


ATLAS Distributed Computing Operations Q1 2012

- Resources utilization
- LFC, FTS, CVMFS
- CERN
- Issues
- Automation

Alessandro Di Girolamo
CERN IT/ES


- By ADC activity

Wall Clock consumption All Jobs in seconds (Pie Chart) (Sum: 782,644,994,031)



Wall Clock consumption All Jobs in seconds (Time Stacked Bar Graph)


109 Days from Week 48 of 2011 to Week 12 of 2012




Maximum: 9,360,006,644 , Minimum: 3,811,385,831 , Average: 7,114,954,491 , Current: 7,619,200,754

- By Tier 0/1/2/3

Wall Clock consumption All Jobs in seconds (Pie Chart) (Sum: 782,644,263,331.061)




Wall Clock consumption All Jobs in seconds (Time Stacked Bar Graph)
109 Days from Week 48 of 2011 to Week 12 of 2012




Maximum: 9,360,006,644 , Minimum: 3,811,385,831 , Average: 7,114,954,491 , Current: 7,619,200,754

WC Consumption for Successful and Failed Jobs (Pie Graph) (Sum: 530,827,415,781)



2011-12-01 -> 2012-03-19

WC Consumption for Successful and Failed Jobs (Pie Graph) (Sum: 198,871,517,057)



Global Wallclock efficiency: 91%


Most frequent failures (in terms of number of failures, not WC):

- 25% LostHeartbeat / batch issues
- 20% storage issues (! but some more are hidden in not easily to understand exit codes in Athena errors & in “others”)
- 20% Athena errors
- 10% Brokering
- 5% DQ2
- 5% TRF
- 15% others

MC12: with high pileup expected in 2012, there will be jobs which need $\text{vmem} > 3.5\text{GB}$: since ATLAS has observed that some batch system kills jobs when exceeding vmem , we ask sites to have limits on vmem of 4GB. SWAP should be set to $2 \times \text{RAM}$ size. Physical memory requirement not changed (2 GB).

Of course ATLAS should define the task with the right requirements, brokering hungry jobs to properly configured queues.





throughput for multi athena jobs



- Test machine:

- 8 core Xeon E5530 @ 2.4 Ghz with hyperthreading
- 24 GB of memory (24GB swap)
- RAW to ESD data with low pileup data from 2011 (on local HD)
- up to 24 jobs can run in parallel with almost 100% CPU utilization, with more jobs machine start swapping. Athena is able to run in event loop within about 1GB of physical mem per job: new high pileup MC samples may need to full address space of 4GB but can fit into less real memory during event loop

Nothing special to report



- Trigger rate up to 400 Hz + 125 Hz delayed:
 - Delayed stream RAW will not be processed at the Tier0 but will be sent to T1 tapes
 - Few random run will be checked
 - Processing will be done at the end of the year
 - 400 Hz data export very similar to data11
 - Details still under approval



Last week Cedric Serfon
DDM Ops presentation

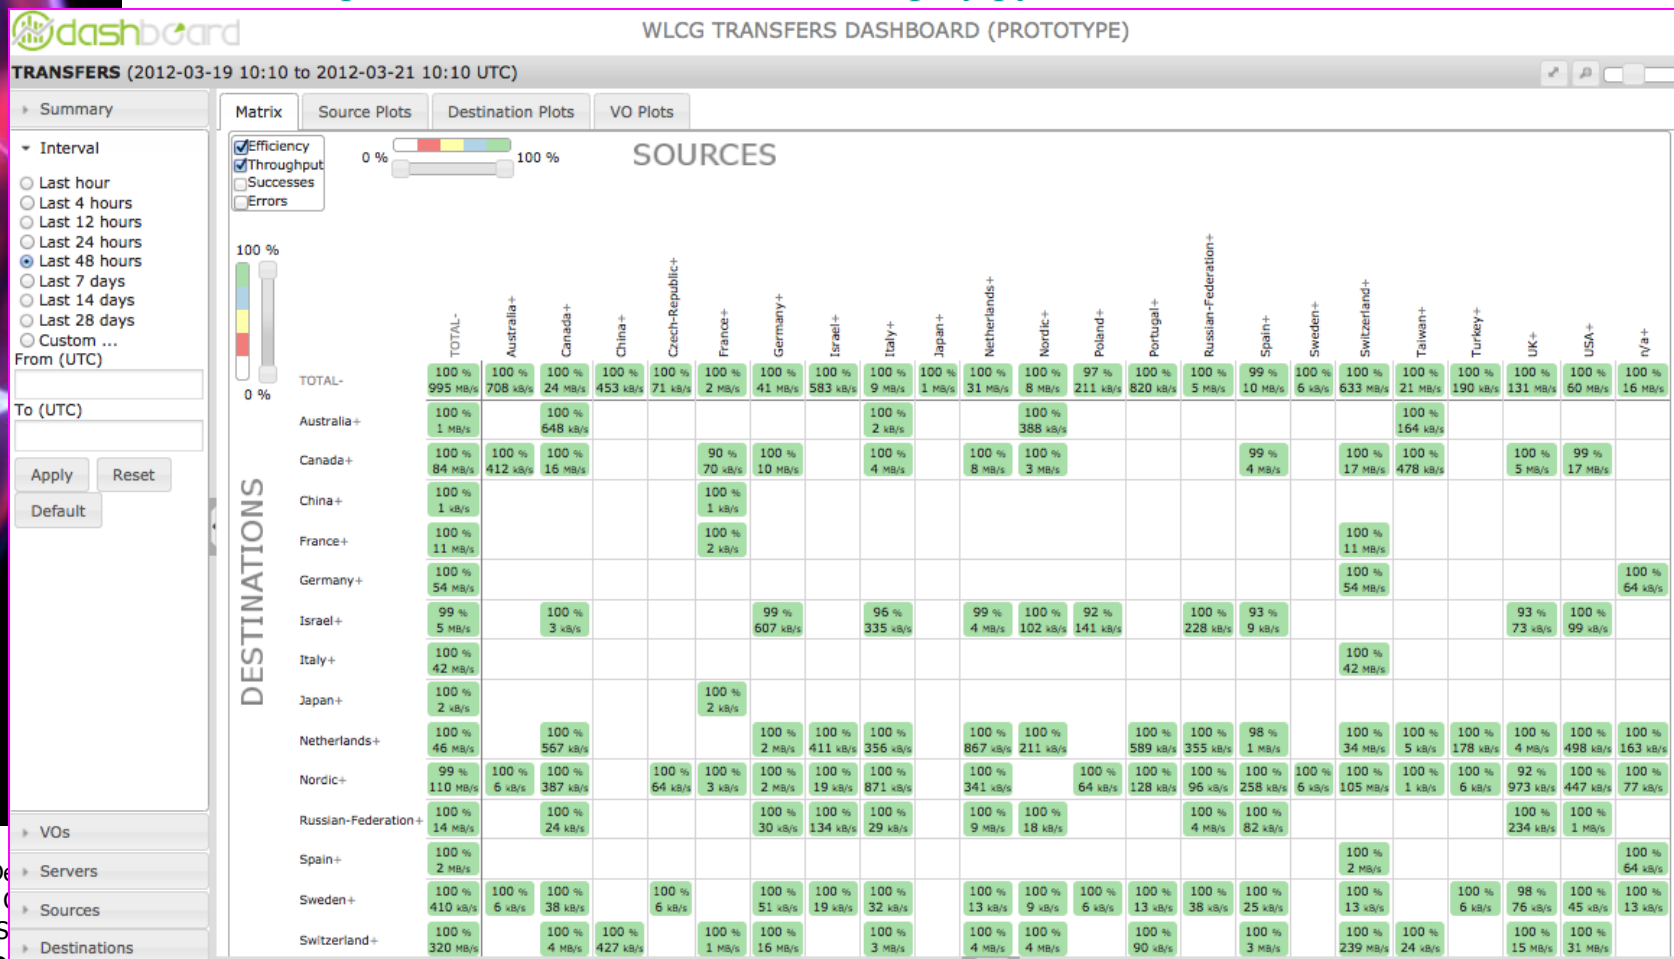
LFC migration - Schedule

- At last S&C week, only NL cloud was consolidated.
- Since then many LFCs consolidated :
 - TW cloud : Nov. 2nd-Nov. 3rd 2011
 - DE cloud : Nov. 14th-Nov. 16th 2011
 - IT cloud : Nov. 29th-Dec. 1st 2011
 - UK cloud : Feb. 14th-Feb. 16th 2011
 - ES cloud : Feb. 28th-Feb. 29th 2012
- 158M replicas hosted on CERN LFC.
- Next week, FR cloud.
- Remaining LFC to migrate :
 - 3 LFCs at T1s : BNL, TRIUMF, NDGF. The 2 latters are mySQL.
 - 8 LFCs on US T2s.

- Smooth procedures
- Many people involved: cloud squads, DBAs (at T1s and CERN), LFC devs. Special Thks to F.Furano and M.Blaszczyk
- LFC load under control:
 - LFC registration moved on the Panda Server for some clouds
 - Plan to move to IP authentication for trusted hosts (PandaServer, DDM SS)
 - Need latest LFC clients

- FTS 2.2.8 functionality tested by CMS & ATLAS in December/January
 - Overwrite: e.g. avoid adding _DQ2timestamp suffix to file after failed transfers
 - Gridftp checksum calculation (→ SRM less storages)
 - Auto-discovery of gridftp endpoints (like e.g. EOS gridftp)
 - Rollout tracked at T1SC meeting
 - ~60% of FTS servers already upgraded, other Tier1s have planned/are planning
 - Issue observed on Tier0 FTS Server
 - “feature” coming with the new globus library (new EMI rel):
 - Gridftp resume transfers in case of particular failures conditions: some gridftp server configured to send out resume and performance markers each second (! default should be 5 seconds !), causing trouble to the node where FTS server runs.
- ✓ Fix already provided in a nonEMI release, installed on CERN pilot and Tier0 FTS server. Will be included in the next EMI rel.

- With FTS2.2.8 we can now use the new aggregated FTS monitor
 - <http://dashb-wlcg-transfers.cern.ch/ui/>
 - Presentation at the last T1SC meeting
<https://indico.cern.ch/conferenceDisplay.py?confId=181610>



- ~ 90 out of ~150 sites now with CVMFS
 - <http://dashb-atlas-ssb.cern.ch/dashboard/request.py/siteview?view=cvmfs#currentView=cvmfs&fullscreen=true&highlight=false&rowsPerPage=-1>


Deployment takes longer but overcomes first technical limitations (disables WNs)

Site Name	CVMFS progress
CERN-PROD	cvmfs
CESGA-EGEE	nfs
CNR-ILC-PISA	nfs
CSCS-LCG2	nfs
csTCDie	nfs
CYFRONET-LCG2	nfs
DESY-HH	nfs
DESY-ZN	cvmfs
EELA-UTFSM	nfs
EFDA-JET	nfs
FMPH-UNIBA	nfs
FZK-LCG2	nfs
GoeGrid	cvmfs
GR-01-AUTH	cvmfs
GRIF-IRFU	cvmfs
GRIF-LAL	cvmfs
GRIF-LPNHE	cvmfs
HEPHY-UIBK	nfs
HG-02-IASA	gpfs
HG-03-AUTH	lustre
HG-04-CTI-CEID	lustre
HG-06-EKT	lustre
HU_ATLAS_Tier2	cvmfs
HU_ATLAS_Tier2_Install	cvmfs
IEPSAS-Kosice	nfs
ifae	cvmfs
IFIC-LCG2	cvmfs
IL-TAU-HEP	cvmfs

Showing 1 to 150 of 150 entries DB query took 0.0148 s

History Legend: afs, lustre, cvmfs, Unknown, nfs, panfs, gpfs

- Tier0 (30KHS06, ~3k job slots) can overspill into grid_atlas (public share, ~3k job slots) since October, and viceversa GRID jobs on Tier0 since December
 - ✓ During HI: up to 5k Tier0 jobs running in parallel when needed.
 - ✓ After stop of data taking GRID MC11 boosted using Tier0 not-anymore-dedicated resources
- Tier0 -> grid_atlas overspill semi-automatic mechanism consolidated
- GRID -> Tier0 still need LSF batch reconf.

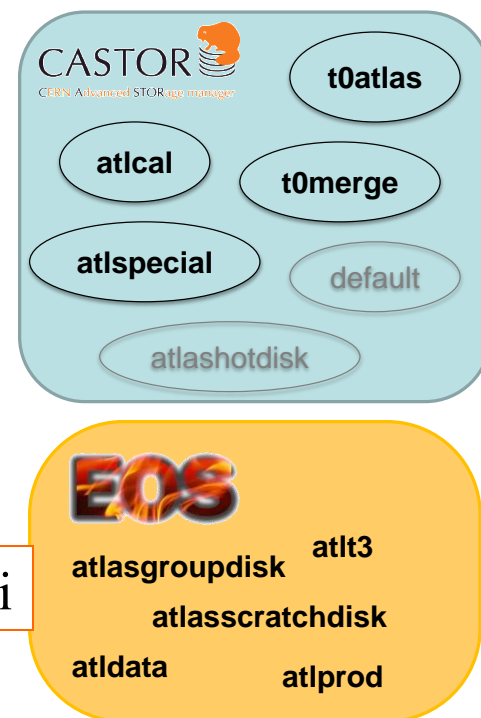



- (part of) ATLAS resources that were previously dedicated have been merged into public share
 - ✓ Usage optimized, much less idle resources!
- 2012 pledges (other ~30 KHS06) will be put in public, Tier0 will overspill:
 - ! Good monitor of share and priority is vital not to occur in problems with slowness of Tier0 reconstruction:
 - Issue addressed during CERN-IT / ATLAS meeting

- CASTOR pool decommissioning ongoing
 - Only pools needed for Tier0 – tape backend- will stay
 - Pool size already increased to be ready for data taking: t0merge will probably be merged with t0atlas
- EOS for all GRID/local activities



Migration: ATLAS



from Luca Mascetti

✓ EOS disk servers now open to outside CERN for xrootd to be used by Panda&DDM

– via HTAR

! Mapping users
EOS+CASTOR+LX BATCH:
now a mixture of gridmap file
and kerberos:

! Other sites may have the same issue

! ATLAS VOMS “nickname” = CERN account: Use it!

! Please agree on an attribute to be used by all VOs!

21 March 2012

Presented at the ATLAS sw week

<https://indico.cern.ch/getFile.py/access?resId=0&materialId=slides&contribId=30&sessionId=3&subContId=3&confId=169695>



Basic Idea, KISS

from Dan van der Ster
and Rod Walker

- The most simple model possible:
 - Configure VMs at CloudSigma which join a condor pool with the master at CERN (one of the pilot factories)
 - I/O over the WAN with CERN SRM (copytools are lcg-cp for inputs and lcg-cr for outputs)
- Created a new PanDA queue *HELIX* which is basically a clone of CERN-PROD and is in brokeroff status
- Assign a few *real* MC production tasks to HELIX



Results: Overall Experience


- The overall experience was uneventful with very few failures. This was a pleasant surprise!
- Finished 6 1000-job MC tasks over ~2 weeks.
- Complete results at:
 - <https://twiki.cern.ch/twiki/bin/view/Main/CloudSigmaResults>

What's missing:

- Automate deployment
- Dynamic provisioning
 - Billed for idle CPU while no tasks
 - ! Need an orchestrator to grow/shrink cluster on demand
- Node monitor
- \$\$\$

- Last 3 months: 438 GGUS team ticket (~ 540 21Sept-22 Dec, 530 21Jul-20Sept)
- ADC SiteExclusion:
 - each time a site is manually excluded from ADC activities (due to long scheduled intervention or due to problems lasting many hours) its status is tracked in savannah adc-site-status


ADC Site Exclusion Savannah, Tickets Quarterly



Total number of ADC
SiteExclusion Savannah each 3
months

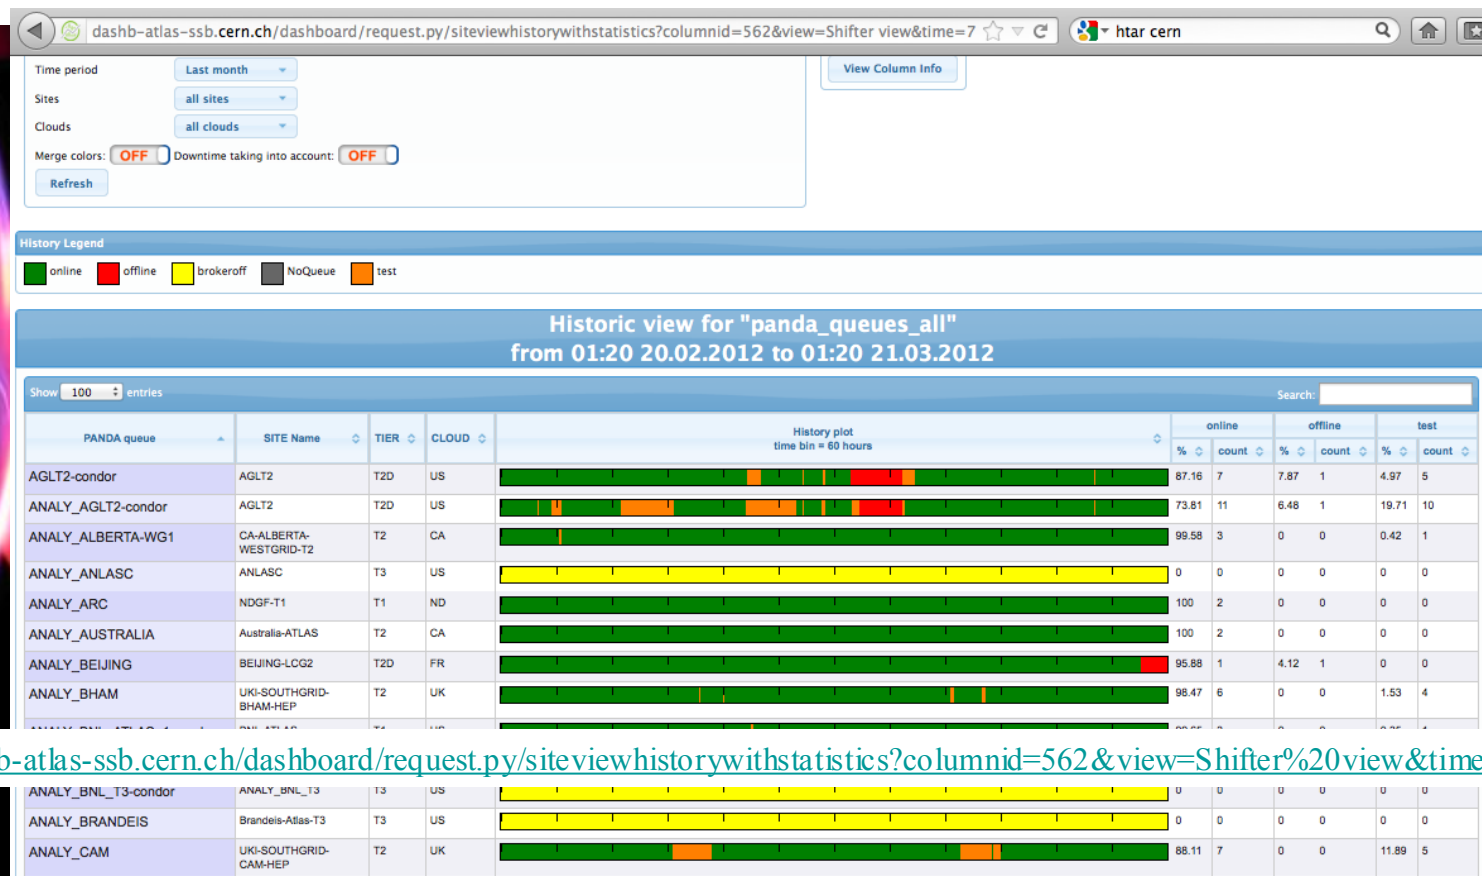
21 March 2012

Number of ADC SiteExclusion Savannah per Squad



- Why: Reduce manpower for basic Operations
 - minimize the repetition of the “same action” by humans
 - use the help of the Experts and the Shifters more efficiently
 - More training will be needed to make them able to monitor also the engines that takes automatic decisions.
- ✓ “Switcher” takes published downtimes into account to stop analysis/production activity
 - ✓ In production since one month
- ✓ HC now checking both analysis and production
- More work ongoing:
 - Set offline pandaqueues when storage is not available
 - Blacklist DDMEndpoints if e.g. SAM test for specific spacetoken availability of the last Nhour below defined threshold


Monthly reviewed and reports requested at ADC weekly: Site usability for analysis activity



<http://dashb-atlas-ssb.cern.ch/dashboard/request.py/siteviewhistorywithstatistics?columnid=562&view=Shifter%20view&time=720>

! Would be useful if the ADC activity (e.g. DataAnalysis, DataTransfer, DataProduction) could be injected in SAM/NAGIOS so that:

- ! sites can interface their internal monitoring system to directly fetch their status
- ! Availability/usability numbers for WLCG management more complete

- 
- ! Data Taking is now starting
 - Trigger rate up to 400 Hz + 125 Hz (delayed stream)
 - ✓ ATLAS is efficiently using all available resources
 - ✓ Ongoing a consolidation of tools to automate/improve Operations (cvmfs, switcher) and
 - ✓ Better publication of site usability for ATLAS
 - ✓ We learned a lot from last year and we are ready for this 2012 year!



COMPUTER PROBLEM SOLVING FLOWCHART

