Anomaly detection using QMonit

some ideas and plans

Munich ATLAS Belle II Computing meeting

Eric Schanet

February 15, 2021

QMonit

- **A lightweight, low-latency job monitoring using aggregated queue information**
  - Main data comes from PanDA server (number of jobs in certain state per queue).
  - Pulls queue metadata from various other sources.
  - Low-latency: granularity of 10min, kept for 7 days, then subsequently downsampled.
  - Lightweight: runs as simple cron-jobs on a machine at CERN's lxplus.

QMonit

Essentially a collection of python scripts, held together by a bunch of run-scripts

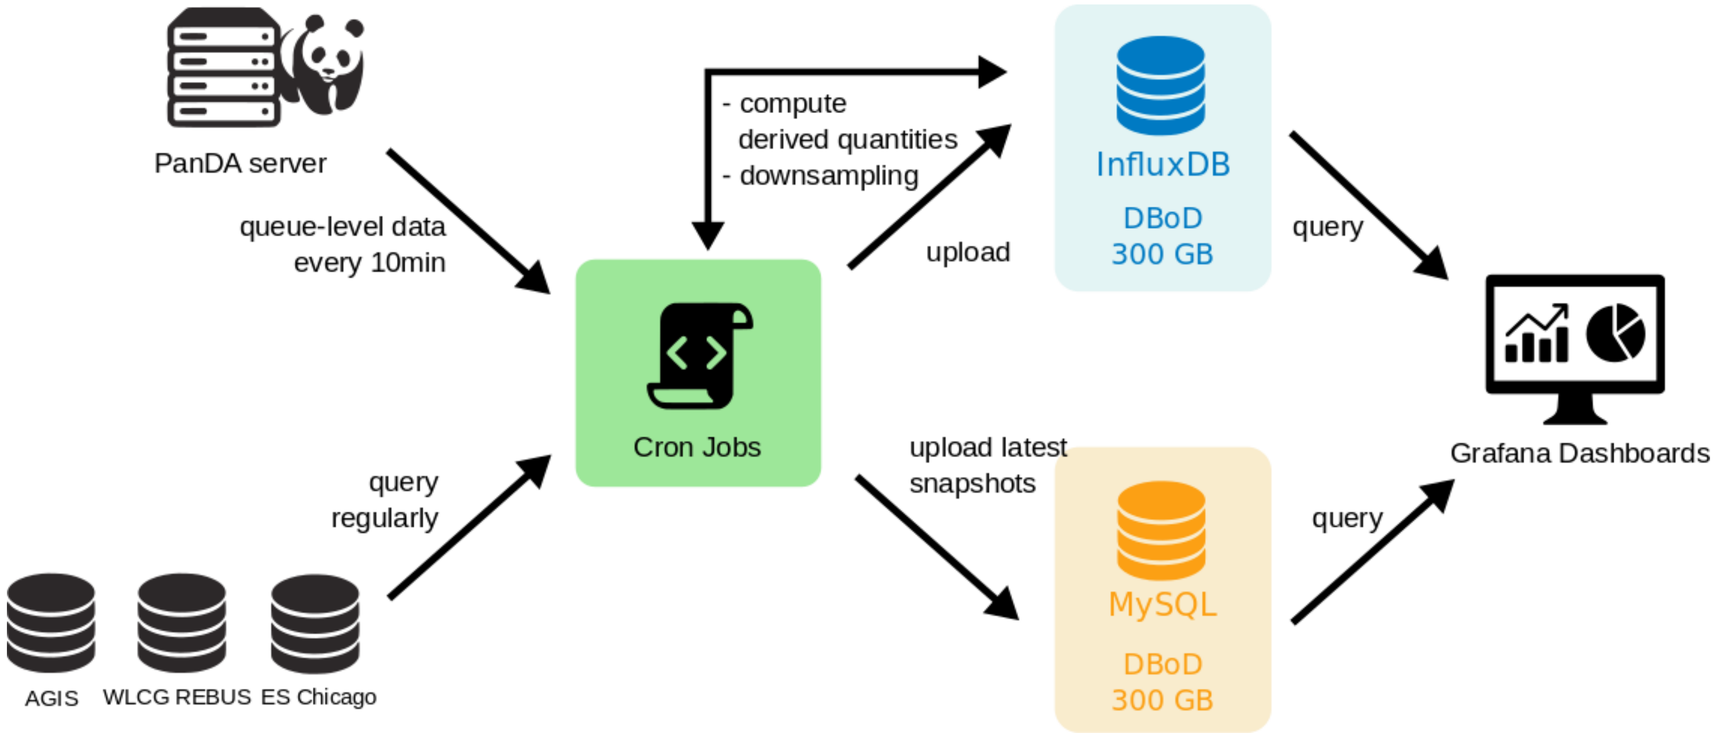

Job monitoring

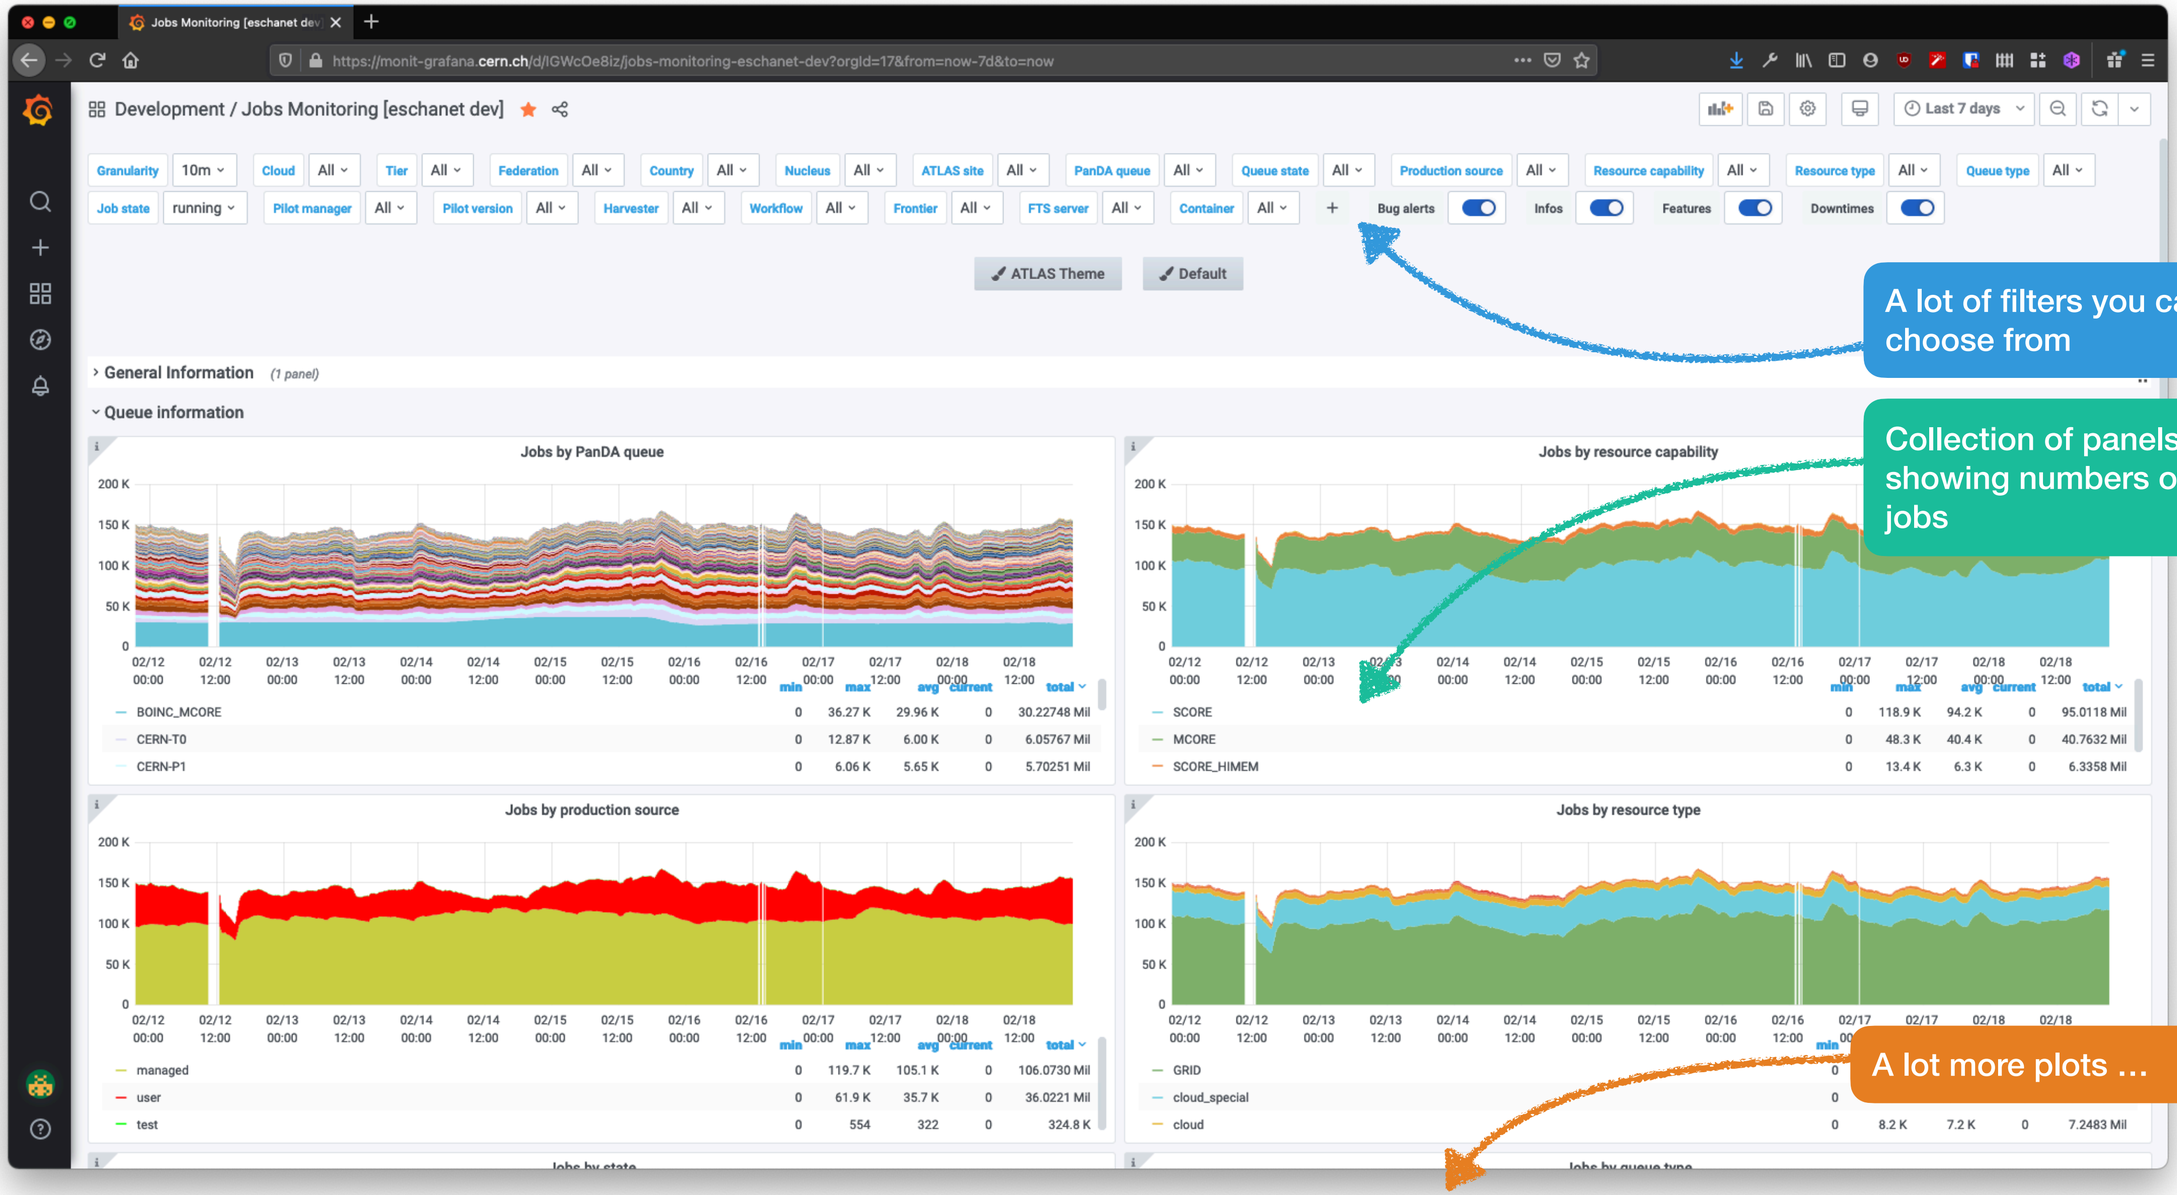

Quite some things to see on there …

- **Large number of metadata fields**
  - Allows to correlate numbers of jobs in different states with different fields.
    - E.g. see when a specific harvester instance goes down, etc…
- **The information is there, but …**
  - You still need someone to actually look at these plots.
  - Anomalies affecting an entire federation/cloud are easy to see
    - Anomalies in single queues or sites not so much
- **Need something more automatic for anomaly detection**

Some special HPC stuff going on

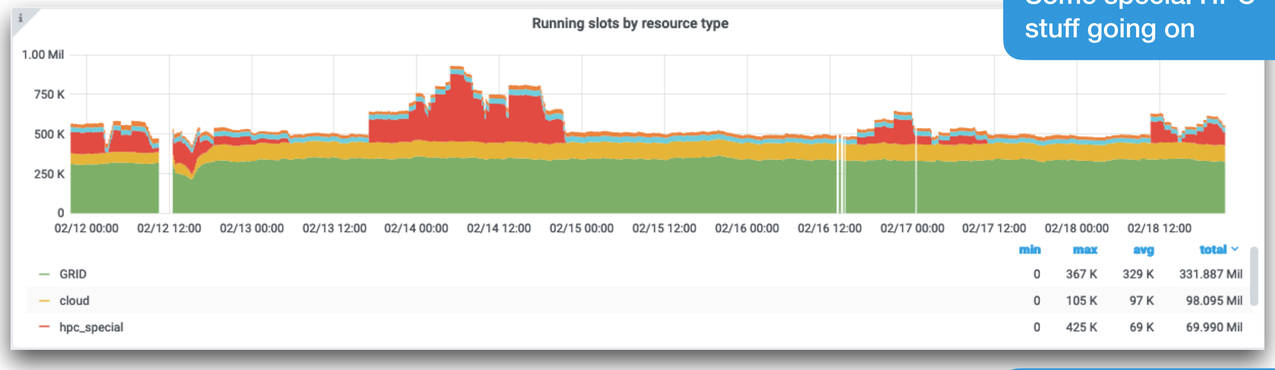

Something's off in UK cloud

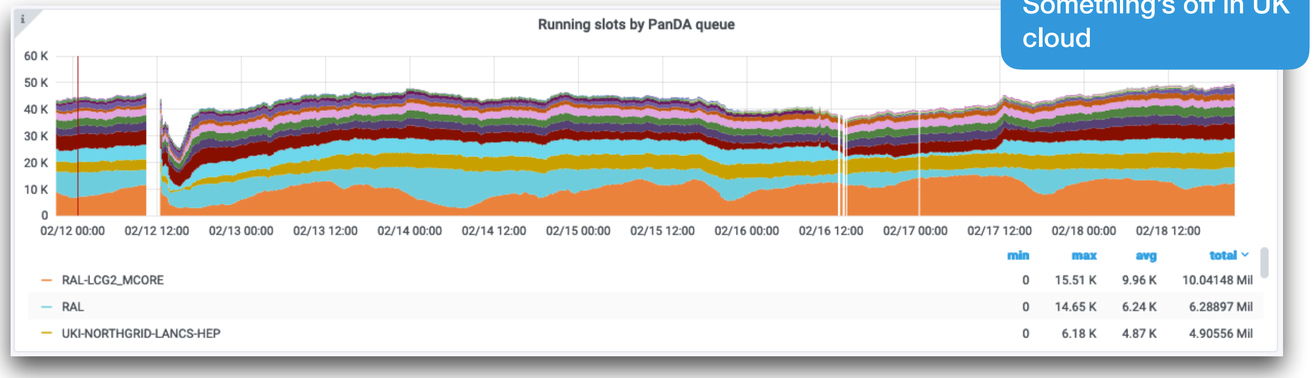

MPPMU also doesn't look so good

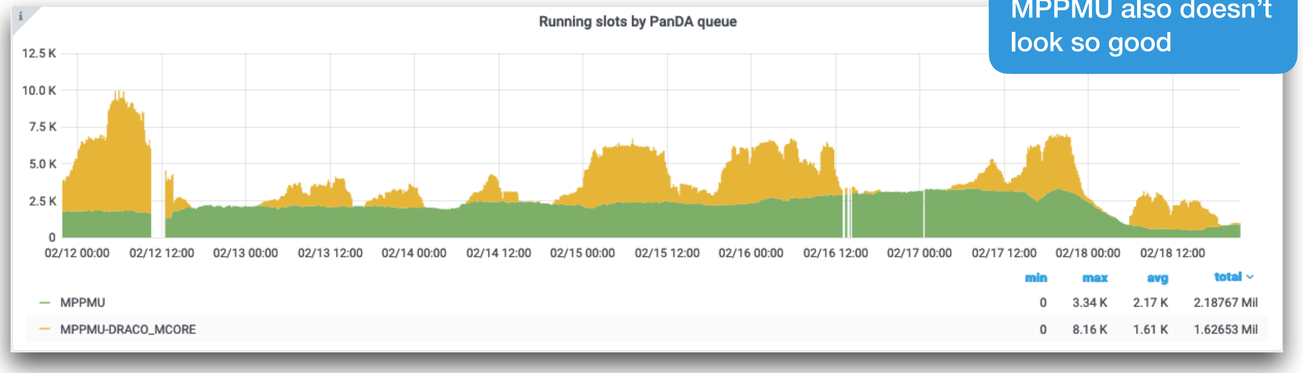

Anomaly detection

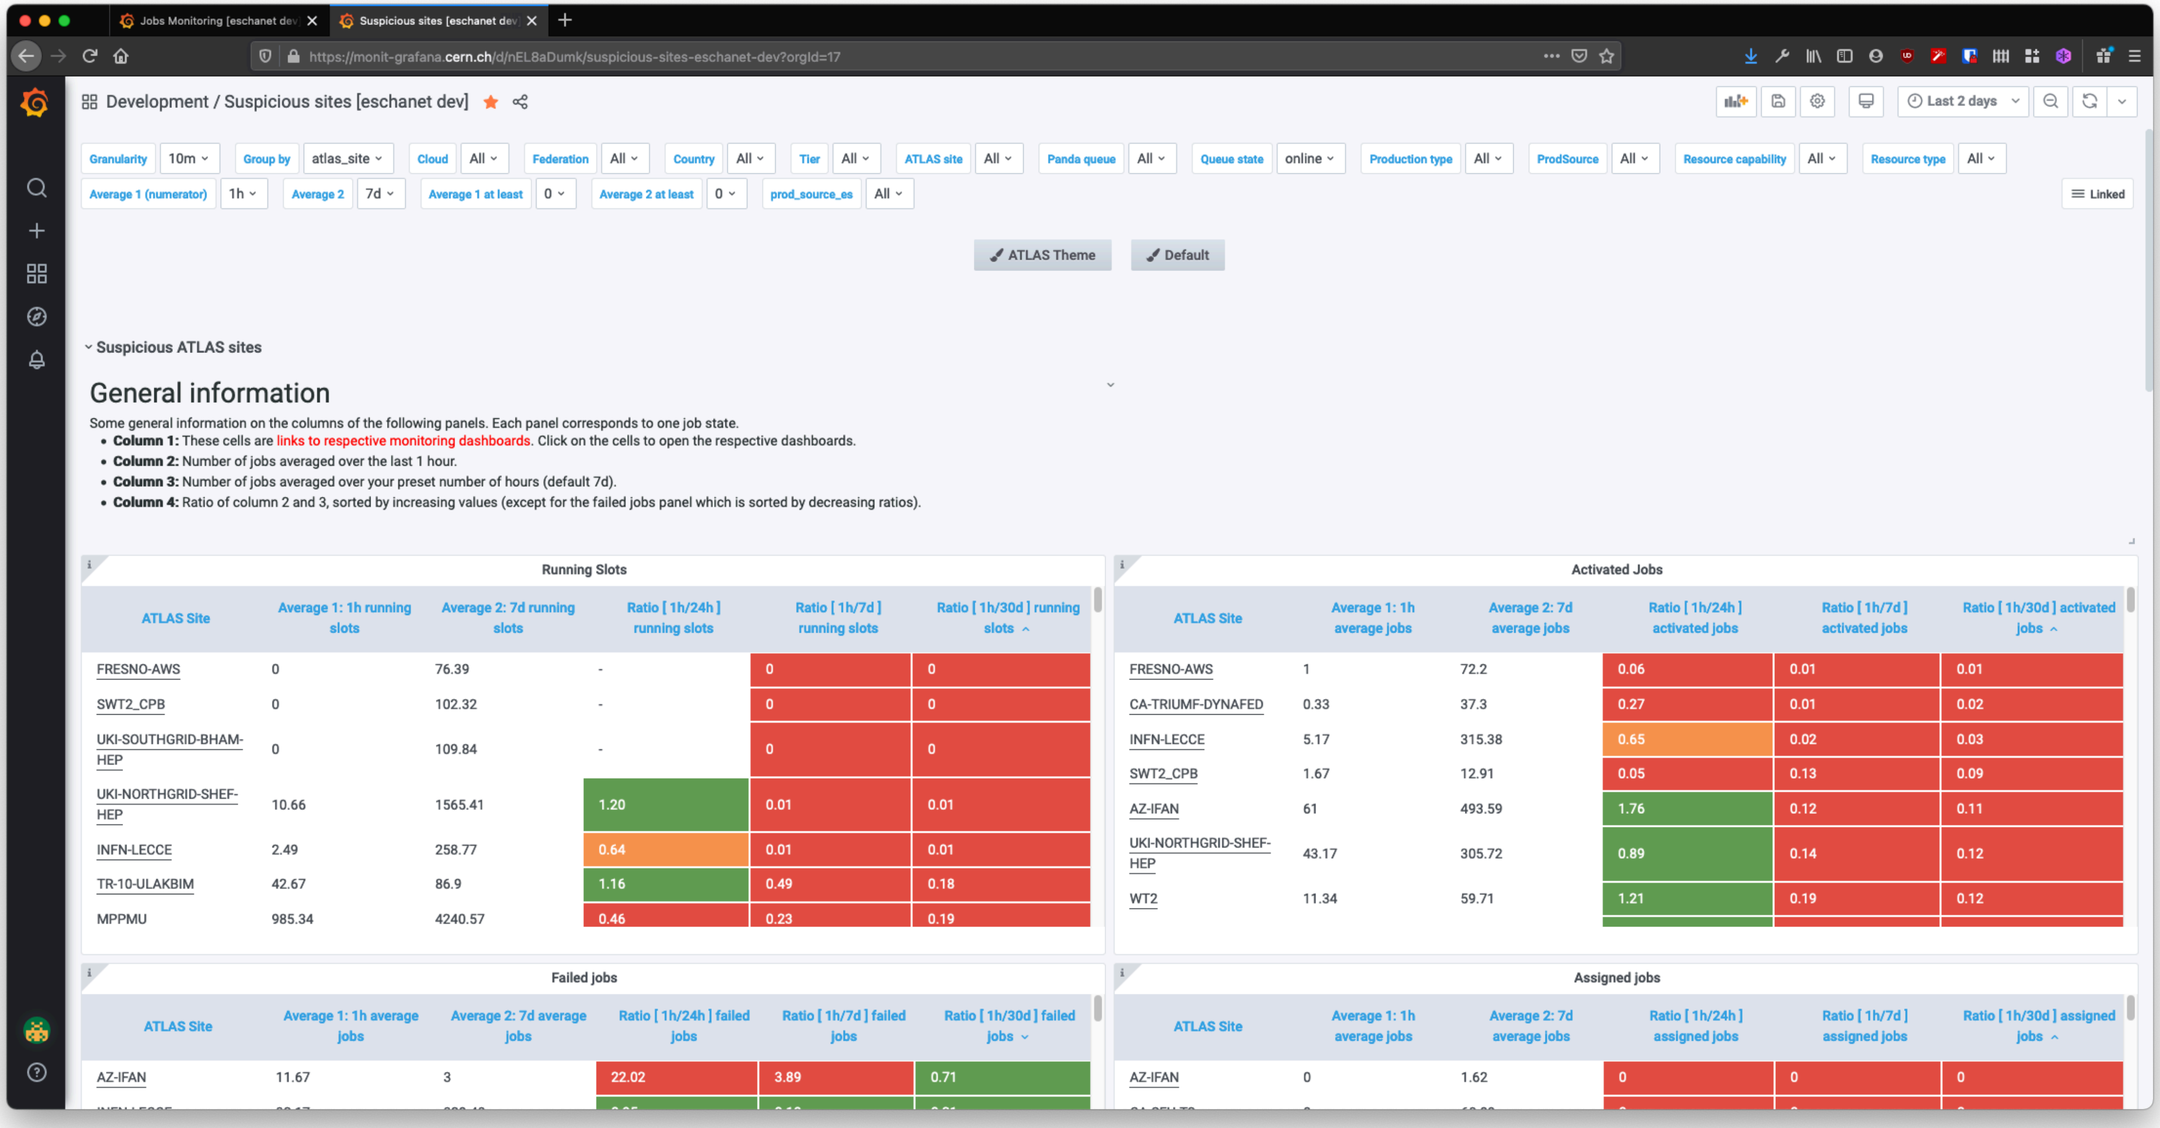

Anomaly detection 2.0?

- **The suspicious sites dashboard is very basic**
  - Just looking at and comparing some moving averages
    - For different relevant job states.
    - For different selections and aggregations of queues (using the ample metadata), which user can specify.
  - Users/shifters still need to look at the actual dashboard
  - At the moment my lowest MA is averaging over 1h (to avoid false positives)
    - sort of negates the 10min granularity …

- **Can we think of something more advanced than this?**
  - Obvious: implement automatic notifications for users/shifters in case of a detected anomaly
  - Could treat it as a supervised or unsupervised problem (labelling anomalies possible, but tedious)
  - Algorithm should be able to handle time-series data in some form.

  - Want to play around with the data? https://cernbox.cern.ch/index.php/s/NysI664NVfcgsHc