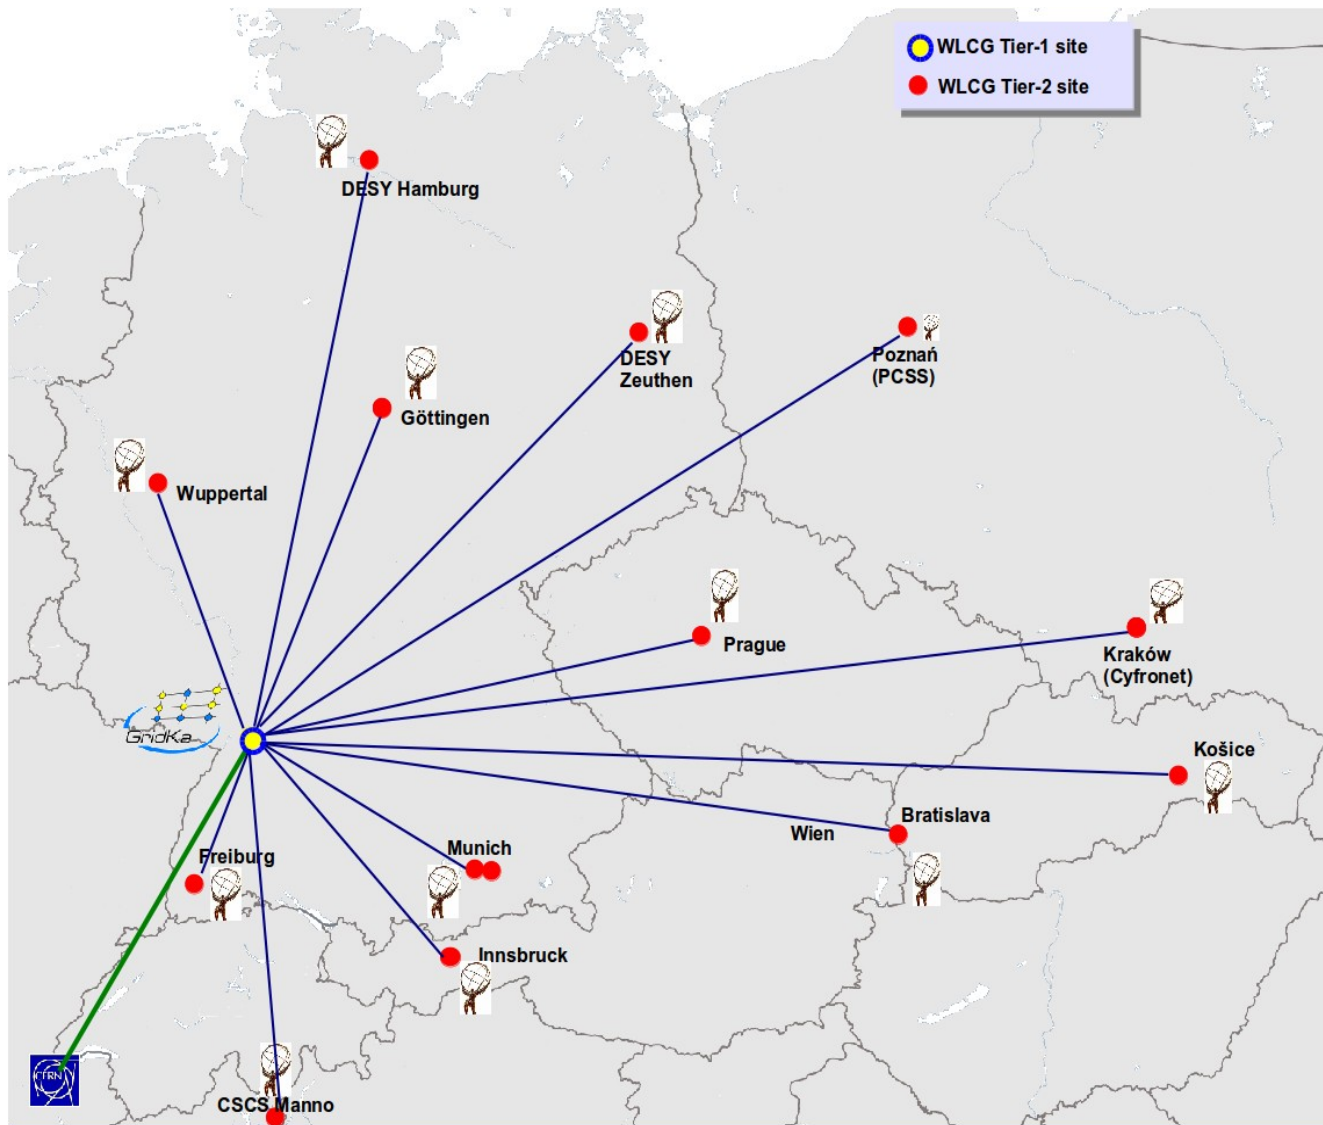


ATLAS DE Cloud

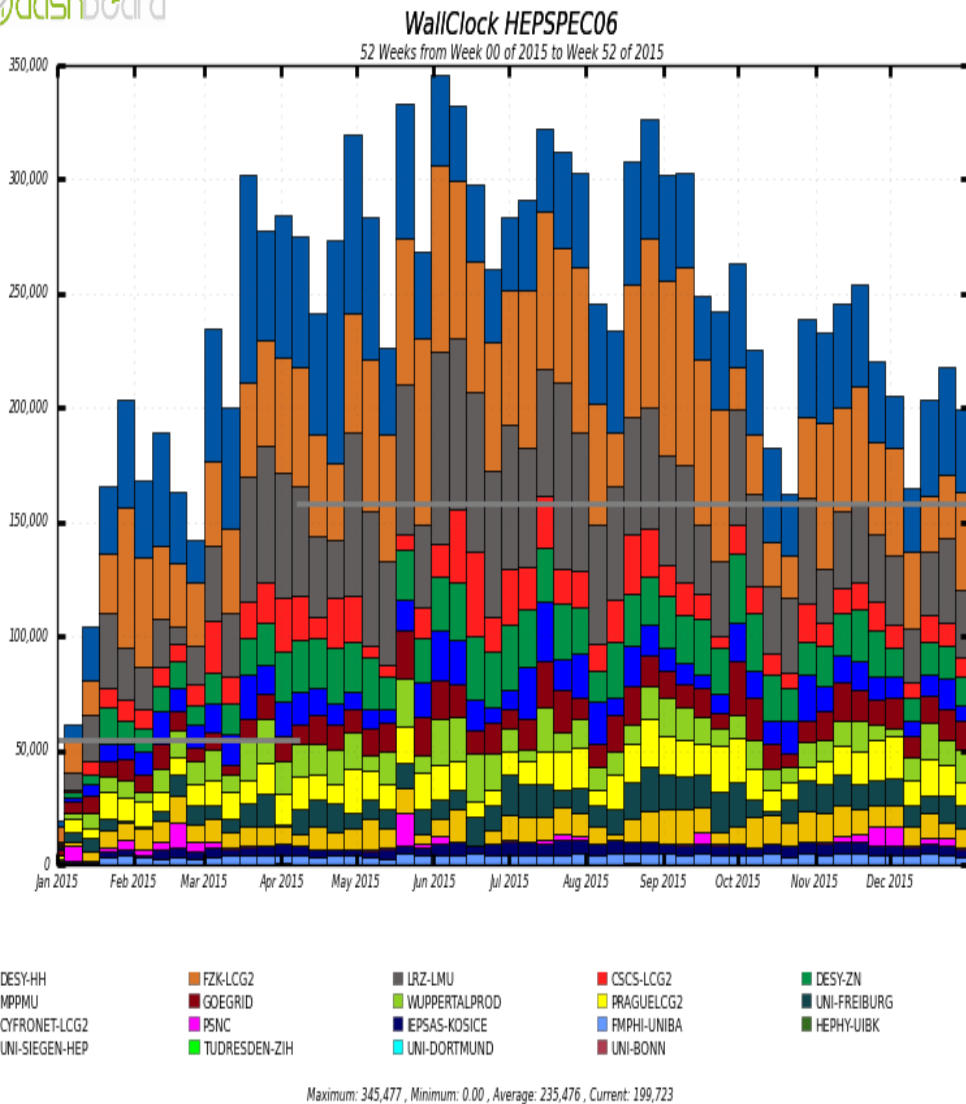
- Overview
- news/changes
- problems/issues

Overview




- Heterogeneous cloud and funding:
 - 15 T1 & T2 sites in 6 countries (AT, CH, CZ, DE, PL, SK)
- Combined resources for ATLAS:
 - 150 kHS06 CPU
 - 15 PB Disk
- Substantial un-pledged resources at T2s (T3-part + HPC)
- Several 'large' T3 sites
- Combined contribution to ATLAS computing ~15%
→ about author share

ATLAS jobs DE cloud Jan - Dec 2015



WallClock HEPSPROC6 (Sum: 2,136,242,334)
DESY-HH - 17.74%



- FZK-LCG2 - 19.83% (423,597,934)
- DESY-HH - 17.74% (378,962,368)
- LRZ-LMU - 17.33% (370,207,111)
- DESY-ZN - 7.58% (161,948,095)
- CSCS-LCG2 - 5.66% (120,915,355)
- WUPPERTALPROD - 4.83% (103,235,367)
- GOEGRID - 4.83% (103,192,041)
- PRAGUELCG2 - 5.10% (108,910,557)
- MPPMU - 5.00% (106,716,032)
- IEPASAS-KOSICE - 1.67% (35,708,400)
- UNI-FREIBURG - 4.19% (89,404,406)
- CYFRONET-LCG2 - 3.77% (80,641,709)
- FMPHI-UNIBA - 1.56% (33,318,272)
- PSNC - 0.81% (17,322,471)
- HEPHY-UIBK - 0.10% (2,162,217)
- UNI-DORTMUND - 0.00% (0.00)
- TUDRESDEN-ZIH - 0.00% (0.00)
- UNI-SIEGEN-HEP - 0.00% (0.00)
- UNI-BONN - 0.00% (0.00)

FZK, Desy-HH, LRZ very close, substantial opportunistic capacity at Desy & LRZ (HPC)

DE cloud organization

- Cloud squad
 - ~10 people from KIT, Desy, Freiburg, Munich, Wuppertal, Cracow/PL
 - most of them also active in T1/2 support or ADC ops/dev/shifts
- weekly vidyo meetings of squad (8-12 people, 20-40 mins)
 - T1 report, DDM & production issues, ATLAS ticket & functional test review
- monthly vidyo meetings squad + sites (15-20 people, ~90 mins)
 - same topics as weekly – more extended discussion + site reports
- yearly F2F meeting
 - next in 3 weeks in Freiburg

DE cloud funding

- Complex and heterogeneous funding structure for computing hardware
 - many countries, states, agencies
 - continuous struggle for funding
 - so far managed reasonably well to provide stable share of pledged resources
 - unclear how well we can keep up in future, especially T2s
- funding for cloud operation manpower included in central DE HEP budget
 - 3 year funding periods
 - competition with analysis, detector, upgrade manpower, ...
 - currently ~3 FTE

Opportunistic resources

- 2015: CPU usage often 200% of pledge sum
- WLCG & Grid sites:
 - most T2s have substantial CPU add-ons (local users, other groups)
 - little at GridKa-T1 – dominated by the 4 LHC expts → hard competition
 - substantial T3s at other institutes/computing centers used when possible
- Other :
 - HPC sites in Munich:
 - SuperMuc, Hydra contribute substantially since ~1.5 years (see R.Walker talk)
 - promising new setup in Freiburg
 - potentially further sites, but we need incentive/acknowledgement of effort

DE cloud news/changes - 1

- dcap → xrootd transition for dcache sites
 - long standing issue – ATLAS aims to reduce # of protocols
 - dcap obvious candidate, dcap dev frozen, xrootd protocol well supported
 - have changed LRZ-LMU & MPPMU in Nov/Dec 2015 to use xrootd for analysis direct IO
 - seems to work well, no problems observed
 - plan to change default access protocol for analysis direct IO for **all DE dCache sites** to xrootd in **next weeks**
 - requires publicly accessible xrootd door (independent of FAX)
 - to be checked/setup
 - AGIS entries for Panda queues need to be updated accordingly
 - will be done by cloud-team

NB: don't confuse with FAX, this here is just xrootd protocol for IO

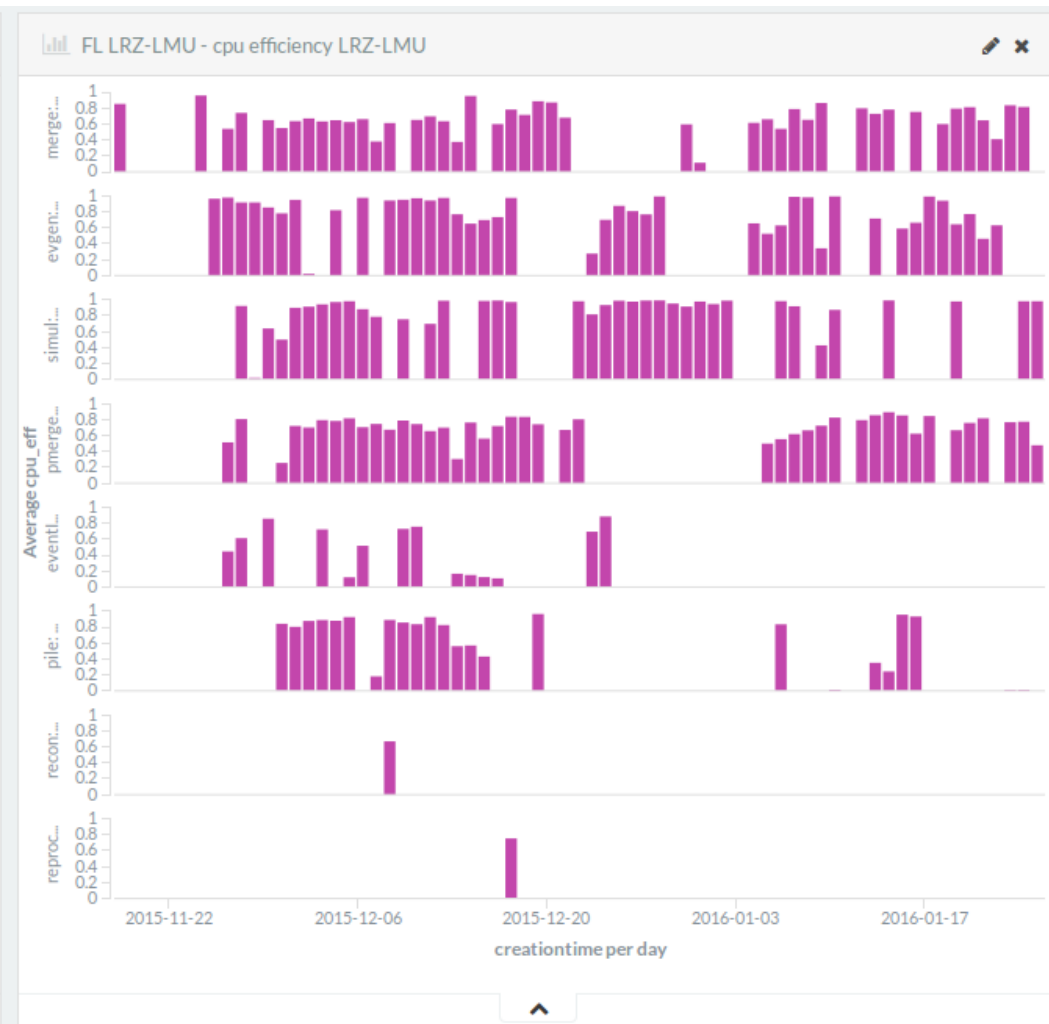
DE cloud news/changes - 2

- simple federation test
 - LRZ-LMU – MPPMU ~500 m apart (funding-wise light-years)
 - ANALY queue at MPPMU configured to use data at LRZ-LMU SE
 - works fine,
 - will do vice versa setup soon
- starting to exploit analytics info for sites/cloud
 - example setup provided by Federica
<http://cl-analytics.mwt2.org:5601/#/dashboard/FL-LRZ-LMU>
 - job efficiency, CPU efficiency, IO volume, ...
 - by Job-type, User, ...
 - easy to get overwhelmed ...

Analytics example: Jobs/queue last 2 months



Analytics example: Cpu-eff per job-type last 2 months



DE cloud problems/issues

- No big issues, most services work reasonably smooth and we saw substantial improvements:
 - HC testing, dark data dumps, lost-file recovery
- A couple of remaining problems/pains:
 - desperately waiting for approval functionality in Rucio (as we had in Datri)
 - many sites with large LOCALGROUPDISK and many users
 - WLCG/SAM tests
 - problems/failures often hard to diagnose, judging relevance, getting feedback, ...
 - most sites focus on HC tests
 - sometimes bad surprises when confronted with bad WLCG availability figs
 - some unclear responsibilities
 - e.g. who takes care of full SCRATCHDISK?
 - FAX services
 - stability & support

LOCALGROUPDISK monitoring service – open for other clouds

[home](#)



BERGISCHE
UNIVERSITÄT
WUPPERTAL

Welcome to the ATLAS German (and Dutch) Cloud LOCALGROUPDISK monitor!

Select site to view:

DE cloud:

- [CSCS-LCG2](#)
- [CYFRONET-LCG2](#)
- [DESY-HH](#)
- [DESY-ZN](#)
- [FMPHI-UNIBA](#)
- [FZK-LCG2](#)
- [GOEGRID](#)
- [HEPHY-UIBK](#)
- [IEPSAS-KOSICE](#)
- [LRZ-LMU](#)
- [MAINZGRID](#)
- [MPPMU](#)
- [PRAGUELCG2](#)
- [TUDRESDEN-ZIH](#)
- [UNI-BONN](#)
- [UNI-FREIBURG](#)
- [UNI-SIEGEN-HEP](#)
- [WUPPERTALPROD](#)
- [FZK-LCG2_LOCALGROUPTAPE](#)
- [PRAGUELCG2_LOCALGROUPTAPE](#)

**Overview of DE sites
with LOCALGROUPDISKS
(and LOCALGROUPTAPES)**

Choose your site




NL cloud:

- [NIKHEF-ELPROD](#)
- [TECHNION-HEP](#)
- [IL-TAU-HEP](#)
- [WEIZMANN-LCG2](#)

home / usage per user

Usage per user

user	used space	# datasets	avg size
Arwa Bannoura	25.96 TB (28548082830367 B)	322	82.57 GB (88658642330.3 B)
Phillipp Tepel	18.42 TB (20247765503316 B)	268	70.36 GB (75551363818.3 B)
tcorneli	11.17 TB (12283237909954 B)	1022	11.19 GB (12018823786.6 B)
Kathrin Becker	10.09 TB (11096689568576 B)	479	21.58 GB (23166366531.5 B)
Daniela Boerner	9.45 TB (10392053709476 B)	1207	8.02 GB (8609820803.21 B)
Manuel Neumann	5.97 TB (6560935852954 B)	56	109.11 GB (117159568803.0 B)
Jan Kuechler	4.16 TB (4575362835707 B)	92	46.32 GB (49732204735.9 B)
Dominic Hirschbuehl	3.83 TB (4206807064468 B)	7	559.70 GB (600972437781.0 B)
Sebastian Fleischmann	1.68 TB (1848556883082 B)	35	49.19 GB (52815910945.2 B)
Julia Fischer	1.56 TB (1720411404022 B)	2179	752.97 MB (789541718.229 B)
Marcello Barisoni	1.34 TB (1468450167443 B)	14	97.69 GB (104889297674.0 B)
Wolfgang Walkowiak	769.25 GB (825973447564 B)	116	6.63 GB (7120460754.86 B)
Andrey Khoroshilov	756.50 GB (812283565349 B)	208	3.64 GB (3905209448.79 B)
Dominik Duda	526.91 GB (565764230920 B)	18	29.27 GB (31431346162.2 B)
Sebastian Fleischmann	152.89 GB (164166433723 B)	7	21.84 GB (23452347674.7 B)
Cedric Serfon	134.37 GB (144282349860 B)	3	44.79 GB (48094116620.0 B)
Sopio Patarai	20.65 GB (22174659289 B)	1	20.65 GB (22174659289.0 B)
Philipp Anger	2.41 GB (2582453849 B)	9	273.65 MB (286939316.556 B)
Sum of space usage	95.4 TB		
Dataset count	604		
Avg dataset size	16.26 GB		



Choose a user

Federating DE cloud ?

- ATLAS/WLCG proposal to move to regional/national centers
- In DE cloud only few small sites → candidates to get Cache-only
 - Most in ~1 PB storage league
- Not obvious (to me) how sites could be federated
 - ~15 T2s, ~10 FAs, dCache & DPM systems
- Aren't we just moving central operations effort to regional federations support?
 - Possibly multiplying it that way ...
- We fully agree that current operations effort cannot be maintained long-term
 - but need full picture: sites, cloud squads & central operations
- Let's see concrete examples ...