



MONIT

Monitoring CERN Data Centres and WLCG Infrastructure

Alberto Aimar (CERN IT/CM)
for the MONIT team

History

Widely-used monitoring use cases at CERN had own solutions and teams

LEMON and SLS for Data Centres and IT services

- Hardware, operating system, and services
- Data Centres equipment (PDUs, temperature sensors, etc.)

WLCG Monitoring dashboards and tools

- Jobs information, data transfers and sites availability
- Used by WLCG, experiments, sites and users

Mission

Provide a Unified Monitoring as a Service (MONIT)

- CERN Data Centres and the WLCG collaboration
- Dashboards, Alarms, Search, Archive

Collect, transport, store and process metrics and logs for services, applications and infrastructure

Migrate data, users and features to a new Monitoring infrastructure

Change Technologies

- Moving to established open source technologies
- Reduce in-house components when possible
- Collaboration with open source projects:
 - Popular/trendy inside our community
 - Development, send bug reports...



MONIT Architecture

Data Sources
100 types
5M doc/min

Transport
72h buffer
separate topics

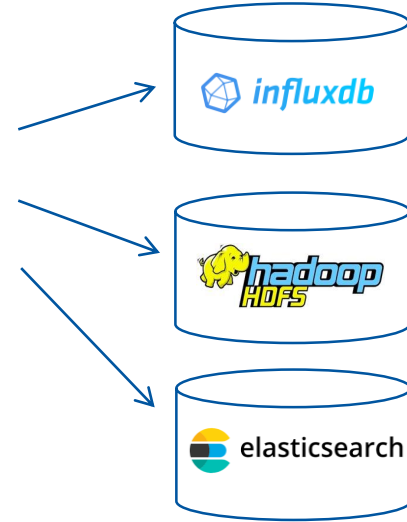
Storage
~3TB/day
compressed

FTS
Rucio
Panda
Condor
OpenStack
Data Centre
Logs
Alarms
Xrootd
Batch
VOfeed
AGIS
Network
Puppet
Metrics

AMQ
JDBC
HTTP
AVRO



Processing
20 jobs



Kafka and Spark

Apache Kafka as rock-solid core data pipeline

- Distributed message broker
- Enables stream processing
- 72h retention policy, replica factor: 3

Apache Spark for stream and batch data processing

- Enrichment: join information from several sources (e.g. network topology)
- Aggregation: over time (e.g. summary statistics for a time bin) or other dimensions (e.g. cumulative metric for a same service)
- Correlation: detect anomalies and failures correlating multiple sources
- Batch processing: reprocessing, compression, historical data, periodic reports



kafka



Storage

ES for data exploration and discovery

- 2 large instances: 30 nodes
- Raw data, including metrics and logs
- Data kept for 1 month

InfluxDB for time-series dashboards

- > 20 instances, from 8GB to 128GB memory per instance
- Data kept for ~5y, down-sampling
- Aggregated (1w/raw, 1M/10m, 5Y/1h)

HDFS for long-term archive

- Compressed JSON or Parquet
- Raw data, metrics and logs



elasticsearch



Visualization and Access

Grafana for dashboards

- Open-source platform for dashboards
- Support multiple back ends: ES / InfluxDB
- Templates, ad-hoc filters, auto-completion
- Advanced queries, organizations, ACLs, Alarms
- Users can create and admin their own dashboards

Kibana for raw data exploration

- Data discovery and logs search

Swan for analytics (notebooks)

- Use data in HDFS and run on the Spark infrastructure



Grafana



Kibana



Monitoring the WLCG Infrastructure

Monitoring of WLCG services at CERN

- Transfers (FTS and Xrootd) and Distributed Data Management
- Job Monitoring and Accounting
- Sites Availability
- + new sources or fields (LHCOPN network, IPV6 info, etc.)




Data from the WLCG and Experiments

- Same data sources of old WLCG Dashboards, without exceptions and hard-coded fixes
- Normalize, enrich and aggregate data (same naming, added topology, hourly bins, etc.)
- Model changed to have users and experts able to develop dashboards and reports

Same technology and dataflow for all WLCG and Experiment sources

- Data collected, transported and stored for search, plots and reports
- Standard dashboards, alarms and processing
- Dashboards by default, can be changed for specific roles/tasks (shiffters, experts, sites, etc.)

WLCG Transfers (5 Years)



Group by **dst_cloud** ▾

Binning **auto** ▾

Activity Analysis Input + Data Brokering + Data Consolidation + Data Rebalancing + Deletion + Express + Functional Test + Production Input + Production Output + Recovery + Staging + T0 Export + T0 Tape + User Subscriptions + default ▾

Source country **All** ▾

Source site **All** ▾

Destination country **All** ▾

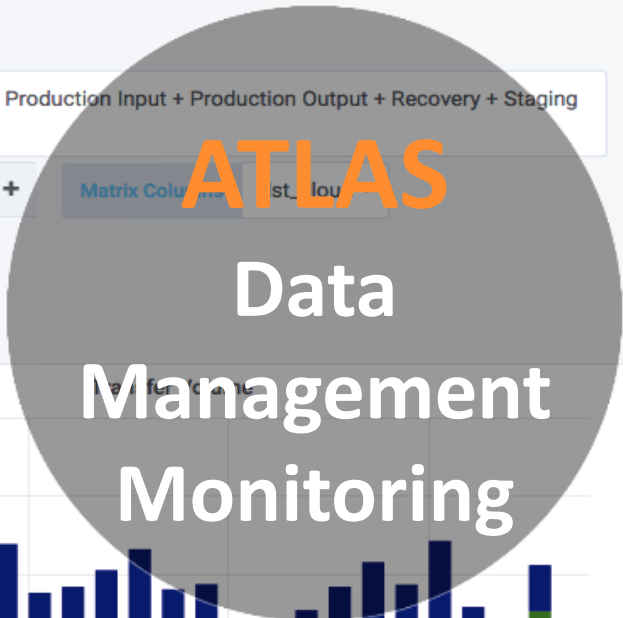

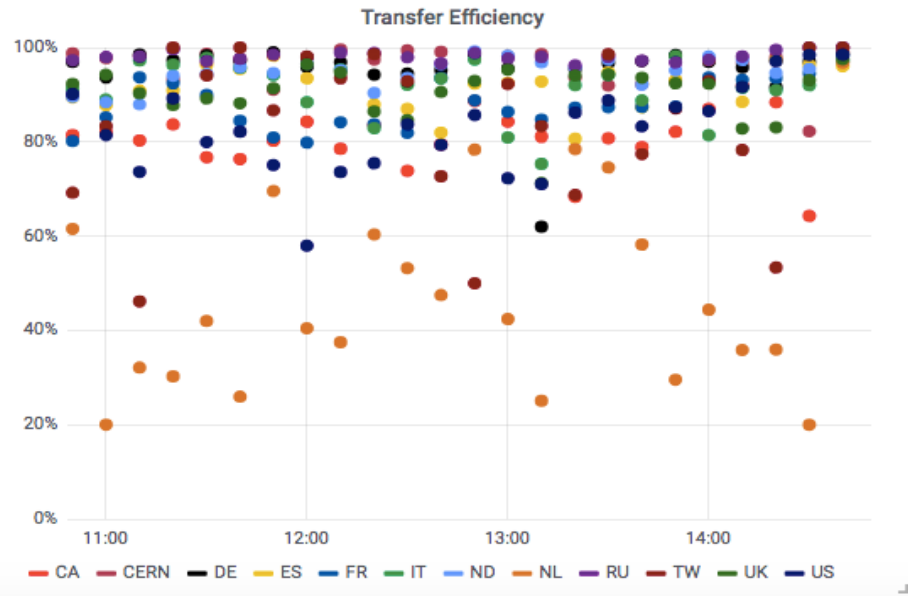
Destination site **All** ▾

Filters +

Matrix Columns **dst_cloud**


Matrix Rows **src_cloud** ▾

Transfers ▾




Group by CMSPrimaryDataTier Tier All Type All CMS Job Type All Primary Data Tier All Binning 1h Filters +

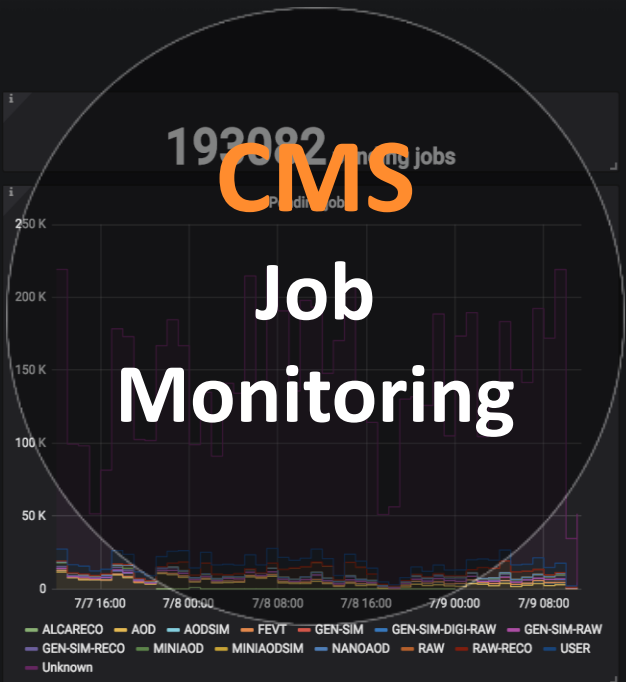
615596 completed (last 24h)




60037 running jobs




193982 pending jobs




Total completed jobs




Running jobs




Currently pending jobs




Completed cores



Running cores



Pending cores



Group by resourcesreporting ▾

Site All ▾

Cloud All ▾

Country All ▾

Fed All ▾

JobType All ▾

Resources Reporting All ▾

Cores All ▾

Resource Type GRID ▾

Groups All ▾

Input Data All ▾

Event Serv All ▾

Input Proj All ▾


Output Proj All ▾

Binning 1h ▾




Overall Status ▾

Submitted jobs




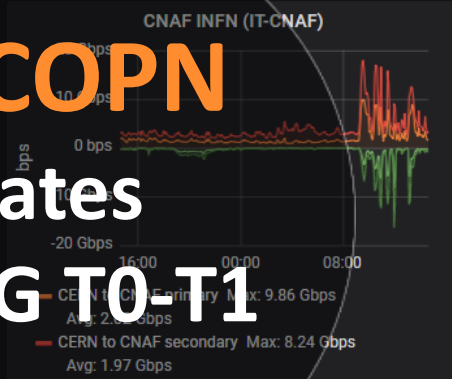





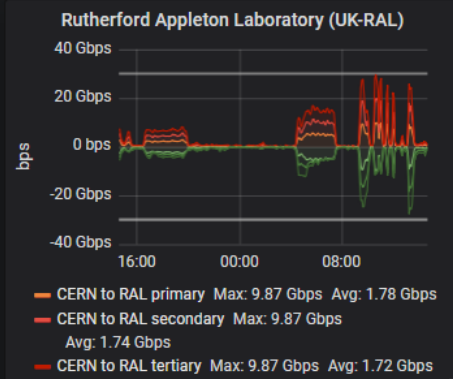



Analysis Data Processing Group Production MC Reconstruction MC Simulation Others

Running jobs



Analysis Data Processing Group Production MC Reconstruction MC Simulation Others



LHCOPN
Rates
WLCG TO-T1

Monitoring Data Centres



Pre-configured data flow for IT Service Managers

- Data transported and stored
- Standard alarms and processing
- Dashboard for host and services metrics



Collectd deployment for CERN Data Centres

- Core technology chosen for all DC metrics
- Installed on all “puppetized” hosts and VMs of the DC
- Service Manager add/develop plugins


DC Logs

DC Metrics:

- OS metrics
 - Hardware metrics
 - Application metrics
- > 38k machines
> 80kHz rate

Monitoring of IT services at CERN

- Support for custom metrics and logs
- Used to monitor main IT services
e.g. batch, cloud, storage



Average Number of Jobs Running Concurrently

100544


Total Job Runtime (Last Week)

2876 years

Average Efficiency (Last week)

81%

Running Slots By Experiment



CERN BATCH
Full
Overview

CPU Delta (7 days)

190

Memory Delta (7 days)

978 GiB


Disk Delta (7 days)

6.72 TiB


Efficiencies by DataCentre (Last Week)

Metric ▾	Min	Max	Avg	Current
wigner	55.85%	86.39%	74.37%	77.32%
meyrin	61.28%	87.41%	81.05%	82.57%



Cloud resources



Openstack services stats



Resource overview by time




CERN CLOUD





OpenStack Overview

Average VM boot time

- p50 without DNS: Avg: 44 s
- p99 without DNS: Avg: 2.3 min
- p50 with DNS: Avg: 8.3 min
- p99 with DNS: Avg: 9.5 min



instance All ▾



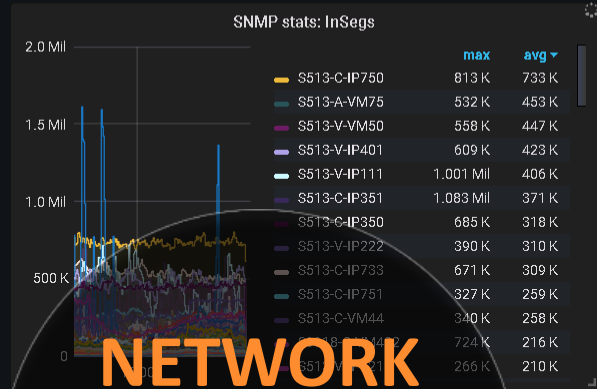


EOS
STORAGE
EOS Traffic
by Experiment







Environment All Top Hostgroup All Hostgroup All Host All LanDB Service Name All Filters +

Dashboard Row (1 hidden panels) Settings Trash

Dashboard Row



Dashboard Row



Alarms (WIP)


☐ Cooling infrastructure ▾




	min	max	avg	currer
Power	224 kW	446 kW	313 kW	359 kW
Apparent Power (right-y)	283 kVA	533 kVA	393 kVA	436 kV




	min	max	avg	current
Power	60 kW	212 kW	148 kW	156 kW
Apparent Power (right-y)	53 kVA	244 kVA	179 kVA	187 kVA




	min	max	avg	current
Power	7 kW	38 kW	19 kW	21 kW
Apparent Power (right-y)	9 kVA	47 kVA	25 kVA	30 kVA



	min	max	avg	current
Power	14 kW	14 kW	14 kW	14 kW
Apparent Power (right-y)	20 kVA	20 kVA	20 kVA	20 kVA



	min	max	avg	current
Power	7 kW	7 kW	7 kW	7 kW
Apparent Power (right-y)	11 kVA	11 kVA	11 kVA	11 kVA



	min	max	avg	current
Power	19 kW	19 kW	19 kW	19 kW
Apparent Power (right-y)	25 kVA	25 kVA	25 kVA	25 kVA

DC
OPERATIONS
 Cooling
 Infrastructure



Live view

Channel Percentage (fernandl) alert
Status - FAILURE

Flume Channel (ALL)
Status - OK

Flume Channel (MONIT)
Status - OK

InfluxDB (m-wlcg)
Status - OK

Spark job (ATLAS JM RAW)
Status - OK

Spark job (DDM)
Status - OK

Spark job (FTS)
Status - OK

Spark job (Network Protocols)
Status - OK

Spark job (XRootD Alice)
Status - OK

Test notification
Status - FAILURE

alarm.batch_vm_error_state
Status - OK

df rootme sent bytes used value
Status - FAILURE

exception.batch_vm_error_state
Status - FAILURE

flume_channel_full
Status - FAILURE

load_load_relative_shortterm
Status - FAILURE

swap_dev_vda2_percent_used_value
Status - FAILURE

tail_base_count_filesystem_error_value
Status - OK

tail_base_count_vm_kill_value
Status - OK



Lessons Learned

Migration to open source technologies

- Robust products, easier to find resources, to support users and share knowledge
- Requires adapting our requirements and leaving custom and tailored solutions
- Well-known in the community and outside

Building a reliable data pipeline

- Decouple producers from back ends, and have a lot of buffers
- Rely on technologies that scale (Kafka, Spark, HDFS, etc.)

Need several storage systems, for different use cases

- InfluxDB: only time-series data, SQL-like language, built-in aggregations, care with high cardinality in number of series, needs some schema design
- Elasticsearch: Good for discovery and search of logs, complex engine that requires special care with configuration
- HDFS: For batch or high latency scenarios, reports and analytics

Summary

CERN IT moving to common monitoring infrastructure (MONIT)

- CERN Data Centres, hardware and service monitoring
- WLCG transfers, job monitoring, sites availability

Adopted established open source technologies

- | | | |
|------------------------------|-----------------------|-----------------------|
| • Transport and Buffering | Flume, Kafka | 90K docs/s |
| • Processing and Aggregation | Spark, InfluxDB | |
| • Storage and Search | ES, HDFS, InfluxDB | 3 TB/day |
| • Visualization and Access | Grafana, Kibana, Swan | 1500 registered users |

Providing a solid standard monitoring dataflow

- Very easy to add data sources, dashboards and alarms
- Support infrastructure, coaching, tutorials for IT service managers & Experiment experts
- Used beyond the DC and WLCG migration ex. IPV6 WG, perfSonar, AMS (see talk later)

Reference and Contact

Contact

cern.ch/monit-support

Documentation

cern.ch/monit-docs

Dashboards (CERN SSO login)

cern.ch/monit

Backup Slides

MONIT Operations

Standard CERN IT Service

- Based on Openstack VMs (CentOS7)
- Puppet for service configuration
- Separate production/QA clusters

MONIT of MONIT

- Collectd for infrastructure metrics
- MONIT alarms



openstack®



Data Centre

IT SERVICES OVERVIEW

HOST SYSLOG (WITHIN CERN ONLY)

HOST METRICS

ALARMS HISTORY

LHCOPN

WLCG

WLCG TRANSFERS 30 DAYS

FTS TRANSFERS

WLCG TRANSFERS 5 YEARS

XROOT TRANSFERS

T0 - T1 TRANSFERS

T1 - T2 TRANSFERS

ATLAS

CMS

MONIT

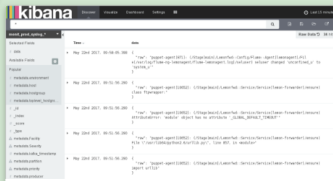
monit.cern.ch

Services




MONIT-GRAFANA

Use Grafana for plotting time series metrics coming from




MONIT-TIMBER

Use Timber to explore logs stored in Elasticsearch and interactive data discovery



MONIT-KIBANA

Use Kibana to explore metrics stored in Elasticsearch and



SWAN

Use Swan to generate more advanced reports and