

Table of Contents

| | |
|----------------------------------------------------------------------------------------|----------|
| Week of 111024..... | 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..... | 5 |
| Thursday..... | 7 |
| Friday..... | 9 |

Week of 111024

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(Alessandro, Alex, Iouri, Jacob, Jan, Luca, Maarten, Maria D, Maria G, Mattia, Raja);remote(Burt, Dimitri, Gareth, Giovanni, Gonzalo, Jhen-Wei, Michael, Onno, Rob, Ulf).

Experiments round table:

- ATLAS reports -
 - ◆ Physics
 - ◇ data taking
 - ◇ HLT reconstruction in progress.
 - ◆ T0/Central services
 - ◇ High load on ADCR (instance 1) over weekend
 - Luca: investigating the matter together with ATLAS DBAs and developers; the latest high load was related to large deletion activities and is gone since ~2 pm
 - ◆ T1 site
 - ◇ RAL: srm down, Saturday morning (3 am UTC). GGUS:75597. Converted to ALARM. The problem for ATLAS was fixed 30 min after the alarm ticket was sent. RAL had a problem with DB behind CASTOR, but ATLAS was not affected since 6:30 on Saturday until 4 am on Sunday, when all transfers started to fail. RAL declared DT until 13:00 on Monday and was off for ATLAS activities . The problem

was fixed at 20:45 on Sunday. Back in ATLAS activities. Warning DT until 13:00 today. GGUS is still open. Many thanks.

- Maria D: the cause of an incident is not always reflected in a ticket's solution
 - when further details are foreseen, it may be desirable not to verify the ticket yet, such that it can be updated later
- ◇ INFN-T1: LFC down, Saturday ~10 am UCT. ALARM GGUS:75601 [↗](#). The problem was fixed in less 1.5 h after submitting the ticket. Many thanks.
 - Giovanni: issue was due to memory leak, the LFC had been up for 1 year; cured by a restart
- ◇ IN2P3-CC : srm down, Sunday 7 am, GGUS:75609 [↗](#), converted to ALARM. The site was offline until the the problem was fixed at 18:00 (there are no details on what caused the problem in GGUS). GGUS is still open. Many thanks.
- ◇ INFN-T1: one stucked FTS job blocking new transfers. GGUS:75524 [↗](#). Now only one channel has a very long queue (queue has the high number of fts jobs submitted). The limit for the number of files that can be simultaneously transferred has been increased to 30. Solved
 - Giovanni: FTS problem not understood, but it is gone for now
 - Alessandro: we investigated the matter with the help of Stefano Antonelli at CNAF and wonder if the problem of long queues may have started after the T2D channels were added? If so, other T1 might be similarly affected. Could CNAF look into monitoring the queue lengths? The problematic files have been transferred in the meantime.
- ◇ Transfers PIC->GRIF-LAL GGUS:75429 [↗](#). We are still seeing the problem, GRIF-LAL has problem for transfers from other sites. There is no reply from GRIF-LAL.
- ◆ T2 sites
 - ◇ ntr
- ◆ Other business
 - ◇ priority for ALARM GGUS was changed from top to less urgent when ticket was updated.
 - should be fixed now: GGUS:75575 [↗](#)
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Data taking ongoing
 - ◆ CERN / central services
 - ◇ CMSR Oracle instance spontaneous reboot problem, GGUS:74993 [↗](#), kept open to follow up with increased logging information, still waiting for a solution
 - ◆ T0 and CAF:
 - ◇ cmst0 : very busy processing data !
 - ◆ T1 sites:
 - ◇ [T1_TW_ASGC]: general file access problem, GGUS:75377 [↗](#), possibly a software problem?
 - ◇ [T1_FR_CCIN2P3]: overload in dCache affecting transfers form/to IN2P3. GGUS:75397 [↗](#). Ongoing.
- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ Experiment activities
 - ◇ Reconstruction and stripping at CERN
 - ◇ Reprocessing at T1 sites and T2 sites
 - ◆ T0

◇ CERN : Running a lot more jobs now, but not fully clear if the fairshare system has been fixed

◆ T1 sites:

◇ IN2P3 : (GGUS:75610) : srm problem on 23 Oct. Fixed after alert from LHCb.

◇ RAL : SE in unscheduled downtime.

◆ T2 sites:

◇ Manchester : Pilots aborted (GGUS:75614)

Sites / Services round table:

- ASGC - nta
- BNL - ntr
- CNAF - nta
- FNAL - ntr
- KIT - ntr
- NDGF - ntr
- NLT1
 - ◆ ~1 hour ago 12 SARA pool nodes were restarted to pick up new host certificates
 - ◆ communication problem between dCache head node and pool nodes reported on Fri persists, dCache developers have been asked to look into it
- OSG
 - ◆ 3h normal maintenance window tomorrow, should essentially be transparent
- PIC - ntr
- RAL
 - ◆ during the weekend all 4 CASTOR instances were down due to problems with the 2 Oracle RAC setups: both suffered nodes crashing and not automatically rebooting; the cause lay in corruption of a disk array area used for backups; the situation was corrected and a restart of all nodes then restored the service; for now the FTS channels and the batch system are throttled, but those limitations are expected to be lifted later today
- CASTOR/EOS
 - ◆ eosatlas updated between 10 and 11 am; deletions then failed due to misconfiguration of added node; fixed, but looking further into unexpected errors that were observed by ATLAS (e.g. connection refused)
 - ◇ Iouri: ~10% of the transfers failed
- dashboards - ntr
- databases - nta
- grid services - ntr

AOB: (MariaDZ) Drills for 9 real GGUS ALARMS, for tomorrow's MB attached at the end of this twiki page.

Tuesday:

Attendance: local(Alessandro, Alex, Iouri, Jan, Luca, Maarten, Maria D, Mattia, Pepe, Raja);remote(Burt, Gareth, Giovanni, Gonzalo, Jhen-Wei, Michael, Rob, Rolf, Ronald, Ulf, Xavier).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central services
 - ◇ CERN-PROD_DATADISK transfer failures. GGUS:75632 verified: due to a scheduled replacement of the SRM server in CERN_PROD from 9-10 UTC. No errors since then, all failed transfers staged successfully meanwhile.

- ◊ CERN-PROD_DATADISK -> JINR 1 transfer is constantly failing with "Source file/user checksum mismatch". Savannah:88115 [↗](#), the file seems to be corrupt.
- ◆ T1 sites
 - ◊ Taiwan-LCG2 production job failures in the morning: host credential (dpm25) has expired. GGUS:75662 [↗](#). Host certificate has been replaced at ~8am, thanks!
 - ◊ Iouri: IN2P3 came back OK after their scheduled downtime
- ◆ T2 sites
 - ◊ ntr
- CMS reports -
 - ◆ LHC / CMS detector
 - ◊ Data taking ongoing. Till 16:30 set-up of injection.
 - ◆ CERN / central services
 - ◊ CMSR Oracle instance spontaneous reboot problem, GGUS:74993 [↗](#), kept open to follow up with increased logging information, still waiting for a solution
 - ◆ T0 and CAF:
 - ◊ cmst0 : very busy processing data !
 - ◆ T1 sites:
 - ◊ [T1_TW_ASGC]: general file access problem, GGUS:75377 [↗](#), possibly a software problem?
 - ◊ [T1_FR_CCIN2P3]: overload in dCache affecting transfers from/to IN2P3. GGUS:75397 [↗](#). Ongoing.
 - ◊ [T1_IT_CNAF]: Failing transfers from and to CNAF. GGUS:75675 [↗](#). There is a problem with the STORM storage backend. Experts having a look.
 - Giovanni: the problem is fixed, was due to a GPFS bug
 - ◆ Others:
 - ◊ Yesterday we spotted out a problem with downtimes tracing in the SSB. There is a bug which affected, at least CNAF, which was marked as being in Unscheduled Downtime for some days, which was not the case. Savannah:124240 [↗](#).
 - Mattia: the Dashboard team are looking into it
- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ Experiment activities
 - ◊ Reconstruction and stripping at CERN
 - ◊ Reprocessing at T1 sites and T2 sites
 - ◊ Starting to actively synchronise files on LHCb Tier-1 SEs with expectation (LFC, DIRAC catalogs)
 - ◆ T0
 - ◊ CERN : GGUS ticket submitted regarding fairshare (GGUS:75663 [↗](#)) as requested yesterday.
 - Alessandro: will the increase in the LHCb share affect other VOs?
 - Raja: possibly, but it will be reset tomorrow
 - ◆ T1 sites:
 - ◊ IN2P3 : "Scheduled" downtime to update Chimera server, but batch queues were not drained. There were also problems with LHCb jobs at IN2P3 even before the downtime officially started.
 - Rolf: LHCb had told us the announcement was sufficient for LHCb to prepare for the downtime; please open a ticket to discuss what LHCb expect to happen in such cases; AFAIK, nothing happened prior to the official start of the downtime, but if there was anything wrong, please open a ticket for that

- ◆ T2 sites:
 - ◇ GRIF in downtime until end of this week. Update appreciated on how soon it will be back - used for LHCb reprocessing.

Sites / Services round table:

- ASGC
 - ◆ no update on Condor-G issue affecting CREAM jobs (see FNAL report)
- BNL - ntr
- CNAF
 - ◆ at-risk downtime on Thu from 9 to 12 for ATLAS LFC upgrade to 1.8.0
- FNAL
 - ◆ there are tickets open against a few T1 CREAM CEs that do not work OK for the Condor-G pilot factory at FNAL
 - ◇ Maarten: there is a known issue in Condor's use of CREAM CE leases, I will forward details offline; the upshot is that the sites may be unable to do anything about such errors
- IN2P3
 - ◆ scheduled downtime was mainly to put more RAM into dCache servers; running jobs were canceled
 - ◆ also some CVMFS work was done
 - ◆ an issue with a CREAM CE for CMS led to an unscheduled downtime, now fixed: certificate DNs containing a '/' in the final CN did not work
- KIT - ntr
- NDGF
 - ◆ on Nov 4 starting at 18:00 for 10 h both OPN links to NDGF will be cut; also NLT1 may be affected by the intervention, but they would have a backup path via KIT; more news expected tomorrow
- NLT1
 - ◆ 1 out of 2 tape libraries had a problem leading to reduced performance, now fixed
- OSG
 - ◆ scheduled maintenance in progress, still 2.5 h to go
- PIC - ntr
- RAL
 - ◆ the site services have been ramped up slowly to full capacity without incident
- CASTOR/EOS
 - ◆ after yesterday's EOS update for ATLAS a problem was discovered: clients effectively have more permissions than they should, allowing files to be stored outside designated areas and with unclear ownership; an emergency update was applied, but then rolled back after it had led to crashes; the developers are on it
- dashboards - ntr
- databases - ntr
- GGUS/SNOW - ntr
- grid services - ntr

AOB:

Wednesday

Attendance: local(Mattia, Luca, Yuri, Maarten, Jamie, Jarka, Jan, Alessandro, Steve, Pepe, Raja);remote(Burt, Giovanni, Gonzalo, Jeremy, Ulf, Michael, John, Rolf, Pavel, Kyle, ASGC, Ron).

Experiments round table:

Wednesday

- ATLAS reports -
- T0/Central services
 - ◆ CERN-PROD_DATADISK transfer failures yesterday: Source error, failed to contact on remote SRM (srm-eosatlas.cern.ch). GGUS:75690 verified: the SRM daemon was running out of file descriptors (is running on new HW since yesterday). Hotfix applied, service restarted, seems to be OK.
- T1 sites
 - ◆ TRIUMF-LCG2_DATADISK ->INFN-MILANO-ATLASC transfer failures (timeouts). GGUS:75437 updated. Manual transfer checks are very slow, traffic TRIUMF->MILANO (and CERN Ixplus) goes through the research network (1G bandwidth), fully saturated.
- T2 sites
 - ◆ Last minute: related to calibration centre at Great Lakes / US. Since 09:00 this morning no calibration data to this centre. Informed expert but no reply. As transfers continuing failing submitted GGUS:75739 [↗](#) (half an hour ago)

- CMS reports -
- LHC / CMS detector
 - ◆ Data taking (high PU runs ongoing). Till Sunday!
- CERN / central services
 - ◆ CMSR Oracle instance spontaneous reboot problem, GGUS:74993, kept open to follow up with increased logging information, still waiting for a solution. **Asked for more information.**
- T0 and CAF:
 - ◆ cmst0 : very busy processing data !
 - ◆ GGUS:75743 alarm ticket -> Problems with T1 Transfer Pool - seems that transfer pool had overload and most transfers timed out in FTS. Ticket closed as soon as we checked all ok.
- T1 sites:
 - ◆ [T1_TW_ASGC]: general file access problem, GGUS:75377, possibly a software problem? **Site reports the file is OK in CASTOR. We will run again on the file to check that the error is still reproducible. Then decide how to proceed....**
 - ◆ [T1_FR_CCIN2P3]: overload in dCache affecting transfers from/to IN2P3. GGUS:75397. Ongoing. **The priority has been increased and this is mandatory to be solved soon. There is a large backlog of transfers (90TB) from IN2P3 to CMS Tier-2s, growing since two weeks (6 Oct), which is not being digested properly, even if they made changes on the FTS configuration:** See attached plot
Also, it seems that the proxy is expired and incoming transfers are failing as well. This needs to be looked at with high priority [Rolf - someone from local CMS support is working on this. Not aware of how much progress has been made but work is ongoing.]
 - ◆ [T1_IT_CNAF]: Failing transfers from and to CNAF.GGUS:75675. There is a problem with the STORM storage backend. Experts having a look. **Closed**
 - ◆ [T1_IT_CNAF]: CREAM reporting dead jobs @ CNAF T1 as REALLY-RUNNING. GGUS: 75648.
- Others:
 - ◆ On Monday we spotted out a problem with downtimes tracing in the SSB. There is a bug which affected, at least CNAF, which was marked as being in Unscheduled Downtime for some days, which was not the case. Savannah:124240. [Mattia - two downtimes for CNAF, confused monitoring. Will be updated by hand this time and working on longer term solution.]

- ALICE reports [↗](#) -
 - ◆ NTR

- LHCb reports -
 - Experiment activities
 - ◆ Reconstruction and stripping at CERN
 - ◆ Reprocessing at T1 sites and T2 sites
 - T0
 - ◆ CERN : Fairshare continues to have strange spikes. For example : Running jobs at CERN
 - T1 sites:
 - ◆ IN2P3 : GGUS ticket (https://ggus.eu/ws/ticket_info.php?ticket=75703) opened about downtime procedures. On hold for discussion at next T1 services coordination meeting.
 - ◆ SARA : Problem running in filling mode. Being followed through GGUS (https://ggus.eu/ws/ticket_info.php?ticket=75698)
 - T2 sites:
 - ◆ Aborted pilots at JINR (https://ggus.eu/ws/ticket_info.php?ticket=75710)
 - ◆ root (the package) problem at ECDF (https://ggus.eu/ws/ticket_info.php?ticket=75671)
-

Sites / Services round table:

- FNAL - ntr
- CNAF - this morning CNAF has joined LHCONe network and operations has gone smoothly
- PIC - ntr
- NDGF - update on networking trouble; peering should take care of everything. Same fibre has links for NDGF, NL-T1 and KIT. Backup link to KIT (CNAF) will carry traffic of all during downtime of 10h. Will probably be "loaded". Issue when 1 fibre can cut of 3 T1s at once. Have to ask networking people at these sites to know for sure. Data transfers will be somewhat slower when all going through CNAF
- BNL - ntr
- RAL - ntr
- KIT - ntr
- IN2P3 - nta
- ASGC - ntr
- NL-T1 - ntr
- GridPP - ntr
- OSG - maintenance yesterday had no issues

- CERN Grid: incident CMS VOMRS service not syncing correctly to VOMS. Holding up new users - on IT SSB. Intervention on Monday - would like to update web servers for T0 export to SL5 - has been on T2 service for some months; transparent, easy rollback. Agent node upgrades will happen in November stop. Ale - ATLAS fine.
- CERN Storage - new EOS; will try to upgrade ATLAS this afternoon.

AOB:

Thursday

Attendance: local(Yuri, Jamie, Alex, Pepe, Mattia, Raja, Jan);remote(Michael, Gonzalo, Ulf, Giovanni, John, Burt, Elizabeth, Jhen-Wei, Rolf, Andreas, Ronald).

Experiments round table:

- ATLAS reports -
- T0/Central services
 - ◆ CERN-PROD_ATLASDATADISK transfer failures: "source file doesn't exist". Savannah:88159 solved: on the 24th from 16:50-17:50 ~4500 files became invisible because the RPM update confused the failover process. All hidden files have been restored yesterday

at ~6pm.

- ◆ ATLAS-AMI-CERN DataBase at SLS showing no availability since ~8pm. The problem came from the CERN apache web redirector. Fixed this morning.
 - ◆ LSF share of ATLAS. Started large HLT production, jobs in queue couldn't start. Expert checked with LSF expert. Issue maybe related to LHCb - share for LHCb was increased day before yesterday and should yesterday have gone back to normal. ~1k slots not available. Soon after HLT jobs began running. In future it would be great if we could receive announcement of such changes - suddenly a large # jobs running and not understood why. Alex - noted
 - T1 sites
 - ◆ FZK downtime completed, all transfers continue.
 - T2 sites
 - ◆ AGLT2 (ATLAS Great Lakes T2) <- T0 export (calibration data) failures. GGUS:75739 solved: the number of cached files, not showing in tokens, increased to a point where all the pools were full or nearly full. AGLT2 team has freed up 1TB minimum on all pools. No failures anymore.
-
- CMS reports -
 - LHC / CMS detector
 - ◆ Data taking (high PU runs ongoing). Till Sunday!
 - CERN / central services
 - ◆ CMSR Oracle instance spontaneous reboot problem, GGUS:74993, kept open to follow up with increased logging information, but no other spontaneous reboot seen since then...
 - T0 and CAF:
 - ◆ cmst0 : very busy processing data !
 - T1 sites:
 - ◆ [T1_DE_KIT]: Poor Data Transfer Quality from T1_DE_KIT to Tier-2 Sites. GGUS:75778. Seems that there is a limitation in the KIT firewall which impacts T2/T3 transfers...
 - ◆ [T1_FR_CCIN2P3]: overload in dCache affecting transfers form/to IN2P3. GGUS:75397. New channel IN2P3-CMSSLOW created to deal with slow transfers to T2s, to help digesting the backlog. The configuration was updated to avoid transfers to timeout. CMS has started warning some US sites which uses an abusive number of files per FTS job. These sites are now blocking this CMSSLOW channel. They are reacting and changing their config. So far, the backlog is getting reduced (91.5 TBs queued 4 days ago - 65 TBs atm). Yesterday, the proxy expired and also affected incoming transfers to IN2P3. **This was fixed**. However, there are other transfer errors from/to other T1s which need to be understood as well (will comment on the ticket).
 - ◆ [T1_FR_CCIN2P3]: Custodial Subscription hasn't been approved in days. Savannah:124303.
 - ◆ [T1_IT_CNAF]: CREAM reporting dead jobs @ CNAF T1 as REALLY-RUNNING. GGUS:75648. JobIds provided and experts having a look.
 - ◆ [T1_TW_ASGC]: general file access problem, GGUS:75377, possibly a software problem? Yesterday we copied the file to FNAL and run interactively with the same CMSSW version, and the test was successful. Now, ASGC is trying the same interactive check.
 - ◆ [T1_TW_ASGC]: transfer of the dataset /ZbbToLL_M-30_7TeV-mcatnlo-photos/Summer11-PU_S4_START42_V11-v1/AODSIM from T1_TW_ASGC to some T2 sites. This transfer is stuck for a few days. Savannah:124305.
 - ◆ [T1_US_FNAL]: Permission denied at FNAL to transfer HI data to T2_CH_CERN. Savannah:124301. Suspect of using incorrect Role. (confirmed - will be changed)
 - ◆ transfers CERN->FNAL failing, misconfig of FTS in FNAL, not recognising EOS as valid ep BUG:124308 [↗](#)
 - Others:

- ◆ On Monday we spotted out a problem with downtimes tracing in the SSB. There is a bug which affected, at least CNAF, which was marked as being in Unscheduled Downtime for some days, which was not the case. Savannah:124240.
 - ◆ Migration from LCG-CE to CREAM not caught by nagios/dashboard. Savannah:124289. At level of dashboard-sam portal.
-

- ALICE reports [↗](#) - NTR
 - LHCb reports -
 - Experiment activities
 - ◆ Reconstruction and stripping at CERN
 - ◆ Reprocessing at T1 sites and T2 sites
 - T0
 - ◆ CERN : Spike of failed jobs this morning between 6AMUTC and 10AMUTC. These were jobs accessing d0t1 storage and the problem seems to have gone away now. [Jan - CC wide network intervention for 15' around 08:00] affected ~800 jobs
 - T1 sites:
 - ◆ First stage of reprocessing almost over.
 - T2 sites:
 - ◆ Minor problems at JINR, RHUL. Being followed by GGUS tickets
-

Sites / Services round table:

- PIC - ntr
- NDGF - it now seems all sites will have connectivity during network problems next week. .NL-T1 also backup connected to IN2P3. Warning downtime just in case.
- CNAF - intervention of ATLAS LFC e-p completed smoothly
- RAL - JANET will do emergency maintenance 0900 - 1700 BST on light path - no interruption expected
- FNAL - ntr other than Savannah reported earlier. Will look into other issues asap
- ASGC - CASTOR unstable. Next Wed sched downtime for power construction in d/s. 00:00 - 11:00 CERN time
- IN2P3 - ntr
- KIT - nta
- NL-T1 - ntr
- OSG - ntr

- CERN storage: we have had EOS ATLAS update at 14:00 to roll out latest stable version. Seems to have been successful. On CASTOR side quite a few upcoming mostly transparent interventions, will put these on SSB. Across all CASTOR instances. Also to latest CASTOR release

- CERN Grid - VOMRS / VOMS issue from yesterday now ok following intervention (Steve)

AOB:

- This Sunday clocks "fall back" in most European countries (but not in the US - that happens Nov 6): <http://www.timeanddate.com/time/dst/2011b.html> [↗](#)

Friday

Attendance: local(Eva, Nilo, Yuri, Maarten, Jamie, Pepe, Mike, Jan, Jhen-Wei, Raja, Alex);remote(Burt, Michael, Xavier, Mette, Giovanni, Gareth, Rolf, Onno).

Experiments round table:

- ATLAS reports -
 - T0/Central services
 - ◆ 2 CERN ATLAS central service machines voatl139, voatl161: swap full issue caused by apache. Fixed: httpd restart helped to resolve.
 - ◆ GGUS system unavailable this morning for ~1.5h (update doesn't work, ERROR 9130, then 10099). Elog:30859. NGI_IT was under maintenance?
 - T1 sites
 - ◆ TRIUMF still some issues with the transfer failures caused by the saturation of the local network. GGUS:75437 in progress.
 - ◆ Taiwan-LCG2 production job failures: "Get function can not be called...". GGUS:75785 solved: the software bug fixed.
 - T2 sites
 - ◆ INFN-MILANO-ATLASSC: SRM down, SAM tests and transfers fail. GGUS:75799 assigned.
-
- CMS reports -
 - LHC / CMS detector
 - ◆ Data taking (high PU runs ongoing). Till Sunday!
 - CERN / central services
 - ◆ CMSR Oracle instance spontaneous reboot problem, GGUS:74993, kept open to follow up with increased logging information, but no other spontaneous reboot seen since then...
 - T0 and CAF:
 - ◆ cmst0 : very busy processing data !
 - T1 sites:
 - ◆ [T1_FR_CCIN2P3]: overload in dCache affecting transfers from/to IN2P3. GGUS:75397. The backlog is getting slowly reduced (91.5 TBs queued 5 days ago - 64 TBs atm). [Maarten - is it clear how this backlog developed? Pepe - backlog generated as IN2P3 had a few days with dCache problems.] [Rolf - it appears that part of the backlog due to fact that transfer speed is not sufficient for some T2s.]
 - ◆ [T1_IT_CNAF]: CREAM reporting dead jobs @ CNAF T1 as REALLY-RUNNING. GGUS:75648. JobIds provided and experts debugging the issue with the blah developers. It seems blah had missed to update the status of these jobs. Anyway it seems that all these jobs had been killed by the batch system because they reached the batch system memory limit(2.5GB)...
 - ◆ [T1_IT_CNAF]: 1 corrupted file at CNAF. Savannah:124310. This is the only replica of the file. It might need an invalidation. DataOps taking a look to it.
 - ◆ [T1_TW_ASGC]: general file access problem, GGUS:75377, possibly a software problem? Yesterday we copied the file to FNAL and run interactively with the same CMSSW version, and the test was successful. Now, ASGC is trying the same interactive check. [Jhen-Wei - run locally tests twice and both failed. Try to retransfer file and run local tests again. Can show log file later]
 - ◆ [T1_TW_ASGC]: transfer of the dataset /ZbbToLL_M-30_7TeV-mcatnlo-photos/Summer11-PU_S4_START42_V11-v1/AODSIM from T1_TW_ASGC to some T2 sites. This transfer is stuck for a few days. GGUS:75780. Seems solved atm. [One problematic disk server in pool and all failed files were there. Disabled and try to restage all files to other disk servers]
 - ◆ [T1_TW_ASGC]: CMS Fall11 jobs have problems opening some files. GGUS:75784. There were problems staging the files. Now, that's been forced. CMS needs to certify this is no a problem anymore...
 - ◆ [T1_US_FNAL]: Files staged on disk at FNAL, but cannot be transferred to CSCS. Savannah:124312. Files are visible well from FNAL to PhEDEx. Need to understand why

this is not processed (it could eventually do..).

- Tier-2s:
 - ◆ Business as usual...
 - Others:
 - ◆ On Monday we spotted out a problem with downtimes tracing in the SSB. There is a bug which affected, at least CNAF, which was marked as being in Unscheduled Downtime for some days, which was not the case. Savannah:124240.
 - ◆ Migration from LCG-CE to CREAM not caught by nagios/dashboard. Savannah:124289 --> The 'latest result' page is ok, however the historical plots same CE for both flavours displayed there, if the same CE name was used for a transition LCG-CE<->CREAM-CE. To be fixed.
-

- ALICE reports [↗](#) -
 - ◆ NTR
-

- LHCb reports -
 - Experiment activities
 - ◆ Prompt reconstruction and stripping at CERN and Tier-1 sites.
 - ◆ 1st round of reprocessing at T1 sites and T2 sites almost over (some tail going on - especially at GridKa)
 - ◆ Next round of reprocessing to start next week.
 - T0
 - ◆ Going through stripping backlog slowly.
 - T1 sites:
 - ◆ Possible problem with GridKA SE - following up offline with site admin (Xavier) to understand what is happening.
 - T2 sites:
 - ◆ Aborted pilots at a few sites (Weizmann, Lancaster)
-

Sites / Services round table:

- FNAL - ntr
- BNL - ntr
- KIT - no issues; on Tuesday a public holiday and many people will be away also Monday so "best effort".
- NDGF - ntr
- CNAF - ntr
- RAL - have declared an at risk on Tuesday 1 November in morning for some networking work that should be transparent
- IN2P3 - next week Nov 1 public holiday in France and bridge on Monday. Rajah - what level of support will be there - we plan to start reprocessing and a lot of data is at IN2P3. Rolf - Monday is a normal working day and service on Tuesday is like the one on Sundays; the engineer on duty will be a dCache specialist.
- NL-T1 - we have a problem with MSS. Heavily used by several VOs simultaneously. Holding store and restore queues to recover backlog. Queues will be reopened at 17:00 CEST
- ASGC - ntr

- CERN Dashboards - problem reported by CMS being looked into although ticket not updated. Will give relevant person a poke
- CERN grid - an update of FTS T0 webservers on Monday. Should be transparent.

- CERN Storage - series of updates; DB updates Mon/Tue, n/s on Tue, expected to be transparent. Glitch on srmpublic and srmalice - pure monitoring issue. Trying to get out of castorpublic for these

tests.

AOB:

-- JamieShiers - 14-Sep-2011

This topic: LCG > WLCGDailyMeetingsWeek111024

Topic revision: r13 - 2011-10-28 - JamieShiers



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