

Unified Monitoring Architecture for IT and Grid Services

CHEP 2016 San Francisco
10-15 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

Current Monitoring

Data Centres Monitoring

Data mgmt and transfers

Job Monitoring

Infrastructure Monitoring

Data Sources

Metrics Manager

Lemon Agent

XSLS

ATLAS Rucio

FTS Servers

DPM Servers

XROOTD Servers

CRAB2

CRAB3

WM Agent

Farmout

Grid Control

CMS Connect

PANDA WMS

ProdSys

TZero

Nagios

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

Data Sources

Metrics Manager

Lemon Agent

XSLS

ATLAS Rucio

FTS Servers

DPM Servers

XROOTD Servers

CRAB2

CRAB3

WM Agent

Farmout

Grid Control

CMS Connect

PANDA WMS

ProdSys

TZero

Nagios

VOFeed

OIM

GOCDB

REBUS

Transport

Flume

AMQ

Kafka

Storage & Search

Hadoop HDFS

ElasticSearch

InfluxDB

Other

Processing & Aggregation

Spark

Hadoop Jobs

GNI

Other

Data Access

Kibana

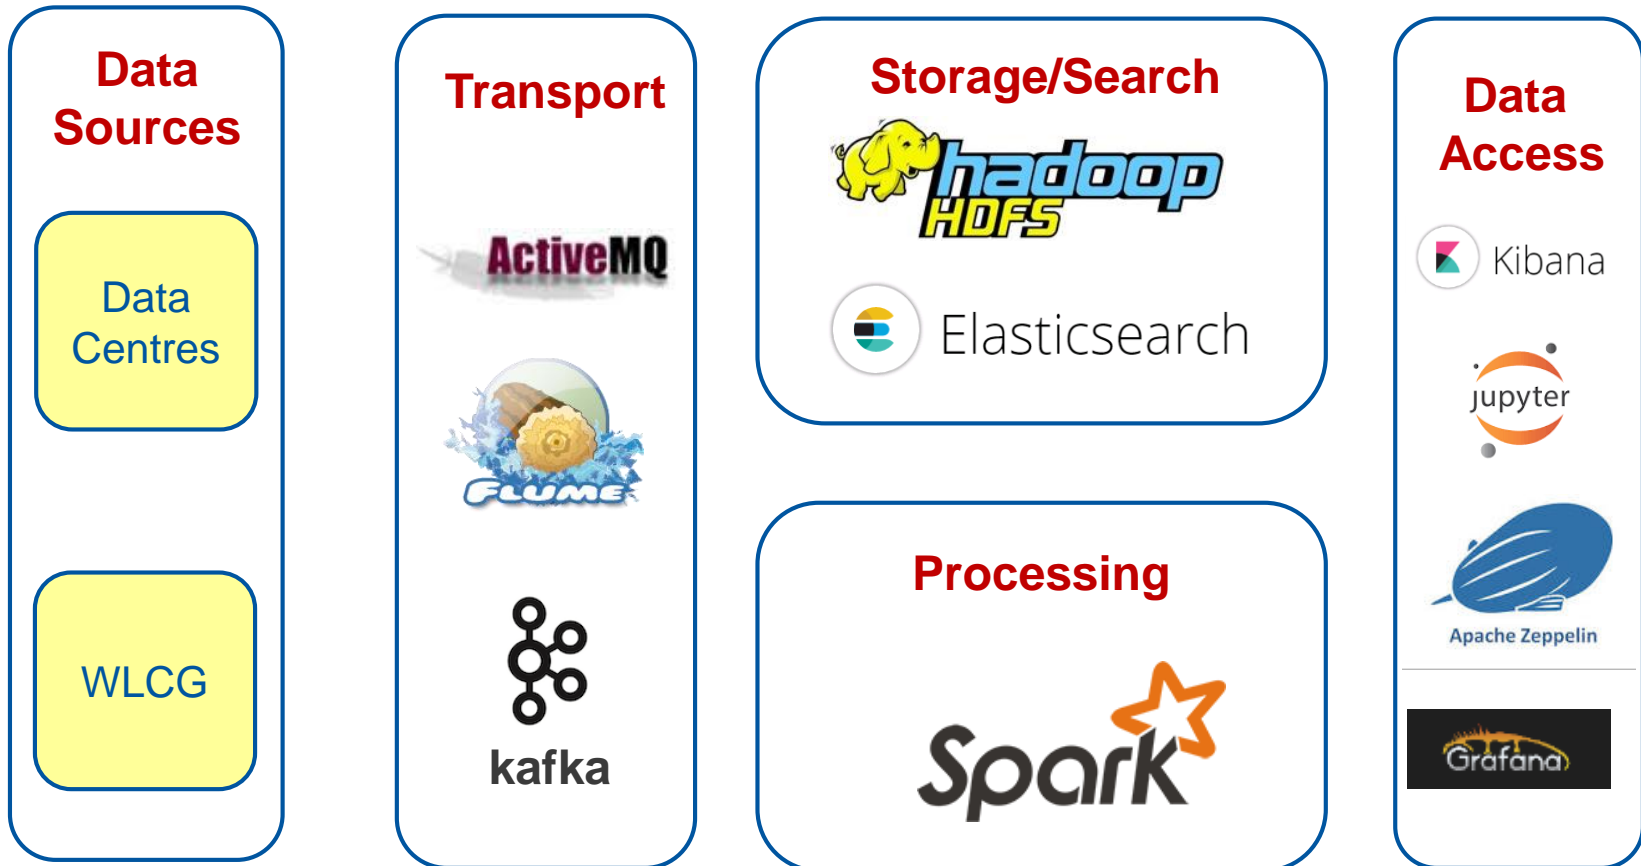
Grafana

Jupyter

Zeppelin

Other

Unified Monitoring Architecture



Conclusions

Working on a common monitoring infrastructure for CERN Data Centres and WLCG

- Moving all data via new transport (Flume, AMQ, Kafka)
- Data in ES and Hadoop (and soon some in InfluxDB)
- Doing aggregation and processing via Spark
- Using only standard features of ES, Kibana, Spark, Hadoop
- Introduce notebooks (e.g. Zeppelin) and data discovery
- Still work in progress, already available

Helping and training on:

- Interfacing new data sources
- Developing custom processing, aggregations, alarms
- 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