

Table of Contents

Week of 121119.....	1
Daily WLCG Operations Call details.....	1
WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments.....	1
General Information.....	1
Monday.....	1
Tuesday.....	3
Wednesday.....	4
Thursday.....	6
Friday.....	7

Week of 121119

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(Simone, Maarten, AlexeyS, AleDiGGi, JeromeB, Wei-Jen ASGC, Mascetti, Guido, Dawid, Mike, Maria), Remote(Michael BNL, Philippe LHCb, Gonzalo PIC, Onno NL-T1, Lisa FNAL, Tiju RAL, Christian NDGF, Alexander NL-T1, Pavel KIT, Ian Fisk CMS, Kyle OSG, Rolf IN2P3, Stefano CNAF)

Experiments round table:

- ATLAS reports -
 - ◆ ATLAS General
 - ◇ FTS3 tests: ATLAS has now started testing the UK FTS3 dev (<https://lcgfts-dev.gridpp.rl.ac.uk:8443/glite-data-transfer-fts/services/FileTransfer>) instance for Functional Test. For now UKI-SOUTHGRID-RALPP_DATADISK is using it, today QMUL will be added. More news soon.
 - ◇ Reprocessing still ongoing, today deletion of already validated ESD and AOD has started.
 - ◆ T0
 - ◇ Contention on ADCR due to repeated attempts to read redo blocks from Oracle's redo log files. Reason: logical corruption in the first partition of a Primary Key's index. index partition rebuilt, problem fixed
 - ◇ Castoratlas upgrade: green light on the proposed date/time (Thu 1pm)
 - ◆ T1

- ◊ Reduced load on KIT DATATAPE from friday afternoon
- ◆ T2
 - ◊ UKI-SCOTGRID-GLASGOW_DATADISK: problem with disk servers, token blacklisted for uploading/writing
- CMS reports -
 - ◆ LHC / CMS
 - ◊ NTR
 - ◆ CERN / central services and T0
 - ◊ Tier-0 will continue to run at near maximum capacity for the next month
 - ◆ Tier-1:
 - ◊ We've overloaded KIT tape. Balance between CMS and the local operators
 - ◊ Problems with transfers to and from IN2P3
 - https://ggus.eu/ws/ticket_info.php?ticket=88597, CMS waiting for a reply.
 - Maria D: is there an issue with the ticketing system per-se (Savannah-GGUS bridge)? Ian: no.
 - ◆ Tier-2:
 - ◊ NTR
- ALICE reports [↗](#) -
 - ◆ KIT: Fri late afternoon all new CREAM CEs reported 444444 jobs waiting for ALICE, GGUS:88557 [↗](#) opened, fixed Fri evening
- LHCb reports -
 - ◆ Reprocessing: new set of runs launched during the WE (55000 files)
 - ◆ Prompt reconstruction: 7,000 files waiting for reconstruction after this WE good LHC performance (CERN + few Tier2s)
 - ◆ MC productions at T2s and T1s if resources available
 - ◆ **New GGUS (or RT) tickets** [↗](#)
 - ◆ T0:
 - ◆ T1:
 - ◊ GridKa: still having significant problems with FTS transfers to GridKa. Investigations still ongoing. (GGUS:88425 [↗](#)). Error:
`globus_gass_copy_register_url_to_url transfer timed out. The timeouts are related neither to a source (happens from CERN + 5 Tier1s) nor to a gridftp server (same file succeeds to same server a few minutes later).`
 - Normally gridFTP v2 should be the default. If v1 is chosen, this already indicate an issue. As hint, the configuration of KIT disk servers negotiating v1 should be investigated.
 - ◊ NL-T1: spikes of data upload failures from NIKHEF to SARA (very short duration), coincides with spikes of uploads (SRM overload?)
 - NL-T1: in case something needs investigation a GGUS ticked should be filled. Philippe: no need to investigate now.
 - ◆ Philippe: DBII publishes various "default" values when some information can not be properly published (e.g. 4444 or 9999). Can this be more consistent? Maarten: various defaults mean different things and there is a logic behind it (normally to discourage clients using that service). The site should be ticketed if those values appear.

Sites / Services round table:

- ASGC: ntr
- BNL: ntr
- FNAL: ntr
- IN2P3: ntr

- KIT: ntr
 - NDGF: ntr
 - NLT1: ntr
 - PIC: ntr
 - RAL: ntr
 - OSG: ntr
-
- Storage/Grid Services/Databases/Dashboards: nrt
-
- GGUS: File ggus-tickets.xls is up-to-date and attached to page WLCGOperationsMeetings. Final ALARM drills are attached to this page for tomorrow's MB. There were 8 real ALARMS in the last 5 weeks.
 - ◆ Maria D: There is an ATLAS alarm on LSF@CERN open since long time (weeks). AleDiGGI: People are working actively on it, but the issue is still there. May be will close as unsolved and will reopen at the next re-occurrence.

AOB: none

Tuesday

Attendance: Local(Simone, GiuseppeB CMS, Maarten, AleDiGGi, ASedov, Wei-Jen ASGC, Jerome, Mike, LucaM), Remote(Michael BNL, Xavier KIT, Stefano CNAF, Lisa FNAL, Kyle OSG, Philippe LHCb, Tiju RAL, Anders NDGF, Ronald NL-T1, Jeremy GridPP)

Experiments round table:

- ATLAS reports -
 - ◆ ATLAS General
 - ◇ data taking -- SPS magnet need to be changed, stop for 24h.
 - ◇ Reprocessing: data deletion has started.
 - ◆ T0/T1s:
 - ◇ INFN-T1 transfer failing due to filesystem full, but SRM still publishing free space . GGUS:88553 [↗](#)
 - Stefano: CNAF will both add space this week and improve the configuration to report a proper error. Details in GGUS.
- CMS reports -
 - ◆ LHC / CMS
 - ◇ NTR
 - ◆ CERN / central services and T0
 - ◇ Tier-0 will continue to run at near maximum capacity for the next month
 - ◆ Tier-1:
 - ◇ We've overloaded KIT tape. Balance between CMS and the local operators
 - ◇ Problems with transfers to and from IN2P3 GGUS:88597 [↗](#) under investigation
 - ◆ Tier-2:
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ Reprocessing: NTR
 - ◆ Prompt reconstruction: NTR
 - ◆ MC productions at T2s and T1s if resources available

- ◆ New GGUS (or RT) tickets
- ◆ T0: problem of ghost CREAM jobs understood, but no patch available yet. It doesn't apply only to CERN, we have many tickets open on this topic. Thanks to Ulrich & Co for having dugged into that with the CREAM developers! Let's hope it is fixed soon. In the mean time sites are requested to clean up their redundant jobs...
 - ◇ Simone: does the problem affect a particular release of CREAM? Philippe: apparently not, the problem has always been there on some extent.
- ◆ T1:
 - ◇ Issue with CPU power estimate at some sites (Tier1s and Tier2s). It makes the remaining CPU work estimate by the pilot unreliable, therefore we disabled the filling mode. Still problematic at some sites even for the initial job. It is necessary to have an agreement on how to estimate the CPU power for a job slot, in particular in case hyperthreading is on. This is a (too long) standing issue. GDB?
 - ◇ Simone: the issue is being tackled by the GDB with the MACHINEFEATURES initiative. Philippe: this still does not cover some scenario.
 - ◇ Need a dedicated GDB discussion. Also, it should be brought to the attention of WLCG Ops Coordination.

Sites / Services round table:

- ASGC: FTS problems this morning. Problem due to DB backend which shut down unexpectedly. Issue resolved.
- BNL: ntr
- FNAL: ntr
- IN2P3: ntr
- KIT: ntr
- NDGF: tried to upgrade dCache but some issues were found after the process (storage stopped accepting ATLAS vomms proxies). Upgrade was rolled back.
- NLT1: ntr
- PIC: ntr
- RAL: emergency shutdown due to electrical problem. No time estimate for the end of the downtime.
- OSG: since Nov 13, queries from OSG BDII timing out against bdi214.cern.ch (OK against all other CERN BDII machines). GGUS:88527 [↗](#) under investigation.
- GridPP: ntr

- Storage: reminder about CASTOR upgrade for CMS and LHCb tomorrow (transparent).
- Grid Services/Databases/Dashboards: nrt

AOB: MariaD: thursday there will be a WLCG SCM. Please report to MariaD any important GGUS tickets which needs attention.

Wednesday

Attendance: Local(Simone, GiuseppeB CMS, Ueda, AleDiGGi, ASedov ATLAS, Jerome, LucaM , LucaC, Wei-Jen ASGC, Mike, Maarten), Remote(Michael BNL, Stefano CNAF , Philippe LHCb , Burt FNAL, Kyle OSG, Pavel KIT, Alexander NL-T1, Rolf IN2P3, Christian NDGF, Gareth RAL)

Experiments round table:

- ATLAS reports -
 - ◆ Vidyo did not work properly yesterday.
 - ◇ Simone: we need more info about this (ticket number, details, etc ...) in order to be able to escalate it.
 - ◆ T0/T1s:

- ◇ RAL is still recovering, blacklisted: would be interesting to know in details the ETA for FTS, DB (and Frontier), storage and WN. Also GOCDDB was down: what's about new downtimes?
- ◇ INFN-T1 Discrepancy between SRM published free space and real available space on disk. Site blacklisted, cloud IT brokeroff. GGUS:88553

• CMS reports -

- ◆ LHC / CMS
 - ◇ waiting for VdM scan
- ◆ CERN / central services and T0
 - ◇ Tier-0: unexpected failure in a merge this morning. Experts investigating probably a file corruption
 - ◇ had a short interruption, around noon, of the srm-cms.cern.ch SRM endpoint, probably related to the CASTOR intervention?
 - LucaM: the SRM problem was due to 2 name servers being problematic and affecting the overall SRM. This affected all VOs.
- ◆ Tier-1:
 - ◇ T1_UK_RAL still recovering from power problems, more news in the afternoon
 - ◇ Problems with transfers towards IN2P3
 - https://ggus.eu/ws/ticket_info.php?ticket=88576 under investigation, waiting for some feedback
 - Rolf: the issue is still under investigation. More news tomorrow
- ◆ Tier-2:
 - ◇ NTR

• ALICE reports -

- ◆ NTR

• LHCb reports -

- ◆ General: pilot filling mode disabled in order to avoid problems with CPU limit (until understood)
- ◆ Reprocessing: NTR
- ◆ Prompt reconstruction: NTR
- ◆ MC productions at T2s and T1s if resources available
- ◆ New GGUS (or RT) tickets
- ◆ T0: NTR
- ◆ T1: NTR... Best wishes to RAL!

Sites / Services round table:

- Tomorrow is thanksgiving. Happy turkey!
- ASGC: ntr
- BNL: ntr
- FNAL: ntr
- IN2P3: ntr
- KIT: ntr
- NLT1: ntr
- PIC: ntr
- OSG: ntr
- NDGF: follow up from yesterday. NDGF tried to upgrade dCache to 2.4.1 and reverted to 2.3
- RAL: here is the description of what happened yesterday and the current status. Some planned work going on at the power system (low risk) yesterday. While the work was taking place, some mistake was done and the equipment powered by the UPS got too much power. Many network switches and fiber channel switches are lost. The fiber channel switches in particular serve the DB cluster. The

equipment is being powered on slowly (focus on DBs). A number of services are already back. Priority to FTS but for this the behind needs to be working. No assessment made about DB behind CASTOR and no assessment about batch. GOCDB was restored in r/o mode during the morning using a backup DB. The master DB system hosting GOC is now back, probably will be able to switch in rw mode soon. Frontier is not powered on yet. Need to investigate if queries to frontier do time out or are rejected immediately (te second is obviously the desirable behavior).

- CERN Grid Services: Scheduled Downtime of ce203 tomorrow.
- CERN DBs: a parameter applied on ATLAS online DB yesterday to fix a problem observed last week. Undergoing work to fully understand the problem occurred last week which made the standby stop and some corruption in indexes. Suspecting an hardware problem. A parameter preventing lost writes is being put in place
- CERN Storage: tomorrow CASTOR upgrade for ATLAS and Alice (transparent). Today's CASTOR upgrade for CMS and LHCb was transparent.

AOB: none

Thursday

Attendance: Local(Simone, Ueda, Fernando, Wei-Jen ASGC, AlexeyS ATLAS, Mike, Maarten, Jerome, Giuseppe CMS, LucaM), Remote(Philippe LHCb, WooJin KIT, Gareth RAL, Rolf IN2P3, Christian NDGF, Paolo CNAF, Onno NL-T1).

Experiments round table:

- ATLAS reports -
 - ◆ Vidyo: yesterday we reported about some problem we observed on Monday afternoon, even if we were not sure if the problem was on Vidyo itself, on Tandberg, or some others. We want to clarify here that what we wanted to know yesterday was if someone else experienced similar issues.
 - ◆ T0/T1s:
 - ◇ RAL : FTS and Frontier is back. Waiting for Castor before cloud white listing.
 - ◇ INFN-T1 Site published correct disk information, so now automatic ATLAS tools will handle the situation (blacklisting and cleaning).
 - ◇ KIT : DATATAPE was excluded from T0 Export as was requested by site on WLCG meeting Nov 16. Cache situation is now better. Is it OK to include it back?
 - KIT will come back with a reply after asking the expert.
- CMS reports -
 - ◆ LHC / CMS
 - ◇ waiting for VdM scan
 - ◆ CERN / central services and T0
 - ◇ Tier-0: express run reconstruction failures yesterday, due to a mis-configuration after the T0 reboot for Oracle update. Problem solved overnight.
 - ◆ Tier-1:
 - ◇ T1_UK_RAL still recovering from power problems
 - ◇ SOLVED: Problems with transfers towards IN2P3, GGUS:88576
 - ◆ Tier-2:
 - ◇ NTR
- ALICE reports -
 - ◆ NTR
- LHCb reports -

- ◆ General: some sites are reporting wrong information on their queues in the BDII (queue length with 999999 or SI00 reported as 0). LHCb will submit GGUS tickets to all of them asking for a fix of the information. This concerns mostly Tier2s. List of LHCb queues characteristics.
- ◆ Reprocessing: slowing down as converging with prompt processing
- ◆ Prompt reconstruction: NTR
- ◆ MC productions at T2s and T1s if resources available
- ◆ New GGUS (or RT) tickets
- ◆ T0: NTR
- ◆ T1:
 - ◇ Best wishes to RAL!
 - ◇ Still 30-40% errors in transfers from Tier0/1 to GRIDKA (transfer timeout without a single byte transferred). This is not a showstopper but a nuisance and the tree can hide the forest...
 - ◇ Peak of FTS transfer failures to/from IN2P3 around 23:00 UTC last night. Recovered rapidly...

Sites / Services round table:

- ASGC: ntr
- KIT: ntr
- RAL: CASTOR is still in declared in downtime (even if the service looks up there are things being checked; transfers seem to be running) and will be declared up in 1h or so. Batch still down, working on it. Problem in ATLAS 3D databases (working on it). Some services are running in low redundancy at the moment, so a hardware failure will be more noticeable; more spares are being put in place (for example for power supplies).
 - ◆ question from RAL: sites lost the possibility to resubmit SAM tests themselves for experiments (was possible some months ago and still is can be done for OPS). It is intentional
- IN2P3: ntr
- NDGF: ntr
- CNAF: ntr
- NL-T1: ntr

- CERN Storage: transparent update of CASTOR. Yesterday night there was an issue for for EOSCMS, which was restarted. No failures, low efficiency for some minutes.

- GGUS: There will be a GGUS Release next Wednesday 2012/11/28 according to the standard algorithm, i.e. last Wednesday of each month.

AOB: none

Friday

Attendance: Local(Simone, AlexeyS ATLAS, Fernando, LucaM, Wei-Jen ASGC, Giuseppe CMS, Jerome, Mike, Maarten), Remote(Philippe LHCb, Xavier KIT, Onno NL-T1, Rolf IN2P3, Boris NDGF, Ronald NL-T1, Gareth RAL, Gonzalo PIC)

Experiments round table:

- ATLAS reports -
 - ◆ T0/T1s:
 - ◇ RAL : Thanks for bringing up Tier1! Four disk servers are still down, so ATLAS would like to ask to extend Warning in GOCDB until they are back.

- ◇ KIT : Still waiting for site approval to include DATATAPE to T0 Export (reported yesterday).
 - From Xavier (KIT): experiment can start again using tape at KIT. They should be aware that the tape I/O capacity at the moment is globally limited to 500MB/s (globally = inclusive for all experiments for read and write).
- CMS reports -
 - ◆ LHC / CMS
 - ◇ waiting for VdM scan
 - ◆ CERN / central services and T0
 - ◇ NTR
 - ◆ Tier-1:
 - ◇ T1_UK_RAL almost recovered
 - ◆ Tier-2:
 - ◇ general power cut in the Rome computer center. Recovering.
- ALICE reports [↗](#) -
 - ◆ LHCONE configuration problem in Italy affected 2 important T2 sites Legnaro and Bari from ~18:00 CET yesterday to ~11:00 today, causing them to get almost completely drained of jobs.
- LHCb reports -
 - ◆ General: some sites have set a wall clock time limit == CPU time limit. This makes any matching using CPU work requirement useless as the wall clock limit will always fire first (unless the job has an efficiency > 1). Sites don't seem to understand the issue... Not a major problem in most cases as our job efficiency is good, but it could be degraded by events outside our control (machine too heavily loaded or too high overcommitment of the machine (slots/cores).
 - ◇ Maarten: this is indeed wrong. There is clear documentation fo sites about how to configure WCT and CPU time limits properly to avoid those issues. See: <https://wiki.egi.eu/wiki/Tools/Manuals/TS190> [↗](#)
 - ◆ Reprocessing: ramping up again after RAL restart and sites fixing their BDII publication of queues.
 - ◆ Prompt reconstruction: idle as LHC was not delivering
 - ◆ MC productions at T2s and T1s if resources available
 - ◆ New GGUS (or RT) tickets
 - ◆ T0: NTR
 - ◆ T1:
 - ◇ Jobs running fine at RAL. Some failures at the beginning (CVMFS cache filling most likely)
 - ◇ GRIDKA transfers still in the 20% failure range (job timeout after 1 hour)

Sites / Services round table:

- ASGC: ntr
- NL-T1: ntr
- NDGF: ntr
- PIC: ntr
- KIT: hardware problems with storage cluster for ATLAS (faulty disks, with 74TB of ATLAS data). A list of inaccessible files is being prepared and will be communicated to ATLAS. KIT is trying to recover the data, but only on monday we will be able to know if it will be possible.
- IN2P3: FTS failures in the night between Wed and Thu were investigated. The issue is caused by a bug in SRM, which is hit when using large proxies (larger than 16KB, due to multiple delegations) for authentication. After some time dCache freezes and needs to be restarted. A Nagios probe in In2P3

will be able to spot the issue and restart automatically, while waiting for a dCache patch (dCache deva are aware of the problem and a ticket has been submitted)

- RAL: all systems basically now available. Main exceptions: some disk servers down this morning (a rack of 8, 4 ATLAS, 4 CMS), look OK now, checking before putting them back online. Only half of the tape drives available (no sufficient power because many power supplies were lost in the incident); the number of drives should anyhow be sufficient for the experiments activities. Batch system slowly coming back in capacity.

AOB: none

-- JamieShiers - 18-Sep-2012

This topic: LCG > WLCGDailyMeetingsWeek121119

Topic revision: r25 - 2012-11-23 - SimoneCampana



Copyright &© 2008-2020 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](#)