Data Analytics in ATLAS

Ilija Vukotic,
Rob Gardner, Lincoln Briant

US ATLAS Distributed Facilities Meeting, Clemson March 14, 2016

ATLAS EXPERIMENT

Year ago

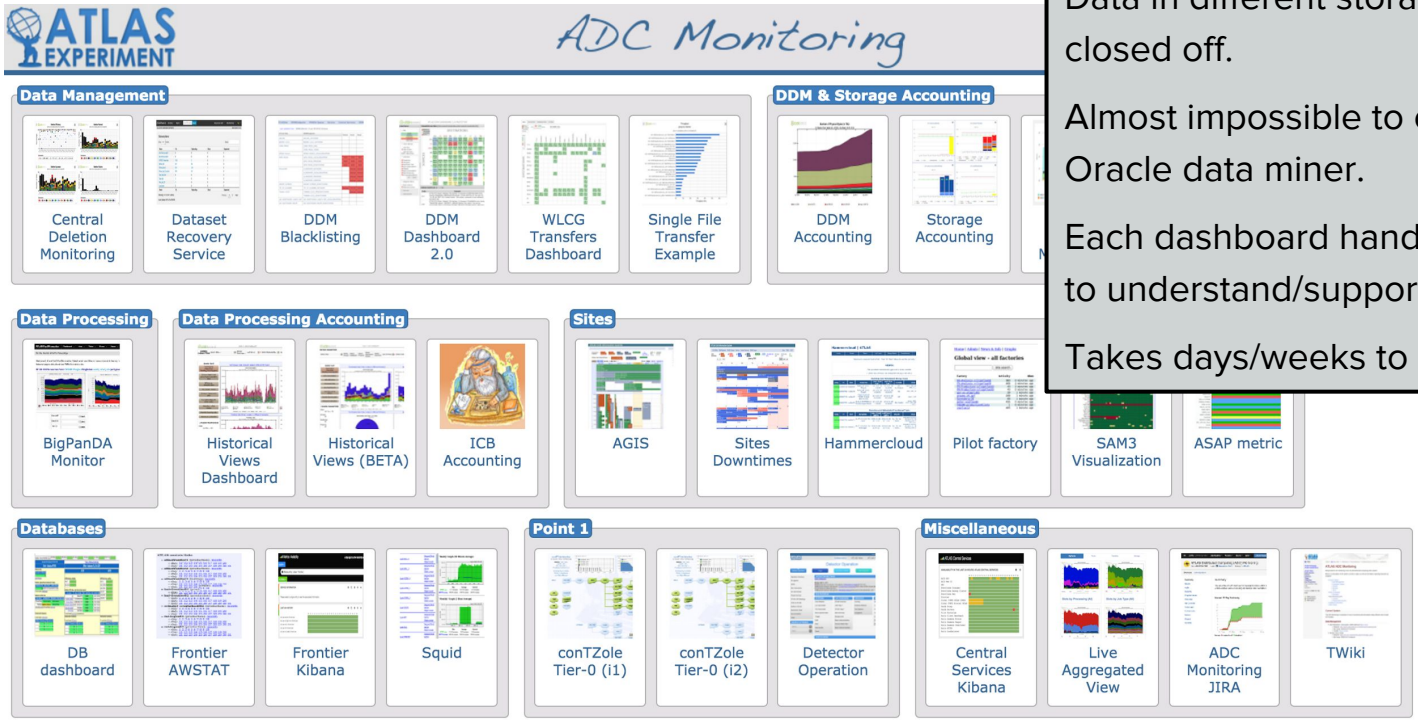

Data in different storages. Mostly Oracle, access closed off.

Almost impossible to combine. Nobody uses Oracle data miner.

Each dashboard handmade, each different. Hard to understand/support.

Takes days/weeks to get a plot added.

ATLAS analytics today - Architecture

Main functions

- Acquisition, filtering and upload of data sources into a repository
- Hadoop cluster for analysis of multiple data sources to create reduced collections for higher level analytics
- Serve repository collections in multiple formats to external clients
- Makes collected sources available for export by external users
- Host analytics services on the platform such as ElasticSearch, Logstash, Kibana, etc.

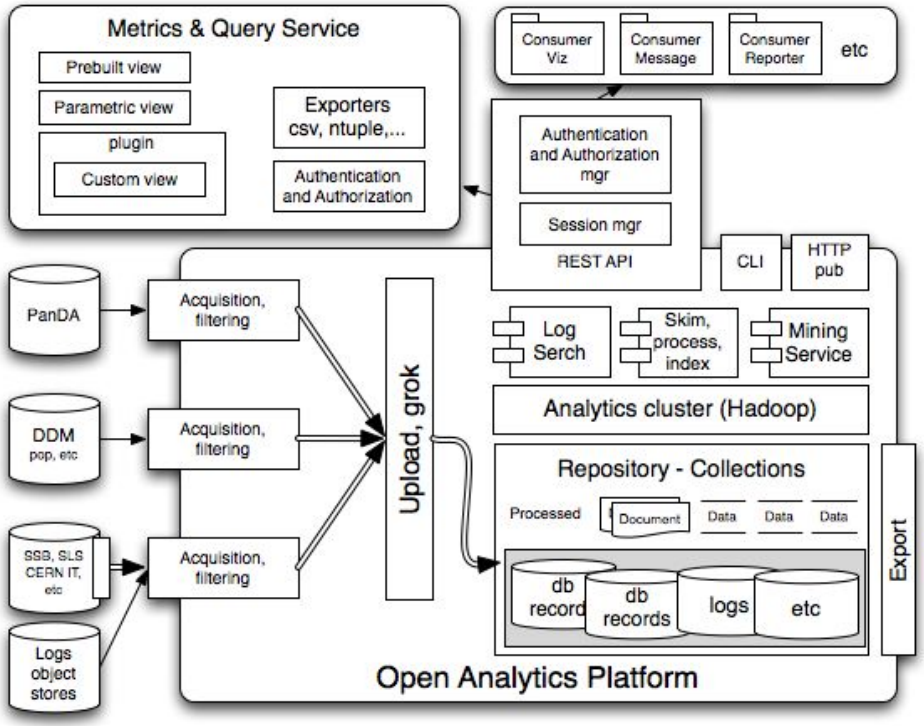

Analytics Platform: Resources & Sources

- ***lxhadoop and analytix*** clusters
  - runs base load map-reduce jobs
- ***voatlasanalytics*** cluster
  - 5 VM nodes
  - HDFS IO operations
  - Runs
    - flume collectors
    - custom made python collectors (from AMQ)
    - Sqoop jobs
    - ElasticSearch
    - Kibana
- All clusters see the same data
- All have pig, Spark

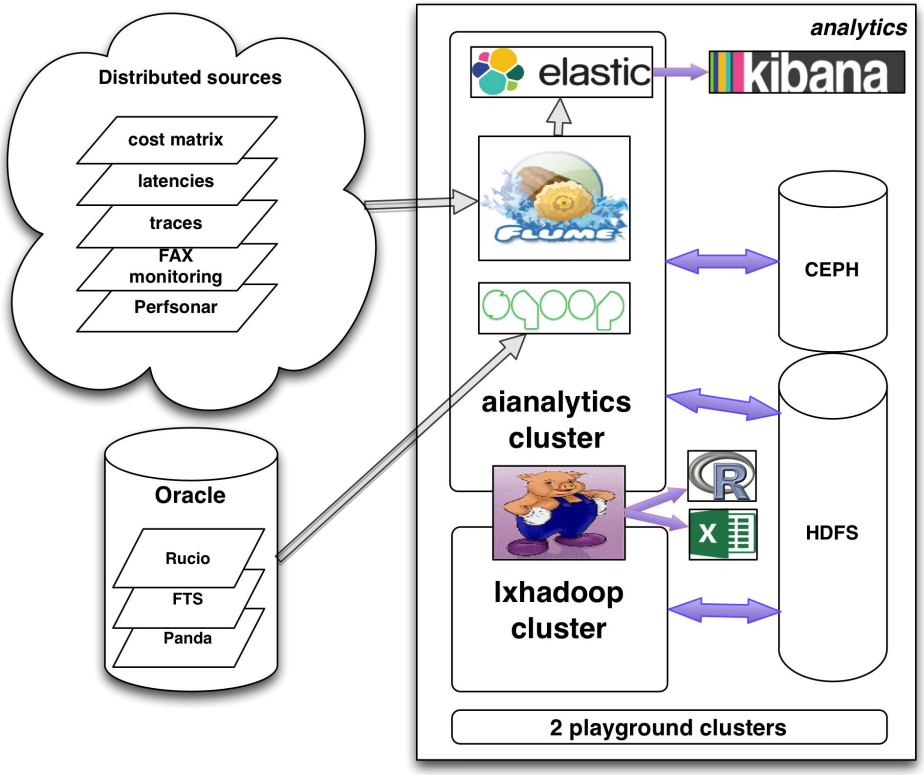

Central Flume Collector

- Listens for JSON messages (from multiple WLCG or ATLAS services)
- “events” multiplexed into different memory channels based on header content, and sent to log files and/or HDFS for analysis and/or ElasticSearch for indexing
- Currently collects from the CostMatrix service, traces, & FAX redirector summaries

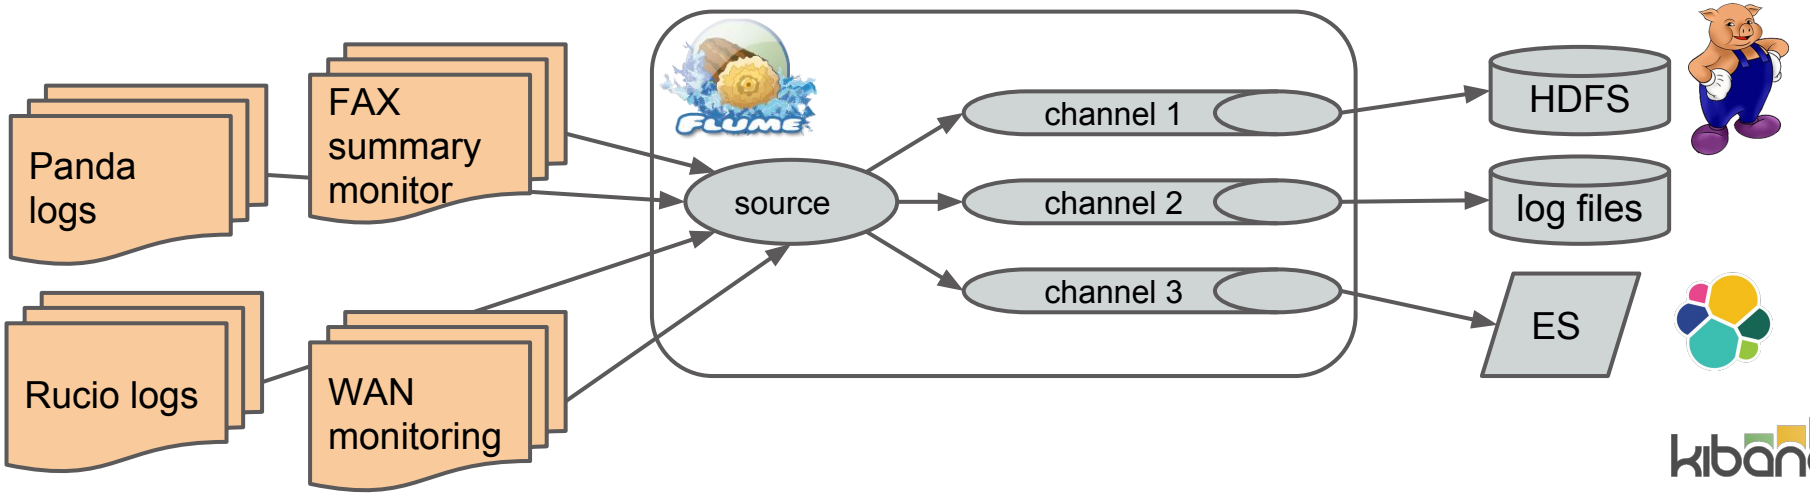

elastic @ CERN

1 x head node:

- 2 core VM
- 4 GB RAM
- 1 TB storage

4 x data node:

- 8 cores VM
- 16 GB RAM
- 1 TB storage CEPH io1

Authentication through SSO

Runs on ***voatlasanalytics*** cluster:

- No in-box storage.
- CEPH duplicates data on top of ES sharding.
- Not enough memory/core.

In need of new, stronger CERN managed cluster. We are not alone ...

elastic @ CERN

20+ other requests for the centrally managed ES at CERN.

As the biggest user we are collaborating with CERN IT in defining the new service.

Two biggest hurdles:

- Hardware - need nodes very different from the standard CERN nodes: less CPU, more memory, local disk - preferably fully SSD
- Authentication, role-based access control. Elastic's own package (Shield) to expensive. Looking for the open source version

Just now we are starting to test a brand new test cluster from borrowed machines:

- 3 master nodes (4GB, dual core, vm)
- 2 search nodes (8GB, for the time being)
- 23 data nodes, of two different types:
  - 4 virtual machines, with 80 GB, 8GB, 4 cores
  - 19 virtual machines, with SSD cache on top of the spinning disks, with 300GB of disk per node, 32GB, 16 core

elastic @ THE UNIVERSITY OF CHICAGO

2 x head nodes:

- 1 core VM
- 4 GB RAM
- 1 TB storage

4 x data nodes:

- 48 GB RAM
- 24 cores (HT)
- 3 TB storage

- Used for MWT2 analytics.
- Currently more than 1.2 billion docs in 3300 shards.

Contains:

- Condor history logs
- FAX costs
- FAX monitoring

Serious upgrade considered

authentication - username/password

elastic @ CloudLab

1 head node
5 data nodes

- 20 cores per node (2 CPUs)
- 256 GB RAM
- 2x10 Gb/s Ethernet card
- 40 Gb/s Infiniband
- 2x1 TB disk drives in each node.

Very easy to add more nodes.
Can change for nodes with
8x1 TB disks + 12x4TB

simple firewall protection

- CloudLab is flexible, scientific infrastructure for research on the future of cloud computing.
- Provides control and visibility all the way down to the bare metal.
- Provisioning an entire cloud inside of CloudLab takes only minutes.

Sites:

- Utah / HP
- Wisconsin / Cisco
- Clemson / Dell
- GENI

Deployed at Clemson University.
Great results and great support.

elastic @ CloudLab & THE UNIVERSITY OF CHICAGO

Clemson

1 indexing node
5 data nodes

- 20 cores per node (2 CPUs)
- 256 GB RAM
- 2x10 Gb/s Ethernet card
- 40 Gb/s Infiniband
- 2x1 TB disk drives in each node.

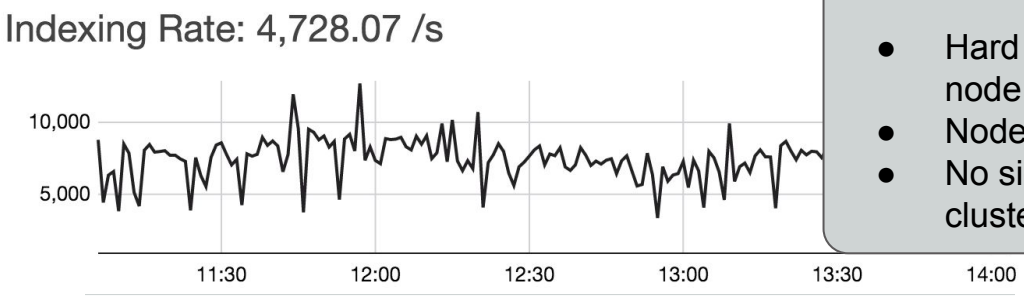

University of Chicago

3 master nodes (VMs 2 cores, 8 GB RAM)
1 client node (VM 2 cores, 16 GB RAM)
5 data nodes

- 24 cores per node (2 CPUs)
- 48 GB RAM
- 1 Gb/s Ethernet card
- 3x1 TB disk drives in each node.

Lessons learned:

- Hard to find VM configuration that works well for a data node
- Nodes with SSD's are order of magnitude faster
- No significant penalty in running geographically distributed cluster

Supporting Map Reduce & Search

Hadoop-based collections

1. PanDA Job Archive (2TB)
2. PanDA State change logs (0.5TB)
3. PanDA Logs (16 GB)
4. FAX cost matrix, traces (4GB)
5. Rucio (42 TB)
6. xAOD usage monitoring

Elastic Search Indices

1. PanDA Job Archive (Sqoop import)
2. PanDA State change logs (pig reprocessing + import)
3. FTS (AMQ + python)
4. Fax redirectors monitoring (Flume)
5. Rucio logs imports (streamed using Logstash)
6. xAOD monitoring
7. PanDA Logger (Flume)
8. Network Data (AMQ + python collectors)
9. ESNet counters and flow data
10. Frontier (starting this week)
11. Full FAX monitoring (starting in 2-3 weeks)

Analytics Infrastructure - lessons learned

- Need really good hardware - lot of RAM, SSD caching, large disks
- ES backups
  - while most of the data could be re-indexed, some exists only in ES.
  - small but important indices (.kibana) easily backed up every night at Amazon S3.
  - all indices backed up on remotely mounted CEPH storage at UC
- It is clear that Kibana accessible data have much more use. Index as much as possible of the hadoop data.
- Non-negligible learning curve (MR, pig, java, spark, jython) - need a lot of documentation, support, education.
- Very important to well document pig (specially UDFs) and spark codes for the different analysis needs.

Covered use cases

- Rucio
  - Error monitoring
  - Activity tracking
  - Tracking a file/dataset
- Usage of beyond-pledge resources
- Per cloud performance metrics
- Data formats popularity
- Network analytics service
- xAOD usage monitoring / analysis
- FAX monitoring jobs accessing data over WAN
- FAX redirectors monitoring
- Monitoring storage resources
- PanDA task duration analysis
- Tier-0 performance

Use cases still to be covered

- WAN performance analysis
- RTT jobs monitoring
- Geant production log analysis
- Addition to bigpanda (PanDA web frontend)- replace the slowest Oracle searches
- Frontier servers - tracking slow queries, rogue codes/hosts.

This list grows faster and faster...

Immediate questions to answer

- is derivation production successful?

- do people run private filtering on AODs?

- how much private production the users run on their own?

- do majority of users use DAODs?

- does a small subset of users consume most of the cpus for analysis which is not in our computing model?

- how many cpus are needed to make 95% of users happy - i.e. to run analysis on DAODs?

- how many users would like to run parallel or complex jobs (eg, high memory, task chains...)?

- is the current system (PanDA) good enough to give a high throughput to DAOD analysis and prevent the users running non-standard jobs to take over the resources?

PandaJobs efficiencies

Use to monitor jobs for batches of inefficient jobs.

Investigate reasons for low CPU eff.

Compare efficiency of the same task on other queues.

After drill-down switch to “Discover” mode to go down to a single job level.

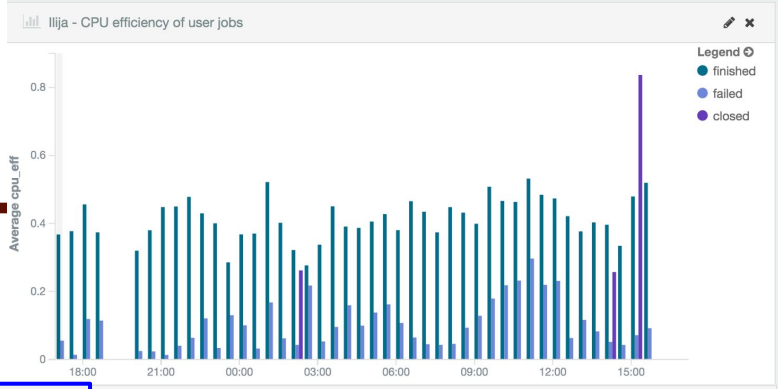

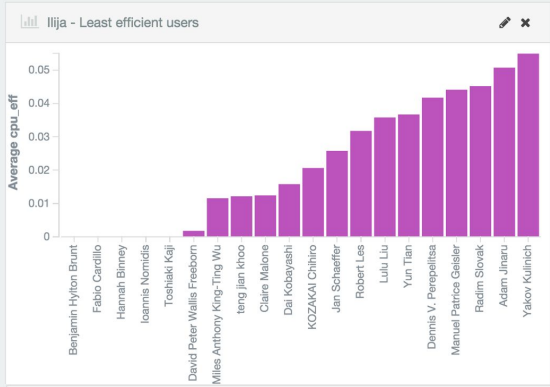

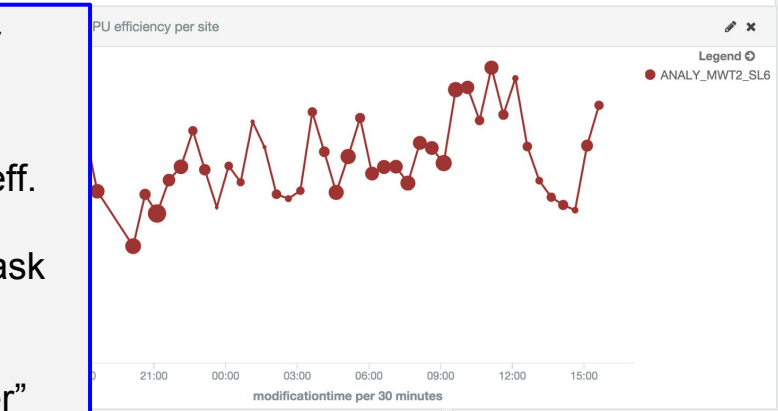

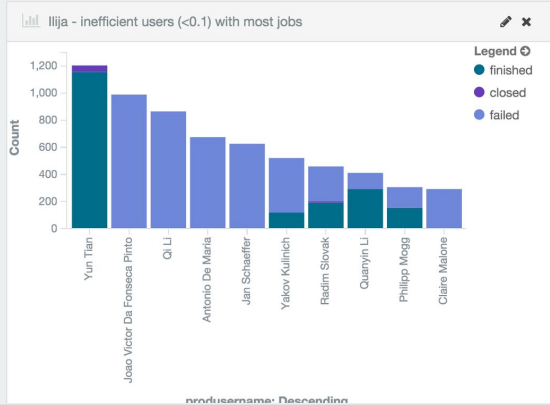

Ilija - CPU efficiency characheristics

filters	Count	Average Average VMEM	Average Average RSS	Average avgswap	Average attemptnr	Average currentpriority	Average wall_time [h]	Average ninputdatafiles	Average inputfilebytes
cpu_eff:<0.1	11,108	418.63MB	124.91MB	19.52KB	1.875	870.812	1.407	6.158	7.32GB
cpu_eff:>0.1	19,651	1,022.5MB	391.43MB	129.52KB	1.262	685.655	1.826	3.275	4.85GB

A network analytics service for ATLAS

We use ES as a network analytics service.

Data collection and warehousing

- Throughput, latencies, and packet loss data come from OSG network datastore, collected from AMQ, enriched with AGIS mapping info.
- FAXcost - measures rate of single file transfer between 65 FAX endpoints and 20 largest ANALY queues.
- FTS - rate of transfer of each single file, sizes, waiting times, reasons for transfer
- ESNet data (LHCONE)
- Stored in Hadoop and indexed into ES
- Developing prediction of future network performance (see the following talk)

Serving the data:

- very simple REST and python interface
- ES should be able to deliver searches in <100ms @ 100Hz

File formats popularity

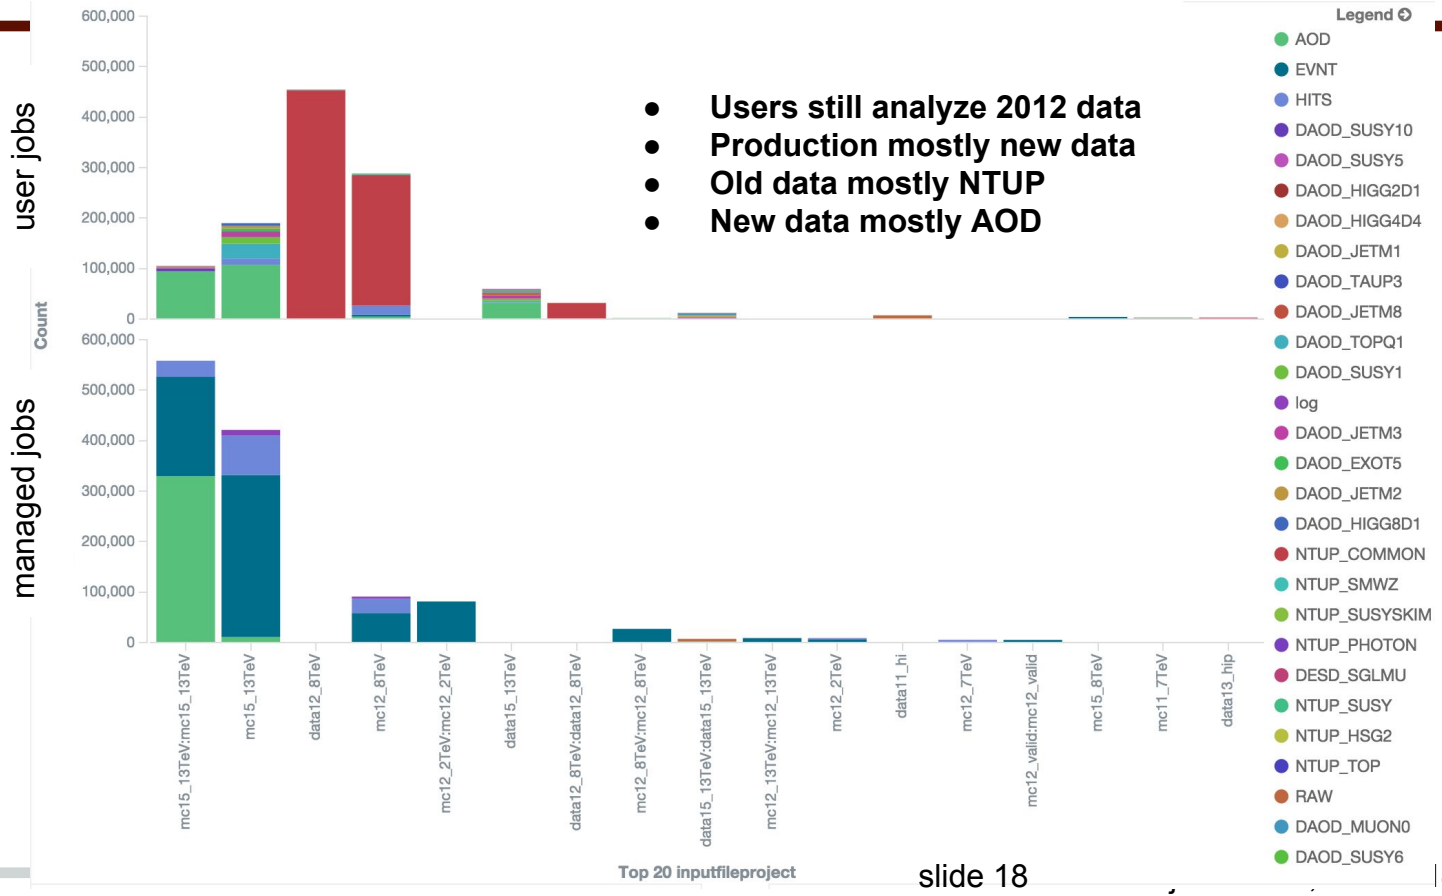

- Users still analyze 2012 data
- Production mostly new data
- Old data mostly NTUP
- New data mostly AOD

PandaJobs

Input data

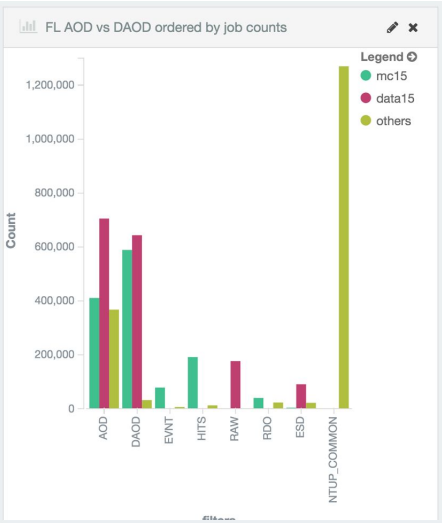

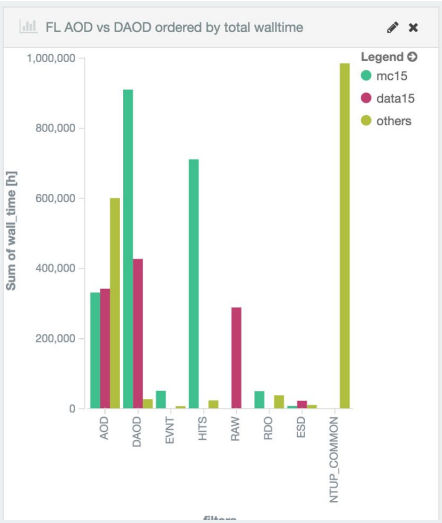

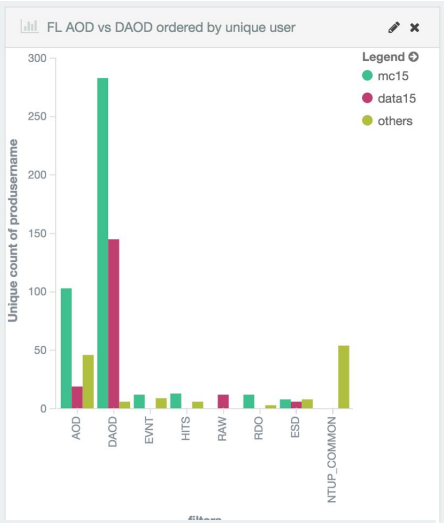

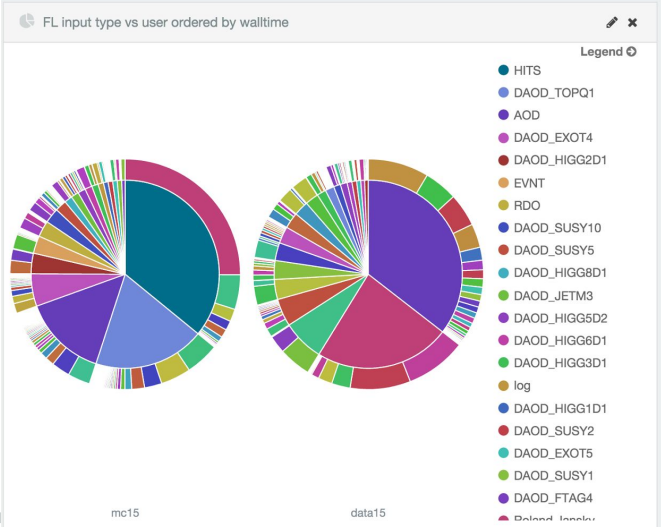

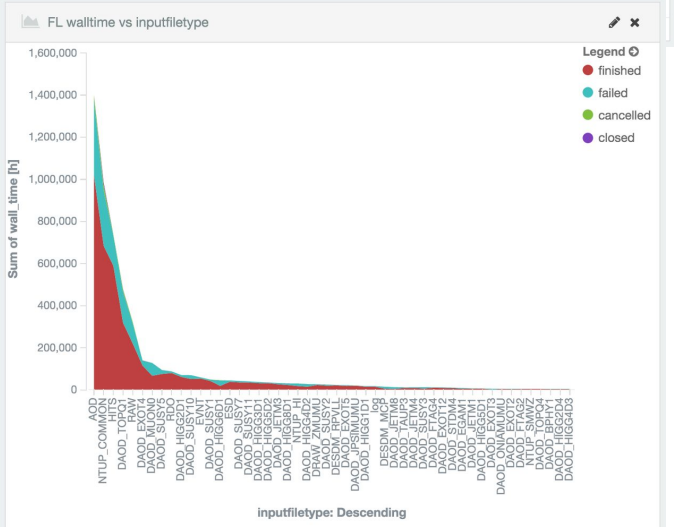

Rucio logging

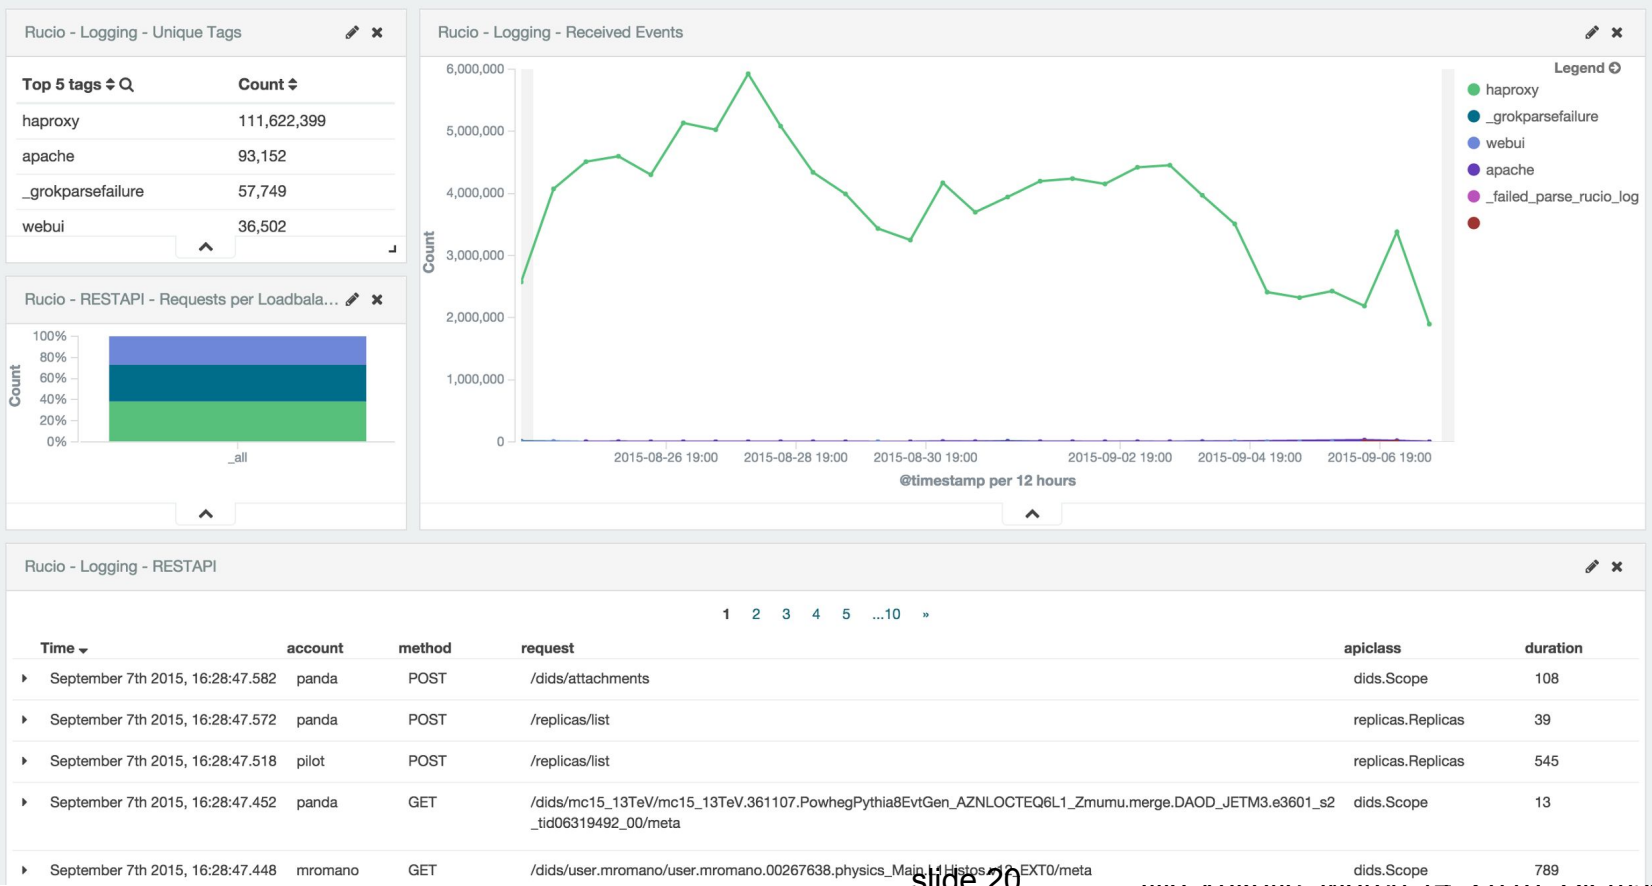

Rucio

Error Tracking

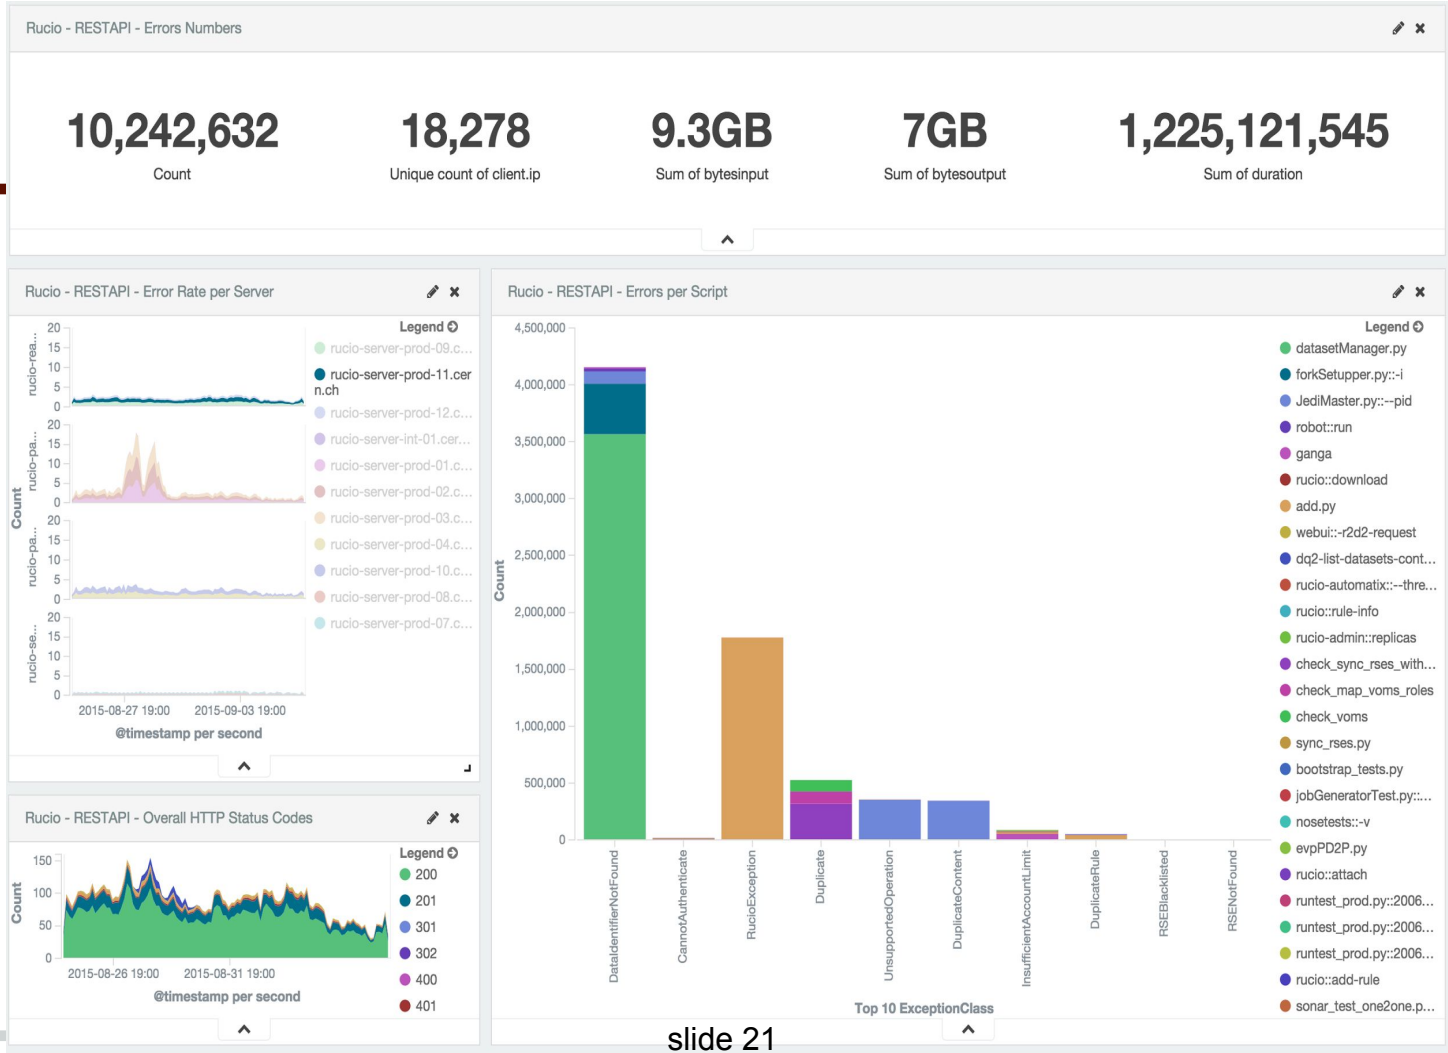

Rucio

Account Activity

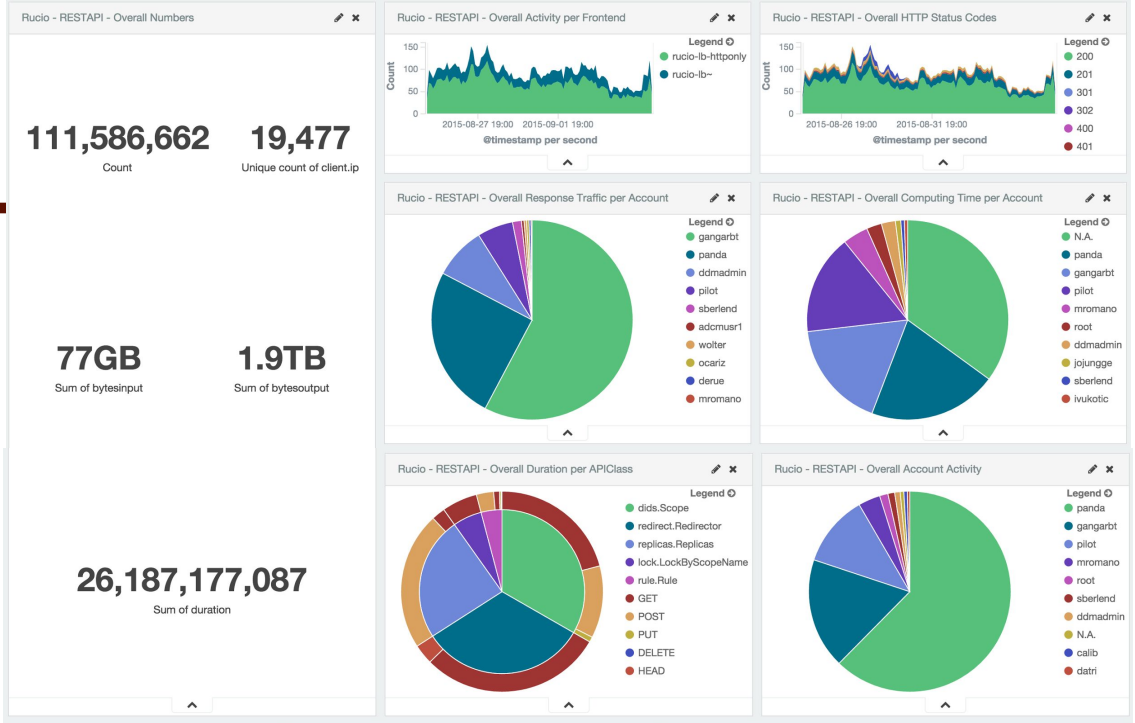

FAX

overflow monitoring

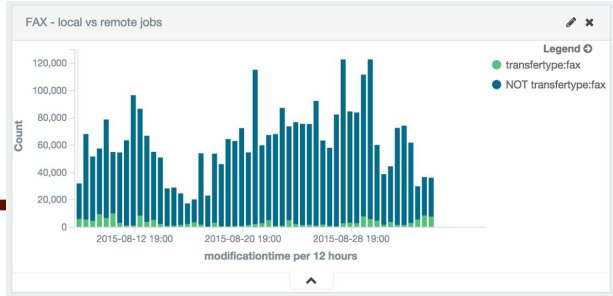

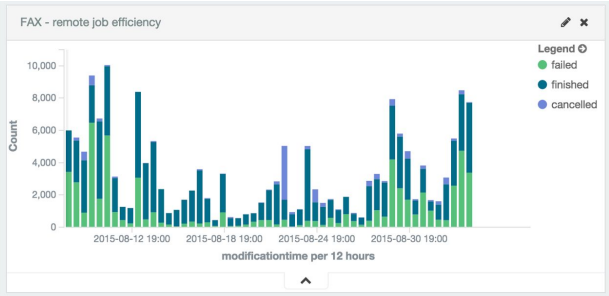

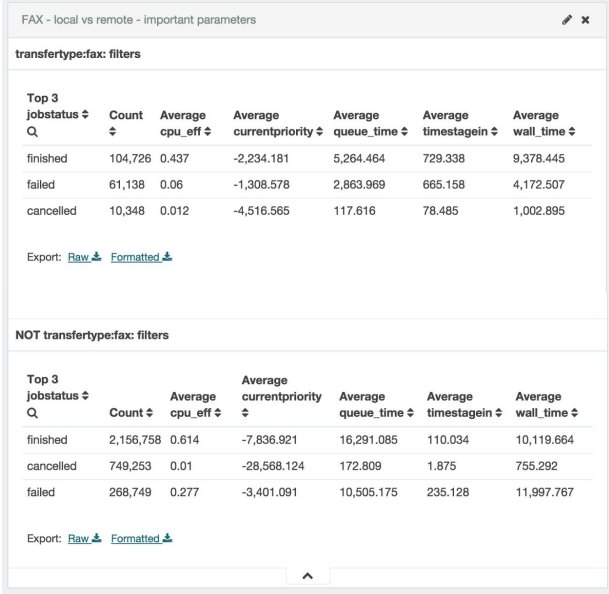

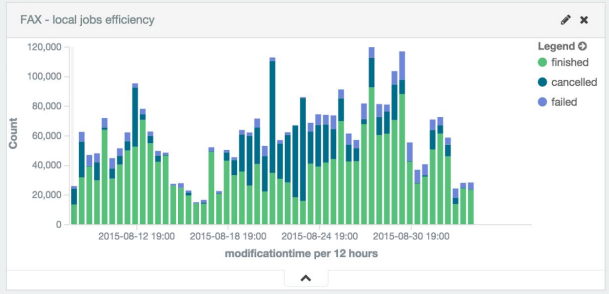

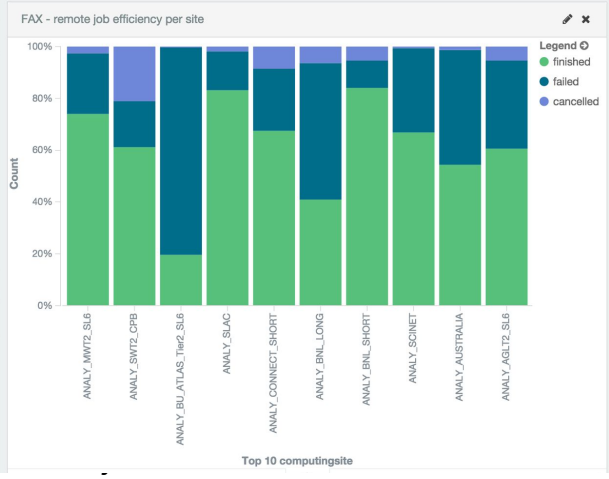

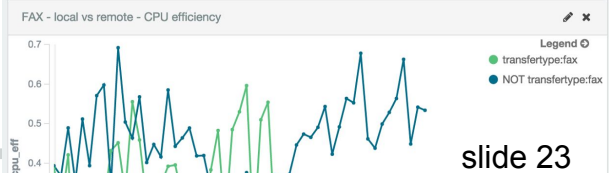

Analytics control center

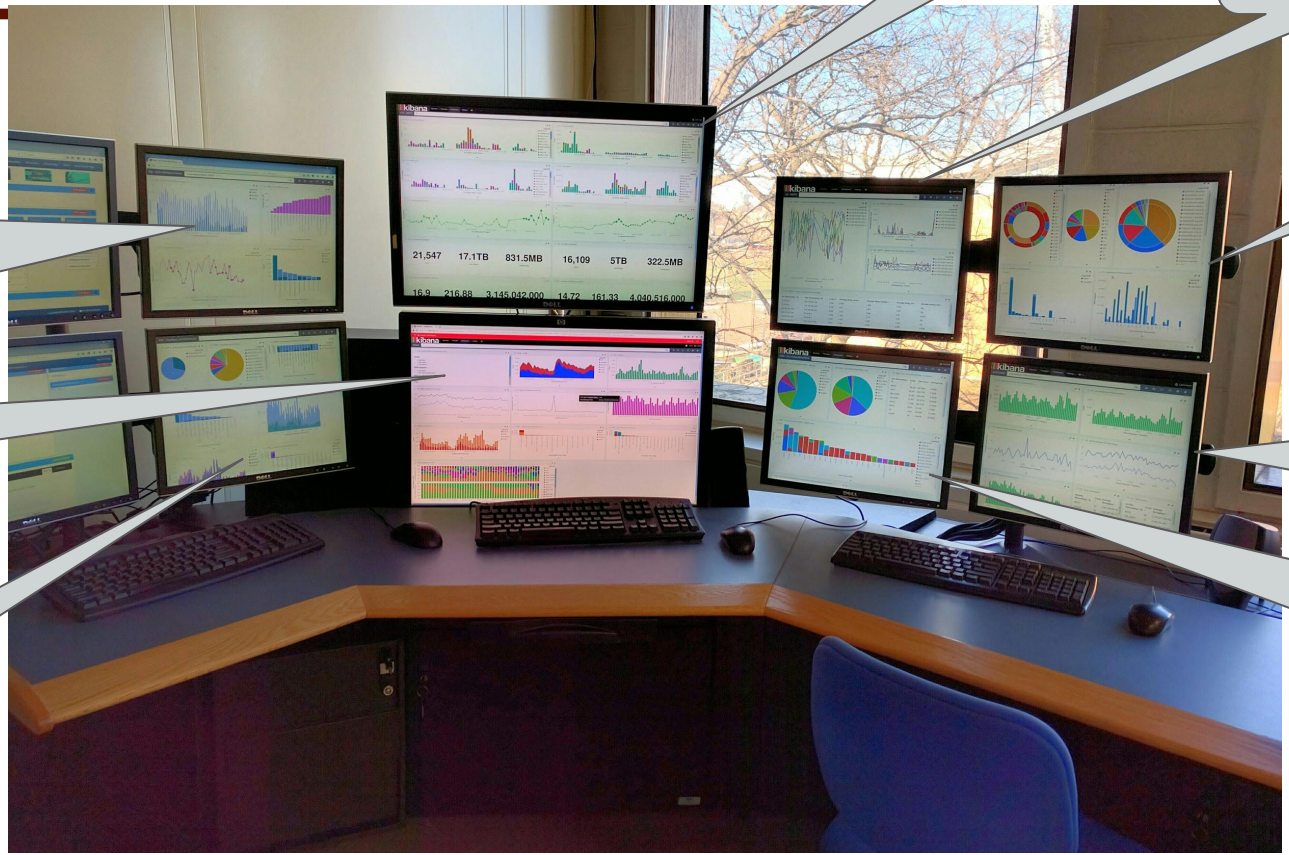

Next step - real analysis

While extremely useful ES+kibana are not intended for the “real” analysis

There are a lot of new frameworks that could be used: pig, Spark, Jupyter, R, SciPy, …

There are a number of young people that want to master the new tools, explore.

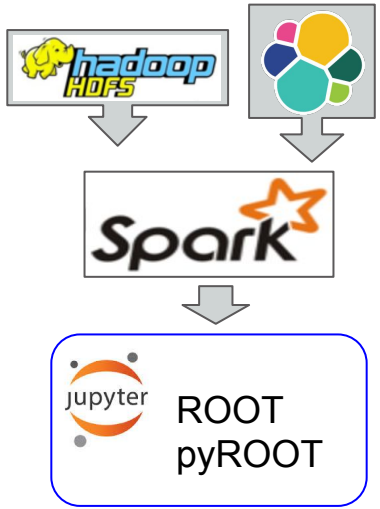

Example data flows

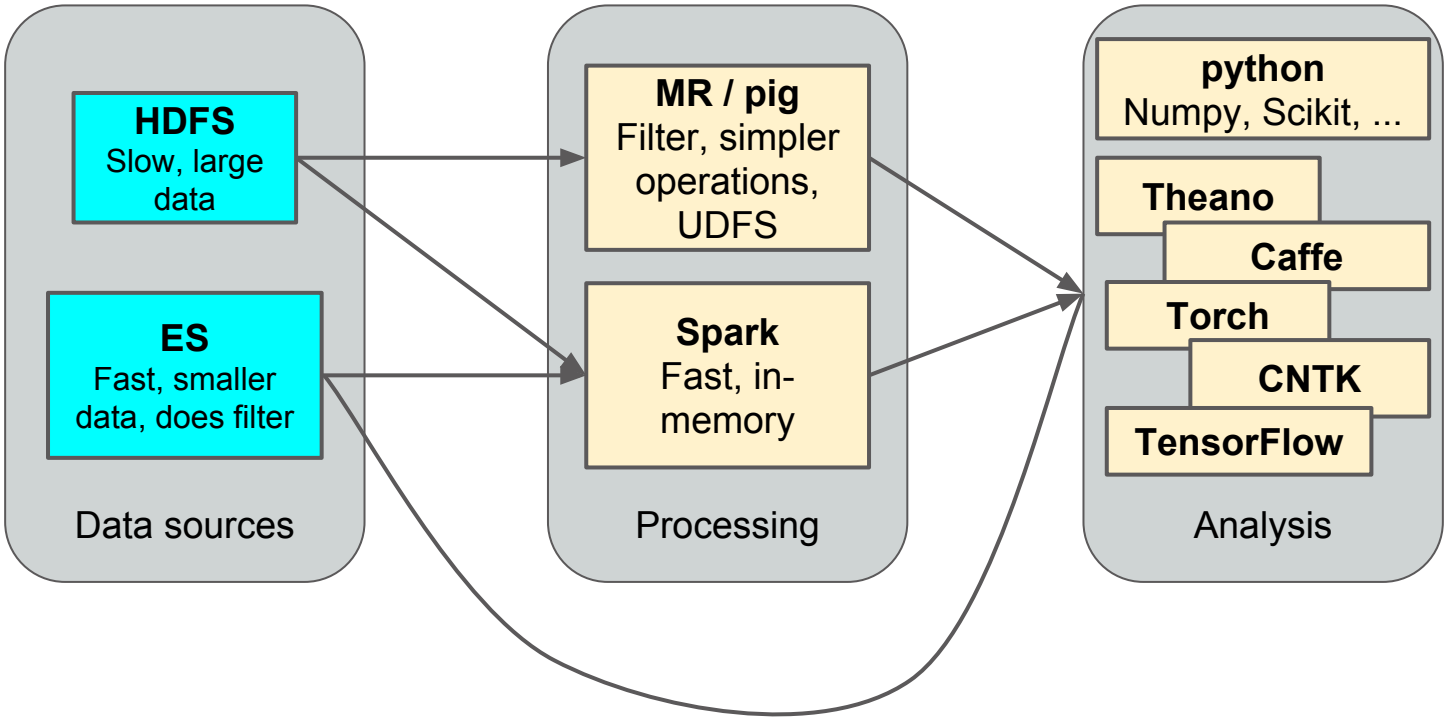

Example data flow

From Mario's S&C talk:
Machine learning for network awareness

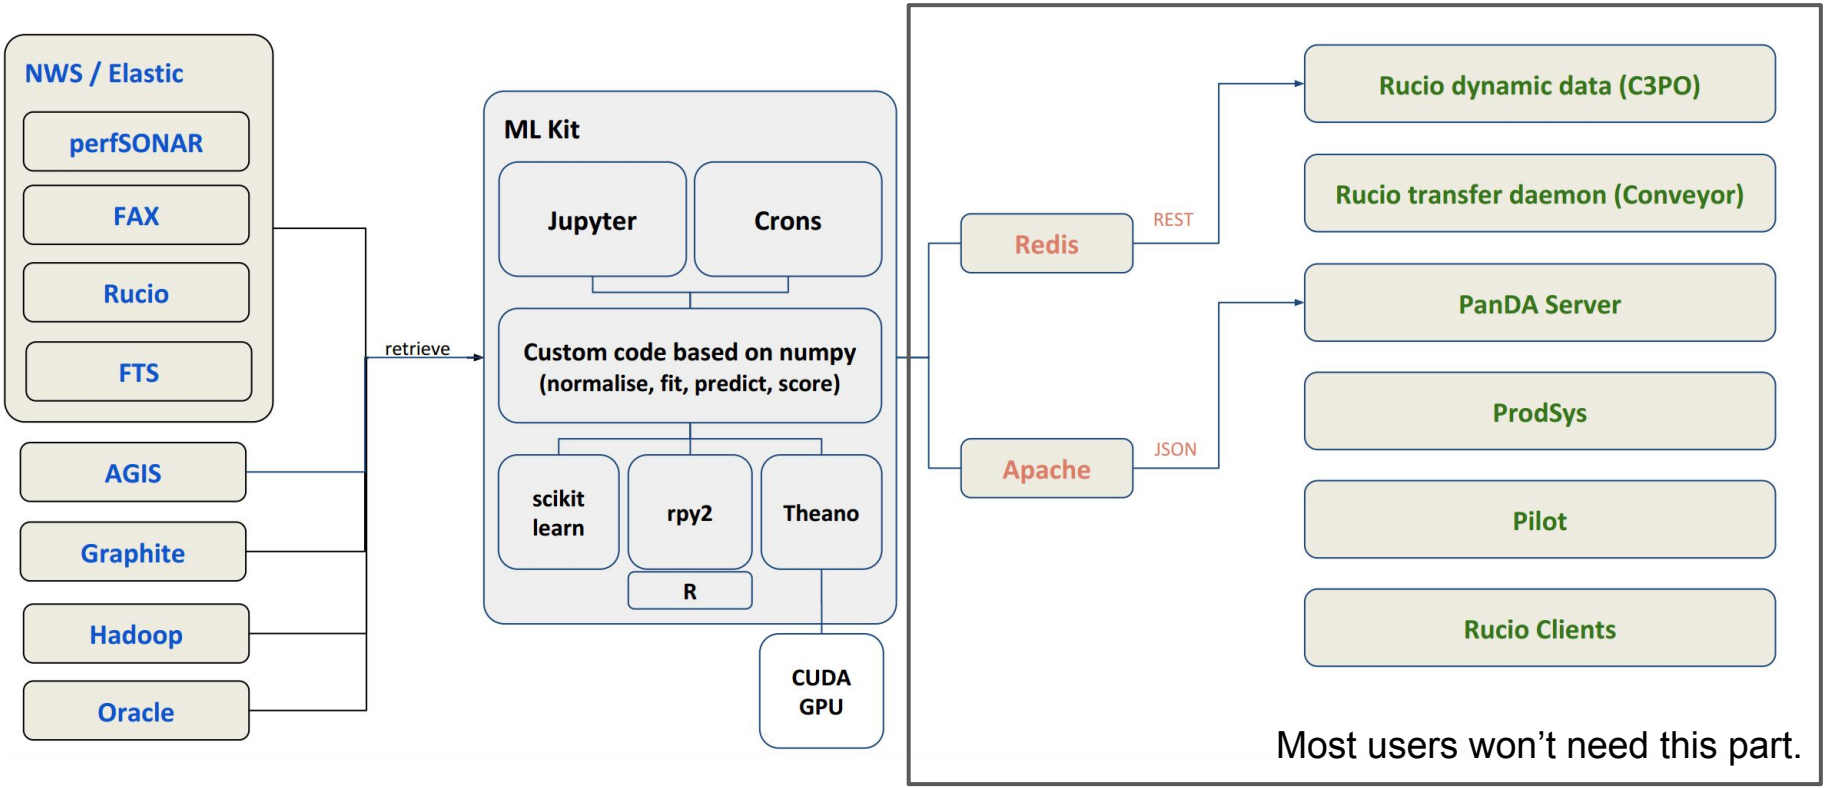

Conclusions

The new stack of Big Data tools (Hadoop, Flume, Sqoop, Logstash, pig, Spark, ES, Kibana) provide great platform for ATLAS ADC analytics tasks:

- horizontally scalable
- performant
- simple to develop for
- easy to use for an analyzer
- fast to make custom dashboards, GUIs
- can be used as a full search service

For future reference

Analytics TWiki page

egroup: atlas-adc-data-analytics

Tutorials:

- part 1: Platform description, data indexed, importing, simple data analysis
- part 2: Sqoop, Flume, Pig, ES indexing
- part 3: Using Kibana

Spares

Data sources

PanDA - a data-driven workload management system for production and distributed analysis processing

Rucio - a Distributed Data Management system used to manage accounts, files, datasets, and distributed storage systems.

FAX - Federated ATLAS storage system using XRootD protocol. Provides a global namespace, direct access to data from anywhere.

PerfSONAR - a widely-deployed test and measurement infrastructure that is used by science networks and facilities around the world to monitor and ensure network performance.

FTS - File Transfer Service - the lowest-level data movement service doing point-to-point file transfers.

xAOD - primary analysis data product.

xAOD accesses analysis

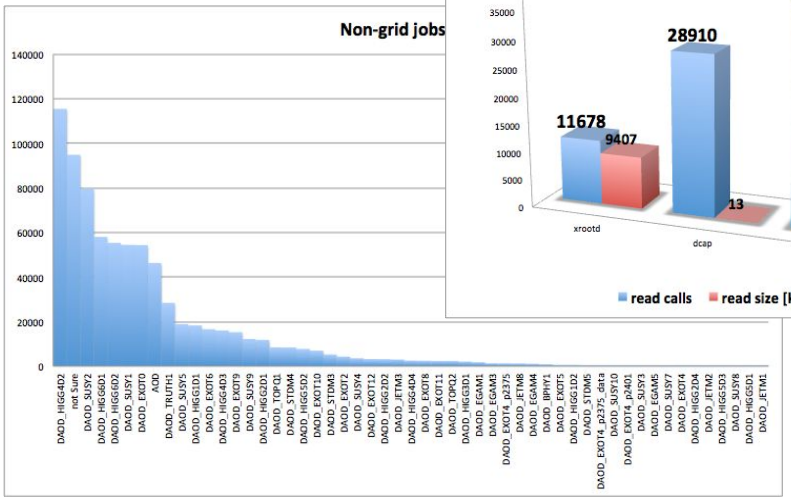

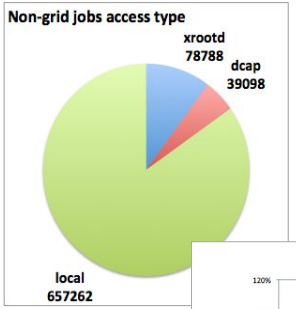

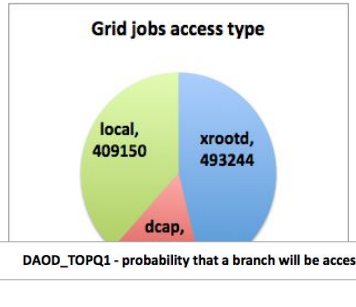

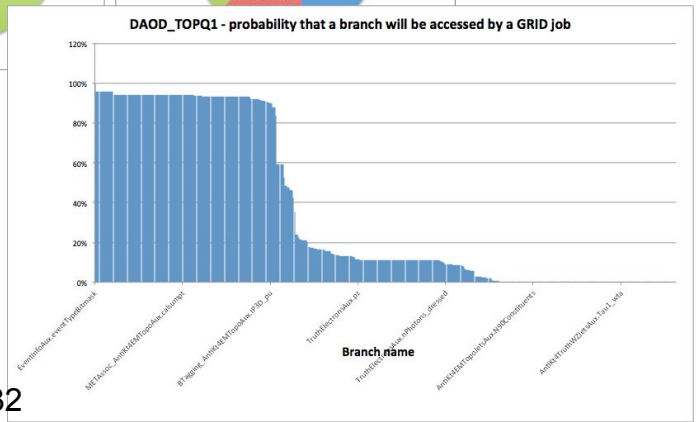

Rucio Account Usage

Rucio - RESTAPI - Account Usage - Activity over Time

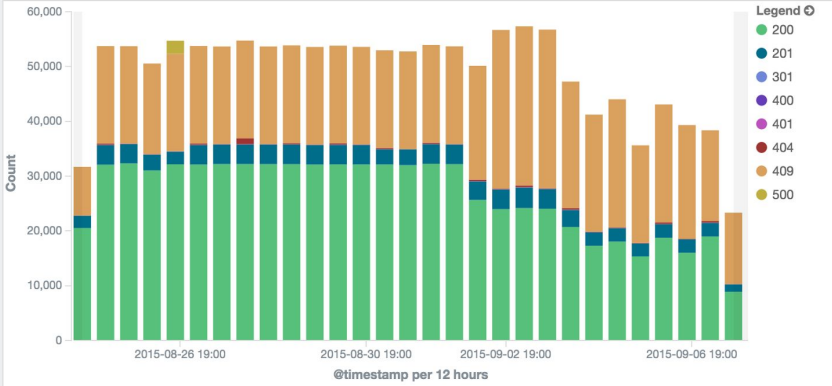

Rucio - RESTAPI - Overall Numbers

1,421,600	9
Count	Unique count of client.ip
669.8MB	2.3GB
Sum of bytesinput	Sum of bytesoutput

573,202,003

Sum of duration

Rucio - RESTAPI - Account Usage - Resources

Top 10 request	Top 0 scriptid	Count	Sum of duration	Sum of bytesoutput	Sum of bytesinput
/accounts/pilot/identities	check_map_voms_roles	24,145	1,306,217	15.4MB	11.2MB
/accounts/phys-higgs/identities	check_map_voms_roles	15,082	694,607	9.5MB	7MB
/accounts/perf-muons/identities	check_map_voms_roles	11,786	543,156	7.6MB	5.6MB
/replicas/list	monitor_client::-s	6,229	3,316,796	34MB	35.9MB
/replicas/list	rucio::list-file-replicas	6	932	4.8KB	2.5KB
/replicas/list	nosetests::-v	2	184	10.6KB	808B
/replicas/list	rucio::list-dataset-replicas	1	133	25.7KB	480B
/accounts/phys-hi/identities	check_map_voms_roles	5,587	292,116	3.6MB	2.6MB
/accounts/phys-gener/identities	check_map_voms_roles	5,285	285,139	3.3MB	2.5MB
/accounts/perf-flavtag/identities	check_map_voms_roles	4,682	244,070	3MB	2.2MB

Export: Raw Formatted

Rucio - RESTAPI - Time spent per Script per API Class

FAX redirector network monitoring