

ATLAS BigPanDA Monitoring and Its Evolution

Siarhei Padolski¹, Torre Wenaus¹ and [Tatiana Korchuganova](#)²

For the ATLAS Collaboration

¹ Brookhaven National Laboratory

² Tomsk Polytechnic University


The 7th International Conference
"Distributed Computing and Grid-
technologies in Science and Education"

ATLAS

ATLAS is one of two general-purpose detectors at the Large Hadron Collider (LHC).

It investigates a wide range of physics, from the search for the Higgs boson to extra dimensions and particles that could make up dark matter.


1PB/sec from all sub-detectors reduces to 1GB/sec raw data which is saved to Data Center.



ATLAS Distributed Computing

- **150** computing centers
- **200** PB storage
- **250k** job slots
- **2M** jobs/day
- **1500** users

PanDA = Production and Distributed Analysis



Monitoring requirements

Main goal - the rapid identification of failures and monitoring of progress of distributed physics analysis and production.

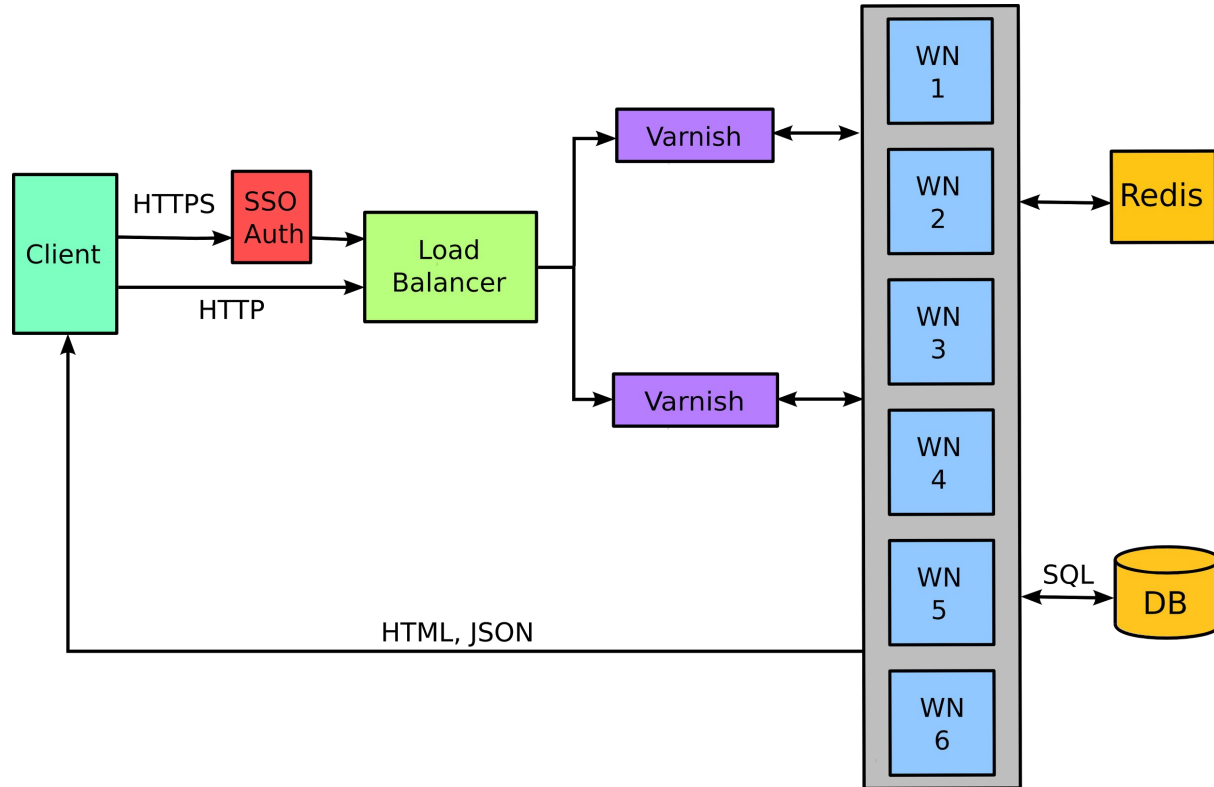
- Access to the real-time and short-term-history (up to 6 months) aggregated information
- 4 distinct user behavior patterns: distributed computing systems operators, shifters, physicist end-users and computing managers
- Drilling-down from high level summaries to detailed diagnostics data
- Data processing time is limited by the web browsing experience

BigPanDA monitoring

- Using Model View Controller provided by Django framework
- Various DB backends (Oracle or MySQL)
- Data processing algorithms splitted between DB and frontend
- Advanced plots generation on client side using d3.js library
- Dynamic data delivery (Ajax)



BigPanDA monitoring



Available views

Common views:

- Jobs;
- Tasks;
- Errors;
- Users;
- Sites;
- Incidents;

Features:

- Attribute summary
- Sorting by essential attributes

Job attribute summary		Sort by count , alpha
atlasrelease (1)	Atlas-20.7.6 (398)	
attemptnr (2)	1 (386) 2 (12)	
cloud (1)	WORLD (398)	
computingsite (9)	BNL_PROD (7) CA-SCINET-T2 (25) CPPM (25) FZK-LCG2 (274) INFN-T1 (1) MWT2_SL6 (27) pic (10) UKI-LT2-RHUL_SL6 (25) UNI-FREIBURG (4)	
eventservice (1)	ordinary (398)	
homepackage (1)	AtlasDerivation/20.7.6.4 (398)	
inputfileproject (1)	data15_13TeV:data15_13TeV (398)	
inputfiletype (1)	AOD (398)	
jeditaskid (1)	8737370 (398)	
jobsetid (1)	7945 (398)	
jobstatus (2)	failed (12) finished (386)	
minramcount (1)	1-2GB (388)	
nucleus (1)	FZK-LCG2 (398)	
outputfiletype (1)	TAG (398)	
priorityrange (2)	500-599 (388) 900-999 (10)	
processingtype (1)	merge (398)	
prodsourcelabel (1)	managed (398)	
producername (1)	atlas-dpd-production (398)	
reqid (1)	7945 (398)	
specialhandling (7)	ddm:rucio,hc:CA (25) ddm:rucio,hc:DE (278) ddm:rucio,hc:ES (10) ddm:rucio,hc:FR (25) ddm:rucio,hc:IT (1) ddm:rucio,hc:UK (25) ddm:rucio,hc:US (34)	
transformation (1)	Reco_tf.py (398)	
workinggroup (1)	GP_PHYS (398)	

Job list Only the most recent 100 jobs are shown. Remove the limit and sort by PandaID, time since last state change, ascending mod time, descending mod time, priority, attemptnr, ascending duration, descending duration												
PanDA ID Attempt#	Owner Group	Request Task ID	Transformation	Status	Created	Time to start d:h:m:s	Duration d:h:m:s	Mod	Cloud Site	Priority	Maximum PSS	Job info
2894682604 Attempt 2	atlas-dpd-production GP_PHYS	7945 8737370	Reco_tf.py	finished	2016-06-20 20:27	0:0:08:53	0:5:24:20	06-21 03:17	WORLD MWT2_SL6	560	1717.76	
	Job name: data15_13TeV.00276329.physics_Main.merge.r7562_p2521_p2685.2892021198 #2 Datasets: In: data15_13TeV:data15_13TeV.00276329.physics_Main.merge.AOD.r7562_p2521_tid07697857_00 Out: data15_13TeV.00276329.physics_Main.merge.TAG.r7562_p2521_p2685_tid08737370_00											
2892021198 Attempt 1	atlas-dpd-production GP_PHYS	7945 8737370	Reco_tf.py	failed	2016-06-18 11:54	0:3:59:30	0:4:18:51	06-20 20:25	WORLD CA-SCINET-T2	560	1557.94	Executable error 1137: Put error: Error in copying the file from job workdir to localSE Dispatcher error 101: job recovery failed for 48 hours jobdispatcher: job recovery failed for 48 hours exe: Put error: Error in copying the file from job workdir to localSE
	Job name: data15_13TeV.00276329.physics_Main.merge.r7562_p2521_p2685.2892021198 #1 Datasets: In: data15_13TeV:data15_13TeV.00276329.physics_Main.merge.AOD.r7562_p2521_tid07697857_00 Out: data15_13TeV.00276329.physics_Main.merge.TAG.r7562_p2521_p2685_tid08737370_00											

Available views

Dashboards:

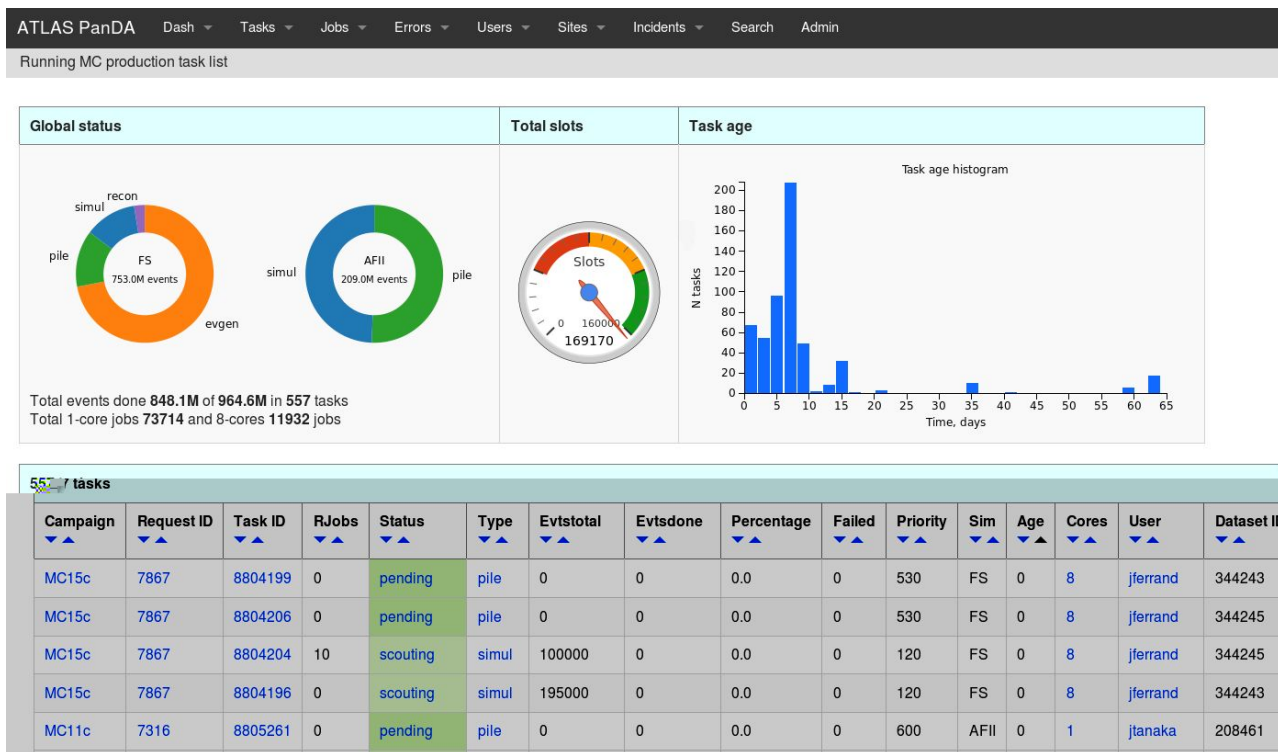
- Jobs: region and cloud views;
- Tasks and jobs summary for working groups;
- WORLD Cloud:
 - Jobs statuses and distribution;
 - HS06*sec summary for nuclei and satellites

Cloud / Site summary of production jobs - Region view		For a description of region view see below																
Cloud	Status	nJobs	nPilots	defined	waiting	assigned	throttled	activated	sent	starting	running	holding	merging	transferring	finished	failed	cancelled	% failed
All clouds		506231	36890	376	0	8351	0	118378	53	7355	96395	2350	13920	31053	204782	21779	1439	9
CA		27382	1880	64	0	757	0	1823	0	3	4707	298	338	1744	13454	4169	25	23
CERN		65245	3326	0	0	278	0	23720	0	2035	16245	120	481	4473	14383	3465	45	19
DE		46596	3504	0	0	815	0	3894	17	1279	9251	239	1574	3556	24151	1785	35	6
ES		9150	763	0	0	596	0	700	0	7	1884	71	88	562	5058	176	8	3
FR		49817	3673	0	0	524	0	5971	2	28	11759	253	105	3137	26236	1737	65	6
IT		20044	1493	16	0	250	0	3617	0	143	5343	105	182	1322	8444	597	25	6
ND		37539	813	0	0	104	0	21806	0	2548	3255	134	1050	940	6623	1041	38	13
NL		27428	1211	8	0	271	0	3110	1	1	4371	119	334	2231	16617	355	10	2
RU		10152	563	0	0	159	0	1411	0	1	3133	61	229	1025	3969	154	10	3
TW		9792	1460	125	0	219	0	448	0	345	1826	243	161	1162	4979	276	8	5
UK		50684	6128	88	0	661	0	4655	2	778	14044	322	1861	4249	22306	1666	52	6
US		152402	12076	75	0	3717	0	47223	31	187	20577	385	7517	6652	58562	6358	1118	9
CA Cloud, Sites	Status	nJobs	nPilots	defined	waiting	assigned	throttled	activated	sent	starting	running	holding	merging	transferring	finished	failed	cancelled	% failed
CA all sites		27382	1880	64	0	757	0	1823	0	3	4707	298	338	1744	13454	4169	25	23
Australia-ATLAS	online	271	35	0	0	0	0	2	0	0	25	0	0	0	241	2	1	0
Australia-ATLAS_MCORE	online	676	55	0	0	25	0	92	0	0	95	2	0	72	377	13	0	3
Australia-ATLAS_MCORE_TEST	test	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Available views


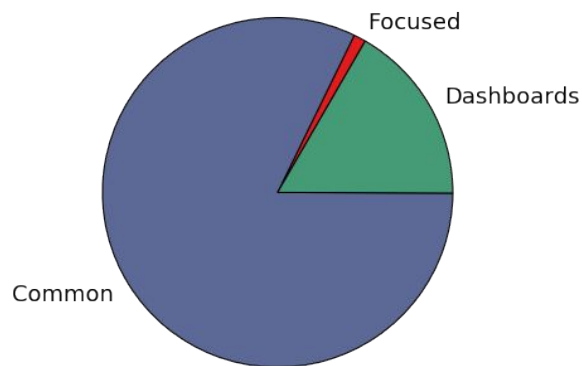
Focused views:

- 'Active' MC production tasks



Statistics of using

- 13k queries / day
- 1500 unique sessions / day



Recent improvements and new features

Performance issue:



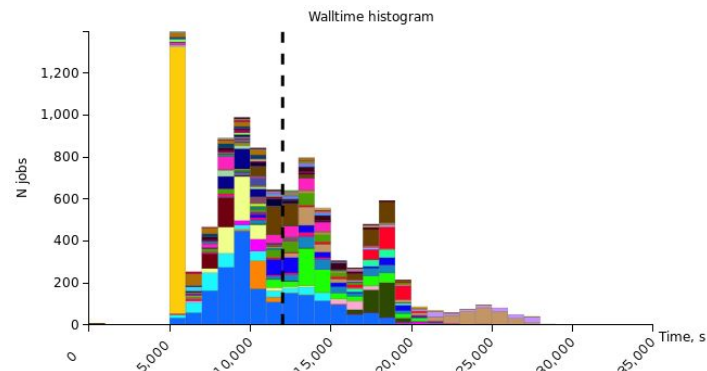
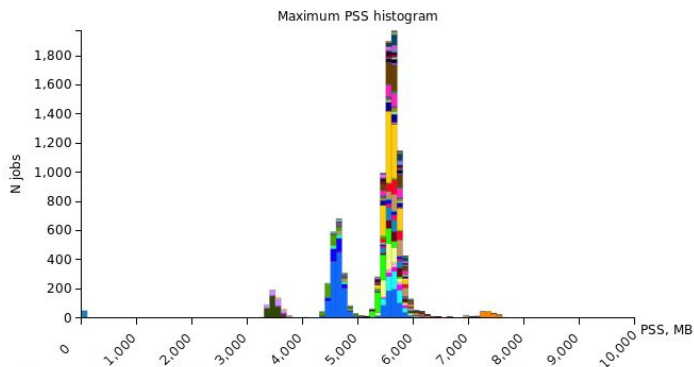
- Idea is to moving the filtering and data aggregation procedure from monitoring server side to Oracle DB server side
- Current results: decreasing loading time of jobs view from **~1 min** (30k jobs limit, last 12 hours jobs) to **~10 sec** (without limit, for last 12 hours = ~800k jobs) [thanks to Gancho Dimitrov and DBA]

It's still in research and development mode.

Recent improvements and new features

Analytics supplementation:

- Advanced plots and diagrams



Future plans

- to continue moving filtering and data aggregation procedure to the Oracle DB server side for other essential views
- to supplement more analytics
- to elaborate monitoring of Event Service

Summary

- BigPanDA Monitoring System is in production since the middle of 2014 and developing continuously (287 commits done, 121 requests received from beginning of 2016);
- It provides a comprehensive and coherent view of the tasks and jobs executed by the PanDA system, from high level summaries to detailed drill-down job diagnostics.
- It satisfied needs of groups of users including distributed computing systems operators, shifters, physicist end-users, computing managers and accounting services in web-based analytics and system state views.