



Streaming Service @CERN IT-DB-SAS

Jose C. Luna
Manuel Martin Marquez

Why?

Growing interest @CERN

In many occasions for central data backbone, buffering, and data enrichment.
Then sending this data to other platform (elastic, hadoop, other databases).

Examples

- Security Team integrated platform
- Central Monitoring of linux hosts, both physical and virtual in data centers
- Central Monitoring of windows servers
- New Accelerator Logging
- ...

What we currently offer

Managed Kafka Clusters

Dedicated kafka clusters, client retains admin access.

Virtual Machines, persistence in Netapp => moving to ceph

Kerberos and certificate support for clients.

We provide:

- Configuration: Integrated in central CERN IT (puppet based configuration)
- Monitoring: central IT pipeline (collectd+influx+grafana). Also testing kapacitor.
- Upgrades
- Expertise

Eg: Security-Team

Just started last year in test mode:

- 60K IN messages per second
- Aggregated bandwidth of around 200 MBps. Clients 20MB/s in 80MB/s out
- 10 Brokers, around 100 partitions per broker. -> 4VCPU / 8GB
- 72 hours retention policy. ~12TB persistence total as of now

Type of content: connections crossing firewall, dns requests, commands executed in shared linux clusters...



Evolution Plans

Shared Kafka Clusters

Substitute admin access through web/api

Small sized customers

Popularize streaming

Build solutions on top. Eg: IoT

IoT

- LoRa -> MQTT (internal) -> Shared Kafka
 - Offer also integration with influx and then grafana

Other

Integrations (kafka connect)