
ADC Analytics Platform at MWT2



David Jordan, Ilija Vukotic, Judith Stephen, Lincoln Bryant,
Farnaz Golnaraghi, Fengping Hu

US ATLAS SLAC Face-to-Face Computing Facilities Meeting
2022.12.1



ENRICO FERMI
INSTITUTE



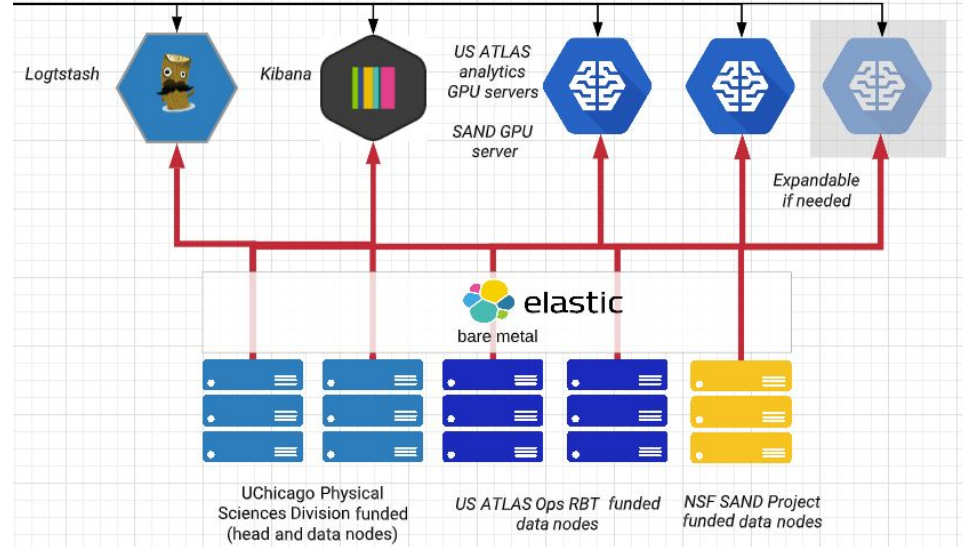
Cluster Hardware

- 18 Physical Machines

- 3 Head nodes
- 14 Data Nodes
- 1 Web Server/Ingress (Kibana)
 - Also runs metricbeat

- All machines use SSDs

- 4-8 SATA SSDs on data nodes
- Newest data nodes use m.2s for root disk. Older use SATA SSDs
- 74TB Capacity with ~46TB used. 28TB still free despite constant usage.



Data Collection

- Logstash

- 9 collectors on the River k8s cluster
 - MWT2 and AGLT2 dCache storage data
 - Perfsonar HTTP archiver data
 - UC AF monitoring/accounting
 - Panda memory/CPU reporting
 - ServiceX logs
 - Squid monitoring
 - Varnish
 - StashCache
 - Xcache

- Directly dumping into ES

- Usually python scripts
- ADC data (jobs, task,...) – running on CERN Analytics k8s cluster
- All details on data accesses –from AnalysisBase.
- Rucio data (traces, events,...) – runs on CERN Spark cluster
- ATLAS Alarm&Alert Service – backend for the service
- GeoModel – tracks revision history. Runs in CI/CD at CERN GitLab.
- VP – from dedicated VP k8s cluster at CERN



What has been collected and indexed?

- All details on every Panda Task(100+ variables) and Job (120+ variables)
- Any opened xAOD files as well as which branches and how many events
- Pilot reports of memory, cpu usage, etc. of each job, sampled at regular intervals
- All data in rucio and any changes to that data.
- Conditions queries made through squids, what server served it, how fast, was it cached, etc.
- Of the big sources, the only thing we don't have are full, real-time logs (these will be CERN only).



Data Center Relocation

- Wanted to move the cluster with minimal downtime
- Borrowed equipment from the University to help transition
 - 3 head nodes, 8 data nodes, new Kibana/ingress node
- Two phases of data migrations
- Spanned the cluster across 2 subnets



Relocation Challenges

- Moving Kibana instance
 - At first, tried creating a brand new one.
 - Created headaches trying to determine how to migrate to the new one
 - Decided to make a hard cut to the new one
 - Same hostname, configuration, certs, etc.
- Internal ES certificate issues
 - Entire cluster broke after moving one of the head nodes
 - Internal cert CA was tied to head node IP address
 - Cluster ran fine with machine turned off
 - Created new certs and CA for the whole cluster



Upgrading to 8.x

- Going from 7.14 to 8.x
 - Needed to upgrade to the latest 7.x version (7.17) first
 - All indices needed to be indexed for ES 7.x
 - All software is recommended to upgrade together
 - Elasticsearch, Kibana, Metricbeat
 - Each Elasticsearch major version has its own repo
- Multiple day process
 - Upgrade to 7.17
 - Reindex indices for 7.x (7.17)
 - Upgrade to 8.3



Upgrade Challenges

- Documentation
 - Breaking changes list was not exhaustive
 - Had to fix and restart until ES did not fail
 - Must do a full cluster restart for major version updates
- Old security-6 index
 - ES 8 complained about an outdated index
 - Elasticsearch service was running, but cluster could not initialize
 - Made in ES 6, but has since been reindexed to 7.x
 - Backported to ES 7, created a blank security-7 index and migrated
 - Deleted the security-6 index



Upgrade Challenges cont.

- Kibana.index removed in 8
 - We used this to use a custom dashboards and spaces
 - Needed to migrate our .kibana-dev index to .kibana
- Users now needed to be assigned roles in Kibana
- A few harmless errors
 - AWS “access denied”
 - Java ScriptEngineManager
 - Bugfix version mismatch
 - Testing if head nodes got 8.3.1.
 - Data nodes got 8.3.2



Closing Notes

- There is a puppet module
 - <https://forge.puppet.com/modules/puppet/elasticsearch>
 - <https://forge.puppet.com/modules/elastic/elasticsearch> – older
- Elastic docs are generally very verbose and good
 - <https://www.elastic.co/guide/index.html#viewall>
- MWT2 does full cluster restart upgrades
 - Much more admin friendly
 - Shut off order: Kibana -> data nodes -> head nodes
- Looking to upgrade to the latest version in December



Extras

analytics.mwt2.org

The screenshot shows a web browser window with the following elements:

- Browser Tab:** "AF ATLAS - Home"
- Address Bar:** "analytics.mwt2.org"
- Navigation Menu:** "Admin", "Services", "Projects", "Email", "Account" (all with dropdown arrows)
- Header Logo:** ATLAS EXPERIMENT Distributed Computing Analytics
- Main Content:**
 - ATLAS EXPERIMENT Distributed Computing Analytics logo
 - ATLAS Distributed Computing Analytics Platform
 - Frontier
- Footer Logos:** University of Edinburgh, Enrico Fermi Institute, and iris hep SSL

ATLAS Alarm And Alert System – AAAS

A web frontend and REST API

- Simplifies and standardizes alarm management
- Has integrated heartbeats monitoring with automatic alarm creation
- Incommon user signup, easy to manage subscriptions, preferences.



Home Docs Alarms Heartbeats Subscriptions

Profile

ATLAS Alarm And Alert

This platform provides a simple way to store generated alarms, handles users subscriptions and generates alerts.

Components

Database

Elasticsearch

REST API and Web frontend

Node.js + express + pug

Deployment

Docker, K8s, Helm (soon)

Authentication



Globus InCommon

Authorization

API key

Mail

Mailgun



ATLAS Alarm And Alert System – AAAS

Already a lot of alarms.

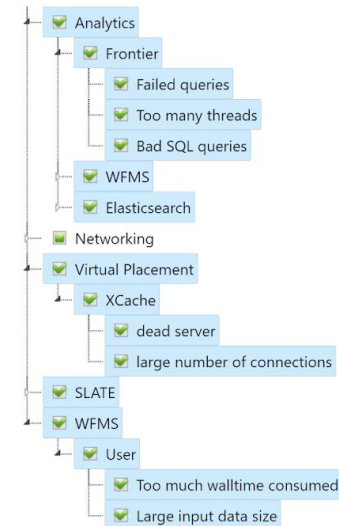
Alerts are grouped and delivered in used defined intervals.

Filtering on Tags reduces volume of emails.


Still some developments to do:

- Annotations
- Alarm acknowledgements
- Alarm redirection

Alarms



Heartbeats



Current Subscriptions

Category	Subcategory	Event	Tags
Analytics	WFMS	indexing	*
Analytics	Frontier	Failed queries	*
Analytics	Frontier	Too many threads	*
Analytics	Frontier	Bad SQL queries	*
Analytics	Elasticsearch	status	red
Networking	Perfsonar	high packet loss	MWT2
Networking	Perfsonar	indexing	*
Networking	Perfsonar	complete packet loss	MWT2
Networking	Perfsonar	Firewall issue	*
Networking	Perfsonar	large clock correction	*
Networking	Perfsonar	bad owd measurements	*
SLATE	Squid	server down	*

GeoModel monitoring


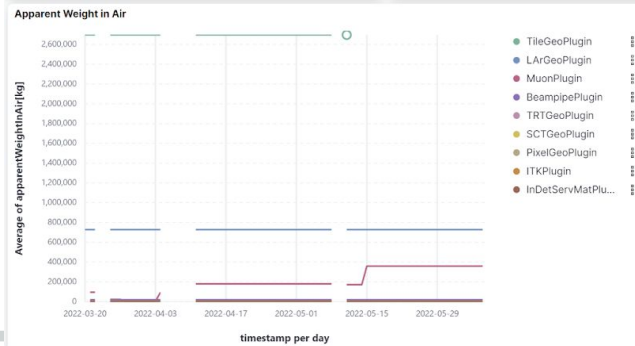
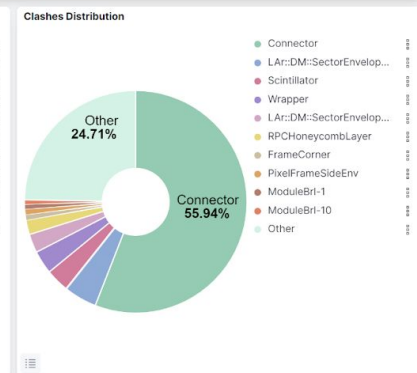
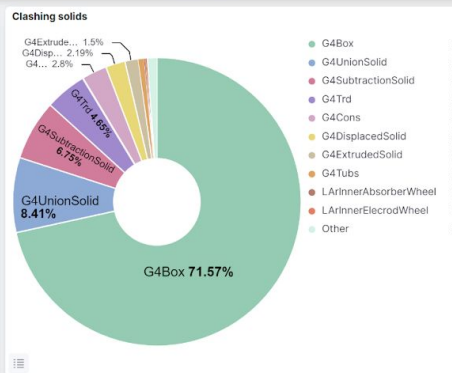
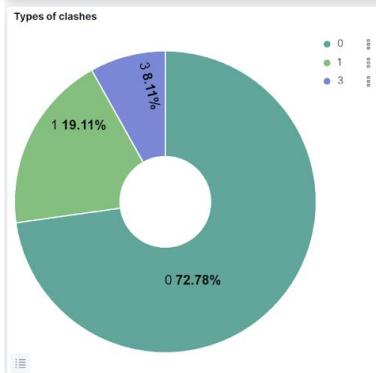

ATLAS GeoModel developments are handled in one GitLab at CERN.

Once per day, and on each merge, model gets rebuilt and series of tests are run.

All results get sent do UC ES.

Visualizations are integrated back into GeoModel web pages.

Collaboration with Riccardo Maria Bianchi & Marilena Bandieramonte



Frontier

Similar to all the analytics projects, if you leave it for a year, don't count on it working.

Software packages change, tokens expire, k8s clusters get upgraded,...

Frontier was ready for a full refresh. Moved from CERN GitLab to GitHub/DockerHub, redone all GitOps, deployments moved to FluxCD.

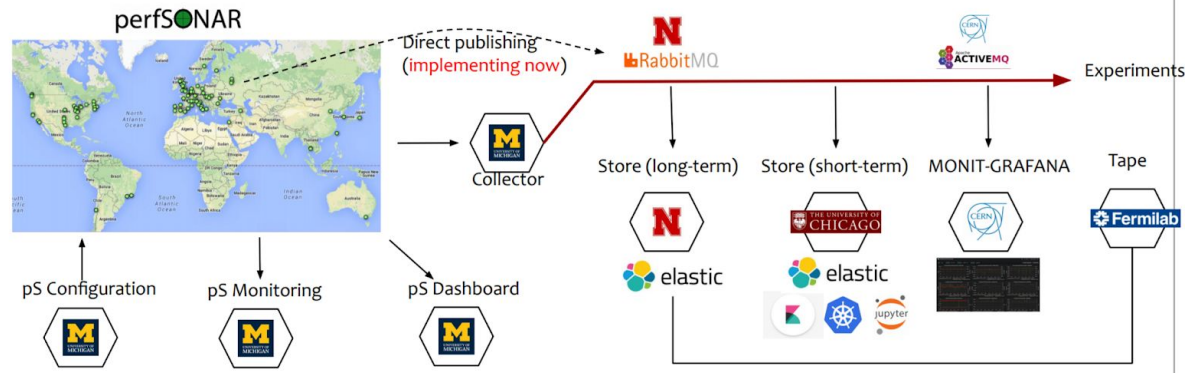
Julio who was running Logstash at all the Frontiers left and was replaced by Jeyhun.

The screenshot shows the 'FRONTIER ANALYTICS' web interface. On the left is a blue sidebar with navigation options: 'Discover', 'Visualize', and 'Caching Efficiency'. The main content area is split into two panels. The left panel, titled 'ElasticSearch Filters', contains search filters for 'Task ID', 'Search by time range' (with date pickers), 'Choose Queries type' (set to 'All Queries'), and 'The parquet File name' (set to 'name.parquet'). A large blue button labeled 'Extract Frontier logs' is at the bottom of this panel. The right panel, titled 'Parquet Files List', shows a table of files with columns for 'Parquet File Name', 'size (MB)', 'Created', 'Delete', and 'Downl'. Below is a table with 6 rows of data.

Parquet File Name	size (MB)	Created	Delete	Downl
18628864_pPboverLay	59.55	Tue Feb 1 01:46:45 2022	[Delete]	[Download]
23470217_pPboverLay_new	56.61	Sat Jan 29 00:53:07 2022	[Delete]	[Download]
27726954_ilija	53.07	Fri Jan 28 23:55:51 2022	[Delete]	[Download]
28193236	25.10	Wed Feb 23 10:06:33 2022	[Delete]	[Download]
28193236_2nd	25.10	Wed Feb 23 13:09:43 2022	[Delete]	[Download]
28605719	40.17	Mon Apr 11	[Delete]	[Download]

Network monitoring

- We collect both PerfSonar and ES-Net data.
- PerfSonar pipeline was too long and complex and with 100 M documents per week, we need it very optimized.



- The new PerfSonar data store is Open distro Elasticsearch.
- It comes with a HTTP publisher.
- In process of testing logstash collector for push sites. (me, Shawn McKee)
- Developing pull collector for pS sites we don't control (Thomas Shearer of UM).
- Networking alarms and alerts constantly optimized (Petya Vasileva, Marian Babik).