

Table of Contents

Week of 120709.....	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.....	5
Friday.....	7

Week of 120709

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(Torre, Luc, Giuseppe/CMS, Eva, Xavier, Claudio, Maarten, Eddy, MariaDZ, Dirk);remote(Michael/BNL, Stefano/LHCb, Boris/NDGF, Kyle/OSG, Jhen-Wei/ASGC, Paolo/CNAF, Lisa/FNAL, Rolf/IN2P3, Tiju/RAL, Dimitri/KIT).

Experiments round table:

- ATLAS reports -
 - ◆ CERN CENTRAL SERVICES, T0
 - ◇ CERN-PROD_TZERO ALARM for slow bsub submission Sat 7:13, ticked update at 10:35 that hot fixes were applied, fixed the problem. GGUS:83947
 - ◇ IT removed from T0 export Sun pm in anticipation of scheduled INFN-T1 storage outage commencing Mon am
 - ◇ ES removed from T0 export today in anticipation of scheduled PIC outage commencing tomorrow
 - ◆ T1
 - ◇ BNL: short lived problem with BNL transfers Sat pm, resolved by site before any ticket, by restarting the SRM and the dCache admin domains. BNL is analyzing the problem.
 - ◇ PIC T0 data export restored Sat pm with transfers stable after resolution of PNFS saturation problem. GGUS:83923

- ◇ PIC transfer failures Sat night, dCache pool rebooted spontaneously and had some trouble remounting the disks, fixed Sun am, data available again about 15:30 Sun pm. GGUS:83952 [↗](#)
- ◆ T2
 - ◇ NTR
 - ◆ Michael/BNL: performed analysis of the issue above. Problem occurred while adding space to a group with the foreseen dcache admin command. In previous versions this operation was working without problems. In 1.9.12-10 threads do not seem to terminate correctly, which has been reported as bug to the dcache developers.
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ successful TOTEM run on Friday and Saturday. Back to standard pp collisions on Saturday evening.
 - ◇ This evening there will be a high pileup test
 - ◆ CERN / central services and T0
 - ◇ LSF stuck for a few hours on Saturday morning (GGUS:83948 [↗](#)). T0 operations were affected. Backlog was recovered during the day
 - ◇ Glitches in Castor availability this morning, probably due to the (transparent) intervention
 - ◆ Tier-1/2:
 - ◇ GGUS:83486 [↗](#) (FTS delegation problem): currently no problems but keeping here until sw is fixed
- ALICE reports [↗](#) -
 - ◆ Job failures due to all EOS-ALICE diskservers going offline; as a cure the EOS team restarted the central node around 13:00.
 - ◆ Global failure on the transfers from Point 2 to the Computer Center since ~14:20. DAQ sent an operator alarm for CASTOR. Apparently the headnode was running out of threads, hence having a flapping behavior. The CASTOR team are working on it.
 - ◆ Xavier: investigating CASTOR issue: seem to run out of threads, but not clear yet why. ongoing...
- LHCb reports -
 - ◆ Users analysis and Reconstruction at T1s
 - ◆ MC production at T2s
 - ◆ T0:
 - ◆ T1:
 - ◇ IN2P3: spikes of failed production jobs with fixed frequency (24h, at around 21:00 every day). Jobs fails with the message "Job has reached the CPU limit of the queue" (GGUS:83985 [↗](#)). Under investigation.

Sites / Services round table:

- From Onno/NL-T1 by email:
 - ◆ For debugging the SARA SRM instability, dCache developer Dmitry had provided a custom built BoneCP library with a hardcoded Postgresql connection limit of 120 instead of 30. On Friday night we installed it. Since then the SrmSpaceManager works fine, without hanging.
 - ◆ dCache developer Gerd Behrmann found in the dCache code a possible deadlock bug: when a certain operation that opens a DB connection, calls another routine that tries to open a second connection. When there are 29 connections already open, the first new connection succeeds but the second fails, leaving the SrmSpaceManager hanging. A high load triggers this deadlock since it makes queries slower and connections last longer.
 - ◆ Second issue: this weekend the SARA compute cluster scheduler crashed because jobs got

stuck in an infinite retry loop in a test queue which was not properly configured. We've removed the queue and changed the config to have such jobs expired.

- Michael/BNL - ntr
- Boris/NDGF - downtime today for ongoing upgrade of central srm
- Kyle/OSG - question about GGUS ticket switch. MariaDZ: new soap interface not yet included in this GGUS release.
- Jhen-Wei/ASGC - ntr
- Paolo/CNAF - GPFS (ATLAS) intervention ongoing - queues closed yesterday for ATLAS
- Lisa/FNAL - ntr
- Rolf/IN2P3 - ntr
- Tiju/RAL - ntr
- Dimitri/KIT - ntr

- GGUS: (MariaDZ)
 - ◆ Please note the Did you know?...[☞](#) text with today's GGUS release (it concerns the ticket category "Change Request").
 - ◆ The SNOW weekly release introduced a new functionality (no relation with GGUS) that didn't accept ticket updates. Following <https://cern.service-now.com/service-portal/view-incident.do?n=INC145017>[☞](#) this functionality was disabled on the SNOW side but a GGUS cron job was still trying to push all the updates that were failing since this morning and until then. We shall rediscuss this workflow with SNOW developers. More at the daily WLCG meeting shortly.
 - ◇ Maarten: should check if this problem may also apply to other GGUS back-end systems
 - ◆ File ggus-tickets.xls is up-to-date and attached to page WLCGOperationsMeetings. We had 5 real ALARMS last week, a total of 12 since the last MB, 2 weeks left until the next one.
- DOEgrids: (MariaDZ) This is for OSG: Would you please process host requests in <https://pki1.doegrids.org:8100/ca/processReq?seqNum=88906>[☞](#) and 88907, 88908. Now they are assigned to the right VO OSG:CMS,
 - ◆ Kyle: will follow up..

AOB:

Tuesday

Attendance: local(Luc/ATLAS, Giuseppe/CMS, Xavier, Eddy, Dirk);remote(Stefano/LHCb, Boris/NDGF, Michael/BNL, Gonzalo/PIC, Kyle/OSG, Xavier/KIT, Jhen-Wei/ASGC, Lisa/FNAL, Tiju/RAL, Rolf/IN2P3, JT/NL-T1, Paolo/CNAF).

Experiments round table:

- ATLAS reports -
 - ◆ CERN CENTRAL SERVICES, T0
 - ◇ NTR
 - ◆ T1
 - ◇ INFN-T1 downtime ongoing. Out of T0 export.
 - ◇ PIC excluded from T0 export in prevision of tomorrow downtime.
 - ◆ T2
 - ◇ NTR
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ Three High Pileup fills overnight: max pileup was 66
 - ◇ MD until Wed 16, then TOTEM physics fill
 - ◆ CERN / central services and T0

- ◇ Tier0 crashing on early express files from high-PU fill, some module looking for trigger results and not finding it
- ◇ Experts analyzing
- ◆ Tier-1/2:
 - ◇ GGUS:83486 (FTS delegation problem): currently no problems but keeping here until sw is fixed
- ALICE reports -
 - ◆ NTR
- LHCb reports -
 - ◆ Normal activity at Tier1 (users analysis and Reconstruction)
 - ◆ No MC production running
 - ◆ T0:
 - ◇ Problems in order to access the VO boxes at CERN. Get in touch with Arne Wiebalck and Benedikt Hegner: yesterday some readonly AFS volume has been accidentally deleted. As there are no AFS callbacks for deleted readonlys, the machines still try to access and get stuck. We are following the issue with helpdesk.
 - ◆ T1:
 - ◇ PIC: downtime started (until friday morning) so we banned the site and related SEs

Sites / Services round table:

- Boris/NDGF - ntr
- Michael/BNL - ntr
- Gonzalo/PIC - ntr
- Kyle/OSG - acknowledged answer from Maria below
- Xavier/KIT - ntr
- Jhen-Wei/ASGC - ntr
- Lisa/FNAL - ntr
- Tiju/RAL -ntr
- Rolf/IN2P3 - ntr
- JT/NL-T1 - ntr
- Paolo/CNAF - ntr
- GGUS: (MariaDZ) Savannah:130205 is the ticket for the workflow revisit due to yesterday's side-effects of SNOW failing to accept ticket updates in the morning. To answer Kyle's question the new GGUS SOAP interface is for a "Not yet scheduled" release and Soichi is in the loop indeed as he submitted Savannah:127763#comment6.

AOB:

- Xavier: Increased number of threads in ALICE Castor in response to yesterday's problems. EOS: still investigating messaging problem which lead to problem for ALICE yesterday.

Wednesday

Attendance: local(Luc, Giuseppe, Luca, Xavier, Eddy, Ignacio, Dirk);remote(Boris/NDGF, Stefano/LHCb, John/RAL, Kyle/OSG, Jhen-Wei/ASGC, Lisa/FNAL, Pavel/KIT, Rolf/IN2P3, Gonzalo/PIC, Michael/BNL, Ron/NL-T1).

Experiments round table:

- ATLAS reports -

- ◆ CERN CENTRAL SERVICES, T0
 - ◇ CERN-PROD_LOCALGROUPDISK full. Blacklisted. savannah #130248
- ◆ T1
 - ◇ INFN-T1 fully back in business after downtime.
 - ◇ PIC ongoing downtime.
 - ◇ RAL. Disk lost in Castor. Recovery process ongoing.
- ◆ CALIB_T2
 - ◇ INFN-ROMA1 Downtime due to Unexpected network outage.
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ MD, TOTEM physics fill in the afternoon
 - ◆ CERN / central services and T0
 - ◇ NTR
 - ◆ Tier-1/2:
 - ◇ GGUS:83486 (FTS delegation problem): currently no problems but keeping here until sw is fixed
- ALICE reports -
 - ◆ NTR
- LHCb reports -
 - ◆ Normal activity at Tier1s (users analysis and Reconstruction)
 - ◆ Few new MC production launched today and running at Tier2s
 - ◆ T0:
 - ◇ Some aborting pilots yesterday due to an AFS volume get corrupted -> now fixed
 - ◇ Today lots of failed pilots (GGUS Ticket: https://ggus.eu/ws/ticket_info.php?ticket=84126): on various WN the home directory were set in read-only mode. Now fixed and people said they are still investigating why it happened.
 - ◆ T1:

Sites / Services round table:

- Boris/NDGF - ntr
- John/RAL - failed disk server - looking into it
- Kyle/OSG - ntr
- Jhen-Wei/ASGC - ntr
- Lisa/FNAL - ntr
- Pavel/KIT - ntr
- Rolf/IN2P3 - ntr
- Gonzalo/PIC - ntr
- Michael/BNL - ntr
- Ron/NL-T1 - ntr

AOB:

Thursday

Attendance: local(Luc, Maarten, Eddy, Giuseppe, Massimo, Alex, MariaDZ, Dirk);remote(Boris/NDGF, Stefano/LHCb, Jhen-Wei/ASGC, Paolo/CNAF, Ronald/NL-T1, Rolf/IN2P3, Gareth/RAL, Lisa/FNAL, Kyle/OSG).

Experiments round table:

Thursday

- ATLAS reports -
 - ◆ CERN CENTRAL SERVICES, T0
 - ◇ CERN-PROD_LOCALGROUPDISK full. Blacklisted. savannah #130248. CAT team contacted for cleanup.
 - ◇ CERN FTS problem (Transfer T0->RAL hanging) GGUS:84154 [↗](#), INC:145995 [↗](#).
 - ◇ ContZole monitor problem in data quality transformation bug#95989. Understood.
 - ◆ T1
 - ◇ INFN-T1 transfer failures GGUS:84152 [↗](#) (problem on GPFS). Stable since this morning.
 - ◇ PIC ongoing downtime.
 - ◇ RAL. Disk lost in Castor. Recovery process ongoing.
 - ◆ Gareth: disk server was recovered after triple failure and went back in production a few hours ago. Data checksums have been confirmed. There should be no data loss.
 - ◇ Maarten: was the root cause for the problems understood? Gareth: this was an older disk server soon up to retirement. In fact multiple disk failures during recovery are not that rare due to the recovery load.
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ LHC is trying to make a successful fill for TOTEM, planned for today
 - ◆ CERN / central services and T0
 - ◇ NTR
 - ◆ Tier-1/2:
 - ◇ PIC is down until Friday for dCache update
 - ◇ GGUS:83486 [↗](#) (FTS delegation problem): currently no problems but keeping here until sw is fixed
 - ◆ Massimo: Is Totem data in a different pool or somehow distinguishable for the service? Giuseppe: Same basic flow, but identified by different dataset names.
- ALICE reports [↗](#) -
 - ◆ CERN: high, inefficient tape reading activity observed this morning; being investigated by the tape experts; ALICE submitted all requests at the same time as usual (to allow for the tape queue to be optimized).
- LHCb reports -
 - ◆ Normal activity at Tier1s (users analysis and Reconstruction)
 - ◆ MC production running at Tier2s
 - ◆ T0:
 - ◇ Still failing pilots (GGUS Ticket: https://ggus.eu/ws/ticket_info.php?ticket=84126 [↗](#)) still investigating what is changing the permission of home directories. Some suspect to be investigated.
 - ◆ T1:

Sites / Services round table:

- Boris/NDGF - ntr
- Jhen-Wei/ASGC - ntr
- Paolo/CNAF - transfers and storm problems still under investigation (see tickets)
- Ronald/NL-T1 - early afternoon nikhef DPM has been unavailable due to interference from a unexperienced user who has been informed. Now back to normal.
- Rolf/IN2P3 - ntr
- Gareth/RAL - ntr
- Lisa/FNAL - ntr
- Kyle/OSG - ntr

- Alex: WMS update 16 (urgent security update) will take place this afternoon
- GGUS: (MariaDZ) On holiday as of next week for a month. The developers can be contacted via GGUS tickets. In the rare case when the interface is down, tickets can be created via email to helpdesk@ggusNOSPAMPLEASE.eu. Moreover, GGUS is now integrated into the KIT on-call service as per Savannah:113831#comment46 [↗](#)
 - ◆ new soap interface: aim for september release

AOB:

Friday

Attendance: local(Luc, Guiseppe, Luca, Xavier, Eddy, Dirk);remote(Michael/BNL, Saverio, Andreas+Xavier/KIT, Stefano/LHCb, Boris/NDGF, Lisa/FNAL, Onno/NL-T1, Jhen-Wei/ASGC, Jeremy/GridPP, Paolo/CNAF, Kyle/OSG, Rolf/IN2P3).

Experiments round table:

- ATLAS reports -
 - ◆ CERN CENTRAL SERVICES, T0
 - ◇ CERN-PROD_LOCALGROUPDISK full. savannah #130248. 200 TB/500T free. Solved.
 - ◇ CERN FTS problem (Transfer T0->RAL hanging) GGUS:84154 [↗](#), INC:145995 [↗](#). No reply yet
 - ◆ T1
 - ◇ INFN-T1 transfers stable GGUS:84152 [↗](#).
 - ◇ PIC ongoing downtime to be finished at 16:00 UTC.
 - ◇ NDGF-T1 transfer failures to MCTAPE due to staging problem GGUS:84207 [↗](#)
 - ◆ CALIB_T2
 - ◇ NTR
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ TOTEM fill yesterday, 0.058 pb-1, CMS efficiency 96.87%
 - ◇ Today high-lumi fill, Van der Meer scan foreseen on Monday night
 - ◆ CERN / central services and T0
 - ◇ T0 is processing now the High Pileup runs taken at the beginning of the week
 - ◆ Tier-1/2:
 - ◇ PIC is down until Friday for dCache update
 - ◇ GGUS:83486 [↗](#) (FTS delegation problem): currently no problems but keeping here until sw is fixed
- ALICE reports [↗](#) -
 - ◆ CERN: yesterday's tape issue was solved by a configuration fix on the CASTOR side.
 - ◆ Andreas/KIT: Alice jobs are since 48h pulling 3 GB/s into WN (most data from IN2P3) and risk to overload firewall. Alice is informed and promised yesterday to reduce number of jobs, but so far no improvement visible. Please reduce job number soon
 - ◆ Rolf/IN2P3: current 30% ALICE at IN2P3 are very short lived. As this is unusual for ALICE: please check
- LHCb reports -
 - ◆ Normal activity at Tier1s and Tier2s
 - ◆ T0:
 - ◆ T1:

Sites / Services round table:

- Michael/BNL - ntr
- Andreas+Xavier/KIT - ntr
- Boris/NDGF - inquired about staging problem, but no news yet from the site
- Lisa/FNAL - ntr
- Onno/NL-T1 - ntr
- Jhen-Wei/ASGC - ntr
- Jeremy/GridPP - ntr
- Paolo/CNAF - ntr
- Rolf/IN2P3 - ntr
- Kyle/OSG - ntr

AOB:

- SIR received from IN2P3 about shared area problem on July 24th - linked in service incident twiki (link above)
- PIC outage is scheduled finish this afternoon - as we are close to the weekend the experiment contacts would appreciate a brief update of the status

-- JamieShiers - 09-Jul-2012

This topic: LCG > WLCGDailyMeetingsWeek120709

Topic revision: r28 - 2012-07-13 - MaartenLitmaath



Copyright &© 2008-2021 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