

Table of Contents

Week of 120409	1
Daily WLCG Operations Call details.....	1
WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments.....	1
General Information.....	1
Monday - No meeting - CERN closed for Easter.....	1
Tuesday.....	1
Wednesday.....	4
Thursday.....	6
Friday.....	8
GEANT incident status report by LCG network team.....	10

Week of 120409

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 - No meeting - CERN closed for Easter [↗](#)

Attendance: local();remote().

Experiments round table:

- ATLAS reports -
- CMS reports -
- ALICE reports [↗](#) -
- LHCb reports -

Sites / Services round table:

AOB:

Tuesday

Attendance: local(Alessandro, Jan, Jarka, Maarten, Mark, Philippe, Xavier E, Zbyszek);remote(Gonzalo, Ian, Jeremy, Jhen-Wei, Lisa, Michael, Rob, Rolf, Ronald, Tiju, Ulf, Xavier M).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central Services
 - ◇ CERN-PROD 0.1% of the srmbbringonline request IDs are lost. Under investigation GGUS:81032
 - ◆ T1s/CalibrationT2s
 - ◇ BNL-ATLAS SRM issue GGUS:80998 (Friday ~ 12:00 CET), site immediately on the problem.
 - Michael: it was due to the PNFS manager failing for the 2nd time, possibly due to running on the same node as Chimera; looking into separating them
 - ◇ NDGF-T1 SRM deletion issues GGUS:81013 (Saturday early morning) -- not a site issue LFC permission were set not correctly, not allowing the deletion. Fixed.
 - ◇ FZK-LCG2 (both errors reported Saturday evening ~ 20:00 CET)
 - GGUS:81021 some files on tape taking long(er than expected, few hours) to be recalled
 - Xavier M: GGUS:81021 has nothing to do with staging from tape; issue also seen for LHCb, looking into it
 - GGUS:81023 one WN had high failure rate with a specific sw release: this WN has 48 cores and NFS as sw area and the same number of jobs (usually doesn't happen at FZK, where jobs are mixed), thus most probably the sw setup (cmt) is too slow. It will have more attention probably on Tuesday.
 - ◇ RAL-LCG2 (problem reported Saturday early morning to the destination site, redirected to RAL Sunday morning) GGUS:81011 FTS server configuration should add also 2 US SRM endpoints (in their *->Tier2D channels) that actually now disappeared from the BDII (ticket opened to OSG for this GGUS:81012)
 - ◇ T0 exports to FZK-LCG2 DataTape failing. Reported 6am Monday (had been occurring since ~3pm Sunday), failures ceased an hour later. Site investigating what happened. GGUS:81039
 - ◇ DE cloud agent error during transfer:No space left on device, seen at most DE sites. Reported 1:42 this morning, failures decreased by 5am, being resolved by 8am (/var full on new FTS front end machine). GGUS:81053
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ Reasonable smooth running
 - ◇ Switched to the physics primary datasets
 - Ian: T1 sites will see custodial and reconstructed data coming
 - ◆ CERN / central services and T0
 - ◇ Generally quiet
 - ◆ Tier-1/2:
 - ◇ Problem reported on pilot factory for jobs to PIC, but it's a central service issue and not a site problem.
- ALICE reports -
 - ◆ NTR
- LHCb reports -
 - ◆ Prompt Reconstruction continued over Easter
 - ◆ Stripping jobs for (both Re-Stripping 17b & Pre-Stripping 18) are 99.9% complete (last few files going through)
 - ◆ MC simulation at Tiers2 ongoing
 - ◆ Had issues with ONLINE farm - Dirac Removal & Transfer agents hanging and thus slowing the distribution of data. Investigations ongoing.

- ◆ Due to different Trigger settings, quite a few of these early runs are taking a long time to process/strip which results in **long** jobs. We have identified the problem and are in the process of stopping the current production, marking bad runs and creating a new one.
- ◆ There is another issue with the 3-4 hour delay between a CondDB update and it being propagated to the WNs. Short term, will put in a 6 hour delay between new Conditions and creating jobs but a permanent solution is in progress.
- ◆ Finally, after investigating issues with slow file access at GridKa and IN2P3, we would like to make a request for leg-cp to handle the dcap protocol (currently uses GridFTP only). This would be preferable for both sites (fewer GridFTP connections) and us (faster transfer).
 - ◇ Maarten: please open a GGUS ticket for that request; it probably cannot be implemented very soon
- ◆ **New GGUS (or RT) tickets** [↗](#)
 - ◆ T0
 - ◆ T1
 - ◇ Gridka: Some nodes using CVMFS are being validated.
 - ◆ T2

Sites / Services round table:

- ASGC - ntr
- BNL - nta
- FNAL - ntr
- GridPP - ntr
- IN2P3 - ntr
- KIT
 - ◆ short LHCb LFC downtime tomorrow for upgrade to 1.8.2-2, exact times still to be decided
- NDGF
 - ◆ 4h network maintenance tomorrow evening starting at 16:00 UTC, will affect ATLAS; concurrent tape maintenance and an extra pool may be added for ATLAS
- NLT1 - ntr
- OSG
 - ◆ looking into the 2 ATLAS T3 missing in the BDII
 - ◇ Michael: we are working with PDSF-NERSC on the ATLAS services they would need to publish
- PIC - ntr
- RAL - ntr
- dashboards - ntr
- databases
 - ◆ on Sun between 07:00 and 20:00 CEST the ATLAS conditions replication to KIT was slow; the cause seemed to be a network problem on the KIT side; the replication latency grew to 5h for KIT and 2h for the other sites; KIT was first isolated from the others, before the issue went away by itself; the DB was not loaded
 - ◇ Xavier M: no evidence observed at KIT
- grid services
 - ◆ ATLAS LFC performance investigations proceeding; there exists a client-side patch for picking a random node behind the alias (workaround for glibc misfeature)
 - ◇ Alessandro: we are interested in that patch, but right now we would rather understand why the machines in the alias always look full, even when the load is supposed to be low
 - ◇ Philippe: LHCb used to have a problem with sessions not being closed properly - maybe something similar here? DB experts may also shed some light on the matter
 - ◇ Alessandro: we now have ~400 concurrent sessions and a limit of 500 - how do things look from the DB perspective?

◇ Zbyszek: we will look into that

- storage
 - ◆ LHCb disk server: controller replaced again, more news tomorrow
 - ◆ LHCb ticket GGUS:81019 [↗](#): thread exhaustion due to DB problem, hotfix applied, please check
 - ◆ EOS slowness reported by ATLAS rather seems to be issue on ATLAS side
 - ◇ Alessandro: confirmed, that is why we did not bring it up here
 - ◆ ALICE CASTOR disk space only has about 400 GB free!
 - ◇ Maarten: probably on purpose, but will alert experts

AOB:

Wednesday

Attendance: local(Eva, Jarka, Maarten, Mark, Ricardo, Torre, Xavier);remote(Burt, Gonzalo, Ian, Jhen-Wei, Michael, Paolo, Pavel, Reda, Rob, Rolf, Ron, Tiju, Ulf).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central Services
 - ◇ CERN-PROD Tier 0: Failing LSF jobs due to no AFS token. Reported 14:30 yesterday. At 21:00 a user with ~800 jobs filling AFS pool space was found, removed. Problem still seen today, affects ~7 nodes (so not a crisis). GGUS:81096 [↗](#)
 - ◇ CERN-PROD 0.1% of the srmbringonline request IDs are lost: ATLAS DDM patched GGUS:81032 [↗](#)
 - ◆ T1s/CalibrationT2s
 - ◇ T0 exports to FZK-LCG2 DataTape failing. Reported Monday, recurred yesterday morning, but OK for 24 hours. Closed this am. Origin of problem not known. GGUS:81039 [↗](#)
 - ◇ Power failure (twice) at TRIUMF-LCG2. CA cloud taken offline ~1-9am this morning. Cloud now back online, sites under test. GGUS:81100 [↗](#) GGUS:81111 [↗](#)
 - ◇ ATLAS functional test failures at IN2P3-CC due to missing files that were deleted during cleanup? Investigating. GGUS:81095 [↗](#)
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ Reasonable smooth running
 - ◇ Switched to the physics primary datasets
 - ◇ Preparing for low luminosity running
 - ◆ CERN / central services and T0
 - ◇ One black hole node seen in Tier-0. Ticket submitted GGUS:81134 [↗](#)
 - ◇ Unable to submit to FTS at CERN between 8-9. Restored for all sites except CNAF
 - Ian: will bridge the CMS Savannah ticket to GGUS or open a new ticket as needed
 - ◇ Ian: some CMS users have VOMS problems since this morning (INC:120113 [↗](#), INC:120148 [↗](#))
 - Maarten: probably related to the DB password changes, which may also explain the FTS problems observed (see central services report below)
 - ◆ Tier-1/2:
 - ◇ A few files lost at FNAL during tape migration. Seen by PhEDEx central agents
- ALICE reports [↗](#) -
 - ◆ Disk SE at KIT fails the tests since yesterday afternoon; under investigation

◇ the cause was a router configuration issue at KIT; tests again OK since 18:00 CEST

- LHCb reports -
 - ◆ Stopped the previous productions that were running over the bad runs and created a new one to run over the data from last night as a test
 - ◆ Jobs going through now - all T1s should be seeing Production jobs going through in the next 24 hours
 - ◆ There is still a known problem with big events slowing down the Stripping, but Reconstruction jobs should behave normally now
 - ◆ 6 hour delay has been put in to avoid the CondDB issue reported yesterday
 - ◆ MC simulation at Tiers2 ongoing
 - ◆ **New GGUS (or RT) tickets** [↗](#)
 - ◆ T0
 - ◇ Any more news on the RAID controller? (GGUS:80973 [↗](#))
 - Xavier: backplane was replaced to no avail; we will try brute-force methods like "dd" to copy the data off the disks, else give up after 2 days
 - ◆ T1
 - ◇ GridKa have updated their LFC and it has been marked online again.
 - ◇ Mark: it would be good if KIT could switch LHCb to CVMFS ASAP, now that the tests are OK
 - Pavel: will forward message to colleagues
 - ◆ T2

Sites / Services round table:

- ASGC - ntr
- BNL - ntr
- CNAF - ntr
- FNAL
 - ◆ so far one of the 3 files lost was looked at: it had a size of 156 GB (sic) and suffered a timeout on transfer; a second attempt then ran concurrently for a while, which is a recipe for trouble; this could point to a bug in PhEDEx
- IN2P3 - ntr
- KIT
 - ◆ ATLAS transfers may have suffered from FTS disk space problem and/or issue with 1 dCache pool
- NDGF
 - ◆ tonight's fiber maintenance should be transparent, as all affected pools were moved from the OPN to the general internet, to be moved back next week
- NLT1 - ntr
- OSG - ntr
- PIC - ntr
- RAL - ntr
- TRIUMF
 - ◆ TRIUMF was hit by 2 site-wide power failures in a single day within ~8 hours interval and lasting several hours, causing significant downtime. Ground fault trips are under investigation. We will circulate a more detailed report in the coming days.
- dashboards - ntr
- databases - ntr
- GGUS/SNOW
 - ◆ around lunchtime CEST today a problem was detected with GGUS attachment forwarding to 3 external ticketing systems, viz. for NGI_GRNET, NGI_PL and NGI_IT; being investigated
 - ◆ as of the April 25 release there will be regular tests also for normal tickets, including

attachments

- grid services
 - ◆ **CERN FTS & VOMS** database passwords were updated this morning for about 80 oracle accounts. Please report any problems.
- storage
 - ◆ ALICE tape recalls were slow due to low priority of the corresponding account ("aliproduct"), now increased

AOB: (MariaDZ) Up-to-date file ggus-tickets.xls is attached to page WLCGOperationsMeetings. There are 6 real ALARMS subject to drills for the next MB of April 24th, the last one was on Easter Sunday (LHCb).

Thursday

Attendance: local(Eva, Jarka, Maarten, Maria D, Mark, Torre, Xavier);remote(Andreas M, Gareth, Gonzalo, Ian, Jhen-Wei, Lisa, Michael, Paolo, Rob, Rolf, Ronald, Ulf).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central Services
 - ◇ CERN-PROD Tier 0: Failing LSF jobs due to no AFS token. Problem still seen, affects ~7 nodes. GGUS:81096 [↗](#)
 - ◆ T1s/CalibrationT2s
 - ◇ ATLAS functional test failures at IN2P3-CC due to missing files deleted during cleanup: files used by FT are being restored. GGUS:81095 [↗](#)
 - ◇ 'No space left on device' for FTS at FZK. 3:25 Thu am. /var full again as loglevel was set to DEBUG. Cleaned. GGUS:81163 [↗](#)
 - ◇ Taiwan-LCG2 Castor tape service down, in unscheduled downtime 18:00 yesterday until 11:00 today. Disk storage elements OK.
 - ◇ Some Tier 0 export file transfer problems to TRIUMF following their power cut until FTS fully recovered. Looks OK now. GGUS:81111 [↗](#)
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ Reasonable smooth running
 - ◇ Updating CMSSW release version
 - ◆ CERN / central services and T0
 - ◇ SUM error on ce203.cern.ch (GGUS:81188 [↗](#))
 - ◆ Tier-1/2:
 - ◇ Job robot failures observed at T1_TW_ASGC
 - Jhen-Wei: would be due to CASTOR issue (see ASGC report)
- ALICE reports [↗](#) -
 - ◆ NFS problem at KIT caused job failures for a few hours; now OK again
- LHCb reports -
 - ◆ New Production going through. Some site issues (see below) but generally significantly better than previously
 - ◆ Fix has been validated for the Conditions DB issue - Will now check local DB install and if not complete will download via web server squid/proxy
 - ◆ MC simulation at Tiers2 ongoing
 - ◆ **New GGUS (or RT) tickets** [↗](#)
 - ◆ T0
 - ◇ Found some nodes with bad HltTCK areas in CVMFS (GGUS:81181 [↗](#))

◆ T1

- ◇ GridKa: Transferred to CVMFS. Had a lot of failures overnight due to empty caches.
 - Mark: it is expected for a new site to need a while before all the necessary software is available on all the WN
- ◇ Seem to be improving now and have increased the timeout on setup to compensate
 - Maarten: I thought files injected at stratum 0 should be available on the WN fairly soon and without hassle?
 - Mark: new files make it fast, triggered by cache misses, while updates to existing files are not checked for very frequently - their latency can be a few hours; this behavior seems reasonable, one just needs to use the system accordingly

◆ T2

Sites / Services round table:

- ASGC
 - ◆ yesterday evening CASTOR went down due to a storage controller failure; unscheduled downtime until this morning; the DB had to be restored from archive; 1 hour of activity was lost; ATLAS was not affected, for CMS it is not clear yet
- BNL - ntr
- CNAF - ntr
- FNAL - ntr
- IN2P3
 - ◆ ATLAS and especially LHCb are submitting lots of pilots that each run less than 200 seconds, which leads to a bad efficiency of the batch system
 - ◇ Mark, Torre: will look into it
- KIT - nta
- NDGF - ntr
- NLT1 - ntr
- OSG - ntr
- PIC - ntr
- RAL
 - ◆ this morning about 8 disk servers suffered network connection drops, not yet understood
 - ◆ 2 new FTS front-end machines were added into the alias; their host certificates are signed by the new UK CA, which has been part of IGTF since autumn last year, but may not yet have been deployed everywhere; may have caused transfer job submission failures at some sites; new machines now removed from alias until issues are understood and fixed
- dashboards - ntr
- databases - ntr
- GGUS/SNOW
 - ◆ see AOB
- storage
 - ◆ LHCb disk server: 1 hard drive has an electrical problem, the others are working; 8 out of the 11 file systems are now being recovered, while the remaining 3 will be tried tomorrow

AOB: (MariaDZ pasting info from GGUS developer Guenter Grein) The problem with attachments to GGUS tickets and the interfaces mentioned yesterday .pl, .it, .gr started with the release on 20th of March. 5 attachments haven't been transferred to other ticketing systems, 3 for Italy and 2 for Poland. We believe the reason for it was the new version of BMC Remedy. The processing of threads in Remedy seems to have slightly changed. Unfortunately none of the external ticketing systems processing attachments offers a test instance. So we couldn't test this before.

Friday

Attendance: local(Eva, Jarka, Maarten, Mark, Ricardo, Torre, Xavier E);remote(Alexander, Gareth, Gonzalo, Ian, Jeremy, Jhen-Wei, Joel, Kyle, Lisa, Michael, Paolo, Rolf, Thomas, Xavier M).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central Services
 - ◇ CERN network problem ~4-7pm yesterday, which I'm sure we'll hear about in this meeting, had many temporary effects.
 - see network services report below
 - ◇ CERN-PROD Tier 0: Failing LSF jobs due to no AFS token. Problem still seen, affects ~7 nodes, awaiting resolution. GGUS:81096 [↗](#)
 - Ricardo: will check (resolved shortly after the meeting)
 - ◆ T1s/CalibrationT2s
 - ◇ TRIUMF-US networking problems since yesterday evening, being worked by BNL/ESnet/TRIUMF. GGUS:81111 [↗](#) GGUS:81213 [↗](#)
 - Michael: we are continuing the investigations this morning; the problem is caused by network issues at or close to TRIUMF
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ Problem with 1 online DB node yesterday. Emphasized again that while the DB is protected by a single node failure, the CMS clients are not and we cannot generally start the run without all nodes up.
 - Ian: may take a while before all the client code has been updated to time out and reconnect automatically
 - ◇ Would like to discuss with IT-DB how we communicate that work has begun on the online DB
 - Eva: it was not an intervention, the node rebooted itself because the number of processes had grown too high
 - Ian: we would always like a phone call in addition to the e-mail that can easily be overlooked in the midst of the hectic situation that typically results from DB problems
 - Eva: OK, will inform colleagues; was already done outside working hours
 - ◆ CERN / central services and T0
 - ◇ Reports of GEANT networking problems yesterday, which manifested themselves in a variety of ways. Timing out connections to the Baltics and Scandinavia. Wide spread PhEDEx agents unable to communicate back. Team ticket submitted, and the situation is improved.
 - see network services report below
 - ◆ Tier-1/2:
 - ◇ Job robot failures observed at T1_TW_ASGC
 - Jhen-Wei: will check, should not be due to CASTOR
 - ◇ Problems with transfers to IN2P3
 - Rolf: ticket?
 - Ian: CMS contact at IN2P3 is looking into it
- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ New Production going through smoothly now after yesterday's interruption.

- ◆ MC simulation has completed for the moment. Waiting on a Stripping fix before more tasks are submitted.
- ◆ **New GGUS (or RT) tickets** [↗](#)
- ◆ T0
 - ◇ GEANT network problem: All jobs started reporting as stalled and had configuration service authentication issues. A fix was applied for the CS and after the network recovered, all jobs back to normal
 - see network services report below
 - ◇ Will there be an incident report covering the issues?
 - Maarten: will ask for that
 - Joel: Thu evening around 19:00 the IT Service Status Board had not yet been updated about the GEANT problems, we then called the operator
 - Maarten: will contact network team and point this out
 - ◇ CVMFS/Broken file problem: Fix in CVMFS client prepared and is currently being rolled out to CERN WNs.
 - ◇ Joel: a new disk server incident response procedure has been agreed between the CASTOR team and LHCb
- ◆ T1
 - ◇ RAL: Minor SRM glitch this morning but recovered OK
 - ◇ IN2P3: SAM jobs found as a possible cause of a lot of do-nothing pilots. Fix is ready and will be rolled out ASAP.
- ◆ T2

Sites / Services round table:

- ASGC
 - ◆ CASTOR checks ongoing
- BNL - ntr
- CNAF - ntr
- FNAL - ntr
- GridPP - ntr
- IN2P3 - ntr
- KIT - ntr
- NDGF - ntr
- NLT1 - ntr
 - ◆ Joel: there was a CVMFS issue at SARA for which SARA would need to contact the CVMFS developers
 - ◆ Alexander: will inform colleagues
- OSG - ntr
- PIC - ntr
- RAL
 - ◆ an at-risk warning until 11:00 UTC was declared for the storage and batch services after yesterday's disk server network disconnections; some kernel patching may have improved the situation
- dashboards - ntr
- databases
 - ◆ the network problem had a side effect on the ATLAS replication to T1 sites, probably due to a bug: the capture process restarted itself from the situation of 7 days ago and needed a number of hours to arrive at the most recent data; the backlog has been cleared
- grid services
 - ◆ **CERN CvmFS Clients** A new (2.0.13) cvmfs client is under going testing on 10% of CERN batch. All being well, sites will most likely be asked to update since it hopefully resolves a bug that LHCb hit. GGUS:81181 [↗](#).

- ◆ **Question to GridKa from CERN** From tuesday it was stated you are validating CvmFS. GridKA currently accounts for 40% of all data downloaded. <http://cvmfs-stratum-one.cern.ch/reports/volume-24hours.txt>. Is the validation normal operation? Contact SteveT.
 - ◇ that traffic may have been due to the ramp-up for LHCb jobs; should go down in that case; cache behavior might have been suboptimal
- networks
 - ◆ see report below
- storage
 - ◆ LHCb disk server: 80% recovered, 10% ongoing, ~10% probably lost

AOB:

GEANT incident status report by LCG network team

On Thursday 12th morning, SWITCH reported to Geant an unusual packet loss in one of the uplinks (CERN was not informed of the impact). The problem had a low impact, until 15:00 when Geant engineers did some operation to diagnose the issue. At this point the connectivity degraded to a level that European sites started having problem to reach CERN. The problem was unnoticed on our side because it was unidirectional, i.e. only packets from CERN to Europe were discarded, especially the big ones (>1000B). At ~18:00 IN2P3 contacted the CERN NOC to report the issue, which looked to us like a central firewall overload due to the increasing number of pending open connections. Finally the problem was pinpointed to affect only R&E institutes in Europe; at 19h15 Geant isolated the faulty hardware and rerouted the traffic to a safe path.

The incident is **not** yet fully resolved and further information will be posted in the IT Service Status Board entry:

- <http://itssb.web.cern.ch/service-incident/connectivity-issues-geant/12-04-2012>

-- JamieShiers - 22-Mar-2012

This topic: LCG > WLCGDailyMeetingsWeek120409

Topic revision: r14 - 2012-04-13 - MaartenLitmaath



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