

Notes on the CERN-IT MonIT infrastructure used by ATLAS

DB et al. – from 11 September 2019 onwards

Introduction

ATLAS used successfully the monitoring and accounting dashboards developed by CERN-IT together with ATLAS people during LHC Runs 1 and 2. They covered the ATLAS needs in the ADC fields of workload and data management, in addition to providing monitoring tools for the site infrastructure (the Site Status Board and derived views).

During Run 2 it became evident that long-term support for the existing dashboards, largely home-made, could not be guaranteed for the medium and long-term future. The MonIT group designed and developed an infrastructure based on open source tools widely used for many modern BigData applications, including a data transport mechanism based on Kafka, a data storage for time series based on InfluxDB, a global data store in HDFS and graphical displays with Grafana. Some other tools are also available but not mainstream in this context.

ATLAS started testing the new infrastructure in 2016. Since 2017 additional ATLAS specific effort was put in 2 directions:

- Produce a replacement for the old DDM dashboard – mostly an ATLAS development (Thomas Beermann) on the basis of the MonIT infrastructure
- Produce a replacement for the old Job dashboard – initially mostly done by the MonIT group with ATLAS input (for historical reasons), then taken over by ATLAS (Aleksandr Alekseev)

It took much longer than expected but the system reached a decent performance in Spring 2019 and we decided to phase out the old dashboards during the Summer.

Current status

The entry point for all ATLAS dashboards in the MonIT infrastructure is:

<https://monit-grafana.cern.ch/d/000000301/home?orgId=17>

There are 3 groups:

- Production dashboards are mostly stable and can/should be used by shifters for monitoring and managers for monitoring and accounting
- Development dashboards are used to test future versions of the production dashboards and for close-to-production new dashboards
- Playground dashboards are free for all who want to develop new dashboards or new views or combinations of data

Now we have a globally working system; nevertheless, there are several shortcomings that should be addressed in the near future. Here we discuss mostly the infrastructure supporting production dashboards for Data and Job management.

DDM dashboards

A few different dashboards have been developed, mostly by Thomas Beermann, to cover different needs. The data collectors takes information from Rucio at frequencies ranging from 5 minutes to once/day, depending on the data type. Data are stored into HDFS and, in aggregated form, InfluxDB, then displayed using Grafana. ATLAS people are by now more or less happy with the information available; apart from the graphical representation (see later), the main problem people report is:

- The synchronisation of the selection options does not work in the same way in dashboards that take inputs from more than one source (InfluxDB and HDFS usually), or one has to wait a long and unknown time before this synchronisation happens

Job accounting dashboard

For this dashboard the data collection procedure was developed by the MonIT team. It scans the PanDA database every 10 minutes and collects 2 kinds of information:

- Transient, such as the number of active jobs (waiting, running, finalising) “now”, with all their attributes; this info cannot be reconstructed *a posteriori* if lost on the way;
- Persistent, such as the number of jobs started or completed in this time bin, with all their attributes; this info can be back-filled if lost for any reason.

The records pass through Kafka and are augmented with information (mostly) from AGIS with the relation between PanDA queues, sites, federations, countries, tiers etc.

Apart from the graphical representation (see later), the main problems people report are:

- Often a few bins are left empty or underfilled, or even overfilled, because of breakdowns of the data flow. Usually shifters and/or the CRC report the fault through a SNOW ticket and the problem is fixed by a manual intervention by the MonIT group, but these problems should not happen in the first place or get fixed promptly and automatically before people notice;
- Because of the long and varying transit time through the MonIT infrastructure, all information relative to the last few (at least 3) hours is to be considered unstable and incomplete; hence this dashboard can be used reliably for accounting but not for monitoring purposes, differently from the old Job dashboard;
 - A separate dashboard was in the meantime developed for short-term monitoring purposes. It is site-oriented, aggregating information for all jobs on each site, but keeping the 10-minute time bins and displaying information immediately as it has its own data collector from the PanDA DB and its own instance of InfluxDB; the only commonality with the MonIT infrastructure is the use of Grafana – in the development area for the time being:
https://monit-grafana.cern.ch/d/IGWcOe8iz/jobs-monitoring-eschanet-dev?orgId=17&var-retention=10m&var-queue_state=online&from=now-24h&to=now
- Because of the data aggregation, it is not possible to select one or a few tasks and see the plots only for the jobs that belong to those tasks;
- Data are collected every 10 minutes and aggregated by default in 1-hour bins, but the 10-minute bins are no longer available (anyway not so useful, see above);
- Aggregation into 1-day bins works correctly, but 7-day bins seem to start on random days of the week (one never knows before trying). Same for 30-day bins, which are anyway not very useful – much more useful would be to have monthly bins.

Some of these shortcomings are currently filled by the ElasticSearch+Kibana system hosted by the University of Chicago, which has its own set of problems, and another part by the BigPanDAMon monitoring that can also provide some task-based histograms. It would be much better to have a single environment providing all necessary views.

Grafana problems

- Grafana is advertised as a system for BigData display but it is very slow unless the number of open plots and the time range are both small; otherwise one should be very patient, at times more than with the old dashboard.
- Every time a selection is changed, grafana tries to display all plots. This is extremely inconvenient when one wants to change several options at once and then “go”, especially when a long time range was previously selected.
- It cannot display matrices or other plots with multiple clickable contents (the old dashboard was able to do it).
- It cannot display bar charts with horizontal bars (the old dashboard was able to do it, with clickable contents).
- It is not possible to customise the display of the plots (labels, title, font sizes, aspect ratio, legend) to make them suitable for slides and reports. The only option is to dump the contents to a CSV file and use another display tool.
- There is no way to force the display of a line (the pledges in our case) to be always in front of the histogram.
- The last bin in any histogram can be selected to have a correct display OR correct statistics, but not both: see for example <https://monit-grafana.cern.ch/d/000000696/job-accounting-historical-data?orgId=17&from=1568246400000&to=1568332799000&fullscreen&panelId=6> (correct statistics but not displaying the last bin) and compare it to <https://monit-grafana.cern.ch/d/000000696/job-accounting-historical-data?orgId=17&from=1568246400000&to=1568332800000&fullscreen&panelId=6> Both plots are for 24 hours, the first one from 00:00:00 to 23:59:59 (same day), the second from 00:00:00 to 00:00:00 the following day. The first plot displays only 23 bins but the 24th exists (you can hover with the pointer and see the values); the second plot displays all 24 bins but there is a 25th invisible one that contributes to the statistics. The reason is probably in the incorrect treatment of time bins in grafana. One can work around this problem but it is far from evident that such problems should exist in public display tools.
- There is no way to have 7-day (weekly) bins starting always on the same day of the week (Sunday or Monday) and have week numbers on the time axis.
- There is no way to have monthly bins, rather than 30-day bins, starting always on the 1st of each month and having the month number, or name, on the time axis.
- The generation of an exportable picture from a dashboard reloads all the data (can take a long time) and takes no account of the display options applied to the legend, like showing or hiding some of the entries.

Suggestions

- Find a replacement for grafana that would not have all problems reported above.
 - Or find a way to configure grafana to overcome all reported issues – not found so far.
- Find a way to collect and keep for a long time (“the duration of the experiment”) the job records that can be used for monitoring and accounting (in aggregated formats) and can be accessed within a coherent display framework.
-