

Revised ATLAS pilot monitoring

Peter Love
16 January 2018

Introduction

There are two main FAQs asked by site admins in regards to ATLAS jobs on their Grid resource

1. why are we not submitting enough jobs?
2. why are we submitting too many jobs?

<http://apfmon.lancs.ac.uk>

[HOME](#) [Cloud](#) [Reports](#) [Help](#)

ATLAS pilot factory monitor

- apfmon was intended to answer these questions
- it mostly succeeded but the questions keep coming so it can be improved

Factory	activity	total	recently	when	version
ATLAS-Atlas1-066		-	3	1 min ago	1.0.0
atlas009		-	0	1 day ago	2.4.6
pipanda002-simpl		-	115	seconds ago	2.3.1-scalefix-2
pipanda009-simpl		-	6	2 mins ago	2.3.1-scalefix-2
pipanda013-simpl		-	4	seconds ago	2.3.1-scalefix-2
pipanda020-ipv6dev		-	0	3 hrs ago	2.4.9
pipanda067-peterdev		226	3	2 mins ago	2.4.6
pipanda079-daviddev		59	19	seconds ago	2.4.9
pipanda125-tickle		3904	34	seconds ago	2.4.9
pipanda126-ptest		1397	19	seconds ago	2.4.9
pipanda135-vac		-	25	seconds ago	2.3.1-scalefix-2
pipanda146		2672	20	seconds ago	2.4.9
pipanda147		2409	21	seconds ago	2.4.9
pipanda149		2643	43	seconds ago	2.4.9
pipanda150		3181	36	seconds ago	2.4.9
pipanda151		2984	35	seconds ago	2.4.9
pipanda152		3228	22	seconds ago	2.4.9
pipanda153		3374	22	seconds ago	2.4.9
pipanda154		6436	21	seconds ago	2.4.9
pipanda155		6897	24	seconds ago	2.4.9
pipanda156		4017	72	seconds ago	2.4.9
pipanda157		4710	3	seconds ago	2.4.9
pipanda158		3066	49	seconds ago	2.4.9
HL-Lumi13-ibover		-	40	seconds ago	2.4.13

Some recent developments

- requesting feedback -

- Previous study (CHEP2016) showed a level of empty pilots which has negative impact on site batch systems.
- Identifying these required joining a couple of database tables and was time consuming to analyze.
- Recently made changes to wrapper and pilot code in order to get truthful information regarding whether a wrapper job runs a payload or not.
- Some new views to be added to apfmon to highlight the problem to ATLAS shifters and site admins.
- Other issues may be better displayed too.

Preview of new views


- Implemented using Plotly Dash libraries
- Data workflow a little different:
 1. Pilot factory logs stored in AWS S3, now with payload field
 2. Analyzed using AWS Athena which is a SQL query service for direct analysis of S3 data
 3. Query results imported by Dash webapp
- Any new analysis ideas welcome!

DEMO

Queue comparison

queues displayed

Queue comparison (30 days)




match apf queue	total jobs	empty jobs	empty percentage
<input type="text" value="Search"/>	<input type="text" value="Search"/>	<input type="text" value="Search"/>	<input type="text" value=">=0"/>
ANALY_ROMANIA07-7675	121651	111887	91.79
NIKHEF-ELPROD-HighMem-13712	103558	100355	96.91
NIKHEF-ELPROD-HighMem-13710	87505	85014	97.15
IN2P3-LAPP-prod_MCO	47064	44026	93.54
UKI-SCOTGRID-DURHAM_MCORE-10955	28613	26200	91.57
RO-07-NIPNE-creamce07-atlasmcore-9835	22540	22283	98.77
CA-SFU-T2_MCORE-14640	19453	17908	92.06
CYFRONET-LCG2-all-prod-CEs-14784	18520	17698	95.56
DESY-HH_TEST-14571	15313	15296	99.89
CYFRONET-LCG2-all-prod-CEs_MCORE-14833	15131	14147	93.5
UKI-LT2-Brunel_ipv6_TEST-12587	13566	13534	99.76
UKI-SCOTGRID-GLASGOW_MCORE-12706	12742	12734	99.94
UKI-SCOTGRID-GLASGOW_MCORE-10895	12742	12187	99.95

Queue comparison

queues displayed

Queue comparison (30 days)



match_apf_queue	▼ total_jobs	empty_jobs	empty_percentage
<input type="text" value="Search"/>	<input type="text" value="Search"/>	<input type="text" value="Search"/>	<20
CERN-PROD_HI-14817	179056	7048	3.94
ANALY_CERN-14799	139490	20596	14.77
ANALY_CERN-14805	96444	17788	18.44
ANALY_INFN-NAPOLI-B	65478	10425	15.92
DESY-ZN-lcg-cream-atlas-pbs-12268	56981	7587	13.31
CERN-PROD_T0_4MCOE-12275	52969	5402	10.2
CERN-PROD_T0_4MCOE-12274	52051	5902	11.27
DESY-ZN-lcg-cream-atlas-pbs-12266	49074	8030	16.1
ANALY_CERN-14793	48760	9382	19.24
pic-htcondor-score-14001	44146	5764	13.06
ANALY_RRC-KI-T1-4736	34866	6231	17.87
ANALY_RRC-KI-T1-14047	34805	6931	19.91
ANALY_FMPHI-UNIBA-5135	31148	4082	13.04

CERN-PROD-14614

