Using Big Data for Anomaly Detection

Javier Cacheiro, CESGA

jlopez@cesga.es

Outline

- Using Big Data to collect Metrics & Logs
- Use case: Anomaly Detection
  - Detecting SSH attacks using log analysis
  - Detecting SLURM job anomalies using metrics

Phases

Phase 1: Measure → Collect & Store

Phase 2: Understand → Analyze & Visualize

Phase 3: Control → Monitoring

Phase 4: Improve → Anomaly Detection

Metrics

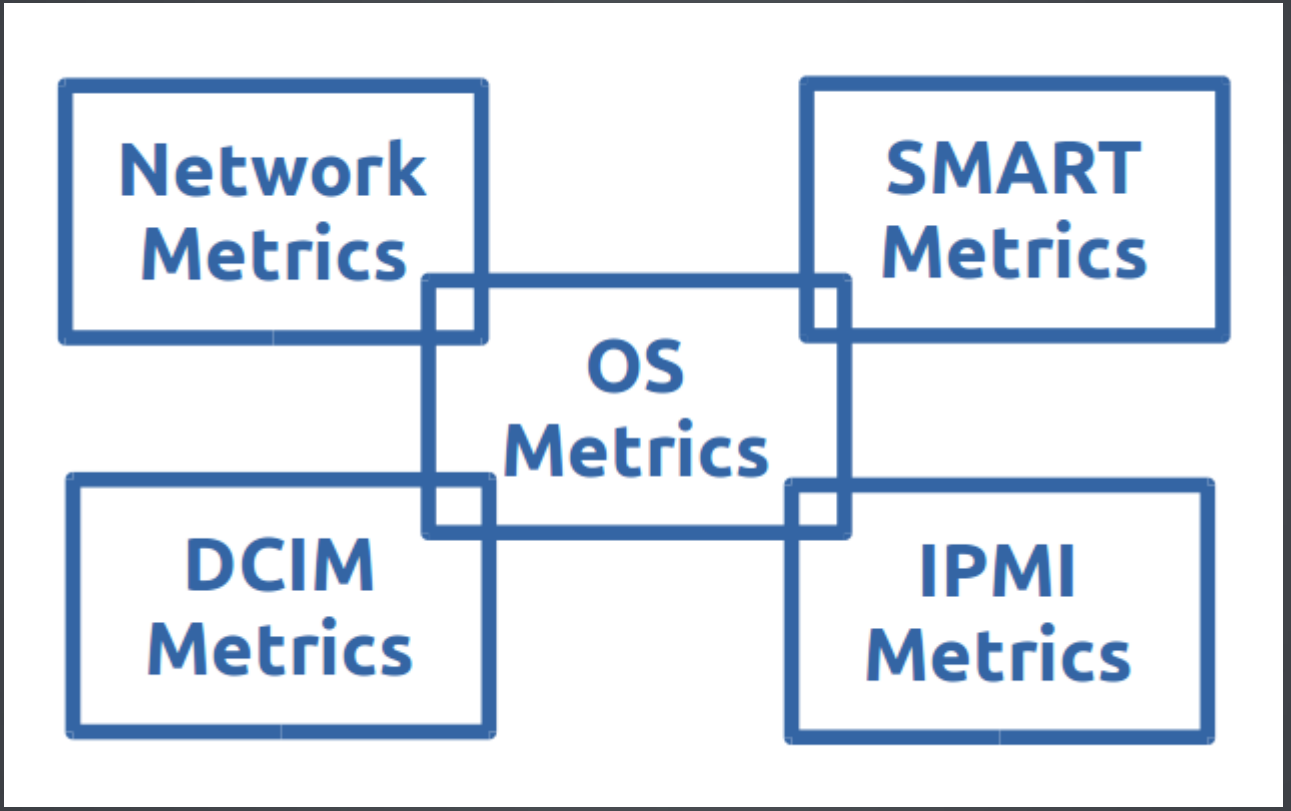
Network Metrics
SMART Metrics
OS Metrics
DCIM Metrics
IPMI Metrics

Logs

etc

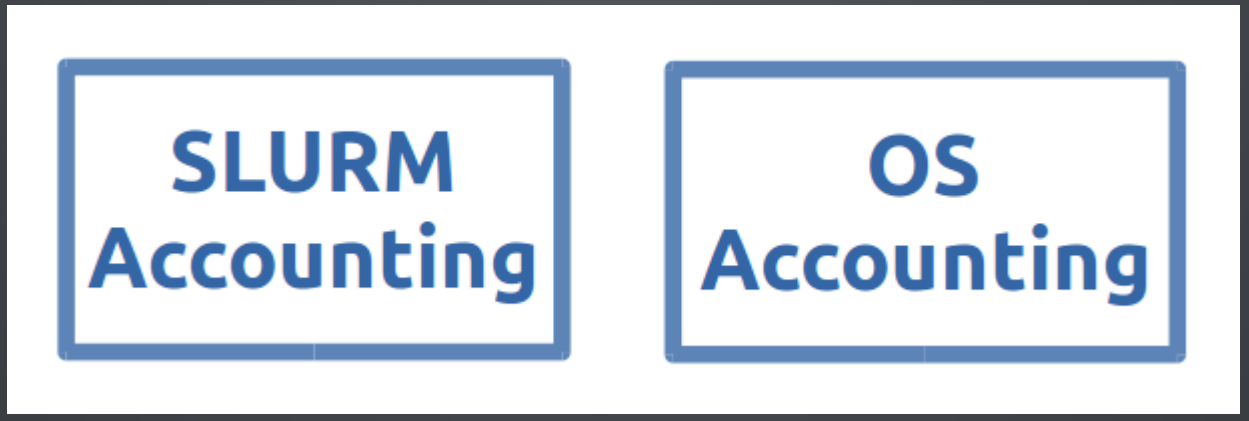

33487 metrics

10 Million time series

Metrics Collection Architecture

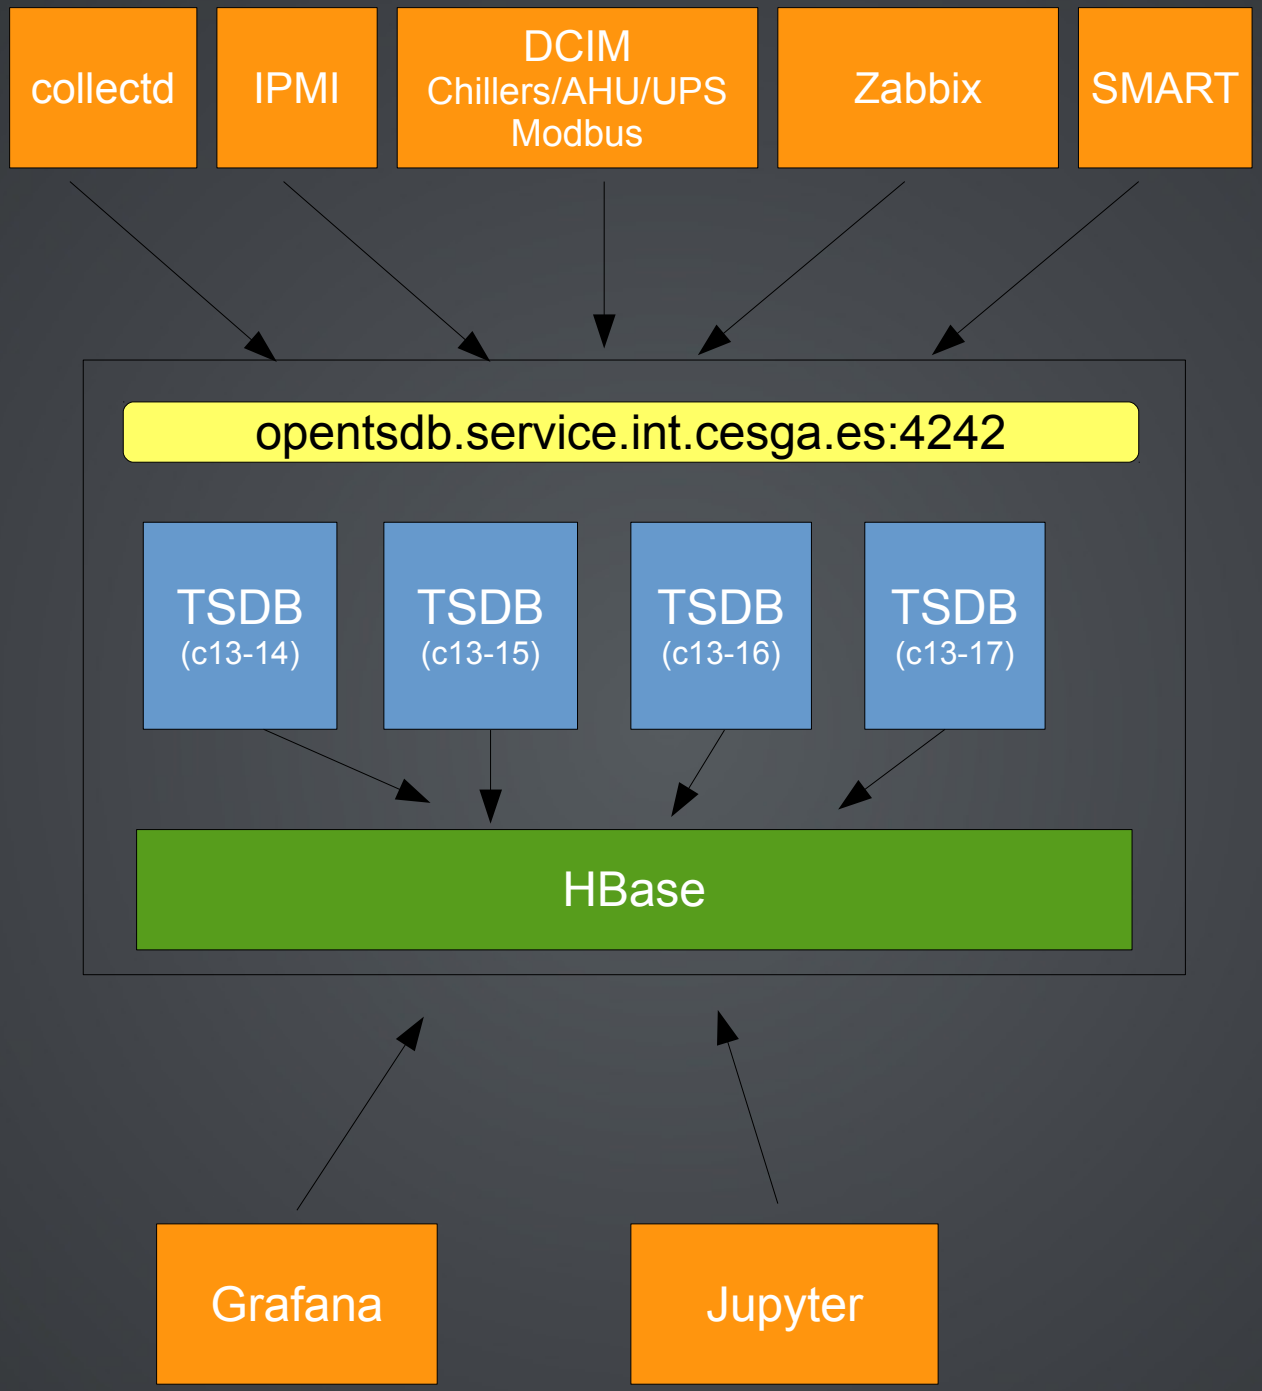
collectd
IPMI
DCIM
Chillers/AHU/UPS
Modbus
Zabbix
SMART
opentsdb.service.int.cesga.es:4242
TSDB
(c13-14)
TSDB
(c13-15)
TSDB
(c13-16)
TSDB
(c13-17)
HBase
Grafana
Jupyter

DCIM Collection Architecture

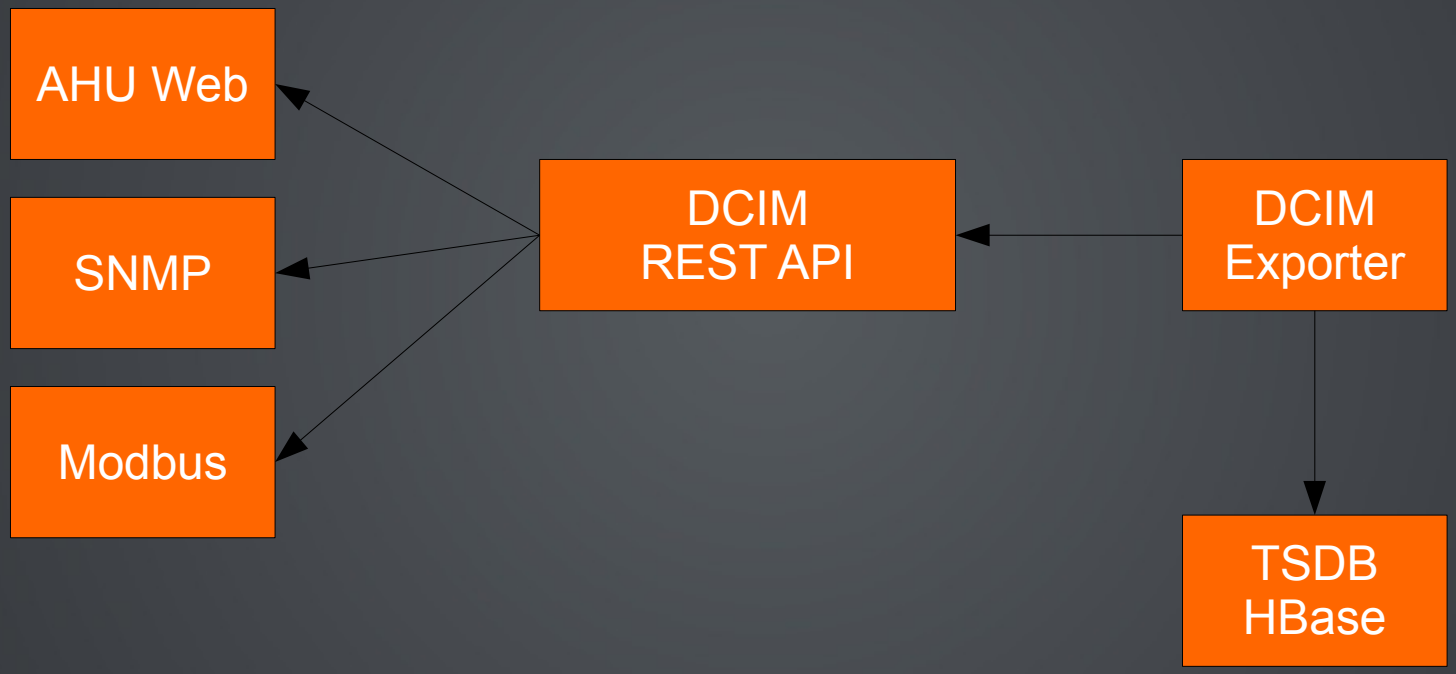

Log Collection Architecture

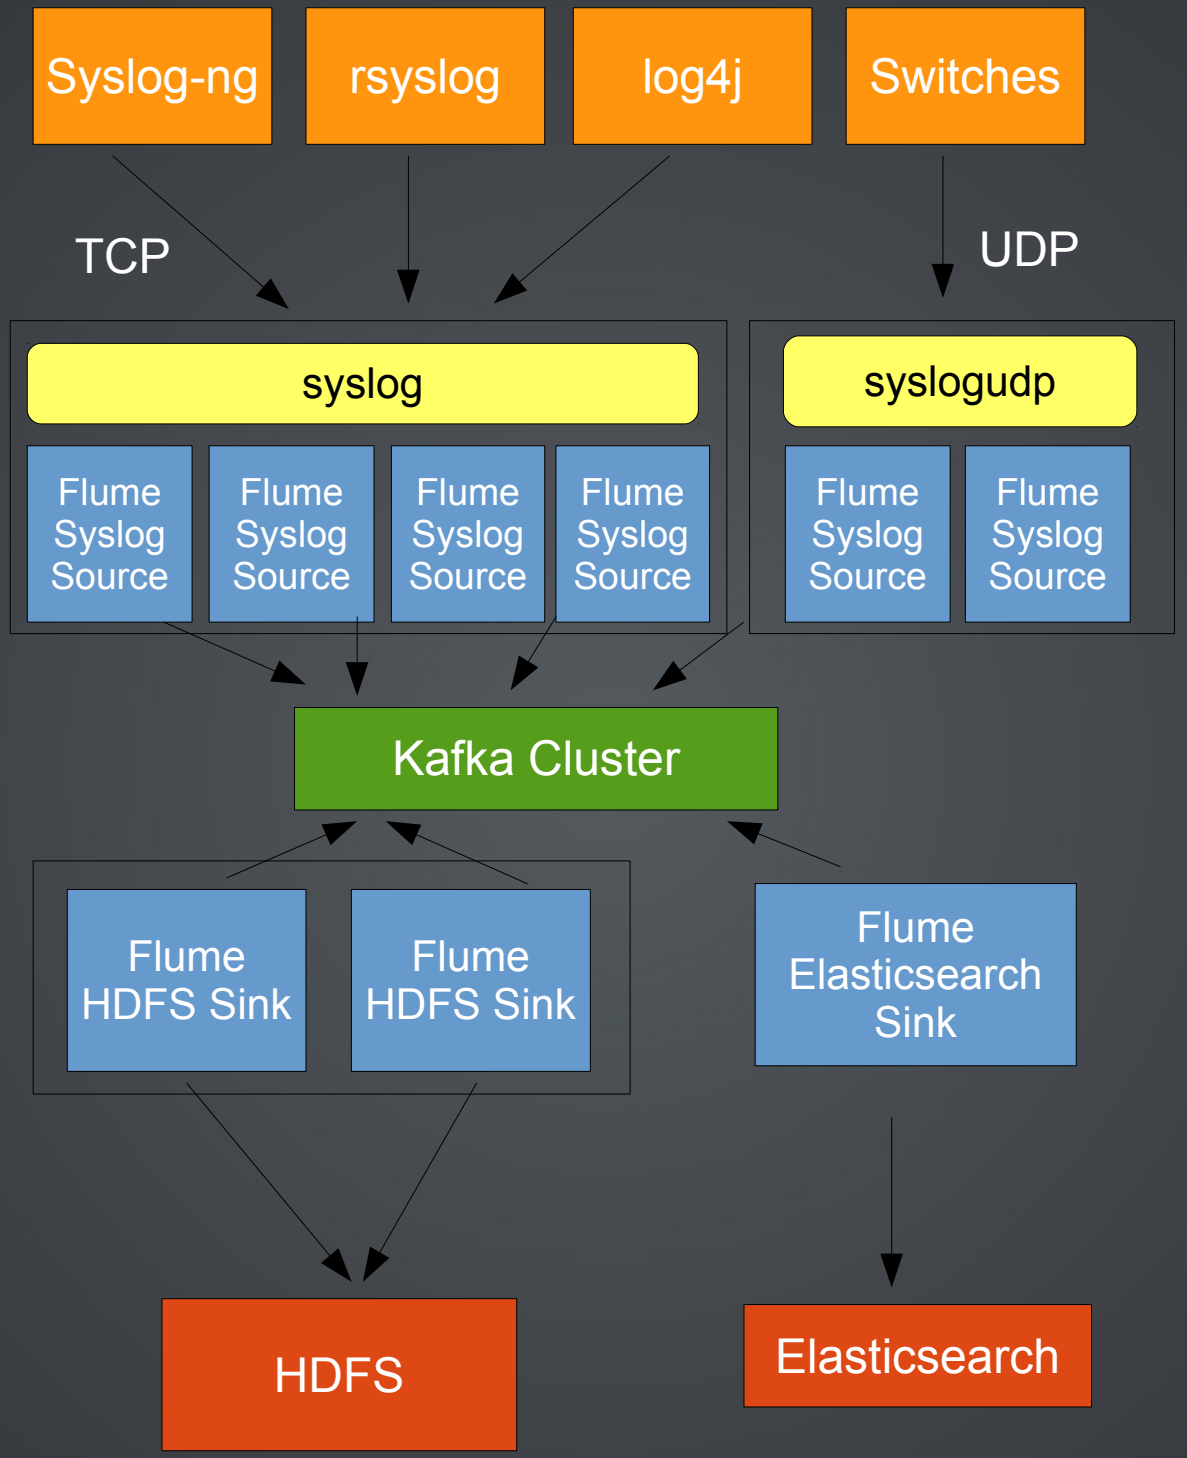
Syslog-ng
rsyslog
log4j
Switches
TCP
UDP
syslog
syslogudp
Flume Syslog Source
Flume Syslog Source
Flume Syslog Source
Flume Syslog Source
Flume Syslog Source
Flume Syslog Source
Kafka Cluster
Flume HDFS Sink
Flume HDFS Sink
Flume Elasticsearch Sink
HDFS
Elasticsearch

Phase 2: Understand

Dashboards

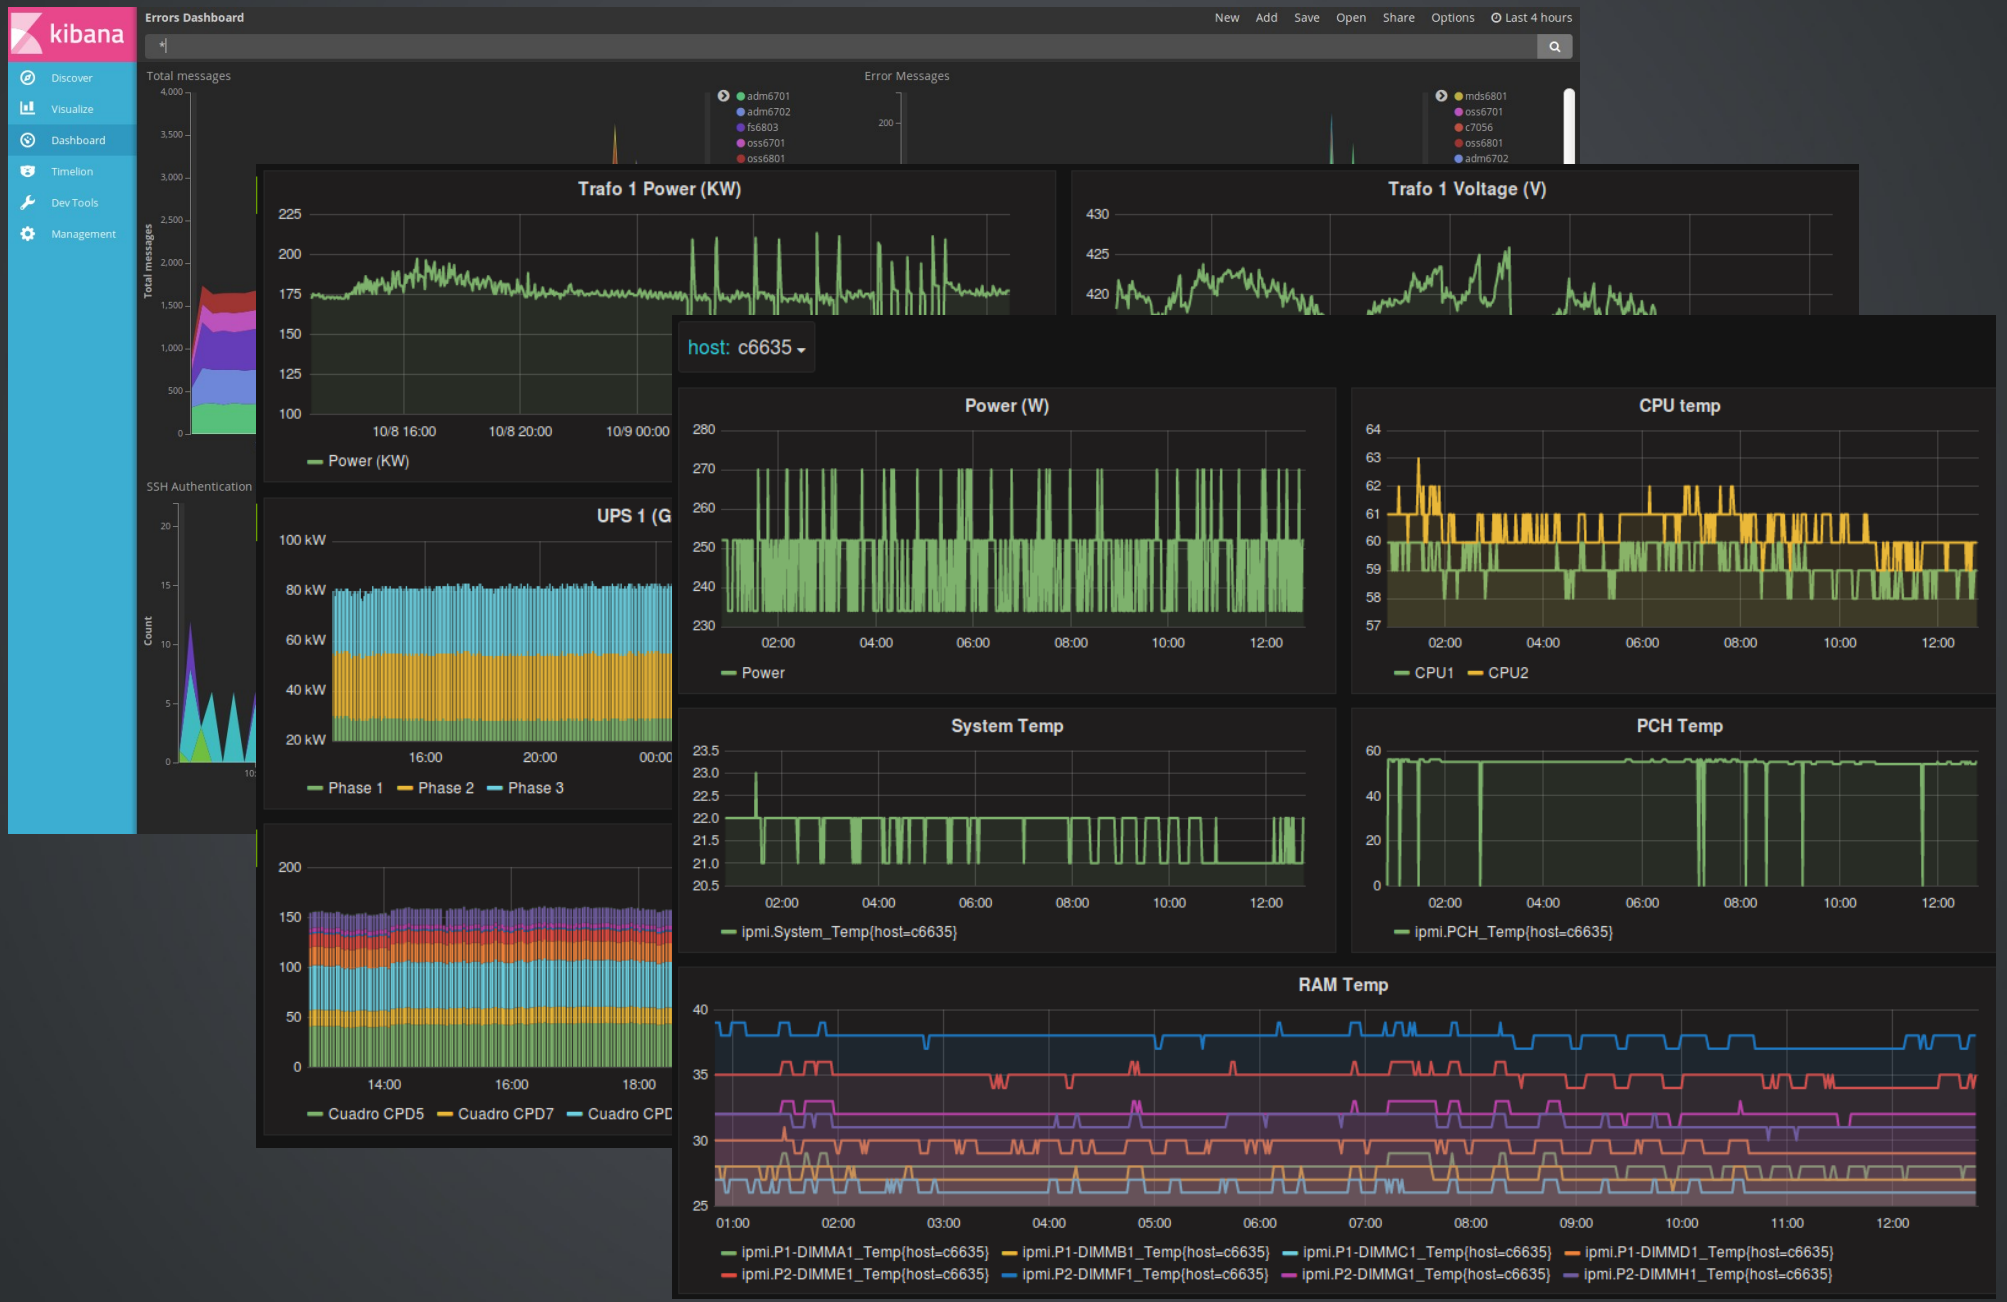
kibana
Errors Dashboard
New Add Save Open Share Options Last 4 hours
Discover
Visualize
Dashboard
Timelion
Dev Tools
Management
Total messages
Error Messages
SSH Authentication
Trafo 1 Power (KW)
Power (KW)
Trafo 1 Voltage (V)
UPS 1 (G
Phase 1 Phase 2 Phase 3
Cuadro CPD5 Cuadro CPD7
host: c6635
Power (W)
Power
CPU temp
CPU1 CPU2
System Temp
ipmi.System_Temp{host=c6635}
PCH Temp
ipmi.PCH_Temp{host=c6635}
RAM Temp
ipmi.P1-DIMMA1_Temp{host=c6635}
ipmi.P1-DIMMB1_Temp{host=c6635}
ipmi.P1-DIMMC1_Temp{host=c6635}
ipmi.P1-DIMMD1_Temp{host=c6635}
ipmi.P2-DIMME1_Temp{host=c6635}
ipmi.P2-DIMMF1_Temp{host=c6635}
ipmi.P2-DIMMG1_Temp{host=c6635}
ipmi.P2-DIMMH1_Temp{host=c6635}

DCIM

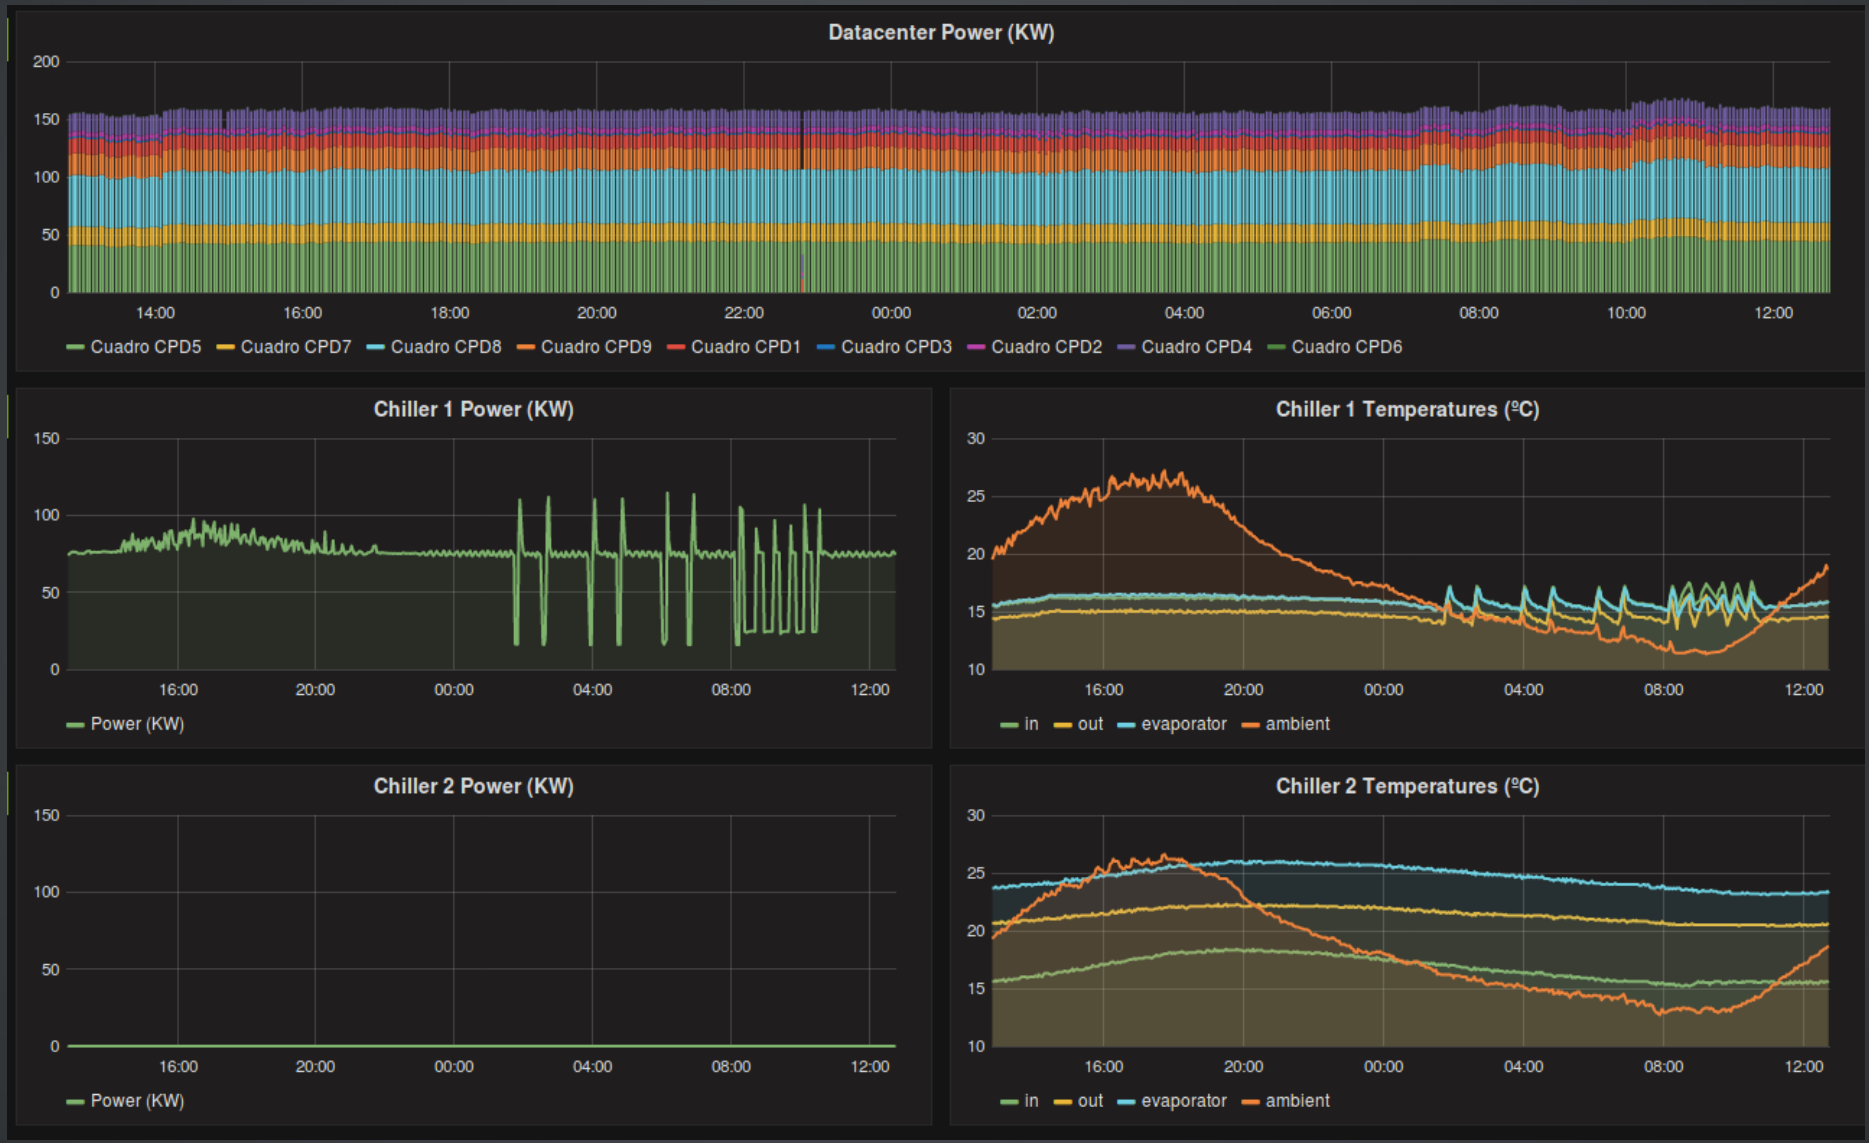
Datacenter Power (KW)
200
150
100
50
0
14:00
16:00
18:00
20:00
22:00
00:00
02:00
04:00
06:00
08:00
10:00
12:00
Cuadro CPD5
Cuadro CPD7
Cuadro CPD8
Cuadro CPD9
Cuadro CPD1
Cuadro CPD3
Cuadro CPD2
Cuadro CPD4
Cuadro CPD6
Chiller 1 Power (KW)
150
100
50
0
16:00
20:00
00:00
04:00
08:00
12:00
Power (KW)
Chiller 1 Temperatures (ºC)
30
25
20
15
10
16:00
20:00
00:00
04:00
08:00
12:00
in
out
evaporator
ambient
Chiller 2 Power (KW)
150
100
50
0
16:00
20:00
00:00
04:00
08:00
12:00
Power (KW)
Chiller 2 Temperatures (ºC)
30
25
20
15
10
16:00
20:00
00:00
04:00
08:00
12:00
in
out
evaporator
ambient

IPMI


host: c6635
Power (W)
280
270
260
250
240
230
02:00
04:00
06:00
08:00
10:00
12:00
Power
CPU temp
64
63
62
61
60
59
58
57
02:00
04:00
06:00
08:00
10:00
12:00
CPU1
CPU2
System Temp
23.5
23.0
22.5
22.0
21.5
21.0
20.5
02:00
04:00
06:00
08:00
10:00
12:00
ipmi.System_Temp{host=c6635}
PCH Temp
60
40
20
0
02:00
04:00
06:00
08:00
10:00
12:00
ipmi.PCH_Temp{host=c6635}
RAM Temp
40
35
30
25
01:00
02:00
03:00
04:00
05:00
06:00
07:00
08:00
09:00
10:00
11:00
12:00
ipmi.P1-DIMMA1_Temp{host=c6635}
ipmi.P1-DIMMB1_Temp{host=c6635}
ipmi.P1-DIMMC1_Temp{host=c6635}
ipmi.P1-DIMMD1_Temp{host=c6635}
ipmi.P2-DIMME1_Temp{host=c6635}
ipmi.P2-DIMMF1_Temp{host=c6635}
ipmi.P2-DIMMG1_Temp{host=c6635}
ipmi.P2-DIMMH1_Temp{host=c6635}

System Metrics

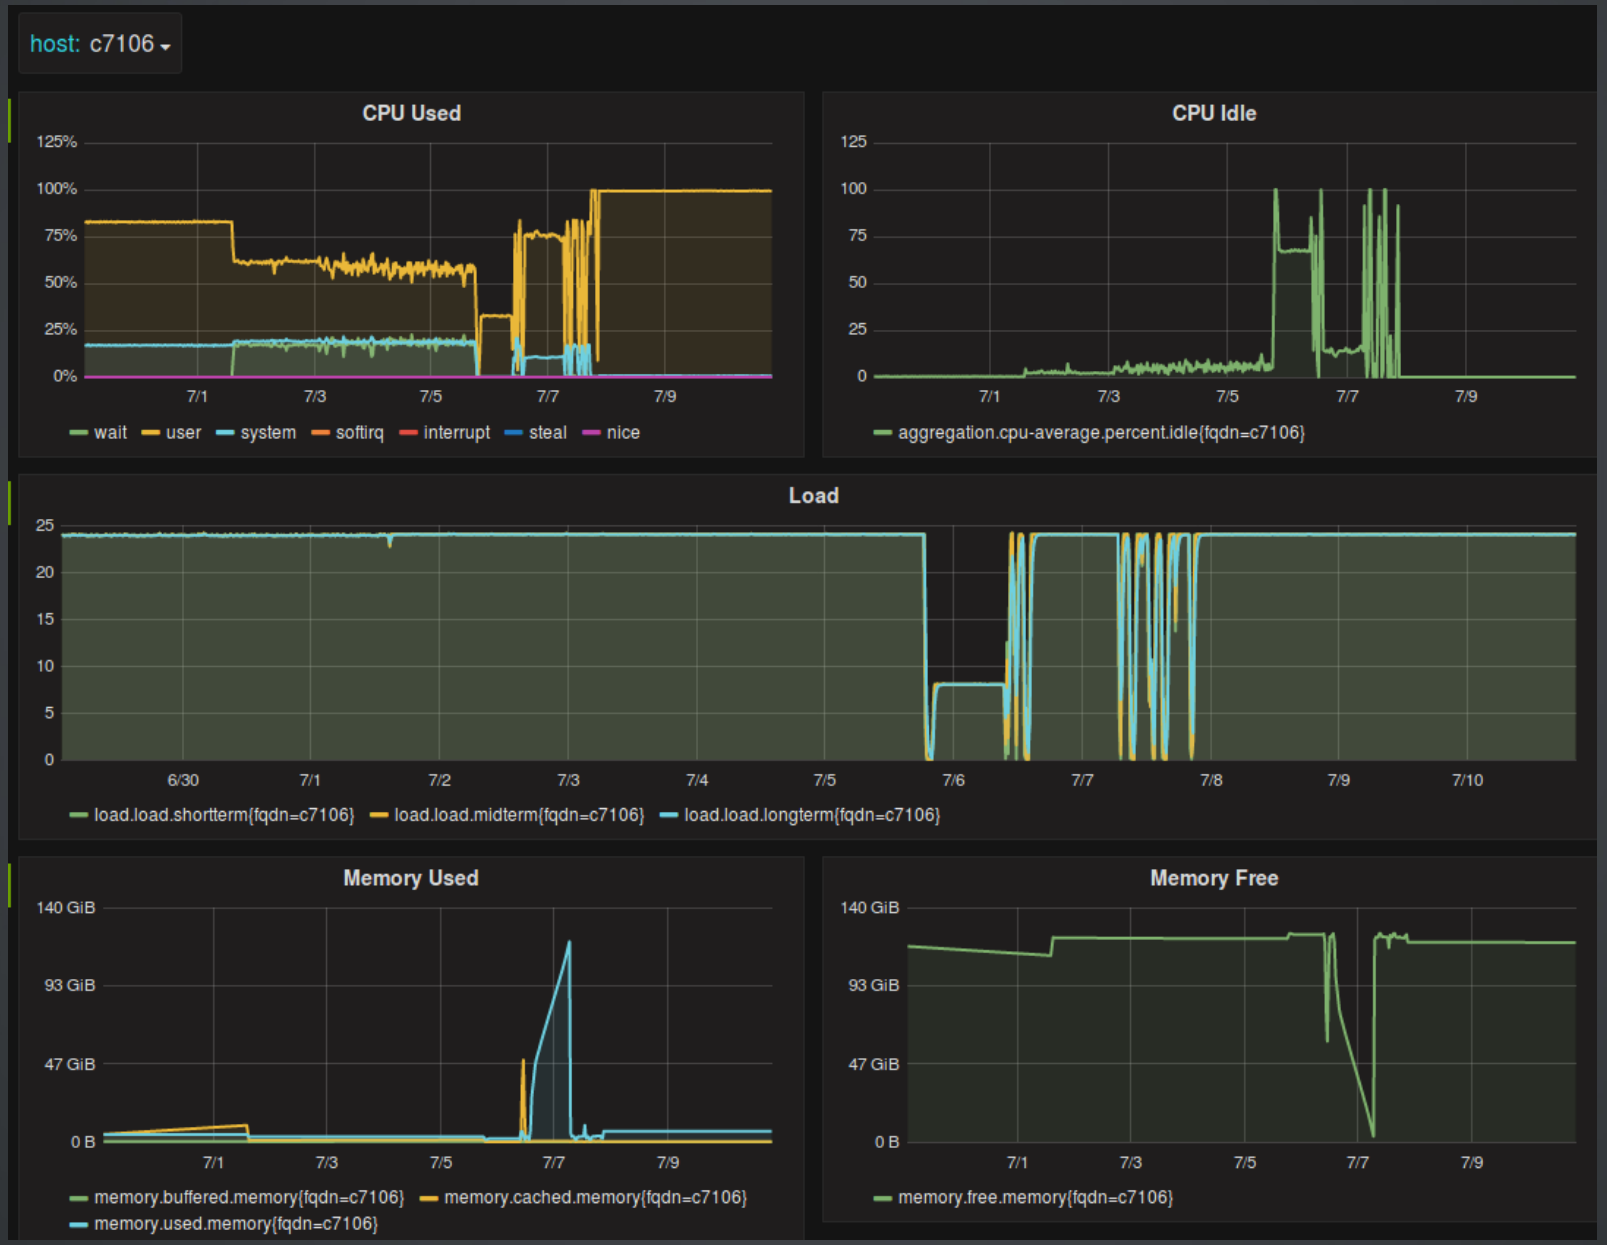

SMART

Current Pending Sectors

	max	total
smart.Current_Pending_Sector.raw_value{host=c14-5.local}	1.707 K	40.968 K
smart.Current_Pending_Sector.raw_value{host=c13-1.node.int.cesga.es}	11	264
smart.Current_Pending_Sector.raw_value{host=c13-3.node.int.cesga.es}	10	240
smart.Current_Pending_Sector.raw_value{host=c13-13.node.int.cesga.es}	5	120
smart.Current_Pending_Sector.raw_value{host=c13-2.node.int.cesga.es}	4	96
smart.Current_Pending_Sector.raw_value{host=c13-4.node.int.cesga.es}	3	72
smart.Current_Pending_Sector.raw_value{host=c13-11.node.int.cesga.es}	2	48
smart.Current_Pending_Sector.raw_value{host=c14-17.local}	0	0
smart.Current_Pending_Sector.raw_value{host=c14-13}	0	0
smart.Current_Pending_Sector.raw_value{host=c14-15.local}	0	0

Reallocated Sector Count

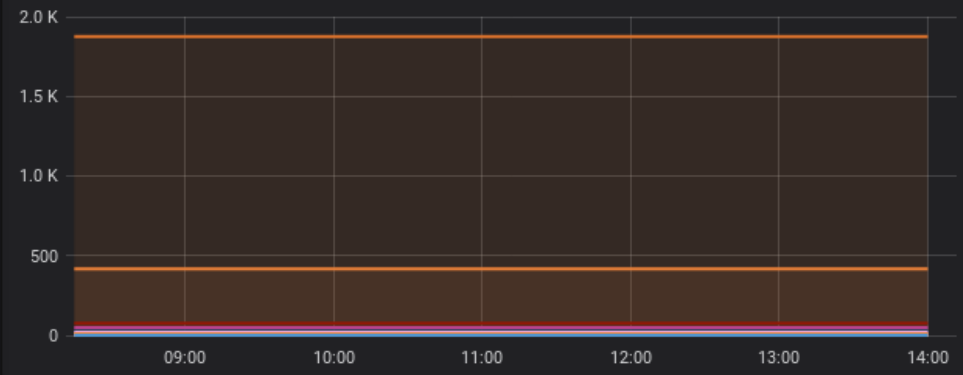

	max	total
smart.Reallocated_Sector_Ct.raw_value{host=c14-5.local}	1.876 K	45.024 K
smart.Reallocated_Sector_Ct.raw_value{host=c13-16.node.int.cesga.es}	417	10.008 K
smart.Reallocated_Sector_Ct.raw_value{host=c13-8.node.int.cesga.es}	76	1.824 K
smart.Reallocated_Sector_Ct.raw_value{host=c13-17.node.int.cesga.es}	49	1.176 K
smart.Reallocated_Sector_Ct.raw_value{host=c13-11.node.int.cesga.es}	28	672
smart.Reallocated_Sector_Ct.raw_value{host=c14-14.node.int.cesga.es}	20	480
smart.Reallocated_Sector_Ct.raw_value{host=c13-3.node.int.cesga.es}	20	480
smart.Reallocated_Sector_Ct.raw_value{host=c14-8.local}	9	216
smart.Reallocated_Sector_Ct.raw_value{host=c13-13.node.int.cesga.es}	7	168
smart.Reallocated_Sector_Ct.raw_value{host=c13-2.node.int.cesga.es}	5	120

Offline Uncorrectable

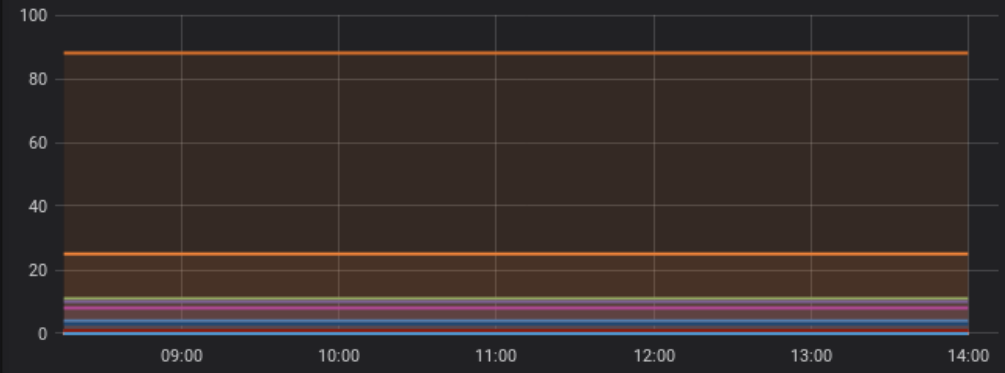

	max	total
smart.Offline_Uncorrectable.raw_value{host=c14-5.local}	88	2.1120 K
smart.Offline_Uncorrectable.raw_value{host=c13-16.node.int.cesga.es}	25	600
smart.Offline_Uncorrectable.raw_value{host=c13-1.node.int.cesga.es}	11	264
smart.Offline_Uncorrectable.raw_value{host=c13-3.node.int.cesga.es}	10	240
smart.Offline_Uncorrectable.raw_value{host=c13-17.node.int.cesga.es}	8	192
smart.Offline_Uncorrectable.raw_value{host=c13-2.node.int.cesga.es}	4	96
smart.Offline_Uncorrectable.raw_value{host=c13-11.node.int.cesga.es}	3	72
smart.Offline_Uncorrectable.raw_value{host=c13-8.node.int.cesga.es}	1	24
smart.Offline_Uncorrectable.raw_value{host=c13-15.node.int.cesga.es}	1	24
smart.Offline_Uncorrectable.raw_value{host=c14-17.local}	0	0

Logs


kibana
Errors Dashboard
New
Add
Save
Open
Share
Options
Last 4 hours
*
Discover
Visualize
Dashboard
Timelion
Dev Tools
Management
Total messages
Total messages
Date
adm6701
adm6702
fs6803
oss6701
oss6801
c13-5
mds6801
c6601
Error Messages
Count
Date
mds6801
oss6701
c7056
oss6801
adm6702
c7228
fs6801
c7057
c7058
c7225
c7226
c6618
c7049
c7021
c7051
c6606
c6942
c6614
c6903
c7230
c6620
c6621
c6622
c6623
c6624
SSH Authentication Failures FT2
Count
Date
fs6801
fs6802
fs6803
c6601
adm6701
SSH Authentication Failures BD
No results found

jupyter Plot_Job_Metric-INTERACTIVELY Last Checkpoint: hace una hora (autosaved)

Logout

File Edit View Insert Cell Kernel Widgets Help

Trusted Python 2

Run Code

job_id 1292278

metric memcpy-socket0

frequency 15min

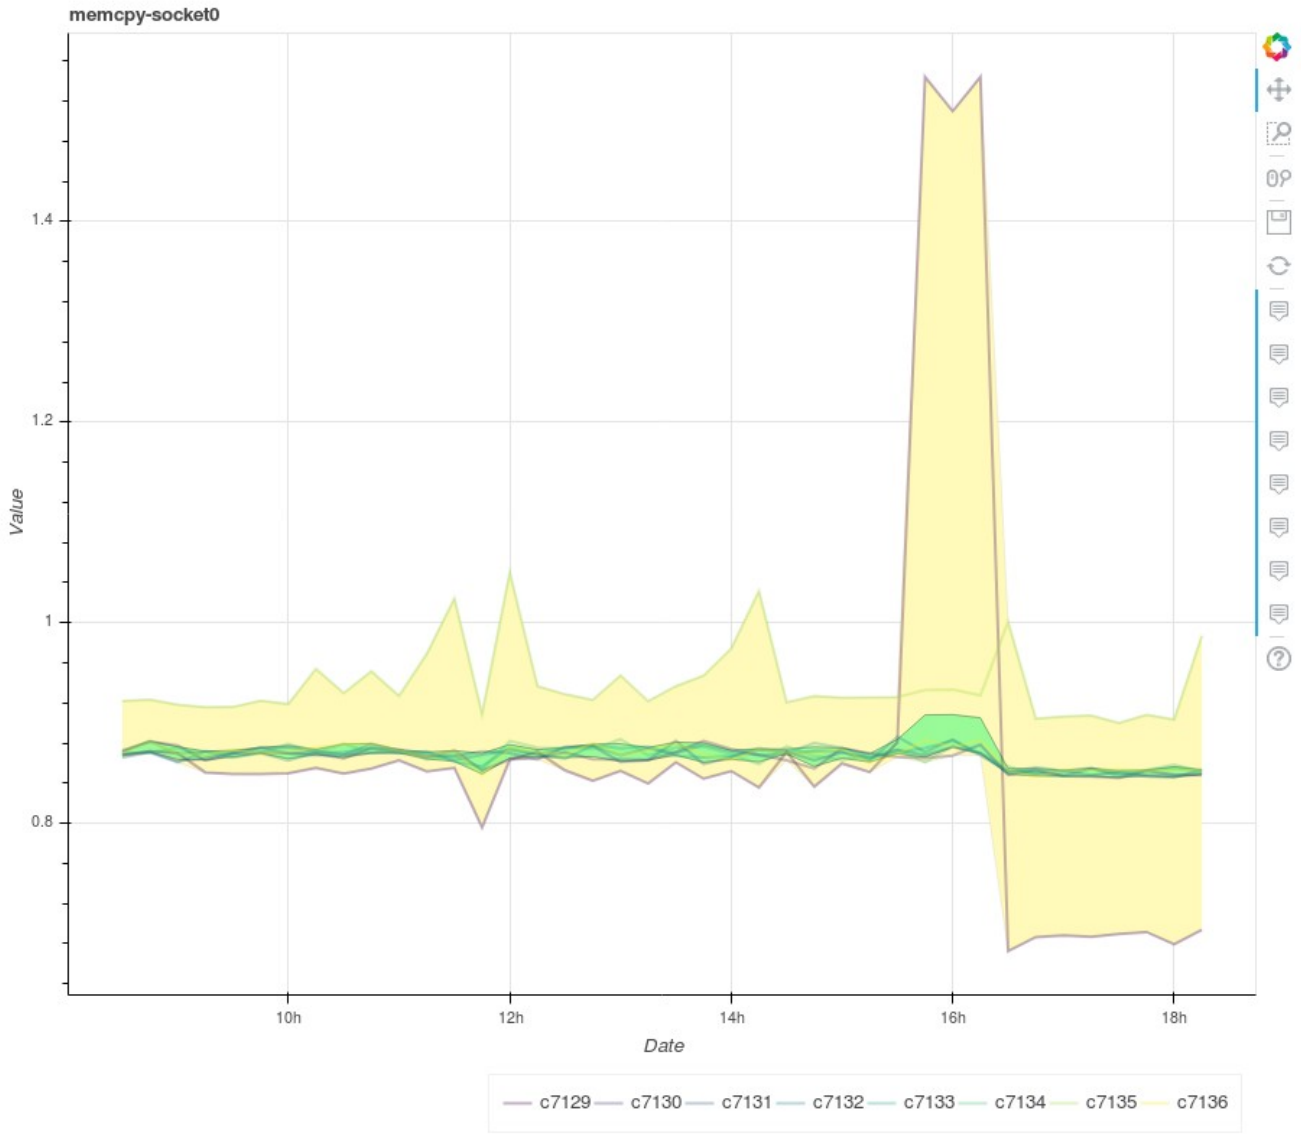

Bokeh App

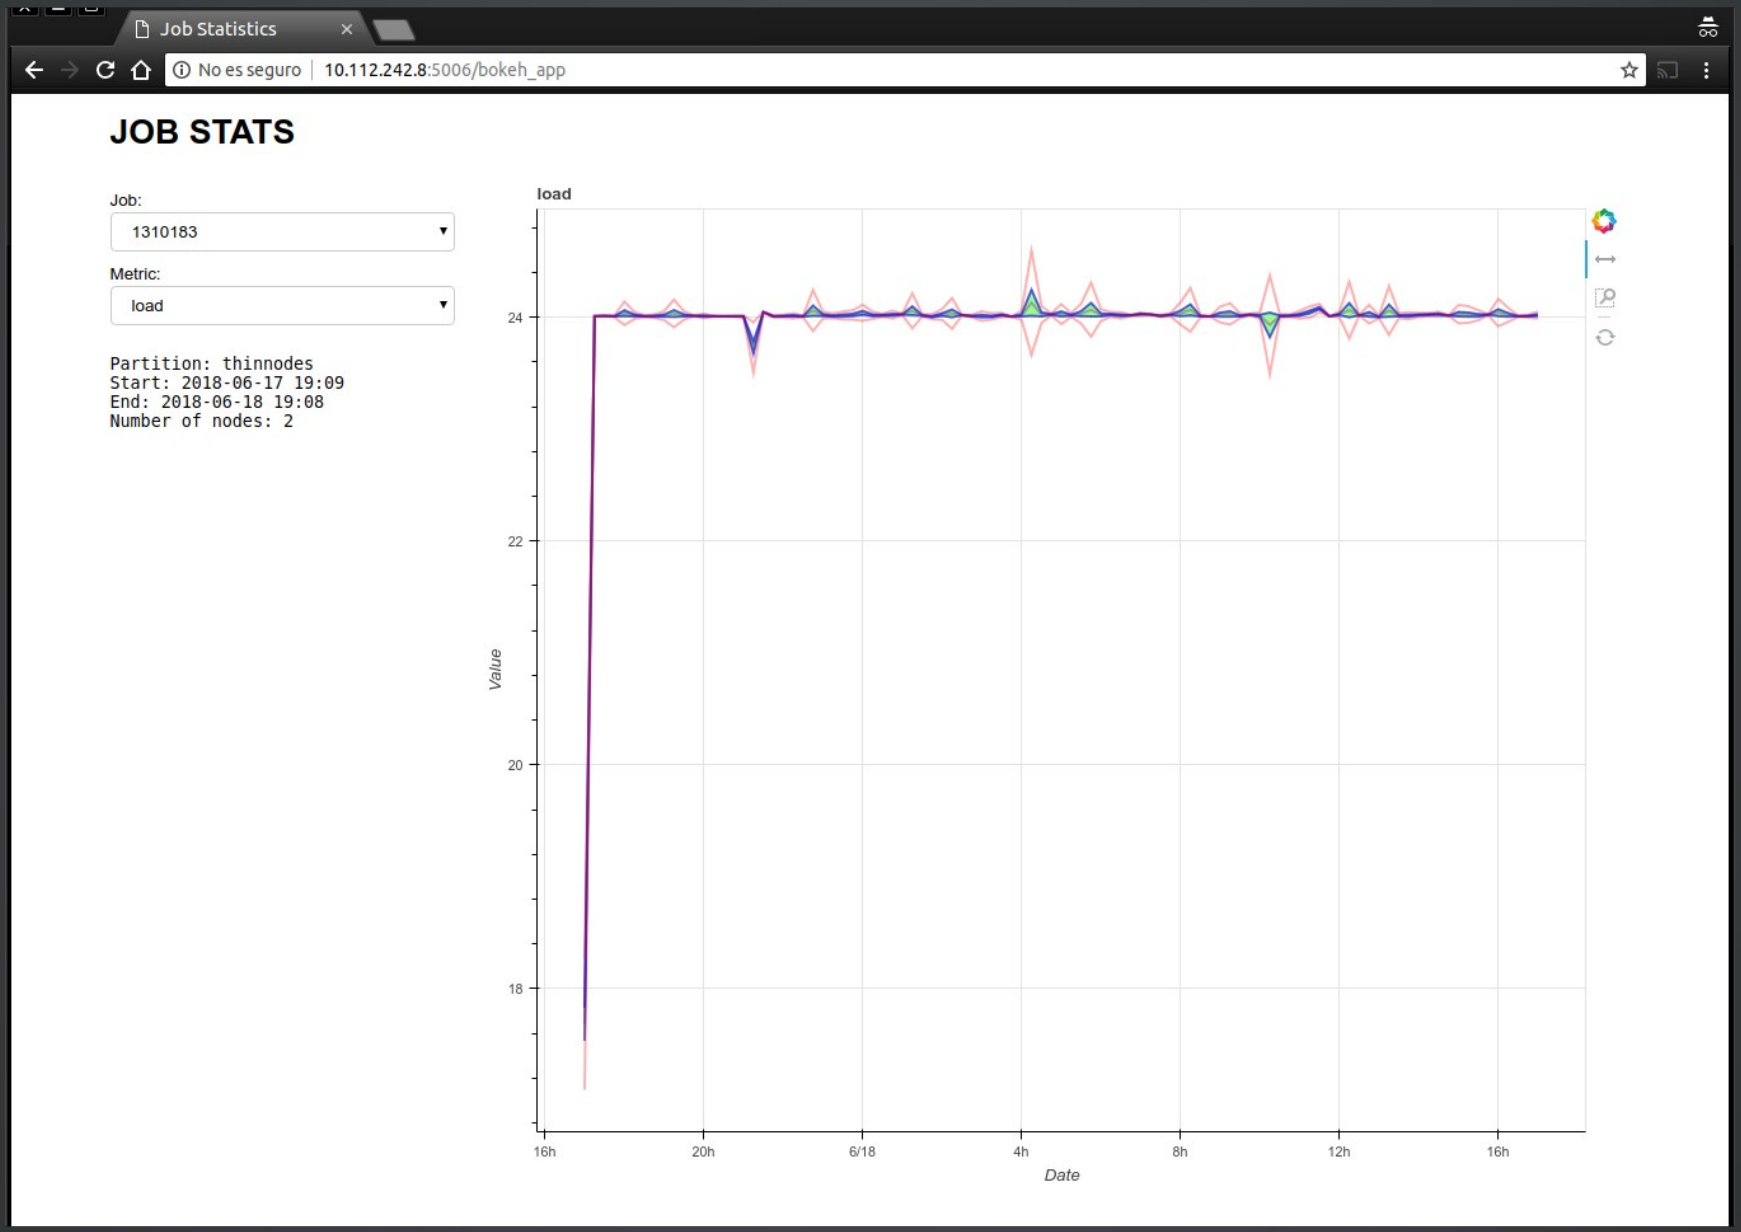
Job Statistics
No es seguro | 10.112.242.8:5006/bokeh_app
JOB STATS
Job:
1310183
Metric:
load
Partition: thinnodes
Start: 2018-06-17 19:09
End: 2018-06-18 19:08
Number of nodes: 2
load
24
22
20
18
Value
16h
20h
6/18
4h
8h
12h
16h
Date

Phase 3: Control

Log Alerts: Icinga

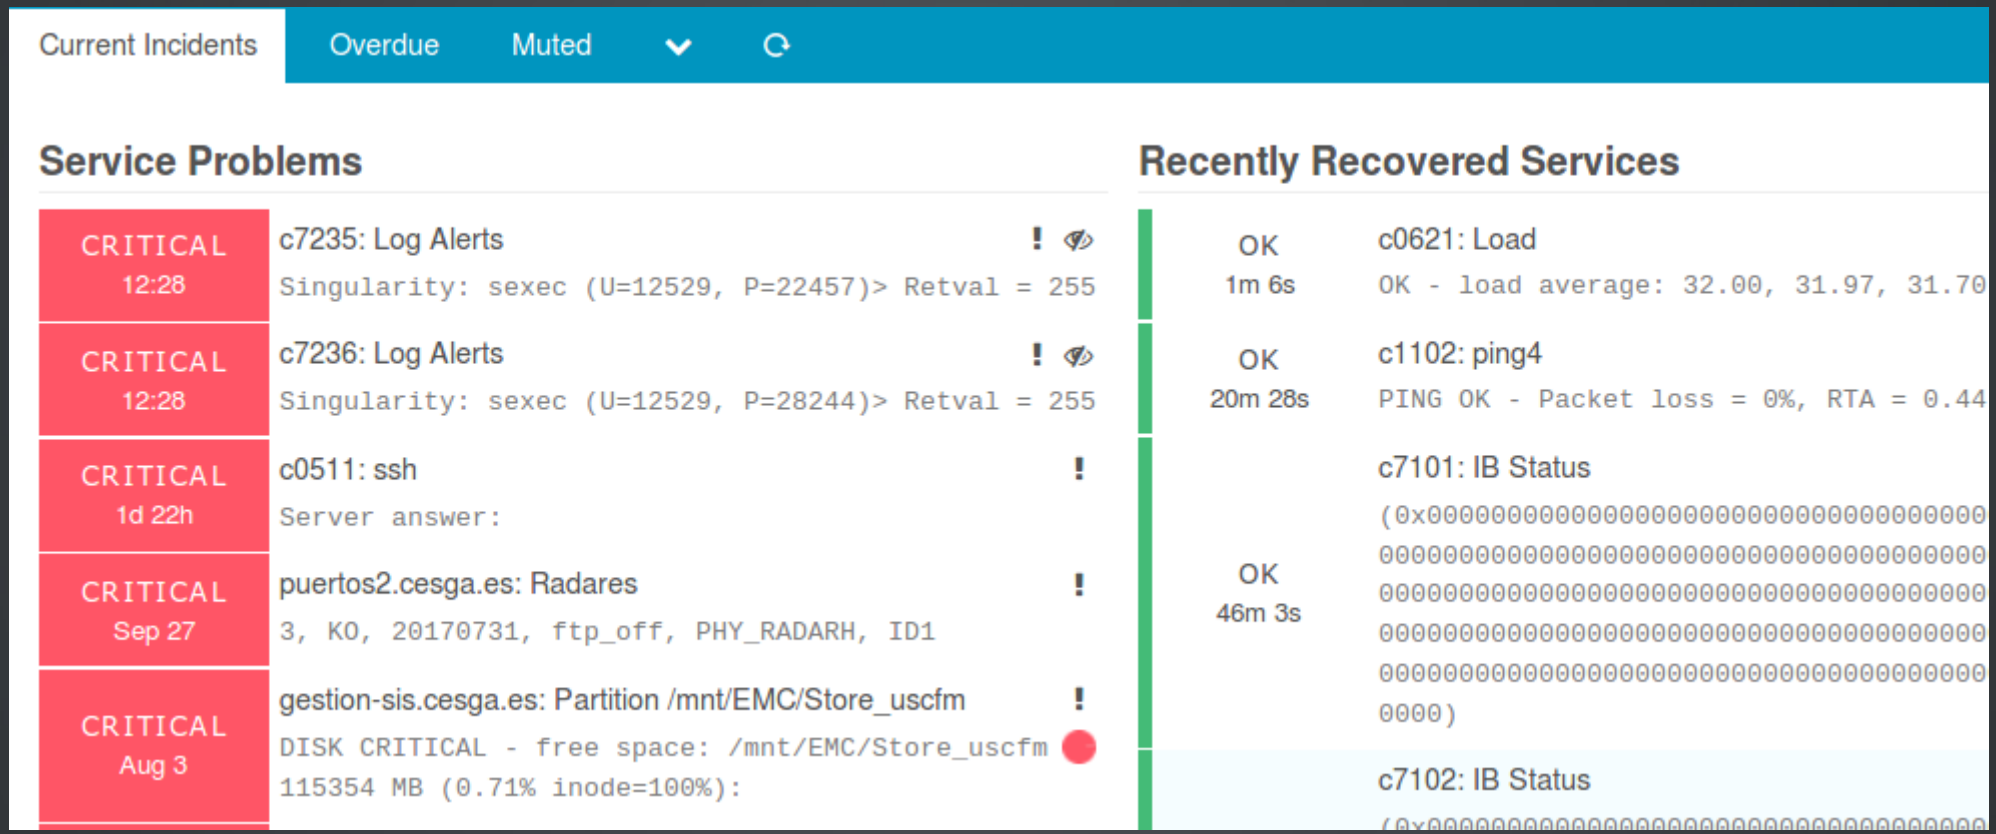

Metric Alerts: Bosun

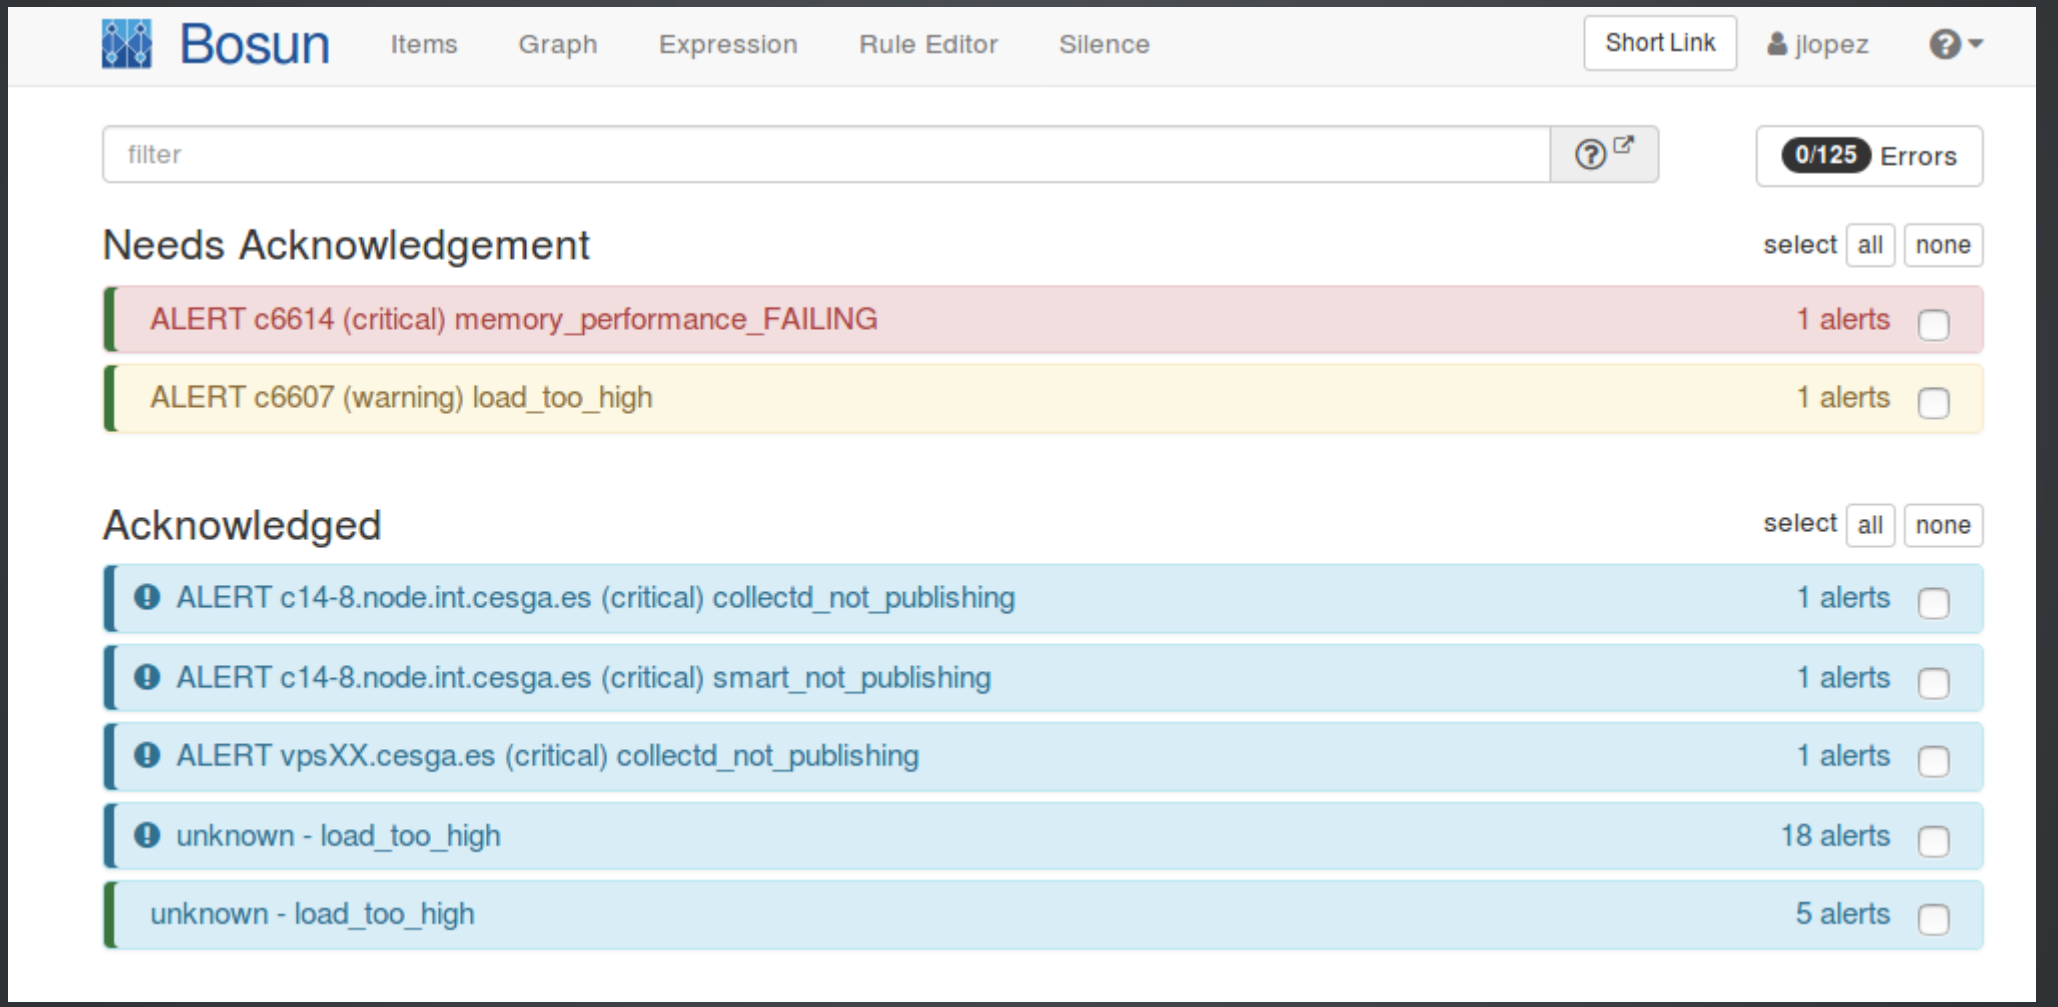

Phase 4: Improve

Anomaly Detection

Types of anomalies in time series

- Outliers
- Change points
- Anomalous time series

Generic Anomaly Detection

- Generic frameworks for Anomaly Detection
- They trigger too many false alerts
- In our experience it is better to trigger specific use cases

SSH Attack Detection

Problem: Daily our public servers are scanned and attacked

Detection: Correlate real-time SSH connection information to detect attacks

Objective: Automatically update router configuration to stop the attacks

SSH Attack Detection

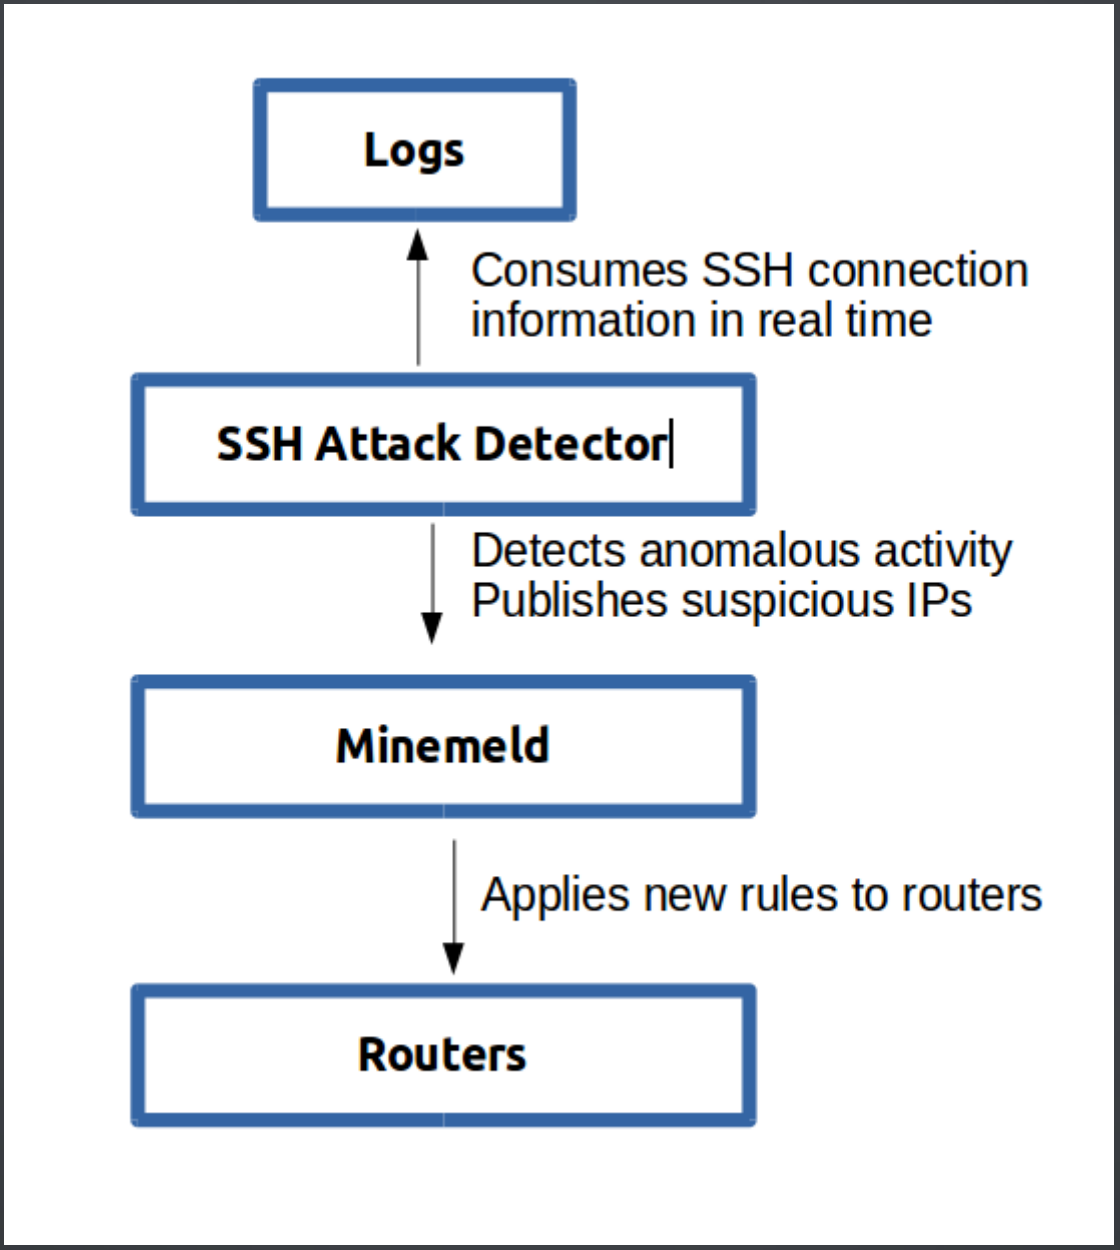

Results

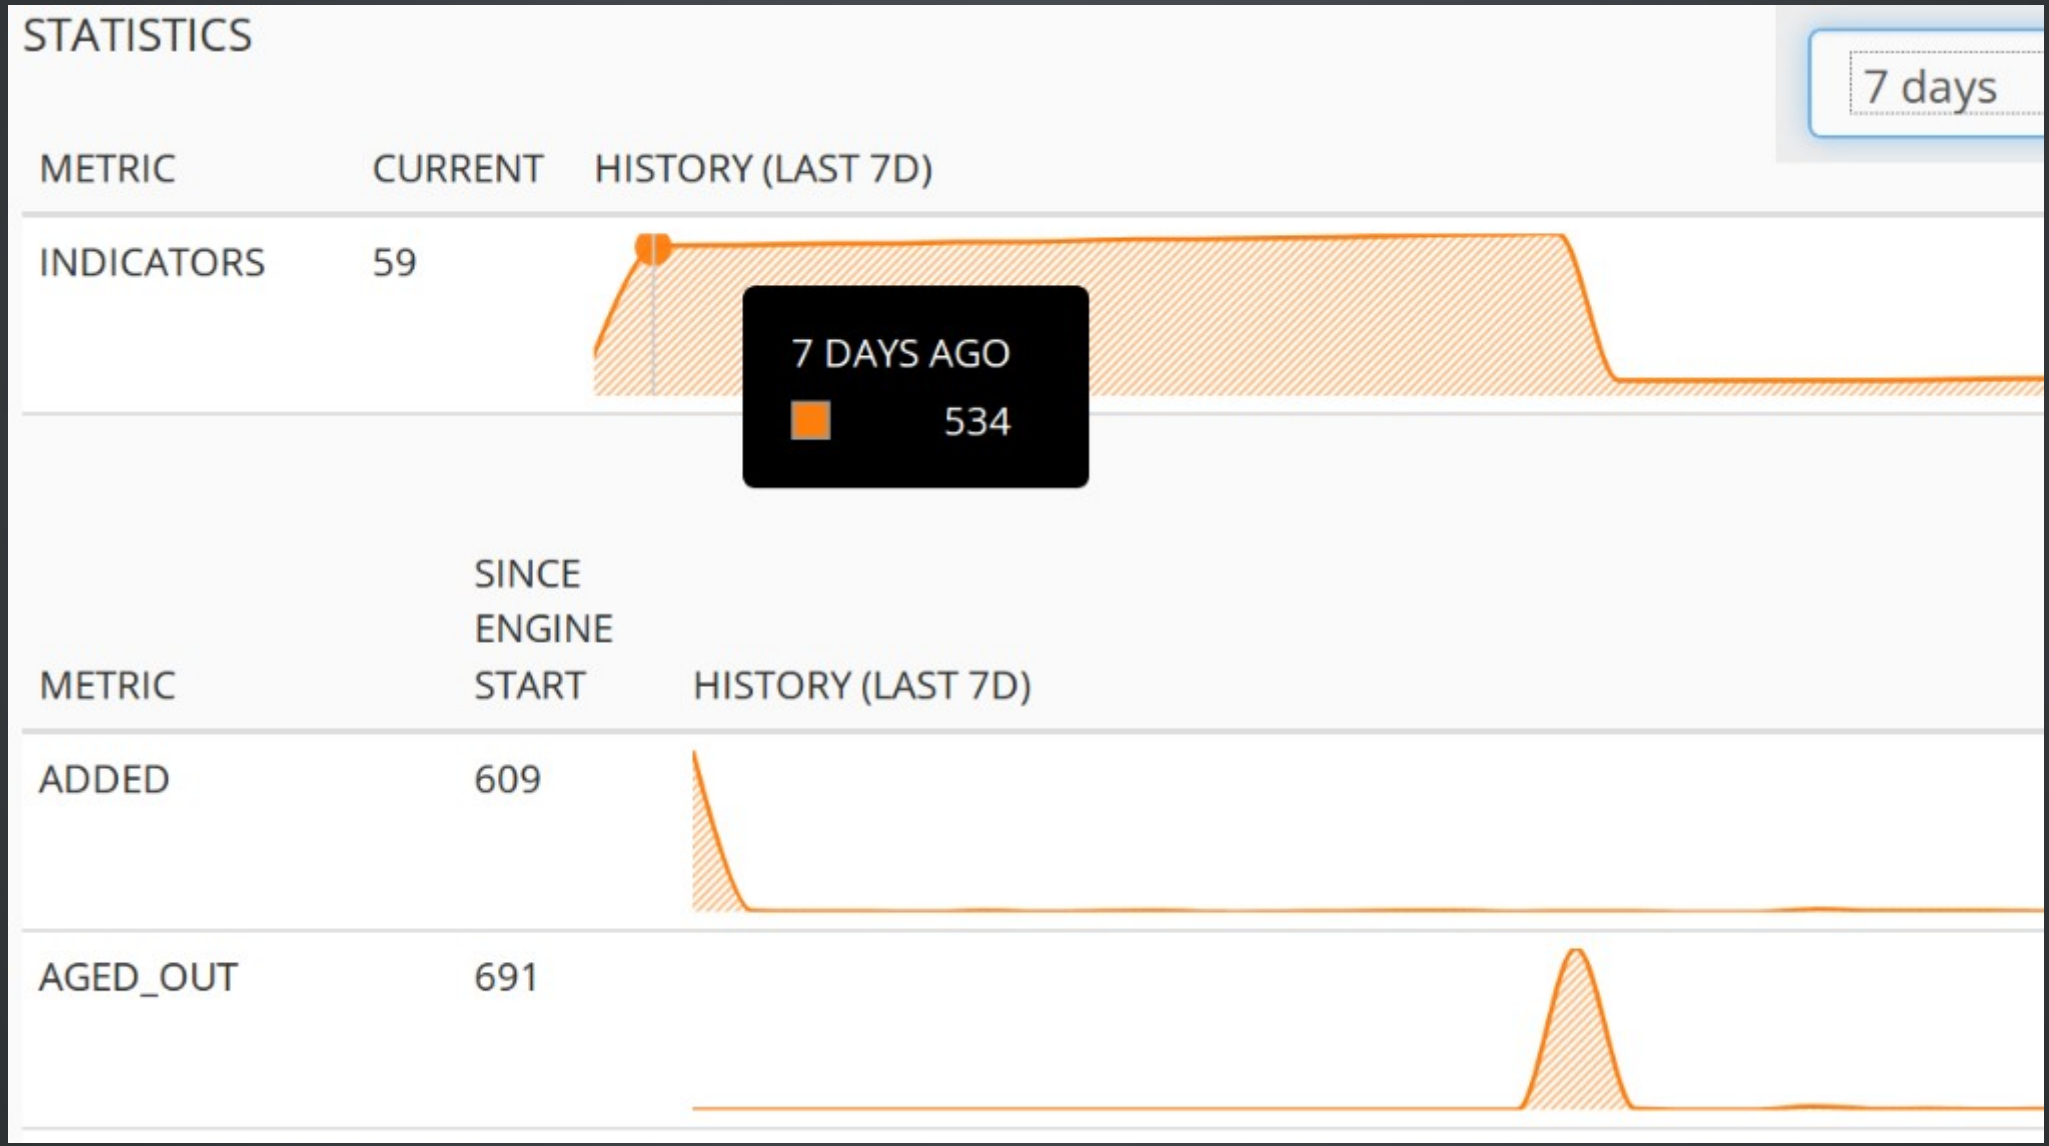
STATISTICS
7 days
METRIC
CURRENT
HISTORY (LAST 7D)
INDICATORS
59
7 DAYS AGO
534
METRIC
SINCE ENGINE START
HISTORY (LAST 7D)
ADDED
609
AGED_OUT
691

Server Anomalous Performance

Problem: Parallel jobs are cancelled because some of the nodes have poor performance. Computation is lost.

Detection: Analyze & visualize server metrics to spot the anomalous node

Objective: Automatically detect low performance nodes

Analyze & Visualize

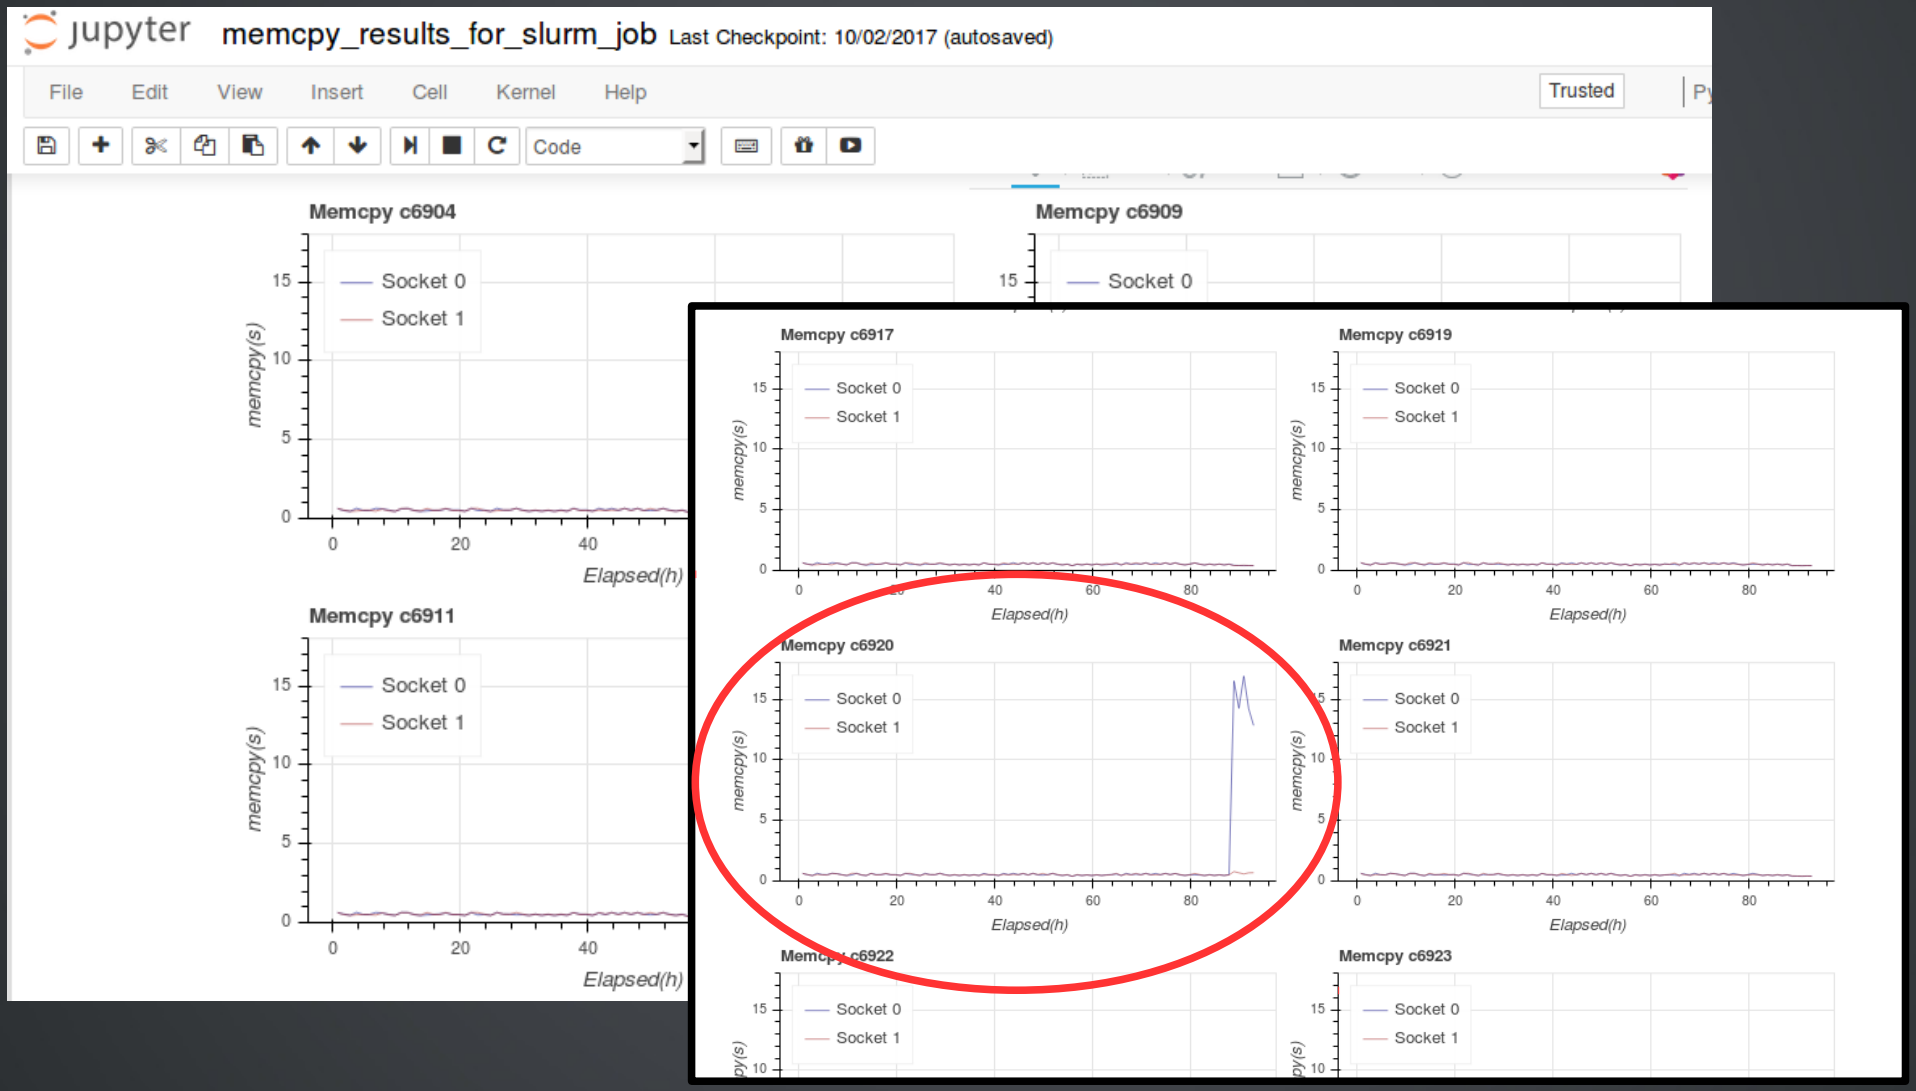

Anomalous Performance Detection

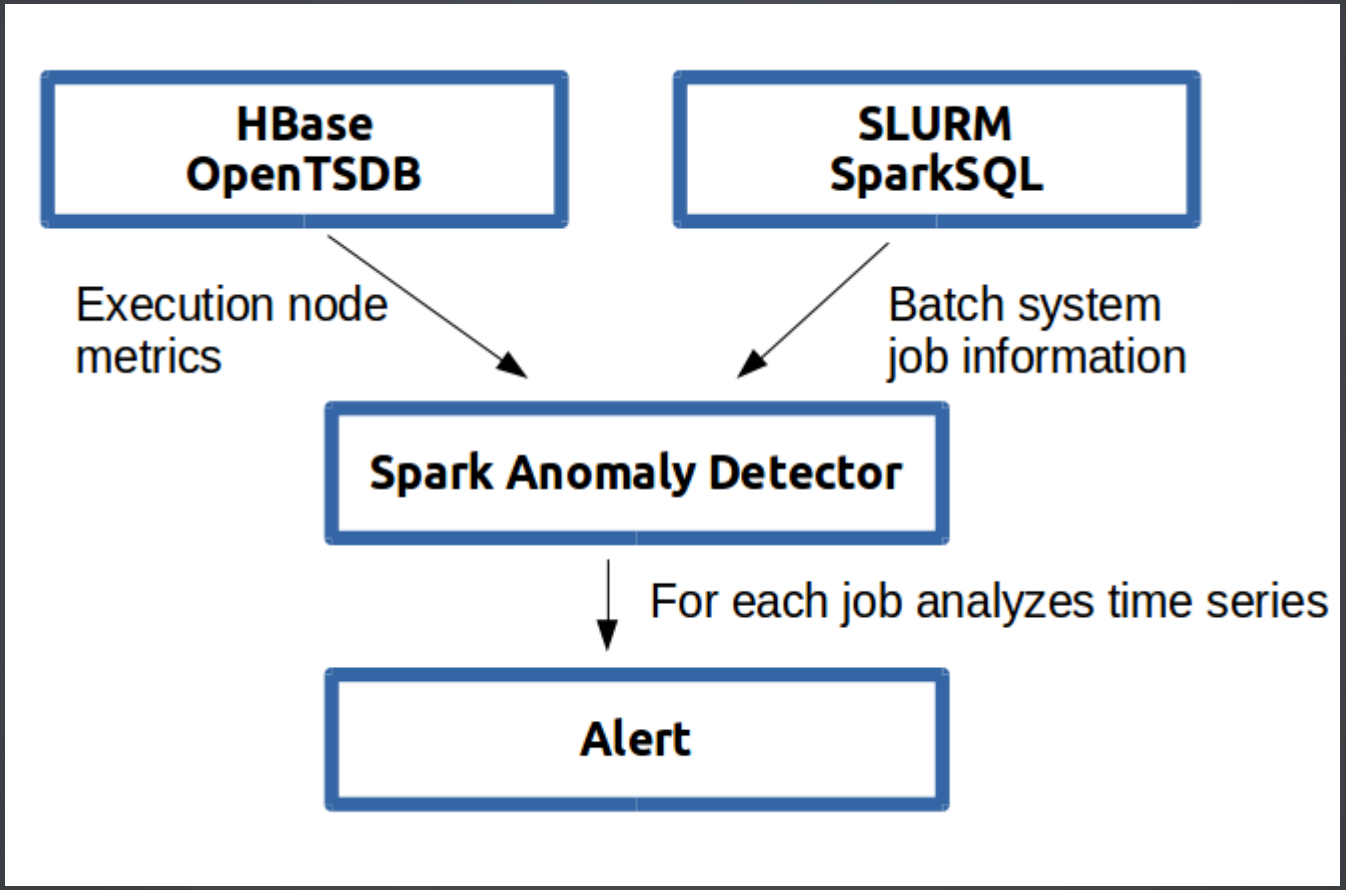

Conclusions

- No longer needed to delete old data
- Generic anomaly detecion systems generate too many alerts
- Target specific use cases to maintain number of alerts low

THANKS!

[jlopez [at] cesga [dot] es]