

# BigPanDA monitoring

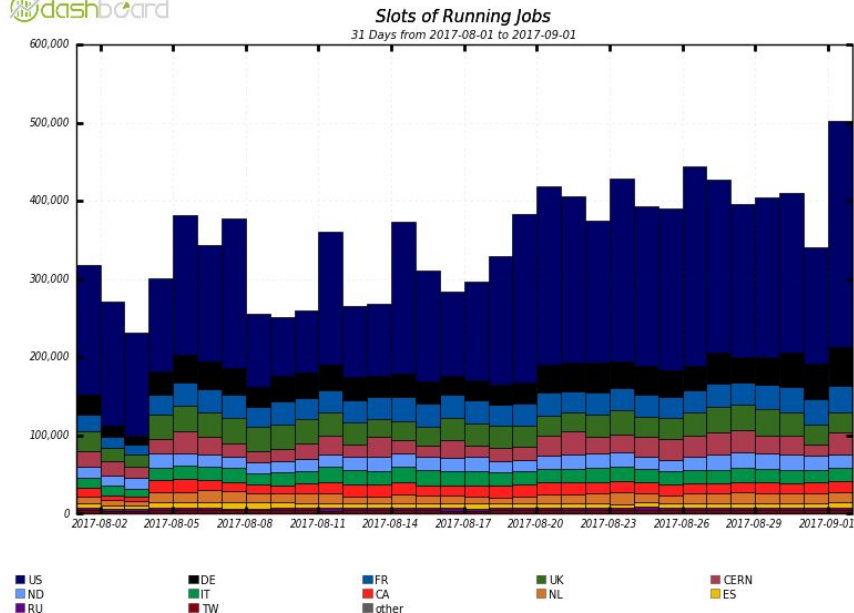
---

A.Alekseev, A.Klimentov, T.Korchuganova, S.Padolski,  
T.Wenaus

# Outlines

- PanDA Workload Management
- Overview of monitoring architecture
- Key views and features
  - Jobs
  - Task
  - Errors
  - User
  - Campaign report
- Summary

# PanDA Workload Management



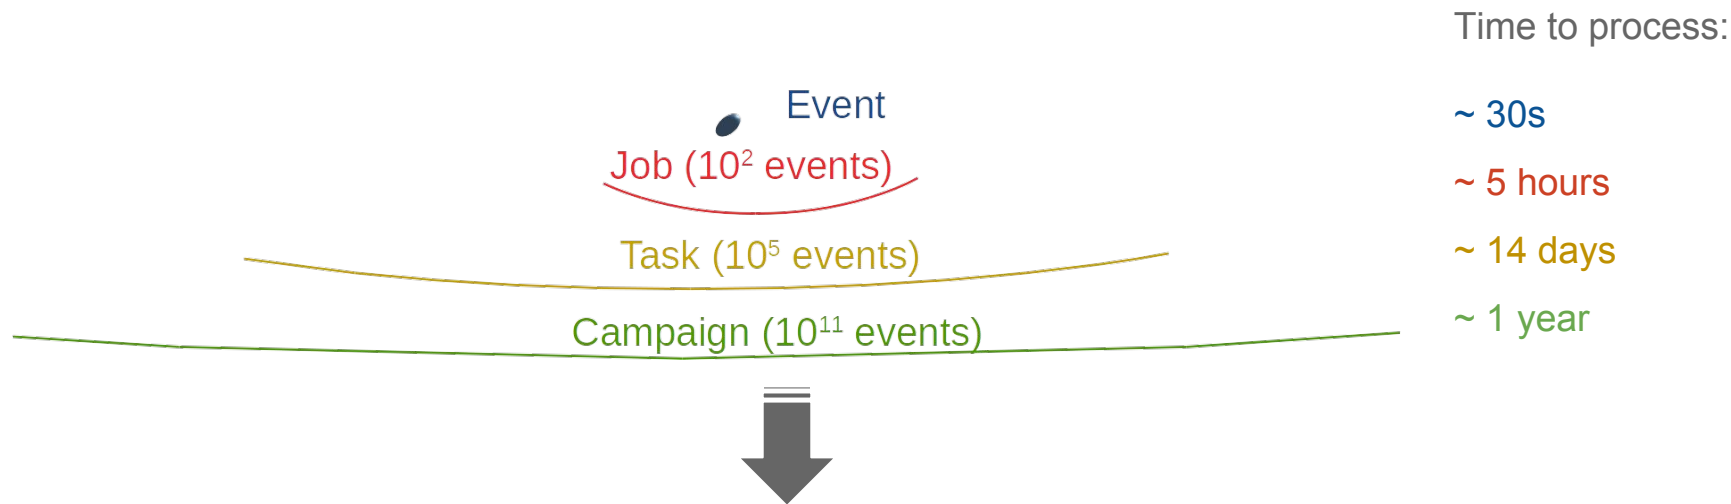
PanDA (Production and Distributed Analysis) Workload Management System is managing the workflow for all data processing on over 200 data centers and opportunistic resources such as commercial clouds, supercomputers, volunteer machines.

~250k simultaneous jobs

~2M jobs a day

~3k users

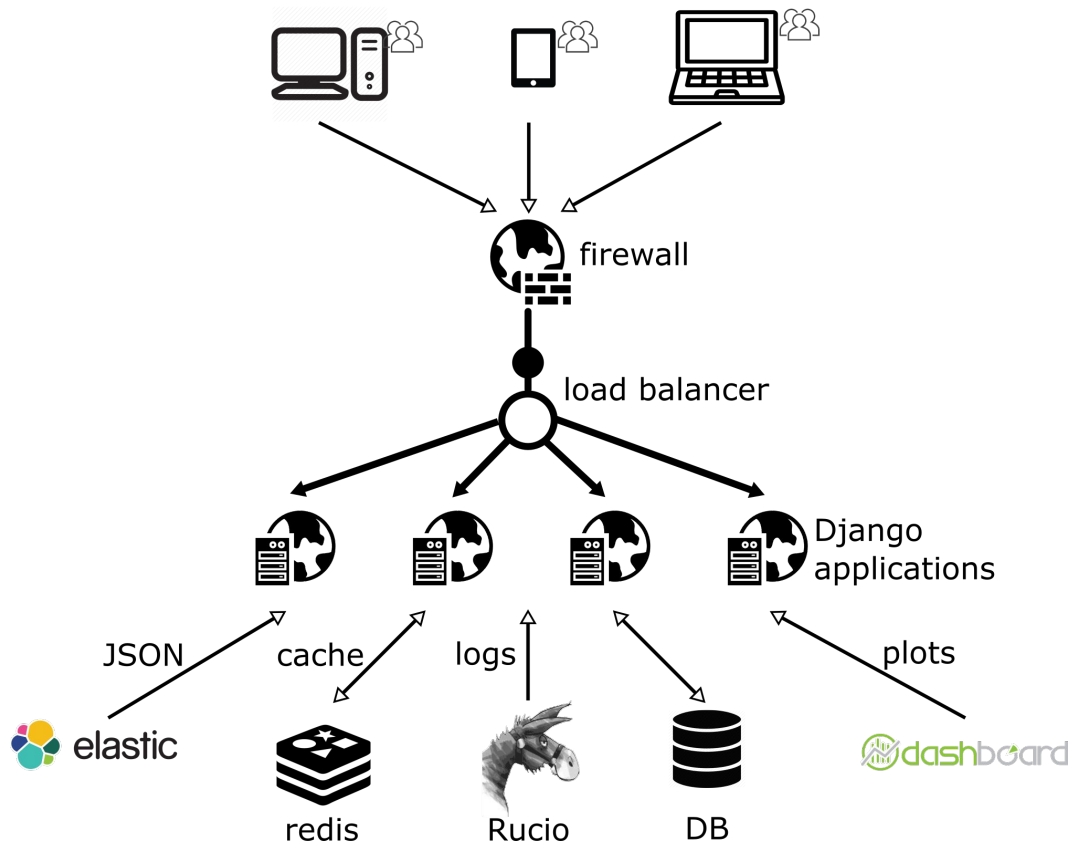
# PanDA processing scales



- Multilevel granularity of data to aggregate and display
- Providing the real-time and long-term-history (up to 2 year) of aggregated information
- Drilling-down from high level summaries to detailed diagnostics data

Goal of monitoring system is the rapid identification of failures and monitoring of progress of a distributed physics analysis and production.

# BigPanDA monitoring architecture



# BigPanDA monitoring system

- Using Model View Controller approach provided by Django framework
- Various DB backends (Oracle or MySQL)
- Gaining data from various sources (ElasticSearch, Rucio, Dashboard service)
- Data processing algorithms splitted between DB and application backends
- Responsive Web designing with Foundation
- Advanced plots generation on client side using d3.js library or matplotlib on server side
- Dynamic data delivery (Ajax)
- Proactive caching



Foundation



# Jobs Page

specified parameters

total number of jobs matched to specified parameters

ATLAS PanDA Dash Tasks Jobs Errors Users Sites Incidents Search Admin Prodsys Services VO Help

PanDA jobs, last 12 hours. Params: limit=10000 Warning: limit 20000 per job table Task has 0 jobs in table Jobsarchived, limit is set to 20000 Total jobs found ≈ 1470000 aipanda105 08:51:18, Reload Login

30013 jobs in this selection

Job modification times in this listing range from 2017-09-25 20:50:36 to 2017-09-26 08:50:36.

Job current priorities in this listing range from 1000000 to -1000000. See priorityrange in the job attribute summary to see how priorities are distributed.

Job attribute summary Sort by count, alpha	
atlasrelease (41)	Atlas-17.2.2 (7) Atlas-17.2.8 (428) Atlas-17.6.0 (10) Atlas-19.2.3 (176) Atlas-19.2.4 (142) Atlas-19.2.5 (5533) Atlas-2.6.4 (8) Atlas-20.1.5 (20) Atlas-20.1.8 (2) Atlas-20.20.10 (9) Atlas-20.20.6 (5) Atlas-20.20.7 (1) Atlas-20.20.8 (486) Atlas-20.7.5 (600) Atlas-20.7.6 (84) Atlas-20.7.7 (393) Atlas-20.7.8 (120) Atlas-21.0 (18) Atlas-21.0.13 (2) Atlas-21.0.14 (95) Atlas-21.0.15 (9358) Atlas-21.0.19 (708) ... more
attemptnr (59)	1 (16519) 2 (10147) 3 (1207) 4 (348) 5 (267) 6 (83) 7 (34) 8 (46) 9 (21) 10 (29) 11 (6) 12 (22) 13 (76) 14 (90) 15 (5) 16 (2) 17 (25) 19 (1) 20 (1) 21 (2) 22 (25) 23 (2) 24 (1) 26 (2) 30 (1) ... more
cloud (13)	CA (1204) CERN (1146) DE (1556) ES (360) FR (1406) IT (322) ND (663) NL (942) RU (467) TW (153) UK (2477) US (3002) WORLD (16291)
computingsite (381)	AGLT2_LMEM (45) AGLT2_MCORE (56) AGLT2_SL6 (139) ALCF_Theta (826) ANALY_AGLT2_SL6 (134) ANALY_AGLT2_TEST_SL6-condor (1) ANALY_AGLT2_TIER3_TEST (5) ANALY_ARNES (18) ANALY_ARNES_DIRECT (78) ANALY_AUSTRALIA (15) ANALY_AUSTRALIA_TEST (2) ANALY_BEIJING (6) ANALY_BHAM_SL6 (5) ANALY_BNL_CLOUD (2) ANALY_BNL_LONG (307) ANALY_BNL_LONG_CONTR (16) ANALY_BNL_MCORE (2) ANALY_BNL_SHORT (180) ANALY_BNL_SHORT_CONTR (38) ANALY_BNL_Test_2_CE_1 (4) ANALY_BU_ATLAS_Tier2_SL6 (10) ANALY_CAM_SL6 (23) ... more

# Jobs Page

Overall error summary			
Category:code	Attempt list	Nerrors	Sample error description
<a href="#">ddm:200</a>	<a href="#">jobs</a>	3	Could not get GUID/LFN/MD5/FSIZE/SURL from pilot XML
<a href="#">exe:-1</a>	<a href="#">jobs</a>	67	LRMS error: (257) Job timeout

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	Job info
<a href="#">3315291794</a> Attempt 1	<a href="#">sgonzale AP_SOFT</a>	<a href="#">11341 11081388</a>	<a href="#">POOLtoEI_tf.py</a>	activated	2017-04-01 06:07:53	0:0:52:59		2017-04-01 07:00:43	<a href="#">WORLD SARA-MATRIX online</a> no active blacklisting rules defined	880	
	Job name: <a href="#">data_evind.00311365.physics_Main.eventIndex.r9264_p3083_i14.3315291794</a> #1										
	Datasets: In: <a href="#">data16_13TeV:data16_13TeV.00311365.physics_Main.merge.AOD.r9264_p3083_tid11038617_00</a> Out: <a href="#">data_evind.00311365.physics_Main.eventIndex.log.r9264_p3083_i14_tid11081388_00</a>										



# Task Page

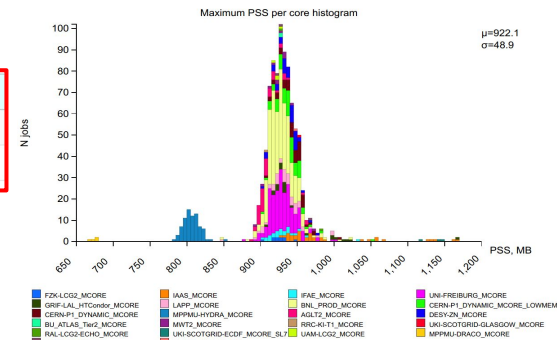
11016615 task: mc15\_13TeV.410533.PowhegPythia8EvtGen\_A14\_ttbar\_hdamp258p75\_allhad\_NoC\_bfil.simul.e5884\_s2726

Task ID	Request	Type	WorkingGroup	User	Destination	Campaign	Task status	Nevents   used	HS06*sec Expected Total done failed	Ninputfiles   finished   failed	Average maxpss	Created	Modified	Cores	Priority	Parent	Tracker
11016615	11397	prod	AP_TOPQ	dhirsch	UNI-FREIBURG	MC15:MC15c	running	980000   952000 (97%)	7782180000 4930938703 4421823312 509115391	980   952 (97%)   25 (2%)	8053	2017-03-23 08:50:37	2017-04-01 02:50:35	8	420	11009029	JIRA

States of jobs in this task (merge jobs excluded) [Show all jobs](#)

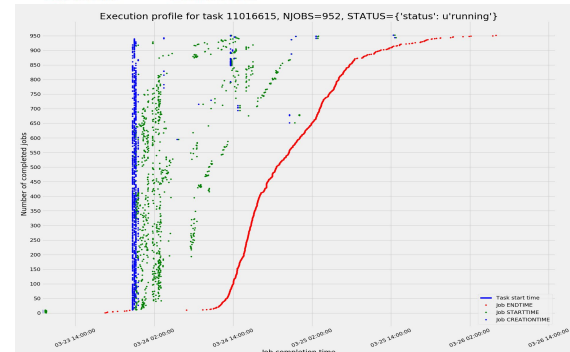
defined	waiting	pending	assigned	throttled	activated	sent	starting	running	holding	transferring	finished	failed	cancelled	merging	closed
					1			2			952	25			

jobs belonging to the task and its statuses



## Task Chain Gantt Diagram

Format: Hour Day Week Month Quarter							Mar 2017 - Week 12							Mar 2017 - Week 13							Apr 2017 - Week 14						
	Step	Duration	% Comp.	Start Date	End Date	Status	20	21	22	23	24	25	26	27	28	29	30	31	01	02	03	04	05	06	07	08	09
11009029	evgen	1.5 Days	100%	22-03-2017 09:26	23-03-2017 20:30	finished																					
11016615	simul	10.4 Days	86%	23-03-2017 08:57	02-04-2017 17:37	running																					
11016617	recon	8.1 Days	94%	24-03-2017 16:38	01-04-2017 18:49	running																					
11016620	merge	7 Days	100%	25-03-2017 06:17	01-04-2017 06:48	running																					
11016623	merge	0 Days		25-03-2017 07:17	25-03-2017 07:17	registered																					



# User page

Inspired by history

## Tasks:

**10641439**, eventservice:eventservice, tasktype:prod, tasktype:prod(ramcount:1-2GB), **10861821**, tasktype:prod(ramcount:1-2GB,reqid:11095), username:Federico Guillermo Diaz-Capriles, tasks

## Jobs:

jeditaskid:**10641439**(eventservice:not2), jobs, computingsite:ORNL\_Titan\_MCORE, jeditaskid:**10861821**(jobstatus:finished), jeditaskid:**10838379**(produsername:Federico Guillermo Diaz-Capriles)

## Other:

globalshares, runningdppdprodtasks, runningmcprodtasks(reqid:10208), users, users


## 100 tasks, sorted by jeditaskid

ID Parent	Task name TaskType/ProcessingType Campaign Group User Errors Logged status	Task status Nfiles	Input files finish% fail% Nfinish Nfail	Total/Remaining events	Modified	State changed	Priority	Nucleus	Cloud
10926905 10926899	mc15_13TeV.387730.MGPy8EG_A14N23LO_TT_stau_800_590.merge.e4485_a766_a777_r6282 prod/merge MC15:MC15a AP_SUSY mcfayden JIRA RequestID:4295 Errors	done 10	100% 10	20000 / 0	2017-03-19 10:21:26	2017-03-19 10:21:26	820	IN2P3- CC	WORLD

## Job list 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	Job info
3273108718 Attempt 1	mcfayden AP_SUSY	4295 10926895	Sim_tf.py	finished	2017-03-10 07:23:30	0:0:23:06	0:9:43:42	2017-03-10 21:18:24	WORLD FZK-LCG2_HIMEM	440	
Job name: mc15_13TeV.387730.MGPy8EG_A14N23LO_TT_stau_800_590.simul.e4485_a766.3273108718 #1											
Datasets: In: mc15_13TeV:mc15_13TeV.387730.MGPy8EG_A14N23LO_TT_stau_800_590.evgen.EVNT.e4485/ Out: mc15_13TeV.387730.MGPy8EG_A14N23LO_TT_stau_800_590.simul.HITS.e4485_a766_tid10926895_00											

# Errors page



## Overall error summary

Category:code	Job list	Errors	Sample error description
<a href="#">taskbuffer:113</a>	<a href="#">jobs</a>	2186	merge job failed <a href="#">more information here</a>
<a href="#">exe:68</a>	<a href="#">jobs</a>	2126	Fatal error in athena logfile: "Logfile error in log.generate: "Sherpa_i FATAL ../src/GenModule.cxx:35 (StatusCode GenModule::initialize()); code 0: genInitialize()" <a href="#">more information here</a>

jediTaskID ▾	PandaID ⚙	Error message
<a href="#">12205171</a>	<a href="#">3622681009</a>	Fatal error in athena logfile: "Logfile error in log.generate: "TTbarWtoLeptonF... ERROR No t or tbar quarks were found decaying to W in a (presumably) ttbar event! Event is rejected. Event dump follows.""
<a href="#">12204846</a>	<a href="#">3622461806</a>	Fatal error in athena logfile: "Logfile error in log.RAWtoESD: "MdtSubdetectorMap ERROR Rod ID not found !""
	<a href="#">3622461812</a>	Fatal error in athena logfile: "Logfile error in log.RAWtoESD: "MdtSubdetectorMap ERROR Rod ID not found !""
<a href="#">12204841</a>	<a href="#">3622640003</a>	Fatal error in athena logfile: "Long FATAL message at line 2408 (see jobReport for further details)"
	<a href="#">3622640001</a>	Fatal error in athena logfile: "Long FATAL message at line 2416 (see jobReport for further details)"

# Campaign reports

## MC16a Campaign Report(▽)

Built on 26 Sep 2017 08:23:22 (UTC)

Overall events processing summary						
Input Events	Simulated	HITS Merged	Reconstructed	AOD Merged	Derived	Derivation Merged
2,027,790,000	2,015,282,490 (99.38%)	29,192,150 ( <sup>A</sup> )	2,025,296,000 (99.88%)	2,025,420,000 (99.88%)	2,027,290,000 (99.98%)	2,056,700,000 (101.43%)

Overall summary  
of events belongs  
to MC16a  
campaign

Events processing summary																
	pending	defined	assigned	waiting	activated	sent	starting	running	holding	transferring	merging	finished	failed	cancelled	throttled	closed
Evgen																
Evgen Merge																
Simulation											1,000	2,015,282,490	1,038,850	13,000		71,900
HITS Merge												29,192,150	1,800			
Reconstruction												2,025,296,000	18,000			2,000
AOD Merge												2,025,420,000	288,000			10,000
Derivation												2,027,290,000	500,000			
Derivation Merge												2,056,700,000				

# Conclusion

- BigPanDA Monitoring system is in production since the middle of 2014 and developing continuously;
- It serves 17k requests daily;
- 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 (~15 views are available);
- It is used beyond ATLAS