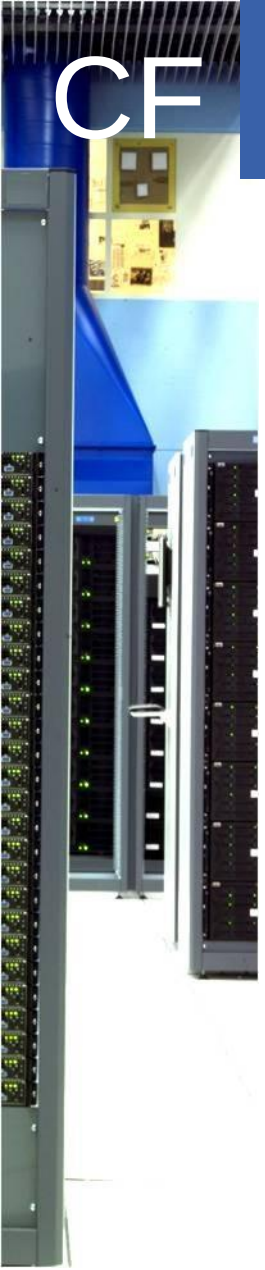

Agile Infrastructure Monitoring

pedro.andrade@cern.ch
CERN IT-CF

HEPiX Spring Meeting 2014

Introduction

Implement a shared monitoring architecture in IT

- Based on a common components tool-chain
- Delivered under a collaborative effort
- Simple to adopt and easy to scale horizontally

Adopt open source tools

- For each architecture block look outside for solutions
- Large adoption and strong community support
- Fast to adopt, test, deliver (and replace!)

Integrate with new infrastructure tools

- OpenStack, Puppet, etc.

Introduction

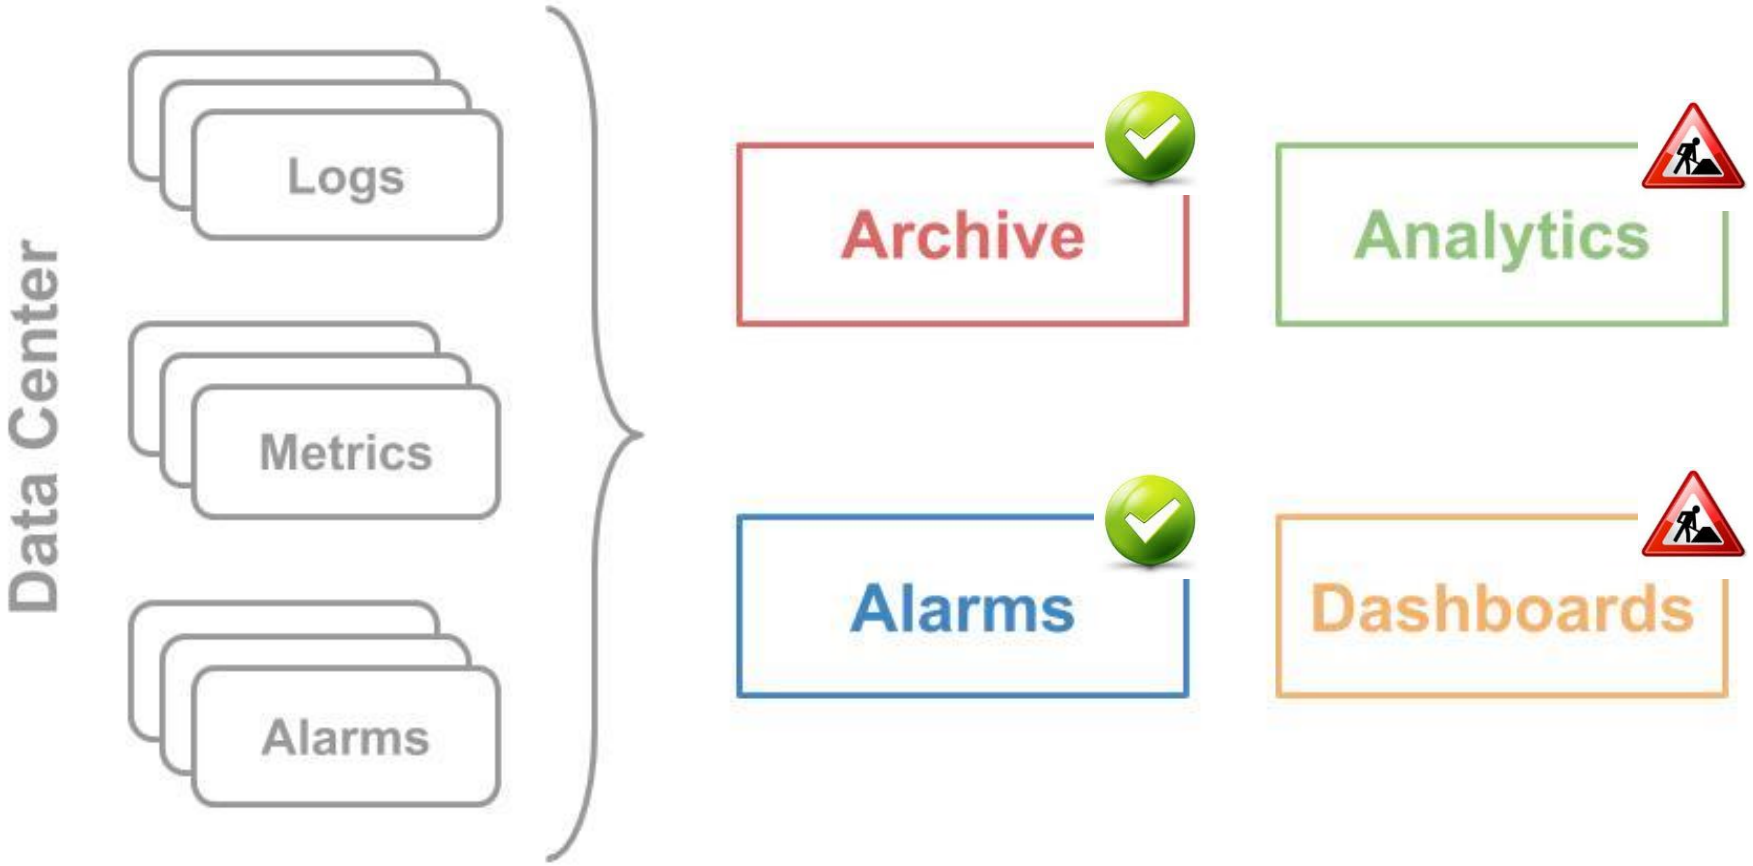

Archive

Long term archival

Store data for offline processing (MapRed)

Future data replay to other tools

Archive (technology)

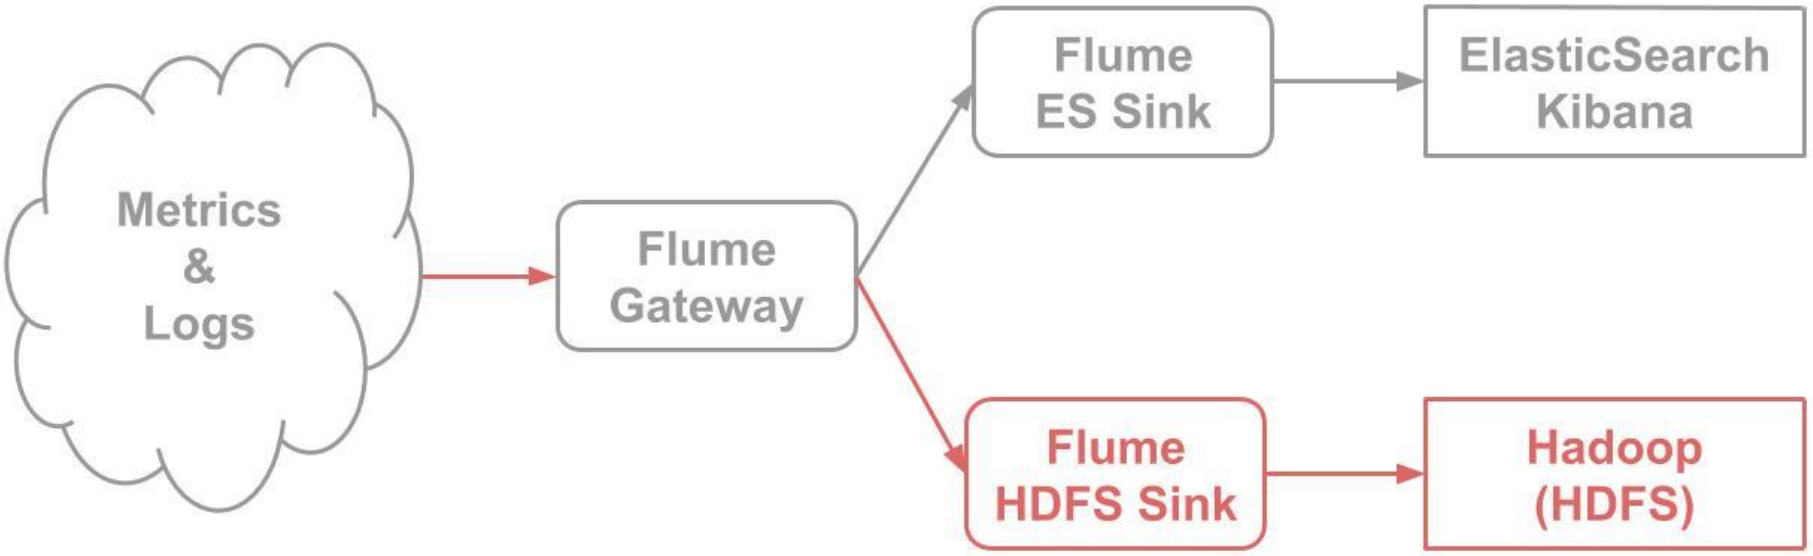

Archive (feedback)

Data stored by hostgroup/cluster

- Good for today's use cases
- No TTL defined so far

Daily aggregation of data

- Using PIG jobs to aggregate data per month
- To avoid having too many small files

Out of the box flume integration

- Needs some tuning to correctly size flume tiers
- Comes with many plugins… saving a lot of time
- Allows "in-flight" data processing (interceptors, handlers)

Archive (deployment)

Hadoop HDFS

- 5 nodes running CDH4.3, being upgraded to CDH5
- Collecting ~50 GB/day since September 2013 (13 TB)

Flume Gateway

- 20 nodes (VMs) running Flume 1.4
- Aggregating data from all nodes, writing to HDFS sink
- File channel based on CEPH, ~8 hours buffer

Flume HDFS sink

- 10 nodes (VMs) running Flume 1.4
- Aggregating data from gateways, writing to HDFS cluster

Dashboards

Dynamic creation of dashboards

Global and application specific views

User friendly !!!

Dashboards (technology)

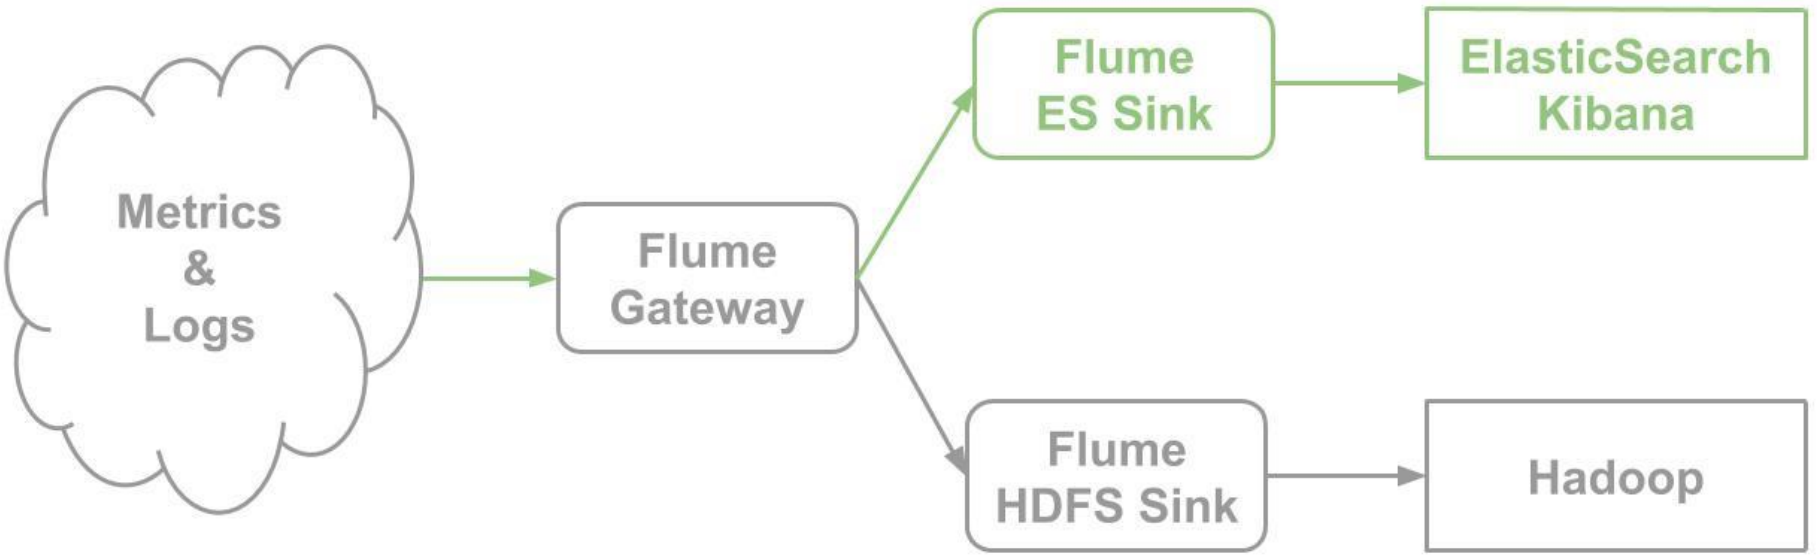

Dashboards (technology)

Dashboards (feedback)

ElasticSearch is easy to deploy and manage

- Using upstream package and puppet module
- Smooth upgrade to new major release (1.1)
- Full cluster start-up time still a bit long
- Easy to set data lifetime (TTL on indexes)
- Many nice plugins: head, elasticHQ, etc.

Kibana fully integrated in ES and growing

- Excellent to plot/search logs, works for metric data
- Not yet exploiting the new ES aggregation framework
- Difficult to create/share private dashboards

Dashboards (deployment)

Flume ES sink

- 10 nodes (VMs) running Flume 1.4
- Aggregating data from gateways, writing to ES cluster
- New ES sink using HTTP API with Flume 1.5

ElasticSearch (host metrics and syslog)

- 2 master, 2 search, 16 data, running ElasticSearch 1.1
- Storing and indexing ~50 GB/day since Summer 2013
- Daily indexes, 10 shards, 2 replicas, 1 month TTL

ElasticSearch (other instances)

- Smaller self-service deployments for specific needs

Analytics

Support different types of analytics use cases

From real-time streaming to offline processing

Provide right tools for the right use case

Analytics (technology)

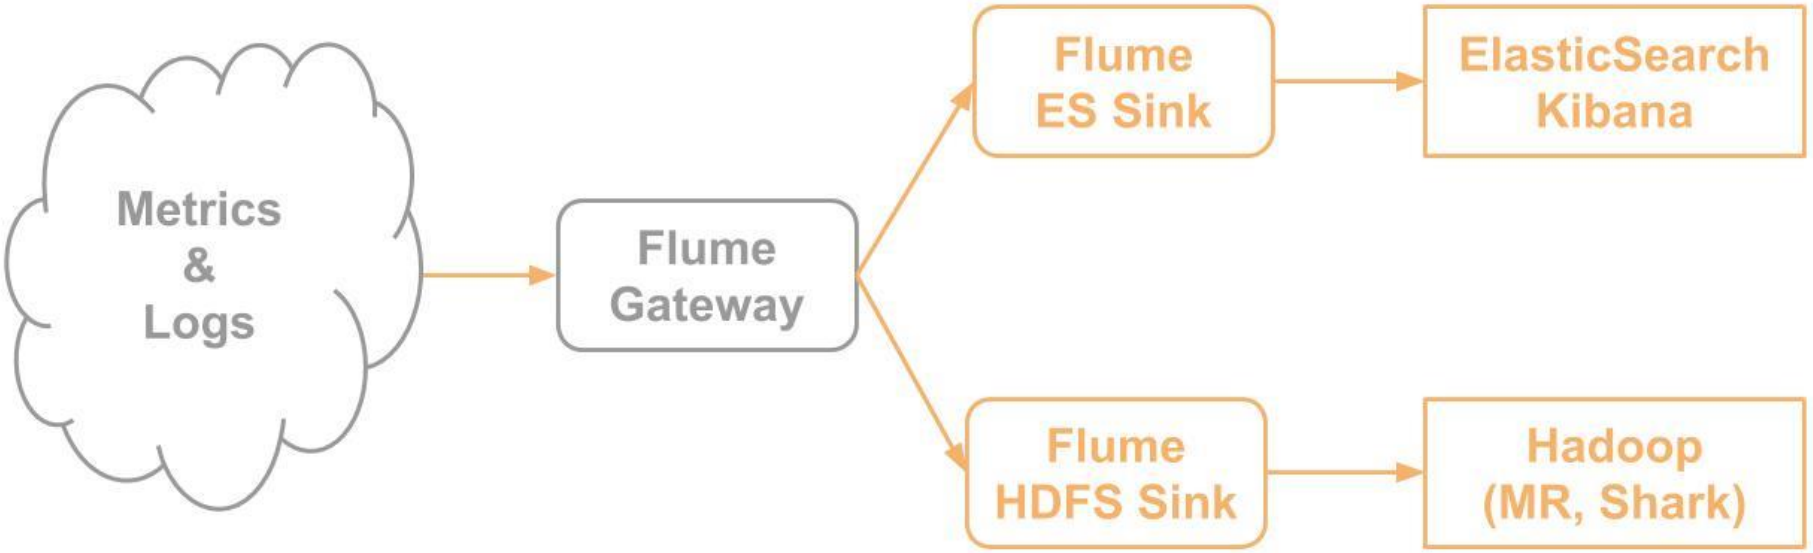

Analytics (feedback)

Offline: Hadoop MapReduce jobs

- E.g. power consumption and CPU load correlation

Querying and plotting: ES/Kibana

- Easy to create charts, search data, filter data, etc.
- Considering to write a simple CLI to query ES

Real-time streaming: Hadoop Spark/Shark

- E.g. data centre metrics, security logs, etc.
- Only initial discussion and prototypes
- Spark is an engine for large-scale data processing
- Shark is a distributed SQL query engine

Alarms

Fast and reliable delivery of alarms

Delivery of notifications to multiple channels

Integration with CERN incident system

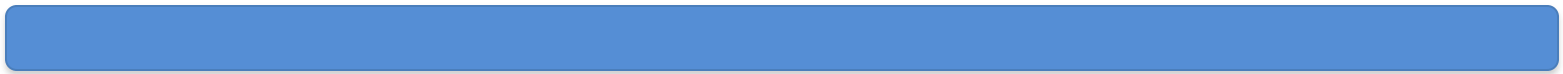

Alarms (technology)

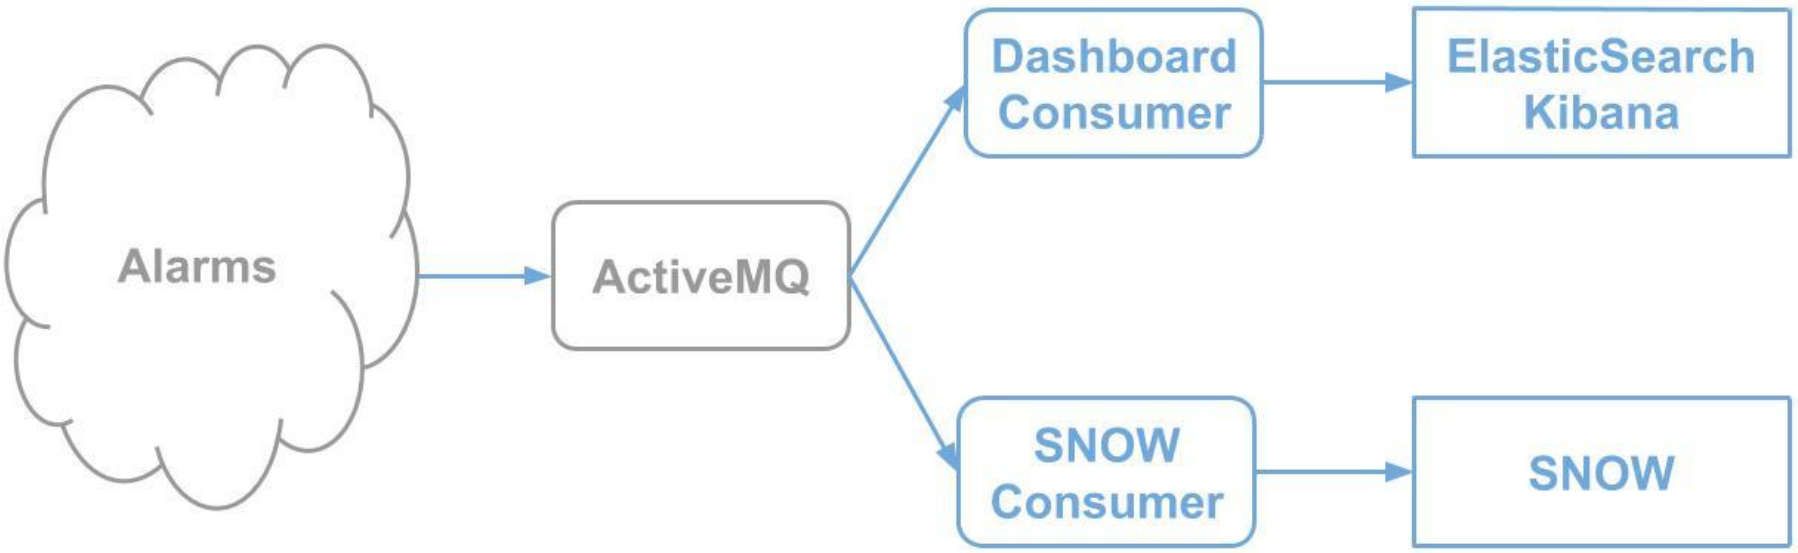

Alarms (feedback)

Alarms have more (special) metadata

- E.g. assignment group, incident creation, etc.
- Easy to configure alarms via puppet
  - assign “ipmi_wrong” alarm to “HW repair” team
  - disable creation of SNOW incident for “tmp_full”

Alarms can be created as incidents in ServiceNow

- Assigned to the right support unit (avoid useless hops)
  - Service incidents are directly assigned to Service Manager
  - Hardware incidents are directly assigned to HW Repair team

Special producer for no_contact alarms

- Based on node heartbeat metrics

Alarms (deployment)

General Notification Infrastructure (GNI)

- High-available messaging consumers in 4 VMs
- Handling all alarms for puppet nodes since April 2014
- Average of 500 SNOW incidents created per week

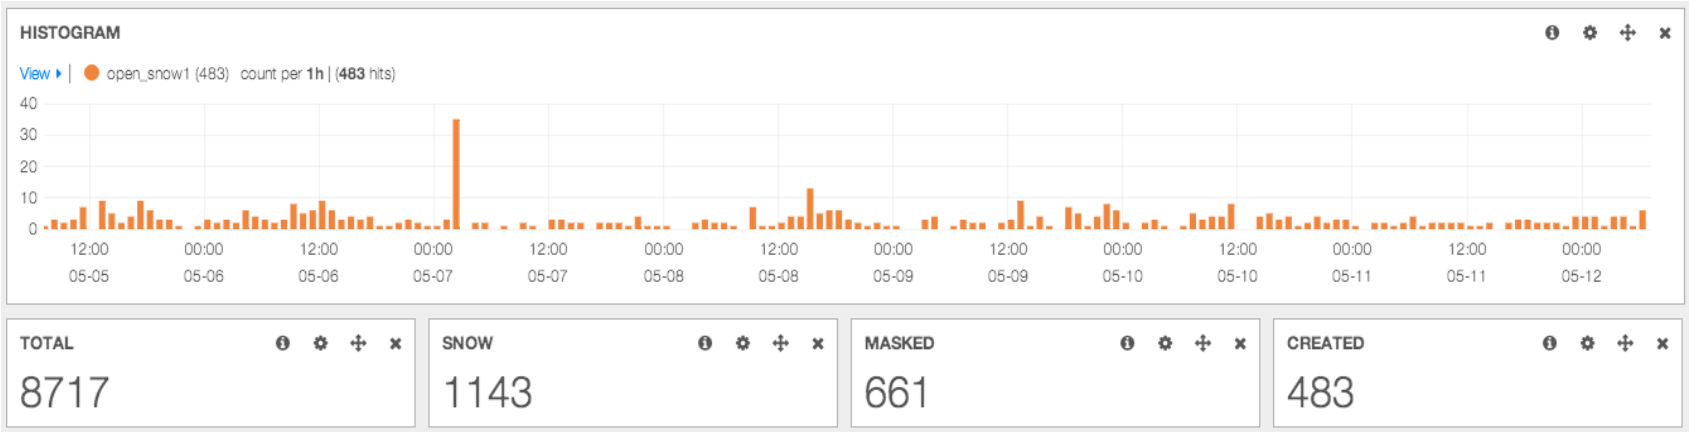

Next Steps

Complete work on dashboards

- More tests with ES/Kibana for metrics and logs
- Integration of host metrics and service metrics

Look into streaming analytics

- Based on Hadoop Spark and Shark
- Test functionality and deployment possibilities
- How to efficiently handle processed data ?

Provide monitoring services as PaaS

- Monitoring services easy to instantiate
- Integration with OpenStack Heat

Thanks !

Questions ?

itmon-team@cern.ch

http://cern.ch/itmon