Monitoring Evolution at CERN

pedro.andrade@cern.ch

Monitoring Wishes

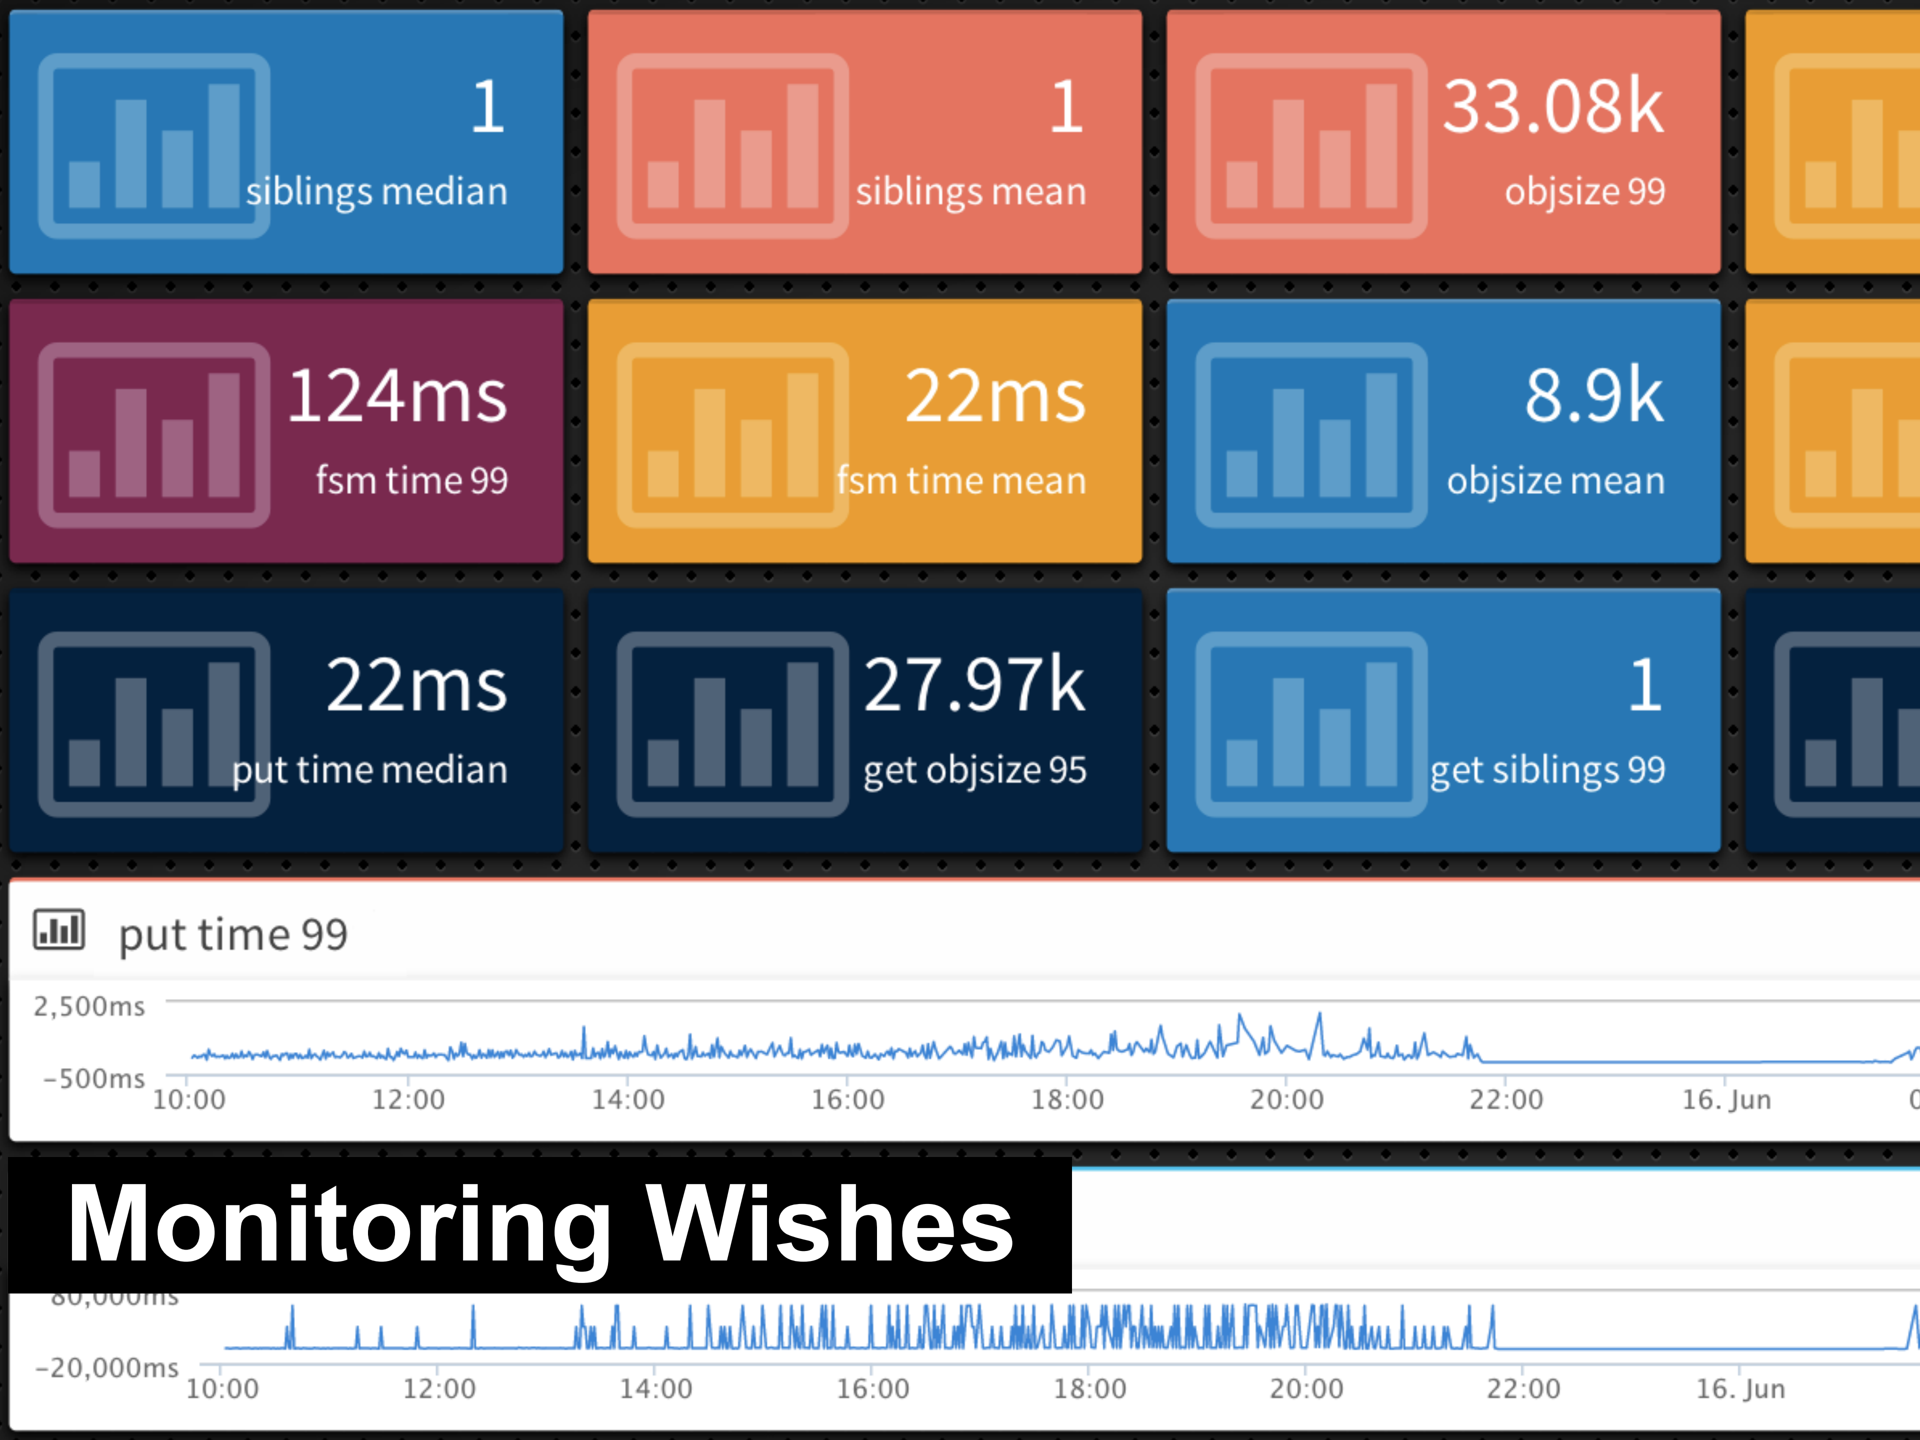

1
siblings median

1
siblings mean

33.08k
objsize 99

124ms
fsm time 99

22ms
fsm time mean

8.9k
objsize mean

22ms
put time median

27.97k
get objsize 95

1
get siblings 99

put time 99

it would be nice to have a dashboard with my hosts and services… which I could share…

the 10-years report of my service write operations per week shows that…

for sure I can do something useful with this temperature and humidity data…

Monitoring Reality

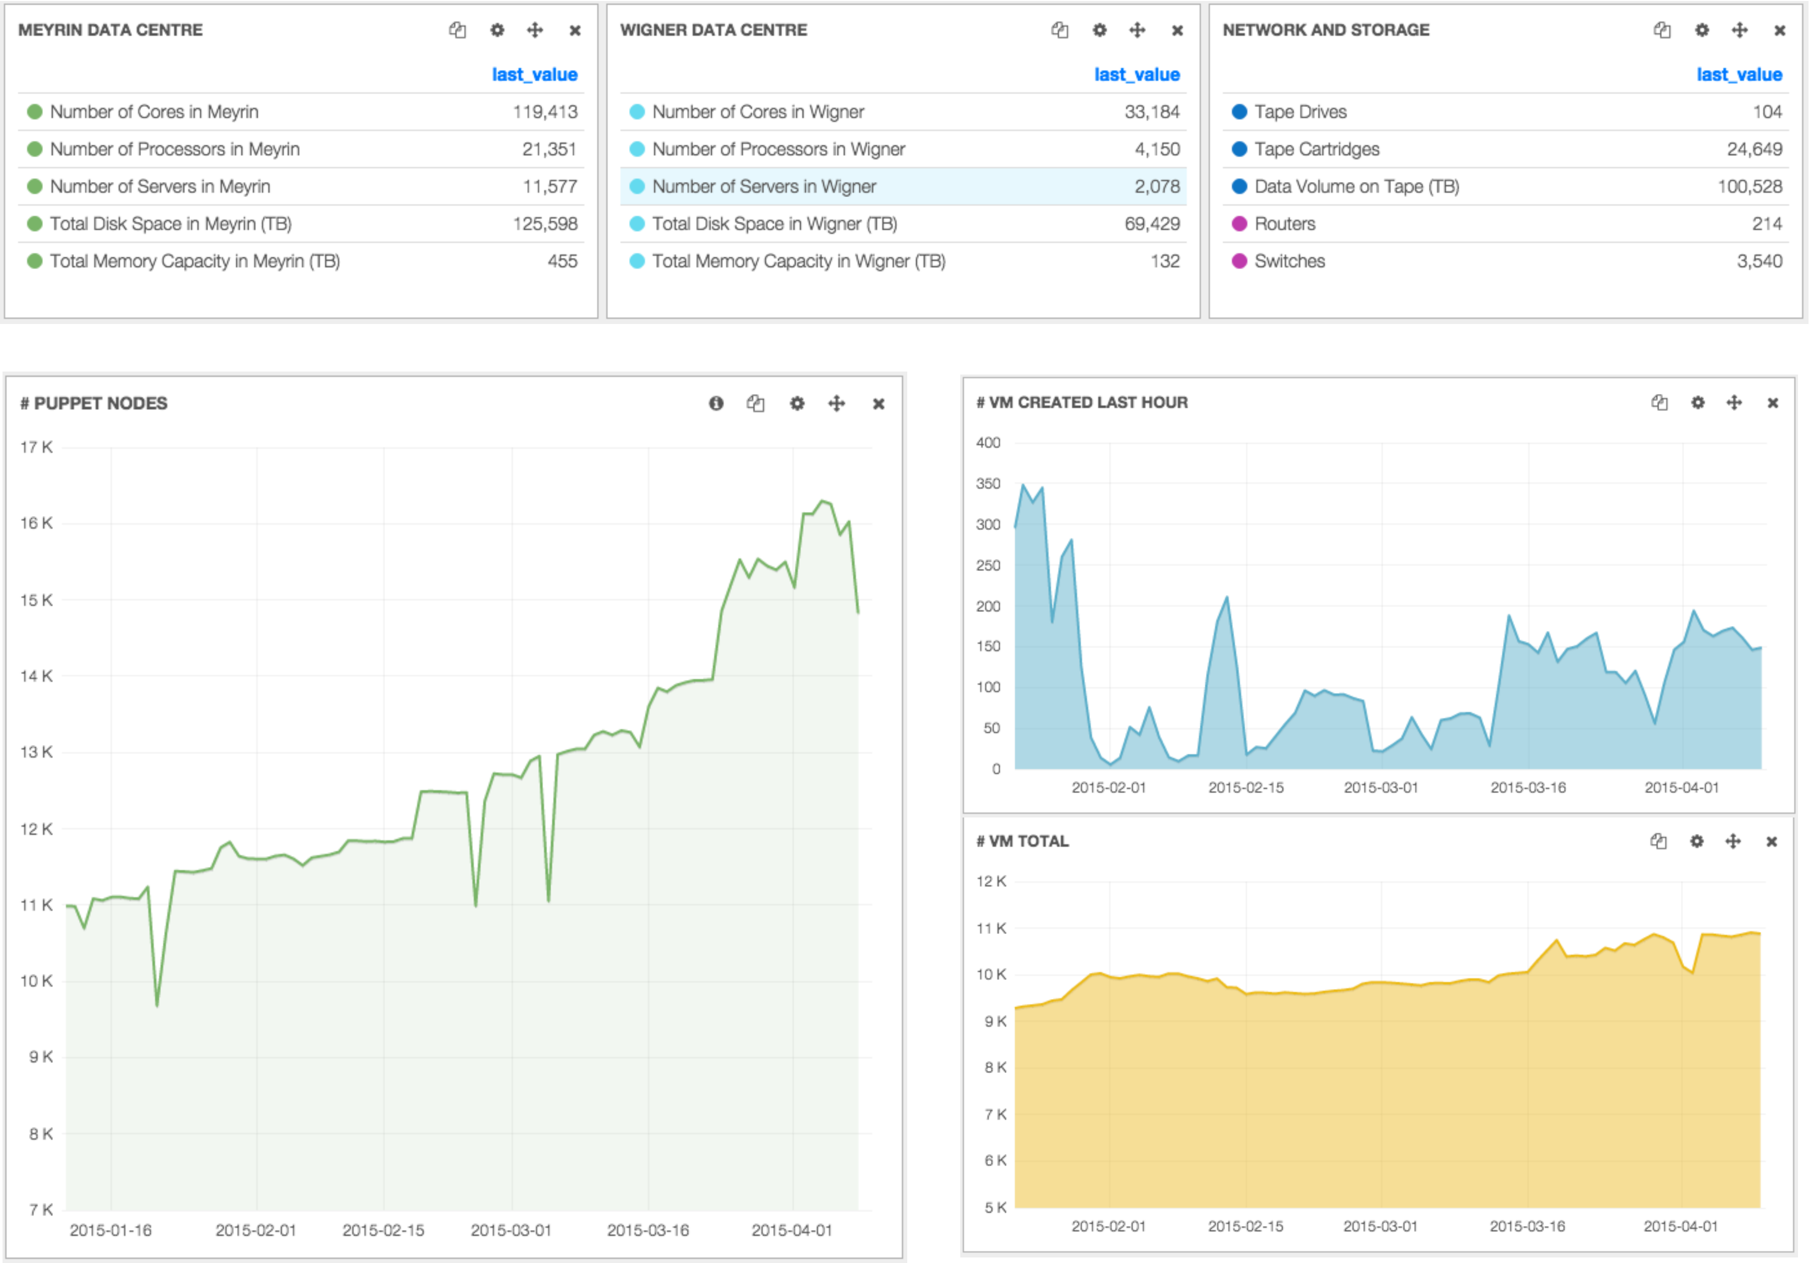

MEYRIN DATA CENTRE

	last_value
● Number of Cores in Meyrin	119,413
● Number of Processors in Meyrin	21,351
● Number of Servers in Meyrin	11,577
● Total Disk Space in Meyrin (TB)	125,598
● Total Memory Capacity in Meyrin (TB)	455

WIGNER DATA CENTRE

	last_value
● Number of Cores in Wigner	33,184
● Number of Processors in Wigner	4,150
● Number of Servers in Wigner	2,078
● Total Disk Space in Wigner (TB)	69,429
● Total Memory Capacity in Wigner (TB)	132

NETWORK AND STORAGE

	last_value
● Tape Drives	104
● Tape Cartridges	24,649
● Data Volume on Tape (TB)	100,528
● Routers	214
● Switches	3,540

PUPPET NODES

VM CREATED LAST HOUR

VM TOTAL

Old tools need replacement

… not scaling to current needs

… old code difficult to maintain

… only providing static displays

Identify

what you want to do with the data

Simplify

the data processing workflow

Rely

on existing technologies & standards

Identify

what you want to do with the data

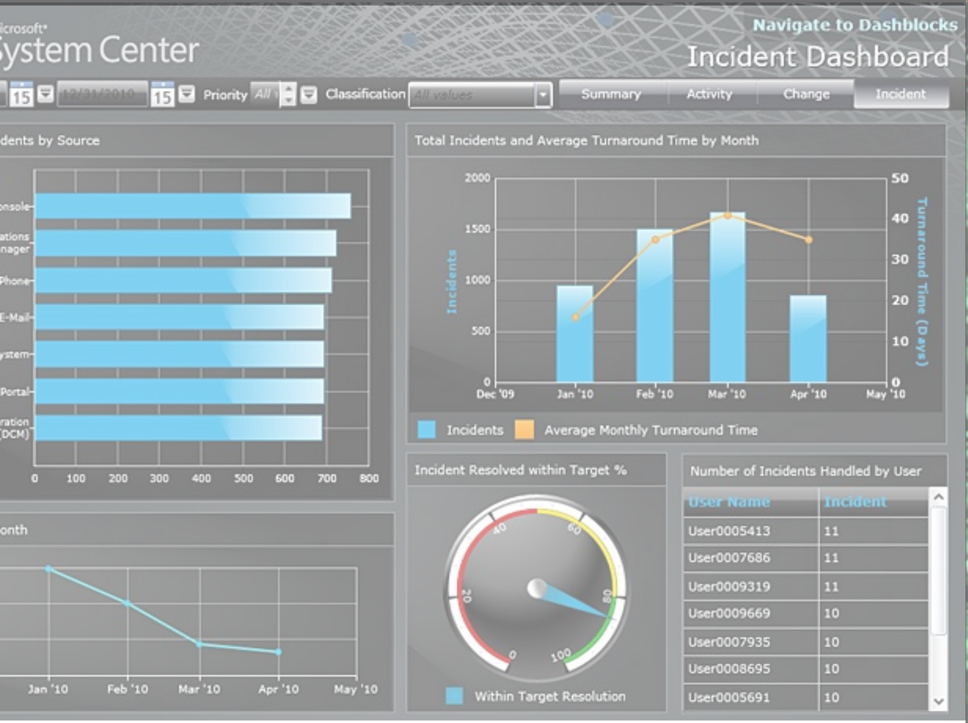

User Name	Incident
User0005413	11
User0007686	11
User0009319	11
User0009669	10
User0007935	10
User0008695	10
User0005691	10

Alerts

to get notified of problems affecting nodes/services

Archives

to create a data lake of monitoring events for offline analysis

Displays

to plot relevant monitoring data in realtime

Streaming

to programmatically process monitoring data

Simplify

the data processing workflow

Based on the lambda architecture data processing principles: stream & batch

Capable of sustaining large quantities of timestamped events: metrics and logs

Common to all monitoring producers, easy to replace individual components

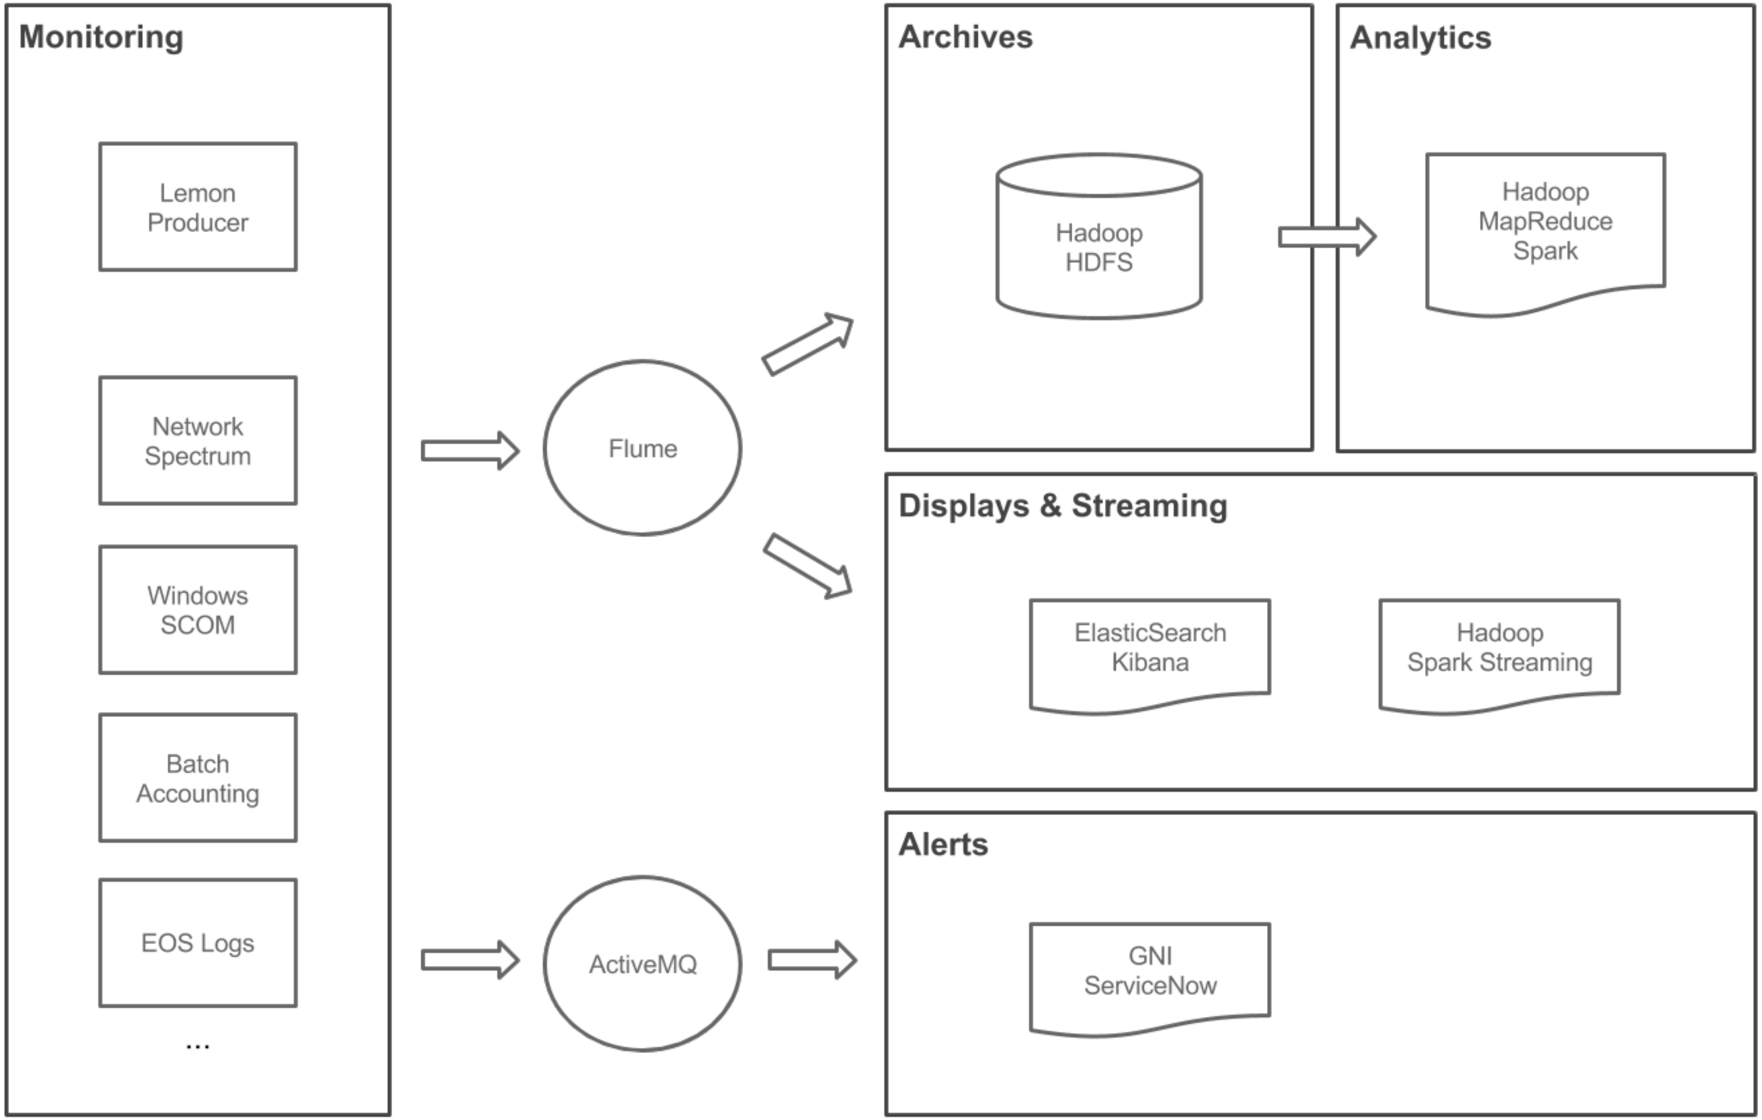

Monitoring

Lemon Producer

Network Spectrum

Windows SCOM

Batch Accounting

EOS Logs

...

Flume

ActiveMQ

Archives

Hadoop HDFS

Analytics

Hadoop MapReduce Spark

Displays & Streaming

ElasticSearch Kibana

Hadoop Spark Streaming

Alerts

GNI ServiceNow

Rely

on existing technologies & standards

- technology watch
- follow upstream releases

graphite

ActiveMQ

hadoop HDFS

elastic

Spark

FLUME

kibana

- osquery
- sensu
- kafka
- grafana
- hue
- …

hadoop
HDFS

elastic

Spark

FLUME

what

Distributed service for collecting large amounts of data, robust and fault tolerant, horizontally scalable, multi-tier deployment, many ready to use input and output plugins

@cern

Running Flume 1.5.0

- scalable 3 tiers deployment
- one flume agent in each DC node (~15k)
- 50 VMs aggregating data, writing to HDFS

what

Distributed framework for large data sets processing, distributed fault-tolerant filesystem for low cost hardware, suitable for applications with large data sets

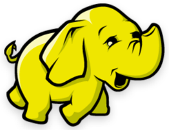

@cern

Running HDFS 2.5.0 (CDH 5.3.1)

- dedicated 16 nodes cluster
- data being archived since 2013
- currently storing ~32 TB data

@cern what

Distributed search and analytics engine, real time acquisition, data indexed in real time, automatically balanced shards and replicas, schema free, document oriented

Running ElasticSearch 1.4.4

- 3 masters, 4 search, 20 data nodes
- 1 week raw data, 1 year history
- 10 shards/index, 2 replicas/shards

what

Distributed large-scale data processing engine, improvement of MapReduce, support for streaming, SQL, machine learning, and graph processing

Spark

@cern

Running Spark 1.2.0 (CDH 5.3.1)

- prototyping on different use cases
- data centre temperatures and power
- detection of nodes “no contact”

Modular lambda architecture

Appropriate technologies

=

Useful monitoring solutions for CERN

IT Service Managers and Experiments

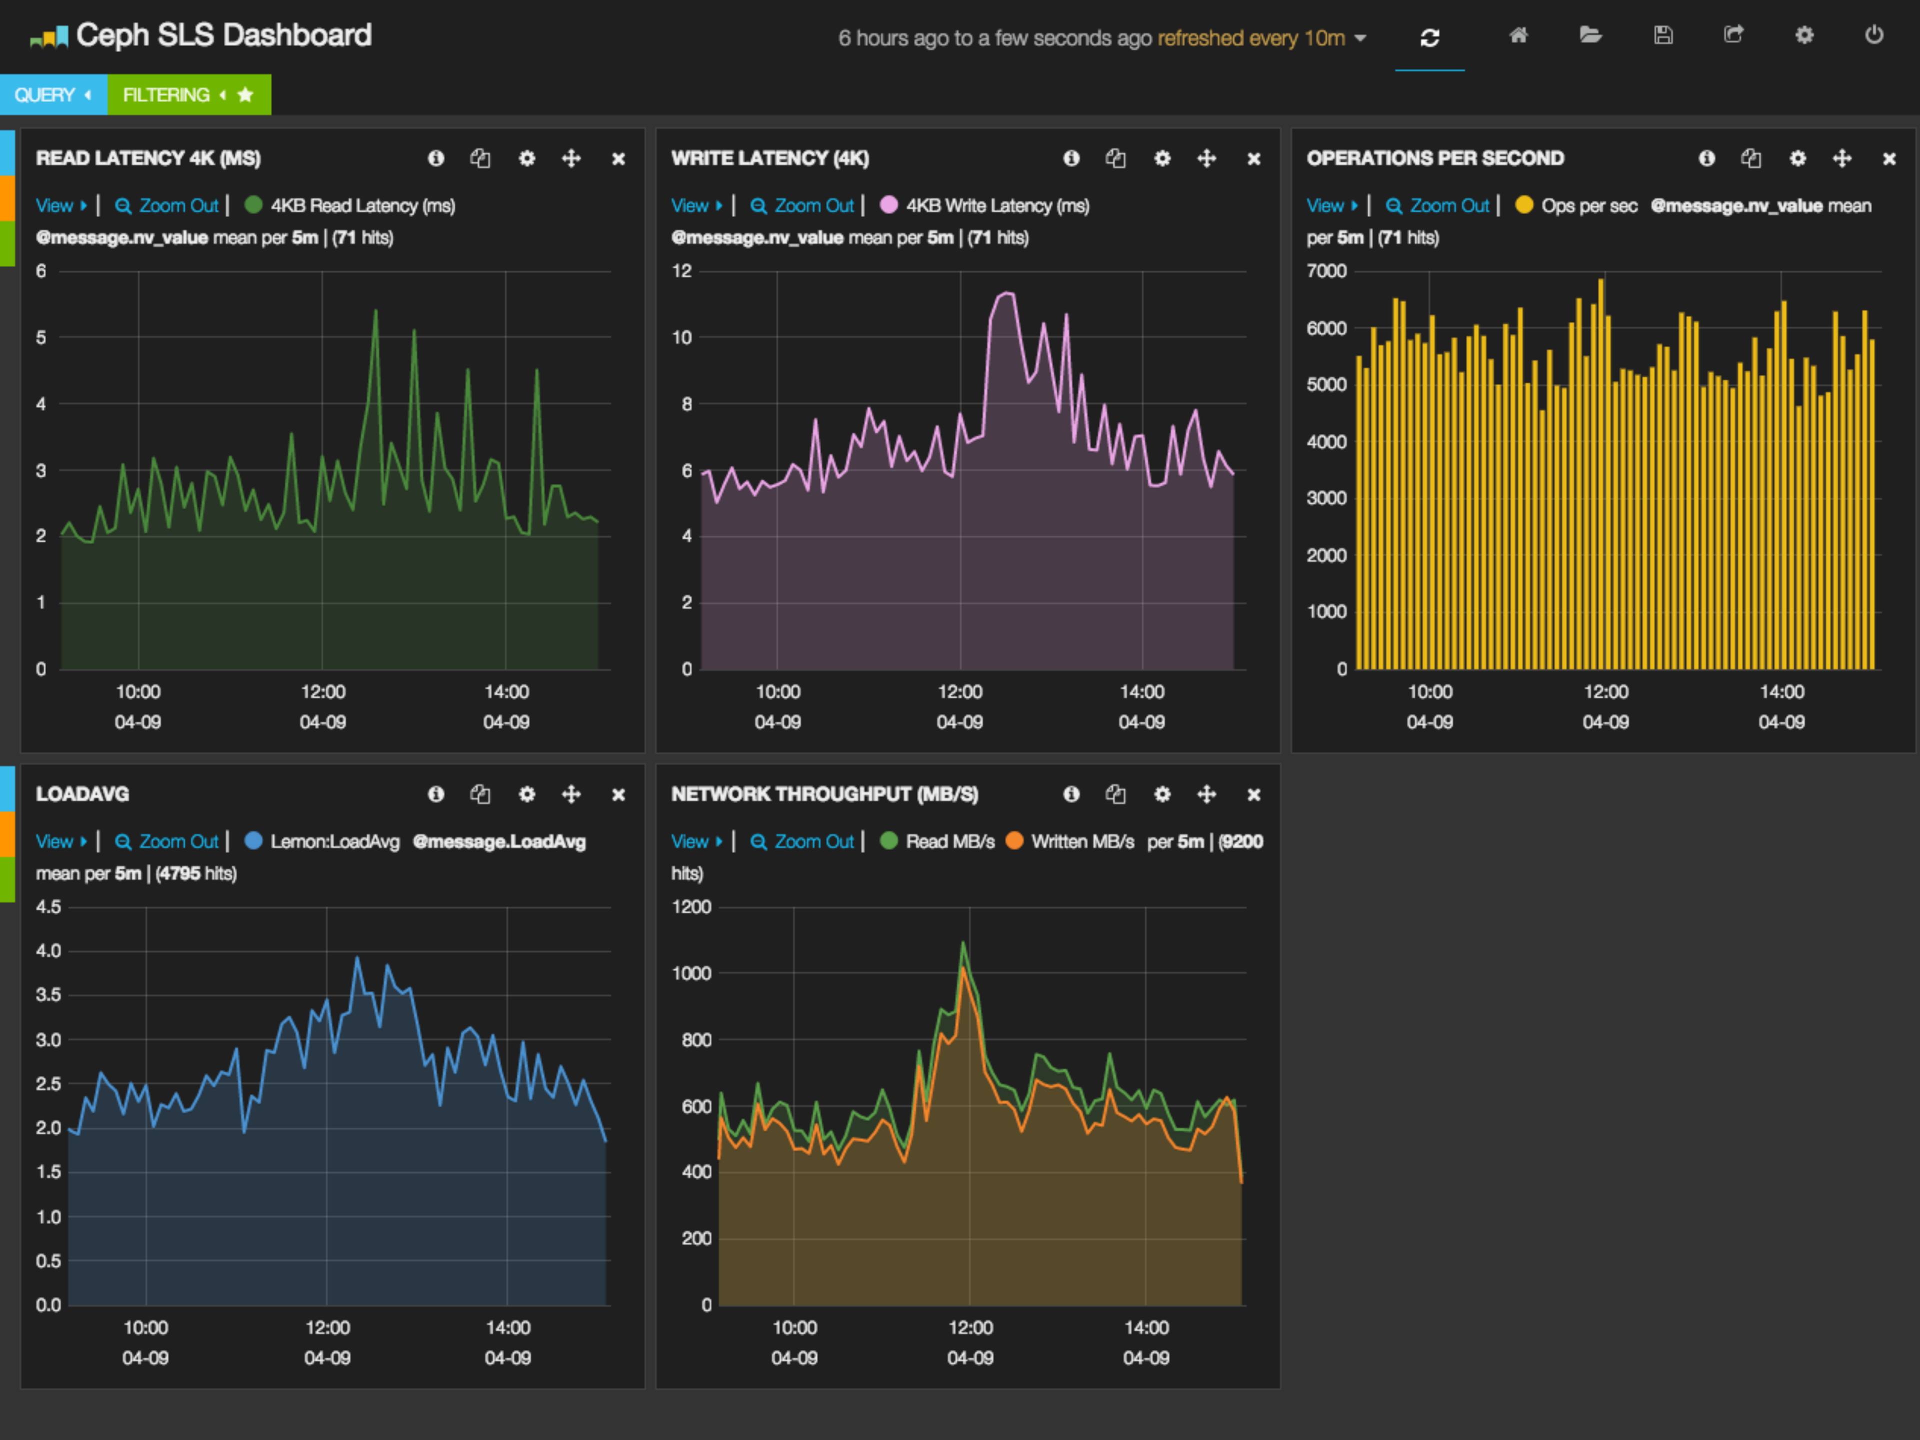

Ceph SLS Dashboard

6 hours ago to a few seconds ago refreshed every 10m

READ LATENCY 4K (MS)

View | Zoom Out | ● 4KB Read Latency (ms) @message.nv_value mean per 5m | (71 hits)

WRITE LATENCY (4K)

View | Zoom Out | ● 4KB Write Latency (ms) @message.nv_value mean per 5m | (71 hits)

OPERATIONS PER SECOND

View | Zoom Out | ● Ops per sec @message.nv_value mean per 5m | (71 hits)

LOADAVG

View | Zoom Out | ● Lemon:LoadAvg @message.LoadAvg mean per 5m | (4795 hits)

NETWORK THROUGHPUT (MB/S)

View | Zoom Out | ● Read MB/s ● Written MB/s per 5m | (9200 hits)

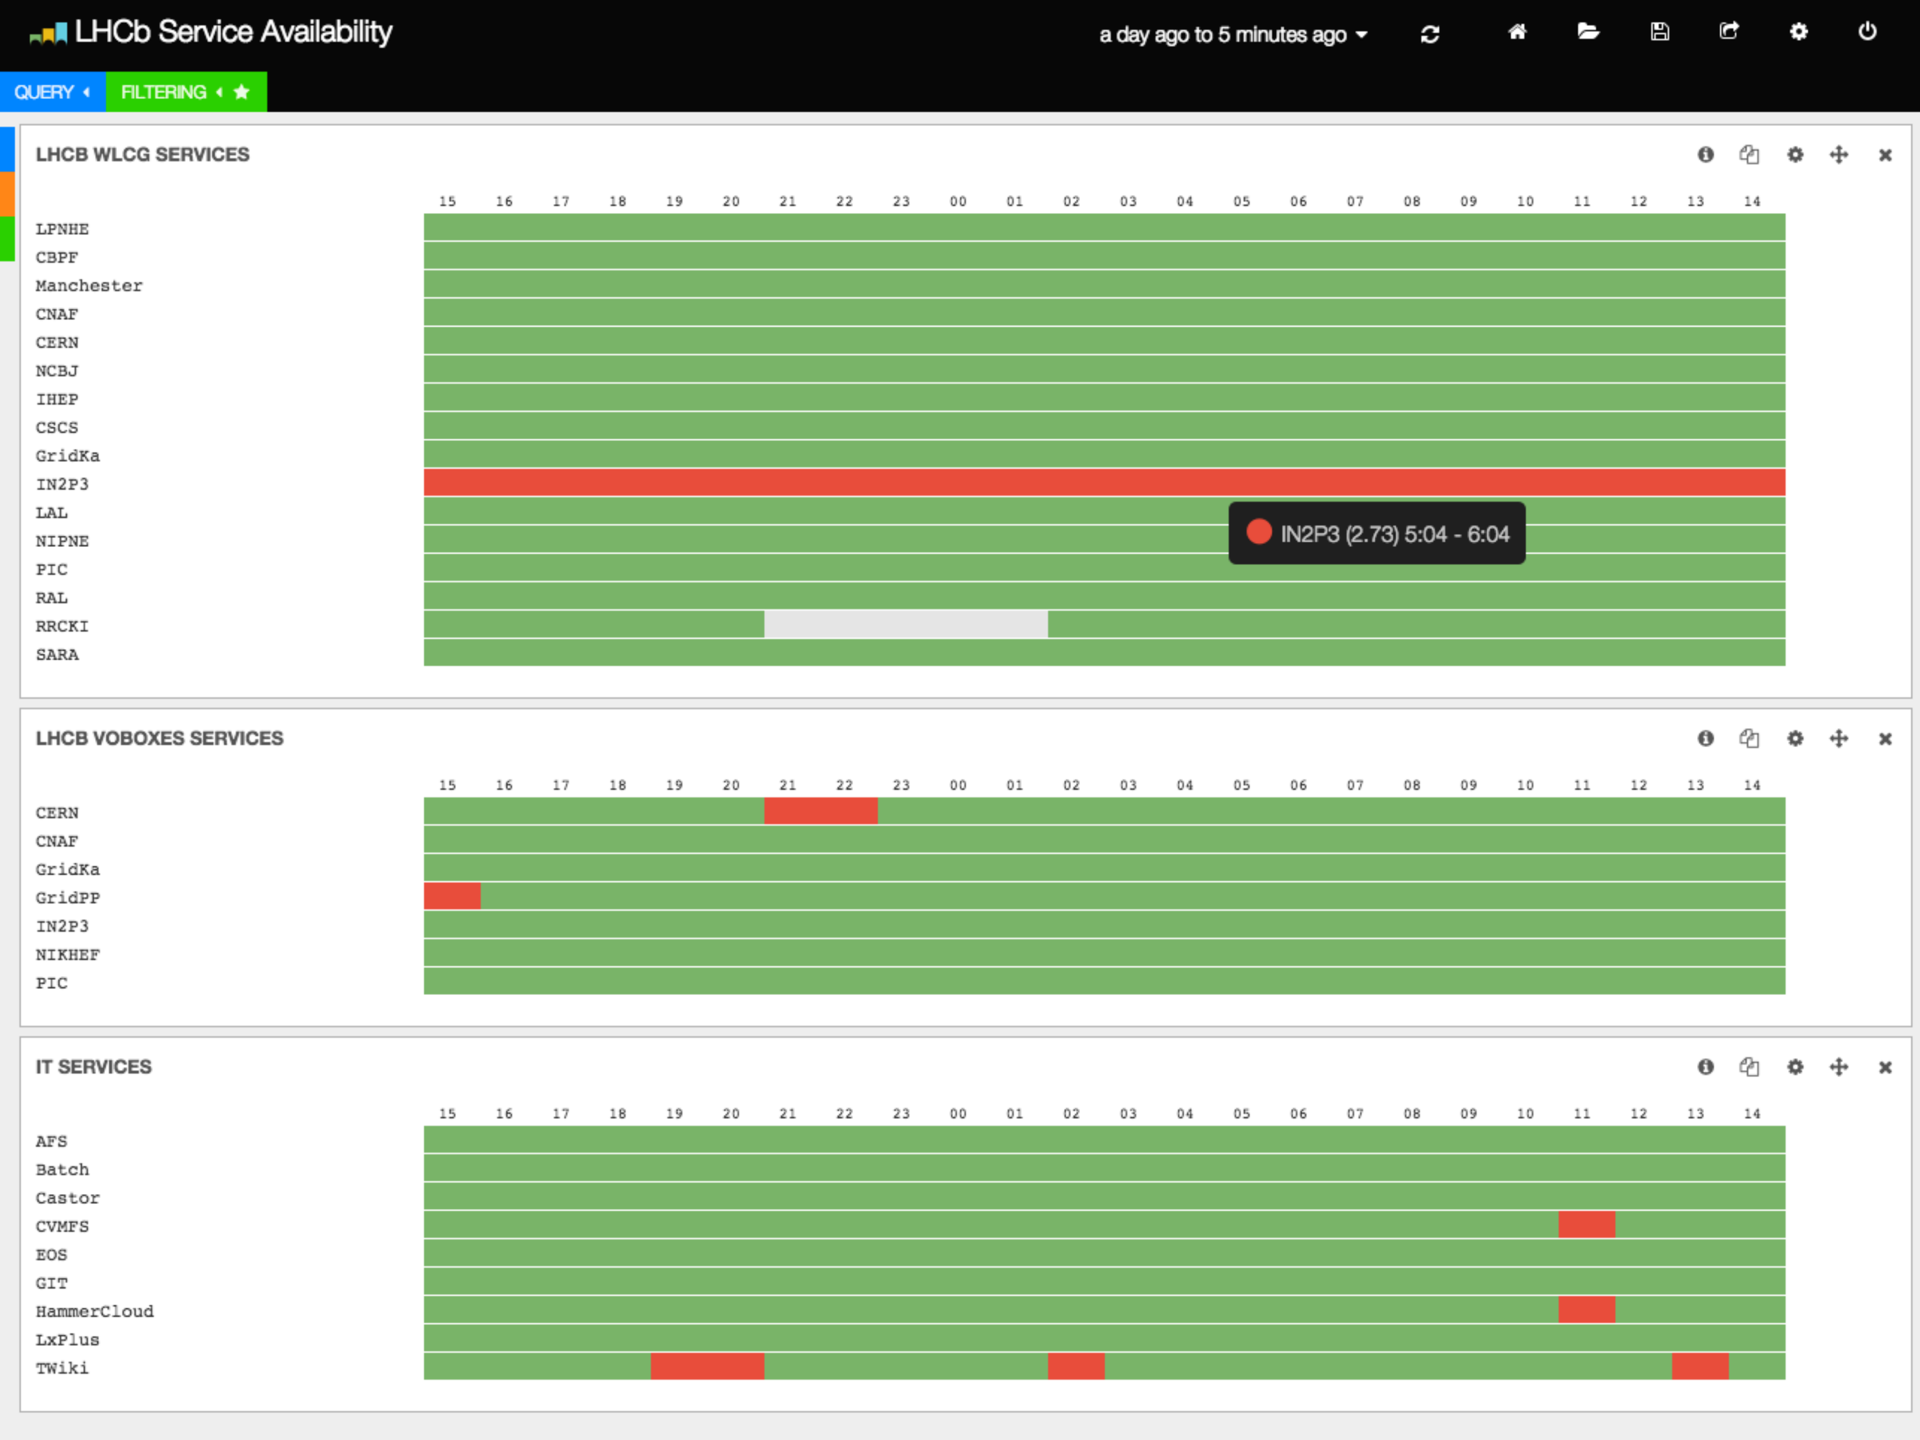

LHCb Service Availability
a day ago to 5 minutes ago
QUERY
FILTERING
LHCB WLCG SERVICES
15 16 17 18 19 20 21 22 23 00 01 02 03 04 05 06 07 08 09 10 11 12 13 14
LPNHE
CBPF
Manchester
CNAF
CERN
NCBJ
IHEP
CSCS
GridKa
IN2P3
LAL
NIPNE
PIC
RAL
RRCKI
SARA
● IN2P3 (2.73) 5:04 - 6:04
LHCB VOBOXES SERVICES
15 16 17 18 19 20 21 22 23 00 01 02 03 04 05 06 07 08 09 10 11 12 13 14
CERN
CNAF
GridKa
GridPP
IN2P3
NIKHEF
PIC
IT SERVICES
15 16 17 18 19 20 21 22 23 00 01 02 03 04 05 06 07 08 09 10 11 12 13 14
AFS
Batch
Castor
CVMFS
EOS
GIT
HammerCloud
LxPlus
TWiki

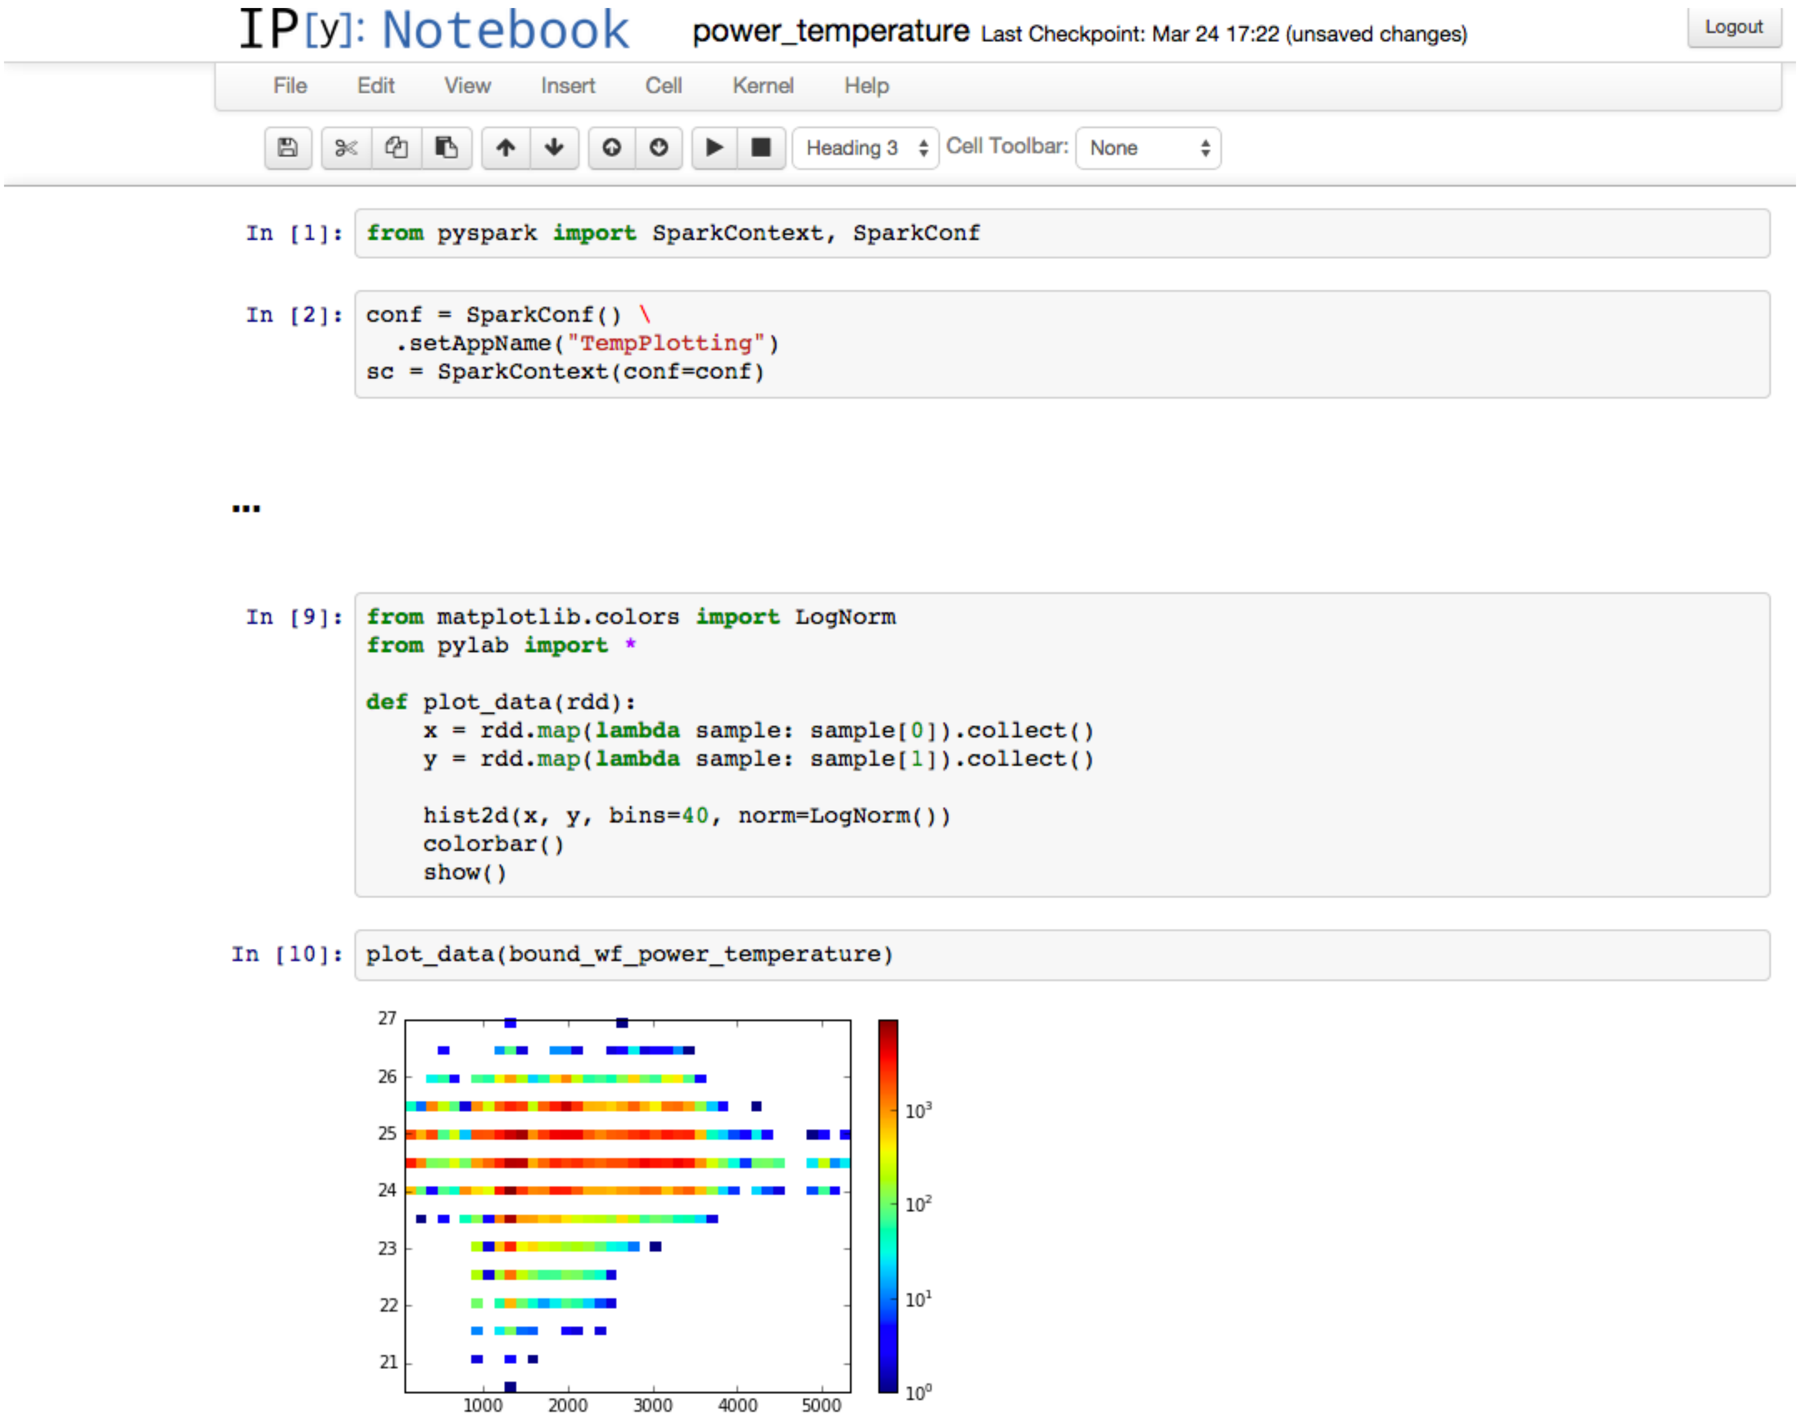

In [1]:

```
from pyspark import SparkContext, SparkConf
```

In [2]:

```
conf = SparkConf() \
  .setAppName("TempPlotting")
sc = SparkContext(conf=conf)
```

...

In [9]:

```
from matplotlib.colors import LogNorm
from pylab import *

def plot_data(rdd):
    x = rdd.map(lambda sample: sample[0]).collect()
    y = rdd.map(lambda sample: sample[1]).collect()

    hist2d(x, y, bins=40, norm=LogNorm())
    colorbar()
    show()
```

In [10]:

```
plot_data(bound_wf_power_temperature)
```

Clear architecture & modular technologies

Different solutions for alerting, streaming, displays, and archive

Successfully deployed and running in production for +1 year

Thank you !!!

http://cern.ch/itmon

itmon-team@cern.ch

CERN

www.cern.ch