

Operational Intelligence



Panos Paparrigopoulos on behalf of the Operational Intelligence initiative



Our Mission

- A cross-experiment effort aiming to streamline computing operations:
 - Improve resource utilization by reducing the time needed to address operational issues
 - Minimize human effort for repetitive tasks by increasing the level of automation
 - Build a community of technical experts: critical mass to have impact on concrete and common issues while setting up sustainable tools.
- Our mission:
 - Identify common projects
 - leverage common tools/infrastructure
 - Collaborate, share expertise, tools & approaches
 - Across experiments
 - Across teams (operations, monitoring, developers)

Operations Today

human
machine



Chat,
meetings,
emails,
jira

ATLAS/CMS: 100+ people involved
in Computing Operations
(50+ FTEs/experiment)!
In 1 year:
> 1k GGUS tickets for ATLAS, > 2k for CMS

Visualization / Monitoring



Processing




logging



Data sources

Systems,
components
services

Data Providers



Actions



Alerts



Operations Tomorrow

human
machine



Frontend: aggregated views, suggestions, collects feedback

Visualization / Monitoring

Backend: Fetches, stores, filters, and analyses information about alerts, issues and solutions



Actions/alerts

Actions

logging



Data sources

Systems, components services

Data Providers



Ongoing Efforts

What we are doing to succeed:

- Develop tools to automate computing operations exploiting state-of-the-art technology and tools
- Run a technical forum, experiment-agnostic to:
 - bring people together
 - discuss ideas, brainstorm together, share experience and code

We have identified the areas where shared development can occur:

- Sites
- Workflow Management
- Data Management

And we are also trying to provide some shared infrastructure:

- A common k8s cluster for all those services to be deployed.
- A framework which can be used to develop new tools





Sites Optimisation

- There have been some efforts by sites that are developing tools very much in line with the OpInt project.
- There is some very interesting work on cloud anomaly detection going on at CERN.
- The idea is to automatically detect anomalies in the Openstack infrastructure and alert users.
- An enchantment on Grafana annotations has been developed in order to automatically tag such anomalies at Grafana dashboards, helping people to correlate activity in graphs with anomalies.
- The above are all efforts of individual sites. One of our goals is to try and push for more collaborative projects, across multiple sites.



Sites Optimisation

- We are also keeping an eye into what big companies from the industry do to automate their computing centers and reduce operational costs and environmental impact.
- Of course in a diversified environment like WLCG these holistic strategies cannot always apply.
- The past years we have moved into a more unified processing pipeline in our sites, something which creates possibilities for collaborative efforts.

Workflow Management

- Workflow management is a complex work that could benefit by optimisation and smart tools.

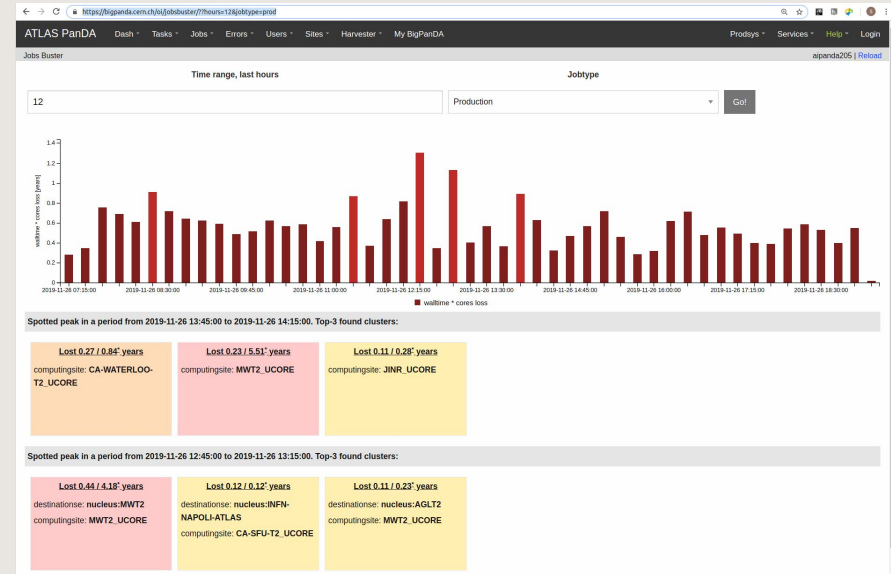
- ATLAS** has developed a tool called Jobs Buster which tries to spot operational problems errors in ATLAS jobs and display results in the BigPanDA monitoring page.

<http://cern.ch/go/8qwC>

- Using NLP Jobs Buster tries to find the common denominator between failed jobs

- CMS** has developed a similar system. The CMS Operator console.

<http://cern.ch/go/z76x>



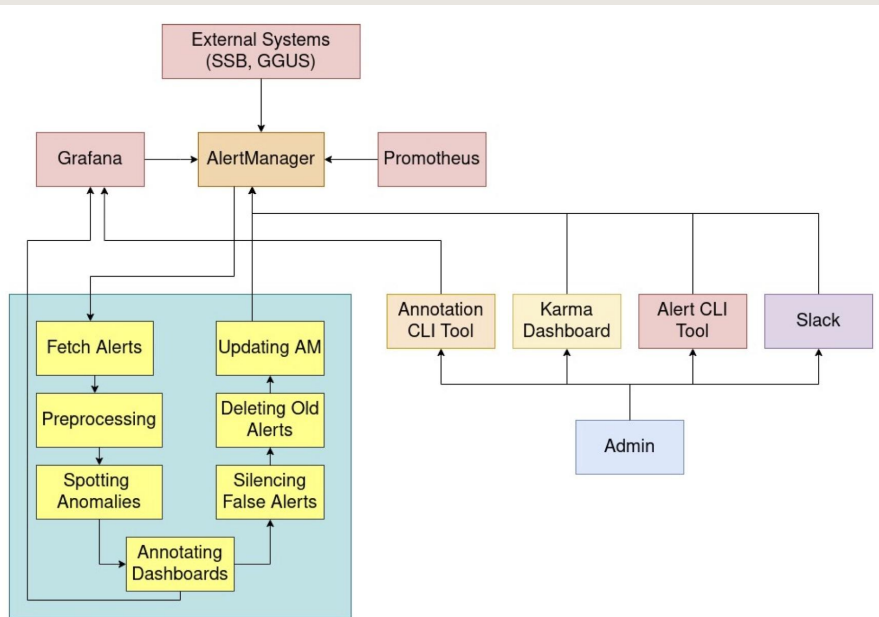


Data Management

- There are currently two efforts in **CMS** and **ATLAS** to classify FTS error logs and try to provide useful information to DDM shifters.
<http://cern.ch/go/ZjMl>
- That's a very good example of collaboration between the two experiments, by sharing ideas and code. FTS is a commonly used tool with common logs and require almost identical operational effort on both experiments.
- Currently the CMS version has moved to production. FTS logs are being analysed with NLP algorithms and the results are being displayed in a Grafana dashboard.
<http://cern.ch/go/BX6l>
- The idea is that shifters can see the evolution of those clusters and give us feedback on how to improve our classification.
- Gathering feedback is one of the most challenging tasks we have to face. Even though operations are not optimal they are solidified and it is not always easy to introduce new layers and systems.

Alert manager (AM)

- Of course there are activities that take place between all the aforementioned areas.
- CMS i.e. developed an intelligent layer in their infrastructure to detect, analyze and predict abnormal system behaviors using the alerts produced by the infrastructure. This allows the operation teams to focus more on finding solutions for the source of alerts rather than searching, filtering and collecting the alerts



- The alert manager fetches the existing alerts, pre-process them and try to spot any anomalies.
- It then annotates the corresponding Grafana dashboard, were false alerts can be silenced and automatically removed when resolved. This info is then being feeded back to the AM.
- SSB and GGUS are now being integrated to AM. Tickets will be feeded in AM, the same way as alerts do and will be automatically annotated to Grafana dashboard etc.
- This will provide useful insights about when outages happen and how they affect the productivity reported by various systems in CMS dashboards.

Collaborative approach

- Most of those initiatives of course come from sites or experiments but are not bound in their borders
- Most of them can be used by multiple experiments and can attract collaborators from all sides
- The OpInt forum brings people together and makes everybody think of the bigger picture
- Except that we also develop tools to enable the collaboration.
- Namely the OpInt FW and the shared K8S cluster are two of those efforts.



The OpInt FW

- OpInt FW is the shared development platform for most of the OpInt projects.
- We have identified the common requirements across most of our projects and decided to create shareable/extendable components that everyone can use.
- This way we try to enforce reusability and minimise the development effort as much as possible. There is no need for people to reimplementing the same things.
- Anyone who wants to do something similar can clone the framework's example-app and start using our code base.





The shared k8s cluster

- Having a common space to deploy our applications is in line with our cross-experiment goals.
- Inside a K8S cluster all the different services can securely communicate with each other and things like authentication can be implemented once and used by the whole cluster.
- We have created a cluster in the WLCG space and we will start with deploying the FTS log monitoring project that is developed for CMS.
 - A new shared monitoring dashboard is needed in order to open this tool to our whole collaboration.



Doma and Oplnt

- Most of the projects that DOMA includes (data lakes, smart caching etc) are projects that are pushing grid technologies to the edge improving our infrastructure.
- We are sure that there is a lot of overlapping between the DOMA projects and the ideas that are discussed in our forum.
- We would like to invite you to present any parts of your projects that you think are relevant to what you do in our forum and find the common ground. We can definitely benefit from that and maybe we can also offer some help!

Questions?

