

Table of Contents

Week of 101011.....	1
Daily WLCG Operations Call details.....	1
WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	3
Wednesday.....	6
Thursday.....	8
Friday.....	11

Week of 101011

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](#)

WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments

VO Summaries of Site Usability				SIRs, Open Issues & Broadcasts			Change assessments
ALICE	ATLAS	CMS	LHCb	WLCG Service Incident Reports	WLCG Service Open Issues	Broadcast archive	CASTOR Change Assessments

General Information

General Information				GGUS Information	LHC Machine Information
CERN IT status board	M/W PPSCoordinationWorkLog	WLCG Baseline Versions	WLCG Blogs	GgusInformation	Sharepoint site - Cooldown Status - News

Monday:

Attendance: local(Eddie, Roberto, Yuri, Stephen, MariaDZ, Jamie, Maria, Alessandro, Jan, Steve, Carlos, Patricia, Lola, Zbyszek, Simone, Harry);remote(Michael, Jon, Gareth, Federico, Rob, Christian Sottrup, Rolf, Ron, Alessandro Cavalli, Dimitri).

Experiments round table:

- ATLAS reports -
 - ◆ Ongoing issues
 - ◇ CNAF-BNL network problem (slow transfers) [GGUS:61440](#).
 - ◆ Oct 9-11 (Sat,Sun,Mon)
 - ◇ Express stream reprocessing was started at ~15:30 CEST Oct 8. Some problems with SW release were found around 20:00 CEST and fixed after 01:00am CEST on Oct 9. The reconstruction step is done. Merging step is in progress. Data replication to CERN and T2s will be started today.
 - ◇ CERN-PROD problem accessing data from ATLASHOTDISK pool in CASTOR via xrdcp on Oct.10. [GGUS:62917](#) solved. [Jan - started as TEAM ticket but nearly ALARM. Newly setup service class not working as expected - replication on first access. Didn't think xrdcp would be used which doesn't go through stager. Ignacio triggered replication for all files in this service class. Another problem - load not taken into account (separate protocol) so likely to always fall on the same server.

Most likely requires a code change. Simone - if there is a code change would like to know. In one week a reprocessing. Can be alleviated by using WN-based caching.]

- ◇ BNL MCDISK file transfer failures 10/09: connection to SRM failed. GGUS:62905 [↗](#)
solved: restart of pnfs and srm services helped.
- ◇ TAIWAN-LCG2 job failures with python load related errors. GGUS:62842 [↗](#)
verified: timeout while sourcing python in shared system fixed by adjusting nfs configuration.
- ◇ IN2P3-CC still some file transfer failures with source error: locality is unavailable on Oct.9. GGUS:62782 [↗](#) in progress updated. Under investigation.
- ◇ issues with several T2s, some of them are blacklisted.

- CMS reports -

- ◆ Experiment activity
 - ◇ Data taking (hopefully)
- ◆ CERN
 - ◇ GGUS:62696 [↗](#) occasional errors when opening files (via xrootd) on the Tier-0 processing pools. Became evident for jobs opening a large number files. For valid service classes it is reporting them as invalid. In progress.
- ◆ Tier1 issues
 - ◇ GGUS:62807 [↗](#) SAM Tests were failing at CNAF since last Wednesday evening, successful since this morning.
 - ◇ Savannah:117216 [↗](#), problems with rereco at T1_TW_ASGC, five files missing. [Gang - looking at stager logs to find if any information on these 5 files.]
- ◆ Tier2 Issues
- ◆ MC production
 - ◇ ongoing
- ◆ AOB

- ALICE reports [↗](#) - GENERAL INFORMATION: Today ALICE is setting up 3 new production cycles. The job profile will be unstable until mid-afternoon.

- ◆ T0 site
 - ◇ Operations required during the weekend because the local PackMan service was not performing as expected. Production stopped during several hours on Sunday morning. System back in production at noon
- ◆ T1 sites
 - ◇ GGUS: 62958. SARA. CREAM-CE seems to be overloaded. (Job status not changing). Site admon contacted via GGUS, experts looking at the service
 - ◇ RAL: Startup of all services this morning.
- ◆ T2 sites
 - ◇ This morning Subatech prevented us about a general failure of all VOBOXES (proxy renewal test in SAM). Problem solved at around 10AM

- LHCb reports - Mostly User jobs. Reconstruction jobs delayed due to a bug in the new release.

- ◆ T0 - none
- ◆ T1 site issues:
 - ◇ SARA: 2 separate issues with ConditionDB: [Zbyszek - apply process aborted one hour ago. Looks like problem with streaming with multi-version dictionary. Not sure if related with issue with SAM tests.]
 - systematically failing SAM tests and affecting production jobs. (GGUS:62896 [↗](#))
 - Database TAG not updated.
 - ◇ RAL: SAM tests for SRM (both DIRAC and gfal unit tests) are failing. Under investigation. (GGUS:62893 [↗](#))
 - ◇ IN2P3 : seg fault (GGUS:62732 [↗](#)) understood: due to newer WN with buggy kernel

◇ RAL: found many files reporting 0 checksum, also not being accessible via ROOT
(GGUS:61532 [↗](#))

Sites / Services round table:

- BNL - comment to report given by Yuri: have observed a couple of issues with data transfer failures at very low rate. Over w/e have gained more experience with cause. Namespace DB has grown very large - postgres now in process of reclaiming old transaction IDs. Wasn't seen before. Requires a tablescan for this. Performed in sections which typically run for 2-4 hours! Some transfer errors result - might see 5% reduction in efficiency. Found ways to reduce impact to tolerable level. Will further reduce effect until postgres has reclaimed sufficient IDs. Will upgrade h/w tomorrow 10-14 Eastern.
- FNAL - ntr
- RAL - LHCb questions - seeing 2 problems; believe they have the same cause - some files have 0 size in castor DB. In both cases timeouts between various castor processes. Not at bottom of it yet. Have CERN castor team looking at it. Have fixup for 0 length files and can / will correct these.
- NDGF - short downtime tomorrow for dCache for 1 hour - all details in GOCDDB
- IN2P3 - ntr
- NL-T1 - CREAM CE issue - guys still working on it. Nothing new to report. LHCb issue - some questions posed to LHCb in ticket - please respond! LHCb streams - maybe caused by config change at CERN? 2nd question - which IP involved? Try creamce2.xxx - maybe this works better.
- CNAF - post-mortem report on problem with CMS almost ready. Sent tomorrow. Problem now solved. Question regarding ATLAS StoRM endpoint. This is out of downtime since a few days but marked as grey in GridView. Why? Ale from ATLAS - don't see the grey. Are you checking OPS or ATLAS tests. Last week you declared 18h downtime for storage. Drained q for CE. For SE were in scheduled downtime so SE green. CE was not in downtime so tests failing - queues closed. In that case you will be accounted as it down.
- ASGC - ntr
- KIT - ntr
- OSG - reminder: BDII maintenance tomorrow morning (ops call time). Expect ~15 minutes degradation.

Services round table:

- network services: ntr
- grid services: ntr
- storage services: castor public upgrade tomorrow (should be transparent).
- databases: ntr

AOB: (MariaDZ) NDGF please reply to the ALARMS' handling at each Tier1 survey as per <https://savannah.cern.ch/support/?116430> [↗](#) . Also NDGF please check GGUS:62074 [↗](#) and re-assign if not for you. Ticket waiting since mid September. NDGF: The site will discuss Alarm handling this Friday.

Tuesday:

Attendance: local(Alessandro, Edward, Harry, Luca C., Luca D., Maarten, Massimo, Patricia, Roberto, Simone, Steve, Yuri);remote(Alessandro C., Christian, Dimitri, Federico, Gang, Gareth, Jeff, Jeremy, Jon, Marie-Christine, Michael, Rob).

Experiments round table:

- ATLAS reports -
 - ◆ Ongoing issues
 - ◇ CNAF-BNL network problem (slow transfers) GGUS:61440 [↗](#) (any update?)
 - ◆ ATLAS activities

Tuesday:

- ◊ Express stream reprocessing
 - reconstruction step completed Oct.11
 - merging step almost completed Oct.11
 - replication of the output Datasets to Tiers started.
- ◆ T1s
 - ◊ IN2P3-CC to T1s file transfer failures:source errors "asynch.wait","locality is unavalable". GGUS:62782 [↗](#).
 - Oct.11 21:19 site reported that some work is being done on the dcache nodes (the RAM memory is changed).
 - On Oct.12 10:30 this maintenance was over. Some other RAM memory changes in the coming days.
 - ◊ IN2P3-CC ->T1s,T2s still file transfer failures: gridftp_copy_wait: Connection timed out GGUS:62907 [↗](#) (Oct.9) updated
 - ◊ Taiwan issues with slow progress of reprocessing and group production. Reprocessing tasks moved to DE. Savannah:117269 [↗](#).
 - Cloud set to brokeroff at 00:10 Oct.12.
 - Site reply: some stale files left in home directory of production account. All garbage deleted. Nagios check implemented at 00:59.
 - Set to test mode at 10:33.
 - Test jobs finished successfully at 11:33.
 - Back in production at 11:40.
 - ◊ INFN-T1_DATATAPE some file transfer failures due to SRMV2STAGER timeouts. GGUS:62973 [↗](#)
 - filed at 20:18 Oct.11
 - solved at 22:22: the recall process was very slow due to a backlog in the system.Fixed.
 - ◊ NIKHEF-ELPROD_PRODDISK to SARA-MATRIX_MCDISK some transfer failures with globus connection timeout. GGUS:62975 [↗](#)
 - filed at 03:56
 - site restarted disk pools at 08:42, but still some failures reported at 09:49,11:50.
- ◆ added after meeting: production dashboard had to be restarted, cause being investigated
- CMS reports -
 - ◆ Experiment activity
 - ◊ Nothing to report; waiting for the beam
 - ◆ CERN
 - ◊ GGUS:62696 [↗](#) occasional errors when opening files (via xrootd) on the Tier-0 processing pools. Became evident for jobs opening a large number files. For valid service classes it is reporting them as invalid. In progress.
 - Massimo: problem understood, recipe expected in a few days, see associated Savannah item
 - ◆ Tier1 issues
 - ◊ Savannah:117216 [↗](#), problems with rereco at T1_TW_ASGC, five files missing.
 - ◆ Tier2 Issues
 - ◆ MC production
 - ◊ ongoing
- ALICE reports [↗](#) -
 - ◆ GENERAL INFORMATION: Still unstable production due to the changes on the production cycles.
 - ◆ T0 site
 - ◊ No issues to report
 - ◆ T1 sites

- ◇ GGUS:62958 [↗](#). SARA CREAM-CE. Still the same problem reported yesterday. Stopping the production at that site due to the huge number of scheduled jobs (at a certain moment the system might have recovered and all registered jobs have entered in the batch system).
 - CREAM's Registered status for jobs not correctly handled yet by AliEn, fix expected soon
- ◆ T2 sites
 - ◇ Usual operational procedures
- LHCb reports -
 - ◆ Experiment activities: During last night took 2pb-1 data but some datasets were not sent to offline. Production validation did not start yet due to a bug in Davinci option files. Expected to start in the afternoon.
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 0
 - ◇ T2: 0
 - ◆ Issues at the sites and services
 - ◇ T0
 - none
 - ◇ T1 site issues:
 - SARA: Request to open port 1521 of the Oracle ConditionDB listener to subnets of the WN at all T1's as per VO Id Card request (GGUS:62971 [↗](#)). For the time being the SARA DB is banned. All the jobs at SARA/NIKHEF will use the CERN DB.
 - RAL: found many files reporting 0 checksum, also not being accessible via ROOT (GGUS:61532 [↗](#)). Problem will be fixed in the afternoon. LHCb need to understand the root cause of this inconsistency on size and locality. In the mean time, in order to re-enable RAL in the production mask a validation stress test is being setup with Hammer Cloud.
 - RAL: SAM tests for SRM (both DIRAC and gfal unit tests) failing (GGUS:62893 [↗](#)). It was an overload issue due to a concurrent backup of DB.
 - IN2P3 : Still issues with shared area (GGUS:59880 [↗](#) and GGUS:62800 [↗](#))
 - ◆ discussion on need for port 1521 to be opened to all LHCb T1 sites
 - ◇ Luca D.
 - can this be limited to T1 subnets only?
 - how will the list be maintained?
 - put services on OPN?
 - ◇ Luca C.
 - there was some problem with using the OPN
 - port should be changed from default 1521 --> obstruct port scanners
 - discussion will be held at 3D workshop Nov 16-17

Sites / Services round table:

- ASGC - ntr
- BNL
 - ◆ reminder of scheduled dCache maintenance 10-14 EDT
- CNAF
 - ◆ problem with GridFTP servers for ATLAS: since Tue morning the 10-Gbit cards for all nodes fail intermittently, the problem may be with the switch
- FNAL - LHCOPN secondary circuit has been down since Oct 7 - issue with ciena module at USLHCNet - total bandwidth CERN-FNAL is limited to 8 Gbps. Being investigated.
 - ◆ Maarten followed up with OPN expert Edoardo

- ◇ HW problem in the network of one carrier providing links across the ocean
- ◇ they should receive the spare part very soon (Thu)
- ◇ it has been agreed that it is up to the Tier1s to handle issues with their links to CERN, not to the CERN network team: LHCOPN Operational Model Responsibilities

- GridPP - ntr
- KIT - ntr
- NDGF
 - ◆ dCache upgrade went OK
- NLT1
 - ◆ acknowledged CREAM problem at SARA reported by ALICE
 - ◆ 1 dCache pool node at SARA crashed with a HW problem
 - ◆ ATLAS report: NIKHEF GridFTP server had not restarted, please try again
- OSG
 - ◆ BDII maintenance window will be used for machine move
- RAL
 - ◆ HammerCloud tests by LHCb welcome
 - ◆ reminder of scheduled LFC/FTS/3D DB outage on Thu
- Central services
 - ◆ grid services - ntr
 - ◆ CASTOR
 - ◇ castor-public upgraded transparently to 2.1.9-9
 - ◇ that version is foreseen to be used in the HI run
 - ◇ upgrades for ALICE and CMS being prepared
 - ◇ upgrade window from Mon Oct 18 through Wed Oct 20, service at risk
 - ◇ try to take advantage of period without beam, when possible
 - ◇ Tue Oct 19 is OK for ALICE, CMS will discuss and follow up
 - ◆ dashboards
 - ◇ reliability reports of yesterday being recalculated for ALICE, ATLAS and LHCb after bug was fixed
 - ◆ databases
 - ◇ apply process for LHCb at SARA failed yesterday, being investigated
 - ◇ apply process for ATLAS at TRIUMF failed today, being investigated

AOB:

Wednesday

Attendance: local(Alessandro, Dawid, Edward, Jan, Maarten, Maria D, Patricia, Roberto, Steve, Yuri);remote(Alessandro C, Christian, Federico, John, Jon, Marie-Christine, Michael, Onno, Rob, Xavier).

Experiments round table:

- ATLAS reports -
 - ◆ Ongoing issues
 - ◇ CNAF-BNL network problem (slow transfers) GGUS:61440 [↗](#) (updated Tue Oct 12)
 - ◆ T0
 - ◇ CERN ATLAS proddashboard monitoring problem (job view). BUG:73904 [↗](#) filed yesterday. Stopped updating again at around 2:00 UTC of 13th.
 - Dawid: high load observed on DB, various indexes have been added now, speeding up some of the queries by orders of magnitude; please contact the DB team about such issues
 - ◆ T1s

- ◇ INFN-T1 gridftp problem seems solved yesterday evening. GGUS:63030. A number of FT failures with source errors coming from MILANO and LYON. MILANO has announced an unscheduled downtime and was excluded from data transfer.
- ◇ IN2P3-CC still file transfer problems. GGUS:62907, GGUS:62895.
- ◇ TAIWAN-LCG2 file transfer failures with destination SRM_ABORTED error. GGUS:63010 in progress, only one disk server for atlas localgroupdisk token.
- ◇ RAL issue on Atlas Castor instance (10:30 UTC). A corrupt index in the database behind the Atlas Castor stager. An outage for srm-atlas in the GOC DB for ~90min.
- CMS reports -
 - ◆ Experiment activity
 - ◇ technical tests and repairs; waiting for the beam later tonight
 - ◆ CERN
 - ◇ GGUS:62696 occasional errors when opening files (via xrootd) on the Tier-0 processing pools. Not so many errors recently. We are monitoring what happens when increasing the number of files again, hoping it remains stable.
 - ◆ Tier1 issues
 - ◇ SAV:117216, problems with rereco at T1_TW_ASGC, five files missing has been closed by ASGC. We are still investigating what happened to those files.
 - ◆ Tier2 Issues
 - ◇ Nothing to report
 - ◆ MC production
 - ◇ ongoing
 - ◆ AOB
 - ◇ The calendar for site downtimes produced by extracting the information from CIC has troubles. This affects SSB as well in that scheduled downtimes do not appear. Relation between CIC and GOCDB?
 - GGUS tickets can be opened against CIC Portal or GOCDB as needed
 - Alessandro D: problem has been seen to occur when a downtime is modified, leading to a loss of synchronization between GOCDB and CIC Portal
- ALICE reports -
 - ◆ GENERAL INFORMATION: 3 Raw production cycles and 1 MC cycle running.
 - ◆ T0 site
 - ◇ voalice09 (voalice09_xrootd) triggered an alarm this morning. Reported to the AliEn experts
 - machine is out of warranty, migration under preparation
 - ◆ T1 sites
 - ◇ FZK: Instabilities found by SAM in the 2nd VOBOX of ALICE at the site (alice-kit.gridka.de). Also reported by the ALICE expert at that site. Problem not reproducible at this moment, but checking it this afternoon again
 - ◇ IN2P3: the publication of the available resources was wrong for one of the local CREAM CE's. As result the number of queued JA's reached a value much larger than the 'normal' one (6000 instead of the usual 200). Problem associated to the resource BDII. Production stopped at the site
 - IN2P3 site contact for ALICE is in the loop
 - ◆ T2 sites
 - ◇ No remarkable issues to report
- LHCb reports -
 - ◆ Experiment activities:
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0

- ◊ T1: 2
- ◊ T2: 0
- ◆ Issues at the sites and services
 - ◊ T0
 - none
 - ◊ T1 site issues:
 - IN2P3: Files unavailable (GGUS:63024): disk server down?
 - IN2P3: problems in accessing the LHCb_RDST and LHCb_RAW (GGUS:63008). Possibly related with the other ticket.
 - IN2P3: few jobs went to buggy kernel WNs: seg fault. Ticket was closed few days ago, reported internally.
 - RAL: share has been put back to its original value. HammerCloud stress test.
 - SARA: CE SAM test failure observed, but seems OK now

Sites / Services round table:

- BNL - ntr
- CNAF
 - ◆ ATLAS StoRM back-end had a high load due to StoRM bug that is under investigation
- FNAL - ntr
- KIT - ntr
- NDGF - ntr
- NLT1
 - ◆ broken SE node at SARA: new Infiniband card did not help, node will be taken out of the round-robin alias for gsidcap servers (GGUS:63036, solved)
- OSG
 - ◆ BDII maintenance went OK
- RAL
 - ◆ reminder of short outage of LFC/FTS/3D DB servers on Thu
 - ◆ problematic lcg-CE replaced with new CREAM CE
- CASTOR
 - ◆ hotfix applied for ATLAS_HOTDISK
 - ◆ fix will be standard in 2.1.9-10
 - ◆ 2.1.9-9 upgrades for ALICE and CMS HI run okayed for Mon, Tue or Wed, exact time will be decided during that window to fall in a period without beam
- grid services - ntr
- dashboards - ntr
- databases
 - ◆ ATLAS dashboard looks OK now (see ATLAS report)

AOB:

Thursday

Attendance: local(Dawid, Edward, Jan, Maarten, Patricia, Stephane, Steve, Yuri);remote(Alessandro C, Christian, Federico, Foued, Gang, John, Jon, Marie-Christine, Michael, Rob, Rolf, Ronald).

Experiments round table:

- ATLAS reports -
 - ◆ Ongoing issues
 - ◊ CNAF-BNL network problem (slow transfers) GGUS:61440 (updated Oct 13).
 - ◆ T0/CERN

- ◇ ATLAS proddashboard monitoring problem (job view). BUG:73904 [↗](#) update. Still stops updating periodically: last one - Oct.13 ~23:00.
 - Edward: DB looks OK now, but the application itself seems to have some problem; production dashboard is not developed/maintained by CERN-IT but by an ATLAS group (Wuppertal?)
- ◇ ATLAS database dashboard page (ATLR): the Tier-0 Interlock folders table was unavailable at ~midnight. Understood: the change in password fixed at ~7:30 by the ATLAS dba. No bug report, just eLog record.
- ◆ T1s
 - ◇ INFN-T1 transfer failures seems solved. GGUS:63030 [↗](#): Storm bug seems fixed(?), transfer is OK this morning.
 - ◇ IN2P3-CC still file transfer problems. GGUS:62907 [↗](#), GGUS:62895 [↗](#).
 - Stephane: 133k files stuck on 2 pools, migration elsewhere will be tried, else they will be declared lost tomorrow; 3 out of 10 pools removed from production for diagnosis by sysadmins
 - ◇ PIC transfer failures Oct.13 (17:30-18:40) resolved. GGUS:63093 [↗](#), one pnfs server went down, restarted.
 - ◇ TAIWAN-LCG2 file transfer failures due to missing source file/incomplete replica. BUG:73967 [↗](#) solved, the new subscription helped.
 - ◇ RAL scheduled downtime (FTS,LFC,3D services). The cloud was set offline temporarily.
- CMS reports -
 - ◆ Experiment activity
 - ◇ Data taking; ready for high lumi sections
 - ◆ CERN
 - ◇ GGUS:62696 [↗](#) occasional errors when opening files (via xrootd) on the Tier-0 processing pools: on hold
 - ◇ High rate of jobs aborted with Maradona error on CERN CEs have appeared again at CERN Tier1; related to GGUS:61706 [↗](#) --> waiting for Platform fix
 - ◆ Tier1 issues
 - ◇ Files lost still being investigated by DataOps
 - ◆ Tier2 Issues
 - ◇ RAID disk problem at T2-US-UCSD affecting CRAB --> being fixed
 - ◆ MC production
 - ◇ large production ongoing
 - ◆ AOB
 - ◇ Nothing to report
- ALICE reports [↗](#) -
 - ◆ T0 site
 - ◇ ntr
 - ◆ T1 sites
 - ◇ GGUS:62958 [↗](#). SARA CREAM-CE. Site admins solved the issue. ALICE tests will be triggered this afternoon to validate the solution
 - ◆ T2 sites
 - ◇ French T2 sites claiming about the large memory consumption of some analysis jobs at their sites
 - memory limits will be implemented for the next AliEn release (confirmed later in ALICE Task Force meeting)
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ Data taking.

- ◊ Mostly user jobs in the morning, reconstruction of new data to start.
- ◆ New GGUS (or RT) tickets:
 - ◊ T0: 0
 - ◊ T1: 0
 - ◊ T2: 0
- ◆ Issues at the sites and services
 - ◊ T0
 - none
 - ◊ T1 site issues:
 - IN2P3: disk server still down at IN2P3. Affecting merge and user jobs.

Sites / Services round table:

- ASGC - ntr
- BNL - ntr
- CNAF
 - ◆ CNAF-BNL network issue: experts will be available for live analysis of future data replication problems; current investigations ongoing; link will be doubled in ~few months
 - ◆ StoRM bug was not fixed, problem was cured by service restart
 - ◆ CREAM ce07 HW problem solved
 - ◆ question: CREAM tests are not critical for site availability? answer: indeed, this will change per experiment when it declares the corresponding SAM test critical, probably first by ALICE (work in progress)
- FNAL - ntr
- IN2P3
 - ◆ issues acknowledged, work in progress
- KIT - ntr
- NDGF - ntr
- NLT1
 - ◆ 1 disk enclosure at SARA to be replaced, data being migrated, should be finished tomorrow, service will be a bit slower in the meantime
 - ◆ downtime could not be declared in the new GOCDB due to authorization issue (GGUS:63111 [↗](#))
- OSG - ntr
- RAL
 - ◆ DB intervention went OK
 - ◆ 1 disk server for ATLAS_MCDISK crashed, fsck ongoing, more news tomorrow
- databases - ntr
- grid services - ntr
- CASTOR
 - ◆ 1.2 PB being added for CMS
 - ◆ CMS ticket (GGUS:62696 [↗](#)) on hold: waiting for new errors
 - ◆ t0alice shows high activity preventing SLS tests from succeeding --> service appears red in SLS
 - ◊ fixed during ALICE Task Force meeting, cause was a bad firewall configuration on voalice10
- dashboards - ntr

AOB:

Friday

Attendance: local(Edward, Jan, Luca, Maarten, Patricia, Steve, Yuri);remote(Barbara, Federico, Gonzalo, Jeremy, Jon, Kyle, Marie-Christine, Michael, Onno, Rolf, Tiju, Tore, Xavier).

Experiments round table:

- ATLAS reports -
 - ◆ Ongoing issues
 - ◇ CNAF-BNL network problem (slow transfers) GGUS:61440 [↗](#), GGUS:63134 [↗](#).
Failures due to transfer timeouts for the large files (>2GB, often ~4GB).
 - ◆ T0/CERN
 - ◇ ATLAS proddashboard monitoring problem. BUG:73904 [↗](#) in progress, updated.
 - ◇ GOCDDB downtime summary not available. GGUS:63144 [↗](#).
 - annoying, but not a major problem; feature has been requested
 - ◆ T1s
 - ◇ PIC transfer failures Oct.14 (~3:30pm) resolved. GGUS:63118 [↗](#), dCache/pnfs issue fixed.
 - ◇ RAL back to production. Some job failures due to stage-in errors. GGUS:63124 [↗](#).
Files are on a mix of two servers, one of which is now back in production, the other server is disabled.
 - ◇ Taiwan-LCG2 job failures: libpython2.6 access. BUG:74033 [↗](#), NFS issue - new NFS installation is ongoing to fix.
 - ◇ IN2P3-CC. GGUS:62907 [↗](#), GGUS:62895 [↗](#), GGUS:63138 [↗](#). Experts are working on solving this migration/file access issue.
- CMS reports -
 - ◆ Experiment activity
 - ◇ Not much happening, various beam dumps
 - ◆ CERN
 - ◇ GGUS:62696 [↗](#) occasional errors when opening files (via xrootd) on the Tier-0 processing pools: on hold
 - ◇ Scheduled Intervention in the ITCORE database from 18:00 to 18:30 which caused SLS to be dropped for about half an hour. CSP shifter raised the alarm as CASTOR transfer seemed to have disappeared. Not immediate to connect with the IT Service Status Board announcement. Ref: GGUS:63135 [↗](#)
 - Jan: it was a team ticket, downgraded to less urgent; an alarm ticket should only be opened when CMS see errors in their own monitoring
 - experts agreed that IT Service Status Board message should have been clearer about which services would be affected
 - ◆ Tier1 issues
 - ◇ Reprocessing at ASGC still has problems accessing some files, Savannah:117301 [↗](#):
Unable to access two files (RAW)
 - ◆ Tier2 Issues
 - ◇ Nothing to report
 - ◆ MC production
 - ◇ large production ongoing
 - ◆ AOB
 - ◇ Nothing to report
- ALICE reports [↗](#) -
 - ◆ T0 site
 - ◇ xrootd-redirector issue reported yesterday at the end of the meeting: SOLVED. The machine (voalice10) was rebooted to apply the latest security patch. The problem was

the firewall, which started after the reboot. It is not compatible with CASTOR and the calls to the stager were filtered. Issue to be solved at the level of the voalice-xrootd quattor template to avoid this issue next time the machine is rebooted.

- Quattor fixes being investigated
- ◇ Production ongoing with no issues to report
- ◆ T1 sites
 - ◇ NIKHEF: Startup of all the services this morning
 - ◇ SARA: GGUS:62958 [?](#). SARA CREAM-CE. Confirmed by ALICE. The system is performing well after the operations applied by the site
- ◆ T2 sites
 - ◇ No remarkable issues to report
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ Data taking.
 - ◇ Mostly user jobs in the morning, reconstruction of new data to start.
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 1
 - ◇ T2: 0
 - ◆ Issues at the sites and services
 - ◇ T0
 - none
 - ◇ T1 site issues:
 - RAL: staging input files failing (GGUS:63140 [?](#)). RAL share reduced again to 0.
 - problems fixed now, RAL share has been put back to normal value
 - IN2P3: disk server still down at IN2P3. Affecting merge and user jobs. Requested list of files.

Sites / Services round table:

- ASGC - ntr
- BNL - ntr
- CNAF - ntr
- FNAL
 - ◆ yesterday CERN OPN experts kindly notified Jon of a problem they observed with the FNAL-IN2P3 traffic; issue is under investigation
- IN2P3
 - ◆ ATLAS tickets acknowledged, work in progress
 - ◆ LHCb ticket has been updated with list of files currently unavailable; waiting for HW intervention
- KIT - ntr
- NDGF - ntr
- NLT1
 - ◆ data migration off bad disk server will probably be completed during the night
 - ◆ 1 pool node crashed due to the high migration load; restarted, looks OK now
- OSG - ntr
- PIC - ntr
- RAL - ntr

- databases - ntr
- dashboards - ntr
- grid services - ntr

- CASTOR

- ◆ some ATLAS users are using castor-public to circumvent the ACLs put in place on castor-atlas; this route will be closed on Monday, helpdesk will be notified
- ◆ ATLAS robot certificate should have the same privileges as Kors Bos now, please verify

AOB:

-- JamieShiers - 11-Oct-2010

This topic: LCG > WLCGDailyMeetingsWeek101011

Topic revision: r22 - 2010-10-17 - MaartenLitmaath



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