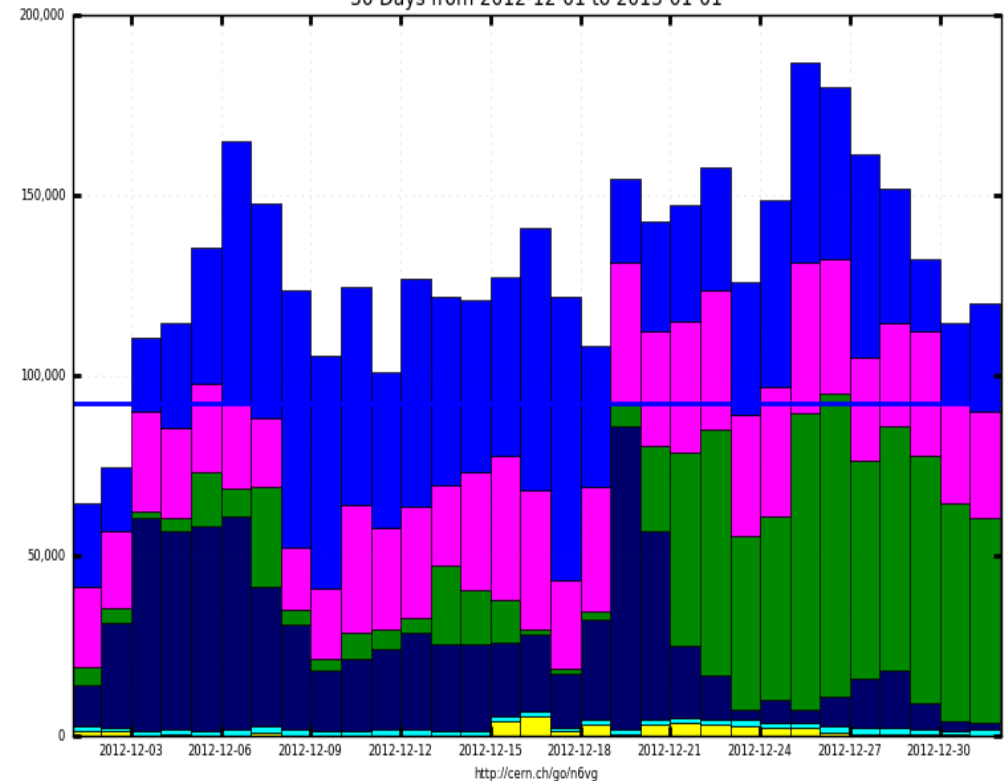


Cloud Status & News Jan 2013

- Happy new year to everybody!
- Traditional smooth and stable running in Dec and Xmas break
- Performance & Stability
 - PSNC & GoeGrid excluded from ANALY_, need to review status
 - other sites mostly fine
- ATLAS T1/2/3 Jamboree 10/11 Dec 12
 - many changes planned / under discussion
 - → slides at the end



WallClock HEPSPROC6 Hours
30 Days from 2012-12-01 to 2013-01-01



■ MC Simulation ■ User Analysis ■ Group Production ■ MC Reconstruction ■ Others
■ Data Processing


Maximum: 186,771, Minimum: 64,707, Average: 130,908, Current: 119,857

Production jobs DE cloud Dec 2012



Wall Clock consumption All Jobs in seconds


30 Days from 2012-12-01 to 2012-12-31



Maximum: 1,256,449,271 , Minimum: 378,474,308 , Average: 897,271,918 , Current: 778,616,410



Wall Clock consumption All Jobs in seconds (Sum: 27,815,429,482)




close to 2012 average (~30B/month)

Analysis jobs DE cloud Nov 2012



Wall Clock consumption All Jobs in seconds
30 Days from 2012-12-01 to 2012-12-31




- PRAGUELCG2
- DESY-HH
- FZK-LCG2
- CYFRONET-LCG2
- CSCS-LCG2
- DESY-ZN
- UNI-FREIBURG
- MPPMU
- LRZ-LMU
- WUPPERTALPROD
- TUDRESDEN-ZIH
- GOEGRID
- PSNC
- HEPHY-UIBK

Maximum: 367,966,730 , Minimum: 148,606,431 , Average: 262,052,151 , Current: 260,361,523



Wall Clock consumption All Jobs in seconds (Sum: 8,123,616,701)
DESY-HH - 16.59%



- PRAGUELCG2 - 16.75% (1,361,047,399)
- DESY-HH - 16.59% (1,347,423,237)
- FZK-LCG2 - 13.13% (1,066,458,011)
- CYFRONET-LCG2 - 10.34% (839,935,368)
- CSCS-LCG2 - 9.70% (787,747,579)
- DESY-ZN - 9.12% (740,576,800)
- UNI-FREIBURG - 7.06% (573,518,338)
- MPPMU - 5.87% (476,831,662)
- LRZ-LMU - 5.52% (448,777,874)
- WUPPERTALPROD - 4.16% (337,583,604)
- TUDRESDEN-ZIH - 0.91% (73,823,793)
- GOEGRID - 0.78% (63,200,469)
- PSNC - 0.08% (6,587,641)
- HEPHY-UIBK - 0.00% (104,926)

close to 2012 average 7,500,000,000

DE site availability Dec 12

Historic view for "panda_queues_all"
from 00:00 01.12.2012 to 00:00 31.12.2012

Show entries

PANDA queue	SITE Name	TIER	CLOUD	History plot time bin = 60 hours	offline		brokeroff		online	
					%	count	%	count	%	count
ANALY_CSCS	CSCS-LCG2	T2D	DE		0	0	0	0	98.17	15
ANALY_CYF	CYFRONET-LCG2	T2	DE		0	0	0	0	97.43	18
ANALY_DESY-HH	DESY-HH	T2D	DE		0	0	0	0	91.52	12
ANALY_DESY-HH_XD	DESY-HH	T2D	DE		0	0	0	0	0	0
ANALY_DESY-ZN	DESY-ZN	T2D	DE		0	0	0	0	98.94	7
ANALY_DRESDEN	TUDresden-ZIH	T3	DE		0	0	0	0	63.04	36
ANALY_FREIBURG	UNI-FREIBURG	T2D	DE		0	0	0	0	99.49	5
ANALY_FZK	FZK-LCG2	T1	DE		1.04	2	0	0	98.06	9
ANALY_FZU	pragueicg2	T2	DE		0	0	0	0	97.36	11
ANALY_GOEGRID	GoeGrid	T2D	DE		0	0	0	0	5.65	6
ANALY_HEPHY-UIBK	HEPHY-UIBK	T2	DE		0	0	99.63	4	0	0
ANALY_LRZ	LRZ-LMU	T2D	DE		0.93	1	0.56	1	91.94	13
ANALY_LRZ11	LRZ-LMU	T2D	DE		0	0	0	0	0	0
ANALY_MPPMU	MPPMU	T2D	DE		0	0	0	0	94.7	9
ANALY_PSNC	PSNC	T2	DE		0	0	0	0	4.14	7
ANALY_wuppertalprod	wuppertalprod	T2D	DE		0	0	0	0	95.43	7

- Sites PSNC & GOEGRID manually set to test mode by cloud squad due to reoccurring problems / instabilities
 - will be put back when HC tests run stable
- Other DE analysis T2 sites well above 90% online




Tier2 report GangaRobot Dec 2012 (table)

input file project: data 01.12.2012 - 31.12.2012

SITE	failed	finished	efficiency	finished jobs					
				runtime(end-start)			cpu time		
				min	mean	max	min	mean	max
ANALY_CSCS	209	5117	0.961	123	160	4390	32	68	114
ANALY_CYF	184	5746	0.969	189	309	6059	51	68	134
ANALY_DESY-HH	113	5495	0.980	124	504	12728	38	91	483
ANALY_DESY-ZN	42	6383	0.993	79	165	1101	34	49	102
ANALY_FREIBURG	43	6770	0.994	123	210	12950	41	63	427
ANALY_FZK	44	5522	0.992	123	303	14661	35	74	106
ANALY_FZU	100	7057	0.986	126	286	5693	43	70	119
ANALY_GOEGRID	505	3026	!-> 0.857	128	240	12180	36	53	72
ANALY_HEPHY-UIBK	58	5332	0.989	191	341	11805	58	71	80
ANALY_LRZ	121	5992	0.980	124	301	6111	48	79	121
ANALY_MPPMU	282	6124	0.956	102	189	4952	43	58	79
ANALY_wuppertalprod	82	5396	0.985	124	346	14963	41	67	115


input file project: mc 01.12.2012 - 31.12.2012

SITE	failed	finished	efficiency	finished jobs					
				runtime(end-start)			cpu time		
				min	mean	max	min	mean	max
ANALY_CSCS	204	5666	0.965	184	410	1959	102	245	465
ANALY_CYF	202	4280	0.955	313	1080	6118	124	208	454
ANALY_DESY-HH	151	4989	0.971	185	812	17887	107	330	748
ANALY_DESY-ZN	35	7267	0.995	184	348	5310	101	170	374
ANALY_FREIBURG	49	7315	0.993	185	410	9160	125	216	2316
ANALY_FZK	65	5714	0.989	185	541	14212	109	257	500
ANALY_FZU	98	7218	0.987	189	544	6765	122	253	510
ANALY_GOEGRID	496	3428	!-> 0.874	190	412	6152	107	179	337
ANALY_HEPHY-UIBK	52	5391	0.990	314	617	12965	155	227	285
ANALY_LRZ	151	6272	0.976	247	517	6701	123	247	430
ANALY_MPPMU	323	6626	0.954	185	383	5313	128	197	312
ANALY_wuppertalprod	78	5576	0.986	185	601	10076	122	234	973



Tier2 report GangaRobotPlot Dec2012 (1)


ANALY_CSCS



CSCS

data
209/5326


1



5326

mc
204/5870

1



5870

no child procs (LSF batch problem), other mixed problems

Federated data access

- Federated data access via xrootd (FAX)
 - lot of activity in Dec, IT cloud fully joined
 - full-dress-rehearsal tests planned for mid/end Jan
 - preparations ongoing
 - still little progress in DE cloud
 - LRZ, RZG, Wup, Prague currently active
 - Desy Gridlab
 - in progress at GridKa
 - FAX service will presumably become ATLAS requirement
 - rather join sooner than later
 - Wiki page for setup of xrootd as proxy service
 - <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/FAXDECloud>
- Federation via http/Webdav as well but independent
 - see <https://twiki.cern.ch/twiki/bin/view/Atlas/WebDavDataAccess>

What is FAX?


Description (from the FAX Twiki):

The Federated ATLAS Xrootd (FAX) system is a storage federation aims at bringing Tier1, Tier2 and Tier3 storage together as if it is a giant single storage system, so that users do not have to think of there is the data and how to access the data. A client software like **ROOT** or **xrdcp** will interact with FAX behind the sight and will reach the data wherever it is in the federation.

Goals (from [Rob Gardner's talk at Lyon mtg. 2012](#)):

- Common ATLAS namespace across all storage sites, accessible from anywhere;
- Easy to use, homogeneous access to data
- Use as failover for existing systems
- Gain access to more CPUs using WAN direct read access
- Use as caching mechanism at sites to reduce local data management tasks

Regional redirectors



Start locally - expand search as needed

T1/2/3 Jamboree Issues

- <https://indico.cern.ch/conferenceDisplay.py?confId=196649>
- DDM plans:
 - Migration DQ2 → Rucio DDM system to start 2013
 - radical namespace change - 2 level subdir created from filename-hash
 - requires renaming campaign - start early 2013 w/ test sites via WebDav
 - with Rucio SpaceToken not really useful any more - single copy per file on site
 - LocalGroupDisk: not clear how to use in future
 - further changes
 - end of LFC by 2014
 - drop SRM for disk-only sites, only needed for tape recall
- Middleware requirements:
 - enforce CVMFS by end of Apr 13 (WLCG reqt??)
 - Squid considered critical service
 - glexec not much progress and not pushed by ATLAS

T1/2/3 Jamboree Issues - 2

- Os & Env
 - SL6: ok for binary prod but not quite ready for analysis (compilation)
 - /etc/machinefeatures foreseen to communicate node specific features
 - to be used by multicore jobs
- perfSonar deployment on all sites requested
 - requires 2 phys nodes in same sub-net as storage pools
 - in DE cloud so far only GridKa, Desy-HH, LRZ (AFAIK)
- Storage access requirements for Sites:
 - xrootd access (straightforward for dCache and recent DPM versions)
 - Federated access: FAX redirector & http/webdav