
Update on the Operational intelligence

— Panos Paparrigopoulos on behalf of the
operational-Intelligence@cern.ch community —

Operational Intelligence - the 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)
 - **Bottom-up** approach

Operational Intelligence - 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

Active ongoing projects:

- OpInt FW
 - Jobs Buster
 - NLP algorithms development and inclusion in the framework
- FTS logs dashboard

What is the Oplnt FW

- Oplnt FW is the shared development platform for most of the Oplnt 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.

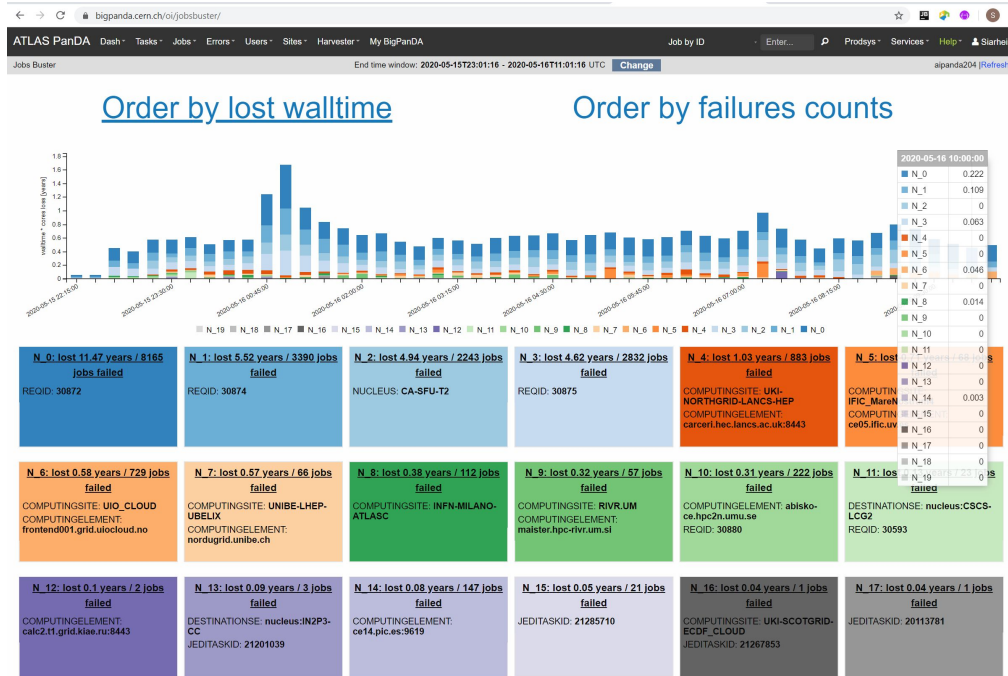
Updates from OpInt FW

- We have implemented a new abstract data loader and used it as a skeleton to create our generic HDFS loader.
- We have been implementing a feedback mechanism to gather feedback of our suggestions/predictions/classifications. This data will be used to identify the success of our algorithms and improve them.
- We are being in the process of porting the CMS operators-console project in the FW.
- Created a documentation platform that we gradually populate with info:
<https://opint-fw-docs.web.cern.ch/>

Updates from NLP

- Following the [requirements](#) we gather we have implemented abstract classes that encode a generic analysis pipeline:
 - Tokenization
 - Vectorization
 - Clustering
- People wanting to contribute just need to provide their custom implementation for the abstract classes and specify how they need to be run
- Expose the different NLP approaches via REST API/ web interfaces:
 - PySpark pipeline already integrated with sample APP demonstration
 - Jupyter notebook with basic implementation has been added to the FW
- We now need to decide on how we will visualise our results and make them available to shifters.

Jobs Buster

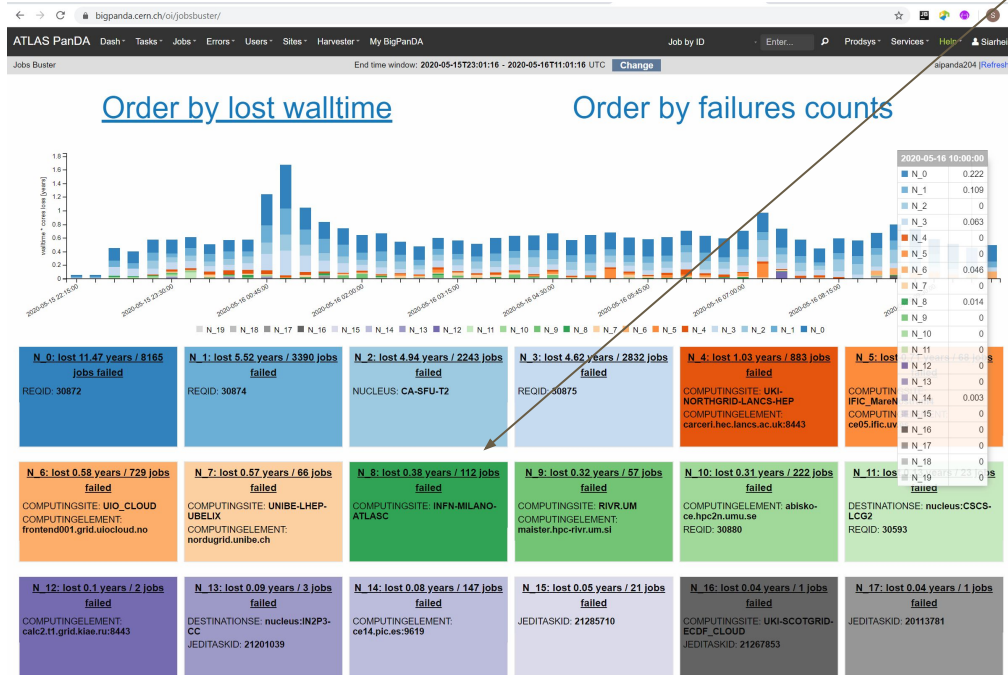


A FW app for spotting operational problems with running jobs

<https://bigpanda.cern.ch/oi/jobsbuster/>

Jobs Buster

Using the FW API to
fetch this info in the
Panda Interface

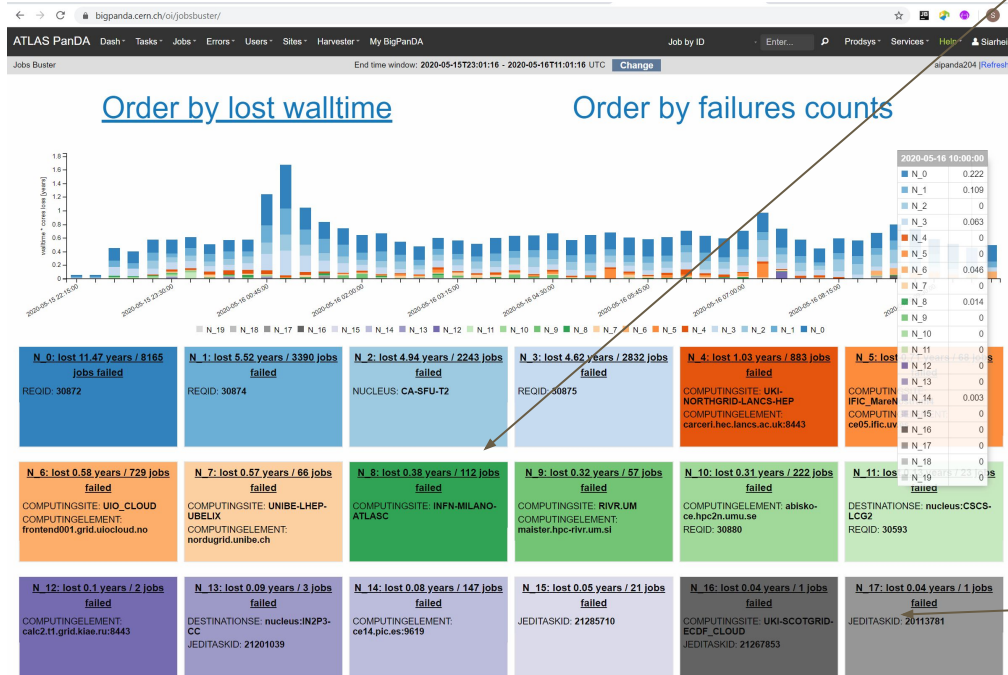


A FW app for spotting
operational problems
with running jobs

<https://bigpanda.cern.ch/oi/jobsbuster/>

Jobs Buster

Using the FW API to
fetch this info in the
Panda Interface



A FW app for spotting
operational problems
with running jobs

NLP Algorithms to
analyse the erros

<https://bigpanda.cern.ch/oi/jobsbuster/>

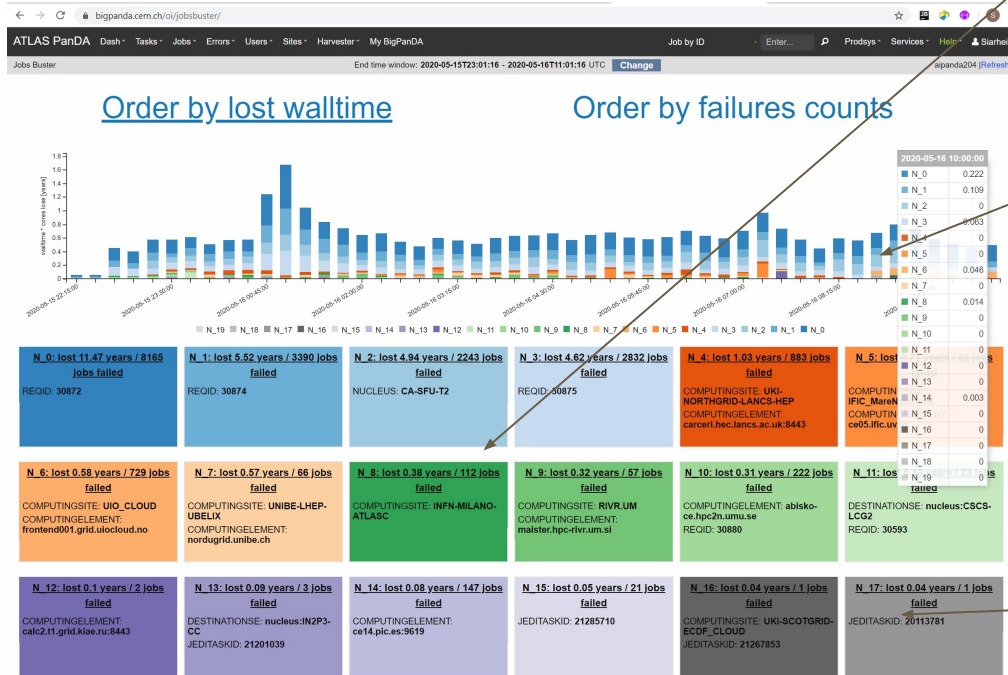
Jobs Buster

Using the FW API to
fetch this info in the
Panda Interface

FW Scheduler to
periodically fetch
errors

A FW app for spotting
operational problems
with running jobs

NLP Algorithms to
analyse the errors



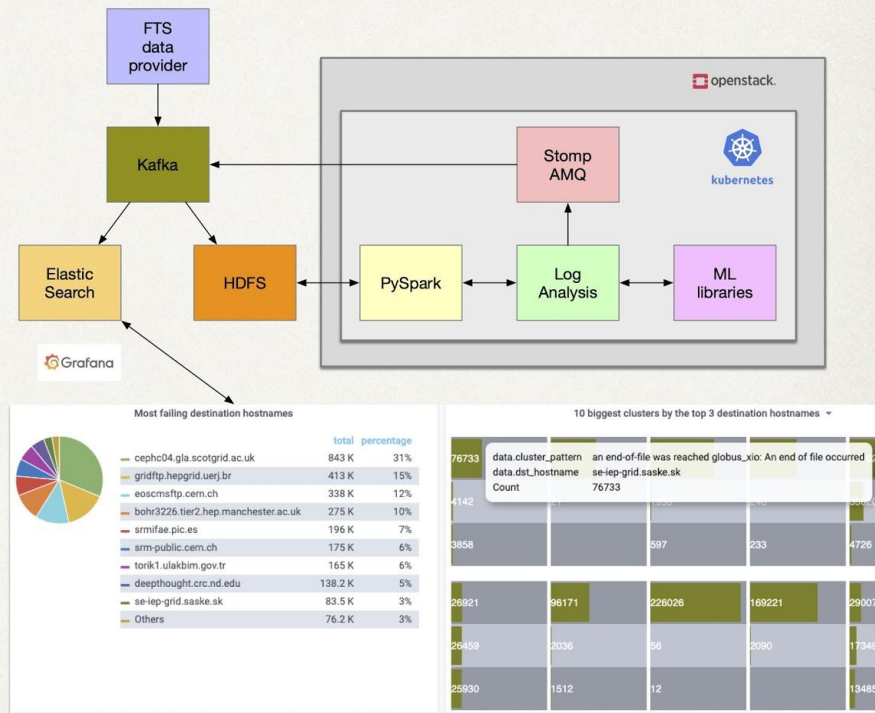
<https://bigpanda.cern.ch/oi/jobsbuster/>

Updates from Jobs Buster

- In assessment by the Atlas Distributing Computing experts
 - Several iterations driven by requests of assessing experts
 - The initial target of finding the “common denominator” approach achieved. Results are not always prominent:
 - A problem during escalation may expand its observables, the “common denominator” changes as well.
 - “Common denominator” not always points into the source of the problem. E.g. a task fails on a particular worker node with no evidence of successful execution on others. Worker node was occupied by failing jobs for some time. Task or Node?
 - Operator always wants to know the reason, not the evidence
- Semantic analysis of logs is the current missing part of Jobs Buster. Not necessarily implemented with DL, we are currently searching for a good approach

FTS Log-analysis workflow

- ❖ FTS logs are available on HDFS
 - ❖ 4 years of logs, 1.7TB on HDFS, data in JSON gzipped format
- ❖ To access and process those logs we used Python and PySpark
- ❖ The codebase relies on Python sklearn library and use DBSCAN clustering algorithm
- ❖ Once data is processed we inject it as JSON records to CERN MONIT infrastructure
 - ❖ this step can be done either in Python/Go/other languages via StompAMQ
- ❖ Visualize data in Grafana



Slide by Valentin Kuznetsov

https://monit-grafana.cern.ch/d/Zx_bXneWz/fts-log-clustering?orgId=11&from=now-7d&to=now

FTS Log-Analysis k8s deployment

- ❖ We wrap the entire codebase into [docker image](#) including access to HDFS and execute it as a cronjob within k8s pod
- ❖ Deployment is done within CMS Monitoring k8s cluster
 - ❖ image size: 3.08 GB (compressed), manifest [file](#)

```
docker run --rm -h `hostname -f` -v /tmp/v/:/etc/secrets -it cmssw/log-clustering /bin/bash
```

k8s status

NAME	READY	STATUS	RESTARTS	AGE
log-clustering-5fc88cdd7-2528d	1/1	Running	0	4h15m

NAME	CPU(cores)	MEMORY(bytes)
log-clustering-5fc88cdd7-2528d	1m	79Mi

Operationally Intelligent Sites

- This is the part of our “manifesto” that we haven’t started tackling yet.
- We had some efforts from single sites doing smart log analysis and trying to reduce their operational needs.
- GDB is a great opportunity to get sites which are interested on board from the beginning.
- We should do our best to unify operations both across Sites and across experiments.

Current challenges / Future steps

- Unifying operations of experiments more (FTS loganalysis could be a shared project in a shared infrastructure).
- Getting more experts on board. OISites can be a challenging project!
- We learned the infrastructure and the tools, we identified potential targets now we need to offer useful “products”.
- If anyone has any ideas on how operations could be simplified we are very happy to listen.

Operational Intelligence - details

- Operational Intelligence [website](#)
- CHEP proceedings [paper](#)
- Github repository: <https://github.com/operationalintelligence>
- Dockerhub: <https://hub.docker.com/orgs/operationalintelligence>
- Indico category: <https://indico.cern.ch/category/11205/>
- Contact us at: operational-Intelligence@cern.ch
- Subscribe at: <http://cern.ch/go/k8zT>