

# Table of Contents

<b>Week of 110425.....</b>	<b>1</b>
Daily WLCG Operations Call details.....	1
WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	1
Wednesday.....	3
Thursday.....	5
Friday.....	8

# Week of 110425

## 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
<a href="#">ALICE</a>	<a href="#">ATLAS</a>	<a href="#">CMS</a>	<a href="#">LHCb</a>	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:

- No meeting - CERN closed.

## Tuesday:

Attendance: local(Eva, Nilo, Ricardo, Ignacio, Maarten, MariaG, Stephane, Ale, Fernando, Dan, Dirk);remote(Michael/BNL, Gareth/RAL, LHCb/Renato, Rob/OSG, Xavier/KIT, Gonzalo/PIC, Ian/CMS, Jon/FNAL, Huang/ASGC, Chrsitian/NDGF, CNAF, Jeff/NL-T1).

Experiments round table:

- ATLAS reports - Fernando
  - ◆ In a nutshell: Physics all day (data11\_7TeV) with short calibration periods and few interruptions
  - ◆ Peak luminosity record broken for a hadron collider
  - ◆ T0
    - ◇ Tier0 manager hung and no heartbeat in T0 monitor (twice over the Easter break). Seems to be related to bpeek option in the task lister. Experts are aware and investigating the problem.
    - ◇ Massive problems reading and writing to ATLAS svcclass/pool T0merge (twice over the Easter break). Caused by a backup job that should have been migrated to a

different node. ([https://gus.fzk.de/ws/ticket\\_info.php?ticket=69992](https://gus.fzk.de/ws/ticket_info.php?ticket=69992),  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=69954](https://gus.fzk.de/ws/ticket_info.php?ticket=69954))

- ◆ T1s
    - ◇ IN2P3-CC: Transfers to/from the site were failing heavily (including T0 export) on Friday between 4AM-8:30AM. Presumably dCache problem - only known solution is to restart the pool manager. ([https://gus.fzk.de/ws/ticket\\_info.php?ticket=69925](https://gus.fzk.de/ws/ticket_info.php?ticket=69925))
    - ◇ Staging from tape broken for TRIUMF-LCG2\_MCTAPE since Monday morning: "Pinning failed: finding read pool failed". ([https://gus.fzk.de/ws/ticket\\_info.php?ticket=69975](https://gus.fzk.de/ws/ticket_info.php?ticket=69975), [https://gus.fzk.de/ws/ticket\\_info.php?ticket=69959](https://gus.fzk.de/ws/ticket_info.php?ticket=69959))
  - ◆ Central Services
    - ◇ Downtime collector stuck - quattor template of the machine had been incorrectly modified and was preventing any activity of the user under which the collectors run. This is why RAL was not automatically re-included in DDM activity for a couple of hours after coming back from their downtime on Thursday ~12:00AM.
  - ◆ Ale: still need better understanding for second castor problem. Ignacio: will follow-up with an incident report once fully understood. Still investigating several possible causes (DB backup, scaling issues with castor file name dump, overwhelmed rsyslog) - most likely candidate is rsyslog problem.
  - ◆ Stephane: VOMS failover after BNL problem did now work as expected. Maarten: not clear as timeout for failover is 3 minutes and people may not have waited sufficiently long. Ale: main message for upcoming VOMS intervention: old service should immediately reject client requests to avoid confusion as also discussed at T1 service coordination meeting last week.
- CMS reports - Ian
    - ◆ LHC / CMS detector
      - ◇ Excellent running over the weekend. Good live time
    - ◆ CERN / central services
      - ◇ Nothing to report
    - ◆ Tier-0 / CAF
      - ◇ Last 24 hours CMS has been averaging 85% utilization of the Tier-0
    - ◆ Tier-1
      - ◇ 2010 Reprocessing launched at Tier-1s
      - ◇ Reprocessed some limited 2011 datasets over the weekend. Good responses by Tier-1
      - ◇ Data Ops has requested a consistency check of the site contacts
    - ◆ Tier-2
      - ◇ MC production and analysis in progress. Reduced effort over the holiday
    - ◆ Other
      - ◇ CRC-on-Duty : Peter Kreuzer as of this evening.
    - ◆ Ricardo: saw swap full at T0 - who is working on this on the CMS side? Ian: current workflow has high memory consumption. David Lang is following this closely, but consumption should go down by 400MB - 1 GB during the next 48h.
  - ALICE reports [↗](#) - Maarten
    - ◆ T0 site
      - ◇ Nothing to report
    - ◆ T1 sites
      - ◇ Nothing to report
    - ◆ T2 sites
      - ◇ Usual operations
  - LHCb reports - Renato
    - ◆ RAW data distribution and their FULL reconstruction is going on at most Tier-1s.
    - ◆ A lot of MC continues to run.

- ◆ T0
  - ◇ Problem SOLVED: Problems staging files out of Tape (72 files). Files requested to be staged yesterday we would have expected them online.
  - ◇ Yesterday, there were files missing to the OFF-LINE due to a hardware failure. Now (earlier) files start to move to OFF-LINE.
- ◆ T1
  - ◇ IN2P3: Problems with software installation, "share" set to zero. Solution is in progress.
  - ◇ RAL: Storage Elements full (RAW and RDST, which use the same Space Token). It was reported that "some tape drives becoming stuck and not working", which seems to be fixed. However, there still a big backlog.
- ◆ T2

Sites / Services round table:

- Michael/BNL - long standing BNL - CNAF network issues closed! Yesterday morning 9:00 hic-up with VOMS tomcat server (probe exists to detect these problems). Service was restarted at 10:00. Oracle RAC for conditions needs to move to other data center: scheduled outage for 4h tomorrow. Maarten: VOMS service was not dead, but did not handle connections anymore. Therefore a modified timeout on connection would not have helped much.
- Gareth/RAL - problems with ATLAS sw server, also with CVMFS (corruption). CVMFS issues should be resolved in latest CVMFS client version. CASTOR: DB backend problems. LHCb castor area became full after tape drive problems. disk cache became full and as knock-on effect also reads from tape were affected as garbage collection cleaned new read from tap before they could be accessed by LHCb . Additional resources have been added and the system is now recovering.
- Rob/OSG - upstream probe problem with sam collector last night. Data will be resent to complete the missing time window.
- Xavier/KIT - FTS Oracle back-end problems on Sat morning - FTS channel restart fixed the problem.
- Gonzalo/PIC - ntr
- Jon/FNAL - ntr
- Huang/ASGC - Reminder: castor in downtime for upgrade until Fri night. Also planned network maintenance with limited connectivity from 10pm to 4 am UTC
- CNAF - CE experienced "out of memory" problems - now back to normal
- Jeff/NL-T1 - ntr
- Mat/IN2P3 - dcache ticket for recent ATLAS problem still open and no answer yet. Operations are monitoring the system and will restart in case the problem re-occurs
- Chrstian/NDGF - ntr
- **CERN VOMS service** The certificate for the LHC voms services on voms.cern.ch will be updated tomorrow on Wednesday around 10:00 CEST April 27th. The current version of lcg-vomscerts is 6.4.0 and was released 2 weeks ago. It should certainly be applied to gLite 3.1 WMS and FTS services. [ Has been put into release for those services a few weeks ago. T1s running FTS services should be sure that they have latest version of RPM.

AOB:

## Wednesday

Attendance: local(MariaG, Peter, Ricardo, Ignacio, Ale, Dan, Eduardo, Maarten, Mattia, Luca, Dirk);remote(Michael/BNL, Jon/FNAL, John/RAL, Giovanni/CNAF, Felix/ASGC, Onno/NL-T1, Marc/IN2P3, Federico:/HCb, Rob/OSG, Christian/NDGF, Foued/KIT).

Experiments round table:

- ATLAS reports - Dan

- ◆ Calibrations / Standalone in the morning.
- ◆ New Runs when injection physics occurred.
  - ◇ 23:33 Run 180309 Physics Ready
  - ◇ 07:10 Run 180309 Stopped 12 pb-1
- ◆ T0/Central Services
  - ◇ Network glitch ~16:45-17:00 (GGUS:70026). Answered in IT Service Status Board before the ticket was routed.
  - ◇ Brief outages in ~all services (observed Panda, T0 ConTZole, DDM Central Cat.) due to above network glitch.
- ◆ T1s
  - ◇ TAIWAN scheduled outage started Tuesday 04:00. Auto excluded in DDM and manually excluded in Santa Claus (RAW data not subscribed to TAIWAN Tape).
  - ◇ SARA analysis pilots problem returns. ~50% analysis jobs fail due to incorrect mapping of the credentials to dteam, hence no dcache read permission (Savannah:120527). Related pilot factory was stopped overnight. This morning SARA fixed the DN->uid mapping in dCache for /atlas/Role=pilot; Now it seems to work.
  - ◇ RAL SCRATCHDISK transfer errors: (GGUS:69863).  
"RAL-LCG2\_SCRATCHDISK has again a lot of errors (1736) with 6 attempts each".  
BDII issue fixed, CASTOR issue is that 6 of 9 servers are full. RAL asked ATLAS to decrease activity or delete some files.
- CMS reports - Peter
  - ◆ LHC / CMS detector
    - ◇ Quite good data taking
  - ◆ CERN / central services
    - ◇ CASTORCMS degradation on various instances since Apr 27, 4AM : was due to rsyslogd being stuck in one the of the CASTOR headnodes. Affected T0STREAMER and CMSPRODLOGS pools (and most likely also T0TEMP pool and transfers to T1s), see GGUS:70050 . After the deaomon got restarted (around 9:30 AM) things started to work again.
    - ◇ CMS Dashboard not reachable around Apr 27, 10:50AM, for 10 minutes, due to NT6R database not being available ( see <http://it-support-servicestatus.web.cern.ch/it-support-servicestatus/INT6Rdatabaseunavailable> )
  - ◆ Tier-0 / CAF
    - ◇ High memory consumption on T0/CMSSW application side. Two immediate actions were taken : (i) increase max-memory on WN to avoid crashes and (ii) decrease the number of events per job, which increased the number of jobs, see <http://lsf-rrd.cern.ch/lrf-lsf/info.php?queue=cmst0>. Will have an application patch today.
  - ◆ Tier-1
    - ◇ 2010 Reprocessing on-going at Tier-1s
    - ◇ Instabilities at PIC, caused by several WNs that were mis-configured (see Savannah:120592)
    - ◇ ASGC downtime (CASTOR upgrade) this week
  - ◆ Tier-2
    - ◇ MC production and analysis in progress (summer11 production)
  - ◆ Other
    - ◇ CMS still waiting for unscheduled downtimes to be included in the SSB, see Savannah:119944
  - ◆ Ignacio: root cause for CMS castor problems (and for similar ATLAS problems yesterday) has been a stuck rsyslog daemon - we will reconfigure for UDP transport instead of TCP to avoid logging problems to affect the service.

- ◆ Maria: why is the production dashboard run on an integration DB server (int6r)? Luca: This was not know to DB team and will be fixed.
  
- ALICE reports [↗](#) - Maarten
  - ◆ T0 site
    - ◇ Nothing to report
  - ◆ T1 sites
    - ◇ Nothing to report
  - ◆ T2 sites
    - ◇ Usual operations
  
- LHCb reports - Federico
  - ◆ Experiment activities:
    - ◇ RAW data distribution and their FULL reconstruction is going on at most Tier-1s.
    - ◇ A lot of MC continues to run.
  - ◆ T0
  - ◆ T1
    - ◇ IN2P3: our jobs hitting the memory limits, being killed. Not clear why only at Asked to increase limits to 5Gb.
    - ◇ RAL: backlog of activities. Staging problems. Disk pools were full.
  - ◆ T2
  - ◆ Mattia: DaVinci problem at IN2P3 was fixed late afternoon. Until then all sam tests failed at all sites. Federico: used old DaVinci - now updated and new tests should go through.

Sites / Services round table:

- Jon/FNAL - ntr
- John/RAL - scratch issue as reported by ATLAS. Ale: did you increase the size of scratch recently (eg new servers)? John: no - rather asking ATLAS to reduce use of scratch - Ale: 6 are full but 3 are still free - maybe distribution is not smooth enough to use available resources. John: will follow-up
- Giovanni/CNAF - ntr
- Felix/ASGC - 2nd day of castor upgrade: head nodes have been upgraded to 2.1.10 (srm 2.10) - status of instance is fine - now moving to disk server upgrades.
- Onno/NL-T1 - ntr
- Marc/IN2P3 - ntr
- Rob/OSG - yesterday afternoon: bdii at cern took several minutes to return data - better this morning. Ricardo: Steve took machine out round-robin for investigation - will update ticket as soon more details are known.
- Christian/NDGF -ntr
- Foued/KIT -ntr
- Eduardo/CERN: network glitch from 16:48-16:50 due to config mistake: part of LGC traffic was lost. Will start using new public prefixes at CERN on 1 June (GPN and LCG). A mail with the details was send to LCG ops list so that sites can update their configuration before the change becomes active.

AOB:

## Thursday

Attendance: local(Ricardo, Steve, Maarten, Dan, Ale, Mattia. Peter, Eva, Fernando, Lola, Jamie, MariaG, Ignacio, Dirk);remote(Michael/BNL, Jon/FNAL, JT/NL-T1, Kyle/OSG, Marc/IN2P3?, Huang/ASGC, Gonzalo/PIC, Gareth/RAL Giovanni/CNAF, Federico/LHCb, Dimitri/KIT).

Experiments round table:

- ATLAS reports - Dan
  - ◆ 23:18 Run 180400 Physics Ready. So far ~16 pb<sup>-1</sup> Collected
  - ◆ T0/Central Services
    - ◇ Outage in one DDM Central Catalog reader for 3 hours last evening. Load balancer redirected queries to another host. Down service returned after a restart.
  - ◆ T1s
    - ◇ TAIWAN scheduled outage continues.
    - ◇ RAL SCRATCHDISK issue (GGUS:69863 [↗](#)): RAL put the servers into drain to attempt rebalance, and decrease # analysis jobslots to decrease load. ATLAS is deleting old data. "Failures of transfers to RAL SCRATCHDISK have now been brought under control"
    - ◇ BNL
      - dcsrm.usatlas.bnl.gov showed expired certificate in FTS transfers (GGUS:70067 [↗](#)).
        - Cert was indeed renewed 12 days ago, so cause is unknown. Restart service solved problem.
      - US LFCs were not updated with the new voms.cern.ch certificate.
        - ~50% Failed transfers,
        - see <https://atlas-logbook.cern.ch/elog/ATLAS+Computer+Operations+Logbook/2>
        - Certs were updated and problem resolved.
      - US T2s were auto-excluded from DA by HammerCloud because they did not have the latest DBRelease replica.
        - (DA client brokerage avoids sites without this dataset).
        - Was the result of a missing FTS stream. Issue was fixed and sites were whitelisted after test jobs started succeeding.
  - ◆ ATLAS requests that sites follow the forthcoming instructions (to be provide by CERN VOMS service report) so that update of the voms certs rpm will be no longer required in future.
    - ◇ VOMSLSCfileConfiguration
- CMS reports - Peter
  - ◆ Long LHC fill on 624 bunches during the night (for CMS : 598 bunches & Instantaneous Luminosity ~ 6.6E32 at beginning)
    - ◇ CMS recorded 16.4 pb<sup>-1</sup>, which represents ~10% of the total integrated luminosity for CMS in 2011 !
    - ◇ CMS also performed high rate L1 trigger tests at the end of last night's fill : measured 84 Khz L1 rate (HLT rate ok at 310 Hz)
  - ◆ CERN / central services
    - ◇ CMS had central PhEDEx agents down for 1 3/4 hour on Wed Apr 27, 17:00 - 18:45 , see [https://sls.cern.ch/sls/history.php?period=24h&offset=null&more=availability&id=CMS\\_PhedEx](https://sls.cern.ch/sls/history.php?period=24h&offset=null&more=availability&id=CMS_PhedEx) . CMS transfer experts investigating.
  - ◆ Tier-0 / CAF
    - ◇ High memory consumption on T0/CMSSW application side : after patch release, jobs are still consuming high memory (2 - 2.5 GB). Increase in max-memory on WN + shorten the jobs helped to avoid crashes
    - ◇ Two related cmst0 LSF tickets that need to be followed up :
      - INC033388 : need to recover of 16 cmst0 hosts that went down lately, hence reducing our resources substantially : lxb8930, lxb8941, lxb8995, lxb8998, lxbrb2111, lxbrg0209, lxbrg1003, lxbrg1008, lxbrg1203, lxbrg2007, lxbrg3802, lxbrg4008, lxbrg4402, lxbrg5801, lxbrl2304, lxbsu0919
      - INC034775 : jobs stay in failed state on worker nodes for days, but bjobs claims the job is still in the RUN state (currently CMS operators must

manually scan bjobs -w then investigate on the worker node itself to try to determine jobs that LSF claims has been running for an excessive time have actually failed)

- Answer by LSF Support to one particular error report : it seems that the job (19566) is still on the machine, LSF is still waiting for it to finish.
- A pstack shows that the job seems to be hung on trying to close the input file. So it looks like NOT an LSF problem and needs further investigation on the CMS side.

◆ Tier-1

- ◇ 2010 Reprocessing on-going at Tier-1s
- ◇ ASGC downtime (CASTOR upgrade) this week

◆ Tier-2

- ◇ MC production and analysis in progress (summer11 production)

◆ Other

- ◇ RQF0009755 : CMS would like to be notified in advance in case of a planed maintenance on CERN/WMS machines, in order to avoid Computing Shift Alarms when the CMS Service Maps show a service degradation (<http://servicemap.cern.ch/ccrc08/servicemap1.html?vo=CMS&site=> and <https://sls.cern.ch/sls/service.php?id=WMS-CMS>). The current degradation seems to be due to the fact that 1 WMS is in maintenance (<https://lemonweb.cern.ch/lemon-status/info.php?entity=wms212&type=host>)
- ◇ Unscheduled downtimes to be included in the SSB (see Savannah:119944) : being currently followed up by Dashboard team
- ◇ Peter: could we get info on maintenance operations on WMSs? Maarten: should be transparent to you. Ricardo: PES will look into this.

• ALICE reports - Lola

◆ T0 site

- ◇ Yesterday there were found around 4k jobs not monitored at CERN (CREAM and LCG) due to the crashing of the CMreport, a Cluster Monitor component. From now on this daemon will be restarted automatically from MonaLisa every time it crashes

◆ T1 sites

- ◇ Nothing to report

◆ T2 sites

- ◇ Usual operations

• LHCb reports - Federico

- ◆ RAW data distribution and their FULL reconstruction is going on at most Tier-1s.

- ◆ A lot of MC continues to run.

◆ T0

◆ T1

- ◇ IN2P3: solved problems with jobs hitting the memory limits

- ◇ RAL: backlog of activities. Staging problems. Disk pools were full. Starting to going through, submission was throttled 2 days, now slowly recovering

◆ T2

- ◆ Mattia: problem with sam tests mentioned yesterday is solved since yesterday afternoon.

Similar problem at IN2P3 - Stefan Roiser is following up.

- ◆ Jeff: Decreasing number of LHCb jobs at NL-T1 - due to LHCb throttling? Federico: No, throttling was only at RAL for specific problems there. Will check for NL-T1.

Sites / Services round table:

- Michael/BNL - problems mentioned above were resolved within 1-2h

- Jon/FNAL - trouble yesterday with CMS jobs which was traced to change of framework default from 32 to 64 bit execution. This change involved changing making a file into a symbolic link, which in fact lead to corruption of the CVMFS cache - many job failures until CVMFS had been restarted. Also wanted to update cvsfs client but s/w did not compile. FNAL/CMS working on these issues.
- Jeff/NL-T1 - On May 10 srm will be down for maintenance for 2h. Next Thu is a national holiday in the Netherlands and site will not attend this meeting.
- Marc/IN2P3 - ntr
- Huang/ASGC - finished upgrade of castor disk servers - remaining services should be done by tomorrow
- Gonzalo/PIC - problem with ATLAS prod jobs for 1.5 days: inefficient/slow copy of input files. many jobs failed due to timeouts. Traced to new version of dcap client. Site is going back to previous version and is in touch to dcache development for fixes to the new client.
- Gareth/RAL - ATLAS scratch: stopped draining of disk servers. Writing is ok now but reading still has timeouts due to small number of disk servers involved (some servers are much bigger in storage capacity and hence attracting most data and users). LHCb issues largely solved. Tomorrow + Monday are holiday in the UK.
- Giovanni/CNAF -
- Federico/LHCb -
- Dimitri/KIT - ntr
- Kyle/OSG - ntr
- Eva/CERN - will deploy latest oracle patches next week on validation DBs - move to production DBs planned for technical stop in 2 weeks.
- Ignacio/CERN - rsyslog config changed to UDP (from TCP) for ATLAS and CMS instances after consulting with experiments - this should prevent the recent problems with ATLAS and CMS instances.
  
- Steve described the VOMS intervention below and is available for any questions by email.
- **CERN Voms** As proposed at the co-ordination meeting I now confirm the date of LHC VOMS migration from SL4 to SL5 as **Tuesday 10th of May**. In particular **voms-proxy-init** will be **unavailable** for CERN hosted VOs from around 07:00 UTC to around 09:00 UTC. Registration processing will be down a little longer. Any updates relating to the process will be given at this meeting. Please contact GGUS or [steve.traylen@cern.ch](mailto:steve.traylen@cern.ch) with questions or comments. VomsInterventions contains the following:
  - ◆ Details relating to failover to BNL and FNAL during the downtime.
  - ◆ Details of any expected changes in the new service (only cosmetic).
  - ◆ A detailed time frame of what will happen.
  - ◆ Details on how to test the SL5 service now.
  
- How to configure "LSC" files on VOMS-aware services and avoid the need to update the `lcg-vomscerts rpm` e.g. on LFC servers:
  - ◆ <https://twiki.cern.ch/twiki/bin/view/LCG/VOMSLSCfileConfiguration>

AOB:

## Friday

Attendance: local(Fernando, Dan, Massimo, Mattia, Michal, Ricardo, Eva, Nilo, Ale, Ian, MariaG, Jamie, Maarten, Dirk);remote(Michael/BNL, Xavier/KIT, Gonzalo/PIC, Giovanni/CNAF, Felix/ASGC, Jon/FNAL, Rob/OSG, Federico/LHCb, Marc/IN2P3, Onno/NL-T1).

Experiments round table:

- ATLAS reports - Dan
  - ◆ Thursday Afternoon and evening ATLAS Physics Runs

- ◊ Collected this year:  $\sim 1/5 \text{ fb}^{-1}$
- ◆ Central Services:
  - ◊ Central catalog outages from  $\sim 13:00$  and  $16:00$ . Seem to be caused by a specific CLI query (which was being used to diagnose the RAL SCRATCHDISK issue). Had to physically reboot machine.
    - <https://atlas-logbook.cern.ch/elog/ATLAS+Computer+Operations+Logbook/24869>
- ◆ T1s
  - ◊ PIC (GGUS:70069): Slow dccp resolved after dcap library downgrade.
  - ◊ IN2P3 ALARM about T0 Export failures (GGUS:70113). GridFTP process stopped on the ccdcatli012 @ 23h30 28/04:11 CEST. Restarted service back at  $\sim 10:30$  CEST.
  - ◊ Testing FTS transfer with overwrite option. Worked to CERN, but not to SARA (GGUS:70128). SARA reports this is disabled (must delete then recopy).
- ◆ Marc?IN2P3: know issue with dcache ftp - connections stay open even after transfer and max connection count was reached. This has been reported already earlier to dcache developers. Maarten: overwrite - should consider change of FTS behavior and avoid dangerous overwrite flag. Ale: discussing with FTS people to be sure of risks and will check offline with Maarten the detailed logs.
- CMS reports - Ian
  - ◆ LHC / CMS detector
    - ◊ LHC fill on 624 bunches since yesterday evening during the whole night + morning
    - ◊ CMS recorded  $\sim 22 \text{ pb}^{-1}$  since yesterday evening 19:40
  - ◆ CERN / central services
    - ◊ The CMSR Database became unreachable around midnight : all the sessions were blocked and one of the instances was completely inaccessible. It was recovered by the IT DB expert about 1 hour later. The root cause was still being investigate on Apr 29, morning.
      - CMS opened an ALARM ticket at 23:59 (21:59 UTC) : GGUS:70114. Since no reaction within 20 minutes and given the emergency, CRC called 75011 around 00:20 (PHY DB piquet phone as specified here). At 00:28 a response to the ALARM ticket was given by IT Computing Operations : "CASTOR Piquet has been called. Please standby.". It is not clear to CMS why the CASTOR Piquet was called in that case ?
      - At 00:40, IT DB expert (Eva Dafonte Pérez) called the CMS CRC and explained the incident. The problem got solved 20 minutes later.
      - Note that the IT DB expert received the ALARM ticket only around 01:45 AM (at first the ALARM ticket was referenced to her by the CMS CRC). So it seems that the ALARM ticket is not the fastest way to reach IT DB experts in case of an emergency.
      - A SIR has been requested to clarify if there were any avoidable delays in the communication chain and to understand why the GGUS ticket did not reach the DB team directly.
  - ◆ Tier-0 / CAF
    - ◊ High memory consumption and crashes on cmst0 due to application are still a worry (despite the patch release + increase in max-memory per WN + reduction of event/jobs)
      - As a consequence, an increasing number of cmt0 nodes became non-responsive, see INC:033388. The issue is being monitored and handled in close collaboration between IT/LSF and CMS.
  - ◆ Tier-1
    - ◊ 2010 Reprocessing on-going at Tier-1s
    - ◊ ASGC downtime (CASTOR upgrade) this week
  - ◆ Tier-2
    - ◊ MC production and analysis in progress (summer11 production)

- ◆ AOB (all related to CMS Computing Shift Monitoring)
    - ◇ On April 27, CMS reported here that CMS Dashboard was not reachable around 10:50AM, for 10 minutes, due to NT6R database not being available ( see <http://it-support-servicestatus.web.cern.ch/it-support-servicestatus/INT6Rdatabaseunavailable> )
      - after investigating back, it turns out this was a misunderstanding between the CMS shifter, who raised an alarm about the above announcement to CMS Computing responsible, which got misinterpreted as a Dashboard production downtime. This was a mistake and we apologize for that
    - ◇ Site Status Board "lense" (Savannah:120367 [↗](#)) : this issue is still pending, namely that CMS would like to recover the old-SSB functionality to set the "lense" symbol to be able to point to open tickets from the SSB and make the trouble shooting more transparent.
    - ◇ WMS monitoring (RQF:0009755 [↗](#)) : as reported yesterday, CMS would like to be notified in advance in case of a planed maintenance on CERN/WMS machines. Until there is an answer to this request (if any), CMS decided to integrate in its Computing Shift instructions that it is expected to have reduced WMS availability for a while, and to raise an alarm only if the WMS SLS status (availability) falls below 50%.
  - ◆ Ricardo: will look resource settings in LSF and get back to CMS with a proposal.
  - ◆ Mattia: the problem with unscheduled downtime in dashbaord ssb should be fixed by fixing the sites topology (now taken from sitedb); the problem with ssb 'lense' will be fixed soon (Savannah:120367 [↗](#)).
  - ◆ Eva: CMS offline DB problem not yet understood - nothing in the logs. The piquet call to CASTOR was a typo in the sysadmin email - in fact the DB piquet was called (correctly).
    - ◇ A first version of the SIR (still in open state) is available at <https://twiki.cern.ch/twiki/bin/view/DB/PostMortem28Apr11>
- ALICE reports [↗](#) - Maarten
    - ◆ Activity ramping up further for Quark Matter conference (May 23-28)
    - ◆ T0 site
      - ◇ Nothing to report
    - ◆ T1 sites
      - ◇ Nothing to report
    - ◆ T2 sites
      - ◇ Usual operations
  - LHCb reports - Federico
    - ◆ RAW data distribution and their FULL reconstruction is going on at most Tier-1s.
    - ◆ A lot of MC continues to run.
    - ◆ T0
      - ◇ CERN: (GGUS:70135 [↗](#)) files can't be staged when trying a replication.
    - ◆ T1
    - ◆ T2
    - ◆ Mattia: sam test failures at PIC: space token name have changed already at PIC (in line with experiment plan) but the tests still use the old name. Should do test with both space tokens until all sites have changed. Federico: Will check with Stefan Roiser.

Sites / Services round table:

- Jon/FNAL -First, more CMS job failures yesterday due to condor incompatibility with new CRAB release. CRAB was fixed. Second, based on yesterday's upcoming voms upgrade: (1) We tested the new sl5 service and it works. (2) We are worried the failover to voms.fnal.gov is not going to work since most sites have taken voms.fnal.gov out of the vomses file because the voms.fnal.gov entry is not widely propagated to the European CEs. (3) We think we can weather the planned outage.

- Michael/BNL - ntr
- Xavier/KIT - ntr
- Gonzalo/PIC - ntr
- Giovanni/CNAF- ntr
- Felix/ASGC- last day of castor upgrade - system looks fine now. Site will be put back online as planned.
- Marc/IN2P3- ntr
- Onno/NL-T1- ntr
- Rob/OSG- ntr
- Massimo/CERN - transparent change for rsyslog configuration (done this week for ATLAS and CMS) will take place for LHCb and Alice on Mon.

AOB:

- From Fernando López Muñoz via email
  - ◆ We have some problems with DE-KIT when our primary LHCOPN link is down. When this occurs, all the transfers between DE-KIT and ES-PIC fails because T1-T1 transfers goes via generic IP network To solve this problem, we have planned and outage next 5th of May between 9:00 to 13:00 CEST in our primary LHCOPN link.
  - ◆ Transfers between T0-T1 will be degraded in the maintenance window because all traffic will be re-routed through 1 Gbps CERN-PIC link. Transfers between T1-T1 will be re-routed to our 10 Gbps Generic IP network
  - ◆ We have created the following GGUS ticket [https://gus.fzk.de/pages/ticket\\_lhcopn\\_details.php?ticket=70073](https://gus.fzk.de/pages/ticket_lhcopn_details.php?ticket=70073)

-- JamieShiers - 19-Apr-2011

---

This topic: LCG > WLCGDailyMeetingsWeek110425

Topic revision: r21 - 2011-04-29 - MattiaCinquilli



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