

Table of Contents

Week of 090223.....	1
WLCG Baseline Versions.....	1
WLCG Service Incident Reports.....	1
GGUS Team / Alarm Tickets during last week.....	1
Daily WLCG Operations Call details.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	3
Wednesday.....	4
Thursday.....	5
Friday.....	6

Week of 090223

WLCG Baseline Versions

WLCG Service Incident Reports

- This section lists WLCG Service Incident Reports from the previous weeks (new or updated only).
- Status can be received, due or open

Site	Date	Duration	Service	Impact	Report	Assigned to	Status
------	------	----------	---------	--------	--------	-------------	--------

GGUS Team / Alarm Tickets during last week

- Tickets: see https://gus.fzk.de/ws/ticket_search.php and select "VO" - this will give all tickets, including team & alarm - or use the following link:
 - ◆ In the GGUS Escalation reports every Monday

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

General Information

- Web archive: <https://mmm.cern.ch/public/archive-list/w/wlcg-operations>
- CERN IT status board: <http://it-support-servicestatus.web.cern.ch/it-support-servicestatus/>
- elogs: <https://prod-grid-logger.cern.ch/elog/CCRC'08+Logbook/>

See the weekly joint operations meeting minutes

Additional Material:

Monday:

Attendance: local(Jamie, Harry, Eva, Miguel, Andrea, Roberto, Andrea, Simone, Jean-Philippe);remote(Michael, Angela, Gareth).

Experiments round table:

- ATLAS (Simone) - was a problem discovered Saturday for FTS at CERN - server for T2s - looks like delegation problem. This FTS server servers channels to ATLAS calibration T2s. Only 2 / 4 of these (channels) show problems. Usual recipe doesn't work - open ticket for this. Miguel - Gavin on it. ATLAS is now testing new version of FTS which solves delegation problem - FTS in PPS - have to wait to see if there is a problem. Next weekend fairly large scale test by trigger-DAQ people - no involvement of T1/T2 but big load on T0. Michael - is there a plan to rerun 10M file test before reprocessing exercise? A: no - reason is that site services have been patched a couple of times which delayed final deployment. Now have stable release but this pushed schedule 1.5 - 2 weeks. Need to prepare tests and distribute data - concurrent activity of trigger-DAQ so no margin to have rerun

before reprocessing. Beginning of April there will be another large scale test, i.e. wait for end of reprocessing.

- CMS reports - (no report - Daniele travelling back from FNAL).
- ALICE (Patricia) - Last Thursday night I put RAL and FZK in CREAM mode. This means, both sites are submitting to the current production CREAM service with very positive results. In fact the system is running all alone. In addition it seems that wms221 (next PPS43 WMS put in production for Alice) is performing quite well. WMS support has asked us to stress it, so I am putting this WMS exclusively at CERN as the unique WMS for Alice use so it can be stress in production conditions.
- LHCb (Roberto) - huge amount of fest-related simulation production (to prepare for next phase). This is LHCb week. GGUS ticket open against CERN re grdiftp server issue last week - lcg_cp fails or was failing Friday evening. Miguel - will check. Roberto - issue at Lyon; problem no longer there - lcg get turls returns wrong url. SRM@CNAF: should be resolved with scheduled upgrade. This Wed another test of fest: transfer to T1s and recons test at T1s may come - to be clarified.

Sites / Services round table

- NL-T1 (JT) - glexec and scas in production at NIKHEF - go ahead and test! Master & slave process within scas service. Every 200s correction - 300s a new slave is restarted internally. Same trick as used by Apache for memory leaks (at lower level...) No need to restart scas. Quite some work to eliminate memory leaks. scas first glite tool to accept multiple connects within same process. Most services don't do this and hence don't see memory leaks grow. New globus which fixes many of these, some in scas code. Hopefully next release will decrease error rate from 0.1% to ... Roberto - q re glexec clients on WN.
- NL-T1 (JT) - A DDN storage device partially crashed and needs a cold reboot and some additional actions. Turned out that rebuild didn't work and needed power cycle.
- FZK Doris) - FTS 2.1 now in production. Would like to ask expts to switch to new instance (FTS-FZK). FTM was changed to point to new instance. Also participating in SCAS pilot - glexec on WNs - only accessible via specific CE (dedicated). See scas pilot home page.
- CERN (Miguel) - config change in routers which now block castor nameservers from Internet. Config done in routers this morning. Should not affect anything but there was a change....
- RAL (Gareth) - upgraded FTS to 2.1. Also brief outage of myproxy for move to different h/w. FTM now available. Address on wiki page: <https://twiki.cern.ch/twiki/bin/view/LCG/LCGFTMEndpoints>. As announced in middle of castor upgrade for ATLAS - ongoing. Simone - different e-p? A: effectively old one upgraded - no config change needed.
- DB (Eva0) - couple of failures on replication from ATLAS online to offline (user mistake) and today on CMS on-off - still working on it. Replication to CNAF disabled during w/e as CNAF unreachable. Affecting all streams (ATLAS, LHCb). Intervention this morning on CMS to apply latest security patches and other DB/cluster patches.

AOB:

- SAM tests for CMS (Andrea) - # jobs submitted during w/e stuck in status 'waiting' - reported to wms.support. (also for OPS tests?)
- Dashboard (JT) - ATLAS & ALICE sft job tests for CE not running on a regular basis. Some 'wierd corner of dashboard code' makes it look as this is happening (but its not...) Visible OK through SAM

portal(!)

Tuesday:

Attendance: local(Julia, Jean-Philippe, Jamie, Harry, Nick, Miguel, Sophie, Patricia, Simone, James, MariaDZ);remote(Gonzalo, John, Daniele).

Experiments round table:

- ATLAS (Simone) - yesterday's upgrades; 1 for ATLAS SS - went smoothly. Problem with 1 of VO boxes (serving Lyon cloud). Config problem - service down for a couple of hours. Now restarted. Other change was dashboard migration. Now all instances run against production DB and both will be merged into one.
- CMS reports (Daniele) - T1s & t2s quite some issues over last days. Closing ticket for FZK about transfers to various T2s. Fixed by downtime last week. New ticket for FNAL - Nebraska gets consistent failures - specific file lost or not transferred?? Or just stuck in stagein. FNAL - issue since Feb16 - 1) general FTS discussion (Akos) 2) transfers in debug Phedex instance to Brazil. Since few days d/s not moving at IN2P3 to tape. No progress from site since Feb19. Transfers to T2 in Fr - connection timeout in gridftp; srm-related problem. Yannick Patois' DN -- FNAL: transfers of specific d/s from T2s back to T1. Good progress - close ticket soon. T2s: closing most tickets; IBHC in Fr - file exists error. More in twiki. No space left at T2 in Poland (Warsaw) - understood, closed. Transfer IN2P3 to Purdue; stagein from HPSS. New ticket from Purdue, some transfer from RAL but link red in phedex. CMS layer so no GGUS ticket .. Wisconsin: transfer issue from IN2P3 - bouncing of ticket Wisconsin <-> FTS. Problem Wisconsin problem.
- ALICE (Patricia) - minor things for Cagliari: s/w area not accessible; 2nd VO box at CNAF for CREAM - s/w area not writeable. CREAM CE progress: 2 sites in production since Friday (RAL & FZK) RAL running ~250, FZK 450-800 jobs - going well. GSI & CNAF should enter today. General ALICE production: 5400 jobs, allows to check how WMS working. Some changes in CNAF WMS - so far so good. Latest version with super-mega-patch. Checking regularly. VO boxes at CERN only point to WMS with these patches and seems ok - no huge backlog. Nick - how many jobs per day? A: will check, < 10K. Sophie - can check too. Sophie - CREAM setup at CERN, problems with package list so will wait for release soon from PPS.
- LHCb (Roberto):
 - ◆ LHCb is trying to generate the FEST input data. Generated so far ~25 million of 100 million events. This is being performed across all sites with the output data being uploaded to CERN. This should complete before Thursday at which point a production to merge the generated data into 2GB files is launched. The merging (a cat of small input MDF files) will all be performed at CERN and should then finish by the end of the week. This data will then be copied from CERN Castor to the pit over the weekend to prepare the full chain of FEST.
 - ◆ Dummy MC production have been requested to be stopped until a problem with the simulation application (Gauss) (causing run away CPU usage) is understood. This was causing many jobs killed by sites existing CPU limits.
 - ◆ New CNAF SRM endpoint for Castor has been successfully tested.
 - ◆ IN2P3: long standing ticket (#45699) concerning wrong status reported is still awaiting for a final solution.

Sites / Services round table:

- PIC (Gonzalo) - comment on issue affecting PIC - coming from LHCb. Quite a lot of jobs arriving to PIC under LHCb SGM acct. Look like normal jobs - not installing s/w! Causing problems - sgms run on dedicated WNs which mount s/w area r/w. Not many WNs - these jobs are accessing SQLite

Tuesday:

through NFS(!!!). Nick - WNs dedicated to LHCb? A: no - shared by other VOs so blocking affects ATLAS and CMS as well. JT: how do you have gridmap file configured? A: will check. vo roles in lcms file. JT - we saw this happen here once before. Something went wrong in VOMS mapping so backed off to map file. Make sure no LHCb people get mapped to sgm this way... GGUS # 46644

AOB: MariaDZ: It's show time folks! End of February: VOs to perform an ALARM ticket test (full round from opening to ticket closing) to Tier1s. savannah ticket #105104 and testing rules.

Monthly USAG meeting this Thursday - reminder sent to mailing lists this morning.

At CERN all mailing lists migrate to e-groups and this migration is "nothing but smooth". e-mail submission to GGUS / notification from GGUS anomalies should best be reported to this meeting and not to mailing lists which might also fail during this migration!

Wednesday

Attendance: local(Jamie, Jean-Philippe, Harry, Ewan, Maria, Simone, Roberto, Antonio, MariaDZ, Angela);remote(Daniele, Jeremy, Gareth, Luca).

Experiments round table:

- ATLAS (Simone) - Only thing is that MC prod and transfers stopped for Taipei as fire in ASGC! (Set cloud offline).
- CMS reports (Daniele) - since last call following up 4 new internal tickets and 2 GGUS. T0: no ticket, some 15% free space (150TB) on CAF. Deleting old data. T1s: ASGC - same as ATLAS! - IN2P3 - both on-site contacts currently unavailable - open more GGUS tickets under these conditions? Tickets to FNAL & CNAF - following up. "Akos Frohner" thread -> problematic credentials. Low speed MC d/s transfers; CNAF - open export errors to Fr T2. T2s: closing ticket to Purdue; [full details in CMS twiki]
- LHCb (Roberto) - However: answering Gonzalo, yes, some SAM jobs do run using SGM credentials on the special WNs set on the site and accessing SQLite file on the shared area causing the problem reported by Gonzalo yesterday. We are going to propose to have SAM jobs running with a new credential (Role=sam) that should have and guarantee the high priority that currently lcgadmin role does guarantee.
Furthermore I have to report a problem we are debugging with Remi Mollon concerning lcg_utils: lcg-getturls (w/o a list of protocols specified as input) does not work properly in some circumstances (e.g. dCache end-points). Debugging that we discovered that still some SRM endpoints like IN2P3 are advertising them selves as srmv1 while we repute that now the default must be SRMv2 everywhere.
- ALICE (Patricia: report submitted after the meeting): wms221 is being stressed now through the Alice production. With 9000 jobs passing through the system yesterday, the stability is good enough to begin the stress tests of the system. wms214 has been put in draining mode to put the whole weight in wms221. Currently Alice is running 7388 jobs through all the sites with special prevalence to wms221. Load of the service going up but still in a normal rate.

Sites / Services round table:

- DB (Maria) - intervention for upgrades of internal net switches of prod clusters on RAC5. Follows similar intervention on integration clusters a couple of weeks ago. Had unexpectedly and under investigation services went to run on a single - hence severe degradation for ~15 minutes - high load seen on ATLAS side. In case one network switch goes down system should be fully redundant so not understood. ATLAS online, offline, compass and downstream capture clusters.

- CNAF (Luca) - Q servers performing SAM tests are on OPN? Have to check. Last Saturday night we suffered network interruption on link other than to CERN and all SAM tests started to fail. DNS ok - maybe no route between some nodes on which SAM tests run and CNAF?
- CERN (Ewan) - in advance of massive network intervention on 18th some LSF queues already draining. Will advance as we get closer... Simone - do we know which services will be affected? A: everything! See <http://it-support-servicestatus.web.cern.ch/it-support-servicestatus/ServiceChangeArchive/ImportantDisruptiveN>
- gLite briefing (Antonio) - <https://twiki.cern.ch/twiki/bin/view/LCG/ScmPps> We received gLite 3.1 pps update 44 which contains new version of CREAM that introduces short-term solution for proxy renewal on WNs. Will be deployed in production today. Fix for bug 44712 - configuration problem in LCMAPS affecting conf file used by glxexec, affecting specifically ALICE almost everywhere. Fix for yaim core, 64bit version of BDII. New glite WN info command on WN. To be called by job wrapper tests. Soon (i.e. today) in production update to WMS 3.1 "megapatch" and fixes to megapatch. New version of CREAM as above.
- RAL (Gareth) - GridView plots all greyed out - why? Other T1s too. Luca - noticed for SE and SRMv1 - phasing out of these tests?
- All of SAM infrastructure is on GPN.

AOB:

Thursday

Attendance: local(Julia, Jean-Philippe, Nick, MariaDZ, Diana, Jamie, Harry, Steve, Roberto, Gavin, Simone, Gareth);remote(Jeremy, Angela).

Experiments round table:

- ATLAS (Simone) - decided that focus for DM transfers now is to be making sure that d/s subscribed make it 100% - shown that many files & many TB can be transferred. Every Thursday functional suite rotates and then look at which d/s didn't make it and why. This is the first Thursday. Only a few problems. All understood! Tail of unfinished transfers TRIUMF - Toronto - Toronto SE suffering overload and hence FTS channel 3 instead of 10. T2 in DE - ran out of disk space - full. Problems with few d/s from PIC and NDGF for tape e-p. Because transfer T0-T1 switched to use FTS in PPS to try out delegation patch last week and PIC & NDGF channels were not served - Gav fixed this morning. Last - Lancaster didn't get all data as concurrent transfers for DPD which had priority. Something like this every Thursday. Functional tests: for MC production which were running a couple of months ago and stopped for migration of DB b/e MySQL -> Oracle will be restarted today. Every day a batch of 200 jobs will be defined for each cloud and end of week stats of how sites did collected. Type of jobs are well behaved - should be no problems related to ATLAS s/w. Report again each Thursday. Angela (FZK) - talking about transfers. Did you notice we changed FTS f.e? We don't see any activity on new instance. Simone - waiting for end of current tests- will start to be used today. Jeremy - pcache suggestion. Any sites implemented? Simone - will check. Brian - could you possibly forward to me info about d/s for Lancaster which didn't complete.
- CMS reports - Apologies for today, I cannot attend (Daniele).
- ALICE -
- LHCb (Roberto) - nothing to report! MC simulation for Fest09 is over - 100M events produced to be injected into online and then into full chain in next week. Requested other 100M events for the next

week by Online people (they need a larger number of L0-accepted events)

Sites / Services round table:

- Network (Gregory Bevillard) - We have a DNS failure caused by ~1000 different hosts in the CC (LCG) which are asking for 0.168.169.202.in-addr.arpa and 0.98.109.140.in-addr.arpa. These queries are failing, this should be related to the disaster in Taiwan.
- FTS (Gav) will set FTS channels inactive

AOB:

Friday

Attendance: local(Jean-Philippe, Nick, Harry, Alessandro, Patricia, Roberto);remote(John, Angela, Jeff).

Experiments round table:

- ATLAS - Some minor problems yesterday. During two hours FTS transfers to TRIUMF failed with a new (to us) error messaged of defective credentials detected - said to be a gridftp problem. There was a 'permission denied' error at Toronto due a known race condition. ATLAS have started MonteCarlo functional tests with one task of 1000 jobs per week per cloud which will be followed by 100 jobs per day per cloud.
- CMS reports -
- ALICE - The VObox CREAM-CE interface is currently being upgraded at several sites. Getting good performance of the CREAM-CE at FZK where they have 2400 concurrent jobs. CNAF and RAL are also running CREAM-CE and two more sites will be added today. The CERN WMS running the 4.3 megapatch crashed on Wednesday night. A bug has been identified and fixed. The upper limit of this WMS appears to be at about 10000 jobs/day but this is much higher than the threshold where backlogs appeared in the unpatched WMS.
- LHCb - Still having jobs stuck in waiting status at sites - sometimes SAM tests, sometimes pilot jobs and either way activity drops at the affected site. There is hope that the WMS megapatch will improve this. There was a swap full condition at 04.00 on the CERN VObox running the Dirac WMS upon which the operator logged in and killed python processes. This procedure needs to be reviewed.

Sites / Services round table:

WMS (NT): the WMS 4.3 megapatch should be released to production next week.

AOB:

This topic: LCG > WLCGDailyMeetingsWeek090223

Topic revision: r12 - 2009-02-27 - PatriciaMendezLorenzo



Copyright &© 2008-2021 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