



NEC'2019



XXVII International Symposium on Nuclear Electronics & Computing

Montenegro, Budva, Becici, 30 September - 4 October 2019

Monitoring and Accounting for the Distributed Computing System of the ATLAS Experiment

Dario Barberis¹, Alberto Aimar², Aleksandr Alekseev³, Pedro Andrade²,
Thomas Beermann⁴, Robert Gardner⁵, Borja Garrido Bear²,
Tatiana Korchuganova³, Luca Magnoni², Siarhei Padolski⁶, Eric Schanet⁷,
Nikolay Tsvetkov², Ilija Vukotić⁵, Torre Wenaus⁶

1) Genoa Univ./INFN, 2) CERN-IT, 3) Moscow State Univ. & Plekhanov Russian Univ. of Economics,
4) Wuppertal Univ., 5) Univ. of Chicago, 6) Brookhaven National Laboratory, 7) LMU Munich



Introduction

- ATLAS used during LHC Run 1 and Run 2 a monitoring and accounting infrastructure for the Distributed Computing (ADC) applications developed ~10 years ago by CERN-IT together with ATLAS members
 - These "old dashboards" started showing aging effects in the last few years:
 - Slowness of data retrieval due to the massive amount of data in Oracle databases
 - Lack of in-depth knowledge for maintenance as original developers left long ago
 - Lack of flexibility and impossibility to develop new views and/or data correlations across different data sources
 - This system worked well enough for general monitoring till the end of Run 2 last year but was evidently in need of a good refurbishing
- Since 2016 the CERN-IT MonIT group started developing a new infrastructure and environment for monitoring and accounting applications base on modern Open Source components
 - ATLAS started implementing "new" dashboards using this infrastructure, for data and workload accounting and global monitoring
- In the meantime the BigPandaMon application was developed for user and task oriented monitoring of the jobs submitted to the ATLAS Grid/Cloud/HPC resources through PanDA
 - This is now the workhorse of user-level job monitoring
- In the recent years many Analytics tools appeared on the market. They can be used for more detailed investigations and to correlate data from different sources
 - The Analytics cluster provided by the University of Chicago allows a more interactive use of monitoring data



Topics

- ADC dashboards in the MonIT infrastructure
- User level job/task monitoring with BigPandaMon
- Analytics cluster at UC and its applications

ATLAS EXPERIMENT *ADC Monitoring*

Data Management

- DDM Transfer Dashboard (current)
- DDM Transfer Dashboard (historical)
- DDM Global Accounting (snapshot)
- DDM Global Accounting (historical)
- DDM Site Accounting
- WLCG Transfer Dashboard
- FTS Transfer Dashboard
- Rucio Dataset Recovery Service
- Rucio Storage Monitoring

Grid Data Processing

- BigPanDA Monitor
- Job Accounting Dashboard
- Kibana expert views
- Harvester Grafana
- Harvester Kibana

Tier-0 and Point 1

- Tier-0 contZole Graph
- Tier-0 contZole Tasks
- Detector Operation

Sites

- AGIS
- Sites Downtimes
- AGIS DDM Blacklisting
- Hammercloud
- Pilot factory
- SAM3 Visualization
- ASAP metric

Databases

- DB dashboard
- Frontier AWSTAT
- Frontier Launchpad Kibana
- Frontier Squid

Miscellaneous

- Central Services Kibana
- Live Aggregated View
- ADC Monitoring JIRA
- ADC Monitoring TWiki
- Grafana at CERN
- Kibana at CERN
- Kibana (prod) at UC
- Kibana (dev) at UC

Entry page for ADC Monitoring:
<https://atlas-adc-monitoring.web.cern.ch>



Topics

- ADC dashboards in the MonIT infrastructure
- User level job/task monitoring with BigPandaMon
- Analytics cluster at UC and its applications



The MonIT Infrastructure

- Monitoring Mission:

- Provide Monitoring as a Service for CERN Data Centre (DC), IT Services and the WLCG collaboration
 - e.g. Dashboards, Alarms, Search, Archive
- Collect, transport, store and process metrics and logs for applications and infrastructure

- Challenges:

- Rate & Volume

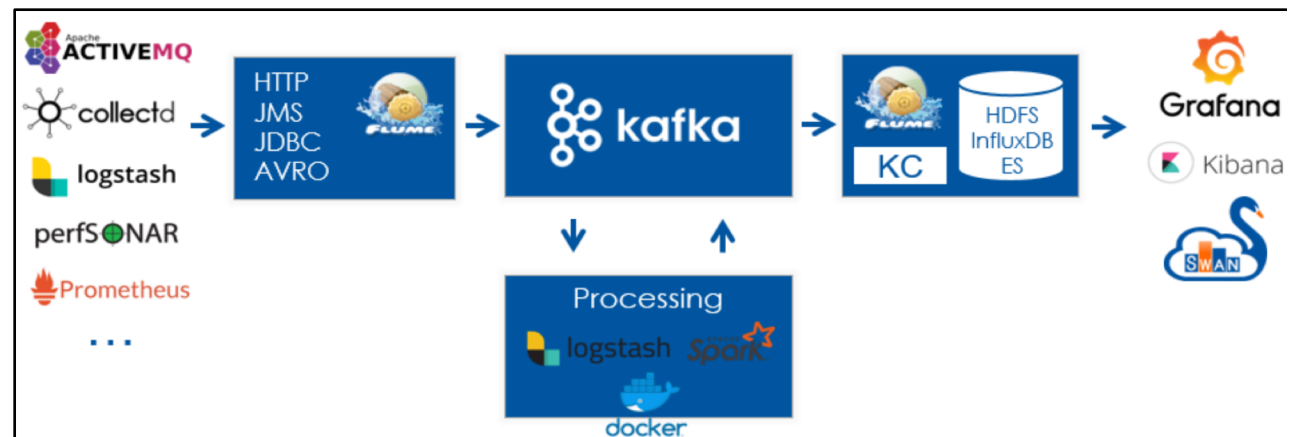
- from ~ 40k machines
- > 3 TB/day
- ~ 100 kHz

- Variety

- > 150 producers

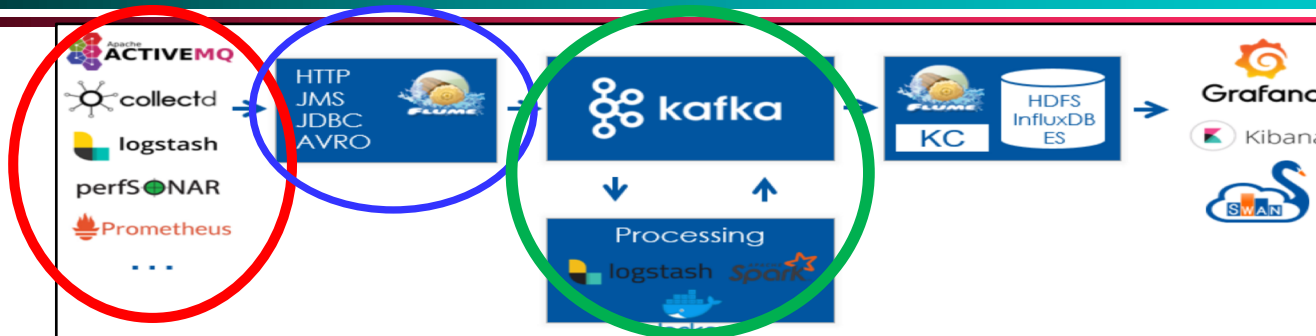
- Reliability

- spikes in rate and volume





The MonIT Infrastructure



Collectors / Collectd



• What is Collectd

"collectd is a daemon which collects system and application performance metrics periodically and provides mechanisms to store the values in a variety of ways"

• Collectd for CERN DC and IT Services

- Default HW/OS plugins on all DC machines
- Additional plugins configured by IT Service Managers
 - Community-plugins or custom-made
- Provide local alarms (~ Nagios Check) and actuators
- Sampling intervals: 1 to 5 minutes

Connectors / Apache Flume



We rely on several Flume features:

- Transactional nature
- Pull/push modes
- Protocol-based agents (sources and sinks)
 - AVRO, JDBC, JMS, HDFS, Elasticserch, HTTP, Kafka
- Interceptor / Morphlines for event transformation and validation

Collectors / Flume & Logstash



- Flume for basic log forwarding
 - Running on all DC nodes for syslog
- Logstash for advanced use cases



- Many read/write plugins
- Allows log parsing and tokenization
- Other lightweight options getting popular
 - Filebeat, Fluentd

Transport / Apache Kafka



- Kafka as rock-solid core of our pipeline
 - decouples producers / consumers
 - enables stream processing
 - resilient (72 hours data retention)



The MonIT Infrastructure



Storage / InfluxDB



- Time Series DB for time series data (metrics)
- Data kept for ~5 years (auto down-sampling)
- >30 instances (from 8 to 128GB mem each)
- Data cardinality matters
- Aggregated write: ~65k points per second
- Gross read (single instance): ~100k pps

Storage / HDFS



- HDFS for long term archive
- Compressed JSON (or Parquet)
- Data kept ~ forever (or by GDPR agreement)
- Works well with batch and data-intensive analytics workflows
 - e.g. Hadoop, Spark


Storage / Elasticsearch



- Distributed search and indexing engine
- Powerful query language to browse and explore
 - Apache Lucene based
- 3 clusters (syslog, service logs and metrics)
 - ~100 TB (total storage size)
- Data kept for 1 month

Visualization

- Grafana for dashboards
 - Users can create their own
- Kibana for data exploration
 - Data discovery and logs
- SWAN for analytics (notebooks)








ATLAS Dashboards

- “MonIT Dashboards” for Data and Job monitoring and accounting (and a number of other applications too)
 - Collect data from the Rucio and PanDA databases in Oracle and other sources
 - Transfer to BigData infrastructure using Kafka - possibility of data enrichment and correlations with other sources of information
 - (e.g. AGIS for the relation between PanDA queues, sites, federations, countries, pledges etc.)
 - Storage in HDFS with aggregation in InfluxDB
 - Display with Grafana
 - Still far from perfect as display tool for us but usable
 - Used by ADC shifters and site admins (monitoring) and computing managers (accounting)
- Three groups of dashboards:
 - Production
 - Development (pre-production)
 - Playground (free for all)



Data Management (DDM) Dashboards


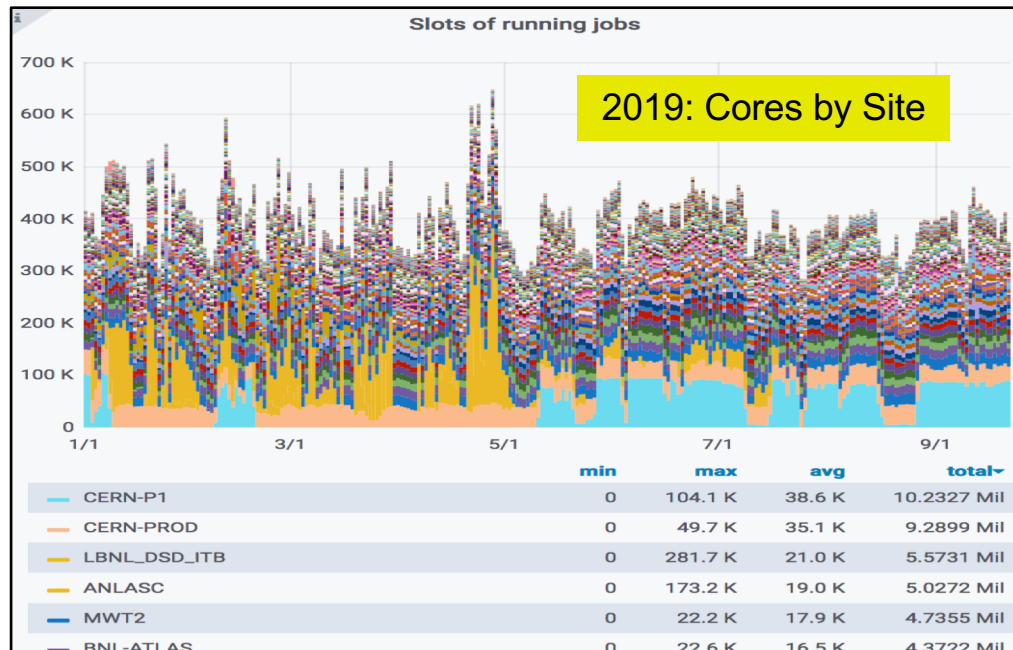
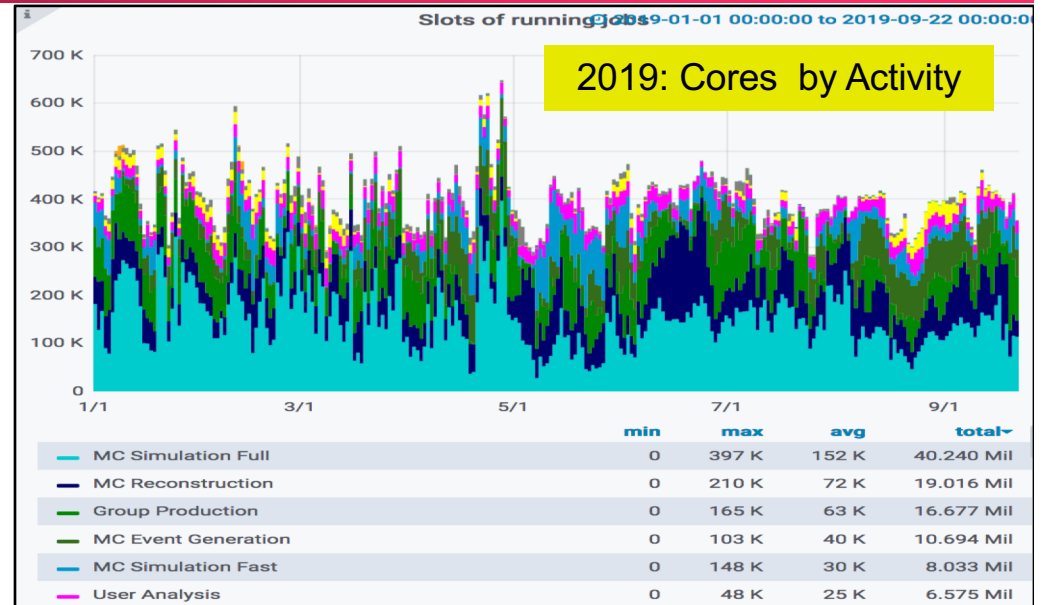
- Traces and events reported constantly by Rucio to the message brokers and further processed in the MonIT infrastructure
- Rucio data are dumped periodically from Oracle to HDFS for site accounting views
 - Then used from there to fill InfluxDB
- Many many views available
 - Historical views and snapshots
 - Data storage volumes
 - Transfer rates and efficiencies





Job Dashboard

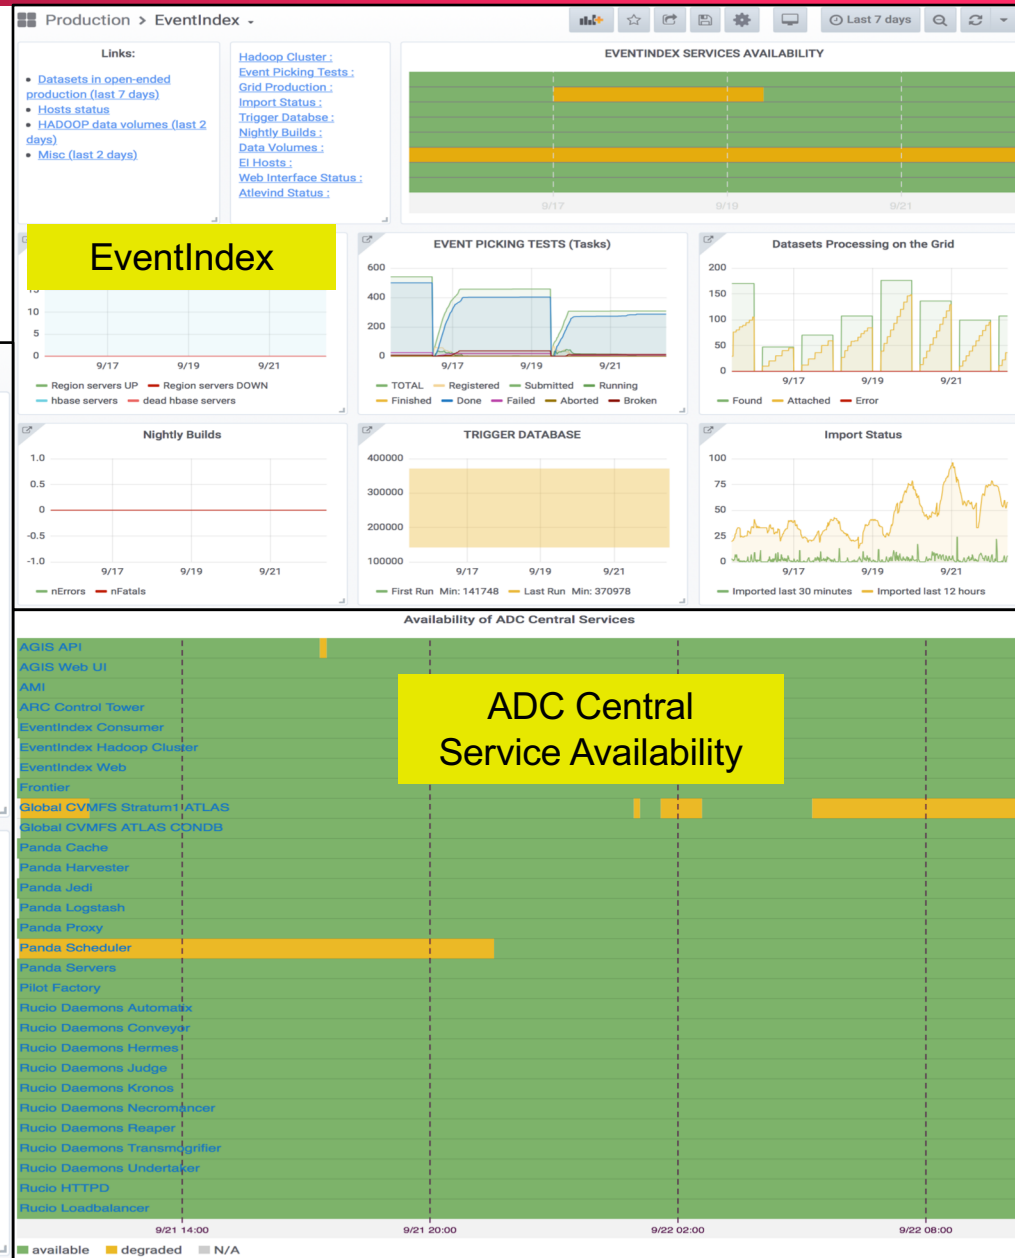
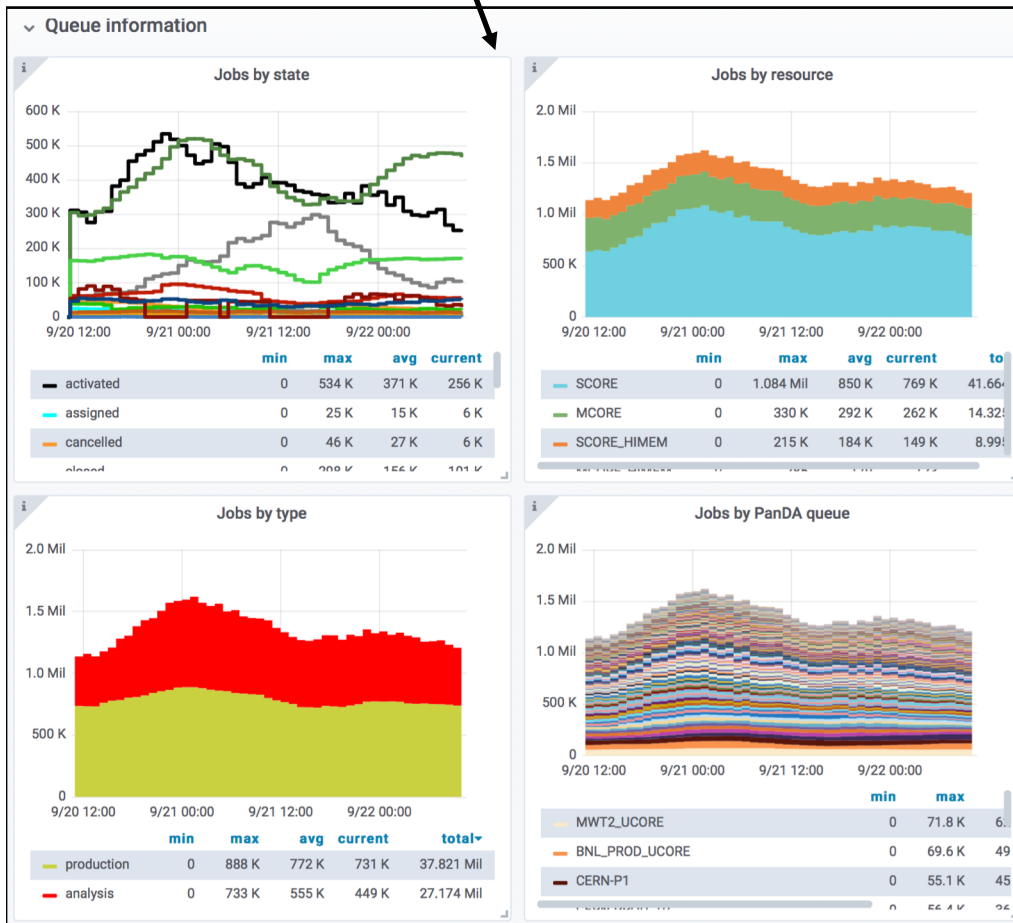
- Data collected from the PanDA DB every 10 minutes and grouped in 1-hour bins
- Transient job stats:
 - Pending, running, finalising jobs
- Permanent job stats:
 - Submitted and completed jobs
- All job parameters on submission and completion
 - Plus site info from AGIS
 - All available for selections and display
- Data imported from start of Run 1





Other Dashboards

- Several lighter dashboards have been implemented for service monitoring
- Also for short-term site oriented job monitoring





Live Page

- Derived collection of statistics and plots
 - Jobs, data storage, transfers, services
 - Inputs from different dashboards
 - Also directly from services and Elogs
- Contains links to information sources for deeper investigations
- Used by shifters and managers for daily monitoring
 - Refreshed hourly

07:16 UTC

Dashboard BigPanda Rucio Transfers HC Storage Pilots CERN Satellites Services Frontiers eLog

Slots Running (06-07 UTC) ☞ : **353k** = **154k SCORE** + **199k MSCORE**
 8h avg: **373k**; 24h avg: **359k**

NJob Fail (06-07 UTC) ☞ : **1.0k (26.4%) PROD** **0.0k (8.0%) ANALY**
 8h avg: **2.0k (10.1%) PROD** **1.3k (10.1%) ANALY**
 24h avg: **2.3k (10.7%) PROD** **1.6k (10.7%) ANALY**

Walltime Fail (4h avg): **1074d (12.5%) PROD** **185d (12.5%) ANALY**
 8h avg: **1166k (9.4%) PROD** **176k (9.4%) ANALY**
 24h avg: **1223k (8.7%) PROD** **214k (8.7%) ANALY**

Removed:src_endpoint=*LAKE*;dst_endpoint=*LAKE*
 Transfer Failure rate (last hour): **1.16 f/s**
 8h avg: **1.50 f/s**; 24h avg: **2.49 f/s**
 Transfer Throughput (last hour): **12.14 GB/s**
 8h avg: **12.21 GB/s**; 24h avg: **14.23 GB/s**
 Transfer Rate (last hour) ☞ : **23.96 f/s**
 8h avg: **25.74 f/s**; 24h avg: **26.30 f/s**





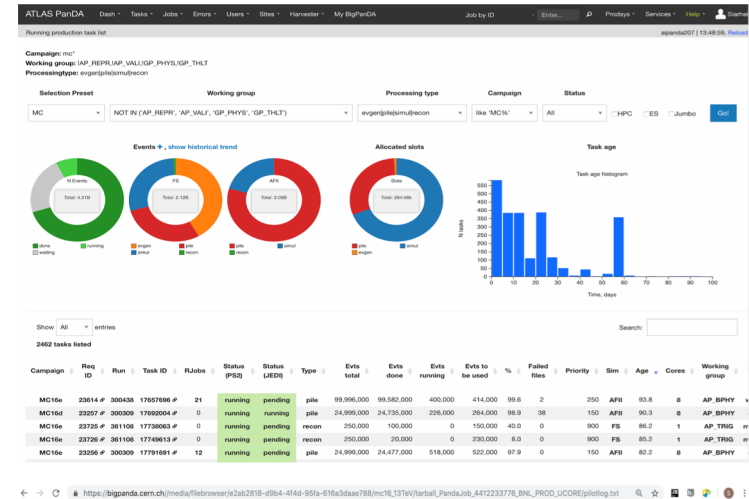
Topics

- ADC dashboards in the MonIT infrastructure
- User level job/task monitoring with BigPandaMon
- Analytics cluster at UC and its applications



BigPandaMon

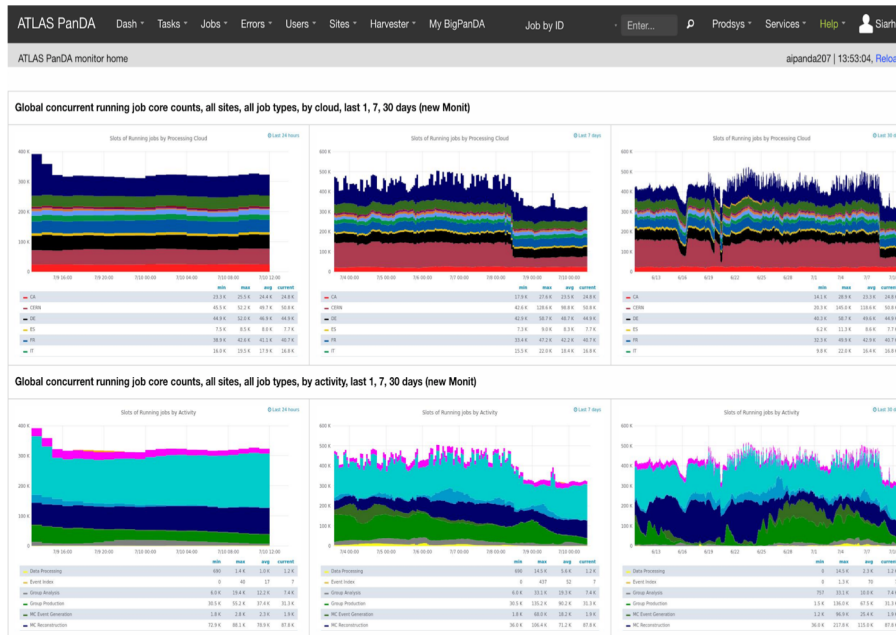
- A window into the PanDA system
- > 100 different views
- From production dashboards to logs
- Covers scope in range $1...10^{11}$ events



```

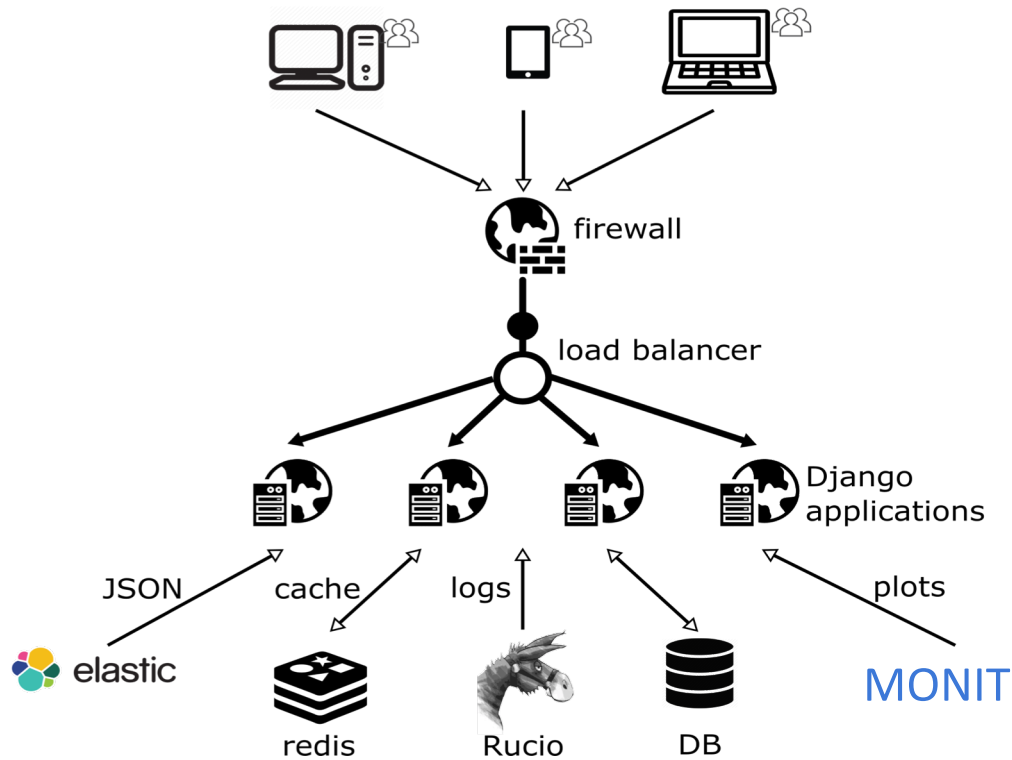
https://bigpanda.cern.ch/media/filebrowser/c2ab2818-d9b4-4f6d-95f9-616a3d8ae788/mc16_13376v1harbal_PandaJob_4412233778_BNL_PROD_UCORE/pilotlog.txt
2019-07-10 01:23:07 |143203|pUtil.py | current dir:
/home/tmp/atlas_jBm5m8Q0/Panda_Pilot_143203_1562721787
2019-07-10 01:23:07 |143203|pilot.py | Pilot options:.....
2019-07-10 01:23:07 |143203|pilot.py | appdir:
2019-07-10 01:23:07 |143203|pilot.py | debugLevel: 0
2019-07-10 01:23:07 |143203|pilot.py | jobrec: False
2019-07-10 01:23:07 |143203|pilot.py | jobRequestFlag: True
2019-07-10 01:23:07 |143203|pilot.py | jobSchedulerId: harvester-CERN_central_A
2019-07-10 01:23:07 |143203|pilot.py | maxJobrec: 20
2019-07-10 01:23:07 |143203|pilot.py | maxNumberOfRecoveryAttempts: 15
2019-07-10 01:23:07 |143203|pilot.py | pilotId: https://aipanda184.cern.ch/condor_logs_2/19-07-
09_22/grid.2148787.16.out
2019-07-10 01:23:07 |143203|pilot.py | pshttpurl: https://pandaserver.cern.ch
2019-07-10 01:23:07 |143203|pilot.py | psport: 25443
2019-07-10 01:23:07 |143203|pilot.py | queueName: BNL_PROD_UCORE
2019-07-10 01:23:07 |143203|pilot.py | rmkdir: None
2019-07-10 01:23:07 |143203|pilot.py | siteName: BNL_PROD_UCORE
2019-07-10 01:23:07 |143203|pilot.py | stageInretry: 2
2019-07-10 01:23:07 |143203|pilot.py | stageOutretry: 2
2019-07-10 01:23:07 |143203|pilot.py | uflag: managed
2019-07-10 01:23:07 |143203|pilot.py | workDir: /tmp/atlas_jBm5m8Q0/Panda_Pilot_143203_1562721787
logFileDir:
.....
2019-07-10 01:23:07 |143203|pilot.py | getSiteInformation: got experiment=ATLAS
2019-07-10 01:23:08 |143203|pUtil.py | Executing command:
/cvmfs/atlas.cern.ch/repo/sw/arc/client/latest/slc6_x86_64/setup.sh:arcproxy -i vomsACValidityLeft
ec=0 output=279137
2019-07-10 01:23:08 |143203|ATLASExperim | Voms proxy verified (279137s)
2019-07-10 01:23:08 |143203|ATLASExperim | Voms proxy verified using arcproxy
2019-07-10 01:23:08 |143203|pilot.py | Collecting WN info from: /tmp/atlas_jBm5m8Q0
2019-07-10 01:23:08 |143203|pilot.py | Got max memory limit: 64000 MB (from queuedata)
2019-07-10 01:23:08 |143203|pilot.py | Executing command: ulimit -a
core file size (blocks, -c) 4194303
data seg size (kbytes, -d) unlimited
scheduling priority (-e) 0
file size (blocks, -f) unlimited
pending signals (-i) 514006

```





Architecture & current usage



17000 json requests a day

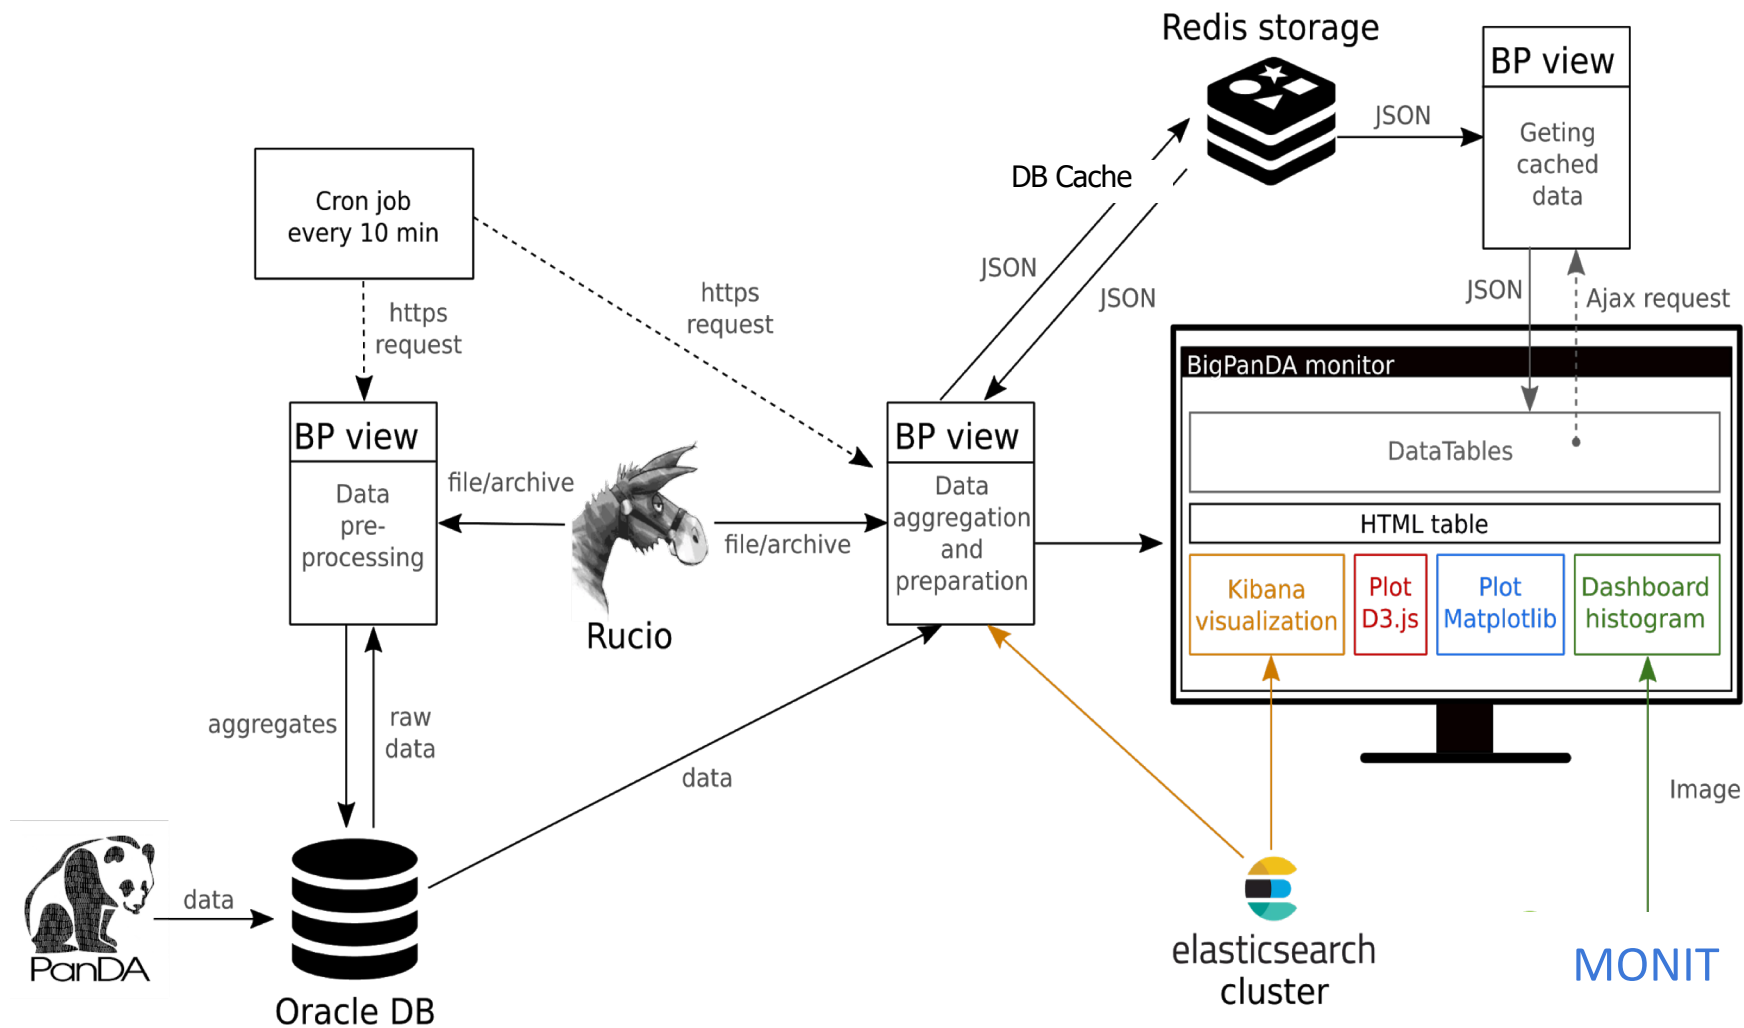


- 6.5 (+3%) user queries a day
- From 1 to 626 pages a day per user
- 1110 monthly active users
- 342 daily users

Is a primary tool ATLAS wide for shifters, experts and ADC in general



BigPandaMon Data-flow diagram





User & task 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:
 jedidtaskid:10641439(eventservice.not2), jobs, computingsite:ORNL_Titan_MCORE, jedidtaskid:10861821(jobstatus:finished), jedidtaskid:10838379(producername:Federico Guillermo Diaz-Capriles)

Other:
 globalshares, runningdpdprodtasks, runningmcpdprodtasks(reqid:10208), users, users

100 tasks, sorted by jedidtaskid

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

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	mcflayden 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

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

11016615 task: mc15_13TeV.410533.PowhegPythia8EvtGen_A14_ttbar_hdamp258p75_allhad_NoC_bfil.simul.e5884_s2726

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	Ninputfiles finished failed	Average maxpss	Created	Modified	Cores	Priority	Parent	Tracker
11016615	11397	prod	AP_TOQP	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

Switch to nodrop mode Prod Task view

Open tasks chain Open Task Profile Open Task Chain Show/hide plots Corresponding dashboard in Kibana Finish Abort

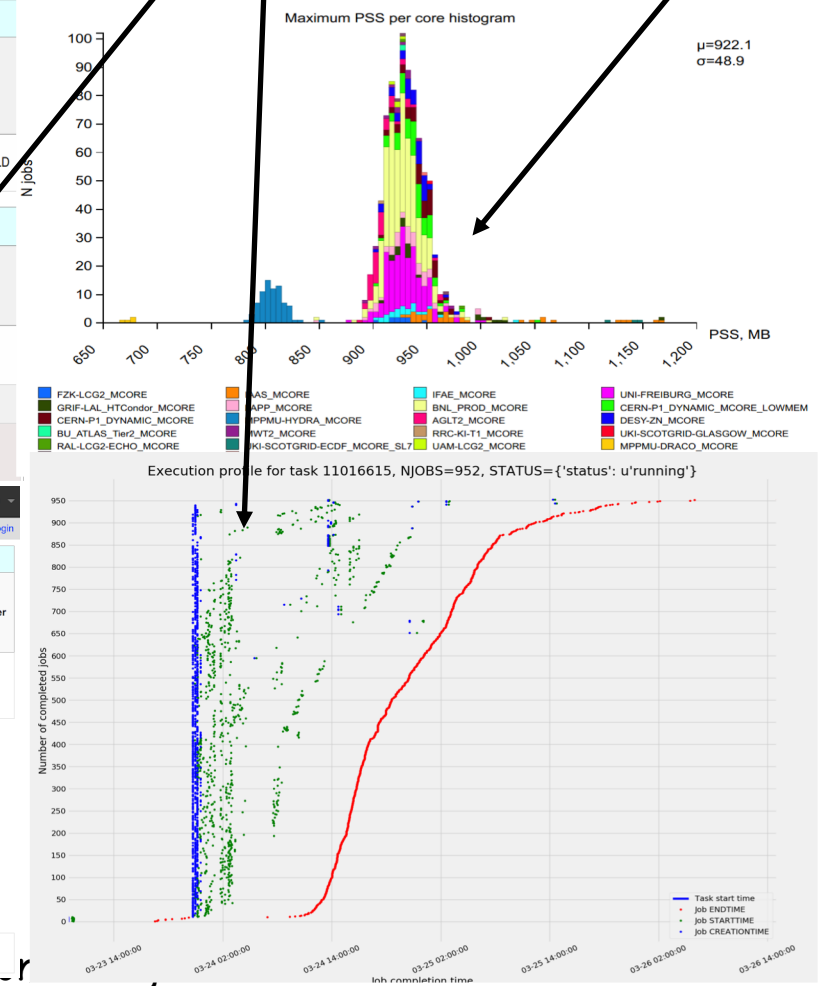
States of jobs in this task (merge jobs excluded) Show all jobs

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

Jump to job parameters, task parameters

View: job list (access to job details and logs) Show ended jobs Show active jobs Show TTC parent task 11009029 child tasks prodsys task page Task brokerage (es-atlas) Job brokerage Action logger (es-atlas) error summary

- User page: all tasks
- Task page:
 - All jobs
 - Exec CPU & memory plots
 - Task exec time





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 2000 **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 (204)	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 (207) ANALY_BNL_LONG_CONTR (46) ANALY_BNL_MCORE (2) ANALY_BNL_SHORT (489) ANALY_BNL_SHORT_CONTR (29) ANALY_BNL_Tier3_CE_1 (4)

Overall error summary

Category:code	Attempt list	Nerrors	Sample error description
ddm:200	jobs	3	Could not get GUID/LFN/MD5/FSIZE/SURL from pilot XML
exe:-1	jobs	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
3315291794 Attempt 1	sgonzale AP_SOFT	11341 11081388	POOLtoEI_tf.py	activated	2017-04-01 06:07:53	0:0:52:59		2017-04-01 07:00:43	WORLD SARA-MATRIX online no active blacklisting rules defined	880	
Job name: data_evind.00311365.physics_Main.eventIndex.r9264_p3083_i14.3315291794 #1											
Datasets: In: data16_13TeV:data16_13TeV.00311365.physics_Main.merge.AOD.r9264_p3083_tid11038617_00 Out: data_evind.00311365.physics_Main.eventIndex.log.r9264_p3083_i14_tid11081388_00											



Topics

- ADC dashboards in the MonIT infrastructure
- User level job/task monitoring with BigPandaMon
- Analytics cluster at UC and its applications



Analytics Cluster at UC

- It provides an interactive environment to:
 - Develop additional or alternative dashboards
 - Investigate correlations between several data sources
- Complementary to the monitoring and accounting infrastructure at CERN

Analytics Infrastructure

CERN

- DB - Oracle, Hadoop, Ingress, Elastic
- Processing - Sqoop, pig, Spark, SWAN, Kubernetes cluster

UC

- Elasticsearch
- K8s cluster, ML platform

Data sources:

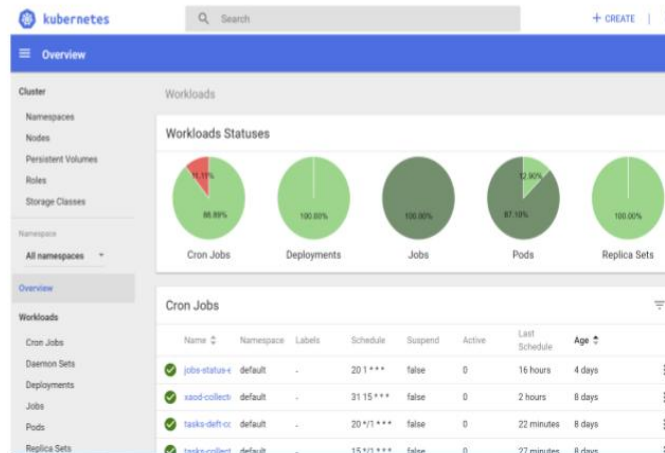
- Oracle (panda, jedi, rucio) CERN UC
- MySQL (BOINC) CERN UC
- AMQ (FTS) CERN UC
- RMQ (PerfSonar) UC
- Pilot (benchmarks) UC
- User codes (xAOD usage) CERN UC
- Frontiers (full logs) UC
- HC jobs (cost matrix) UC
- SLATE UC
- XCache UC

2

Ilija Vukotic, S&C Week, CERN October 1st, 2019

Data collection infrastructure - CERN

- All services dockerized
- All running in a single kubernetes cluster at CERN. Very stable.



Ilija Vukotic, S&C Week, CERN October 1st, 2019

Data collection infrastructure - UChicago

- All services dockerized. All running at UChicago kubernetes cluster.
- Running:
 - logstash collectors
 - rollup jobs (DDM)
 - Alarm & Alert service
 - REST interfaces (benchmarking, cost matrix, xcache backend)

Ilija Vukotic, S&C Week, CERN October 1st, 2019



Analytics Cluster at UC

- It provides an interactive environment to:
 - Develop additional or alternative dashboards
 - Investigate correlations between several data sources
- Complementary to the monitoring and accounting infrastructure at CERN



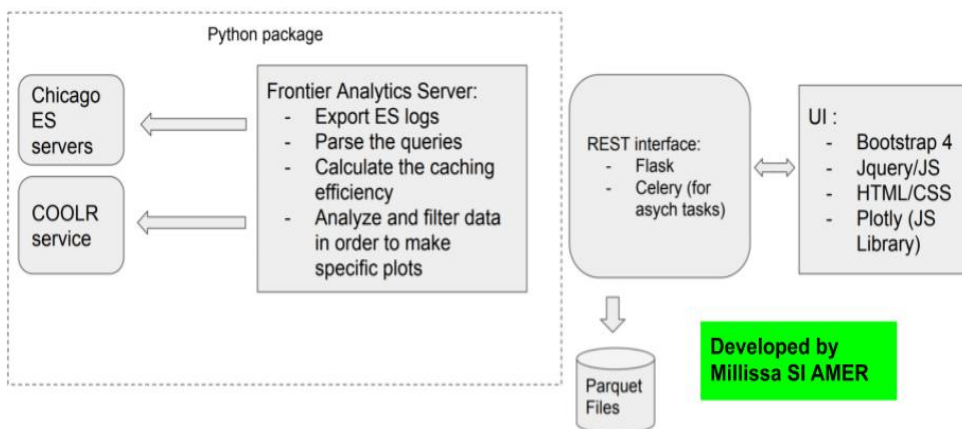
Analytics platforms

While all the data can be analysed using kibana, or JupyterLab notebooks, some projects would benefit from a custom analytics platforms:

- Extra processing and result caching
- Dedicated web site
- Custom searches
- Custom visualizations

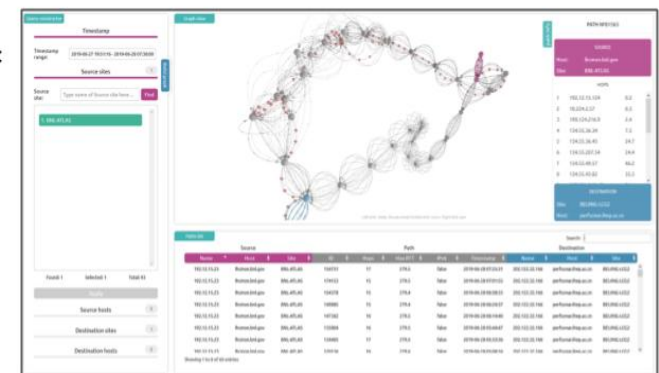
We have a way to run these platforms in UC River2 K8s cluster. Currently we support two - Frontier and Perfsonar analytics platforms. Expect more to come.

Frontier analytics tools



Perfsonar analytics platform

- Made by Евгений Третьяков Сергеевич (MEPhi)
- Django + ES + JavaScript libraries: three.js, d3-force-3d, jquery
- Future (short term)
 - Full production deployment
 - Pagination
 - Adaptive interface Full HD+ (Now only FullHD)
 - Advanced search
- Future (long term)
 - Integration of derived data
 - Other Perfsonar indices





Outlook

- ATLAS Distributed Computing has a coherent set of monitoring and accounting dashboards and interactive tools
- Technologies evolve all the time - we follow them
 - Trying to use Open Source solutions as much as possible
 - Even if at times some home-made parts are inevitable
 - See the low number of display options in Grafana
- The future is in more interactive environments providing the possibility to correlate information from many different sources
 - This is real BigData in action
 - Not just tens of billions of statistically equivalent event records!
- Any way we are quite well set for the start of LHC Run 3



Thank you!

Thank you for listening

Thanks to all authors of this work

This work was funded in part by the Russian Science Foundation
under contract No. 19-71-30008
(research is conducted in the Plekhanov Russian University of Economics)