CHEP 2024

Computing in High Energy and Nuclear Physics
October 19 – 25, 2024

Advanced monitoring capabilities of the CMS Experiment for LHC Run3 and beyond

Federica Legger[1], Nikodemas Tuckus[2], Brij Kishor Jashal[3]
INFN[1], CERN[2] ,RAL[3], TIFR[3]

On behalf of The CMS collaboration

CMS-monitoring-project

CMS Monitoring & analytics

150+
Worldwide LHC
Computing Grid sites
(T0, T1s, T2s and T3s)
(Storage &Compute)

Services utilizing this infrastructure

Statistics and metadata generated by these services

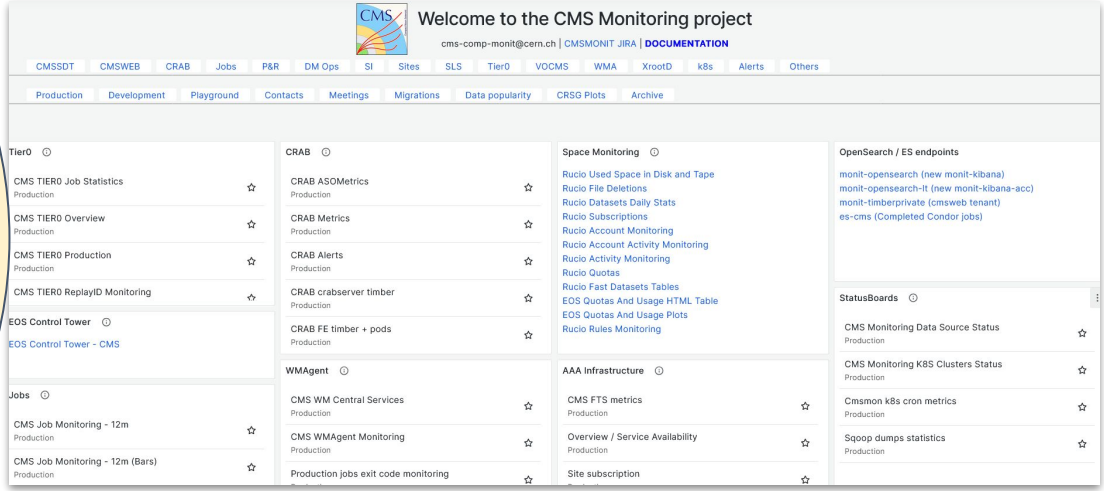

Requirements

Baseline requirement

- Growing datasets and services
- Near real-time monitoring
- Metrics for general and detailed dashboards
- Smart alert system
- Store historical information

Infrastructure and services

- Leveraging open source tools
- Scalable microservices architectures
- CERN-Monit backend.

What do we monitor

"In case you're wondering, yes, these are real names...

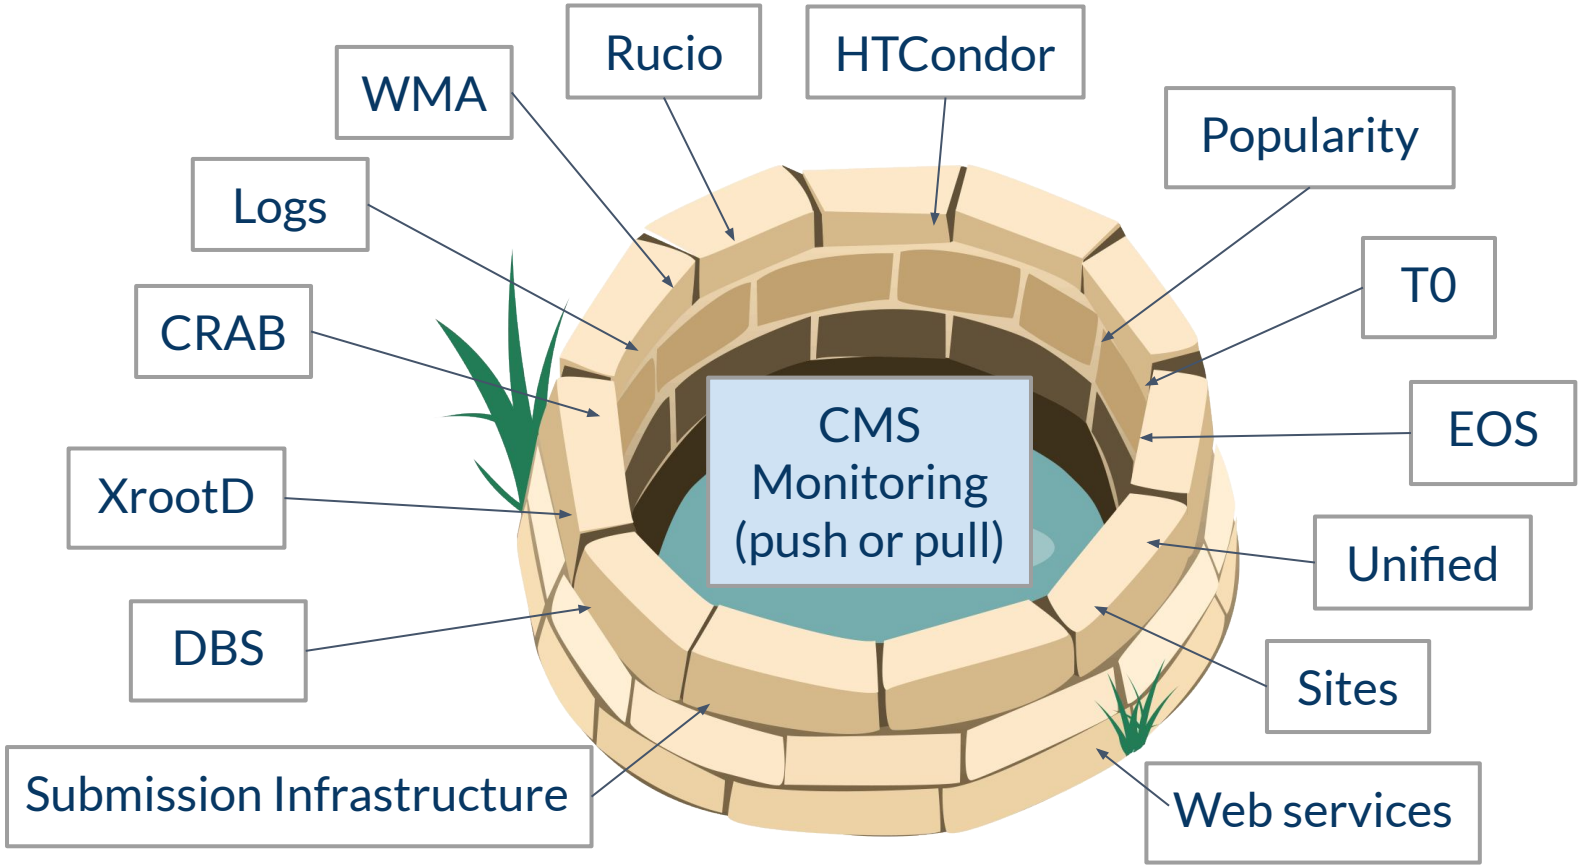

Data-flow: or 'death by a thousand acronyms!'

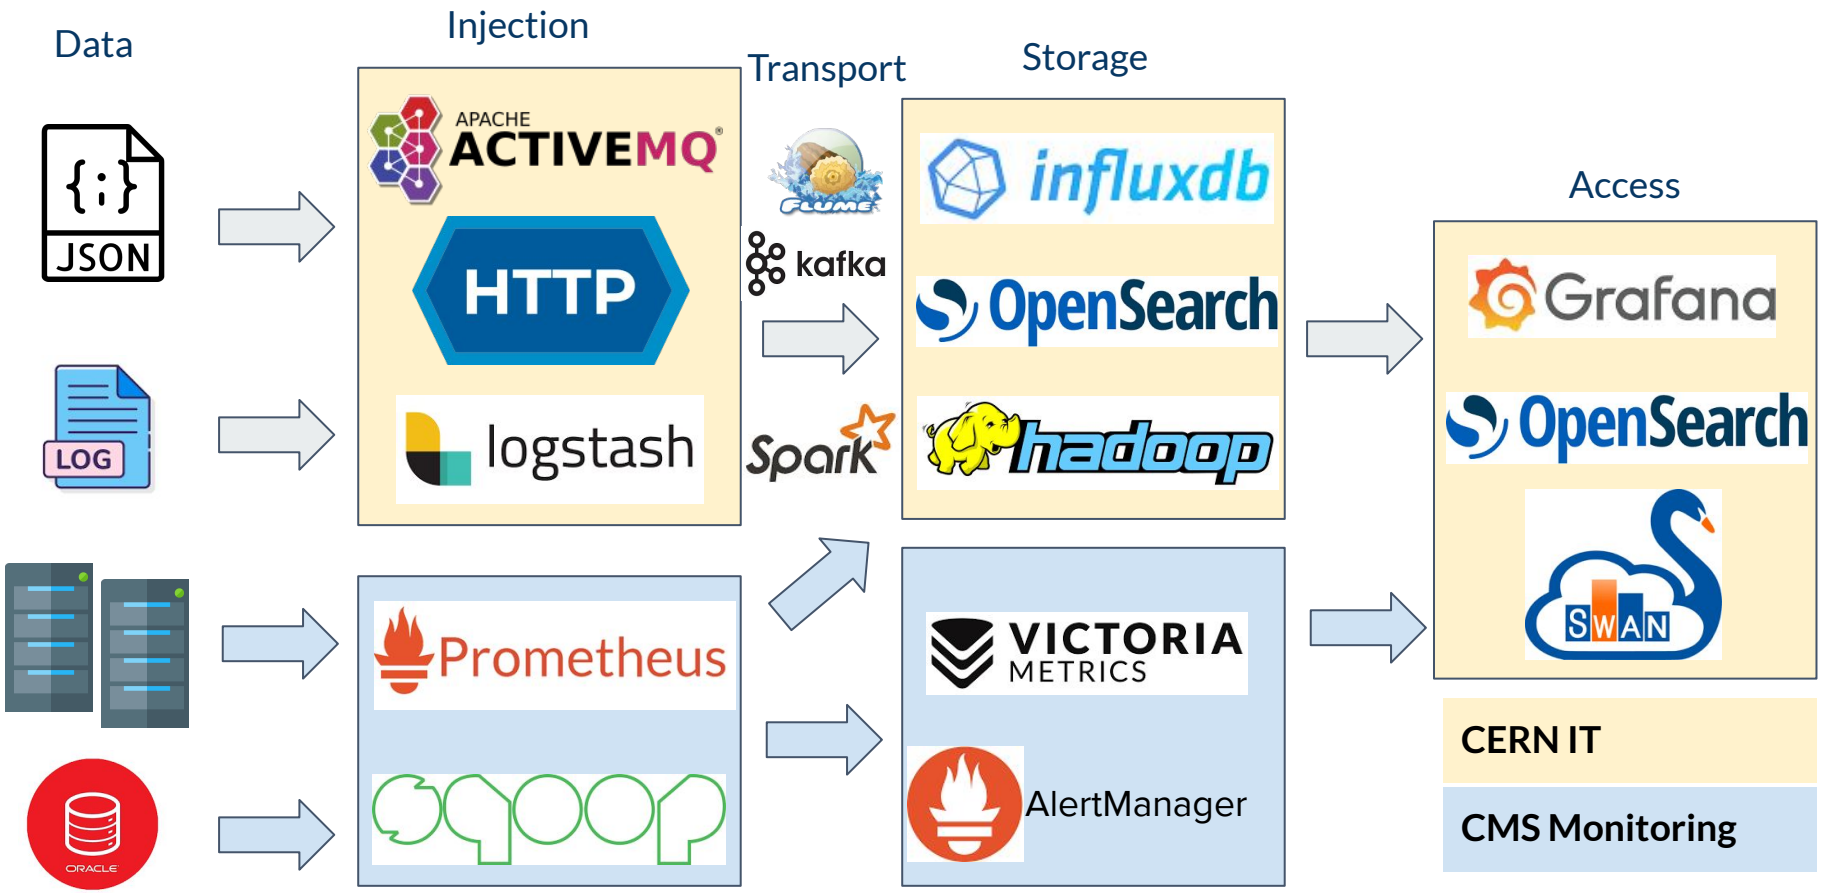

Infrastructure and services

- We manage:
  - 6+1 kubernates (k8s) clusters
  - 150+ Grafana dashboards in production (>750 overall)
  - 3x60 Billion data points
  - 10+ TB of data on OpenSearch
  - 55+ TB of compressed data on HDFS
  - 4-5 M msg/hour to AMQ brokers with 6.5 kHz
- We monitor:
  - 25+ k8s clusters
  - multiple services on VMs
  - 100+ services on k8s
  - 150+ alert rules

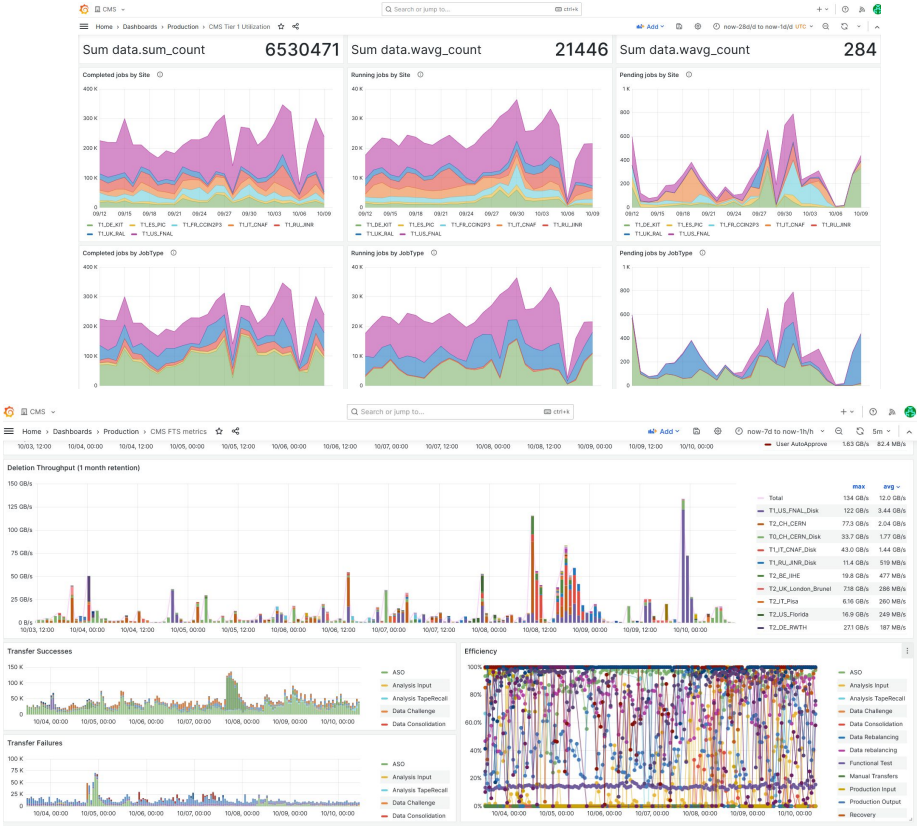

K8s Infrastructure

Robust infrastructure with zero downtime

- Prometheus (High Availability mode), VictoriaMetrics, AlertManager and Exporters: MAIN, HA-1, HA-2
- Single Grafana data-source to thousands metrics of Prometheus and VictoriaMetrics via Promxy
- Long-term time-series metrics storage for aggregated data: (2xVictoriaMetrics) VM-AGG
- Rucio dataset monitoring DM-MONIT
- Data aggregation/ETL batch jobs (Spark, Sqoop, etc.): CRONS

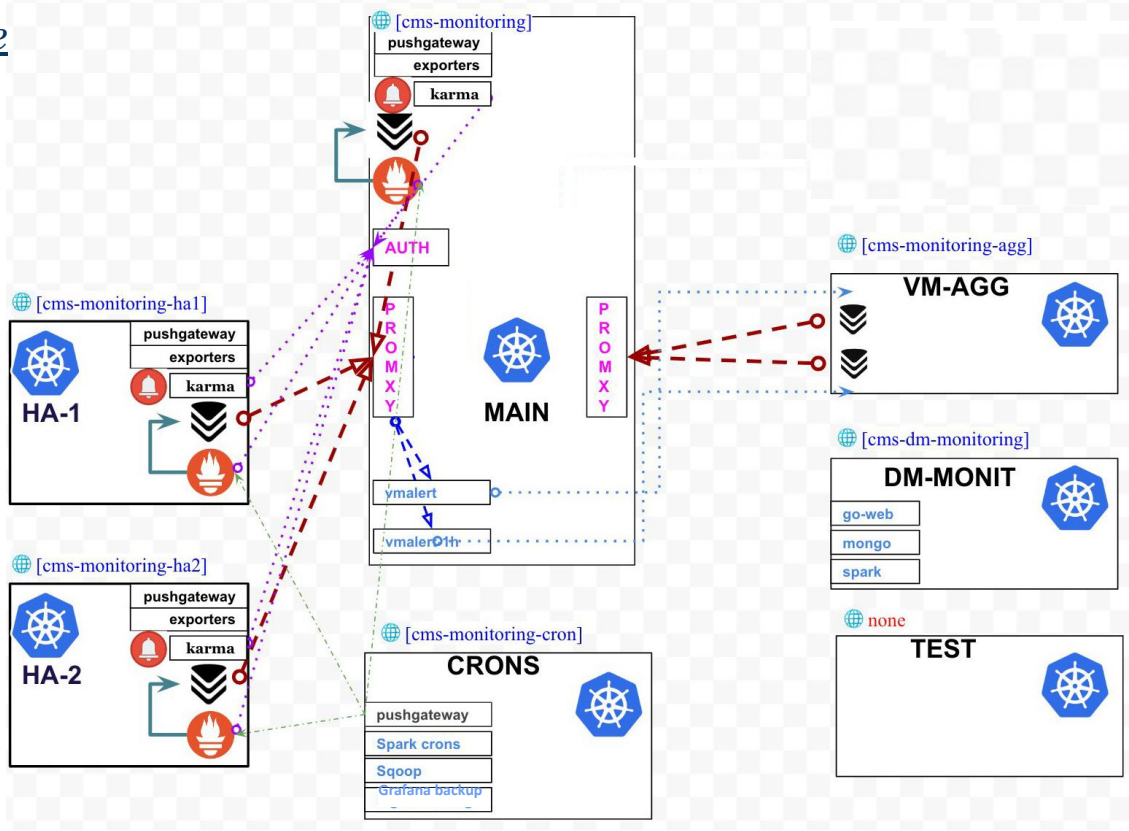

Rucio Dataset Monitoring

- Detailed view (tables) of all datasets (Name, Size, RSE, Creation date, Last access)
- Time-series monitoring of dataset metrics

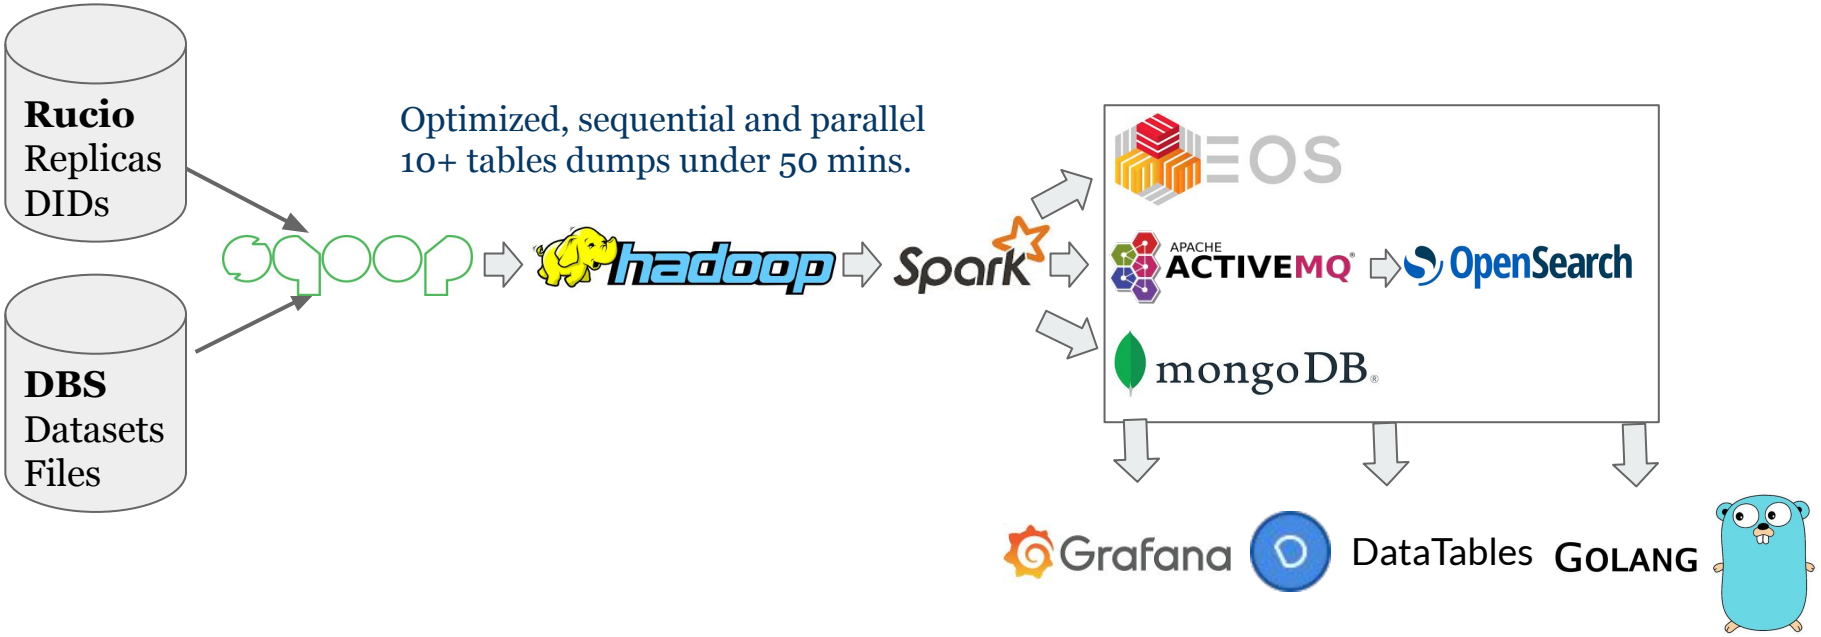

HTCondor job monitoring

The [CMS htcondor ES] (or spider) is a set of scripts which consumes CMS HTCondor ClassAds (Job processing docs), converts them to JSON documents and feeds them via [CMSMonitoring/ActiveMQ] module to CERN MONIT infrastructure.

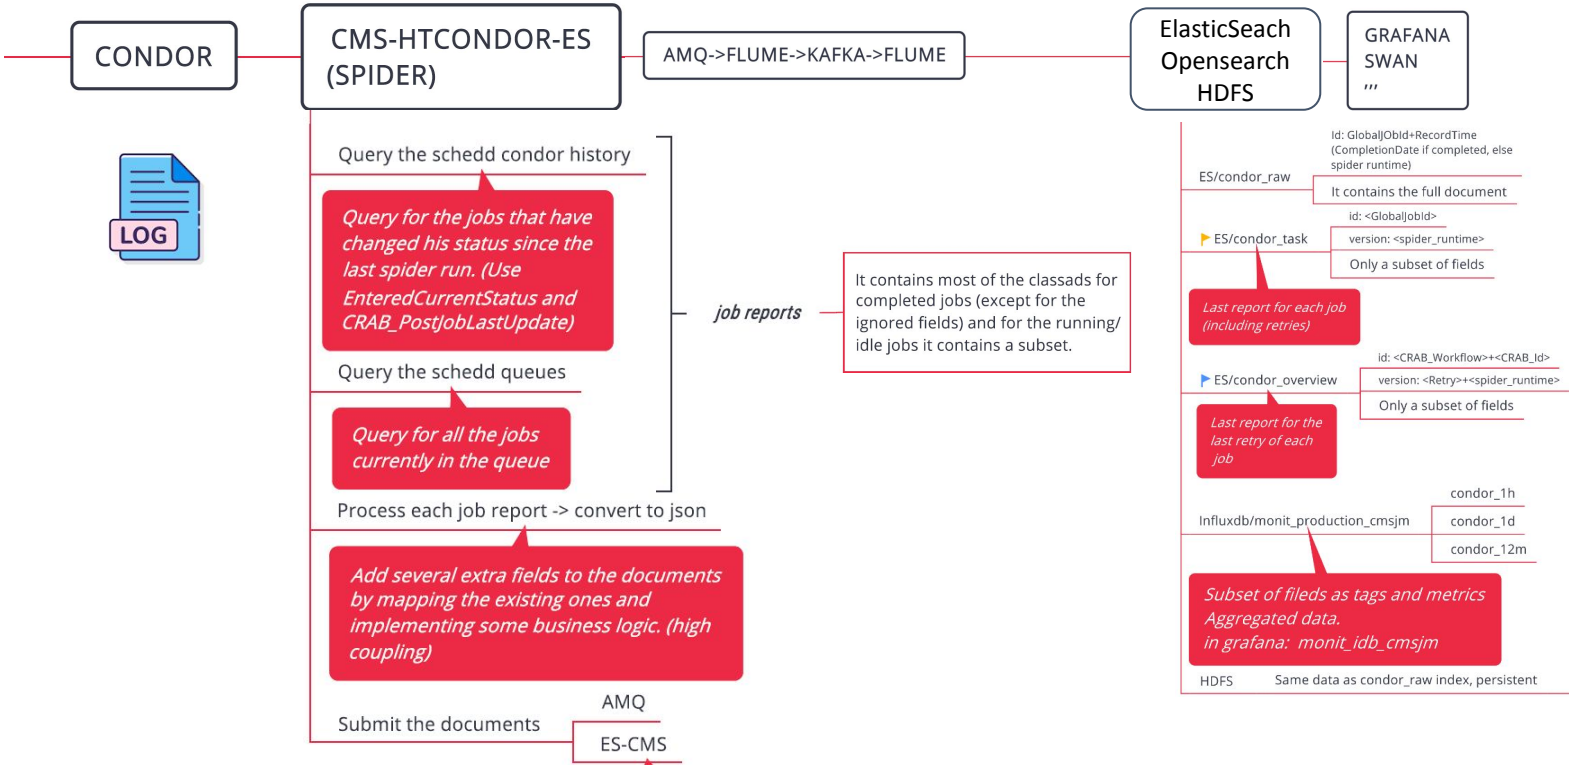

CMSSW popularity service

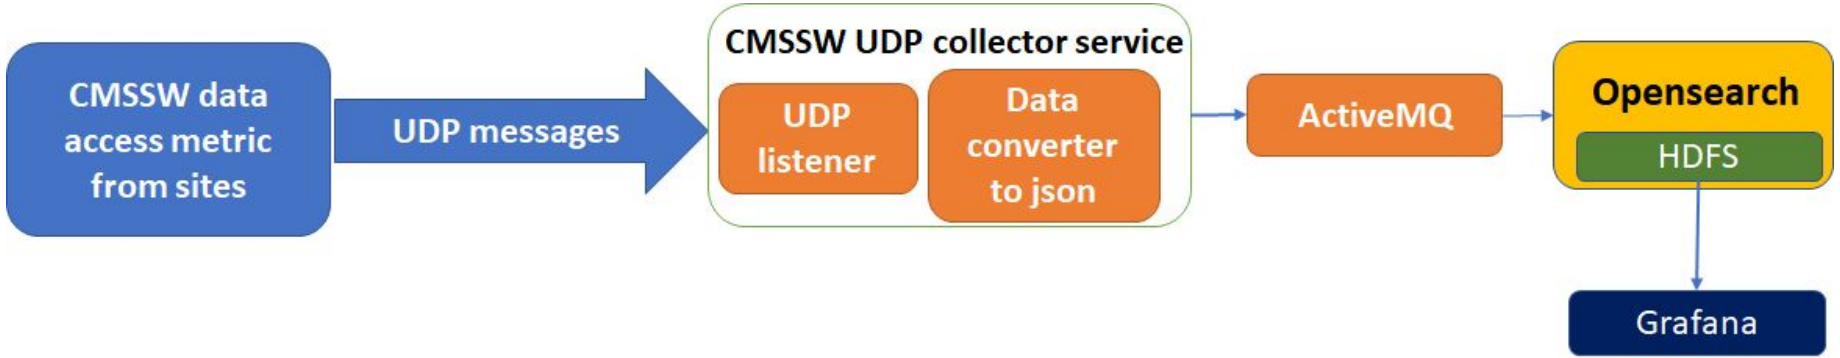

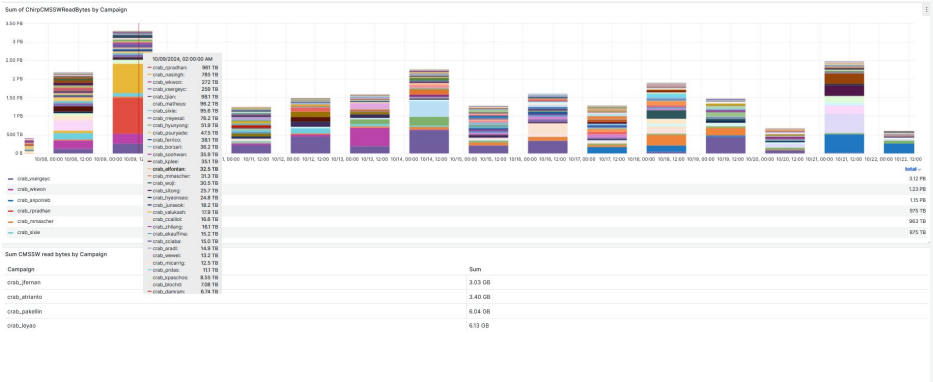

Infrastructure advancements

- Unified configuration management: Collected ownership of data pipelines from various users into a single service account, simplifying the system.
- Database Upgrade: Migrated from MongoDB to OpenSearch for improved data storage and query performance.
- Cronjobs Migration: Cronjobs (e.g sqoop jobs, spark jobs, backup jobs) from virtual machine to Kubernetes.
- Visualization Change: Transitioned from DataTables to Grafana for more advanced, user-friendly data visualization, allowing the creation of summarized charts.

Summary

- CMS Offline and Computing developed a comprehensive monitoring infrastructure
  - Use of common tools and technologies, keeping maintenance and operation tasks to minimal
  - Custom applications for specific use cases
  - System is performing well during Run-3 and evolvi for Run-4 and beyond.

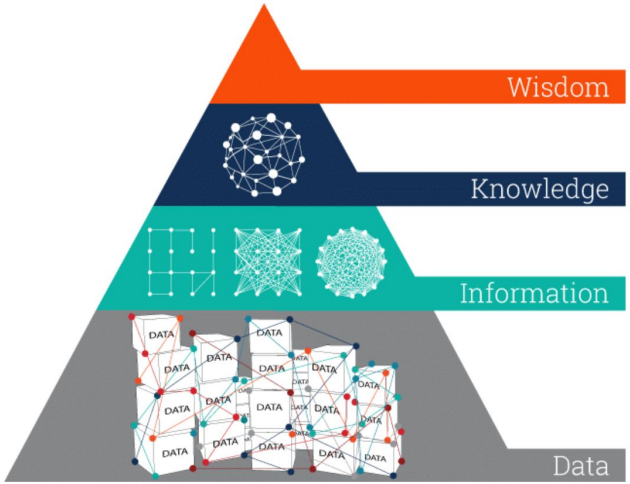

What gets monitored gets improved

Thank you