Unified Monitoring Architecture for IT and Grid Services

CHEP 2016 San Francisco 10-14 October 2016

Edward Karavakis on behalf of the CERN IT-CM-MM monitoring team



Monitoring

Data Centre Monitoring

- Monitoring of DC at CERN and Wigner
- Hardware, operating system, and services
- Data Centres equipment (PDUs, temperature sensors, etc.)
- Used by service providers in IT, experiments

Experiment Dashboards

- Sites availability, data transfers, job information, reports
- Used by WLCG, experiments, sites and users

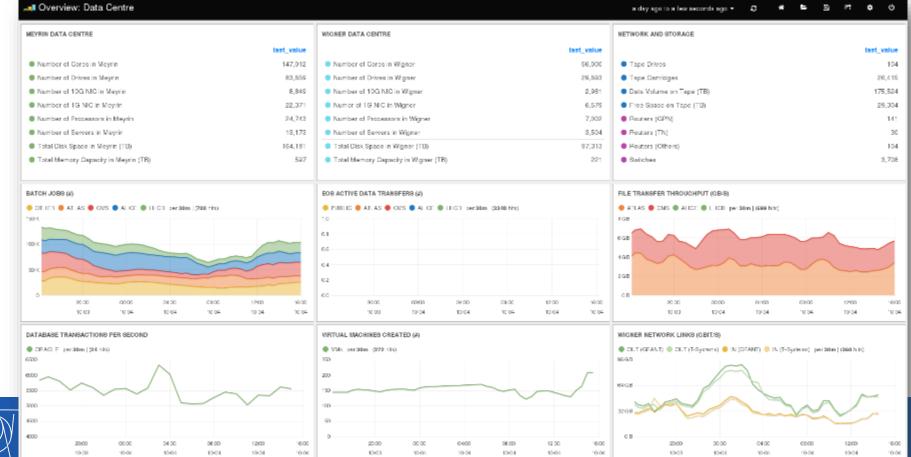
Both hosted by CERN IT, in different teams

Mandate for 2016

Regroup monitoring projects



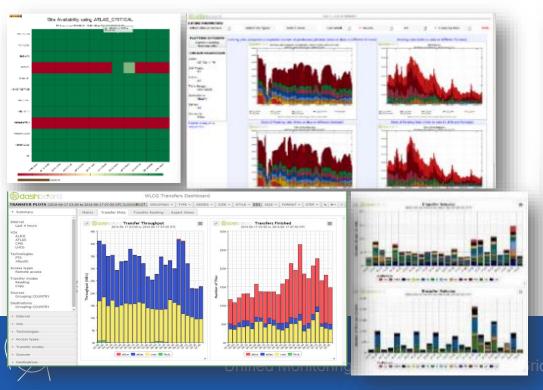
Data Centres Monitoring

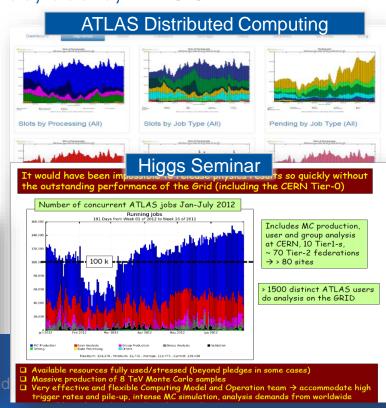




Experiment Dashboards

- Job monitoring, sites availability, data management and transfers
- Used by experiments operation teams, sites, users, WLCG





Current Monitoring

Data Sources Data Centres Monitoring **Metrics Manager Lemon Agent** XSLS **ATLAS Rucio** Data mgmt and transfers **FTS Servers DPM Servers XROOTD Servers** CRAB2 CRAB3 WM Agent Job Monitoring Farmout **Grid Control CMS Connect** PANDA WMS **ProdSys TZero Nagios** Infrastructure Monitoring **VOFeed** OIM **GOCDB**

REBUS

Transport Flume **AMQ** Kafka **Flume AMQ GLED HTTP Collector** SQL Collector **MonaLISA** Collector **AMQ** HTTP GET **HTTP PUT**

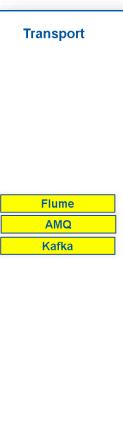
Storage & Search **HDFS ElasticSearch** Oracle **ElasticSearch HDFS** Oracle **ElasticSearch** Oracle **ElasticSearch Processing & Aggregation Spark Hadoop Jobs** GNI Oracle PL/SQL **ESPER** Spark Oracle PL/SQL **ES Queries ESPER**

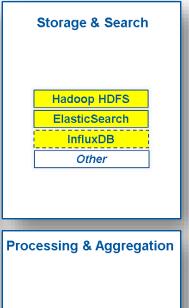
Display Access Kibana Jupyter Zeppelin Dashboards (ED) Kibana Zeppelin Real Time (ED) Accounting (ED) API (ED) SSB (ED) SAM3 (ED) API (ED)

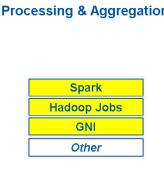


Unified Monitoring













Unified Monitoring Architecture

Data Sources

> Data Centres

WLCG

Transport





& kafka Storage/Search



Elasticsearch

Processing



Data Access





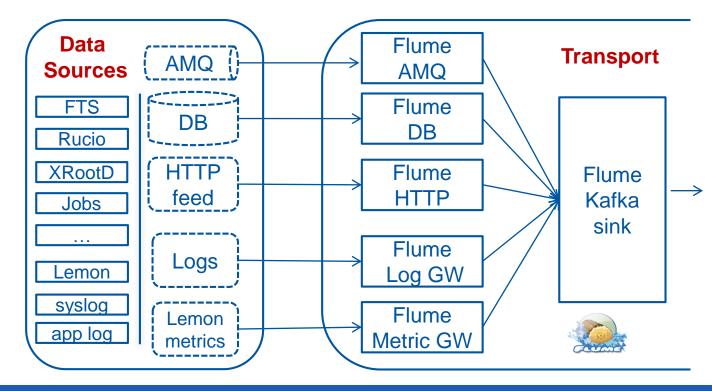




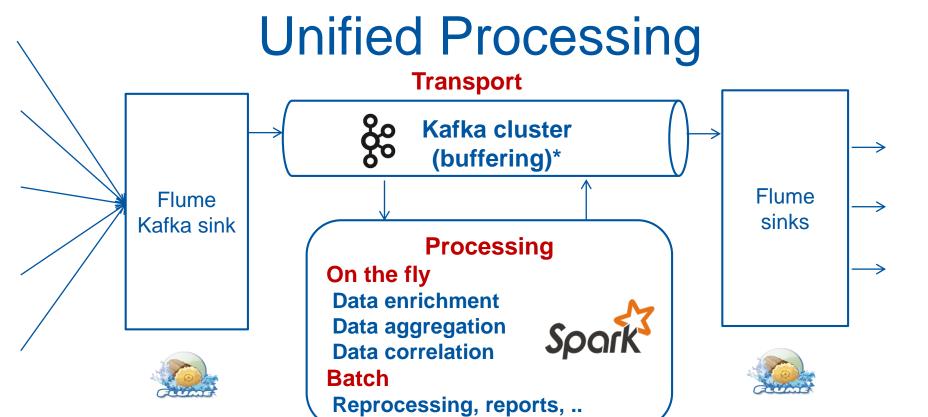


Unified Data Sources

- Data is channeled via Flume, validated and modified if necessary
- Adding new Data Sources is documented and fairly simple



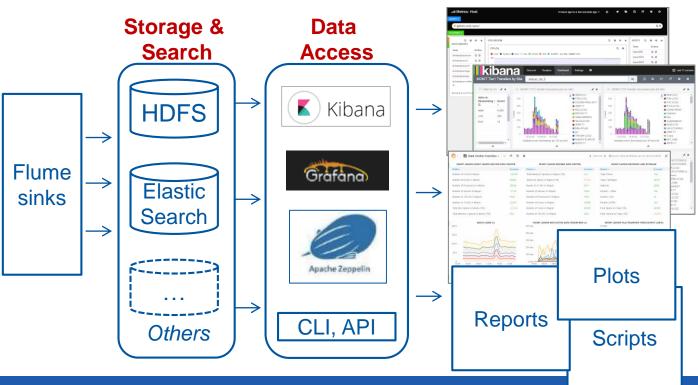








Unified Access



- Multiple data access methods (dashboards, notebooks, CLI)
- Mainstream and evolving technologies

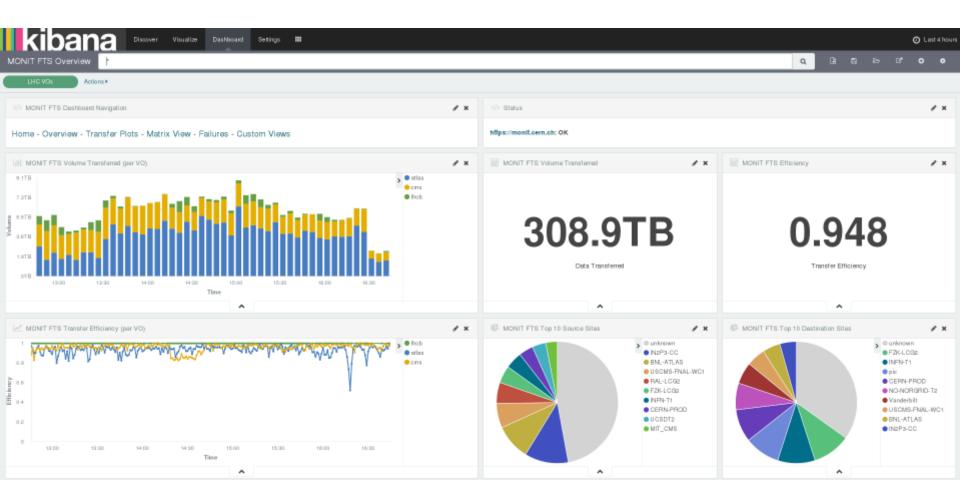


Status since January

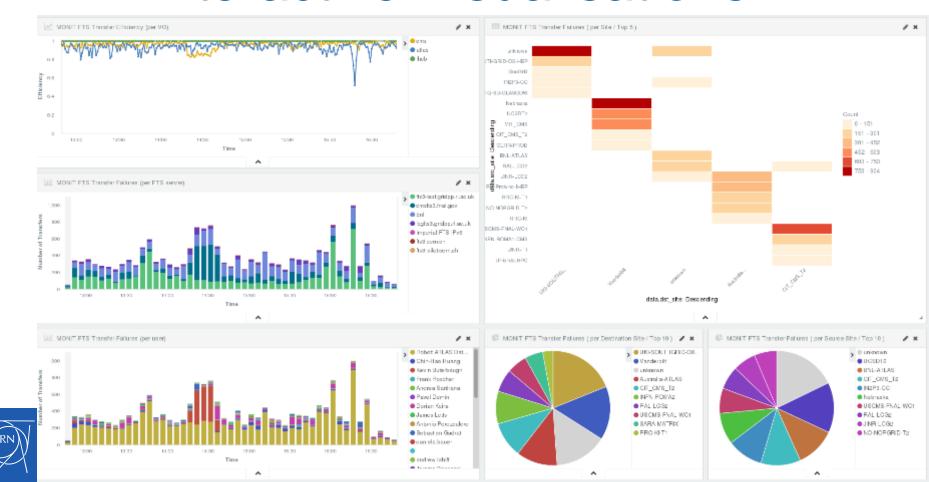
- Data size ~200 GB/day (~500 GB/day at full scale)
- ~200M documents/day with spikes at 10KHz
- Relying on VMs complaint to the CERN IT AI standards (Openstack/Puppet)
- 20 nodes Kafka cluster
- 23 nodes Flume input and output gateways
- 9 nodes Mesos/Spark cluster
- Scaled down version for QA environment
- Centrally provided by IT
 - Dedicated Elasticsearch cluster: ~30 nodes
 - Shared Hadoop/HDFS cluster: ~40 nodes



Interactive Visualisations



Interactive Visualisations



Conclusions

Merged Monitoring projects

Same team, technologies, practices and infrastructure

Unified monitoring infrastructure for CERN Data Centres and WLCG

- Moving all data via new transport (Flume, AMQ, Kafka)
- Data in ES and Hadoop (and soon in InfluxDB)
- Aggregation and processing via Spark
- Visualization with Kibana, Grafana, Zeppelin

Service Proposed

- Collect, process, aggregate, visualise and raise alarms
- Cover metrics and logs
- Operate and scale the infrastructure
- Support interfacing new data sources, custom processing, building dashboards and reports



Reference and Contact

Dashboard Prototypes monit.cern.ch

Feedback/Requests (SNOW) cern.ch/monit-support

Early-Stage Documentation cern.ch/monitdocs



Backup slides



Monitoring Processing Platform

- Reliable and scalable job execution (Spark)
- Job orchestration (Marathon / Chronos)
- Lightweight deployment (Docker)

