MONITORING SYSTEM OF THE AMS SCIENCE OPERATION CENTRE

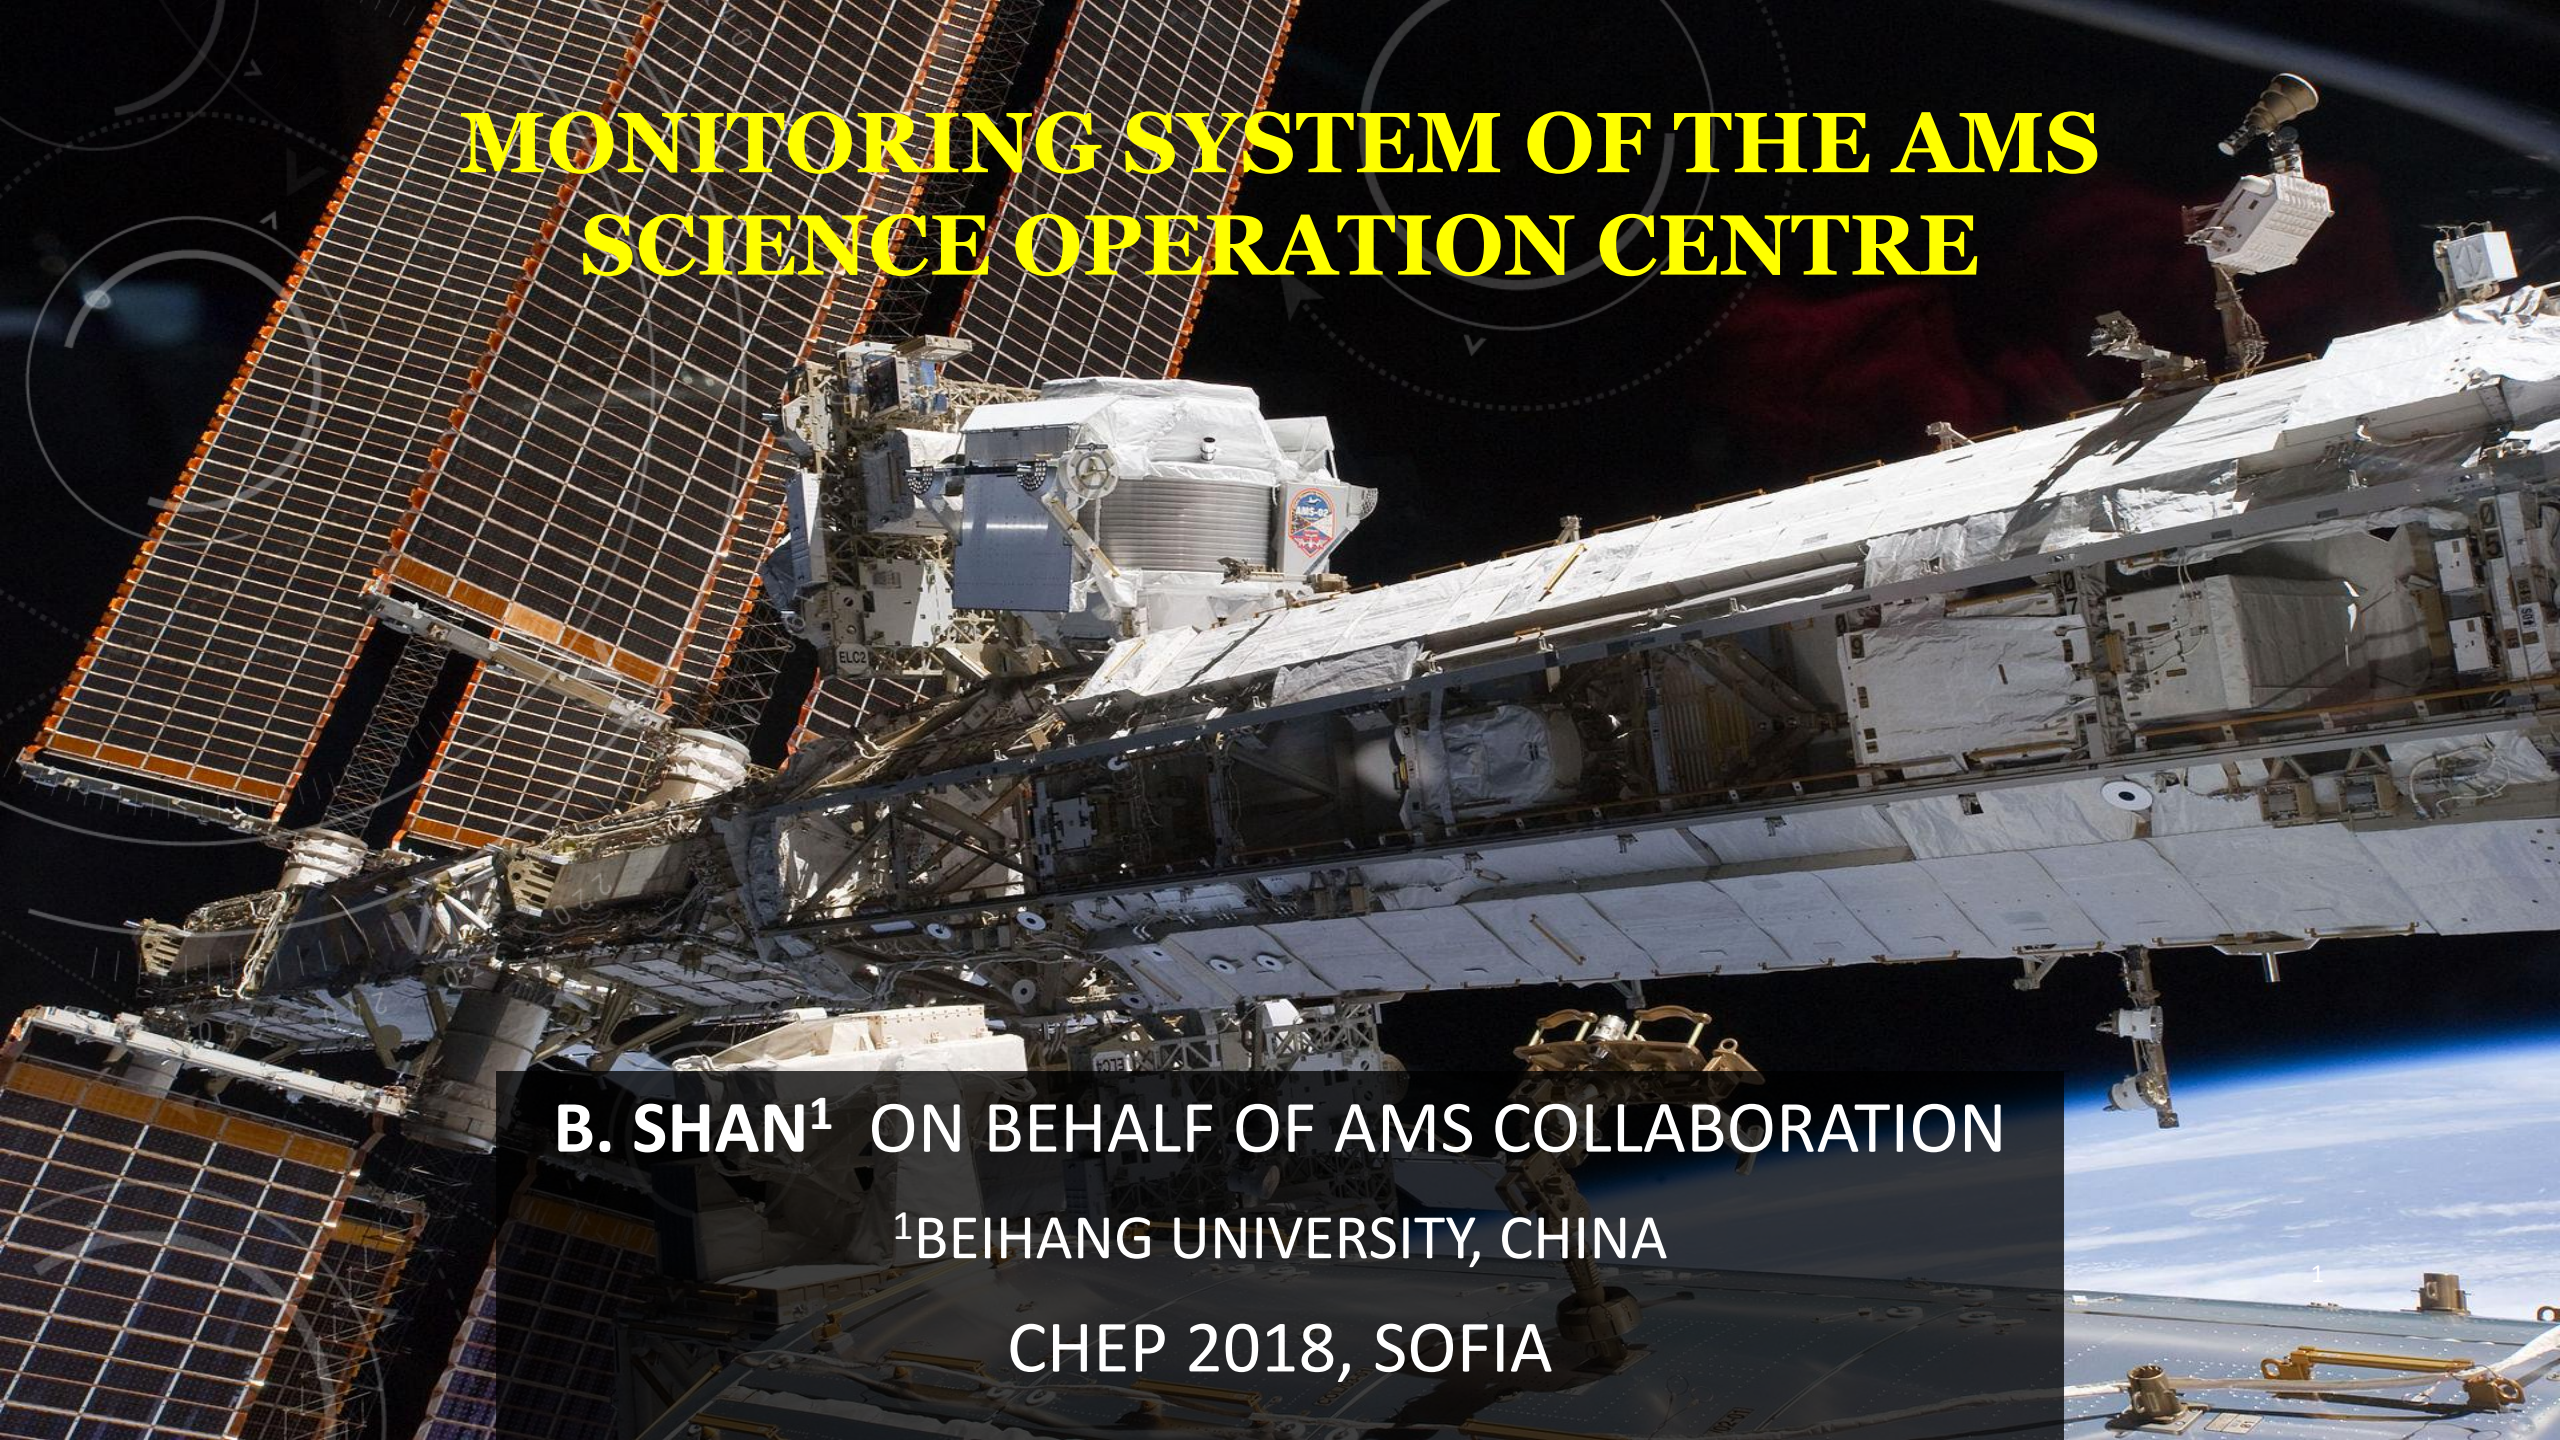

B. SHAN[1] ON BEHALF OF AMS COLLABORATION

[1]BEIHANG UNIVERSITY, CHINA

CHEP 2018, SOFIA

AMS EXPERIMENT

- The Alpha Magnetic Spectrometer (AMS) is a high energy physics experiment on board the International Space Station (ISS), featured:
  - Geometrical acceptance: **0.5 m^2•sr**
  - Number of Read-out Channels: ≈**200K**
  - Main payload of Space Shuttle **Endeavour's last flight** (May 16, 2011)
  - **Installed on ISS on May 19, 2011**
  - 7x24 running
  - Up to now, over **122 billion events** collected
  - Max. event rate: **2KHz**

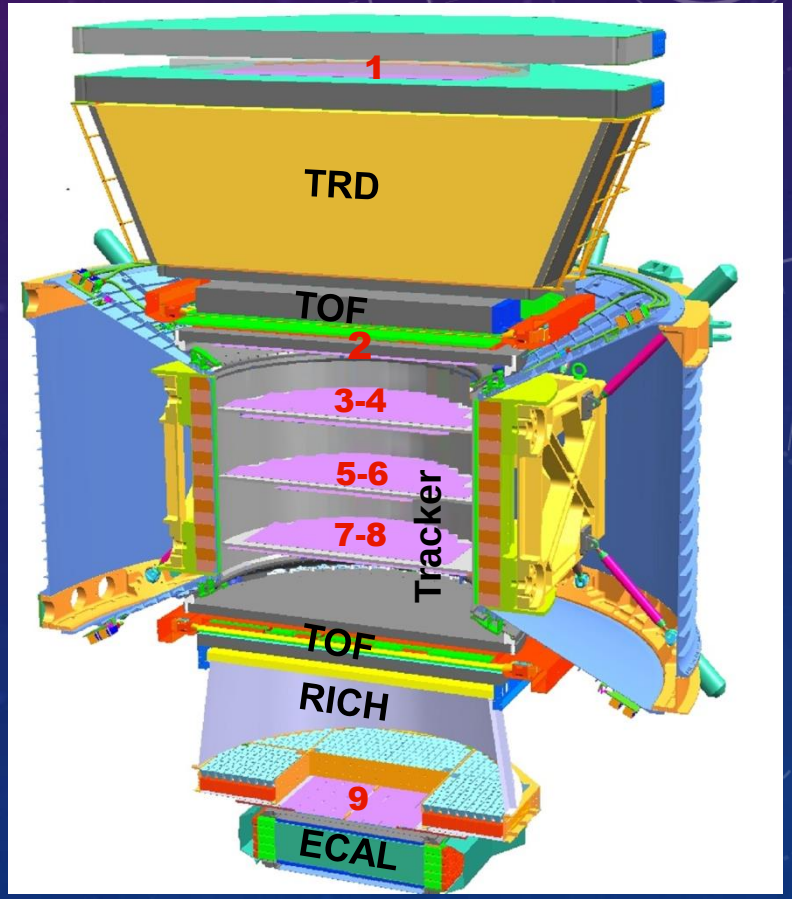

AMS DATA FLOW

- Data transferred via relay satellites to Marshall Space Flight Center, then to CERN, nearly real-time, in form of one-minute-***frames***
- ***Preproduction***: Frames → runs (***RAW***): 1 run = ¼ orbit (~23 minutes)
- ***Standard Production***
  - Runs 7x24 on freshly arrived data
  - Initial data validation and indexing
  - Produces Data Summary Files and Event Tags (ROOT) for fast events selection
  - Usually be available within 2 hours after flight data arriving
  - Used to produce various calibrations for the second production as well as quick performance evaluation

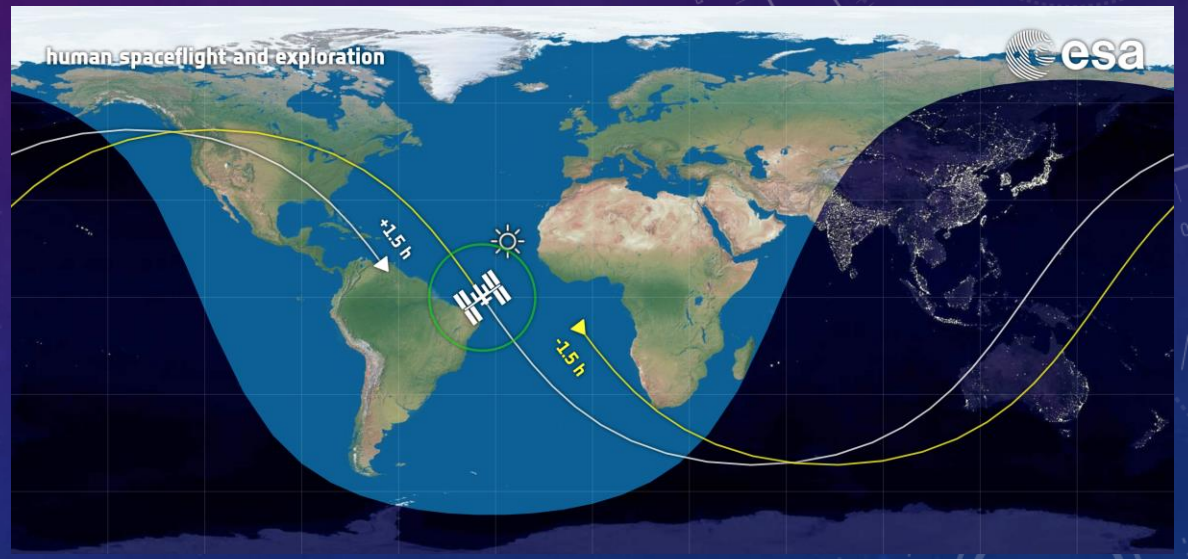

SCIENCE OPERATION CENTER (SOC)

- Processing of the AMS science data for detector evaluation and physics analysis
  - Data production
  - Monte-Carlo simulation
- Maintaining local production farm
  - 20 hosts, 302 cores
  - 400 TB storage
- Operations on the services/resources provided by CERN IT
  - Batch (HTCondor & LSF), EOS, CASTOR, AFS, CVMFS, ELOG, PDB-R, etc.

Monitoring (all of the above) is important and challenging!

CHALLENGES OF SOC MONITORING

- Data taking and processing is 7x24 running
  - For detector performance evaluation
- Long data transmission path
- Widely distributed compute facilities
  - Europe: CERN, RWTH, CNAF, IN2P3, …
  - Asia: SEU, Acad. Scinica, …
  - America: NERSC, ALCF, …
- Various service providers

WHAT WE ALREADY HAVE

- Frame monitor
  - delay of frames
- Production monitor
  - standard production jobs
- Data files list
  - Summary of raw and ROOT files
- NetMonitor
  - Hardware/storage monitoring of local production farm

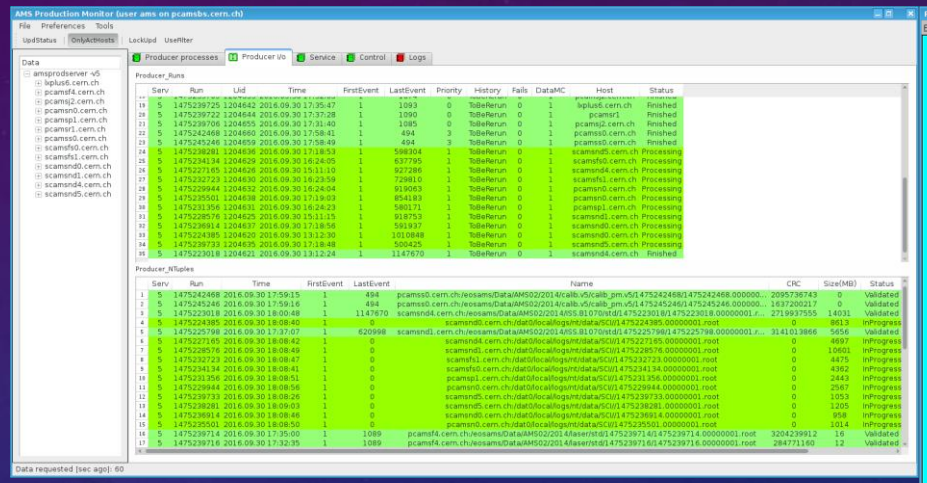

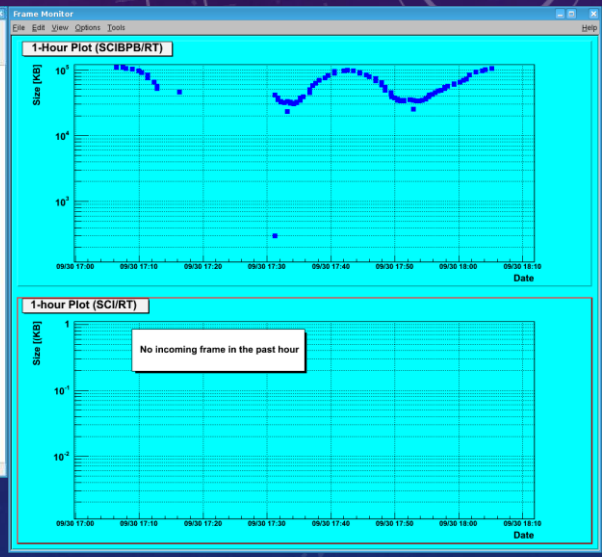

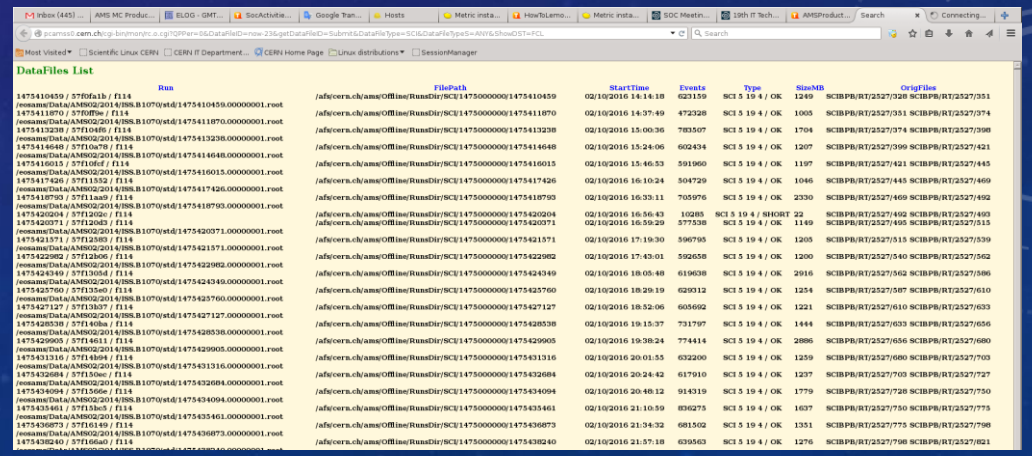

We need more centralized monitoring

WHAT CERN IT PROVIDES

- MONIT Infrastructure [1]

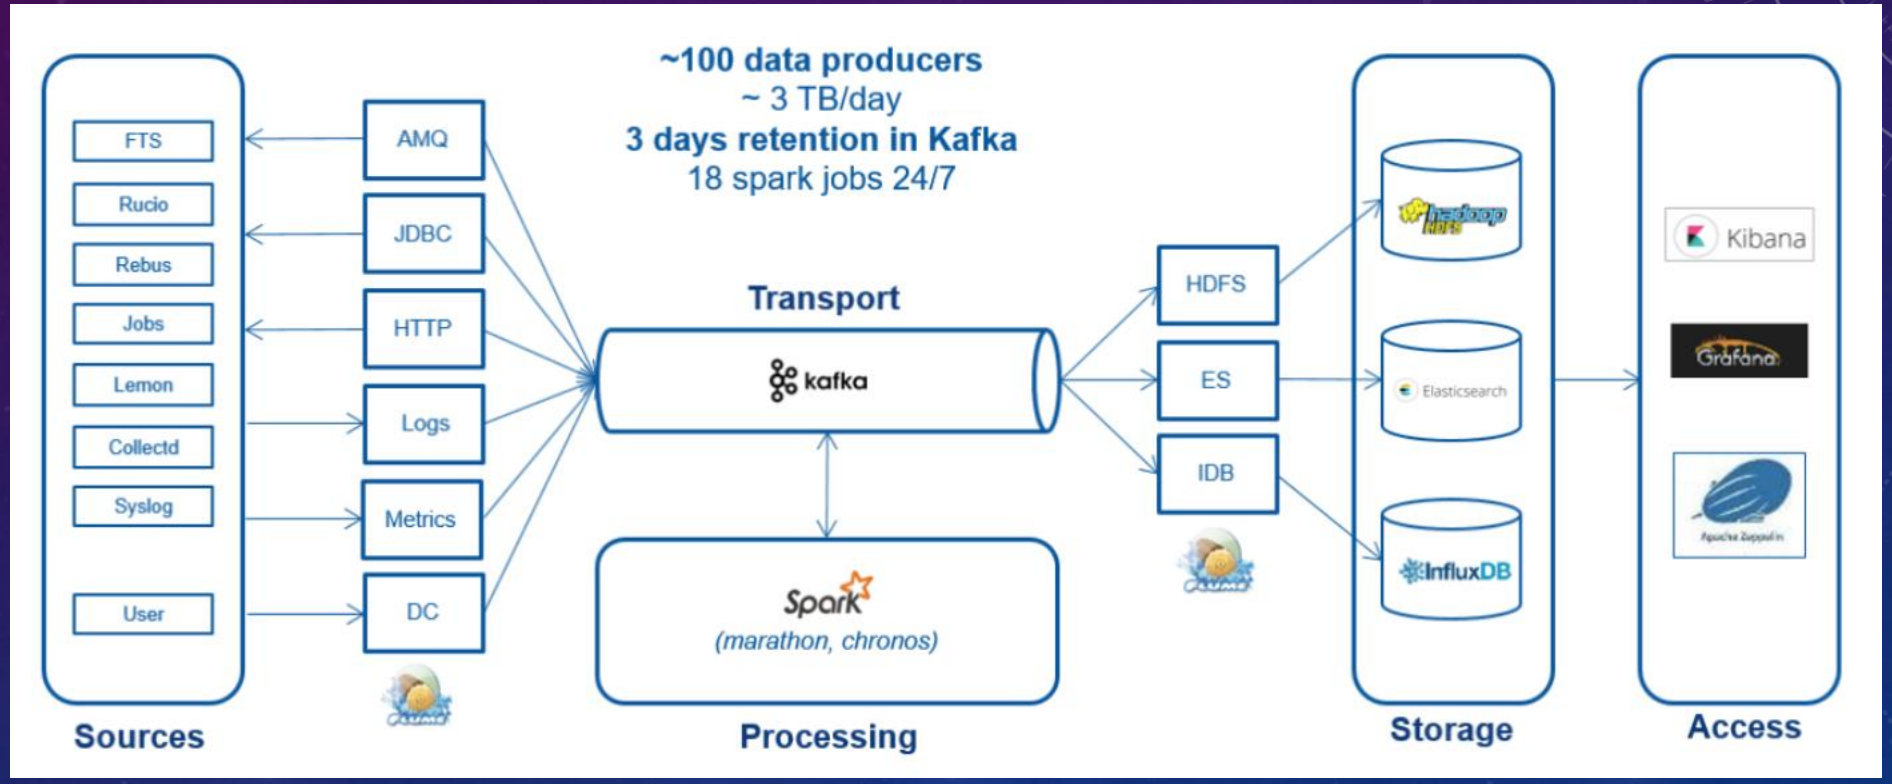

[1] http://monitdocs.web.cern.ch/monitdocs/training/index.html

ROADMAP – CERN IT SERVICES/METRICS

- Sources
  - Taken care of service providers (EOS/batch/FTS/… teams)
- Transport, processing, and storage
  - Taken care by MONIT
- Access
  - Modify dashboards from service providers to concentrate on the services/hosts which affect us
  - Create relevant Alerts with notifications

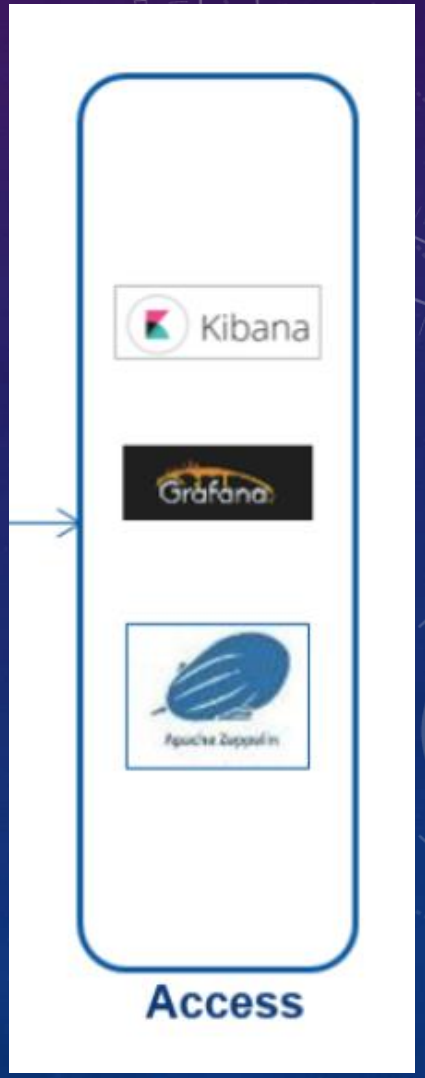

ROADMAP – OUR SERVICES/METRICS

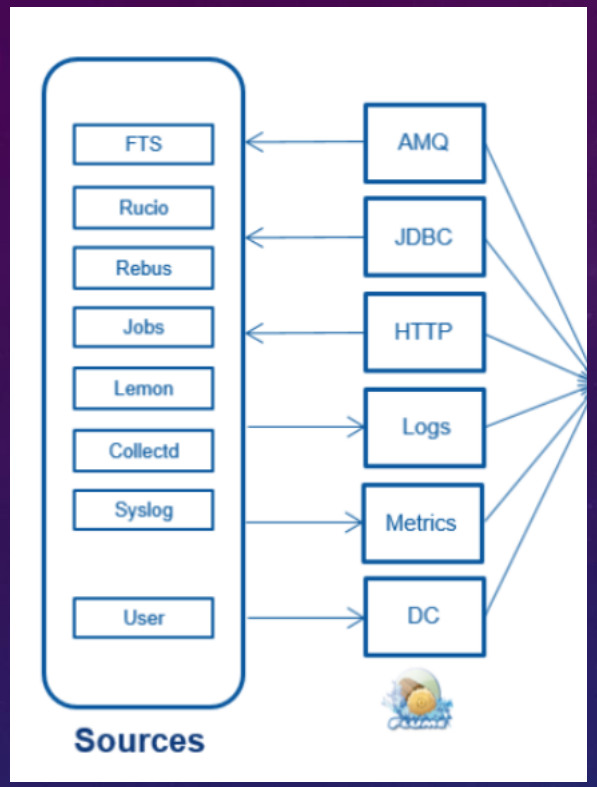

- Sources
  - Modify our existing monitoring tools to send “Metrics” to InfluxDB
  - Use Collectd to gather the standard hardware/storage data of hosts
  - Log files to HDFS
  - And more…
- Transport, processing, and storage
  - Taken care by MONIT
- Access
  - Build dashboards
  - With Alerts

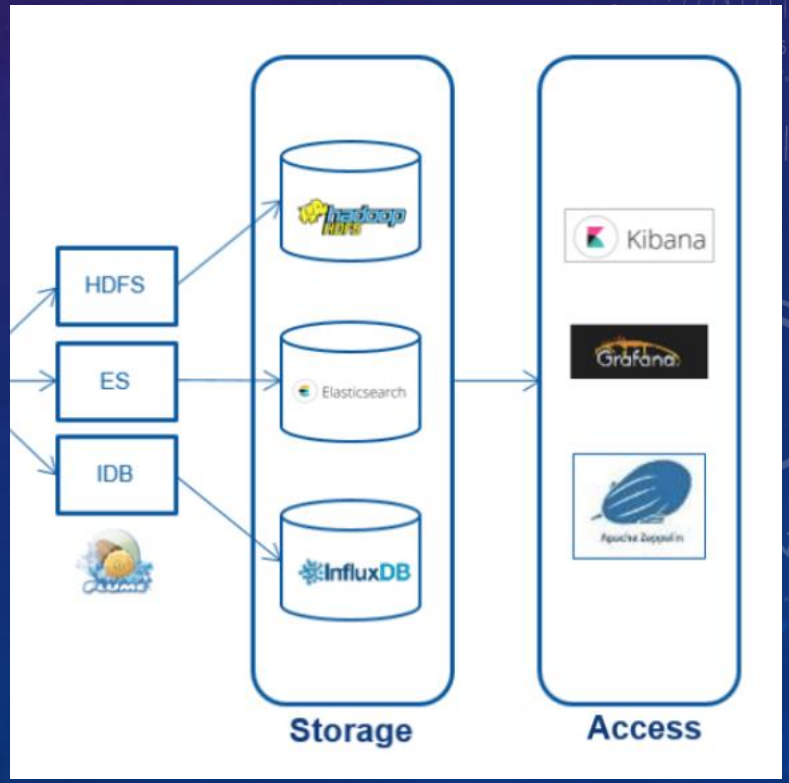

IMPLEMENTATION

- Request a new organization in monit-grafana
- Request the access right for the data sources we concern
  - monit_idb_eos, monit_idb_tape, monit_idb_transfers, monit_idb_collectd_XXX
- Request to create our own InfluxDB data source
  - monit_idb_amssoc
  - Modify/write scripts/programs to send data to the data source
- Modify/write dashboards for data access

EXAMPLE – EOS QUOTA

- Digging data from monit_idb_eos
- Singlestat to show the percentage of byte usage
- Graph to show the data and history of:
  - used_bytes and quota_bytes
  - used_files and quota_files
- Alert to warn when approaching the quotas

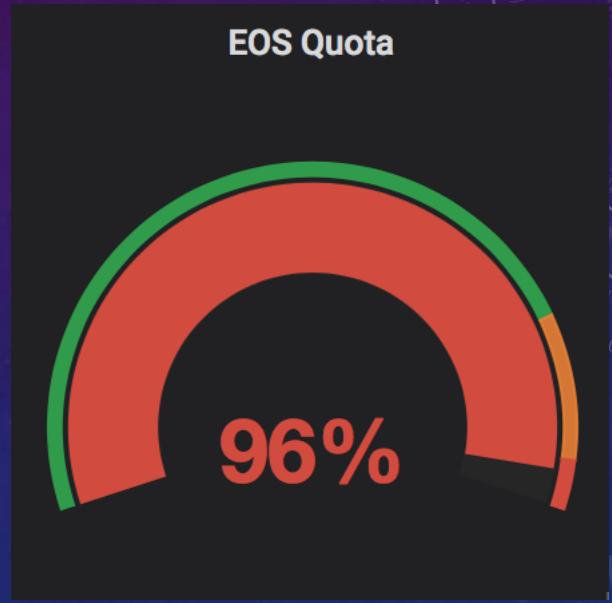

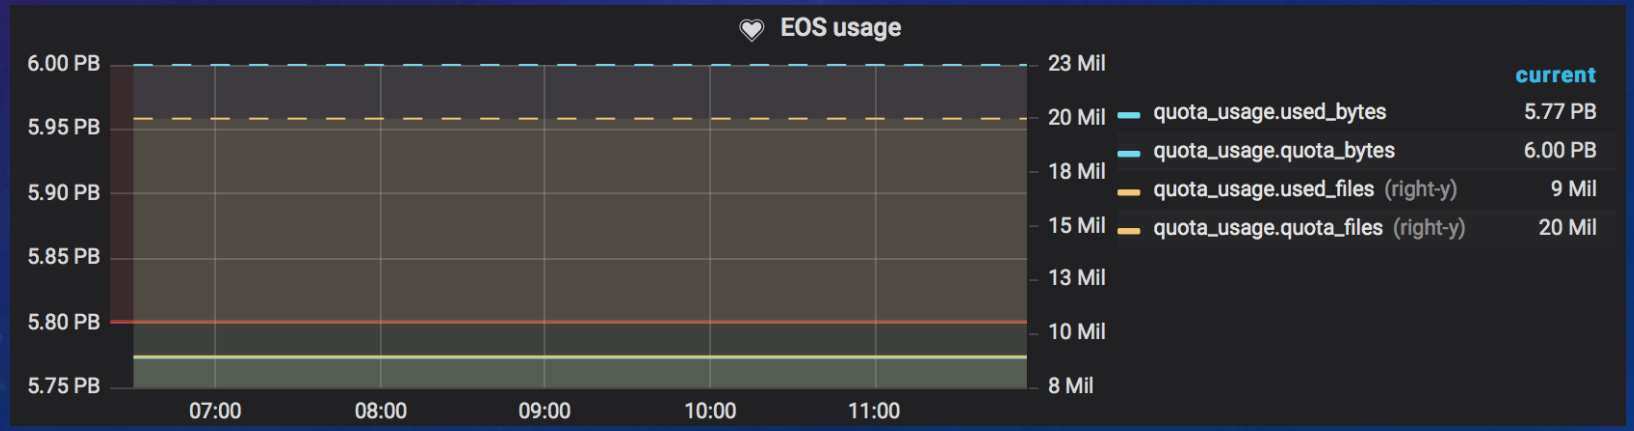

EXAMPLE – HOST LOAD

- Digging data from monit_idb_collectd_load: load
- Graph to show the data and history of the load of specific host(s)

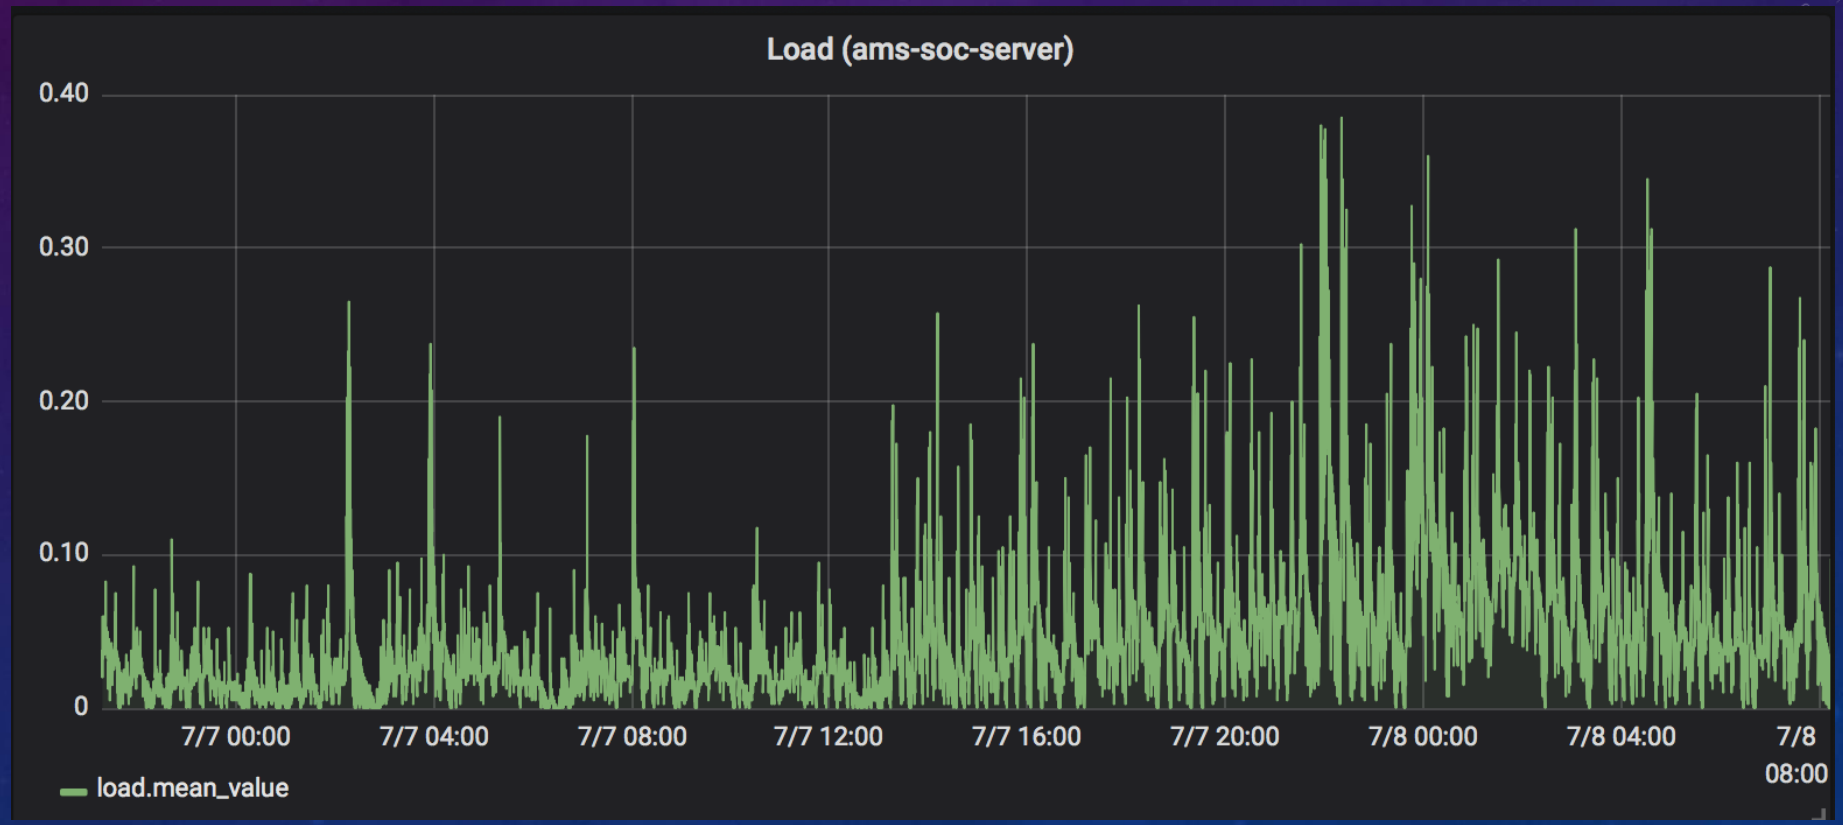

EXAMPLE – STANDARD PRODUCTION DELAYS

- Data format and structure
  - The last frame data arriving time
  - The last frame data collecting time
  - The last raw file validation time
  - The last ROOT file validation time
  - The latest reconstructed run collecting time
- A script running every 15 minutes to get the above data and post to http://monit-metrics:10012/


```
[{
    "producer": "amssoc",
    "type_prefix": "raw",
    "type": "metric",
    "timestamp": 1531077980243,
    "experiment": "ams",
    "frame_arr_delay": 32,
    "raw_delay": 144,
    "root_delay": 2075,
    "frame_delay": 19580,
    "fresh_run_delay": 12811,
    "idb_tags": ["experiment"],
    "idb_fields": ["frame_delay", "frame_arr_delay
", "raw_delay", "root_delay", "fresh_run_delay"]
}]
```

EXAMPLE – STANDARD PRODUCTION DELAYS (CONT.)

Data Source: monit_idb_amssoc | Options | Help | Query Inspector

A	SELECT last("frame_arr_delay") AS "Frame_arrival" FROM "metric" WHERE $timeFilter GROUP BY time($__interval) fill(none)
B	SELECT last("frame_delay") AS "Frame" FROM "metric" WHERE $timeFilter GROUP BY time($__interval) fill(none)
C	SELECT last("raw_delay") AS "RAW" FROM "metric" WHERE $timeFilter GROUP BY time($__interval) fill(none)
D	SELECT last("root_delay") AS "ROOT" FROM "metric" WHERE $timeFilter GROUP BY time($__interval) fill(none)
E	SELECT last("fresh_run_delay") AS "FRESHRUN" FROM "metric" WHERE $timeFilter GROUP BY time($__interval) fill(none)

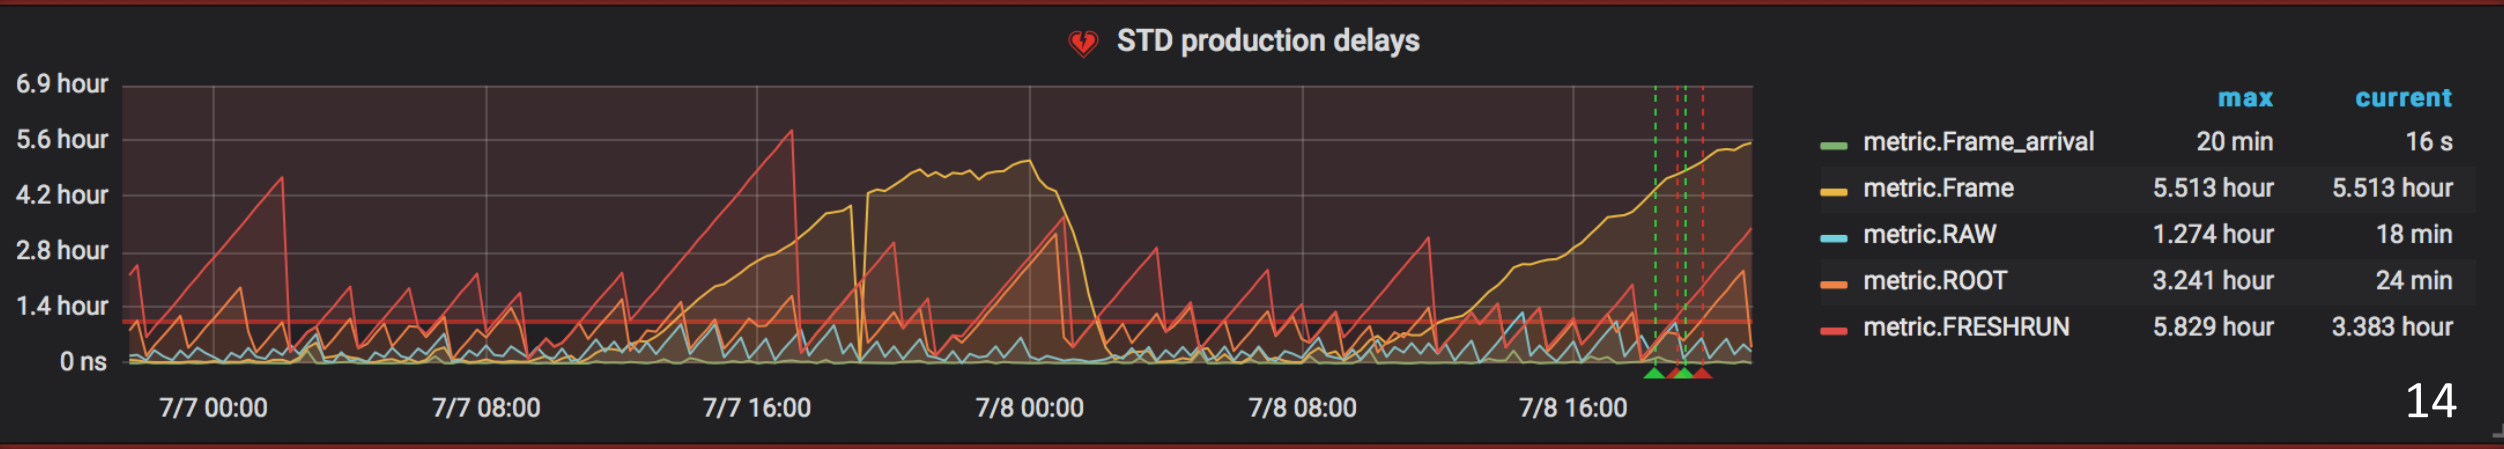

FURTHER WORKS

- Add more data sources:
  - Hardware: CPU/memory/disk/…
  - Batch service: under negotiation
  - MC production monitoring
  - Data from remote computing centers
- Add more alerts and connect to our FE
- Tune the dashboard structure

SUMMARY

- 7x24 data taking and processing brings more challenges for monitoring.
- SOC monitoring tools have been modified for the integration with MONIT infrastructure.
- Metrics Source provides a flexible way for our monitoring data ingestion.
- Grafana dashboard is used for visualization of the monitoring data from our own monitoring tools as well as from the CERN IT public services (EOS, Condor, FTS, etc.)

THANKS TO:

- Alberto Aimar, Simone Brundu, Pedro Andrade (monitoring)
- Herve Rousseau (EOS)
- Ben Jones (batch)
- Xavier Espinal Curull (IT contact)