

Table of Contents

Week of 100222	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.....	4
Thursday.....	6
Friday.....	8

Week of 100222

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 Availability				SIRs & Broadcasts		Change assessments
ALICE	ATLAS	CMS	LHCb	WLCG Service Incident Reports	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 Blogs	GgusInformation	Sharepoint site site - Cooldown Status Status - News News

Monday:

Attendance: local(Ulrich, Miguel, Jamie, Harry(chair), Nicolo, Tim, Jean-Philippe, Alessandro, Patricia, Toberto, Timur, Eva, MariaDZ);remote(Jon(FNAL), Angela(KIT), Rolf(IN2P3), Ron(NL-T1), Rob(OSG), BRIAN+Gareth(RAL), Michael(BNL),Gonzalo(PIC), Roger(NDGF), Gang(ASGC)).

Experiments round table:

- ATLAS reports - 1) eLog outage (~10h) Feb 20 morning. 2) Data distribution of reprocessed data has started with replication T1-T1, T1-CERN, T1-T2. 3) During the weekend 5 Tier-2s showed problems and were taken out from data distribution. All have GGUS tickets. 4) Feb 22. SARA is down. 5) Feb 22. RAL_MCDISK storage issues. 6) AFS Problems at CERN over the weekend caused false alarms that were caught by ATLAS shifters.
- CMS reports - 1) T1 Highlights: Some stuck transfers CNAF, RAL --> FNAL - investigating. CNAF probably due to their Phedex agent. 2) MC ongoing in all T2 regions. 3) SAM test errors at T2_US_UCSD, T2_ES_IFCA, T2_TW_Taiwan. 4) T2_IN_TIFR not visible in BDII - ROC re-running certification
- ALICE reports [ALICE](#) - 1) GENERAL INFORMATION: No new MC cycles during the weekend nor analysis trains. As it was decided during the last ALICE TF meeting the MC production expected before the real-data taking will be low with eventual early physics runs (not confirmed by the Physics groups yet). The weekend production basically consisted on the Pass1 reconstruction still ongoing at the T0 site (with no incidents to report) and pass 5 reconstruction executed at several T1 sites and

NIHAM. Currently some unscheduled user production jobs are also running. 2) T1 sites - CCIN2P3: The new "verylong" queue announced during the last week has been put in production on Friday morning and has also entered the ALICE production setup. This queue was put in production to cope with those jobs requiring a high memory for their execution. (Issue reported during the last week). 3) T2 sites - Hiroshima T2 asking for the testing of their CREAM-CE system. This testing will be performed today.

- LHCb reports - 1) System kept busy during the weekend by some MC productions and data stripping. Not very much extensive though. More production submitted this morning (currently 4K jobs). 2) Reported very low efficiency transferring to CASTOR at CERN on Friday evening due to many gridftp errors as shown in the following (see full report) plot on Friday evening. Preventively a mail sent to Jan informing that we were ramping up the activities and were observing on lemon an increased load on our lhcddata pool. The problem seems to have gone away by itself apart RAL-CERN where we have still failures at the source for the file availability (to be investigated). 3) The ticket (TEAM- urgent) submitted for the previous (CASTOR) issue is still in status progress (set by fnaz CERN ROC supporter this morning). Nobody touched it since its original submission on Friday at 20:37 (CT0000000662919) - discussion clarified that this ticket was routed as normal first to the CERN ROC reaching the internal castor support around 14.00. 4) RAL: most likely some issue with gds378 gridFTP server. 5) Shared area at IL-TAU-HEP, +g+ not found at BIFI and SAM tests failing at egee.fesb.hr

Sites / Services round table:

- IN2P3: Had an SRM problem on Saturday afternoon that was announced but not widely noticed. A server in a rack had a power supply problem that tripped switches in the same rack. Recovered about 18.00.
- NL-T1: 1) Bad SRM performance was improved by tuning of the backend database server. 2) One pool node went down and had to be restarted.
- RAL: (Gareth) 1) Have declared a 1 hour CASTOR outage on Wednesday for a DB configuration - batch will be paused. 2) Had an ATLAS batch problem this morning due to contention on their software server. 3) Had space problems under the ATLASMCDISK space token. Had set 12 to drain to be replaced but their space stays as available so ATLAS filled the space token. 3 of the replacement servers were added this morning. (Brian) ATLAS were filling 100TB in 100 hours - we can remove the empty drained servers before adding the new ones or leave them in to be further emergency space.
- BNL: The Great Lakes ATLAS Tier2 has finally resolved its issues after a storage upgrade and went back into production last night so is being validated by functional tests.
- PIC: 1) Network firewall configuration changes were made last Friday to solve the old problem of intermittent timeouts in NDGF to PIC transfers thought to be due to some NDGF servers not being in the OPN. Hoping to close ticket 54806 tomorrow. 2) PIC is banned by LHCb - what is the status? Roberto said this was the known incompatibility between the root version used by LHCb and the dcache libraries (dcap/gsidcap) at PIC. Will be fixed very soon.
- ASGC: Have added 9 disk servers, 340 TB, into their standby pool for later allocation to ATLAS or CMS.
- CERN: CASTOR suffered the side effects of a double CERN AFS failure over the weekend since local (not grid) stage requests check access to a users AFS home directory. An afs server had a root file system corrupted and also a kerberos server (user authentication) went down and would not restart completely as it found a database inconsistency.

AOB: (MariaDZ) GGUS is automatically interfaced to the savannah CMS computing tracker. Nevertheless the site names used by CMS in savannah are different from GOCDB names and even different from <https://twiki.cern.ch/twiki/bin/view/LCG/TierOneContactDetails> for the Tier1s. As this 'private' naming scheme covers all Tiers, up to Tier3, in order for the savannah-GGUS interface not to break, we strongly encourage CMS to use the GOCDB names.

(MariaDZ) On the general feeling by the experiments that Friday evening TEAM tickets never get treated before the next Monday pm:

1. The LHCb GGUS TEAM ticket opened at 20:33 last Friday is:
https://gus.fzk.de/ws/ticket_info.php?ticket=55739 Please try to put clickable URIs in all experiment reports to reach conclusions faster.
2. One can see that e-group grid-cern-prod-admins@cern.ch was notified immediately, i.e. received email.
3. Farida (GGUS Support Unit: CERN_ROC) assigned the ticket in CERN PRMS at 8:59 this morning which is good. The PRMS category was probably wrong as Jan re-classified it at 13:05
4. As a conclusion I repeat TEAM tickets are not supposed to urge solutions during the weekend.
 Workflow reminder:
https://gus.fzk.de/pages/ggus-docs/PDF/1530_Graph_TEAM_Ticket_Process.pdf and more doc.
 from <https://gus.fzk.de/pages/docu.php#team>

Tuesday:

Attendance: local(Ulrich, Harry(chair), Nicolo, Simone, Eva, Nilo, Jean-Philippe, Lola, Andrea, Julia, Timur, Malik, MariaG, MariaDZ, Roberto, Alessandro, Jamie, Miguel);remote(Jon(FNAL), Michael(BNL), Gang(ASGC), Ronald(NL-T1), Angela(KIT), Jeremy(GridPP), Josep(CMS+PIC), Rob(OSG), Gareth(RAL), Rolf(IN2P3), Alexei(ATLAS)).

Experiments round table:

- ATLAS reports - 1) Reprocessed data replication is in progress. Data consolidation at Tier-1s is done apart from TAG_COMM. 2) Many (2.5k) errors from TRIUMF TAPE ([SRMV2STAGER] jobs failed with forced timeout after 43200 seconds) .
- CMS reports - 1) T0 Highlights: File to be retransferred to IN2P3 not staging because it's on an archived tape - Remedy #CT663457. 2) T1 Highlights: RAL - prompt skimming failures during the weekend, rfio_open errors. CNAF - memory issues with jobs, caused by excessive mem usage by Java (used by dcache SRM client). Fixed by setting a lower heap size in JAVA_OPTIONS env. CNAF --> Nebraska transfer failures with CRL errors from CNAF FTS yesterday (but today the FTS is down for upgrade - will check again after downtime) IN2P3 - SAM test failures for SRM issue (same as last week), now OK. KIT - SRM SAM test unstable since last night, timeout failures. 3) T2 highlights: MC ongoing in all T2 regions. SAM test errors at T2_US_Purdue - dCache pool issues. 4) Weekly Data Ops Plan: Tier-0: processing ongoing data taking (Cosmics atm, beam splash likely late this week or early next week). Tier-1: running backfill at all sites, while waiting for more serious requests. Tier-2: full-scale production of new MC requests. See full report for weekly Facops plans.
- ALICE reports - 1) GENERAL INFORMATION: Pass1 reconstruction still ongoing and in addition there is a significant number of user analysis jobs running in the system. 2) T0 site: Announcement of an intervention on CASTORALICE on Wednesday 24th February to apply an Oracle security patch to the backend Oracle database. This Oracle patching is transparent. It will start at 9:30am Geneva time. The intervention window is 1 hour. 3) T1 sites: CCIN2P3: Yesterday we announced the setup of a verylong queue to cope with very memory demanding jobs. The queue provided by the site was included into the Alice DB in the afternoon and the system has been already running analysis jobs with no incident to report. 4) T2 sites: GRIF-IPNO announced yesterday the setup of a CREAM-CE

available also for the Alice production. However, the information provided by the system in terms of available resources and waiting jobs needs still some tuning before testing the job submission. Issue reported to the site admin.

- LHCb reports - 1) 8-9 MC production requests to be verified prior to submission, others are running fairly smoothly at the good rate of 10K concurrently. Stripping jobs on Minimum bias events is tuning the retention factor, number of events, size of the output and CPU length in order to accommodate jobs in the available queues. 2) GGUS: trapped a short unavailability this morning (#55807). 3) T0 sites issues: LHCb will come up with a request to move their RDST space from T1D0 to T1D1 and reverse the share RDST<--> RAW. 4) T1 sites issues: RAL: most likely some issue with gds378 gridFTP server. 5) T2 sites issues: *g++ not found at INFN-TORINO and pilots not picking up payload at RRC-KI.

Sites / Services round table:

- ASGC: Have put 12 new tape drives online bringing the total to 22 with 8 more to come soon.
- BNL: Have scheduled the FTS upgrade to version 2.2.3 for today [note the upgrade was successfully completed within 2 hours.]
- KIT: 1) Some CMS timeouts from SAM tests probably due to slow transfers blocking some pools. 2) There are instabilities in their batch server and it has been restarted several times. Under investigation.
- RAL (Gareth): Still having problems (since Monday) with ATLAS batch jobs loading their software server. Michael reported BNL had experienced this and moved to a special (Bluearc Appliance) system. Gareth reported they are seeing more writes than reads in this NFS area which is the opposite of the usual behaviour. Simone asked to be told the owners of these jobs since normally only the SGM account would want to write into the software area.
- OSG: A few tickets open all about geographical coordinates in the information system bdi's being inconsistent. Concerns U.Florida and Nebraska while Boston has been resolved. We also saw the GGUS unavailability so need to check we have not lost any tickets.
- CERN CASTOR: Oracle security patches will be applied tomorrow to the CASTORALICE and CASTORLHCb backend databases. These should be transparent interventions.

AOB: (MariaDZ) There was a short GGUS unavailability this morning due to human error. Concretely, a file under development was uploaded by mistake to the production server instead of the training server. This erroneous situation lasted not longer than 5-10 minutes (roughly between 10:15-10:25am CET today) and affected only the ticket page (frontend), not any interface. Hence, no worries concerning OSG tickets etc. Related ticket https://gus.fzk.de/ws/ticket_info.php?ticket=55807 can't be updated because it was 'verified' by the submitter (Joel) which now makes it un-re-openable.

Following yesterday's AOB please note that there is a GGUS Support Unit (SU) going to castor.support. Full SU list on page https://gus.fzk.de/pages/resp_unit_info.php. Please join the next USAG 2010-03-17 where agenda <http://indico.cern.ch/conferenceDisplay.py?confId=85636> item 8 will address precisely such issues.

Wednesday

Attendance: local(Harry(chair), Miguel, Nicolo, Oliver, Ulrich, Steve, Eva, Nilo, Timur, Malik, Eduardo, Simone, Jamie, Patricia, Alessandro);remote(Xavier(KIT), Jon(FNAL), Joel(LHCb), Michael(BNL), Jeremy(GridPP), Gang(ASGC), Onno(NL-T1), Rob(OSG), John(RAL), Rolf(IN2P3)).

Experiments round table:

Wednesday

- ATLAS reports - 1) Since 12.00 transfers from Uni. Geneva to Nikhef are failing since FTS in Nikhef cannot map the two endpoints to a channel. A Ggus ticket has been sent. Onno later reported they had found the probable cause in the channel configuration where '*'-NIKHEF should be '*'-NIKHEF-ELPROD. 2) Germany deployed FTS 23 and ATLAS activated checksum verification for all the German cloud which is working for all sites except dpmc (?) which is running an old version of dpm that needs to be upgraded for checksum support. This will have to be done for every cloud of course.
- CMS reports - 1) T1 Highlights: CNAF - FTS server down for issues during upgrade to 2.2.3, rolled back. PIC - JobRobot failures, MARADONA errors. RAL - Savannah #112903 - rfio_open issues in skimming due to disk server with issues - files already migrated were recalled from tape, other files currently unavailable. 2) T2 highlights: MC ongoing in all T2 regions. UKI-SOUTHGRID-RALPP in downtime. IN2P3-CC-T2 and UKI-SOUTHGRID-BRIS-HEP back in MC production.
- ALICE reports [↗](#) - 1) GENERAL INFORMATION: The Pass1 reconstruction jobs (still going on) and some user analysis tasks represent the current activity of ALICE not exceeding the 1000 concurrent jobs for the moment. Effort concentrated currently on the next AliEn v2.18 version, which is expected to be deployed at all the sites at the end of March. 2) T0 site: CASTORALICE upgrade operation announced yesterday has concluded. However, the DAQ team is reporting several issues after the upgrade: AliEn registration failing (ALICE internal issue). ROOT transfers timeout (first occurrence at 12:17:49, last current occurrence at 13:23:39, the previous transfer was made with success at 03:18). One of such files is:
root://voalice08.cern.ch//castor/cern.ch/alice/raw/global/2010/02/24/06/10000110996026.40.root A test transfer to another path completed OK and apparently the file reported above eventually made it through (it looks OK via rfdir) so it could be an XROOTD-related problem. 3) T2 sites: Waiting for the green light of the site admins at GRIF_IPNO to begin the testing of the local CREAM-CE system which yesterday was suffering of some inconsistencies in the information system.
- LHCb reports - 1) Ramped up MC productions + chaotic user analysis. 2) Having troubles with Castor at CERN. Merging jobs got stalled at CERN because of no CPU being used. Typically on a merging job this means the connections have been lost to the disk server (lhcbdata). Happening since yesterday it does not seem to be correlated with the intervention today. 3) T2 sites issues: CBPF object not found in shared area. CY-01-KIMON shared area issue.

Sites / Services round table:

- KIT: Still having problems with the batch system but workaround is to set the 2-3 bad worker nodes per day offline losing some user jobs. A kernel bug is suspected but the problem is still under investigation.
- BNL: Have completed the upgrade to FTS 2.2.3 which took only 2 hours, completed smoothly and was carried out as a transparent intervention. Michael suggested to add to the installation documentation that a privilege is needed in Oracle to be able to activate the history tracking (fts_history.start_job needs the CREATE JOB privilege) and Oliver (IT/GT) agreed this would be done.
- RAL: 1) Scheduled outage this morning went well. 2) More ATLAS jobs have been allowed to run without overloading their NFS software shared area so this will be opened up further.
- CERN CASTORATLAS: In order to fix a recently identified problem where the xrootd access is failing, it is proposed to install a new version of the xrootd software.

The details of the change are at

<https://twiki.cern.ch/twiki/bin/view/CASTORService/ChangesCASTORATLASXrootSSL>.

The timescale proposed is 14:00 tomorrow (Thursday) and will lead to a small interruption of service (<1 second and only being seen by new connections).

If there are any concerns, please raise them before 17:00 today.

- CERN gLite AFS UI: New versions of the gLite AFS_UI have now been installed. Over the next two weeks they will be rolled to the default configurations. Tomorrow the *new_3.1* and *new_3.2* links will be updated to **3.1.44-0** and **3.2.6-0** respectively. In a couple of weeks time these will become the "current" links with a broadcast and announcement here a few days in advance of the actual change. Exact details of links status in */afs/cern.ch/project/gd/LCG-share*

Link	Target OS	Today Resolves To	Tomorrow Will Resolve To	After 2 Weeks Will Resolve To
sl4	SL4	3.1.38-0	3.1.38-0	3.1.44-0
current_3.1	SL4	3.1.38-0	3.1.38-0	3.1.44-0
new_3.1	SL4	3.1.38-0	3.1.44-0	3.1.44-0
previous_3.1	SL4	3.1.25-0	3.1.25-0	3.1.38-0
sl5	SL5	3.2.1-0	3.2.1-0	3.2.6-0
current_3.2	SL5	3.2.1-0	3.2.1-0	3.2.6-0
new_3.2	SL5	N/A	3.2.6-0	3.2.6-0
previous_3.2	SL5	N/A	N/A	3.2.1-0

Currently "current" points to "current_3.1" and "previous" points to "previous_3.1" . I would like to never point current and previous to a 3.2 releases since they are too vague. They would expire when eventually 3.1 itself expires.

- CERN CA rpm and lcg-vomscert update Upgrade of CA rpms to new version 1.34-1 Upgrade of lcg-vomscerts to new version 5.8.0-1 affecting Quattor managed machines which use these packages
- CERN GLEXEC_wn The next round of scheduled upgrades of quattor managed machines at CERN will contain a new release of the worker node software, as well as the GLEXEC_wn addon. This release is currently in pre-production for testing. The production release is foreseen for 2/3/2010

AOB:

- CERN FTS Support: Have had support issues coming from several sites trying the FTS upgrade. CNAF had to roll back after trouble with the schema change.
- CERN non-CASTOR Alarm Tickets: Simone reported it was not clear to ATLAS how CERN FTS and LFC GGUS alarm tickets would be handled so will make a test tomorrow just before this meeting. The issue arises since the data management piquet of DSS group only handles CASTOR issues while FTS and LFC services are run by PES group which has no piquet service. Perhaps it would be better to use team tickets for CERN FTS and LFC.

Thursday

Attendance: local(Nicolo, Harry(chair), Steve, Andrea, Lola, Luca, Alessandro, Guissepi, Jamie, Jean-Philippe, Timur, Miguel, Malik, Tim, Simone);remote(Jon(FNAL), Michael(FNAL), Ronald(NL-T1), Gang(ASGC), Angela(KIT), Roger(NDGF), Gareth(RAL), Jeremy(GridPP), Joel(LHCb), Gonzalo(PIC), Rolf(IN2P3), Rob(OSG)).

Experiments round table:

- ATLAS reports - 1) LFC registration errors at atlaslfc.nordugrid.org. LFC acls not correctly set for group09 directory, problem FIXED. 2) From/to BNL errors ggus 55936, problem noticed in the same time by BNL experts and ATLAS shifters. 3) TAIWAN power outage (~3am UTC). 4) Some T1-T1 slowness for some channels: NDGF- BNL, BNL-LYON, BNL-PIC. Under investigation but not thought to be BNL related. 5) LFC Lyon not accessible from outside (at least) , ggus 55931 . From ELOG <https://prod-grid-logger.cern.ch/elog/ATLAS+Computer+Operations+Logbook/9843> : LFC in Lyon likes voms.cern.ch but not lcg-voms.cern.ch.
- CMS reports - 1) GridMap still unavailable, following up in GGUS #55917. 2) T1 Highlights: MC pre-production running at FNAL & KIT. PIC & ASGC - JobRobot failures, MARADONA errors. 3) T2 highlights MC ongoing in all T2 regions, including NorduGrid. UKI-SOUTHGRID-RALPP back in MC production. T2_RU_ITEP SAM test failures on CE (CMS SW area not available), now green.
- ALICE reports [↗](#) - 1) GENERAL INFORMATION: As reported yesterday, the Pass1 reconstruction jobs (still going on) and some user analysis tasks represent the current activity of ALICE with around 1200 concurrent jobs running on the system. 2) Concerning the CASTOR issue reported yesterday, experts have been running several tests during the mentioned period with no issues found. For the moment the trace is enabled to track any issue if the transfer problem appears again and they are reported by the DAQ team. 3) The current configuration defined for Prague-T2 into the LDAP configuration of Alice is wrong and it is preventing jobs are ran at this site. Issues reported by the site admin, it will be fixed this afternoon. No actions from the site admin will be required.
- LHCb reports - 1) Drained the merging jobs at CERN (stalled issue), more MC jobs have been submitted yesterday to the grid (4K jobs running). 2) T2 sites issues: EFDA-JET illegal instruction. UKI-SOUTHGRID-BRIS-HEP jobs aborting. 3) Have broadcast to LHCb sites to install an rpm on their worker nodes (if not already done) that is needed for root to work.

Sites / Services round table:

- BNL: Currently having an SRM issue under investigation but expect to be back within 30 minutes.
- KIT: Batch system is stable again.
- PIC: Batch system was stuck last night for about 4 hours due to wrongly configured new worker nodes, Was solved at about 02.00.
- NDGF: Completed a scheduled downtime of 30 minutes. Some (old) issues with tests on a Finnish site related to port numbers.
- ASGC: 2 weeks ago we announced a scheduled downtime to parallelize the DC generator with the power source, after that ASGC should not suffer from scheduled power cut because we can move to the DC generator earlier. The vendor started to test this parallel system early today, seems the vendor made some wrong wiring on control line and didn't find it during the checking. This lead to power failure at ASGC at 3:00 UTC. Most services are already recovered and SIR will be delivered later. About 5 hours total downtime.
- CNAF (email report from Paolo Veronesi): We tried the upgrade to FTS 2.2.3 on 23rd Feb, but we experienced a couple of errors. Details at https://gus.fzk.de/ws/ticket_info.php?ticket=55828. We rolled back to FTS 2.1. The second error reported (ERROR: Some line does not contain the same number of group IDs and group names ! Please check ! Exiting.) has been understood and it was due to a typo in the file users.conf used by yaim. A test instance has been setup today to (re)try the upgrade, with the opportunity to investigate better any other related problem, with no impact in production. If we can upgrade the instance of testing, we will schedule a new upgrade for the production one.

- RAL (email report from Gilles Mathieu): APEL status - 25/02. After our recent outage APEL is currently experiencing a large backlog of record being published/republished. During this time publisher log files may show errors connecting to the APEL central system. We are aware of this and are taking steps to manage the backlog. Please continue to publish normally as we are implementing a flow control system to centrally manage the problem. An AT_RISK downtime has been scheduled for tonight. This will allow us to run the APEL server in a specific configuration during the night.
- RAL: Have now removed the limit of the number of ATLAS batch jobs as their software server is no longer overloaded. The root cause is not really understood.
- IN2P3: SIR for Feb 15 interruption of worker node network connectivity has been posted.
- CERN xroot: the agreed new version was put in at 14.00.
- CERN streams replication: CMS online to offline replica for conditions crashed last night at midnight and was restarted this morning at 9. The crash was caused by long transactions being run by CMS at the time, which is a known streams problem. We have increased memory for streams to avoid (mitigate) the occurrence of this in the future
- CERN LHC VOMS: lcg-voms.cern.ch service certificate update is now complete. Affects in particular those sites running gLite 3.1 services.
It is clear from looking at SAM test results that some sites including key sites are failing due to not upgrading to lcg-vomscerts-5.8.0-1.noarch.rpm.
- CERN AFSUI: As of today:
 - ◆ /afs/cern.ch/project/gd/LCG-share/new -> 3.1.44-0
 - ◆ /afs/cern.ch/project/gd/LCG-share/new_3.2 -> 3.2.6-0
- OSG: Volunteered to help BNL with their SRM issues if needed. OSG grid operations centre is under a production freeze until March 30th.

AOB:

Friday

Attendance: local(Harry(chair), Malik, Eva, Nilo, Alessandro, Timur, Patricia, Andrea, Nicolo);remote(Gonzalo(PIC), Xavier(KIT), Michael(BNL), Jon(FNAL), Gang(ASGC), Rolf(IN2P3), Tore(NDGF), Tiju(RAL), Alessandro(CNAF), Rob(OSG)).

Experiments round table:

- ATLAS reports - 1) TAIWAN recovering from power outage https://gus.fzk.de/ws/ticket_info.php?ticket=55912 - SRM agent restarted. 2) SARA-MATRIX problems in exporting data https://gus.fzk.de/ws/ticket_info.php?ticket=55960 . FIXED. 3) Problems during the night exporting calibration data to INFN-T1_DATATAPE and CALIBDISK endpoint: FIXED. It was due to a problem in one part of a configuration file of the ATLAS Site Services. That part was not properly read, and it has been spotted only after the upgrade of the SS (yesterday, went smoothly otherwise). 4) Disk server problem at RAL <https://prod-grid-logger.cern.ch/elog/ATLAS+Computer+Operations+Logbook/9879>.
- CMS reports - 1) T0 Highlights: Failed jobs on T0 caused by upload of incorrect conditions for Castor (the CMS subdetector). 2) T1 Highlights: FNAL failures in backfill jobs reported yesterday - caused by issue in new ProdAgent version under testing, now fixed going back to previous version. Temporarily enabled T2_CH_CAF-->T1_FR_CCIN2P3 transfers to enable recovery of some datasets

lost in deletion incident in November 2009. IN2P3 - Two skimming jobs stuck waiting long time to access input files. ASGC - Squids down after power cut, restarted. CNAF - Intermittent SAM test failures for StoRM instabilities. 3) T2 highlights: MC ongoing in all T2 regions. T2_TW_Taiwan SAM test failures after power cut but now green again. 4) Services: GridMap still unavailable, following up in GGUS GGUS #55917.

- ALICE reports [↗](#) - 1) GENERAL INFORMATION: Current production dominated by user analysis jobs with more than 4000 jobs currently running. 2) T0 site: Tiny amount (about 50) of jobs running at the CERN resources to finish the pass1 reconstruction which has been going on during the full week. 3) T2 sites: Still large memory consumption observed in several T2 sites. In particular Subatech (Nantes, France) has announced today that several jobs are being killed at the sites because of their large memory consumption. The problem was observed and solved at CCIN2P3 a few weeks ago. However the solution found at that moment cannot be exported to all sites (establishment of a very long queue for Alice use). We are collecting information from the rest of sites and also rates of failing jobs before establishing a strategy to follow before the necessary fixes in AliRoot.
- LHCb reports -

Sites / Services round table:

- PIC: 1) Upgrade of FTS to FTS 2.2.3 is now scheduled for the morning of 9 March. 2) Tuesday 30 March there will be an all-day annual electrical shutdown.
- FNAL: Upgrade to FTS 2.2.3 is planned for next Monday 1 March. There will be a 15-30 minute window of no transfers.
- ASGC: The parallelised power system will be retested on March 3 between 10:00 and 12:00 at Taiwan local time (2:00 - 4:00 UTC time), we are going to announce this in GOCDB soon.
- IN2P3: Confirmation of an outage 1-2 March for maintenance of batch and mass storage.
- RAL: Performed an intervention on an ATLAS disk server that had a memory fault and also on an LHCb disk server with a failed tray.
- CNAF: The CMS problem reported above was due to transient CRL errors on an SRM machine - now solved.

AOB:

-- JamieShiers - 19-Feb-2010

This topic: LCG > WLCGDailyMeetingsWeek100222

Topic revision: r33 - 2010-02-26 - HarryRenshall



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