

Table of Contents

Week of 150727.....	1
WLCG Operations Call details.....	1
General Information.....	1
Monday.....	1
Thursday.....	3

Week of 150727

WLCG Operations Call details

- At CERN the meeting room is 513 [R-068](#).
- For remote participation we use the Vidyo system. Instructions can be found [here](#).

General Information

- The purpose of the meeting is:
 - ◆ to report significant operational issues (i.e. issues which can or did degrade experiment or site operations) which are ongoing or were resolved after the previous meeting;
 - ◆ to announce or schedule interventions at Tier-1 sites;
 - ◆ to inform about recent or upcoming changes in the experiment activities or systems having a visible impact on sites;
 - ◆ to provide important news about the middleware;
 - ◆ to communicate any other information considered interesting for WLCG operations.
- The meeting should run from 15:00 until 15:20, exceptionally to 15:30.
- The SCOD rota for the next few weeks is at ScodRota
- General information about the WLCG Service can be accessed from the Operations Web
- Whenever a particular topic needs to be discussed at the daily meeting requiring information from site or experiments, it is highly recommended to announce it by email to wlcg-operations@cern.ch to make sure that the relevant parties have the time to collect the required information or invite the right people at the meeting.

Monday

Attendance:

- local: Maria Dimou (chair), Oliver Gutsche (CMS), Hervé Rousseau (T0 storage), Nacho Barrientos (T0 Grid Services), Stefan Roiser (LHCb), Maarten Litmaath (ALICE).
- remote: Ulf Tigerstedt (NDGF), Alexander Verkooijen (NL_T1), Chris & Kyle (OSG), Sang Un An (KISTI), Lisa Giacchetti (FNAL), Tiju Idiculla (RAL), Michael Ernst (BNL), Dario Barberis (ATLAS), Salvatore Tupputi (CNAF).

Experiments round table:

- ATLAS reports (raw view) -
 - ◆ Nothing special to report
 - ◆ GGUS was down yesterday, all atlas tickets vanished (but others were still there). A warning was announced in GOCDB, fixed this morning.
- CMS reports (raw view) -
 - ◆ nothing to report
 - ◆ GGUS:115120 [is](#) being actively followed up, thanks.
- ALICE -
 - ◆ CERN: job submissions to CERN CEs very slow or failing on Sat (team ticket GGUS:115238 [is](#))
 - ◇ CREAM and ARC CEs affected, they have e.g. Argus in common
 - ◇ issues were indeed seen on the Argus side
 - ◇ situation OK again at 22:00 Sat evening, but further instabilities occurred.

Maarten added during the meeting that ALICE was also inconvenienced by the GGUS unscheduled downtime. On the slow submission to the CERN CEs, Nacho confirmed that Argus high load was observed at the same time. Adding more capacity shouldn't be the remedy as the capacity is more than enough already. The main developer Andrea Ceccanti from CNAF is now on holiday so the investigation has to wait. It might be a low-level java issue or one related to the Agile Infrastructure or ... ?

- LHCb reports (raw view) -
 - ◆ Data Processing: data "stripping" productions ongoing. User and MC jobs
 - ◆ T0
 - ◇ Discussion with LSF team about wall and cpu time queue lengths, ongoing (GGUS:115027 [↗](#))
 - ◆ T1
 - ◇ IN2P3: getting low on user disk space, some more disk to be installed

Sites / Services round table:

- ASGC: not connected
- BNL: ntr
- CNAF: ntr
- FNAL: Scheduled downtime on August 5th from 8am to 5pm CST to upgrade EOS and other storage services. The WNs and interactive services will have to be restarted.
- GridPP: not connected
- IN2P3: not connected but David Bouvet emailed that *a workaround for critical vulnerability is under deployment*.
- JINR: not connected
- KISTI: ntr
- KIT: not connected
- NDGF: Short downtime tomorrow Tuesday 28/7 at 11am UTC to upgrade dCache. It should pass unnoticed. The problem experienced by ALICE last week is now understood. It was due to an xrootd client bug, attempting negotiation with the transfer source on 3rd party transfer but not checking whether the destination does actually support 3rd party transfers. A work-around by ALICE is in place.
- NL-T1: ntr
- NRC-KI: not connected
- OSG: Last Friday an incident with the OSG GOC ticketing system (couldn't connect to the database) was quickly discovered and solved. No problem with the GGUS-Footprints' interface.
- PIC: not connected
- RAL: There will be an outage next Tuesday 4/8 8:30am-4:00pm UTC to investigate a network router software problem.
- TRIUMF: not connected

- CERN batch and grid services: nothing to add to the Argus-related comment above.
- CERN storage services:
 - ◆ EOSALICE restarted (took ~ 1h) OTG0023387 [↗](#): instance was seeing lots of timeouts, caused by big number of NERSC nodes talking to EOSALICE, they don't have reverse DNS entries, we do not cache the negative hits long enough, forcing the MGM to re-resolve everytime.
- Databases: not present
- GGUS:
 - ◆ Please add comments in <https://its.cern.ch/jira/browse/GGUS-1454> [↗](#) about how to propagate **unscheduled GGUS downtimes** to users more efficiently.
 - ◆ Reminder! New release scheduled for Wed 29th of July, with a new certificate for the alarms, and several SU decommissioned. There will be ALARM tests as usual.
- Grid Monitoring: not present

- MW Officer: Critical vulnerability affecting RedHat 5, 6, 7 broadcasted by EGI CSIRT (https://wiki.egi.eu/wiki/EGI_CSIRT:Alerts/libuser-2015-07-24) which allows local root exploit. This vulnerability is present in the case of access via a local user account, so only UIs where access is given via local passwd file could be affected. A fix for the affected package (libuser) is available for RedHat 6 and 7, but not for Red Hat 5, so sites should either upgrade to a new OS release or apply the workaround described at <https://access.redhat.com/articles/1537873>. According to the CSIRT procedures, sites not upgrading by 31/07 risk site suspension.

AOB:

Thursday

Attendance:

- local: Maria Dimou (chair), Hervé Rousseau (T0 storage), Nacho Barrientos (T0 Grid Services), Stefan Roiser (LHCb), Maarten Litmaath (ALICE), Lorena Lobato (T0 DB), Andrea Manzi (MW Officer).
- remote: Ulf Tigerstedt (NDGF), Kyle Gross (OSG), Sang Un An (KISTI), John Kelly (RAL), Dario Barberis (ATLAS), Christoph Wissing (CMS), David Bouvet (IN2P3), Michael Ernst (BNL), Antonio Falabella (CNAF), Thomas Hartmann (KIT), Andrew Pickford & Dennis van Dok (NL_T1).

Experiments round table:

- ATLAS reports (raw view) -
 - ◆ Nothing to report
 - ◆ RCC-KI-T1 went down yesterday and then recovered overnight.
- CMS reports (raw view) -
 - ◆ Dataset needed by SAM accidentally removed from several sites on Monday
 - ◇ Data re-transferred when necessary
 - ◆ Presently low HammerCloud success rates for various sites
 - ◇ Seems to be more a monitoring issue - not a site problem
 - ◇ Under investigation
- ALICE -
 - ◆ high activity
 - ◇ new record 83k briefly reached on Jul 29
 - ◆ an ongoing issue with the central services has led to bad CPU-over-wallclock efficiencies around the grid
 - ◇ under investigation
- LHCb reports (raw view) -
 - ◆ Data Processing: data "stripping" productions ongoing. User and MC jobs
 - ◆ T0
 - ◇ Discussion with LSF team about wall and cpu time queue lengths, ongoing ([GGUS:115027](https://ggus.cern.ch/ticket/115027))
 - ◆ T1
 - ◇ RAL: LAN network instabilities, disk storages were automatically banned on Wednesday
 - ◇ SARA: unscheduled outage of CEs

Sites / Services round table:

- ASGC: not connected

- BNL: ntr
 - CNAF: On Tue 28/7 there was a problem with a switch, which required a board change. This was done transparently, the intervention was short.
 - FNAL: not connected
 - GridPP: not connected
 - IN2P3: added more disk space for LHCb last Tuesday.
 - JINR: not connected
 - KISTI: ntr
 - KIT: ntr
 - NDGF: ntr
 - NL-T1: ntr
 - NRC-KI: not connected
 - OSG: ntr
 - PIC: not connected
 - RAL: Scheduled downtime recorded in GOCDB for next Tuesday, to check the routers and, hopefully, fix the network problems which affected the site last Tuesday for 15 minutes and on Wednesday for about 5 hrs.
 - TRIUMF: not connected
-
- CERN batch and grid services: The ARC CEs (ce501.cern.ch and ce502.cern.ch) ended up dying as apparently something was (is?) triggering a massive memory leak in Condor. We were forced to reboot both and then, when the daemons were started, a new version of slapd was picked up. Both machines are in QA, meaning that the two of them have the SLC6 testing repositories enabled which basically implies that they already are on 6.7. Apparently this minor upgrade pulled in a new OpenLDAP (2.4.40-5.el6) which essentially makes slapd segfault after a few minutes running. This might be for obvious reasons a problem for BDII and its impact should be evaluated, firstly by trying to reproduce the problem. For the moment we've just downgraded OpenLDAP to version 2.4.39-8.el6. Linuxsoft@CERN are, most probably, already aware.
-
- CERN storage services: ntr
 - Databases: ntr
 - GGUS:
 - ◆ Yesterday's release went well. All test ALARMS were handled according to the agreed workflow. Check [JIRA:1439](#). Release notes [here](#). The [Did you know?](#) article this month reminds of existing GGUS features requested at the WLCG Site Survey last autumn.
 - ◆ The GGUS unscheduled downtime last weekend gave birth to [JIRA:1454](#) so GGUS-VOs find an agreed way to notify the supporters/shifters.
 - Grid Monitoring: not present
 - MW Officer: ntr

AOB:

This topic: LCG > WLCGDailyMeetingsWeek150727

Topic revision: r13 - 2015-07-30 - MaartenLitmaath



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback