affects all VOs, not just ATLAS! Jan - two overlapping issues: T0merge pool - installation of m/c with SLC5. High error rates to NDGF - still investigating this. Stephane - error rate dropped at 23:00 yesterday. Jan - Simone's ticket still stuck in HELPDESK - suggest to open TEAM ticket for things clearly not for HELPDESK. castor.support@cernNOSPAMPLEASE.ch since yesterday goes to helpdesk.
- CMS reports - Not much to highlight. ASGC status and data to distribute to T1 sites. At moment have d/s that could be available within days - ~100TB. Custodality to ASGC? Re-route cosmics to FNAL and still don't know if ASGC will recover in time.. Start running some backfill jobs at ASGC to be sure see site behaving ok. If so will assign custodality to ASGC as planned. If problems -> IN2P3.
- ALICE - Production - nothing special to mention. Report conclusions on VO box registration / myproxy. Meeting this morning: requests for new ALICE registrations -> ALICE responsables, only when they have checked that these are trustable VO boxes will be sent to px.support. Creating now a full list of trusted nodes in production -> FIO. Gavin - we will be requested lists from other nodes too. Concern about enormous historical list of nodes that must be cleaned up!
- LHCb reports - This is FEST week, processed 20 runs so far. Real detector activity ongoing. Still usual MC prod on going, ~100 jobs in system. A lot of sites complain that SAM test checking VOMS certificate were failing. The instance lcg-voms.cern.ch was not in the VO Id card for LHCb (now added). 2 new tickets: T1 - SARA, all transfers failing, looks like TURL resolution timing out. Many T2 sites - lots of jobs aborting due to application error - particularly for sites recently moving to SLC5. T0 - transparent intervention on CASTOR tomorrow (see detail). LSF server node overloaded - temporarily stopped waiting for improved logic. Lyon - long jobs running on short q, bug in timeleft utility. Fixed. Gonzalo - FEST activity. Should we see raw data arriving at T1? Roberto - usual express stream exercise entirely at CERN; full stream implying recons at T1s. Exercise started yesterday, soon will see rawdata coming, ~0.5TB.

Sites / Services round table:

- KIT: 2 interventions, dCache now on golden release; done in 1/2 time allocated. FTS: needed 1 hour more than expected. Both services running fine.
- BNL: 3 core switches replaced yesterday; also about 500 nodes reconnected and other system work. All completed in scheduled window. Had to shutdown T1 services and wanted to know if restart procedures were ok. Found needed <3hours to restart all services and make them available again. As expected(!)
- PIC: ntr

- RAL: in addition to data transfer problems also migrating data to tape problems. Think by downgrading tape server firmware can fix this. Brian - can confirm that PIC is still transferring data over production network? One T1 only using OPN for T0-T1 traffic. Gonzalo - not quite sure of Q but data arrives to PIC via OPN. Only T2 data via standard internet.
- NL-T1: LHCb already mentioned SARA SRM is having problems after last Monday's dCache 1.9.4-4 upgrade GGUS 52958. ATLAS also affected. OPS VO not affected(?) - SAM tests running fine! Working on it - think maybe pin manager. Contacted dCache developers for support.
- IN2P3: see post below about main problem yesterday about cooling caused by human error. Caused some of WN to be temporarily stopped plus some other services. Almost all power for batch farm now back, trying to understand why/how this happened. Modification to CPU coefficient according to HEPSPEC benchmark - published was below, now ok. dCache intervention next Monday (see below).
- CNAF: Planned intervention with at risk for castor upgrade to 2.1.7-27 but postponed until next week. Will announce when scheduled again. (Transparent anyway).
- ASGC: Following action plan continuing importing data from backup of corrupt db. Finished this morning. Stager & some of core services up and running. DLF - has been recreated from scratch & importing schema files(?) Hope to restart castor service tonight. Maria - nameserver: import finished and DB open; setting up backups now.
- by email from Jason/ASGC: Accounting issue at ASGC T1/T2: apel records have ceased updating more than one month for T2 and two weeks for T1. earlier resolution is to apply gap publisher, but latest release tag seems able to cope with out of memory issue effectively, more details refer to https://gus.fzk.de/ws/ticket_info.php?ticket=51229
- by email from Marc/IN2P3: A cooling system incident has occurred yesterday while doing some maintenance work on the heating system. As a result batch system has been closed down for a couple of hours in the afternoon of Tuesday, November 3rd. The services are now available but the calculating power is still reduced. A more detailed report will follow asap.
On another matter, please remember our dCache outage on coming Monday:
 - ◆ An update of dCache (going to version 1.9.5-4 aka "dCache Golden Release", officially supported for the LHC start-up) is scheduled on Monday November 9th 2009, from 8:00AM till 3:00PM. The impacted services are ccsrmt2.in2p3.fr and ccsrmt2.in2p3.fr. This work will only disturb the four LHC VOs (Alice, Atlas, CMS and LHCb).
The dCache resources of the LHC groups will be drained on Sunday November 8th 2009. As a consequence, all the corresponding jobs will remain queued a few hours prior to and during this intervention.
- CERN: had problem on 1/2 of CEs since mid morning. Some SAM tests failing. Package replaced by 64bit version... has been corrected and should be ok now (LHCb - we noticed some test failures). Not sure how/why - CASTOR client now on CEs and has a dependency... Unrelated to upgrade of CEs. Myproxy update done during morning; hoping most sites were using myproxy.cern.ch which as a 15' DNS timeout. Took over 2h until requests died out - some sites using hostname or IP address (!!!??). New service both under new name and old name just as a temporary solution. Should use myproxy.cern.ch.
- DB: in contact with IT-GS for cleanup of dashboard for ATLAS & CMS. Would appreciate regular report on both ATLAS&CMS cleanup status. Intervention ongoing for patching (security) PDB cluster (ALICE & HARP). Planned intervention at RAL on 11 Nov - again security patching of Oracle. Streams - fully operational after BNL intervention.

AOB:

- CERT & production news - gLite 3.1 update 59 for SL4 includes only new lcg info sites: UI, VO box and WN indirectly.

Thursday

Attendance: local(Gang, Jan, Harry, Alessandro, JPB, Nick, Romain, Graeme, Roberto, Andrea, Jamie, MariaG, Stefan, Dirk);remote(Xavier/KIT,Daniele/CMS, Jeremy/GridPP, Gareth/RAL, Ron/NL-T1, Fabio/IN2P3, Jason/AGC, Alessandro/CNAF).

URGENT ANNOUNCEMENT TO ALL SITES

Romain: severe linux vulnerability was published yesterday, which could allow local user acquire root privileges. Exploits are already circulating. All exposed systems (including all worker nodes) need to be urgently patched/rebooted. More details have been sent out today by EGEE broadcast at 15:48.

URGENT ANNOUNCEMENT TO ALL SITES

Experiments round table:

- ATLAS - (Graeme) - SARA: FTS lost all channel configs last night - dutch cloud affected (ticket 52983). LFC heavily degraded: error message "file exists" (ticket 52950). Dutch cloud therefore out of analysis and production. Q: related to LFC DB move this week? Probably, maybe backup of DB was not properly restored including the Oracle sequence values needed by LFC. Ron: last Tue DB migration went according to plan, tables seemed fine. In contact with Akos. JPB: tables may be ok but LFC sequence may be out of sync. Check the value of this sequence for being greater than the max number from existing files.
- CMS reports - (Daniele) - Asian facility meeting this morning: status reports and response to question: how to protect CMS operations from problems like at ASGC in the future. Useful for planning current activities. FNAL transfer not finishing to T2 in France - closed. Still some T0 to FNAL transfer issues. More details on other T2 related issues as usual in the CMS twiki.
- ALICE -
- LHCb reports - (Roberto) - Problem with SARA timing out on turl requests was fixed (verified by LHCb). Castor LHCb instance has been upgraded to 2.1.8-15.

Sites / Services round table:

- Xavier/KIT - ntr
- Ron/NL-T1 - dcache issue with pin manager getting stuck - a new release feature used which was now disabled again. Format error in a file with VOMS configuration due to wrong documentation. Site is in contact with developers. NIKHEF: 5h internet was cut off due to a WAN switch being down (no backup link?). The problem was fixed after 5h. FTS: started with new DB on RAC. Investigating why VO can not be mapped.
- Jeremy/GridPP -ntr
- Alessandro/CNAF: ntr
- Fabio/IN2P3 - next Mon intervention to reconfig the CE: goal - prevent submission to SL4 for CMS ALICE, LHCb (most do only SL5 submission already). Still some problems for ATLAS on SL5 - will be done afterwards. Twiki at CERN seems slower than usual - any change? No know problems at

CERN.

- Jason/ASGC: con call for Castor DB recovery this morning - core castor services are restored. functional tape testing taking place now. SRM schema had to be recreated, even though the DB restore succeeded without problem (needs followup). Site will now recreate space tokens. JPB: did you open SRM - trials from CERN did not show working end-point? Jason: maybe an issue with endpoint publishing in info system. ASGC will follow up. JPB will suggest by email a list of test for the site to validate proper service functioning (SRM,gridftp, etc) before opening the service to users. These will include testing lcg-cp between ASGC and CERN in both direction and the same test between ASGC T1 and T2. Experiments will validate after ASGC has reported success.
- Gavin: reboot campaign to apply security fix: lxplus has been rebooted, CVS will be soon, security fix will be applied into installation templates. VO responsible should reboot VOboxes as soon as template change has been announced. Deadline for all systems to be fixed and rebooted Fri. lxbatch will be drained for reboot - should be done by Fri morning.
- MariaG/CERN: observed high load on BNL: need to throttle massive job submission. ATLAS will go for pilot jobs to control submission rate. BNL will move ConditionsDB via data guard to new HW next week on the 11th (4h outage).
- Gareth/RAL: networking problem of yesterday was tracked down to a wrong router configuration (resolved by 16:00 local time). Still issues with tape system which are being tracked down - CASTOR migrations still slower than expected. New problem on batch system - jobs landing on the wrong node - under investigation.

AOB:

Friday

Attendance: local(Graeme, MariaG, Jamie, Gang, Wei, Jean-Philippe, Andrea, Alessandro, Jan, Roberto, Dirk);remote(Michael/BNL, David/NL-T1, Jeremy/GridPP, Gareth/RAL, Jason/ASGC, Alessandro+Elisabetha/CNAF, Brian).

Experiments round table:

- ATLAS - Graeme - Dutch cloud back in data distribution and production since last night. Jean-Philippe's suspicion about dropped sequence during DB transfer confirmed. Now only two T2s down for other issues. ASGC reported Castor service back yesterday evening. ATLAS put site in test mode. started ok, but now SRM probs (ticket 53060). T0 SRM: high load on DB backend - large burst of transfers, but no visible failures for ATLAS. ATLAS central services being reboot today - already half done.
- CMS reports - Andrea: all CMS VOboxes patched and rebooted.
- ALICE - (Patricia previous to the meeting). Yesterday night all ALICE VOBOXES at CERN were rebooted following the indications of Gavin in order to apply a kernel upgrade. The status of the kernel of the voboxes were reported at the end of the operation. Three VOBOXES were asked to be deprecated: voalice01, voalice02 and voalice03. These nodes have already replacement (kernel upgrade operation was also applied to the new nodes). After the kernel upgrade, voalice13 was put back in production with no mayor incidents to report
- LHCb reports -(Roberto) - all VObox patched/rebooted for LHCB - no issues. Started massive stripping activity against minimum bias sample from MC09 (last summer). Good test for dcache which has shown problems in similar activities before. STORM file listing issue (timeouts) at CNAF

(ggus ticket exists). Many pilots aborting at SARA - being investigated.

Sites / Services round table:

- Michael: concerning imminent sec problems: Do we need a certification process before sites start massive upgrade / reboot campaigns? Additional discussion on level of vulnerability for sites which have SELinux disabled (unlike CERN). BNL tried exploit on their system with a test program and was not able to use the published exploit. David: There may be other vulnerabilities than the SELinux one. BNL requested ATLAS opinion on patch/reboot at BNL? Graeme: Site is free to protect itself and resulting job failures have to be tolerated by the experiment. Micheal: will go ahead with patch application but no time to drain queues. Graeme: ok. Jamie: Will get back to security team for certification message on proposed patches. Intervention: migrate ORACLE (3D) to new hw - next Wed ok? Maria: fine for us. Another issue: running outdated solaris on storage server and will need update some 60 machines. Proposal: start Mon with batches of 6 machines at a time and do two batches per day (one week for all). Will not declare general downtime as only small part of the setup is affected at a time, still ATLAS will see transfer and access failures. Graeme: go ahead - site at risk is ok for ATLAS. Another issue: major network upgrade on Tue - after FTS was rebooted and due to automated update procedure at BNL FTS was upgraded to V2.2. Initial compatibility issue with the existing FTS DB were fixed (Thanks to Akos). (Andrea: similar upgrade to 2.2 also happened at FNAL and FZK). Since the upgrade BNL saw issues with a few sites: eg with RAL new error (file locality unavailable - source error during transfer prep). Also some minor Issue with Bestman sites (needs small config change for some new properties). Question from BNL: Should site stay with FTS 2.2 (Andrea: FNAL went back, FZK stayed). Alessandro: ATLAS uses FTS 2.2 patched (not released to production) at CERN to get checksum support. Jan: you may see problems with CASTOR sites - suggested to move to patched version (or watch for errors very carefully). JPB: some issued of 2.2 may not yet be solved by patch. Michael: BNL will stay on 2.2 and use patch - suggest that ATLAS validate this plan and get back to BNL. Alessandron: will respond by next Thu.
- David/NL-T1: confirm sequence number issue after DB move (was done by DB export/import - not from backup). Currently some SRM trouble with hostnames. Applying rolling kernel patch - should be done by end of today.
- Gareth/RAL: tape migration problems now understood and backlog gone. Batch system problem (job on wrong nodes) still under investigation. Some disk server (from MCDISK) shows hw probs - will replace faulty server RAM and then use checksums to confirm data integrity.
- Jason/ASGC: ATLAS transfers restarted after crash of stager. Reverted SRM config in one head node. Transfers resumed. ATLAS could not yet confirm successful transfers at the meeting but will follow progress closely.
- Additional comments from Jason received after the meeting by email: start from 11 UTC, we have notice the degradation of srm transfer errors, ticket Id 53060 created by shifter and was closed around 13 UTC. it's lately aware that the stager crash earlier and the srm2 cfg seems being revert earlier at one of the nodes which favor CERN settings that actually cause intermittent transfer failures earlier. except for the srm contact error at FTS, simple srmPing also being stuck forever and eventually fail with srm timeout error. have fixed it now. another minor issue observed after Atlas start FT at ASGC is that one of the disk server showing i/o error and have been resolved by reloading the kernel module of fc channel. could be the f/w upgrade perform two days ago affect the access to b/e raid subsystem. after SRM service recover, we resume the tape functional testing and confirm the tape system mechanical problem cause the scsi access become abnormal. the problem have been recovered in an hour and tape service fully operate after.
- Jan/CERN: Just noticed problem with castor LHCb (severe: scheduling error in logs) investigating. MyProxy move next Tue (fine for CMS?). Suggest to replace current HA setup by single box.

AOB:

- The IN2P3 team has produced a Service Incident Report concerning the recent cooling problems. Uploaded to the SIR table on the twiki.

-- DirkDuellmann - 02-Nov-2009

This topic: LCG > WLCGDailyMeetingsWeek091102

Topic revision: r12 - 2009-11-09 - DirkDuellmann



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback