

Table of Contents

Week of 120813.....	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 120813

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(Alexandre, Cedric, Doug, Jan, Maarten, Raja);remote(Burt, Daniele, Dimitri, Jhen-Wei, Marc, Michael, Rob, Tiju, Ulf).

Experiments round table:

- ATLAS reports -
 - ◆ T0/T1
 - ◇ EOS DATADISK space problem. : [GGUS:85057](#) Errors on dashboard and prod jobs fail. Monitoring said 150TB free. Not space quota but #file quota of 10M. Solution: #file quota Raised to 15M. (1.8PB and 10M files implies 180MB/file. We have small log files there too.)
 - Jan: the next version will have better error messages, distinguishing file from space quotas
 - ◇ LSF sub times increased to several minutes: [GGUS:85058](#) Early Sunday - completely dead (eg. bsub from lxplus fails). Solution: 17:20 Ulrich fixed it - apparently an errant user process killed it.
 - Maarten: all 4 experiments were affected; will follow up with the LSF experts
 - the LSF team are looking into further mitigations to prevent query overload from unmindful user scripts, e.g. by replacing the `bjobs`

command with an intelligent wrapper; no timeline yet, but this month looks unlikely

- CMS reports -
 - ◆ LHC / CMS
 - ◇ CMS solenoid went into fast discharge on Friday, ~7 am CERN time. 4-5 days before recovery is complete.
 - ◇ on Saturday, there was a T0 problem with the HLT menu for some runs, now fixed
 - ◇ Physics fill through Monday, 0 Tesla data taking with a special HLT menu and high rates into ZeroBias and AICa PDs
 - ◇ High load expected on the T0, then the plan would be not to have beams until Tue night (TBC)
 - ◇ **for Castor Ops:** follow-up from Friday's WLCG call, to schedule interventions stay in contact with <cms-crc-on-duty@cern.ch>
 - Jan: the CASTOR update is still being discussed, the amount of downtime uncertain, as e.g. DB indices will need to be rebuilt; it seems better to delay the operation until the next technical stop (Sep)
 - ◆ CERN / central services and T0
 - ◇ LSF: 3.5k job saturation observed on Saturday
 - ◇ over the weekend, we saw LSF issues (see ATLAS GGUS), from the T0 Ops:


```

...
Returned stderr: Cannot connect to LSF. Please wait ...
Cannot connect to LSF. Please wait ...
Cannot connect to LSF. Please wait ...
Failed in an LSF library call: A socket operation has failed: Connection r
...
              
```
 - ◆ Tier-1/2:
 - ◇ Only minor Savannahs at the T2 level, nothing major at the T1 level
- ALICE reports [↗](#) -
 - ◆ CERN: low job throughput during the weekend (see ATLAS GGUS)
 - ◆ QM 2012 [↗](#) has started
- LHCb reports -
 - ◆ T0:
 - ◇ GGUS:85062 [↗](#) Significant problems with jobs over the weekend again. Solved now.
 - ◇ GGUS:85069 [↗](#) Ongoing problems resolving turls by jobs. Problematic diskserver(s)? Interestingly these files are (wrongly?) reported as nearline (they should not be having a tape copy).
 - Jan: last week 300 files were reported as lost, which is linked to this incident; a few machines were briefly put in production, then taken out and reinstalled; this caused various problems; a list of files would be helpful
 - Raja: will provide a list
 - ◇ Lost files due to diskserver failure - dealt with within LHCb (<https://lblogbook.cern.ch/Operations/11304> [↗](#)).
 - ◆ T1:
 - ◇ IN2P3 : New file with corrupted checksum seen (re old ticket - GGUS:82247 [↗](#)). It was created on 10 August. Filename : `srm://ccsrm.in2p3.fr/pnfs/in2p3.fr/data/lhcb/MC/MC11a/ALLSTREAMS.DST/00019658/000`
 - ◆ Others
 - ◇ FTS : Switched off checksum checks to allow very old files to be transferred to CNAF (GGUS:85039 [↗](#)). For now do not see significant errors transferring to GridKa - will keep an eye on it.

Sites / Services round table:

Monday

- ASGC - ntr
- BNL - ntr
- FNAL - ntr
- IN2P3 - ntr
- KIT - ntr
- NDGF - ntr
- OSG - ntr
- RAL - ntr

- dashboards - ntr
- storage - nta

AOB:

Tuesday

Attendance: local(Alexandre, Cedric, Doug, Eva, Jan, Maarten, Raja);remote(Ian, Jeremy, Jhen-Wei, Lisa, Marc, Michael, Rob, Ronald, Tiju, Ulf).

Experiments round table:

- ATLAS reports -
 - ◆ T0
 - ◇ Nothing to Report
 - ◆ T1
 - ◇ Taiwan Castor glitch GGUS:85138 [↗](#) , caused some job failures - Solved
 - ◇ Doug: reminder for ATLAS Tier-1 sites to implement atlasgrouptape directory
- CMS reports -
 - ◆ LHC / CMS
 - ◇ Magnet being brought back up
 - ◆ CERN / central services and T0
 - ◇ Nothing new to report
 - ◆ Tier-1/2:
 - ◇ Only minor Savannahs at the T2 level, nothing major at the T1 level
- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ T0:
 - ◇ GGUS:85069 [↗](#) Ongoing problems resolving turls by jobs - files lost in bad diskserver
 - ◇ GGUS:85134 [↗](#) Opened second ticket as requested in above ticket. More files found lost.
 - Jan: both tickets were updated and the first was closed; about half of the file list were part of that ticket and have been recovered; the rest are still being looked into, probably lost

Sites / Services round table:

- ASGC
 - ◆ yesterday around 23:00 CEST the CASTOR stager became unavailable; the node was rebooted and it recovered; the cause of the problem is being investigated
- BNL - ntr

- FNAL - ntr
- GridPP - ntr
- IN2P3
 - ◆ tomorrow is a holiday in France
- NDGF
 - ◆ for some 50 minutes the FTS and site BDII were affected by a loss of IP connectivity to the outside, but it seems to have gone mostly unnoticed
 - ◆ tomorrow there will be short breaks for dCache pool upgrades, downtimes have been declared in the GOCDB
- NLT1 - ntr
- OSG - ntr
- RAL - ntr

- dashboards - ntr
- databases - ntr
- storage
 - ◆ EOS-ALICE nodes are being rebooted transparently to bring the service into a clean state

AOB:

Wednesday

Attendance: local(Alexandre, Jan, Luca C, Maarten, Raja);remote(Alexander, Dimitri, Doug, Gareth, Ian, Jhen-Wei, Lisa, Michael, Rob, Ulf).

Experiments round table:

- ATLAS reports -
 - ◆ T0
 - ◇ Nothing to Report
 - ◆ T1
 - ◇ Nothing to Report
- CMS reports -
 - ◆ LHC / CMS
 - ◇ NTR
 - ◆ CERN / central services and T0
 - ◇ Reasonable number of jobs on the Tier-0
 - ◆ Tier-1/2:
 - ◇ A few pre-staging and tape family requests were submitted for continuing MC production
- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ T0:
 - ◇ GGUS:85134 [↗](#) Lost files marked as bad within LHCb. Much lower failure rate at CERN now. However as a result of pulling the faulty disk servers, we have are down to a low level of storage : only 22TB free now.
 - Jan: will check and update the ticket
 - ◆ T1 :
 - ◇ GridKa : GGUS:85208 [↗](#) Possible srm problems. Turls not being returned for some files.

- Dimitri: will ping our storage experts
- ◇ CNAF : Diskserver rebalancing - slow access to storage causing jobs and transfers to fail / go very very slow since ~Monday. SE banned until the situation improves, backlog of files to be transferred to CNAF building up.

Sites / Services round table:

- ASGC - ntr
- BNL - ntr
- FNAL - ntr
- KIT - nta
- NDGF
 - ◆ 1 site found its dCache update taking more time than anticipated and postponed the remainder to Friday
- NLT1 - ntr
- OSG
 - ◆ there was a network issue affecting the ticket synchronization between the OSG GOC and BNL; still being looked into, but the tickets seem to be OK now
- RAL
 - ◆ our grid test queue now has a few WN running the EMI-2 SL5 release
 - ◇ Maarten: in the coming weeks and months we will ramp up the testing of EMI-1 and EMI-2 WN, esp. by ATLAS and CMS, who depend much more on the WN contents than ALICE and LHCb
- dashboards - ntr
- databases - ntr
- storage
 - ◆ EOS-ATLAS SRM access was blocked since the early morning due to BeStMan running out of file descriptors; restarted (GGUS:85201 [↗](#))

AOB:

Thursday

Attendance: local(Alexandre, Doug, Eva, Jan, Jarka, Maarten, Raja);remote(Burt, Gareth, Ian, Jhen-Wei, Michael, Paco, Philippe, Ulf, WooJin).

Experiments round table:

- ATLAS reports -
 - ◆ T0
 - ◇ Nothing to report
 - ◆ T1
 - ◇ Taiwan - Another Castor glitch GGUS:85138 [↗](#) (reopened) caused failures of ATLAS jobs. Solved
 - ◇ RAL - Network disruption within switch - ATLAS e-mail notification should also go to atlas-adc-expert@cernNOTSPAMch in future.
 - ◇ IN2P3 - low level (< 4%) file transfer failures - GGUS:85253 [↗](#) - (informational ticket)
 - ◆ Central Services -Yesterday the SAM framework production API upgrade to release 17 started. Now the production SAM API is not available for the intervention time being, only data from preproduction is available. ATLAS is relying on SAM tests (and other metrics) to take automatic actions when sites are not properly performing. The intervention date&time was not announced to ATLAS in advance. We would prefer that a 3 day long operation be

announced in advance and agreed upon with the experiments, so that next time ATLAS Operations are not affected.

- ◊ Jarka: the intervention was announced yesterday morning and started right away; the grid-monitoring alias was moved to a temporary preprod-like instance that was found to serve incomplete data, thereby affecting SUM (Site Usability Monitor); the monitoring might be affected for 3 days, the duration of the intervention, which in turn would affect operations e.g. through automatic blacklisting
- ◊ Eva: the preprod instance is affected by an Oracle bug for which we have a case open
- ◊ Maarten: will follow up
 - the planned intervention was scheduled in the GOCDDB, but not communicated to the Dashboard team; the SAM team management has agreed that interventions will be formally agreed in advance between all affected parties from now on
 - the preprod Oracle bug complicated the preparation of the intervention, but did not cause this week's issues
 - the main problem with temporary instance has been fixed, the availabilities are gradually returning in SUM
 - the SAM and Dashboard teams will keep a close eye on the situation and follow up ASAP on issues that may impact operations

- CMS reports -

- ◆ LHC / CMS
 - ◊ NTR
- ◆ CERN / central services and T0
 - ◊ We appear to have lost one Tier-0 Frontier Server or the monitor for a few hours, but it's back now and no job failures were reported
 - Jan: this morning 1 CMS VObox was dead, which took some time to fix; its admins were presumably contacted per procedure
- ◆ Tier-1/2:
 - ◊ High utilization of T1 and T2, but no problems to report

- ALICE reports [↗](#) -

- ◆ NTR

- LHCb reports -

- ◆ T0:
 - ◊ GGUS:85213 [↗](#) Thanks for the extra space.
 - Jan: all the affected CASTOR machines should be back in now; there are 500 TB unused in EOS
- ◆ T1 :
 - ◊ GridKa : GGUS:85208 [↗](#) srm problems resolved. But sudden downtime (?) for ~15 minutes at 2PM as a consequence to restore redundancy.
 - WooJin: the immediate SRM problem was fixed yesterday, but the underlying real problem had to be fixed by the HW vendor today and required a reboot; we expect very few jobs to have been affected
 - ◊ CNAF : Diskserver rebalancing - possibly (?) completed - FTS backlog cleared overnight.
 - ◊ Raja: transfer issues were just observed for the CERN-RAL and/or IN2P3-RAL channel(s); being investigated on our side
 - ◊ Raja: more missing files observed at CERN, a ticket will be opened

Sites / Services round table:

- ASGC

- ◆ this morning the CASTOR instance for ATLAS was inaccessible due to a configuration error; no data was lost
- BNL - ntr
- FNAL - ntr
- IN2P3 - ntr
- KIT - ntr
- NDGF
 - ◆ tomorrow some dCache pools will have short downtimes for finishing their upgrades; should essentially be transparent
- NLT1 - ntr
- OSG - ntr
- RAL
 - ◆ this morning a network switch problem affected one set of WN for 1h; the given ATLAS mailing list will be informed next time (see ATLAS report)
 - ◆ we see transfer failures for LHCb on the IN2P3-RAL channel (see LHCb report); large files intermittently get transferred slowly and time out after 1h
 - ◇ Philippe: IN2P3 storage experts will have a look
- dashboards - nta
- databases - ntr
- storage
 - ◆ the public CASTOR instance will be upgraded on Monday; should be transparent to the LHC experiments

AOB:

Friday

Attendance: local(Alexandre, Doug, Eva, Jan, Maarten, Nilo, Raja);remote(Alexander, Ian, Jeremy, Jhen-Wei, John, Lisa, Michael, Philippe, Rob, WooJin).

Experiments round table:

- ATLAS reports -
 - ◆ T0
 - ◇ Nothing to Report
 - ◆ T1
 - ◇ NDGF - SRM errors - affect staging in of data for MC production (GGUS:85292[?])
 - ◆ CERN EOS - Tier 0 exports to CERN-PROD failing (GGUS:85269[?]) appears to be fixed now.
 - ◇ Jan: there still are files whose sizes are zero in the meta data; they were written during periods of instability; the cause and the cleanup are being investigated
- CMS reports -
 - ◆ LHC / CMS
 - ◇ NTR
 - ◆ CERN / central services and T0
 - ◇ We appear to have lost one Tier-0 Frontier Server again or the monitor for a few hours, but it's back now and no job failures were reported
 - ◆ Tier-1/2:
 - ◇ High utilization of T1 and T2, but no problems to report
- ALICE reports[?] -
 - ◆ NTR

- LHCb reports -
 - ◆ T0:
 - ◇ GGUS:85260 [↗](#) Missing files - actually caused by unmounted file system. Fixed.
 - ◇ Raja: the SLS sensors for LHCb looked bad just before the meeting
 - ◆ T1 :
 - ◇ GridKa : GGUS:85270 [↗](#) SE problems since early this morning due to "Fault in storage subsystem, vendor support is involved to solve it."
 - WooJin: 1 controller had a glitch this morning; another controller should have taken over, but did not; the whole LHCb pool is unavailable; the vendor is still working on the issue, the downtime has been extended, expected to be finished later today, otherwise we will notify LHCb

Sites / Services round table:

- ASGC - ntr
 - BNL - ntr
 - FNAL - ntr
 - GridPP - ntr
 - IN2P3
 - ◆ our dCache experts have looked into the problem with timeouts on long transfers to RAL for LHCb; the cause has not yet been found; for now the timeouts could be increased to reduce the error rate
 - ◇ Raja: we also can lower the number of parallel transfers for that channel in the RAL FTS; have been in contact with Gareth
 - ◇ John: being looked into
 - KIT - nta
 - NLT1 - ntr
 - OSG
 - ◆ instructions for getting a new (USCMS) T2 into the WLCG accounting?
 - ◇ Maarten: details should be on the WLCG web site
 - <http://wlcg.web.cern.ch/getting-started> [↗](#)
 - RAL - ntr
- dashboards
 - ◆ SUM seems to have recovered from yesterday's incident; an agreement was reached on a better procedure for future upgrades
 - ◇ Maarten: see yesterday's notes for further details
 - databases
 - ◆ we received an Oracle patch that should fix the problem seen in the SAM preprod instance; to be deployed on the integration DB on Monday
 - storage
 - ◆ some EOS-ATLAS machines were affected by the `sudo` post-install bug that caused `/etc/nsswitch.conf` to become readable only for root
 - ◇ for SL(C)5 that is fixed in `sudo-1.7.2p1-14.e15_8.3`

AOB:

-- JamieShiers - 02-Jul-2012

This topic: LCG > WLCGDailyMeetingsWeek120813

Topic revision: r11 - 2012-08-17 - MaartenLitmaath



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback