Show 100 entries   Search: ANALY

PANDA queue	SITE Name	TIER	CLOUD	History plot time bin = 62 hours	offline %	offline count	brokeroff %	brokeroff count	online %	online count	NoQueue %	NoQueue count	test %	test count
ANALY_RHUL	UKI-LT2-RHUL	T2	UK		0	0	0	0	99.42	4	0	0	0.34	1
ANALY_BHAM	UKI-SOUTHGRID-BHAM-HEP	T2	UK		1.48	1	0	0	97.58	6	0	0	0.69	3
ANALY_CAM	UKI-SOUTHGRID-CAM-HEP	T2	UK		0.81	1	0.54	1	97.21		0	0	1.2	6
ANALY_OX	UKI-SOUTHGRID-OX-HEP	T2D	UK		0	0	0	0	97	6	0	0	2.76	3
ANALY_SHEF	UKI-NORTHGRID-SHEF-HEP	T2	UK		0	0	0	0	96.89	8	0	0	2.87	5
ANALY_MANC	UKI-NORTHGRID-MAN-HEP	T2D	UK		0	0	0	0	96.12	5	0	0	3.63	2
ANALY_GLASGOW	UKI-SCOTGRID-GLASGOW	T2D	UK		0	0	0	0	95.17	18	0	0	4.58	15
ANALY_RAL_XROOTD	RAL-LCG2	T1	UK		0	0	0	0	93.91	10	0	0	5.85	7
ANALY_LIV	UKI-NORTHGRID-LIV-HEP	T2	UK		0	0	0	0	93.77	11	0	0	5.96	8
ANALY_RAL	RAL-LCG2	T1	UK		0	0	0	0	93.55	12	0	0	6.21	9
ANALY_RALPP	UKI-SOUTHGRID-RALPP	T2	UK		3.25	1	0	0	87.96	19	0	0	8.55	17
ANALY_QMUL	UKI-LT2-QMUL	T2D	UK		0	0	0	0	87.31	11	0	0	12.44	8
ANALY_ECDF	UKI-SCOTGRID-ECDF	T2D	UK		0	0	0	0	82.96	17	0	0	16.79	16
ANALY_LANCS	UKI-NORTHGRID-LANCS-HEP	T2D	UK		0	0	0	0	60.95	11	0	0	38.8	12
ANALY_UCL	UKI-LT2-UCL-HEP	T2	UK		15.33	2	0.54	1	53.54	12	0	0	30.34	9
ANALY_GLASGOW_XROOTD	UKI-SCOTGRID-GLASGOW	T2D	UK		0	0	99.75	3	0	0	0	0	0	0

Showing 1 to 16 of 16 entries (filtered from 74 total entries)   First Previous 1 Next Last

UKI-SOUTHGRID-RALPP (87.99%): downtime due to power outage during Christmas lasted until the 2/1/2013 [1].

We've had an intermittent issue with the CreamCE (apparently) where the all suddenly slow down and starts timing out most connections which usually seems to last for a couple of days before disappearing for a week or so. After several rounds of this and looking for problems with the various CreamCE service, the batch system, job submission by various VOs and a lot of poking around in (and cleaning of) the Cream databases we finally tracked the issue down to the Argus server. We need to do more investigation of the problem there but now should have a quick fix of restarting the Argus services to get things running fine again.

[1] https://goc.egi.eu/portal/index.php?Page_Type=View_Object&object_id=119600&grid_id=0

UKI-LT2-QMUL (87.31%+4.23%=91.54%): 4.23% unavailability is corresponds to two peaks of jobs failing with "Failed to Get LFC replicas" [1] It might be due to random failures from RAL top BDII and the Imperial failover which were overloaded. Or it might be due to the LFC.

QMUL had issues with our squids upgrading and failing (fixed in latest release), and StoRM not surviving a mysql update - that was fixed quickly, but there was some problem at the ATLAS end that meant that Elena had to put us online manually.

In addition to that, we've had some scheduled network at risk on a pair of Mondays that resulted in a short downtime within the at risk period (we hadn't expected this, but knew it was a risk).

[1] http://tinyurl.com/cyngveq

UKI-SCOTGRID-ECDF (82.96%): We have suffered from multiple exclusion instances over the last month due to test jobs failing with CVMFS timeout errors.

We have been pursuing different avenues of investigation for these timeouts and have been in discussion with Jakob Blomer on the issue (see cvmfs-talk mailing list for details). The large majority of job failures occurred on our hyperthreaded nodes - where double the amount of slots for ATLAS jobs are available - and as a precaution we have reduced the number of slots on these nodes

for the time being.

Note that the availability of ATLAS jobs to ECDF resources was not affected during these exclusions (i.e other ATLAS jobs were still running successfully through these exclusion periods).

UKI-NORTHGRID-LANCS-HEP (60.96%): During January Lancaster suffered from problems with the srmv2.2 daemon crashing on our DPM (tracked to a bug in dpm which is being worked on, we have put a watch dog in place to auto-restart the daemon). A downed disk node cost us a couple of days whilst we rebuilt it - jobs died whilst they tried to access data on this server. Atlas were slow to respond to our input- https://ggus.eu/ws/ticket_info.php?ticket=90778. During the same week we had problems with the fetch-crl cron job on our nodes, whacked by a mistake made whilst updating our pool nodes and deploying webdav, which cost us another few days.

UKI-LT2-UCL (53.54%+15.33%=68.87%): downtime due to DPM head node upgrade from glite to EMI-2 [1]. Between the 14-20/01/2013 a number of jobs were cancelled by the system. It is not possible to know the reason because neither panda nor the dashboard reports it [2]

[1] https://goc.egi.eu/portal/index.php?Page_Type=View_Object&object_id=120197&grid_id=0
[2] http://tinyurl.com/d4fwz45

TEST Problems

Cancelled jobs: job gets cancelled without reporting a reason making it difficult to know why they were cancelled. See UCL.