



MONIT/CMS Meeting

15.06.2020



Outline

- Dashboards migration
- Current status
- Major achievements
- Future plans (MONIT internal)
- Future plans (impact on CMS)

Dashboards migration

Objectives of dashboards migration

- Consolidate in common MONIT infrastructure
- Preserve old functionality / Import historical data
- Build an internal community / Contribute upstream

Successful migration of old dashboards

- Activity started in 2017 with MONIT infrastructure deployment
- Major boost in 2018 with creation of CMS Monitoring group
- All dashboard migrated (except WLCG Site Monitoring)
- Many other/new CMS producers integrated in parallel

Old dashboards migration timeline


- *Sep 2017* - WLCG Transfers
- *Jul 2019* - CMS Job/Task Monitoring
- *Mar 2020* - CMS Site Status Board
- *May 2020* - WLCG Rebus (*in collaboration with CRIC*)
- *Jul 2020* - WLCG Site Monitoring

Group by CMS_JobType Tier All Site All Type All CMS Job Type All Primary Data Tier All Binning 1h Filters +

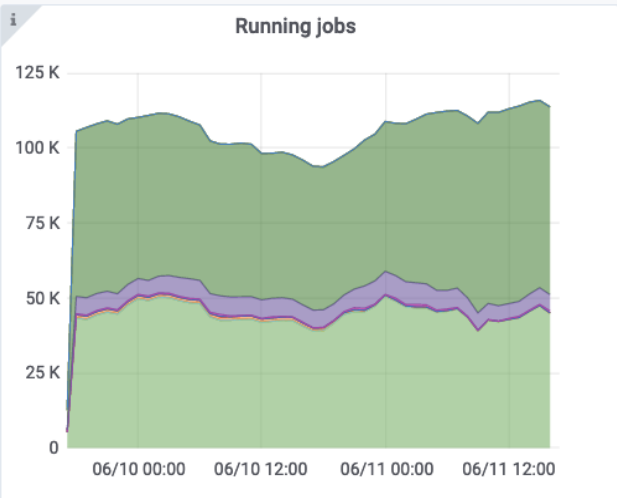
Please read

Instructions on how to use the dashboard at the bottom of this page.
 Plots with 1-hour binning are available from Nov. 10th, 2018.
 Plots with 1-day (and more) binning are available starting from Dec. 4th, 2017.
 Please be patient when you request data from a time period longer than 6 months (queries will go faster if you pick 7 or 30-day binning).

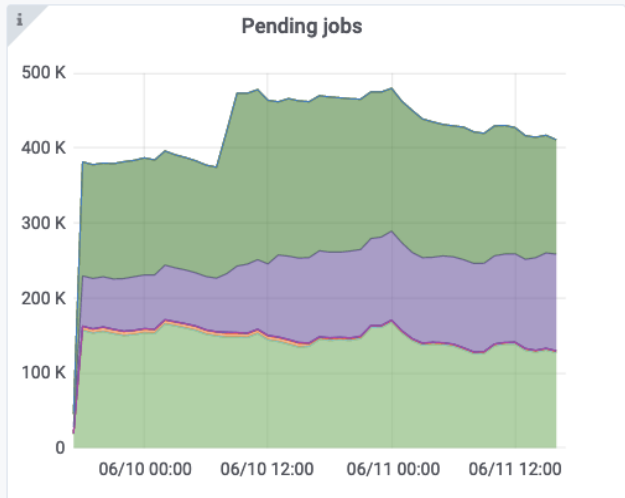
1688327 completed



115890 running jobs



416995 pending jobs



VO cms Country All Federation All Tier 1 Site All Flavour All Profile CMS_CRITICAL Hostname All

Recomputations Start Recomputations End

More Dashboards



Sites

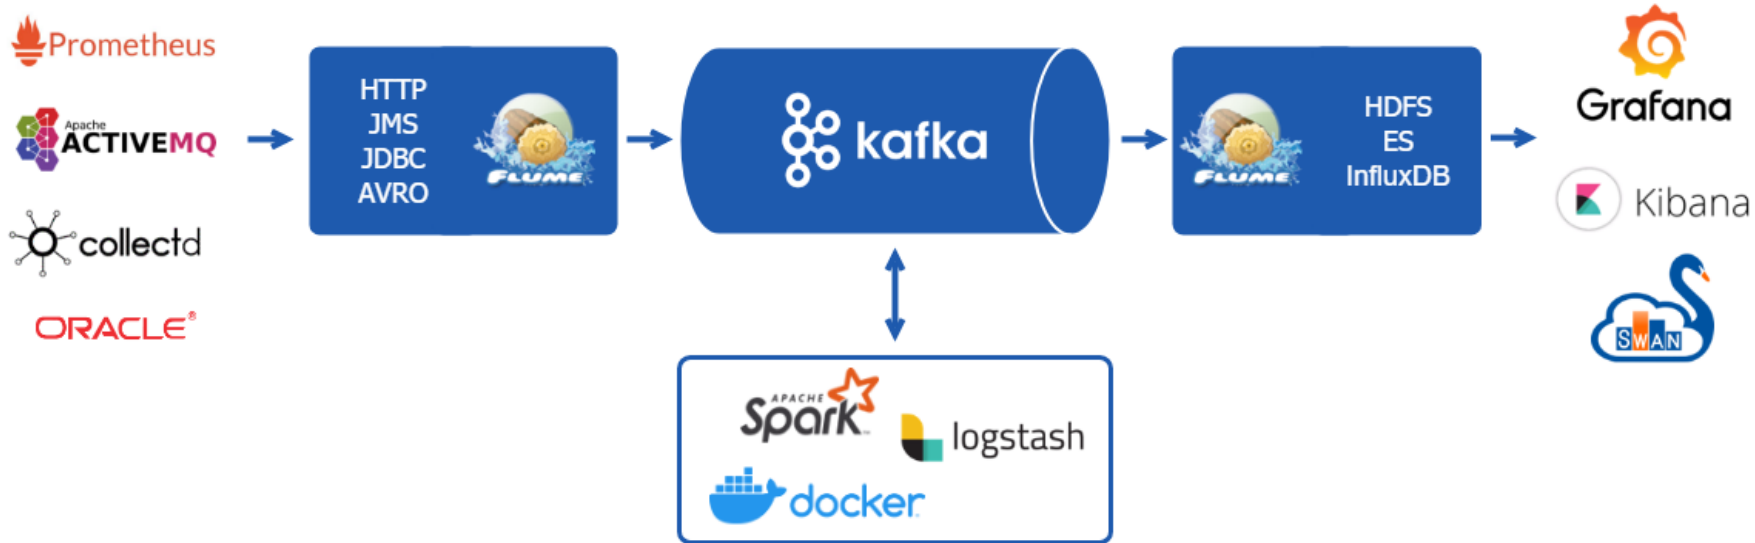
T1_DE_KIT: OK (97.26%) CRITICAL (1.94%) UNKNOWN (0.49%) DOWNTIME (0.28%) NODATA (0.03%)	OK	OK	OK	OK	OK	OK
T1_ES_PIC: OK (99.83%) CRITICAL (0.14%) NODATA (0.03%)	OK					
T1_FR_CCIN2P3: OK (99.90%) CRITICAL (0.07%) NODATA (0.03%)	OK					
T1_IT_CNAF: OK (99.27%) CRITICAL (0.63%) WARNING (0.07%) NODATA (0.03%)	OK	OK	OK	OK	OK	OK
T1_RU_JINR: OK (99.97%) NODATA (0.03%)	OK					
T1_UK_RAL: OK (99.06%) UNKNOWN (0.42%) DOWNTIME (0.49%) NODATA (0.03%)	OK	OK				OK
T1_US_FNAL: OK (98.23%) CRITICAL (1.74%) NODATA (0.03%)	OK	OK	OK	OK		

Endpoints (1 panel)

Tests (1 panel)



Current status

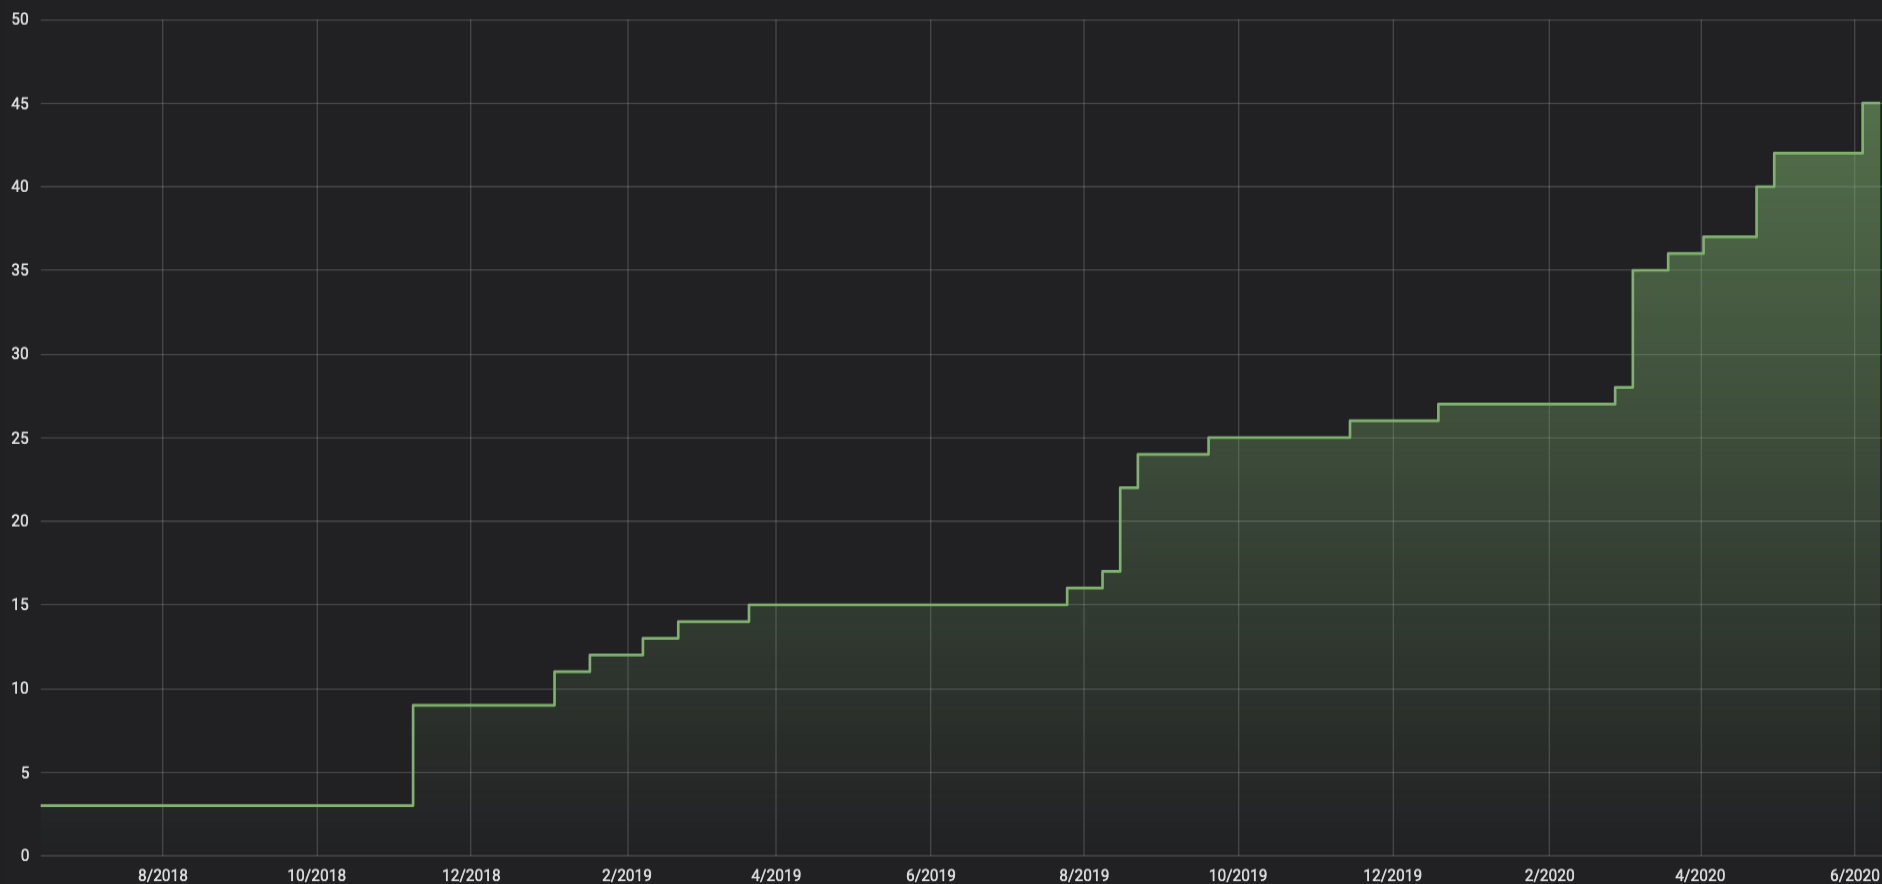


CMS is using all MONIT services

- *Ingestion*: Integration of data via AMQ and HTTP (metrics and logs)
- *Transport/Processing*: Spark data enrichment and aggregation
- *Storage/Visualization*: ES, InfluxDB, and HDFS / Grafana and Kibana

Some statistics of CMS @ MONIT today

- Number of producers: 45
- Number Grafana dashboards: 323 (using 92 data sources)
- Total data volume in ES: ~10 TB (kibana) + ~0.7 TB (timber)
- Total data volume in HDFS: ~27 TB



Major achievements

Migration of Job Monitoring backend to ES

- Better performance, especially on longer time periods
 - e.g. 6 months data (1 day bin): from ~3-5 min to ~2-5 sec
- Easier to change schema (add field, replace field)
- Out-of-the-box correct time literals (1w, 1M)

Improved FTS efficiency

- From CMS proposal and accepted by all other VOs
- More accurate computation algorithm

Implemented new spark job to enrich Rucio data

Grafana as central location for all dashboards

- Easy to add new plots for same/different data sets
- Very useful and complete homepage

Stable infrastructure operations

- Improved procedure for interventions (e.g. Grafana)
- Added Test/QA flow to many CMS producers

Very good collaboration between teams

- Meetings every other week / Participation in training event
- 91 SNOW tickets closed in the last 12 months

Future plans (MONIT internal)

Keep the infrastructure updated

- Upgrade Kafka and ElasticSearch to latest versions

Move services to Kubernetes

- Good candidates are Mesos, Marathon, and Chronos

Evaluate new interesting technologies

- Mainly in the areas of timeseries storage and data transport
- Alternatives for some use cases where InfluxDB and Flume are used

Redesign Monitoring of Monitoring

- Current duplicate workflow is too complex

Future plans (impact on CMS)

Integrate Prometheus metrics in MONIT

- Use MONIT flow for aggregated Prometheus metrics

Redesign old GLED collectors

- Old collectors for XRootD data must be retired/replaced

Migrate XSLs producers

- Move producers from old XML to new JSON format

New data retention in HDFS

- For logs default will be 13 months from 1st Sep
- Will also be introduced for metrics later

Move to CRIC topology

- Replace current VOFeed endpoint
- Affects WLCG Transfers and WLCG SiteMon

