

Table of Contents

Week of 131202.....	1
WLCG Operations Call details.....	1
WLCG Availability, Service Incidents, Broadcasts, Operations Web.....	1
General Information.....	1
Monday.....	1
Thursday.....	3

Week of 131202

WLCG Operations Call details

To join the call, at 15.00 CE(S)T, by default on Monday and Thursday (at CERN in 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 [↗](#)

The scod rota for the next few weeks is at ScodRota

WLCG Availability, Service Incidents, Broadcasts, Operations Web

VO Summaries of Site Usability				SIRs	Broadcasts	Operations Web
ALICE ↗	ATLAS ↗	CMS ↗	LHCb ↗	WLCG Service Incident Reports	Broadcast archive ↗	Operations Web

General Information

General Information			GGUS Information	LHC Machine Information
CERN IT status board ↗	WLCG Baseline Versions	WLCG Blogs ↗	GgusInformation	Sharepoint site ↗ - LHC Page 1 ↗

Monday

Attendance:

- local: Stefan, Manuel, Felix, Luca, Maarten
- remote: Andrei, Michael, Stefano, Lisa, Onno, Alexei, Tiju, WooJin, Pepe, Christian, Rob,
- apologies: Sang-Un, Rolf

Experiments round table:

- ATLAS
 - ◆ Central services
 - ◇ Pandasrv proxy was malfunctioning (production role problem) resulting in failure to write jobs output on Friday morning. Resolved at around 1PM. Under investigation.
 - ◇ ATLARC DB was down, some work is ongoing (<https://cern.service-now.com/service-portal/view-outage.do?from=CSP-Service-Status-Board>)
 - ◇ Experiment Databases not reporting to SLS, misconfiguration was resolved [INC:445465](#) [↗](#) (<https://cern.service-now.com/service-portal/view-outage.do?from=CSP-Service-Status-Board>)
 - ◆ T0/T1
 - ◇ FZK-LCG2 : A problem with the internal DNS caused the certificate update to fail, resulting in two FTS node a01-013-12{8,9} running with a few expired certificates,

which affects transfers to Prague, CSCS and the other sites. Resolved.

(GGUS:99364 [↗](#))

◇ NDGF-T1: One pool problem (GGUS:99277 [↗](#))

◇ INFN-T1: Was set back to production on Friday. Looks OK.

◇ TAIWAN-LCG2: Looks OK.

- CMS

- ◆ ru-PNPI - During the installation of SL6 a severe power incident happened which led to burn-out of power supplies in their disk array. CMS disk space is not available for the moment.

- ALICE -

- ◆ NTR

- LHCb

- ◆ Main activities is Simulation at all Sites.

- ◆ T0:

- ◆ T1:

- ◇ Network interruption at PIC on the weekend

- ◇ GRIDKA: problems with file downloading presumably due to failure getting metadata from the local SRM

Sites / Services round table:

- BNL: NTR

- CNAF: NTR

- FNAL: NTR

- NL-T1: Currently two pool nodes are down because there is hardware maintenance on the attached storage controller. On december 5th there will be maintenance on a power feed. Not yet submitted to GOCDB.

- RAL: NTR

- KIT: NTR

- PIC: Saturday incident, started ~ 6.25pm local time b/c of power supply problem, affected only WAN, fixed around 22.30. All local jobs were continuing to run ok.

- NDGF: dCache upgrade is done and pools are coming back now

- OSG: problem with SAM availabilities and reliabilities where the transferring data was not working correctly. The problem is fixed and should be cleared up shortly

- ASGC: Castor server for CMS for HC and SAM test down, get it back asap.

- KISTI: Scheduled downtime from 4th December 06:00 (UTC) to 09:00 (UTC) for network intervention. The network bandwidth for KISTI-CERN will be 2Gbps after the intervention.

- IN2P3: downtime on December 10th (Major update for network equipment (routers for IPV6) Minor updates for: CVMFS, dCache servers, mass storage system, batch system controller), Consequences: Total network outage in the morning (no Grid Operations portal at that time neither, so no downtime notifications during 2 hours). Batch downtime starts already at December 9th in the evening, back in the evening of the 10th.

- Storage: EOS ALICE DT in the morning for upgrade to latest version. Wednesday Castor upgrade for ALICE & LHCb, this morning ATLAS & CMS happened

- Grid Services; Outage of batch service on Sat morning, not possible to submit jobs. B/c of faulty configuration. Down from 6am - 12am. SIR

AOB:

Thursday

Attendance:

- local: MariaD/SCOD, Felix/ASGC, Maarten/ALICE, Przemek/DB, Robert/Dashboards, Luca/Storage, Alberto/Grid_Services.
- remote: Alessandro/ATLAS, Andrei/LHCb, Sonia/CNAF, John/RAL, Ronald/NL_T1, Lisa/FNAL, Pavel/KIT, Christian/NDGF, Rolf/IN2P3.

Experiments round table:

- ATLAS reports (raw view) -
 - ◆ Central services
 - ◇ ss-sara-ru(atlas-SS07) in ATLAS_DDM_VOBOXES went down around 1:30UTC on Dec. 5th. Recovered manually at 9:00UTC
 - ◆ T0/T1
 - ◇ FZK-LCG2, PIC, TRIUMF-LCG2, SARA-MATRIX : Many transfer failures due to [SOURCE error during TRANSFER_PREPARATION phase: [GENERAL_FAILURE] RQueued] Duration [0]
 - ◇ Appeared around 22:00UTC on Dec. 3rd and disappeared around 8:00UTC at all the Tier1s listed above. See <http://dashb-atlas-data.cern.ch/ddm2/#activity=%285%29&date.from=201312031200&date.in>
 - ◇ (GGUS:99276 [↗](#)) reopened.
 - ◇ PIC fixed this issue by increasing the limit of SRM requests (to 4k).
 - ◇ SARA-MATRIX in downtime 3:00-6:30UTC on Dec. 5th

Responding to Maarten's request, Alessandro said that the dCache problems started 36 hrs ago and disappeared yesterday evening. Still investigating the reasons why.

- CMS reports (raw view) -
 - ◆ Ken Bloom submitted the following: My apologies -- I cannot attend today! Also, filling this before noon, perhaps things will change by 15. Please contact me by email if there are issues that require my action.
 - ◆ Breaking news: CNAF storage is down and queues have been closed. Perhaps we will see many things fail there.
 - ◆ GGUS:99435 [↗](#), some issue at IN2P3 that might actually be a glitch, but unclear. There is not enough info on the GGUS ticket for me to tell (despite my request for some).
 - ◆ GGUS:99382 [↗](#), Savannah tickets weren't getting bridged. Fixed, but no explanation of what the problem really was.
 - ◆ PIC has been having problems with SRM overloads affecting SAM tests and transfers, e.g. GGUS:99405 [↗](#) and SAV:141083 [↗](#). But these seem to be resolved, at least for now.

Sonia from CNAF commented that the unscheduled downtime they are going through now is due to an unexpected extension of a scheduled one they had announced for yesterday. Their current unavailability affects CMS and LHCb. Rolf also committed to enter the investigation conclusions in the IN2P3 ticket.

- ALICE -
 - ◆ NTR
- LHCb reports (raw view) -
 - ◆ Main activities is Simulation at all Sites.
 - ◆ T0:
 - ◇ EOS name space migration to .../lhcb/grid/lhcb/... completed making possible to migrate user files from Castor to EOS storage

- ◆ T1:
 - ◇ Disk failure at PIC (resolved) resulting in temporary Input Data resolution problems
 - ◇ GRIDKA: xrootd configuration problems, using dcap access for the moment
 - ◇ GRIDKA: currently all pilots are failing with direct submission to the CREAM CE
GGUS:99491 [↗](#)
 - ◇ CNAF storage at srm://storm-fe-lhcb.cr.cnaf.infn.it/... is in unscheduled shutdown

There was a discussion on why services, at times, suffer from expired certificates. The answer was that some have a 10-years lifetime, which makes it hard to remember what procedure to follow close to expiration.

Sites / Services round table:

- ASGC: ntr
- NDGF: Work on LHCOPN next Wed 11/12.
- CNAF: nta
- NL_T1: ntr
- FNAL: ntr
- IN2P3: nta
- OSG: ntr
- KIT: Tape system maintenance work on Tue 10/12 at 8hrs UTC for one hour.
- PIC: (From Pepe by email) We are working to fix the issues with our SRM and dcache overload. Apparently, after some reboots and changes on the config, things are stable, but we'll see. Anyway, we are not on our best week! (SIR on the last weekend network incident is being finalized atm).
- CERN:
 - ◆ Dashboards: INFN was appearing down last week for CMS. This also happens now for IN2P3 (a lot) and RAL. Some catalogue files are not synchronised with their master. This makes these files look critical. Robert was advised to open a GGUS ticket against VOSupport(cms), which he did, as there was no CMS supporter connected.
 - ◆ Grid Services:
 - ◇ 10% of batch at CERN are now running CVMFS against a new stratum 1 (v2.1) at CERN. All of WLCG will be migrated transparently at some to be announced point to this new service.
 - ◇ New CEs are available at CERN: ce40[1-8]
 - ◆ DB:
 - ◇ Mon 9/12 @ 10:30am CET work on INT2R integration DB. Connections from outside CERN will be disabled.
 - ◇ Tue 10/12 @ 8am CET work on CMSINTR for Oracle 11 upgrade (to v.11.204?).
 - ◇ Wed 11/12 @ 8am CET CMS INT2R upgrade.
- GGUS: For the Year End period: GGUS is monitored by a monitoring system which is connected to the on-call service. In case of total GGUS unavailability the on-call engineer (OCE) at KIT will be informed and will take appropriate action. Apart from that WLCG should submit an alarm ticket which triggers a phone call to the OCE.

AOB:

This topic: LCG > WLCGDailyMeetingsWeek131202

Topic revision: r11 - 2013-12-05 - MariaDimou



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