

Table of Contents

Week of 111212.....	1
Daily WLCG Operations Call details.....	1
WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	4
Wednesday.....	5
Thursday.....	7
Friday.....	9

Week of 111212

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(Dan, Gavin, Lukasz, Maarten, Massimo, Torre);remote(Andreas, Burt, Gonzalo, Jhen-Wei, Michael, Paco, Rob, Rolf, Stephen).

Experiments round table:

- ATLAS reports -
 - ◆ T0
 - ◇ Friday evening: Permission denied error writing into CERN-PROD_SCRATCHDISK (GGUS:77313). On the 6th of december one user chmod most directories in scratchdisk to mode 750. (EOS allows all users with write-permission to chmod). Permissions fixed. ATLAS will follow up for better solution (POSIX says only owner can chmod).
 - ◇ Sunday: Unable to access job logs on CERN-PROD_SCRATCHDISK (GGUS:77333). "The AXIS engine could not find a target service to invoke." Was caused by wrong URL in EOS info provider. Changed to use the BeStman-style "/srm/v2/server".
 - Massimo: the value was changed 10 days ago - did ATLAS cache the value published previously in the BDII?
 - Dan: if it was changed then, why did we only see it now? we would need to look further into the matter

WLCGDailyMeetingsWeek111212 < LCG < TWiki

- ◇ Monday AM: CERN-PROD_DATADISK transfer to UK sites failed to contact on remote SRM. (GGUS:77342 [↗](#)). "ERROR] failed to contact on remote SRM [http://srm-eosatlas.cern.ch:8443/srm/managerv2]." Wrong SRM URL again?
- ◆ T1 sites
 - ◇ Sunday PM: Site NDGF-T1 has 376 deletion errors in the last 4 hours. (GGUS:77335 [↗](#))
 - ◇ Monday AM: Destination errors to SARA-MATRIX, checksum mismatch (GGUS:77336 [↗](#))
 - ◇ Monday AM: Jobs failing in NDGF-T1 (GGUS:77338 [↗](#)) "Failed in data staging: Metadata of replica and index service differ for srm://srm.ndgf..."
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Shutdown
 - ◆ CERN / central services
 - ◇ Change of UK CA causing problems for CMS User (probably first of many) GGUS:77179 [↗](#)
 - Maarten: so far there does not appear to be a way to handle this easily on the VOMRS side; we may need to handle each case manually; proposal: do the current case and record the necessary steps in the ticket - then we may judge if the necessary effort can be tolerated for this exceptional case; the total number of UK CA users in WLCG currently is 425 (ALICE: 9; ATLAS: 264; CMS: 40; LHCb: 80; DTEAM: 31; OPS: 1) and their certificate renewals probably will be spread over many months to come, so on average there would be at most a few per day to deal with, depending on the affected VO
 - ◇ Database monitoring question GGUS:77324 [↗](#), raised by some issues seen in PhEDEx transfers.
 - ◇ PhEDEx issue was due to a database lock. A certain agent from all sites was banned Saturday night. T0/1/2 agents restored shortly after, T3s still banned. Some sites need to restart their agents manually. SAV:125196 [↗](#) .
 - ◆ T0
 - ◇ Running HI prompt reconstruction.
 - ◆ T1 sites:
 - ◇ MC production and/or reprocessing running at all sites.
 - ◇ T1_TW_AS GC: Migration problems (SAV:125041 [↗](#)) and file access problems (GGUS:77047 [↗](#)). They think it is their tape system but not sure if it is hardware or software.
 - ◇ Problem of throughput from T1_DE_KIT and T1_FR_CCIN2P3 to US T2s still open (GGUS:75985 [↗](#) and GGUS:75983 [↗](#))
 - ◇ Unable to run multicore jobs at PIC SAV:125174 [↗](#)
 - ◆ T2 sites:
 - ◇ NTR
 - ◆ Other:
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ No report

Sites / Services round table:

- ASGC

Monday:

- ◆ downtime for network upgrades from Thu 0:00 to Fri 10:00 UTC
- BNL
 - ◆ also affected by BDII problem detailed in FNAL report
 - ◆ Fri evening there was a failure of 1 storage server, fixed in 1h
 - ◆ HPSS upgrade has started, will take until Fri afternoon
- FNAL
 - ◆ during last 12h the US sites have been missing from the top-level BDIIs at CERN, impacting service discovery and SAM tests
 - ◆ Rob: the last time a similar problem occurred (Nov last year) it was due to network congestion near the OSG BDII that is queried by the top-level BDIIs and the temporary workaround was to increase the timeout
 - ◆ Gavin: the timeout has been increased again (patch from 30s to 120s), while the matter is being investigated further; GGUS:77339 [↗](#)
- IN2P3 - ntr
- KIT
 - ◆ Update on file transfer issues to and from KIT: In the last weeks, Atlas and CMS reported occasional file transfer problems to and from KIT. These problems started when Atlas was switching to "any-to-any" transfers. Datasets which were formerly transferred from KIT to other T1s and then read by T2 from there are now read directly from KIT, thus using the general 10Gb/s IP uplink instead of the LHCOPN. This leads occasionally to an overload of the firewall at this uplink (which is capable of doing about 4 Gb/s), possibly increased by packet retransmits due to long latencies, mostly to American sites. As a workaround, we arranged firewall bypasses for several T2 sites and meanwhile the situation relaxed a bit. However, due to these bypasses, a large fraction of the general IP bandwidth of KIT is now used up by HEP traffic (which is sometimes exceeding by far the former "nominal" values), so this general IP uplink becomes the next bottleneck. KIT has an additional 10 Gb/s uplink to the German part of the LHCONE testbed and we observe much better transfer rates and quality to those T2s which are connected to LHCONE as well, especially to those in the U.S.. So, as a short-term or mid-term solution, it would help a lot, if more US T2s could be connected to LHCONE. KIT is of course still working together with people from German and US network providers and sites to improve the situation.
 - ◆ Michael: some doubts on LHCONE as short-term solution - it currently is a testbed that does not have connections foreseen between US sites and the EU; Edoardo has pointed out that using VLAN technology does not scale; therefore the recent architecture meeting in Amsterdam decided on a tiered structure to interconnect US, European and Asian infrastructures; people are working on that, but it will not be ready before the end of January
 - ◆ Maarten: what can ATLAS do in the meantime?
 - ◆ Michael: we avoid mixing US resources with the DE and FR clouds
 - ◆ Torre: just a matter of configuration, no development needed
 - ◆ Maarten: what can CMS do in the meantime?
 - ◆ Stephen: failing transfers are rerouted; hopefully the changes by ATLAS are already sufficient
 - ◆ Michael: ATLAS have made those changes a few weeks ago - any remaining problems would have other causes; have a look at the detailed traffic statistics
- NLT1 - ntr
- OSG
 - ◆ would like to be able to open an alarm ticket directly when observing a critical problem with top-level BDIIs at CERN (see FNAL report)
 - ◆ Maarten: alarm tickets currently are driven by experiments - your use case may have technical implications, please discuss it with the GGUS developers
- PIC - ntr
- CASTOR/EOS - nta
- dashboards - ntr

- grid services - nta

AOB: (MariaDZ) Real ALARM drills attached at the end of this page.

Tuesday:

Attendance: local(Gavin, Jan, Lukasz, Maarten);remote(Burt, Jeremy, Jhen-Wei, Rob, Rolf, Ron, Stephen, Tiju, Tore, Torre, Xavier).

Experiments round table:

- ATLAS reports -
 - ◆ T0
 - ◇ CERN-PROD: Castor access failures early this morning (GGUS:77386 [↗](#))
 - ◇ CERN-PROD: failed to contact on remote SRM (srm-eosatlas.cern.ch). Believed to be fixed but not propagating? (GGUS:77363 [↗](#), related to Sunday's GGUS:77333 [↗](#))
 - ◇ Problem with BDII publishing from OSG to WLCG/CERN BDII remains open (stopgap fix in place) (GGUS:77361 [↗](#))
 - ◆ T1 sites
 - ◇ SARA-MATRIX scheduled outage today, only brief dip in DDM functional tests thus far
 - ◇ Power surge in Taiwan data center this morning, since recovered except for FTS
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Shutdown
 - ◆ CERN / central services
 - ◇ Change of UK CA causing problems for CMS User GGUS:77179 [↗](#). Two users deleted to allow them to rejoin.
 - ◇ PhEDEx issue was due to a database lock. T3s now also unbanned. Some sites need to restart their agents manually. SAV:125196 [↗](#) .
 - ◆ T0
 - ◇ MC backfill (LHE production) will be starting
 - ◆ T1 sites:
 - ◇ MC production and/or reprocessing running at all sites.
 - ◇ T1_TW_ASGC: Migration problems (SAV:125041 [↗](#)) and file access problems (GGUS:77047 [↗](#)). They think it is their tape system but not sure if it is hardware or software.
 - ◇ Problem of throughput from T1_DE_KIT and T1_FR_CCIN2P3 to US T2s still open (GGUS:75985 [↗](#) and GGUS:75983 [↗](#))
 - ◇ Unable to run multicore jobs at PIC SAV:125174 [↗](#), should be fixed now.
 - ◆ T2 sites:
 - ◇ NTR
 - ◆ Other:
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ No report

Sites / Services round table:

Tuesday:

- ASGC
 - ◆ recovering from power cut, all services back except FTS, which has an issue with its DB
- FNAL - ntr
- GridPP
 - ◆ UK CA problem (GGUS:77179) already seen for DTEAM users, each case treated manually so far
 - ◆ Stephen: might the CA DN string just be edited in the VOMRS/VOMS DB?
 - ◆ Maarten: will ask Steve if such a hack would be an option
- IN2P3 - ntr
- KIT - ntr
- NDGF
 - ◆ dCache head node downtime Thu 09:30-10:30 UTC
- NLT1
 - ◆ SARA downtime going as planned so far
- OSG
 - ◆ because of an RSV DB problem during the weekend, availability records needed to be re-sent to SAM and corresponding availabilities recomputed, all OK now
 - ◆ reopened BDII ticket GGUS:77339 (US sites not in CERN BDIIs) because the root cause has not yet been determined
 - ◆ looking into network performance between CERN and Indiana
 - ◇ throughput from CERN to Indiana is twice the throughput from Indiana to CERN, but even the latter value should be much more than what is needed
 - ◇ during the incident in Nov last year the throughput to CERN was too low
 - ◆ note: only the CERN BDIIs appear to be affected, other top-level BDIIs were found to contain the US sites without having to increase the query timeout
 - ◇ further checks by CERN would be desirable
- RAL - ntr
- CASTOR/EOS
 - ◆ EOS information provider behavior being investigated: info appears OK locally, but does not make it into the top-level BDII
- dashboards - ntr

AOB:

Wednesday

Attendance: local(Gavin, Luca, Maarten, Maria D, Massimo, Stephen, Steve, Torre);remote(Burt, Dimitri, Jeremy, Jhen-Wei, Michael, Rob, Rolf, Ron, Tiju, Tore).

Experiments round table:

- ATLAS reports -
 - ◆ T0
 - ◇ No update since yesterday am: CERN-PROD: failed to contact on remote SRM (srm-eosatlas.cern.ch). Believed to be fixed but not propagating? (GGUS:77363, related to Sunday's GGUS:77333)
 - Massimo: looking into why the correct information does not end up in the top-level BDII
 - ◇ Problem with BDII publishing from OSG to WLCG/CERN BDII remains open (stopgap fix in place) (GGUS:77361)
 - ◆ T1 sites
 - ◇ UK ATLAS has reported the same problem with the change of UK CA already reported here by CMS (GGUS:75996 from Nov 4)

· see CERN grid services report below

- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Shutdown
 - ◆ CERN / central services
 - ◇ Change of UK CA causing problems for CMS User GGUS:77179. Ticket closed. CMS would rather delete users in this situation.
 - see CERN grid services report below
 - ◇ CASTOR reported the loss of 32 disk-only files. They were all test files so not so there is no problem.
 - ◆ T0
 - ◇ MC backfill (LHE production) will be starting
 - ◆ T1 sites:
 - ◇ MC production and/or reprocessing running at all sites.
 - ◇ T1_TW_ASGC: Migration problems (SAV:125041) and file access problems (GGUS:77047). They think it is their tape system but not sure if it is hardware or software.
 - ◇ Problem of throughput from T1_DE_KIT and T1_FR_CCIN2P3 to US T2s still open (GGUS:75985 and GGUS:75983)
 - ◆ T2 sites:
 - ◇ NTR
 - ◆ Other:
 - ◇ NTR
- ALICE reports -
 - ◆ NTR
- LHCb reports -
 - ◆ No report

Sites / Services round table:

- ASGC
 - ◆ now also the FTS has recovered from the power cut
- BNL - ntr
- FNAL - ntr
- GridPP
 - ◆ concerned about ongoing UK CA issue
 - ◇ see CERN grid services report below
- IN2P3 - ntr
- KIT - ntr
- NDGF - ntr
- NLT1
 - ◆ yesterday evening there was a problem with 1 dCache pool node, fixed this morning
- OSG
 - ◆ BDII issue: have exhausted the tests and checks on the OSG side for now, looking forward to clues from the CERN side
 - ◇ Gavin: we are looking into the LDAP query times, which show intermittent behavior; at some point we will involve the network team
- RAL - ntr
- CASTOR/EOS
 - ◆ CASTOR SRM upgraded to version 2.11 for ALICE and LHCb, transparent; ATLAS and

- ◆ CMS will be done tomorrow
 - ◆ CMS EOS upgrade postponed, probably until early January
- dashboards - ntr
- databases
 - ◆ integration DBs for CMS and LHCb were upgraded to Oracle 11g
 - ◆ downstream capture DBs for LHCb were upgraded to 11g, mostly OK, 1 problem for KIT being worked on
 - ◆ downstream capture DBs for ATLAS being upgraded to 11g
 - ◆ a patch has been applied to the CMS offline production DB to fix a library cache locking issue that was observed in the last few months
- GGUS/SNOW
 - ◆ see AOB
- grid services
 - ◆ CERN CVMFS. The CERN stratum one service **cmvfs-stratum-one.cern.ch** will start using a new but "identical" stratum 0 service from tomorrow (Thu 15th) morning at 09:00 UTC. This is expected to be **100% transparent** to clients other than those clients connecting to CERN may see updates to repositories up to 1 hour quicker than before. The other Stratum1s at BNL and RAL will be requested to switch next week if all is well. IT SSB [↗](#)
 - ◆ UK CA migration.
 - ◇ The UK CA is migrating from the 'old' CA to the new '2B' CA:
 - /C=UK/O=eScienceCA/OU=Authority/CN=UK e-Science CA
 - /C=UK/O=eScienceCA/OU=Authority/CN=UK e-Science CA 2B
 - ◇ Note: there also is a '2A' CA, but it appears not to be used for WLCG.
 - ◇ To ease the migration of individual UK users who have not performed the necessary steps themselves, a secondary User-DN-new-CA-DN pair will be added to the identities of existing users with a User-DN-old-CA-DN.
 - ◇ (Un)fortunately a test of the process today at around 09:30 UTC with **CMS** was actually executed with 74 UK CMS members now having additional '2B' identities. CMS VO managers have been informed.
 - ◇ Following any feedback from CMS this will be repeated for all other VOs on Monday.
 - ◇ Note this does not include dteam VO which is now hosted by hellasgrid.gr.

AOB: (MariaDZ) ATLAS user ticket GGUS:77118 [↗](#) was escalated today and got adequate attention by supporters. This note is just to remind people of this useful GGUS functionality.

Thursday

Attendance: local(Eva, Gavin, Lukasz, Maarten, Maria D, Massimo, Torre);remote(Andreas M, Burt, Gareth, Giovanni, Jeff, Jhen-Wei, Mette, Rob, Rolf).

Experiments round table:

- ATLAS reports -
 - ◆ T0
 - ◇ CERN-PROD: SRM service URL issue, progressing on resolving BDII problem. Some instances behind the load balancer not getting correct info (GGUS:77363 [↗](#) and GGUS:77333 [↗](#))
 - ◇ Problem with BDII publishing from OSG to WLCG/CERN BDII (stopgap fix in place) -- no recent updates, can someone please update when there are concrete developments (GGUS:77361 [↗](#))
 - ◆ T1 sites
 - ◇ RAL unscheduled downtime with SRM (DB) problems (preceded by network outages) (GGUS:77470 [↗](#))

- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Shutdown
 - ◆ CERN / central services
 - ◇ From time to time all sites fail SAM visibility test, they are disappearing from BDII
 - Gavin: there are 2 issues that look unrelated
 1. OSG BDII queries take too long
 - ◆ we are improving the monitoring to be able to present evidence to the network team
 2. lcg-bdii.cern.ch and sam-bdii.cern.ch instabilities (GGUS:77452)
 - ◆ random subsets of the nodes are often seen serving partial results for a while
 - ◆ that problem looks new and much more urgent
 - ◆ we have contacted the developer while collecting more evidence and working on mitigations
 - ◆ T0
 - ◇ MC backfill (LHE production) will be starting
 - ◆ T1 sites:
 - ◇ MC production and/or reprocessing running at all sites.
 - ◇ T1_TW_ASGC: Migration problems (SAV:125041) and file access problems (GGUS:77047). They think it is their tape system but not sure if it is hardware or software.
 - ◇ Problem of throughput from T1_DE_KIT and T1_FR_CCIN2P3 to US T2s still open (GGUS:75985 and GGUS:75983)
 - ◆ T2 sites:
 - ◇ NTR
 - ◆ Other:
 - ◇ NTR
- ALICE reports -
 - ◆ NTR
- LHCb reports -
 - ◆ No report

Sites / Services round table:

- ASGC - ntr
- CNAF - ntr
- KIT - ntr
- FNAL
 - ◆ yesterday there was a large network outage, but CMS services appear to have been unaffected
- IN2P3
 - ◆ still looking into the network issues affecting CMS and others; in contact with GEANT and a site in China; the matter is complicated
- NDGF - ntr
- NLT1
 - ◆ yesterday evening and during the night there were crashes of dCache pool nodes due to a new kernel being incompatible with their network cards; all OK now after downgrading
- OSG
 - ◆ waiting for news on pending CERN BDII issue
 - ◆ Maria D: the DOEGrids CA hopes to ramp down its activities next year - how will OSG deal with that and what would be the time line?
 - ◆ Rob: no changes in the next few months; we are trying to get a commercial CA into the IGTF

to replace the DOEGrids CA

◆ details at <https://twiki.grid.iu.edu/bin/view/Security/OSGCATransition2012>

- RAL

- ◆ there were 2 site network outages yesterday evening, between 9 and 10 pm and between 11 pm and midnight
- ◆ during the night there were DNS issues affecting the ATLAS SRM
- ◆ in the morning there was a severe problem with the Oracle RAC (Resource Manager taking all the resources), which led to the decision to migrate the ATLAS SRM DB to new HW today, while for other DBs the migration remains planned for the new year
- ◆ the outage will be turned into an at-risk downtime at 16:00 UTC, we will throttle the FTS for now, ramping it back up tomorrow when all looks OK
- ◆ further details in [GGUS:77470](#)

- CASTOR/EOS

- ◆ CASTOR SRM 2.11 upgrades done for ATLAS and CMS, transparent
- ◆ Jan 11 morning: major name server upgrade, all CASTOR instances unavailable
- ◆ upgrades to Oracle 11g on new HW, ~2 h downtime per intervention:
 - ◇ Jan 16: ATLAS
 - ◇ Jan 17: CMS
 - ◇ Jan 23: ALICE
 - ◇ Jan 24: LHCb

- dashboards - ntr

- databases - ntr

- GGUS/SNOW - ntr

- grid services

- ◆ see CMS report
- ◆ Jan 9: LSF maintenance to move master to new HW, batch down for a few h

AOB:

Friday

Attendance: local(Edoardo, Gavin, Lukasz, Maarten, Massimo);remote(Alexander, Burt, Gareth, Giovanni, Jhen-Wei, Michael, Rob, Rolf, Stephen, Torre, Ulf, Xavier).

Experiments round table:

- ATLAS reports -

- ◆ Computing operations

◇ ATLAS has a heavy need of simulation in the coming weeks. Sent a request to all sites to please ensure full operation over holidays and investigate whether additional resources are available for temporary use

· Maarten: ATLAS can only count on a holiday service level

· Maarten: note that for **new** releases the CVMFS infrastructure currently relies on a machine that is not in the CC yet - if there is a power cut during the holidays, that machine may not be recovered quickly

- there will be a scheduled power cut on Jan 3

- **existing** releases do not rely on that machine

- ◆ T0

◇ CERN-PROD: SRM service URL issue, progressing on resolving BDII problem.

Some instances behind the load balancer not getting correct info ([GGUS:77333](#))

· Massimo: the problem with the info provider looks fixed, the investigations were hampered by the other BDII problem (see below)

- ◇ Problem with BDII publishing from OSG to WLCG/CERN BDII -- no recent updates, can someone please update when there are concrete developments (GGUS:77361 [↗](#))
 - Still seeing info propagation failures in BDII of US software release info to WLCG/CERN BDII. Appears lengthened timeout workaround not fully effective
 - Gavin: BDII developer Laurence Field has analyzed the instabilities and determined the increased timeout to be the cause
 - the timeout applies to all site BDIIs that are queried
 - there always are a fraction of site BDIIs that time out
 - ◆ machines could e.g. be down for maintenance
 - the complete query loop got delayed so much that live records in the BDII were becoming stale
 - stale records are deleted by default
 - a caching mechanism allows such records to be kept with all status attributes reset to "unknown"
 - that mechanism will be switched on now
 - ◆ it was in the pipeline anyway, for all top-level BDIIs in WLCG
 - we need to solve the OSG BDII query performance issue ASAP
 - ◆ **NOTE: see the IN2P3 report below for the resolution !!!**
 - Torre: while this matter is not solved, we have an internal workaround in the ATLAS information system
 - ◆ T1 sites
 - ◇ Taiwan scheduled outage for network maintenance extended 24hrs to tomorrow morning
 - CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Shutdown
 - ◆ CERN / central services
 - ◇ NTR (any update on SAM BDII?)
 - see ATLAS report
 - ◆ T0
 - ◇ MC: LHE production has started (writes to EOS /store/lhe)
 - ◆ T1 sites:
 - ◇ MC production and/or reprocessing running at all sites.
 - ◇ T1_TW_ASQC: Migration problems (SAV:125041 [↗](#)) and file access problems (GGUS:77047 [↗](#)). They think it is their tape system but not sure if it is hardware or software.
 - ◇ Problem of throughput from T1_DE_KIT and T1_FR_CCIN2P3 to US T2s still open (GGUS:75985 [↗](#) and GGUS:75983 [↗](#)). Reportedly getting worse again from KIT. Is there any update?
 - Xavier: updates will be posted in the tickets
 - Michael: ATLAS multicloud production was stopped for the DE and FR clouds, but user subscriptions are allowed to transfer data from any T2 to the ATLASSCRATCHDISK space at KIT; the amounts of such data vary, we do not have a lot of statistics for that
 - ◆ T2 sites:
 - ◇ NTR
 - ◆ Other:
 - ◇ NTR
- ALICE reports [↗](#) -

◆ NTR

- LHCb reports -
 - ◆ No report

Sites / Services round table:

- ASGC
 - ◆ downtime extended to Sat morning (UTC) because the vendor needed more time
- BNL
 - ◆ MSS upgrade to HPSS 7.3.3 proceeding well, tests for ATLAS look OK, expect to be fully operational by the end of the day
- CNAF
 - ◆ upgrade of last CREAM CE to EMI release in progress
 - ◆ memory leak in BLAH identified, fix is available, new version will be released soon
- FNAL
 - ◆ how would BDII caching (see ATLAS report) affect availability calculations?
 - ◆ Maarten: as the state of any affected service would be reset to "unknown", WMS job submission to affected CEs would fail, since the "Production" state is required by default (that could be changed, but would have undesirable consequences)
 - ◇ the SAM team would have to correct availabilities for any affected periods
 - ◆ Burt: only the SAM tests seem affected by such matters, production carries on fine
 - ◇ added after the meeting: user jobs submitted via the WMS would also be affected
- IN2P3
 - ◆ investigating high traffic from worker nodes to AFSDB/KDC servers at CERN, causing CERN firewall overload
 - ◇ no ticket, AFS experts at both sites involved
 - ◇ suspect a bug in AFS client 1.6.0
 - ◇ this is not related to the AFS callback timeout problem discussed a few weeks ago
 - ◇ request rates exceeding 1 kHz
 - ◆ Edoardo: the problem started ~Tue and the firewall performance decreased steadily since
 - ◇ users started complaining today
 - ◇ at 14:00 CET the traffic in question was moved to the bypass
 - also from other sites
 - ◇ the firewall performance immediately jumped back to normal levels
 - ◆ Massimo: is the smooth linear increase of the problem understood?
 - ◆ Edoardo: no
 - ◆ Rolf: we will reboot part of our cluster and downgrade another part to see the effect of each intervention on the traffic
 - ◆ after the meeting: **the time line of this issue agrees with onset of the BDII problem !!!**
 - ◇ strange that only OSG BDII queries seemed to be affected, but the CERN BDIIs look OK again...
- KIT - ntr
- NDGF
 - ◆ tomorrow starting at 23:00 UTC an OPN maintenance will make 1 site in Sweden unavailable for 5 h, affecting ATLAS and ALICE data
- NLT1 - ntr
- OSG
 - ◆ network graphs have been added to BDII ticket GGUS:77339, drops are observed in traffic from CERN to Indiana
 - ◇ see IN2P3 report
- RAL
 - ◆ CASTOR OK since yesterday, FTS ramped up to 100% this morning
 - ◆ CASTOR DB outage for CMS and ALICE between 6 and 9am

- ◆ Thu Jan 5 outage to move CASTOR DB to new HW, will take a few h
- CASTOR/EOS
 - ◆ reminder of upgrades to Oracle 11g on new HW, each taking a few h:
 - ◇ Jan 11 morning: name server, all CASTOR instances unavailable
 - ◇ Jan 16: ATLAS
 - ◇ Jan 17: CMS
 - ◇ Jan 23: ALICE
 - ◇ Jan 24: LHCb
- dashboards - ntr
- grid services
 - ◆ see ATLAS report
 - ◆ CERN VOMS. On monday morning for VOs LHCb, OPS, ALICE and ATLAS additional "2B" CA identities will be added for existing UK members who do not already have a "2B" identity. This is a repeat of the CMS process executed on Wednesday.
 - ◇ Users with a second identity added will **not** receive an email about this addition unlike I had previously suggested.
- networks - nta

AOB:

-- JamieShiers - 28-Nov-2011

This topic: LCG > WLCGDailyMeetingsWeek111212

Topic revision: r22 - 2011-12-16 - MaartenLitmaath



Copyright &© 2008-2022 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