

Table of Contents

ADCOperationsWeeklySummaries2018.....	1
Instructions.....	2
Reports.....	3
11 December - 17 December.....	3
03 December - 10 December.....	3
12 November - 19 November.....	3
05 November - 12 November.....	3
29 October - 05 November.....	3
22 October - 29 October.....	3
15 October - 22 October.....	3
17 September - 24 September.....	4
28 August - 03 September.....	4
21 August - 27 August.....	4
14 August - 20 August.....	4
7 August - 13 August.....	5
31 July - 6 August.....	5
23 July - 30 July.....	5
17 July - 22 July.....	5
10 July - 16 July.....	6
3 July - 9 July.....	6
12 June - 18 June.....	6
4 June - 11 June.....	6
29 May - 4 June.....	7
21 May - 28 May.....	7
14 May - 21 May.....	7
07 May - 14 May.....	7
01 May - 07 May.....	7
17 April - 23 April.....	8
10 April - 16 April.....	8
2 April - 9 April.....	8
19 March - 26 March.....	9
12 February - 19 February.....	9
5 February - 12 February.....	9
29 January - 5 February.....	10
22 January - 29 January.....	10
15 January - 22 January.....	10
8 January - 15 January.....	10
18 December - 8 January.....	11
ADCOperationsWeeklySummaries2017.....	11
ADCOperationsWeeklySummaries2016.....	11

ADCOperationsWeeklySummaries2018

Instructions

- Reports should be written on Monday morning by the CRC summarising the previous week's ADC activities and major issues
- The WLCG SCOD should copy the report to the WLCG Operations Meeting Twiki after 14:30 CET on Monday
- The issues are reported to the WLCG Operations Meetings at 15:00 CET on Monday by the CRC or by Ivan G and/or Tomas J.
- ATLAS-internal issues should be documented daily in the ADCOperationsDailyReports pages.

Reports

11 December - 17 December

Very quiet and stable operations

- BNL FTS upgrade looks good and switching traffic back to that instance
- no campaigns planned over Christmas or New Year
- thanks to all sites for collaboration over 2018 and Merry Christmas from ATLAS Computing

03 December - 10 December

Smooth running after end of 2018 data-taking.

- Production benefits from CERN-P1 and T0 and uses currently with analysis more than 400k slots
- Low level of errors, few due to Arc Control Tower-Panda communication instabilities
- Transfer rate is stable at 15-20GB/s
- Lot of data disk full and big deletion campaign ongoing (400-700 files/h)

12 November - 19 November

- ♦ HI data taking - as expected i.e.
 - ◊ · Limiting factor is writing to tape - 2.5 GB/s
 - In case of higher LHC efficiency, the SFO-CASTOR handshake will be switched off.
 - For the time being all data is saved successfully.
- ♦ DB: ADCR problems on Thursday - solved. GGUS:138318 [↗](#)
- ♦ NDGF Rucio problems reported last week - solved.

05 November - 12 November

- ♦ Production
 - ◊ Continuing switching to unified queues and Harvester for fully dynamic sites (ongoing for FR, CA, US and NL cloud)

29 October - 05 November

- NTR

22 October - 29 October

- Production - 340k / 381k slots
 - ♦ HLT farm in GRID production - 95k slots since Sunday.
- Storage & Transfers - NTR

15 October - 22 October

- Transfers
 - ♦ We have saturated several links over the weekend with a pile-up premixing test workflow (CERN-PIC, IN2P3). The workflow is aborted and the problem should have been mitigated

17 September - 24 September

- Usual activities:
 - ◆ Stable running at 300-400k run slots depending on opportunistic resources (CERN@P1 for instance)
 - ◆ Lately lower transfer traffic than usual (small reprocessing tasks finished)
 - ◆ Some UK moved sites from APF to Harvester submission
 - ◆ Few T1s full: lifetime model will be run soon
- Problems
 - ◆ CERN reboot campaign impacted production a bit on Tuesday evening

28 August - 03 September

- Production:
 - ◆ 295 / 302 k running slots
 - ◆ Moving IT sites from APF to Harvester submission

21 August - 27 August

- Activities:
 - ◆ normal activities
- Problems
 - ◆ T1_DATADISKS getting full
 - ◇ there are few PB of small files to delete - they are now mixed with bigger files so deletion can keep up but it still slows down deletion
 - ◇ this will take longer time
 - ◆ unavailable files at INFN-T1_MCTAPE (GGUS:136823 [↗](#)) - they were on a list of files to be recalled which had not been processed since August 6th
 - ◆ lost files at INFN-T1 tapes - 95 files on datatape and 35 files on mctape - DDM ops will process the list
 - ◆ transfers from pic to all clouds fail with "Transfer canceled because the gsiftp performance marker timeout of 360 seconds has been exceeded, or all performance markers during that period indicated zero bytes transferred" (GGUS:136820 [↗](#))
 - ◇ failures from disks - some dCache pools were overloaded, the maximum number of active movers reduced
 - ◇ failures from tape - a hardware issue with one of tape recall pools
 - ◆ RAL IPv6 problem
 - ◇ during Thursday evening we observed increasing number of queued requests in rucio
 - ◇ on Friday, it turned out rucio is slowed down by slow submission to RAL FTS
 - ◇ sites using RAL FTS moved to CERN FTS (which overloaded CERN FTS)
 - ◇ IPv6 problem solved, RAL switched from CERN FTS Saturday morning
 - ◇ timeouts appeared again on Saturday afternoon and were solved in the evening

14 August - 20 August

- Activities:
 - ◆ normal activities
- Problems
 - ◆ T1_DATADISKS getting full in the last week - there are few PB of small files to delete - they are now mixed with bigger files so deletion can keep up
 - ◆ EOSATLAS namespace powercycle incident (OTG:0045385 [↗](#)) - after the issue was resolved, transfers started working but the deletion was still failing (GGUS:136727 [↗](#))

- ◆ central service monitoring grey (GGUS:136733) - there was an issue with the meter cluster during Thursday night, it was solved during morning, backlog processing finished in the afternoon
- ◆ Taiwan-LCG2 power cut (GGUS:136690) - solved
- ◆ RAL ECHO problems with swapping nodes (bug in Ceph) - solved on Friday
- ◆ transfers from pic to all clouds fail with "Transfer canceled because the gsiftp performance marker timeout of 360 seconds has been exceeded, or all performance markers during that period indicated zero bytes transferred" (GGUS:136778) - started this morning

7 August - 13 August

- 290 - 350 grid jobslots used, peak of 500k jobslots used on HPCs (CORI)
- changes in pilot led to changes in protocol used to access some resources
 - ◆ CERN EOS used via srm by some jobs
 - ◆ fixed planned for this week (change in rucio)
- wrong ST used at PIC for a small fraction of files, issue understood
- TRIUMF tapes blacklisted for writing (for a period of migration)
- T1_DATADISKS getting full, deletion in progress

31 July - 6 August

- 250 - 350 grid jobslots used plus some HPC (Titan) resources are also used
- a new procedure to restart rucio web servers helped to avoid midnight problems
- another 2 sites hit by CREAM CE update issue described in https://ggus.eu/index.php?mode=ticket_info&ticket_id=136074

23 July - 30 July

- variable production between 270-350k slots
 - ◆ almost 100k from CERN-P1 for most of week (but new installation and scalability issues)
 - ◆ problems with some big sites and their local storage
 - ◆ problems with grid CE for CERN-T0 resources (update & bouncycastle)
 - ◆ updates related to the pilot causing job failures
 - ◆ issues with production input file transfer rate
- overloaded rucio readers & problems with automatic midnight restarts (new implementation in pipeline)
- data reprocessing campaign - more inputs transferred from tape
- storage at INFN-T1 unstable and down several times during last week
- automatic downtime synchronization from OIM/GOCDDB to AGIS not working for some SEs (requires manual blacklisting)

17 July - 22 July

- Overall - no major issues and smooth running.
 - ◆ Grid running with around 260k-330k
 - ◇ Mixture of simulation (major part) and some data reprocessing
- Problems:
 - ◆ Quite week with a few smaller issues
 - ◆ Issue at NFN-T1 which had no production at the beginning of the week (now back to normal)
- If comments : crc.shifter@cernNOSPAMPLEASE.ch

10 July - 16 July

- smooth production without major issues, ~ 280-300k job slots in average (+ HPC)
 - ◆ changes in job brokerage (queued vs. running) caused few times small drop in production jobs for some corner cases (e.g. not enough activated jobs for HPC)
 - ◆ few minor issues / glitches with rucio (production input not transferred on Sunday, expired file deletion stuck on Saturday, rucio replica location 1 hour failure)
- transfers to NET2 (BU_ATLAS_Tier2) and huge number of deletions overloading their BeStMan storage (still not completely understood where they comes from)
- one storage server from RAL old CASTOR storage not recoverable (80k files, almost no primary data, already migrated)
- BNL developed & applied procedure to detect / remove ATLAS dark data for their dCache storage
- some information from GOCDB downtime calendar not correctly propagated to our storage endpoint blacklisting (e.g. SRM-less site)

3 July - 9 July

- Overall - smooth running.
 - ◆ Slots running around 300k
 - ◇ Production running smoothly, mostly simulation this week
- Problems:
 - ◆ No major issues during this week, smaller issues
 - ◆ Issue with the pilot at INFN-T1 (GGUS:135986)
 - ◆ Problem at site NIKHEF-ELPROD causing a few errors (GGUS :135978)

12 June - 18 June

- Nicolò Magini reporting as ATLAS CRC for last week - no CRC this week
 - ◆ GGUS:135679 still thought that I was CMS, fixing.
- Stable operation with peaks > 350k slots including Sim@P1 resources during MD/TS
 - ◆ Current reprocessing campaign might cause heavy load on squid infrastructure, looks OK for now, keeping an eye on it.
 - ◆ Currently high job failure rate on INFN-T1_MCORE queue with input read error GGUS:135692 - possibly files on gpfs are not accessible from WNs ATM?
- Transfers: 11 GB/s
 - ◆ Issue with high request load on FZK-LCG2 tape for staging of input for reprocessing GGUS:135651 - ongoing work by FZK and ATLAS people to make the queue more sustainable
 - ◆ INFN-T1 storage issue for few hours on 13/06 GGUS:135658 - fixed
 - ◆ On 15/6 BNL FTS was unresponsive for two hours - came back before ticket submitted.
 - ◆ A couple of issues with transfer submission on Rucio side (problem with urllib dependency, expired proxy delegation, slow daemon), fixed.
 - ◆ According to space monitoring FZK-LCG2_SCRATCHDISK is full of dark data after deletion on Fri 15/6 - either deletion left behind dark data, or space reporting is inconsistent.
- Started tape carousel tests - first round done at BNL, now involving other T1s

4 June - 11 June

- Stable operation at:
 - ◆ Slots: 288k / 330k
 - ◆ Transfers: 8 GB/s

- Starting tests to switch the Spanish Cloud to Harvester (the new ATLAS work load management system)

29 May - 4 June

- Overall production smooth running
- ATLAS concern of EOS stability
 - ◆ Response and followup support excellent. CERN-IT please comment on the general issue.
 - ◆ For example, looking at month of May we have significant outages reported on 4th, 9th, 25th, 31st.
- Friday night Rucio proxy delegation issue caused all FTS transfers to fail
 - ◆ Fixed Saturday morning, issue understood
 - ◆ couple of misguided tickets opened, apologies

21 May - 28 May

- Overall production smooth running
- INFN-T1 temporary storage configuration and tape endpoint issues
- CERN eosatlas crashed on Thursday at 10:10 and was not available until 19:45. There were some delays in transfers of files and production at CERN crashed but not more disturbances found

14 May - 21 May

- Overall smooth running
- Production jobs
 - ◆ Production running fine at 315k jobs and good contributions from HPC
 - ◆ Slightly increased rates of jobs failing, but sites reacting to tickets even during weekend holidays
 - ◆ Recent special productions: completion of data17 reprocessing and production of derivations mostly completed
- Transfers/Storage
 - ◆ Transfer rates increased up to 12 GBps, 35 files/s due to special productions
Increased level of transfer failures 3.7 files/s due to known issues
FZK SCRATCHDISK full, broken pool at PIC
 - ◆ Tape library at Triumpf optimized (added tape drives) for recall of older data17 data

07 May - 14 May

- - ◆ No CRC
 - ◆ Production - 297k / 330k slots
 - ◇ Full steam T0 utilization
 - ◇ Bug in pilot for Titan HPC lead to wrong output guid. Problem fixed and produced files are being corrected
 - ◇ Running out of simulation drained 100k cores in the last 12 hours. Recovering.

01 May - 07 May

- Overall smooth running
 - ◆ for jobs: above 300k running jobs slots
 - ◇ relatively low level of analysis, production running smoothly
 - ◆ and for data management

◇ transfer rate for jobs was increased during the week and reached almost 2,3 PB a day, roll back on Friday to normal level to ensure good performance for T0 data export

- 2 important issues during the week
 - ◆ ATLAS eos crashed on Friday morning due to high load
 - ◇ caused T0 export to go down
 - ◇ high load source understood, alternative being worked on
 - ◆ voms-proxy operational troubles during the week (proxy was not renewed)
 - ◇ interfere with panda on Wednesday May 2: all analysis queues blacklisted
 - ◇ interfere with rucio this weekend: all deletion and transfers stopped on Sunday morning
 - ◇ origin of problems comes from the deployment by CERN-IT of V3 voms package supplied by JAVA from EPEL
 - ◇ problem fixed rolling back to UMD version of voms-proxy

17 April - 23 April

- Production
 - ◆ 304k / 489k slots average usage
 - ◆ Added ~5500 cores to T0 cluster
 - ◆ Decided to attach TAPE end points to user analysis queues

10 April - 16 April

- - ◆ Overall - no major problems.
 - ◆ Production
 - ◇ Jobs: 305k/560k slots
 - ◇ Several generator job failures found (Sherpa, Pythia) - under investigation.
 - ◆ Transfers / Storage
 - ◇ xrootd 4.7 incompatible with cache <= 2.16.48 ([link](#)). Please upgrade your cache version. Workaround is to force using xrootd 4.8.1
 - ◇ Testing usage of T0 resources with SMT - on.
 - ◇ EventService is creating datasets with 800k files which kills the DB. Working on solution.
 - ◇ 200k / 80k files on EOS that might be damaged - The impact assessed to be low. Looking at adding MD5 checksum to the files metadata.

2 April - 9 April

- Activities:
 - ◆ Grid is full with mainly light activities: event generation and simulation
 - ◆ Some event generation tasks have rather high failure rates and/or use large amounts of disk I/O. Being investigated by experts.
- Problems:
 - ◆ ADCR Oracle database serving all ATLAS grid-related services suffered under high load, leading to a couple of outages last week
 - ◆ CERN network outage on Thurs morning did not affect operations too much (HammerCloud blacklisting was disabled during the outage)
 - ◆ Certificate renewal season: one service mistakenly did not have the correct certificate renewed which led to a slight draining of the grid over the weekend
 - ◆ EOS reported potential corruption of files for an 8h period on 30 March - files may be corrupted even if Adler32 checksum is correct.
- Reminder:

- ◆ ATLAS is looking forward to the 2018 pledge deployment. New disk space should be put into the ATLASDATADISK token.

19 March - 26 March

- Activities:
 - ◆ normal activities
 - ◆ ran low on simulation last week while waiting for inputs to be produced which affected sites which can only run simulation
- Problems:
 - ◆ No major issues after CERN network intervention (some job failures as expected)
 - ◆ No other major issues

12 February - 19 February

- Activities:
 - ◆ normal activities
 - ◆ second half of reprocessing started
- Problems
 - ◆ some overlay tasks are causing frontier degradation - cap on number of jobs decreased
 - ◆ rucio overload - sites could see decrease in number of mcore jobs and increase in transferring jobs
 - ◆ Deletion at BNL failed (GGUS:133551 [↗](#)) - configuration updated
 - ◆ Transfers to RAL-LCG2-ECHO fail with "Address already in use" (GGUS:133399 [↗](#)) - fixed
 - ◆ Transfer failures from INFN-T1 via RAL FTS (GGUS:133320 [↗](#))
 - ◆ Transfer from CERN-PROD_datadisk fail with "No such file or directory" (GGUS:133414 [↗](#)) - still being investigated
 - ◆ Transfers from IN2P3 -CC_DATATAPE fail with "Changing file state because request state has changed" (GGUS:133545 [↗](#))
 - ◆ Transfers from TAIWAN-LCG2_DATADISK fail with SRM_INVALID_PATH (GGUS:133546 [↗](#)) - files not stored into the DPM successfully, will be declared lost

5 February - 12 February

- Activities:
 - ◆ normal activities
 - ◆ HepOSlibs update - blas library needed for user analysis - sites were informed about update
- Problems
 - ◆ slow transfers to BNL (GGUS:133295 [↗](#)) - caused by dCache bug [[#9341](http://www.dcache.org)]
 - ◆ timeouts to FZK-LCG2 tapes (GGUS:133332 [↗](#)) - tapes are served by RAL FTS now and limit is set to 100 connections
 - ◆ Transfers to CERN-PROD_PERF-MUONS failed with "Permission denied" (GGUS:133348 [↗](#)) - permissions fixed
 - ◆ Transfers from SARA-MATRIX are failing with "File is unavailable" (GGUS:133407 [↗](#)) - a diskpool on one of nodes wasn't started
 - ◆ Transfers to RAL-LCG2-ECHO fail with "Address already in use" (GGUS:133399 [↗](#)) - site ran out of available ports - fix needs changes on central router, so it will be done this week
 - ◆ Transfer from CERN-PROD_datadisk fail with "No such file or directory" (GGUS:133414 [↗](#)) - we need to find out why the file(s) unavailable
 - ◆ CNAF flooding - problem with lack of communication/response, e.g. GGUS:133320 [↗](#)

29 January - 5 February

- RAL FTS
 - ◆ switched ND, NL and UK clouds to use production instance lcgfts3.gridpp.rl.ac.uk on January 30
 - ◆ unfortunately production service run out of available space on Sunday
- slow transfers to BNL killed after 60 minutes by SRM (< 1MB/s)
 - ◆ transfers of big files killed and restarted
 - ◆ a lot of rebalancing activity
 - ◆ trying to tune site connection limit in FTS (lowered from 700/1000 to 400)
- EOSATLAS crashed on few times affecting job @ CERN (fixed)
- CNAF tape library back online - successful transfer from CERN to CNAF tape

22 January - 29 January

- kernel update & reboot of CERN machines
 - ◆ did not caused any major problems
 - ◆ some of our services were in same availability zone
- CERN-P1 in downtime for cooling maintenance => ~ 40k job slots sorter for our production
- preliminary tests with/without meltdown & spectre patches - 2.5-3% for reconstruction
- RAL FTS server issues ~ 1M transfers stuck
 - ◆ started on Thursday but by that time one US site in downtime was a suspect (GGUS:133066)
 - ◇ was not automatically blacklisted for ATLAS transfers
 - ◇ on Friday we changed ND & IT cloud to use other FTS, today morning also UK
- RAL CASTOR performance issues (necessity to reindex database) after deletion backlog caused by removing 600TB from storage
- tape buffers - currently needs manual blacklisting to throttle transfers it case of full buffer
 - ◆ more transfers to tape between sites because of changes in our policy (tape & disk replica can't be on one site)

15 January - 22 January

- WLCG report
 - ◆ Frontier bug encountered (old bug from 2015) where squids serve a stale error result, workaround is to flush squid cache. Bug being handled.
 - ◆ Case of T2 site running DPM on SL5, offline/unresponsive pending upgrade, being handled now but should remind all T2 about MoU committments.

8 January - 15 January

- - ◆ No CRC
 - ◆ No major issues
 - ◆ Workflow management
 - ◇ Stable operation ~320k slots. Including HLT and T0 resources
 - ◆ Data transfer
 - ◇ CNAF recovery:
 - Pointed to 4 data16 RAW files lost on Castor - under investigation.
 - We continue with exporting data15 RAW from Castor (finished data16 and data17)
 - ◇ RAL-600k-lost-files incident - All files are declared bad.

18 December - 8 January

- No CRC
- Workflow Management
 - ◆ Relatively smooth operation
 - ◆ 20-25th December: evgen very small job with a lot of outputs, created serious troubles to FTS.
 - ◆ Currently - not many score / evgen. Under investigation
 - ◆ Reprocessing finished very smoothly.
- Data Transfers
 - ◆ Relatively smooth operation
- Meltdown and Spectre
 - ◆ Very quick tests with ATLAS simul/reco, degradation for reco is about 5% (Dekstop, no VM, using singularity, Intel(R) Core(TM) i7-3770K CPU @ 3.50GHz4 core/8 HT, 8-core job). Under further investigation.

ADCOperationsWeeklySummaries2017

<https://twiki.cern.ch/twiki/bin/view/AtlasComputing/ADCOperationsWeeklySummaries2017>

ADCOperationsWeeklySummaries2016

<https://twiki.cern.ch/twiki/bin/view/AtlasComputing/ADCOperationsWeeklySummaries2016>

Major updates:

-- PetrVokac - 2017-01-10

Responsible: IvanGlushkov

Last reviewed by: *Never reviewed*

This topic: AtlasComputing > ADCOperationsWeeklySummaries2018

Topic revision: r39 - 2018-12-17 - PeterLove



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