

Table of Contents

Week of 081117.....	1
WLCG Baseline Versions.....	2
WLCG Service Incident Reports.....	3
GOCDB downtime report.....	4
GGUS Team / Alarm Tickets during last week.....	5
Daily WLCG Operations Call details.....	5
General Information.....	5
Monday:.....	5
Tuesday:.....	6
Wednesday.....	8
Thursday.....	11
Friday.....	13

Week of 081117

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
None							

GOCDB downtime report [↗](#)

- report

GGUS Team / Alarm Tickets during last week

- Team tickets
- Alarm tickets

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(Jean-Philippe, Nick, Simone, Julia, Harry, Maria, Jamie, Sophie, Andrea, Roberto, Maria, Markus, Gavin, Patricia, Olof);remote(Michael, Gareth, JT).

elog review:

Experiments round table:

- ATLAS (Simone)
 - ◆ ASGC: still Oracle problem, 70% efficiency at the moment. Jason is aware. (Start with corruption, system degrades, eventually collapses...)
 - ◆ CERN: Problem in SRMGet for the PPS SRM CASTOR instance at CERN. Being followed by Jan. (Fixed apparently - 'iptables'(?))
 - ◆ Planning: the 10M files test foreseen for the end of the week might have to be postponed to next week, since the new ATLAS site services might not be ready and deployed in time. Need confirmation from developers. Some sites have shown interest in having the test this week for various testing of the storage (CERN and RAL). This can be done, using the existing site services. More news during the week. (Wednesday?)
 - ◆ ATLAS production system DB - lock condition on Sunday. Not clear what caused lock or solved situation. P-M awaiting from Gancho & Florbella. Escalate to whom? Maria - IT-DM physics DB services. Expert on-call called Florbella & Gancho. Maria - F&G don't have admin access to these services.
- CMS (Andrea) - in "no running" condition. T0 trying to processed failed recons jobs. T1s: as ATLAS see ASGC problem.
- LHCb(Roberto) - not much to report. Currently running usual dummy MC. 1300 jobs currently running. Transfer load generatro running. No evidence of any problematic endpoint. GGUS ticket about LFC replication to GridKA? No news from GridKA. Problem still there - cannot see new

entries. Maria - no problems at streams level. Ticket assigned to GridKA. Q: how are R/O LFCs chosen? A: normally that defined on WN.

- ALICE (Patricia) - production going up and down continuously. Some validations in code have to be done. Hope for green light for deployment of latest Alien this week. SLC5 - continue tests this week but will stop when production restarts.

Sites round table:

- RAL (Gareth) - downtime scheduled for tomorrow for ATLAS and CMS srms. In GOCDB.
- NL-T1 (JT) - intervention scheduled "at risk" Tuesday (not Wed as stated at meeting) - several Zen host m/cs will be rebooted kernel upgrade. Some things will drop out for 1-2 minutes. Jobs from ATLAS & LHCb (about 100) - thanks!
- CERN - vault power off tomorrow- tape robotics - 1/3 LHC data will be inaccessible. Stage ins delayed.
- Simone: followup on CNAF-BNL transfer tests? Michael - still in progress. Latest news is that BNL doesn't see any problems at NY peering site GEANT-ESNET. Need a host there to test connectivity from POP to either site (CNAF|BNL). Either NY-CNAF or at POP between 2 service providers. Marco - in charge of networking in Italy - sent message proposing a person to setup a host at NY POP. Not as simple as end problem - in or in between nets. Q2 for Jeff - do alarm tickets for SARA & NIKHEF end up at same place. A: yes.

Services round table:

- DB services (Maria) - applying today ATLAS CPU+10.2.0.4 latest bundle. Warned users that no integration or test RACs 8-18 due to vault. Production clusters on RAC6 for test affecting critical area - some downgrading by 50% of # servers.

AOB:

- Q: what do VOs require in terms of direct routing for all tickets for sites that are uncertified (normally "not yet" production) or suspended? Tickets routed to site but go to ROC in addition.
- Xmas activity - plan and requirements. ATLAS - reprocessing controlled from CERN. MC production relies on services at CERN and BNL but more distributed in terms of processing and support. RAL - concern how enough on-call will be provided. Interesting to know requirements. Olof - data ops at CERN have piquet coverage.

Tuesday:

Attendance: local(Gavin, Maria, Jamie, Jean-Philippe, Nick, Harry, Sophie, Julia, Andrea, Maria, Simone, Patricia, Jan);remote(JT, Michael, Gareth, Jeremy).

elog review:

Experiments round table:

- ATLAS (Simone) - yesterday load generator for transfers stopped about 16:00. Runs using acron and problem in daemon running on m/c which runs tests. Restarted - stopped again - Birger contacted Veronique. acron started. Ticket - not fully fixed (Soph'). 3 points:

Monday:

1. ATLAS would like to push for FTS on SLC4 on T1s asap. Not new func but many fixes and monitoring. Is it ready? Nick - have to check if out of cert.
 2. ATLAS would like to push dCache sites - at least T1s - to upgrade to 'fast pnfs' - postgres 8.3 + ?? Michael - modified pnfs code. Doesn't run through full chain of authorisation - saves many CPU cycles. BNL have taken from dcache.org and deployed at BNL. Really dramatically improves pnfs performance! Avoids limit of 1000 entries / dir? No - limited chain of auth.
 3. Found that entries in LFC at CERN and maybe elsewhere which are not writable by ATLAS - dir owned by USATLAS user. Comes from fact that some jobs run via PANDA - pilot submitted at that time by Nercan. Asked LFS support@CERN if possible to provide recipe to grant r/w/e privs on all dirs in LFC owned by ATLAS (/grid/atlas) at DB level - 18M entries - once procedure available and tested should schedule intervention. Test on dump of DB. Sophie - asked for script to developers based on Lana's script for LHCb. Probably ask for same intervention also at T1s - 'solve forever'. Gareth - dcache request, does this apply for T2s as well? A: discussed for T1s, T2s would also benefit. Jeremy - some problems with installations of latest ATLAS install at one site. Switching to WMS? Fault with glite WMS and default umask? Simone - WMS used for quite some time. Which site? Royal Holloway.
- CMS (Andrea) - squid sam test failed at all sites. Frontier server in CERN cc went down due to power tests. CMS SAM tests moving to WMS. Should be completed today after which CMS will be an RB-free zone. (wms201 & 206).
 - ALICE (Patricia) - still not in production. Not even user jobs. Checking with dummy jobs -> central task q if new WMS model is ok. Need DN for WMS at NL-T1 asap. JT - ok. Do ALICE need a 2nd VO box with CREAM CE. A: wrote document -> OPS meeting. T1s specified that a 2nd VO box would be provided. Just for period when LCG and CREAM CEs co-exist.
 - LHCb (Roberto) - there are some stripping activities (still pending production request from DC06) that will take place (T1 to be advertized) concurrently the Particle Gun production will also start and the file merging of FEST09 simulated data for the FEST exercise.
The main issue encountered now with sites is about the LFC service at GridKA (a replication problem) that is under investigation by local DBAs and Eva/Maria. The former claim "Streams for LHCb-LFC from CERN to GridKa is not working properly", the latter claim that Streams are working properly (accordingly their monitoring). The GGUS ticket open is #43533 also discussed at the weekly ops meeting with Angela from GridKA.
A second point is about MC dummy production that pointed out (for some jobs) a wrong normalization of the CPU time turning into jobs brutally killed by LRMS.

Sites round table:

- NL-T1 jobs! And only from LHCb! Did anyone notice during maintenance window a few services didn't come back automatically. Useful exercise - continuing to discover this....
- RAL - still in downtime for ATLAS & CMS CASTOR SRMs.

Services round table:

- SRM PPS problem mentioned yesterday by ATLAS:

Concerning:

CERN: Problem in SRMGet for the PPS SRM CASTOR instance at CERN. # Being followed by Jan. (Fixed apparently - 'iptable'(?))

Today we started the preparations to upgrade the different PPS SRM endpoints to the latest s/w version (2.7-10), by upgrading our own srm-pps.cern.ch validation endpoint.

We found that an obsolete set of iptables rules for RFIOD was still configured, and removed it from all PPS endpoints. Unfortunately this blocked communication between the SRM's and their Stagers, as the removed ruleset actually was needed to allow stager callbacks...

This was understood and fixed at ~14:00, when we deployed updated iptables rules. At the same time we upgraded all the PPS endpoints to 2.7-10 as well.

This should be the end of the report.

At the moment we are still receiving SAM reports for failing tests, both on PPS **and** production endpoints. We are following this up with priority...

[Debugged as problem with # DB servers - not related to previous problem - due to scheduled upgrade yesterday. All should be ok now. Only affected PPS endpoints. Sophie - ticket updated - to be checked.]

- DB (Maria) all int & test services came back at 14:30 after power cut affecting vault. Streams is working to FZK - since August due to problem affecting only FZK updates have latency of about 1h for updates to be available. (Oracle apply not done in real time but with redo logs - about 1h delay wrt other sites).

Hi Roberto,

After talking with Maria, I would like to know what are exactly the problems you are observing with Gridka replication. From the Streams point of view, everything is working fine and the apply is updating the destination database. Since August, the replication to Gridka is not running in real-time due to the propagation problem we found and which is being investigated by Oracle (we decided to maintain Gridka in a separate Streams setup in order to avoid any impact on the other 5 tier1 sites). This means that Gridka only gets the updates every time a switch log file happens, approximately every 1 hour. If you think there is a different problem or you need any further information, please let me know and we can meet in order to discuss this issue.

Thank you in advance. Cheers, Eva

AOB:

Wednesday

Attendance:

local(Simone,Eva,Maria,Flavia,Harry,Jean-Philippe,Sophie,Maria);remote(Michael,Gareth,Jeremy,Jeff).

elog review:

Experiments round table:

- ALICE: Testing the latest WMS before putting them in production. Kolkotta T2 has already provided a CREAM-CE setup to be tested today. AliEn v2-16 quite advanced to be tested tomorrow during the TF meeting

* ATLAS (SC): One current issue that two T0D1 files at CERN cannot be located and may be lost. There were no side effects from yesterday's scheduled CERN power cut. Flavia asked the status of the problematic BNL-CNAF network traffic. Michael reported that BNL, ESNET, Geant and the Italian research network

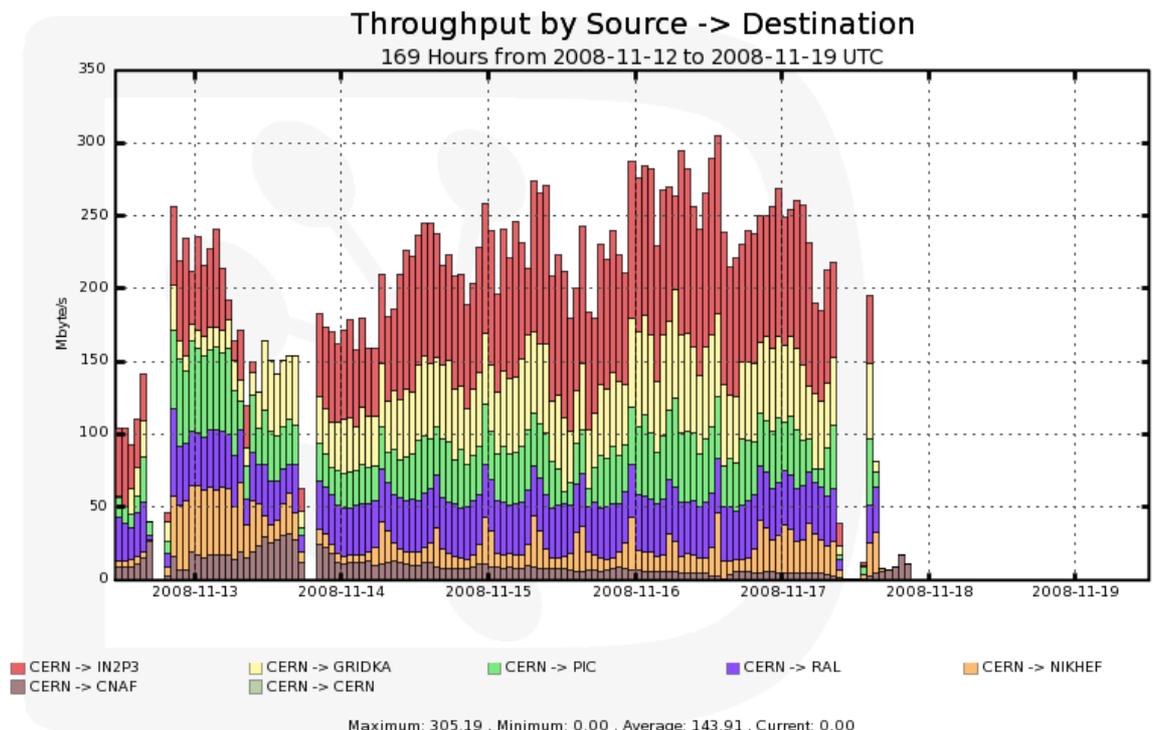
were setting up hosts to use between the sites for the next series of tests which he hoped would happen over the next few hours.

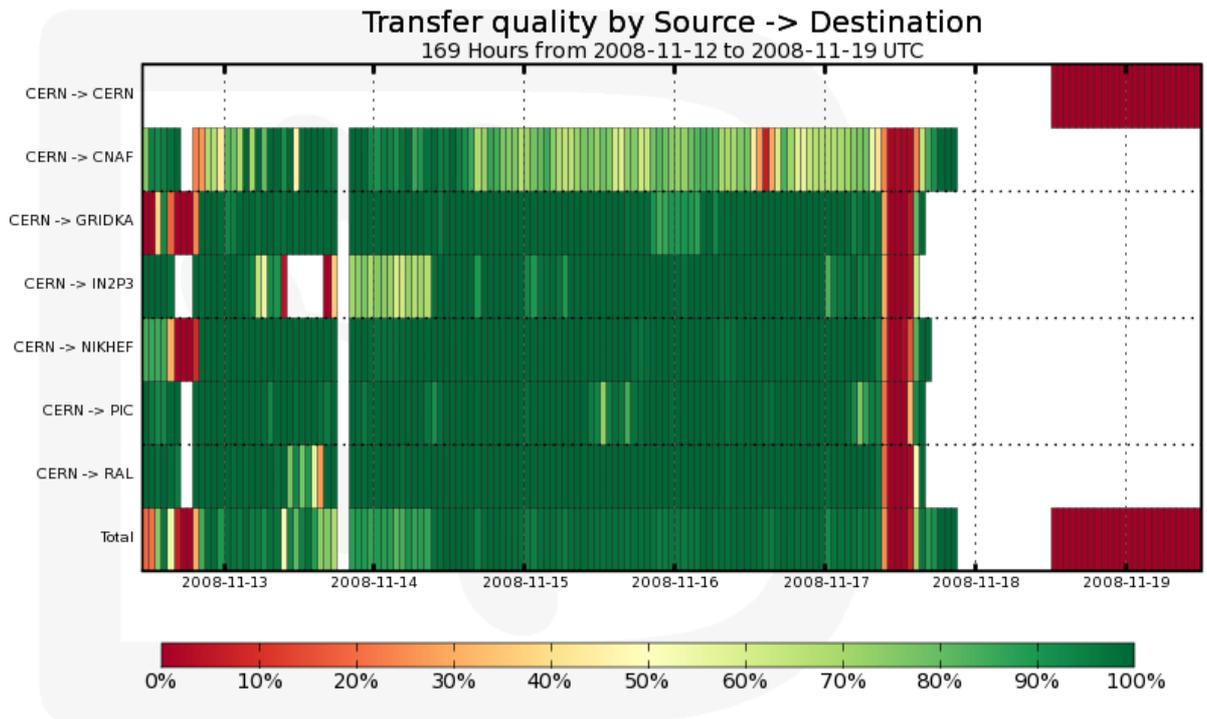
* LHCb - More precisely, concerning this staging in exercise from LHCb: A more precise breakdown per site in terms of staging in rates from MSS (in MB/s) that will be sustained in the coming days is:

- ♦ CERN 47.84
- ♦ CNAF 43.98
- ♦ GRIDKA 40.89
- ♦ IN2P3 64.04
- ♦ NIKHEF 135.42
- ♦ PIC 22.38
- ♦ RAL 30.86

These rates reflect approximately the LHCb needs +/- 50%

- LHCb - The dummy MC simulation is the main activity going on the grid (running ~4K jobs in WLCG) right now LHCb is testing the so called "Particle Gun" production and is running the file merging activity for FEST'09 preparation (few hundred of jobs though). Still DC06 stripping activity for some pending physics production (very small)
- LHCb - The FTS data transfer load generator has been stopped yesterday. (the plots





show the throughput over the last days out of CERN running steadily at about 200MB/s with a very good quality) The stage in exercise has also been stopped. Andrew will investigate on the problems seen at IN2P3 and CNAF. Jos (GridKA) reported that they have some serious problem with their tape system that would fail the exercise.

Sites round table:

- RAL (e-mail from Brian Davies) - During the downtime of the SRM-ATLAS interface to CASTOR at RAL (yesterday) (18th November) the Oracle database used by CASTOR was successfully moved from an Oracle "RAC" to a standalone Oracle database. This will enable further investigation of errors seen within the database. During the meeting Gareth reported that the move to a standalone database was completed and that errors are still being seen in this configuration.
- ASGC (pending more detailed input): we're working on the DB issue now, action planning have been proposed yesterday while we're eager to clarify why we keep hitting with same failures of the datafile, that the error have been observed before applying the action.
keep you posted tomorrow and if problem persist, we could continue with concall that could learn more input from experts with assistant from your coordination. During the meeting ATLAS requested a new conf call and this will be tried for Friday morning.

Services round table:

- Data Management: Sophie reminded sites and experiments to submit gpus tickets for problems. She remarked they had seen an LHCb RAC database server reboot during the CERN critical power tests yesterday (LFC reacted well) and was assured by Maria that this was normal. Half of the RAC servers are on critical power, the others not, so that when non-critical power goes individual servers may stop but applications should recover onto those servers on critical power. There would be degradation but no loss of service. Flavia pointed out that Andrew's tests of gfal-ls showed it can be a very slow operation in dcache sites. Jean-Philippe said this was site dependent with better performance at IN2P3. Flavia will follow up.
- Streams->FZK for LHCb LFC: (Q: what changed so that the replication to FZK apparently 'restarted' (updates to master now propagated OK - maybe the problem was always not waiting for the log file

switch - ~1hour - as noted above):

Hi,

Nothing has changed. The replication to Gridka was always working but in archived log mode instead of real time mode (as the other 5 destination sites).

The problem was observed the first time in May 2008. The propagation to Gridka was aborted with the error TNS connection lost contact being not possible to re-start it anymore. A SR with Oracle was opened but there was not any clue about the cause so in the meantime I re-created Gridka's propagation separately to fix the problem. Later on the propagation was merged with the original Streams setup.

However in July 2008 the problem occurred again. We followed the same procedure and we separated Gridka's propagation from the main Streams setup. Dawid logged the problem in the Streams postmortem wiki page: <https://twiki.cern.ch/twiki/bin/view/PSSGroup/StreamsPostMortem> Gridka's propagation was merged again with the original setup some weeks after.

But in September 2008 we had a new occurrence of the problem. We separated Gridka's propagation and decided to maintain it running separately in order to avoid the impact which this problem was causing on the other 5 replicas. This means that the replication is based since then in archived log mode so the changes are not transferred until the log file is switched, approximately every 1 hour (only for Gridka, the other tier1 sites are in real time replication). Oracle support provided a diagnostic patch, which was installed at Gridka, in order to get more information when this problem is reproduced and we are basically waiting.

Please let me know if there is any other question.

Cheers, Eva

During the meeting Jean-Philippe asked why FZK was different from the other 5 sites pointing out that, with an hours delay in updating the LFC, tests on the presence of new files would fail. Eva did not know the root cause of the difference - something specific to FZK and their LFC/Oracle environment.

- Databases: The power tests at CERN yesterday went as expected. Jeff said that NL-T1 is trying to specify performance requirements for database server upgrades, mostly for their 3D services, and asked if CERN had information such as the average size in bytes returned by conditions DB queries and the number of queries per second. Maria reported this should really come from the experiments and that even at CERN these numbers were hard to get in advance. She agreed to send Jeff some relevant pointers e.g. several talks at the recent DB workshop.

AOB:

Thursday

Attendance: local(Nick,Jan,Simone,Sophie,Harry,Gavin,Andrea);remote(Michael.Gareth,Brian,Jeremy,).

elog review:

Experiments round table:

ATLAS (SC): Site services are not yet ready for starting the 10 million (small) files transfer tests of the new ddm/dq2 and will probably not be ready next week either. Hence ATLAS will separately perform intensive tests to the RAL production endpoint and the CERN PPS using 1000 datasets of 1000 files each. The CNAF STORM front end appears to be down and there were destination problems yesterday at Lyon, since announced to be working, and SARA also working but with no announcement. The transfers to ASGC have

now dropped from 90% efficiency to 20% over a week as predicted. There will be a conference call with them tomorrow. After the CERN power cut on Thursday an ATLAS server had a broken optical link that was detected on the client side before system monitoring so this needs to be improved.

CMS (AS): The midweek global cosmics run is happening as usual. Since yesterday production transfers to FNAL have been failing with a proxy problem.

Sites round table:

RAL: From Matt Hodges following ATLAS request for rapid SL4 FTS deployment: Problems are: (1) Change in third-party (Java) dependencies has broken tomcat. See:

<https://twiki.cern.ch/twiki/bin/view/LCG/FtsKnownIssues21#BouncyCastle> for a workaround and:

<https://savannah.cern.ch/patch/?2643>

There are also dependency problems that have to be resolved by hand, breaking our kickstart scripts. The pros and cons of depending on third-party repositories have been discussed to death, but this is a major con that obviously places a burden on the sites when things break.

(2) Unable to start FTS agent processes; problem previously unknown to developers at CERN, but now reproduced there. Awaiting a fix.

Brian Davies reported they were currently seeing DB issues causing a high failure rate of the ATLAS small file transfers. Their DBA's are investigating.

Gareth asked how they should have flagged yesterdays tape problems in GOCDB when CASTOR and it's disk layer were still working. Possibilities are degraded or at risk ? Nick to follow up at ops meeting. There will be a scheduled downtime for the LHCb LFC on Monday. They plan to turn off their RBs soon - Nick thought LHCb might still need them - Harry will check (Roberto reports they are not needed).

CERN: Jan reported they had successfully upgraded CASTORATLAS and want to do the other experiments next week starting with CMS on Monday.

Services round table:

- FTS: We want to upgrade the SLC4 FTS services to the currently released patch (patch 2116) on Monday morning. This will patch several non-critical issues that were reported in the last stages of PPS testing. Patch 2551 fixing the "error message issue" (currently in PPS) is already installed at CERN. The intervention ShouldBeTransparent.
- FTS: patch 2643 (fixing the problem caused by the name change of the security BouncyCastle library and the path change of the Oracle instantclient libraries) doesn't affect CERN until we explicitly upgrade BouncyCastle or Oracle (we don't use yum). It will affect sites trying to install FTS 2.1 from yum, since you'll automatically get the new BouncyCastle / Oracle versions. Until patch 2643 is released (currently in certification) you will need to apply the workarounds documented here:
 - ◆ https://twiki.cern.ch/twiki/bin/view/LCG/FtsKnownIssues21#Oracle_Instantclient
 - ◆ <https://twiki.cern.ch/twiki/bin/view/LCG/FtsKnownIssues21#BouncyCastle>

AOB: Andrea asked about current CERN flickering bdii response as, in particular, this is causing poor SAM availability results. Sophie replied this is thought to be due to a site bdii publishing too much data but this is intermittent so hard to trace. They will recalculate the site availability numbers.

Jan reported CERN is having many false positive alarms from srmv1 tests at srmv2 endpoints as a by product of SE tests. RAL reported they have the same problem.

Nick announced GocDB will be down next Tuesday for various fixes including rationalisation of service

types. He will publicise this (see tomorrow's AOB). He also invited people to join or send issues to their bi-weekly EGEE-OSG coordination meetings - mail to nick.thackray@cernNOSPAMPLEASE.ch.

Gareth asked the status of the FTS service supporting a hot standby and Gavin said this was still being worked on.

Friday

Attendance:

local(Julia,Sophie,Laurence,Gavin,Maria,Harry,Nick,Roberto,Olof,Ignacio,Jan,Simone);remote(Michel,Michael,Jeff,C

eelog review:

Experiments round table:

LHCb (RS): From yesterdays question the RB are not needed at RAL and LHCb SAM tests against them will be stopped. Transfers into the CERN LHCb_RDST spacetoken are failing with disk full since this is setup as a T1D1 pool. LHCb would like this redefined as T1D0 (request is in a ticket). The stagein exercise at Tier 1 continues with FZK now ready to start tests and CNAF showing problems to recall files from tape. Roberto queried the status of the longterm LFC streams replication failures at FZK and Maria reported that Eva is working on this (an Oracle diagnostic patch has been deployed). Jan asked the status of an LHCb ticket about problems of their SRM in the PPS to which he had replied and was told that the problem had been traced to an LHCb client.

ATLAS (SC): 1) There was a disjointed conf call with CASTOR and DB experts and ASGC this morning which diagnosed their degraded file transfer performance as due to a corrupted stager database block. Before that ASGC had restarted some components but degradation set in again after an hour. They will act on the bad block and a new conf call is scheduled for Monday. 2) The small file transfer exercise testing new site services at RAL and CERN PPS started yesterday. That to RAL was stopped at 16.00 at their request to look at some DB issues. They reported some DB errors but an error rate (ignoring ASGC) of only 2.5%. Investigations continue. The CERN tests developed a problem overnight and Jan reported that the Castoratlas service suffered a severe degradation last night between 17:00 and 18:30 when services were available again at 90%. Service was fully restored at 21:00. A detailed summary of the different events that led to the degradation of the ATLAS Stager Service is available on <https://twiki.cern.ch/twiki/bin/view/FIOgroup/PostMortem20081120>. 3) ATLAS have a long queue for tape migration - reported to be a DB problem. 4) From Monday ATLAS will only use the SLC4 FTS servers at CERN. 5) ATLAS is draining their CNAF tasks/subscriptions in preparation for their announced (see below) 5 day downtime of the MCDISK space token.

Sites round table:

CNAF (By mail from Luca): I would like to notify that next Monday we will close the access to one of ATLAS space token (MCDISK) for ~ 5 days for maintenance operations. The involved srm end-point (storm-fe.cr.cnaif.infn.it) will be up and the other ATLAS space tokens will be available (i.e. DATADISK, USERDISK, GROUPDISK and the "old" DISK).

BNL (ME): Reported on the BNL-CNAF slow networking problem, first seen end October. After several weeks work by an impressive group of people the probable cause is the transatlantic link from Vienna to New York city. An alternative routing via Frankfurt has now been tested giving a fifty-fold improvement. The Vienna-NYC link is owned by the LHC network and there seems to be no automatic detection of this problem. Investigation had to proceed from level to level in the many components of the multi-level architecture supporting the BNL to CNAF link. This model will be difficult to sustain and this will be raised both within the LHCOPN and to the LCG Management board.

Services round table: 1) There will be a transparent (!) upgrade of the SLC4 FTS services on Monday. 2) Also Monday upgrade of CASTORCMS to level 2.1.7-23. 3) A new release of CASTOR SRM (2.7-11) is in preparation. This addresses a segment violation seen only by CMS.

AOB: Roberto asked if ATLAS are seeing successful tape recalls at CNAF but the reply was that this is not currently being exercised.

Follow up from GD group on 'flickering' CERN top=level bdii: Last night (Wed/Thur) was bad for SAM! Both the sam-bdii.cern.ch and lcg-bdii.cern.ch started to return sporadically inconsistent information (none, some, all). This resulted in test failures, the inability to publish some test results, and only a sub-set of the normal tests being launched. The problem was trace to information coming from the GRIF site. For some not well identified reasons, a misconfiguration in one of the subsite BDIIs at GRIF ended up querying itself and thus adding its entries once more at each query. A protection mechanism for such situation will be built into the top-level BDII as a result.

Follow up on ATLAS (and FIO) request for improved monitoring of optical links (problem after Tuesday vault power off) from CS group: The CERN tape infrastructure in the Vault was severely affected by a problem on a F10 router last Wednesday 19th. The problem was reported by IT/FIO at 09:16am and solved at 16:36pm. The issue was due to a 10G optic that was recognized as not genuine when the router rebooted after the power tests. As a result, the traffic in one direction was discarded by the router. As announced at the C5 of the 7/11, we are deploying a feature that detects these kind of errors. At this stage the feature is deployed on a few links to verify it's safe, but unfortunately not everywhere yet. Our apologies for the problems caused.

Follow up on GocDB Service TYPES: Here is the information on the Service Types in the GOC database. The first list is what will be in the GOC database as of next Tuesday. The second list are all those Service Types that will be deleted. All services that are registered in the GOC database with one of the old Service Types will be mapped to one of the new/remaining Service Types. This was all discussed at a couple of the previous Joint Grid Operations Meetings.

Nick.Thackray@cernNOSPAMPLEASE.ch

List 1: Service Types that will be in the GOC database as of next Tuesday

> *The list of services that will remain in GOCDB after today 3pm GMT is:*

- > - CE
- > - gLite-CE
- > - ARC-CE
- > - CREAM-CE
- > - APEL
- > - MON
- > - Site-BDII
- > - Top-BDII
- > - UI
- > - SRM
- > - Classic-SE
- > - Central-LFC
- > - Local-LFC
- > - LFC (to be discarded later)
- > - WMS
- > - RB
- > - VOMS
- > - MyProxy
- > - LB
- > - AMGA

- > - *FTM*
- > - *FTS*
- > - *VO-box*
- > - *RGMA-IC*

List 2: Service Types that will be deleted from the GOC database as of next Tuesday

SE gLite-UI VO PROX TOP-Bdii (change of case) Local-LFC ARC-CE LCFG GRIS RB-CMS classic
RB-ALICE MDS MDS-GIIS GridIce I/O server BDII-Core DATABASE RB-ATLAS RB-LHCB IC
WEB-PORTAL BDII-RB-UI RB-Core WEB-SERVICE SAMAP RB-Testzone BDII-CMS BDII-ATLAS
RLS

-- JamieShiers - 17 Nov 2008

This topic: LCG > WLCGDailyMeetingsWeek081117

Topic revision: r17 - 2008-11-21 - HarryRenshall



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](#)