

Table of Contents

Week of 120806.....	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.....	6
Friday.....	7

Week of 120806

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 (AndreaV, Alexandre, Ken, Marc, LucaM, Philippe, Alessandro, Eva); remote (Ulf/NDGF, Kyle/OSG, Michael/BNL, Jhen-Wei/ASGC, John/RAL, Alexander/NLT1, Marc/IN2P3, Dimitri/KIT, Burt/FNAL).

Experiments round table:

- ATLAS reports -
 - ◆ T1s:
 - ◇ PIC: Issue exporting T0 data to PIC during the night. GGUS:84833. Solved this morning. From PIC: "We had a huge load on ATLAS pools due to data replication to a new hardware. It was solved more than one hour ago but we'll keep the ticket opened a couple of hours until we are sure that current load won't affect data transfers anymore."
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ Not much data this weekend due to series of unfortunate events.
 - ◆ CERN / central services and T0
 - ◇ Had to delay prompt reconstruction of some recent runs because of a problem updating databases that is still being investigated.

- ◇ Some files have been inaccessible from Castor, see INC:151642 [↗](#). [LucaM: files were inaccessible from EOS (not Castor) due to two machines down, now fixed.]
- ◆ Tier-1/2:
 - ◇ ASGC is having downtime today to fix Castor problems. Currently have open tickets for Hammercloud (GGUS:84658 [↗](#)), SUM test (GGUS:84632 [↗](#)), and MC production (SAV:130787 [↗](#)).
- ◆ Other:
 - ◇ Daniele Bonacorsi is the next CRC.
- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ Ongoing issues with RAW export to Gridka. Still many files in "Ready" status. Very low transfer rate. Latest updates on tickets say that CERN see SRM timeouts? (GGUS:84550 [↗](#), GGUS:84778 [↗](#) (Alarm), GGUS:84670 [↗](#)).
 - ◇ [Dimitri/KIT: problems ongoing with connectivity to tape, experts are following up. Marc: could these issues also explain the problem with low space on tape cache? Dimitri: not sure, will ask the experts.]
 - ◇ [LucaM: will do some debugging from CERN too, is this issue also seen at other sites? Marc: no, this is only Gridka.]
 - ◆ T1:
 - ◇ GridKa: Very low space left on GridKa-Tape cache. Possibly related related to above. Ticket opened (GGUS:84838 [↗](#))
 - ◇ IN2P3: Investigating why pilots aren't using CVMFS both here and at the Tier2. This caused some job failures last week.

Sites / Services round table:

- Ulf/NDGF: Alice is writing a lot of data and we are having troubles coping with the data rates. Tomorrow the situation may get worse because of a network intervention, we will only have 2/3 of the writing capacity for three hours.]
- Kyle/OSG: ntr
- Michael/BNL: ntr
- Jhen-Wei/ASGC: Castor downtime is ongoing, scheduled until 4pm UTC (6pm CET). Ale: will switch on a few transfers to Taiwan tonight but will wait for tomorrow morning before switching on T0 exports to Taiwan, after we are sure that everything is back to normal.]
- John/RAL: ntr
- Alexander/NLT1: ntr
- Marc/IN2P3: ntr
- Dimitri/KIT: nta
- Burt/FNAL: ntr

- Philippe/Grid: ntr
- Eva/Databases: ntr
- LucaM/Storage: nta
- Alexandre/Dashboard: SRM voput tests failing for CMS/ASGC. [Jhen-Wei: due to the ongoing Castor intervention.]

AOB: none

Tuesday

Attendance: local (Andrea, Alexandre, Alessandro, Mark, Eva, Philippe); remote (Kyle/OSG, Ulf/NDGF, John/RAL, Xavier/KIT, Marc/IN2P3, Jhen-Wei/ASGC, Michael/BNL, Burt/FNAL, Ronald/NLT1; Daniele/CMS).

Experiments round table:

- ATLAS reports -
 - ◆ T0/1s:
 - ◇ CERN: ALARM ticket slow response of LSF GGUS:84928 [↗](#). "Since 22 on average takes 10s to LSF to answer, with peaks of 3 mins". Problem disappeared after 2hours.
 - [Philippe/LSF: slowdown yesterday evening because LSF hit machines that do not exist in LANDB/CDB/DNS. Ulrich is looking into this now.
 - Alessandro: we are seeing issues with LSF since several weeks, it would be nice to add some monitoring and checks in LSF to block client IPs from submitting too many jobs if they are not recognized as production users (normal users are often submitting too many jobs and affecting production without realizing). Philippe: thanks for this suggestion, will pass it on to LSF experts as temporary solution; however, as discussed two weeks ago, LSF is hitting its limits and we are investigating a possible replacement by an alternative to LSF, as a longer term solution.]
 - ◇ PIC: same issue of yesterday, ticket updated GGUS:84833 [↗](#)
 - ◇ [Alessandro: T0 export to ASGC still disabled, will enable it tomorrow morning.]
- CMS reports -
 - ◆ LHC / CMS
 - ◇ pp physics collisions
 - ◆ CERN / central services and T0
 - ◇ T0 Ops reported problems in accessing 2 files from EOS. First investigation showed that they have only one copy and they are on a faulty diskserver (waiting for a vendor call). Quickly files were made available again, thanks (INC:152249 [↗](#)).
 - **We need a stable solution here. Loosing files in this way is impacting T0 Ops. BTW, apart from this specific case (1 replica only should never happen) can we increase the replication factor from 2 to 3?**
 - [Alessandro: you can probably increase yourslef the replication factor from 2 to 3, but will then be charged a factor 1.5 more.
 - Daniele: thanks, yes we could probably do it ourselves, but we want to discuss and negotiate this with the EOS team first.]
 - ◇ T0_CH_CERN: HC errors, late July ticket, now green, no ticket update though (Savannah:130709 [↗](#) bridged to GGUS:84617 [↗](#), HC 2nd line support)
 - ◆ Tier-1:
 - ◇ busy with a data reprocessing campaign (in the tails now), plus MC prod
 - ◇ ASGC: out of the extended downtime. Felix and colleagues kept CMS constantly informed, much appreciated. Will need to deal with pending open tickets later (for reference: Savannah:130054 [↗](#), Savannah:130589 [↗](#), Savannah:130727 [↗](#) -> GGUS:84632 [↗](#), Savannah:130754 [↗](#) -> GGUS:84658 [↗](#), Savannah:130768 [↗](#) -> GGUS:84677 [↗](#), Savannah:130787 [↗](#), Savannah:130880 [↗](#))
 - ◇ IN2P3: low transfer quality to several sites, pending a reply, bridged to GGUS (Savannah:130877 [↗](#) -> bridged to GGUS:84947 [↗](#)). [Marc/IN2P3: someone is looking at this, transfers already look better now.]
 - ◆ Other:
 - ◇ MC prod and analysis smooth on T2 sites

- ALICE reports [↗](#) -
 - ◆ NDGF: experts at CERN and NDGF are looking into the unexpectedly high write rates reported yesterday and discovered a few issues that are being followed up.
 - ◇ [Ulf/NDGF: we are working on this and found the probable cause. We are realizing however that there has been a communication problem with Alice. It is unfortunate that Alice had already seen this type of error but never reported it to us. The Alice system tries over and over again to write tapes when there are problems, so now our tapes are full of files that are not usable or not necessary. We are in contact with Alice to improve this situation.]

- LHCb reports -
 - ◆ T1:
 - ◇ Ongoing issues with RAW export to Gridka. Still many files in "Ready" status. Very low transfer rate. GridKa still investigating (GGUS:84550 [↗](#), GGUS:84778 [↗](#) (Alarm), GGUS:84670 [↗](#)).
 - ◇ GridKa have added 40TB of front end disk to the Tape system to allow transfers to come through (GGUS:84838 [↗](#)). [Xavier/KIT: actually added 80 TB, not 40 TB. We cannot really tell however what improved the situation, we have experimented with many changes and one worked, but we are now trying to understand which one was the relevant one. We will leave the setup as it is now, so that LHCb can recover from the backlog.]
 - ◇ IN2P3: Re: Pilots not using CVMFS - there was confusion as CVMFS is in /opt at IN2P3 rather than /cvmfs. The problem turned out to be an LHCb config issue.

Sites / Services round table:

- Kyle/OSG: ntr
- Ulf/NDGF:
 - ◆ mainly working on the ALICE issue
 - ◆ downtime at 15h UTC today (OPN to Finland down)
 - ◆ one dcache upgrade in Norway tomorrow
- John/RAL: ntr
- Xavier/KIT: nta
- Marc/IN2P3: nta
- Jhen-Wei/ASGC:
 - ◆ downtime was necessary this morning at 1030 to complete the Castor intervention started yesterday. Database was moved to a new storage; Castor was upgraded to 1.11-9; xrootd was switched on. Transfers look stable since this morning, so we are in production again.
- Michael/BNL: ntr
- Burt/FNAL: ntr
- Ronald/NLT1: ntr

- Philippe/Grid: BDII was restarted this morning to fix issues with BNL and FNAL
- Alexandre/Dashboard: ntr
- Eva/Databases: ntr

AOB:

- Maintenance work on KIT internet connection (originally planned for August 20) is postponed to September 17th same time slot (see also GOCDB:115602 [↗](#)).

Wednesday

Attendance: local (Massimo, Alexandre, Alessandro, Mark, Eva, Philippe, Alexandre, Daniele); remote (Kyle/OSG, Ulf/NDGF, John/RAL, Pavel/KIT, Marc/IN2P3, Jhen-Wei/ASGC, Michael/BNL, Burt/FNAL, Ronald/NLT1;).

Experiments round table:

- ATLAS reports -
 - ◆ T0/1s:
 - ◇ ADCR DB instance 2 (Panda) showed since yesterday high read value (order of hundreds of MB/s, while the average of the past month is 20/40 MB/s). ATLAS DBA and Panda experts are in the loop, it seems that PandaMonitor is the application creating the load.
 - ◇ TAIWAN-LCG2 is now working fine, it has been reincluded in T0 data export.
- CMS reports -
 - ◆ LHC / CMS
 - ◇ pp physics collisions
 - ◆ CERN / central services and T0
 - ◇ **EOSCMS unavailability**: after yesterday's outage (resolved at ~4:30 CERN time), CMS T0 Ops noticed a severe T0 degradation as from SLS, main source seems to be EOSCMS, down since midnight. CMS express jobs stuck since ~1 AM (jobs crucial to provide calibration constants for the current data taking run). T0 basically stopped because we cannot stageout files. Opened a GGUS ALARM ticket (GGUS:84966), being worked on.
 - ◆ Tier-1:
 - ◇ Updates:
 - ASGC: See yesterday. Still giving them some time after resuming operations before we go again through open tickets in detail (for reference: Savannah:130054, Savannah:130589, Savannah:130727 -> GGUS:84632, Savannah:130754 -> GGUS:84658, Savannah:130768 -> GGUS:84677, Savannah:130787, Savannah:130880)
 - IN2P3: low transfer quality to several sites (Savannah:130877 -> bridged to GGUS:84947). We can live with the lower quality in some channels, they would need more network-investigation to improve further. GGUS solved.
 - ◆ Other:
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ CERN: yesterday evening EOS-ALICE came back in production after an upgrade to v0.2, thanks! It now provides a feature that is needed for managing a separate name space within EOS, urgently needed by ALICE to allow also the calibration data to be hosted in EOS: ALICE are looking forward to the deployment of that name space, thanks!
- LHCb reports -
 - ◆ **New GGUS (or RT) tickets**
 - ◆ T1:
 - ◇ GridKa transfers stable overnight. Will keep the tickets open for another 24 hours just to be sure (GGUS:84550, GGUS:84778 (Alarm), GGUS:84670).
 - ◆ Other :

Sites / Services round table:

- ASGC: ntr
- BNL: ntr
- FNAL: ntr
- IN2P3: ntr
- KIT:
 - ◆ LHCb problem (tape reading) gone but not yet understood in detail
 - ◆ LHCb observation (~10 "slow" CPU nodes failing jobs after days of running) being investigated
- NDGF: ntr
- NLT1: ntr
- RAL: ntr
- OSG: recalc availability requested by ticket

- CASTOR/EOS:
 - ◆ EOSCMS fixed by restart. Invertigating the root cause
 - ◆ EOSALICE upgraded (0.2.8)
 - ◆ CASTORALICE: one corrupted file detected. Experiment notified
- Central Services: ntr
- Data bases: CMS@P5 called DB piquet but it was a false alarms (new CMS monitoring in place but needs some tuning)
- Dashboard: ntr

AOB:

Thursday

Attendance: local (Andrea, Alessandro, Alexandre, Cedric, Ulrich, Eva, Mark, LucaM); remote (Marc/IN2P3, Michale/BNL, Woojin/KIT, Gareth/RAL, Jhen-Wei/ASGC, Burt/FNAL, Ulf/NDGF, Ronald/NLT1, Kyle/OSG, Jeremy/GridPP).

Experiments round table:

- ATLAS reports -
 - ◆ T0/1s:
 - ◇ CERN: ALARM ticket of 2 days ago still open, ATLAS still observe LSF slowness GGUS:84928 [\[Ulrich: this is a long standing issue, we are investigating this with the vendor but there is little else we can do. Please check with Gavin who suggested a possible workaround involving parallel submissions. Alessandro: parallel submissions may not be optimal for the T0, where they need to submit some bunches of jobs in a given order. Ulrich: unfortunately there is not much else that can be done now. Alessandro: ok we'll check with Gavin.\]](#)

- CMS reports -
 - ◆ LHC / CMS
 - ◇ pp physics collisions
 - ◆ CERN / central services and T0
 - ◇ EOSCMS: Stephen adjusted replication factor 2->3 for /eos/cms/store/t0temp [Daniele: we discussed this with Jan Iven, we understand that increasing the replication factor should not require a large storage capacity increase.]
 - ◆ Tier-1/2:
 - ◇ NTR
 - ◆ Other:
 - ◇ NTR

- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ T1:
 - ◇ GridKa still stable (GGUS:84550[↗](#), GGUS:84778[↗](#) (Alarm), GGUS:84670[↗](#)). *
[Mark: forgot to mention, the slow computers reported yesterday have been excluded from the Gridka pools, Philippe is in contact with Gridka about this.]
 - ◇ IN2P3: Question about short pilots - due to development testing [Mark: other sites may expect similar issues due to these tests by Dirac developers]

Sites / Services round table:

- Marc/IN2P3: ntr
- Michael/BNL: ntr
- Woojin/KIT: ntr
- Gareth/RAL: had a power failure on an ATLAS software server at 5am, fixed by 10am
- Jhen-Wei/ASGC: ntr
- Burt/FNAL: ntr
- Ulf/NDGF: ntr
- Ronald/NLT1: ntr
- Kyle/OSG: ntr
- Jeremy/GridPP: ntr
- Ulrich/Grid: nta
- Eva/Databases: had a big issue with some databases at noon today, affecting LCGR, CMS offline, online and Active Data Guard. Someone run a parallel query on LCGR that completely blocked the LCGR database and also the storage below (shared with the CMS databases). One LCGR instance went down, the storage and other databases were unusable because of the load. This is being investigated to find out who did this.
 - ◆ [Alessandro: can this be related to the performance issues seen by the dashboard this week? Eva: it is the same LCGR database, but most likely these are unrelated issues.]
 - ◆ [Daniele: how did LCGR affect CMS? Eva: the storage is shared and the storage was blocked by the load on LCGR.]
- Alexandre/Dashboard: ntr
- LucaM/Storage:
 - ◆ the database behind Alice Castor was updated today
 - ◆ we had an issue with one of the new LHCb disk servers, we are now trying to recover files and we are in contact with LHCb as some files may be lost. [Mark: was not aware as there is no GGUS ticket. Andrea: please open a GGUS ticket for this.]

AOB: none

Friday

Attendance: local (AndreaV, Rod, Cedric, Alexandre, Mark, Philippe, Eva, LucaM); remote (Michael/BNL, Ulf/NDGF, Alexander/NLT1, Burt/FNAL, Jhen-Wei/ASGC, Kle/OSG, Marc/IN2P3, Tiju/RAL).

Experiments round table:

- ATLAS reports -
 - ◆ T0/1s:
 - ◇ NDGF : Tape recall issue : GGUS:85037[↗](#)

- CMS reports -
 - ◆ LHC / CMS
 - ◇ pp physics collisions
 - ◆ CERN / central services and T0
 - ◇ EOSCMS stable. Nothing else to report.
 - ◆ Tier-1/2:
 - ◇ Only minor Savannahs at the T2 level, nothing major at the T1 level
 - ◆ Other:
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ T0:
 - ◇ A 'background' of aborted pilots. Jobs are still running though. CREAM losing contact with LSF? [Mark: is this connected to the ATLAS overload problem reported this week? Philippe: will investigate.]
 - ◆ T1:
 - ◇ GridKa:
 - still stable - will close the tickets this afternoon (GGUS:84550 [↗](#), GGUS:84778 [↗](#) (Alarm), GGUS:84670 [↗](#)).
 - Slow computers (GGUS:84988 [↗](#)). Seems like an odd clock speed problem.
 - This morning at ~11am SRM authentication issues seen. Quickly fixed.
 - ◇ CNAF:
 - What look like SRM errors stopping both FTS and internal transfers (GGUS:85041 [↗](#))

Sites / Services round table:

- Michael/BNL: ntr
- Ulf/NDGF: tape issue tonight was fixed in the morning
- Alexander/NLT1: downtime on August 27 for tape maintenance (is in GOCDB)
- Burt/FNAL: ntr
- Jhen-Wei/ASGC: ntr
- Kyle/OSG: problems with USCMS reporting mentioned yesterday are now fixed
- Marc/IN2P3: ntr
- Tiju/RAL: ntr

- Eva/Databases: ntr
- Luca/Storage: ntr
- Philippe/Grid: ntr
- Alexandre/Dashboard: ntr

AOB:

- LucaM: is the CMS magnet down and will this affect the physics schedule? If there is no physics and the schedule is changed, we could take the opportunity to do some Castor upgrades without waiting for the technical stop. Daniele: correct, and the schedule is being discussed at the level of the CMS Physics management. In case some input comes over the weekend, we will send an email to the SCOD and the Castor Ops list over the weekend, anyway we will report at the Monday meeting. Beyond this case, if you need to contact the CMS people attending the WLCG Ops calls also over the weekend (when there is no daily WLCG call), feel free to use the CMS-CRC list <cms-crc-on-duty@cern.ch>: it will reach the Computing Run Coordinator on-duty.

-- JamieShiers - 02-Jul-2012

This topic: LCG > WLCGDailyMeetingsWeek120806

Topic revision: r16 - 2012-08-10 - AndreaValassi



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](#)