IT-SDC
Data Analytics: metrics & plans

Data Analytics WG
10/09/2014

Data Sources

WLCG Monitoring (4 VOs)

- Job processing dashboards
  - 0.8 M - 1.5 M records/day/VO
  - Metrics: job metrics (CPU, wall clock, exit code,...) meta data (time,user, application, site, ...)
- Data Transfer Dashboards
  - FTS, XRootD: ~8M records/day
  - Metrics: transfer metrics (bytes read/write, vector reads,...), meta data (time, user, lfn, ...)
- Site Status
  - Large variety of metrics (direct/derived) to compute site availability/reliability

	ALICE	ATLAS	CMS	LHCb
Site Status	VO SSBs, WLCG Mon	VO SSBs, WLCG Mon	VO SSBs, WLCG Mon	VO SSBs, WLCG Mon
Job Processing	MonALISA	Job Dashboard	Job Dashboard	Dirac
Data Transfers	WLCG Transfer Dashboard	WLCG Transfer & DDM Dashboards	WLCG Transfer Dashboard	WLCG Transfer Dashboard

Already in use for first efficiency studies

Data Analytics Use Cases

Extract value from the large volume/variety of monitoring data

- Profile the site performance w.r.t the various experiment workflows
  - production (MC prod., reco), analysis, local access, remote access, ...
- Compare performance
  - Among sites (eg. CERN-Meyrin Vs Wigner)
  - Among different HW, SW releases, real/virtual cores
- Improve efficiency and reliability of services
  - Find patterns, correlate, predict

Initial investigation

Study job performance CERN Meyrin Vs Wigner

- Contributors (Valentina M., Eddie K., Alessandro DiG., Nicolo' M. Maarten L., Stefan R., Daniele S., D.G.)
- Study distributions and trends of various quantities: Tcpu, Twc, Events, ... and their combination
  - Data analysis still ongoing

By product: extract from Dashboard and make available in ElasticSearch a set of metrics per jobs (currently only for ATLAS)

- NEventsProcessed, WrapCPU, WrapWC, StartedRunningTimeStamp, FinishedTimeStamp
- SchedulerJobId, VOName, GridName, TaskName, TYPE, JOBTYPE, STATE
- Batchid, WNHostName, WNIp, isWigner, isSLC6, isPhysical, isIntel, isAmd
- HammerCloud benchmark workflow id

Data insertion in ES for the past months is ongoing (Kibana test url)

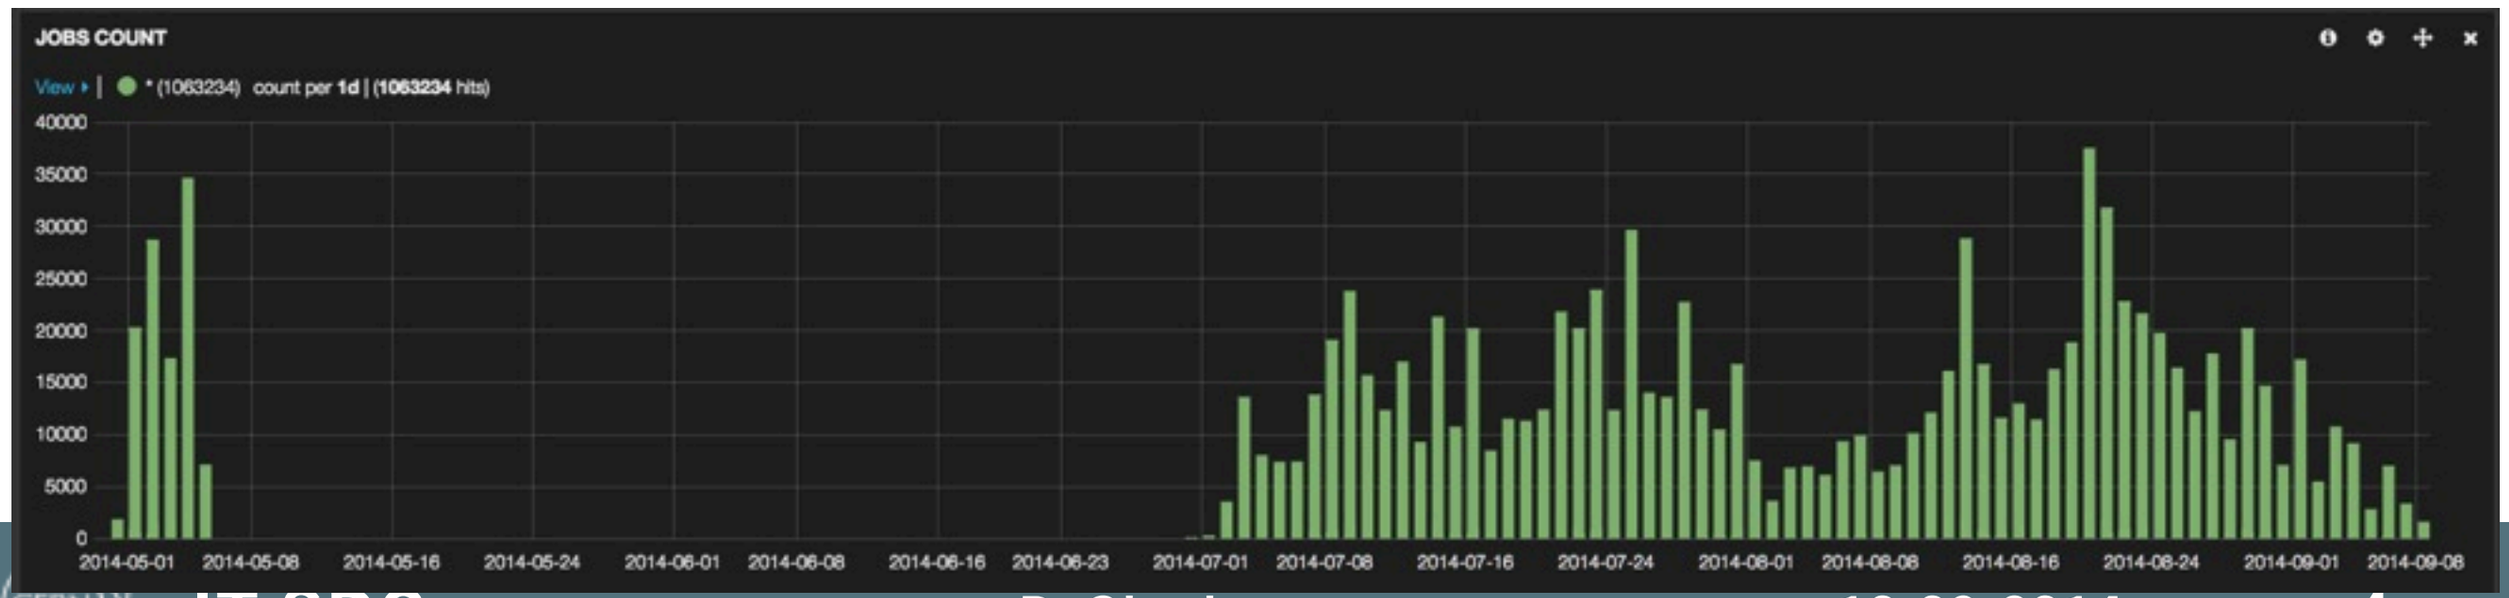

Common needs

Extend data coverage

- Include job monitoring information for other VOs (CMS, ALICE, LHCb)
- Include metrics from batch, storage and network systems
- Join data sources
  - Job info ⊕ File access info ⊕ service metrics (batch, storage, network systems)
  - are keys already available?
    - ✴ Batch system information could be linked by batch-id
    - ✴ What about storage, network, perfSonar, system logs

Reduce operation cost for data harvesting - extraction.

- common data format
- unified procedures for data extraction and aggregation (log parsing, map-reduce, data access API)
- Validate collected metrics
  - Are Nevents properly reported by dashboard? Is CPU Time reported by dashboard compatible with batch monitoring?
- Calibrate metrics (e.g. normalize CPU performance by HEP-Spec)

Analytics for services

WLCG monitoring today is based on well-defined workflow:

- collect monitoring raw data (mainly in OracleDB); extract and aggregate metrics by features; provide views (aka result) to users

Monitoring tools and services need an analytics platform to be built upon:

- scalable with data volume
- reliable and fault-tolerant
- flexible in data processing (e.g. with configurable latency)
- integrated with the most common data analytics libraries

SDC/MI started to work in this direction (twiki)

- HDFS/MapReduce/Avro for batch-processing
- in-memory stream analysis (e.g. Esper)