RIPE NCC

Building a logging pipeline with Open Source tools

Iñigo Ortiz de Urbina Cazenave

NLUUG | Utrecht - Netherlands | 28 May 2015

whoami;

- Iñigo Ortiz de Urbina Cazenave
- Systems Engineer

whoami; groups;

- Iñigo Ortiz de Urbina Cazenave
- Systems Engineer @ RIPE NCC
- RIPE NCC
  - RIR for *Europe, the Middle East, parts of Central Asia*
  - IP and ASN allocation, registration
  - RIPE DB
  - DNS
  - Routing Information Service
  - RIPE Stat
  - RIPE Atlas

What is RIPE Atlas?

v1 & v2: Lantronix XPort Pro

v3: TP-Link TL-MR3020

RIPE Atlas anchor: Soekris net6501-70

What is RIPE Atlas?

Largest active measurement network

https://atlas.ripe.net

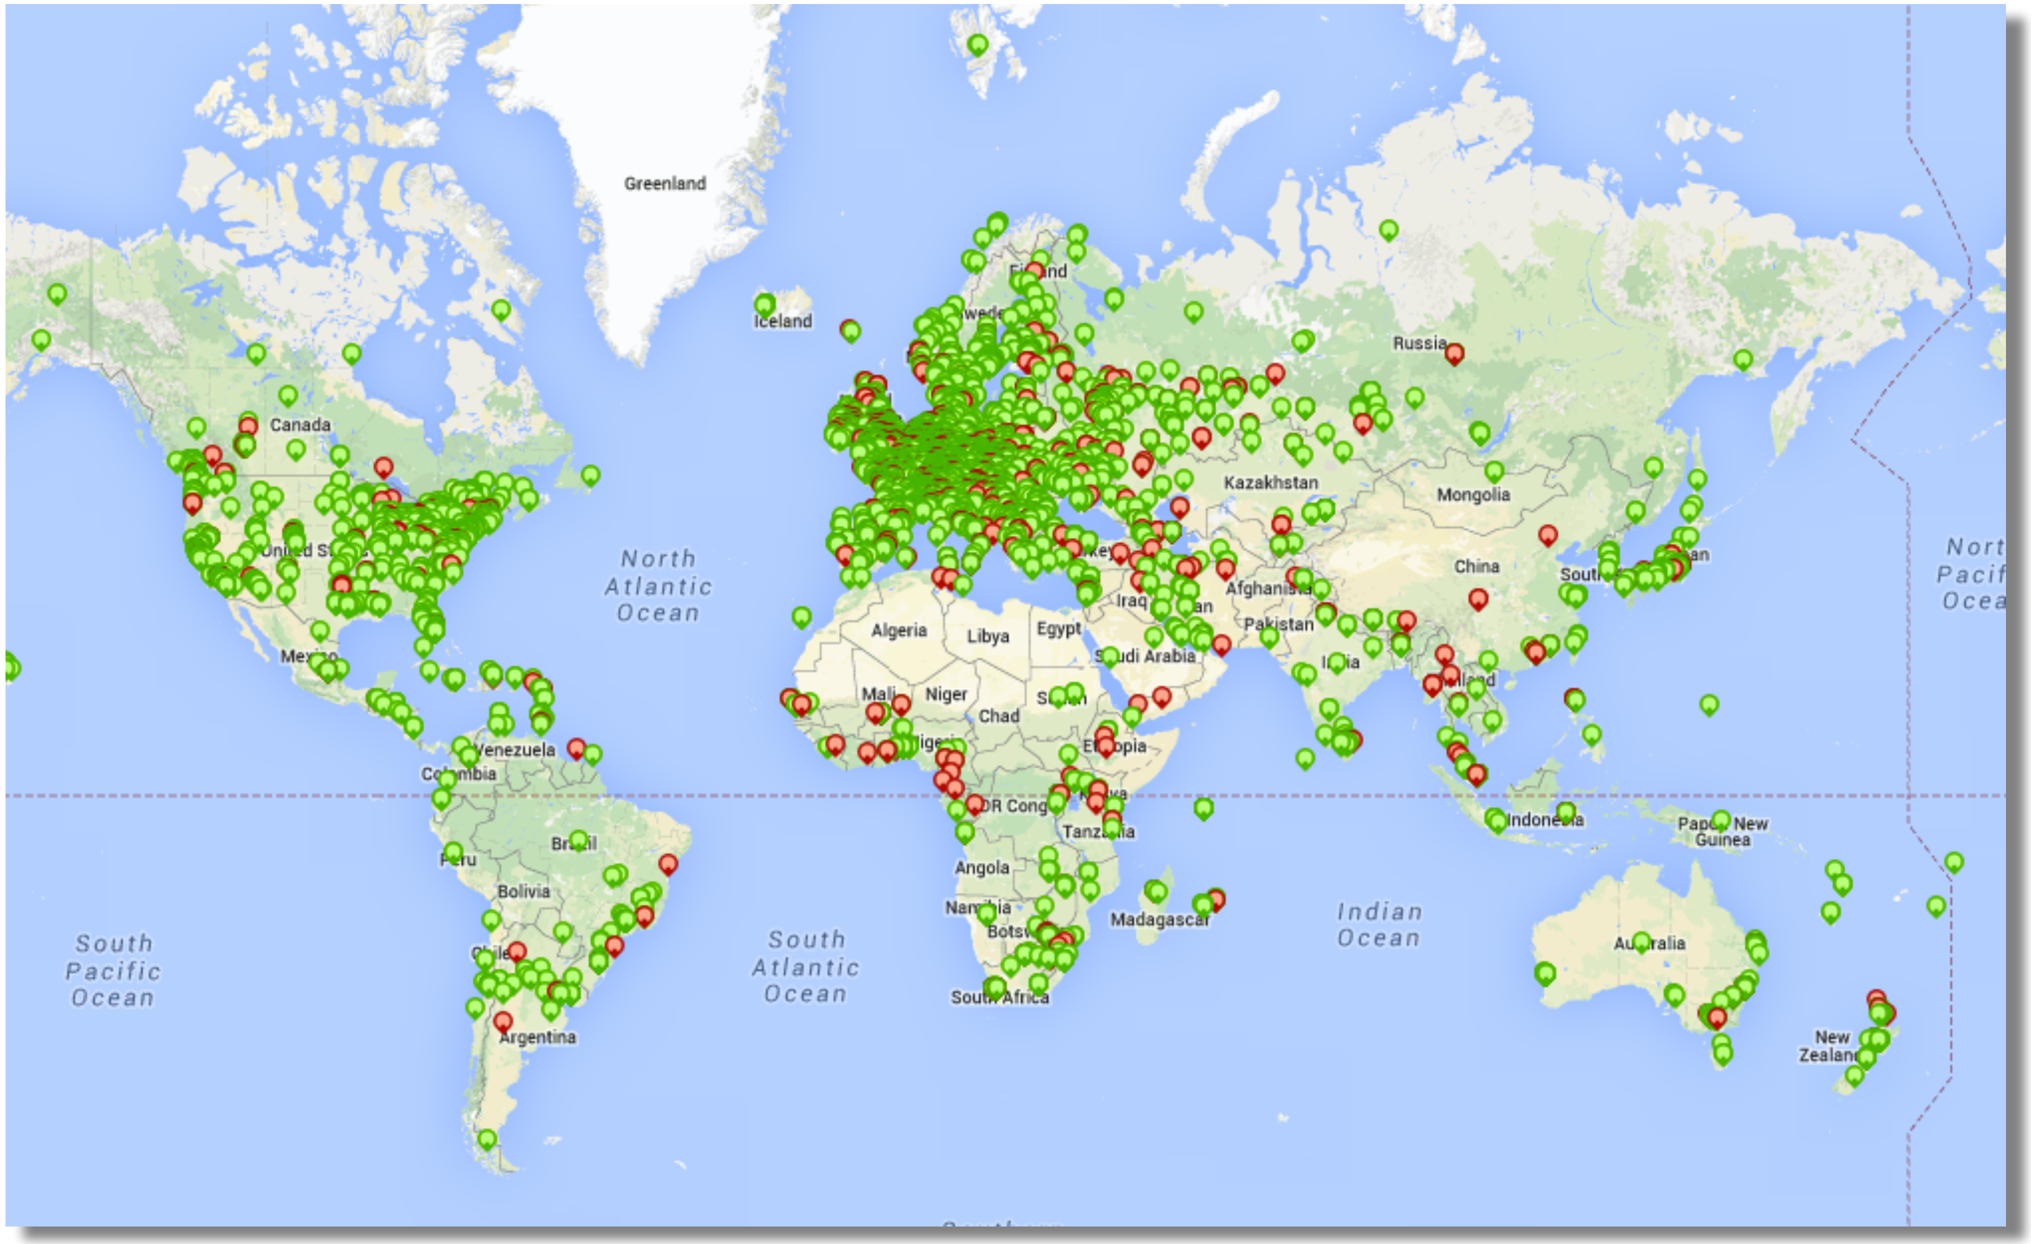

RIPE Atlas in numbers

- ~8200 probes online
- ~21000 users
- ~7800 ongoing measurements
  - ~250 built-ins
  - ~7550 user defined measurements
- ~2300 results per second
- ~40 servers

Infrastructure Overview

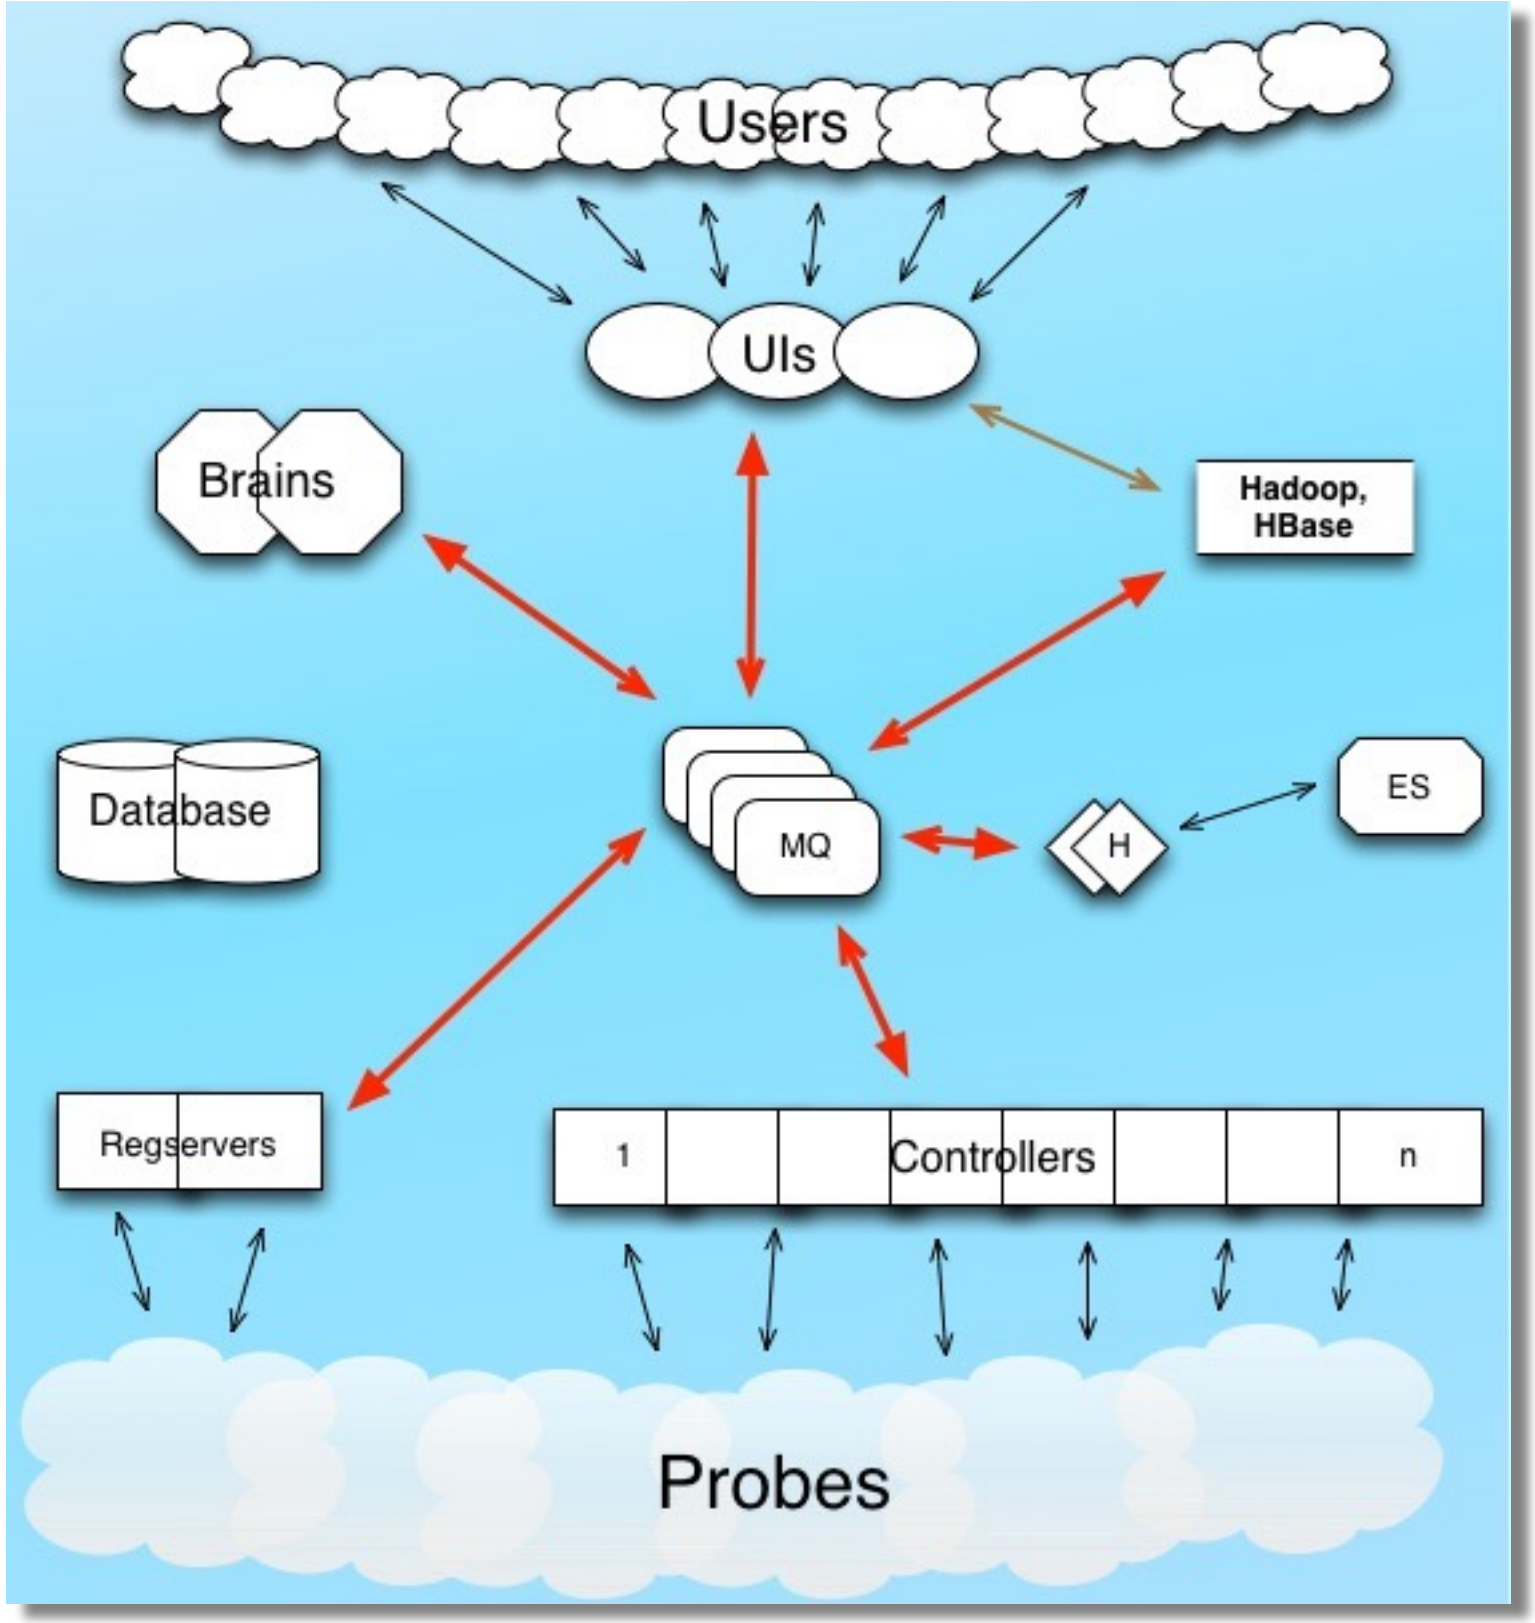

Log sources

- RIPE Atlas httpd access logs
- RIPE Atlas Software (warning, error, critical)
- Hadoop, HBase, zookeeper, Thrift
- Other:
  - Syslog
  - Custom scripts

The problem

- Production
- Collection
- Transport
- Queueing, buffering
- Massaging
- Storage
- Search and analysis

The problem

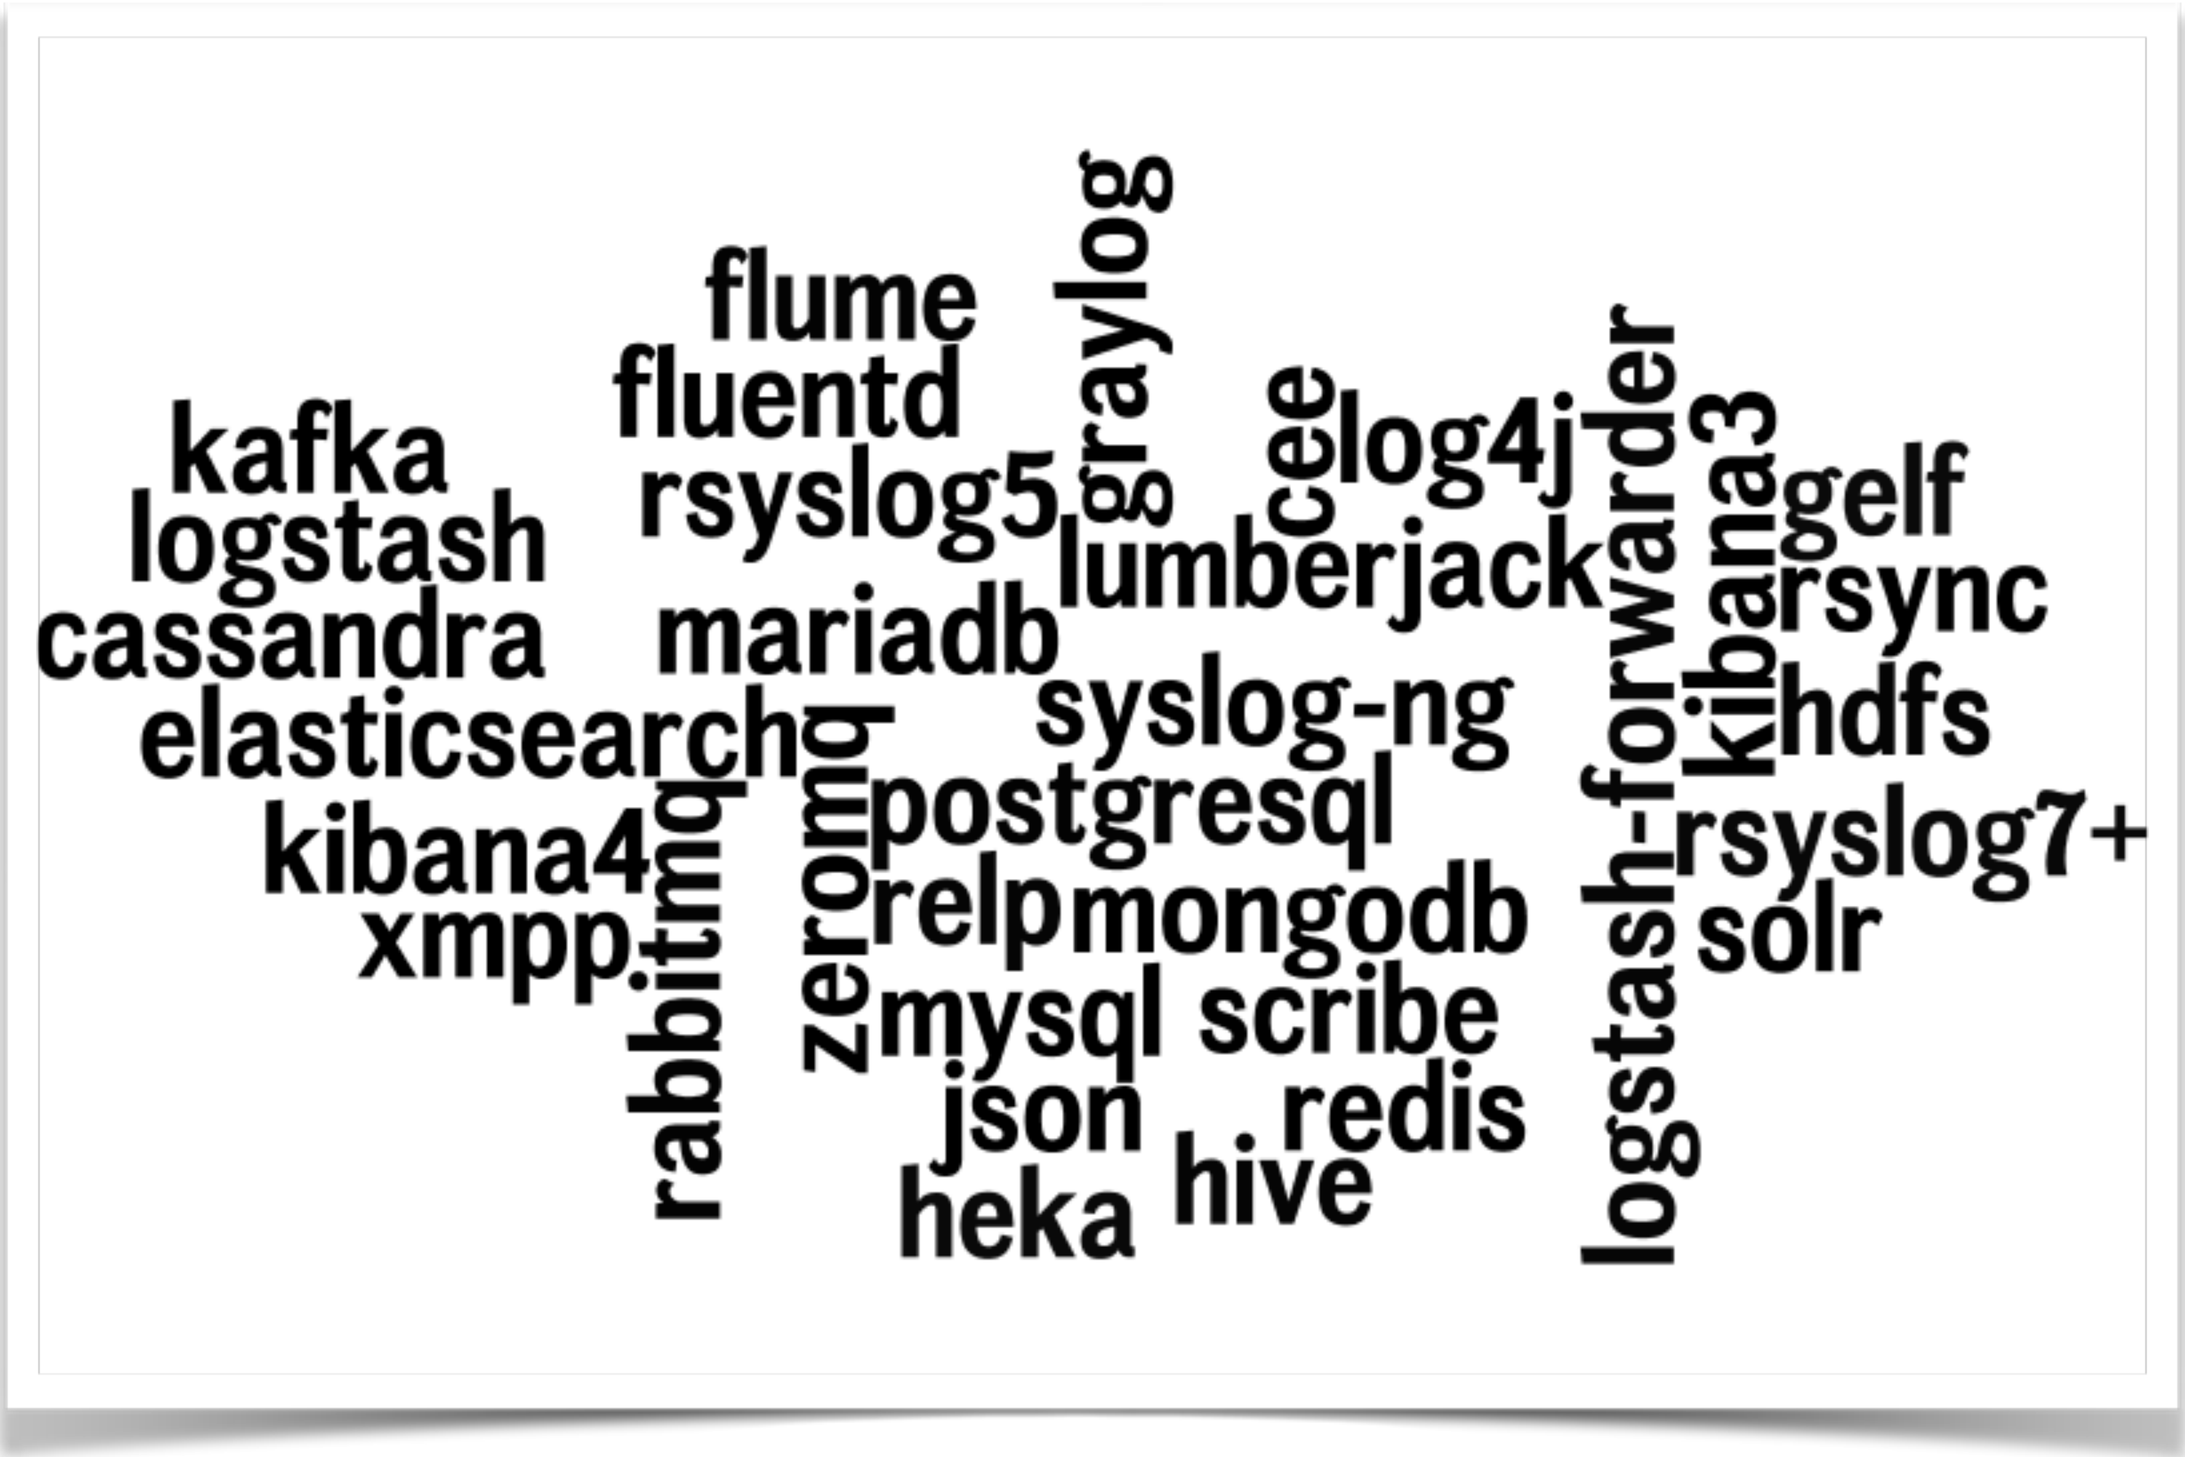
flume
fluentd
graylog
kafka
rsyslog5
cee
log4j
logstash-forwarder
kibana3
gelf
logstash
lumberjack
rsync
cassandra
mariadb
elasticsearch
syslog-ng
hdfs
rabbitmq
zeromq
postgresql
kibana4
rsyslog7+
relp
mongodb
xmpp
solr
mysql
scribe
json
redis
hive
heka

The target

- Timestamps: ns since epoch, RFC3339
- Structured information
- Common representation and semantics across the board
- Robust, scalable, stateless pipeline for events and metrics
- *One stop shop* for logs and metrics

The prototype

- Servers publish events to message brokers
- Workers consume events from queues
  - Perform arbitrary data transformation on *raw* data
  - Store events
- Logging backend supports:
  - Log search
  - Log analysis and visualisation
  - Monitoring dashboards

The toolset

- Heka
  - Collect, transport, enhance, output events
- RabbitMQ
  - Decouple producers from consumers
- Elasticsearch
  - Distributed event store and search
- Kibana
  - UI for search, analysis, dashboards
- Ansible and Git
  - Version control and configuration management

Heka

- Written in go
- Small footprint
- Performant
- Uses protobuf
- Sandboxed execution of custom LUA scripts
- TOML config files
- Sensible internal pipeline

Heka internal pipeline

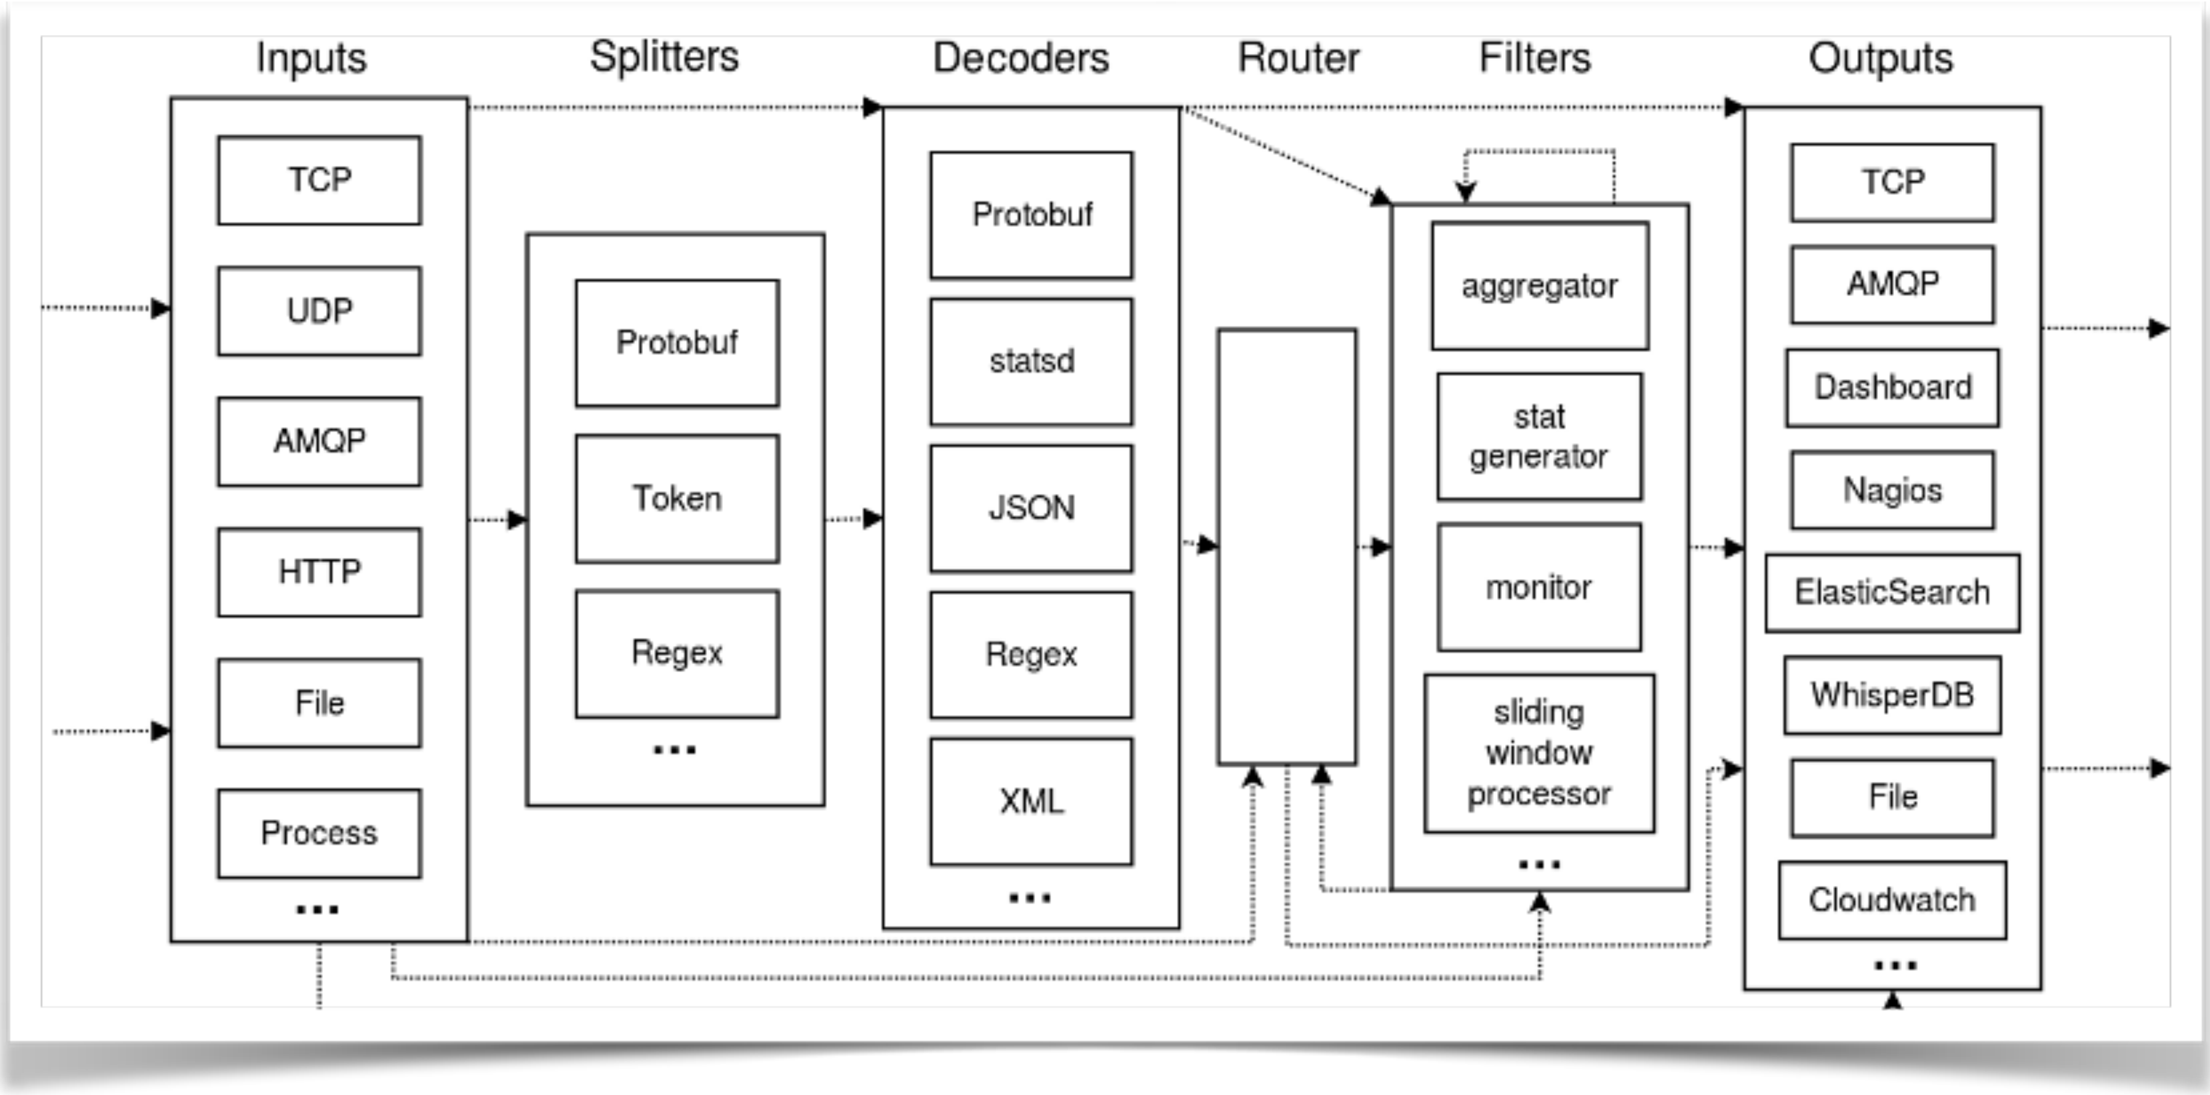

RabbitMQ

- Written in erlang
- Dedicated vhost/exchange/queue per *user*
- Acknowledged messages, persistent when required
- Standalone instances behind LB pool
- Nice flow control features
- ~250 concurrent connections
- ~550 channels
- 8 exchanges
- 12 queues

Elasticsearch

- Distributed free text search engine
- Apache Lucene
- Scalable, fast
- Aggregations (*facets*)
- Battle-tested at GitHub, ebay, The Guardian, bol.com…
- Powerful Query DSL
- Backup and restore (NFS, HDFS)
- Sensible defaults

Cluster description

- ~1B docs, 90+ indices, 550+ shards, ~1K events/sec
- 3x dedicated servers (Dell PowerEdge C5000 chassis)
    - 4x 1TB disks (7.2K RPM), 32GB of RAM
- No dedicated master node
- Dedicated standalone cluster for monitoring and index management
- PXE booted, RAM based root FS

Cluster monitoring

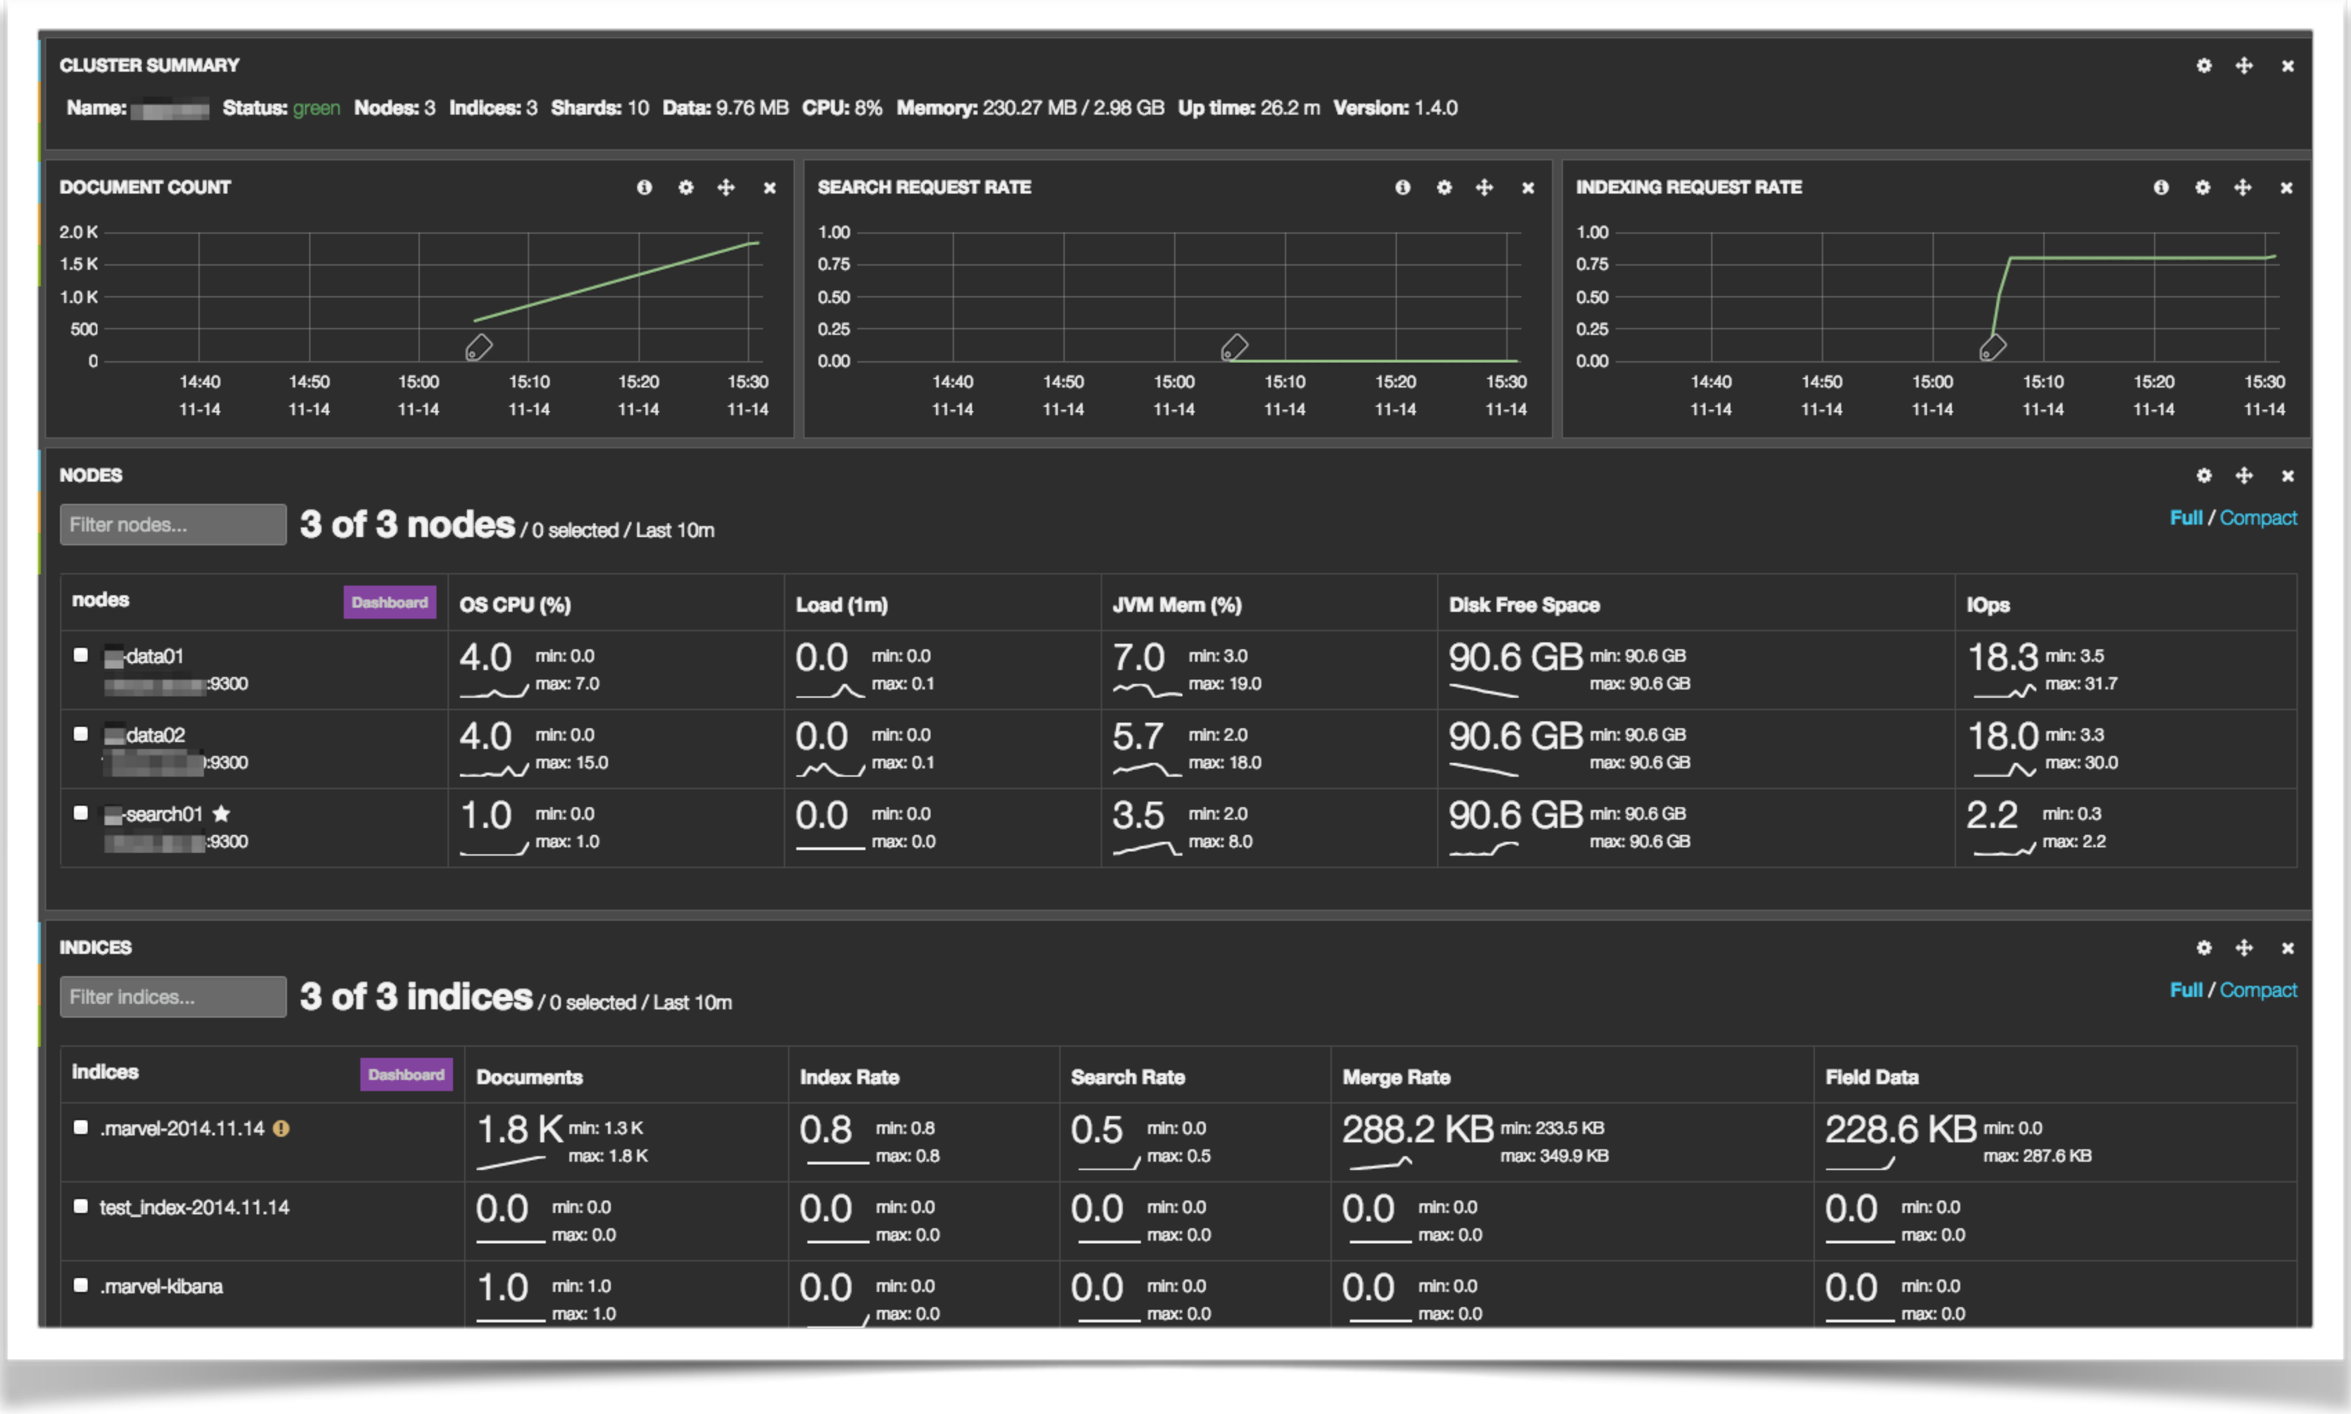

CLUSTER SUMMARY

Name: [redacted] **Status:** green **Nodes:** 3 **Indices:** 3 **Shards:** 10 **Data:** 9.76 MB **CPU:** 8% **Memory:** 230.27 MB / 2.98 GB **Up time:** 26.2 m **Version:** 1.4.0

DOCUMENT COUNT

SEARCH REQUEST RATE

INDEXING REQUEST RATE

NODES

Filter nodes...

3 of 3 nodes / 0 selected / Last 10m

Full / Compact

nodes Dashboard	OS CPU (%)	Load (1m)	JVM Mem (%)	Disk Free Space	IOps
□ [redacted]-data01 [redacted]:9300	4.0 min: 0.0 max: 7.0	0.0 min: 0.0 max: 0.1	7.0 min: 3.0 max: 19.0	90.6 GB min: 90.6 GB max: 90.6 GB	18.3 min: 3.5 max: 31.7
□ [redacted]data02 [redacted]:9300	4.0 min: 0.0 max: 15.0	0.0 min: 0.0 max: 0.1	5.7 min: 2.0 max: 18.0	90.6 GB min: 90.6 GB max: 90.6 GB	18.0 min: 3.3 max: 30.0
□ [redacted]-search01 ★ [redacted]:9300	1.0 min: 0.0 max: 1.0	0.0 min: 0.0 max: 0.0	3.5 min: 2.0 max: 8.0	90.6 GB min: 90.6 GB max: 90.6 GB	2.2 min: 0.3 max: 2.2

INDICES

Filter indices...

3 of 3 indices / 0 selected / Last 10m

Full / Compact

Indices Dashboard	Documents	Index Rate	Search Rate	Merge Rate	Field Data
□ .marvel-2014.11.14 !	1.8 K min: 1.3 K max: 1.8 K	0.8 min: 0.8 max: 0.8	0.5 min: 0.0 max: 0.5	288.2 KB min: 233.5 KB max: 349.9 KB	228.6 KB min: 0.0 max: 287.6 KB
□ test_index-2014.11.14	0.0 min: 0.0 max: 0.0	0.0 min: 0.0 max: 0.0	0.0 min: 0.0 max: 0.0	0.0 min: 0.0 max: 0.0	0.0 min: 0.0 max: 0.0
□ .marvel-kibana	1.0 min: 1.0 max: 1.0	0.0 min: 0.0 max: 0.0	0.0 min: 0.0 max: 0.0	0.0 min: 0.0 max: 0.0	0.0 min: 0.0 max: 0.0

Index management

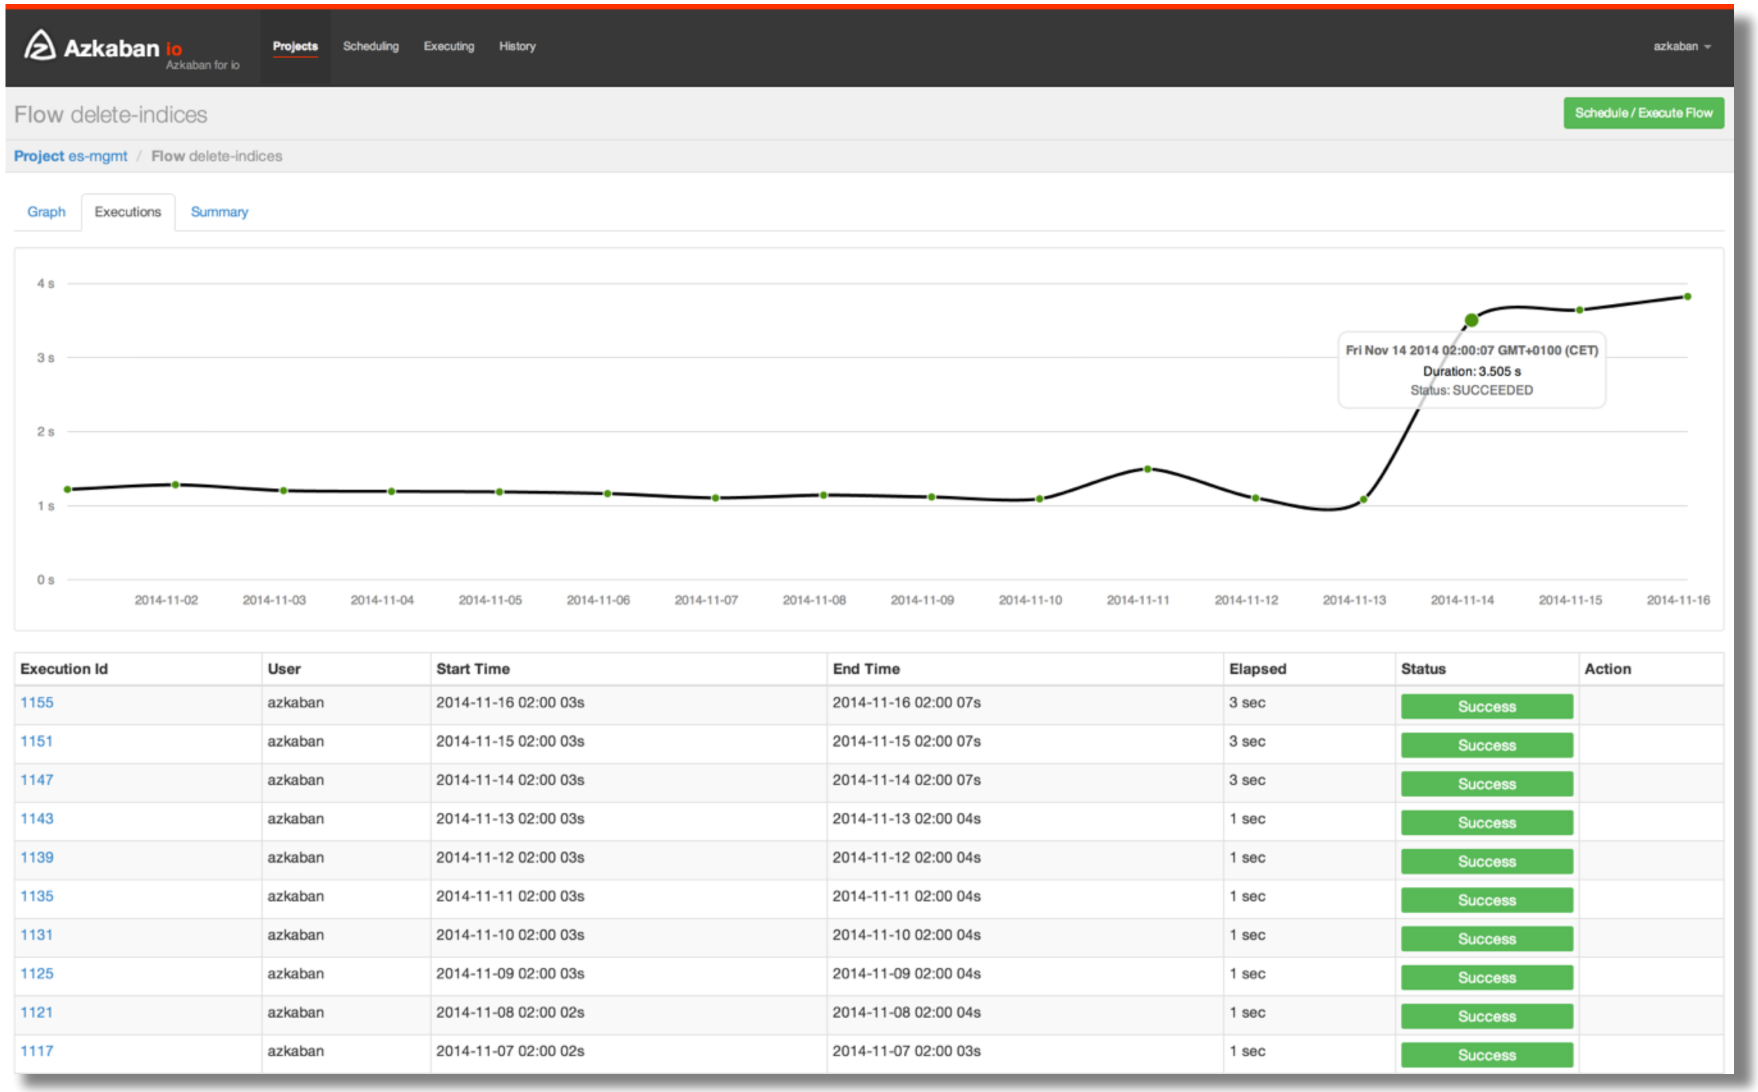

Flow delete-indices

Project es-mgmt / **Flow** delete-indices

Execution Id	User	Start Time	End Time	Elapsed	Status	Action
1155	azkaban	2014-11-16 02:00 03s	2014-11-16 02:00 07s	3 sec	Success	
1151	azkaban	2014-11-15 02:00 03s	2014-11-15 02:00 07s	3 sec	Success	
1147	azkaban	2014-11-14 02:00 03s	2014-11-14 02:00 07s	3 sec	Success	
1143	azkaban	2014-11-13 02:00 03s	2014-11-13 02:00 04s	1 sec	Success	
1139	azkaban	2014-11-12 02:00 03s	2014-11-12 02:00 04s	1 sec	Success	
1135	azkaban	2014-11-11 02:00 03s	2014-11-11 02:00 04s	1 sec	Success	
1131	azkaban	2014-11-10 02:00 03s	2014-11-10 02:00 04s	1 sec	Success	
1125	azkaban	2014-11-09 02:00 03s	2014-11-09 02:00 04s	1 sec	Success	
1121	azkaban	2014-11-08 02:00 02s	2014-11-08 02:00 04s	1 sec	Success	
1117	azkaban	2014-11-07 02:00 02s	2014-11-07 02:00 03s	1 sec	Success	

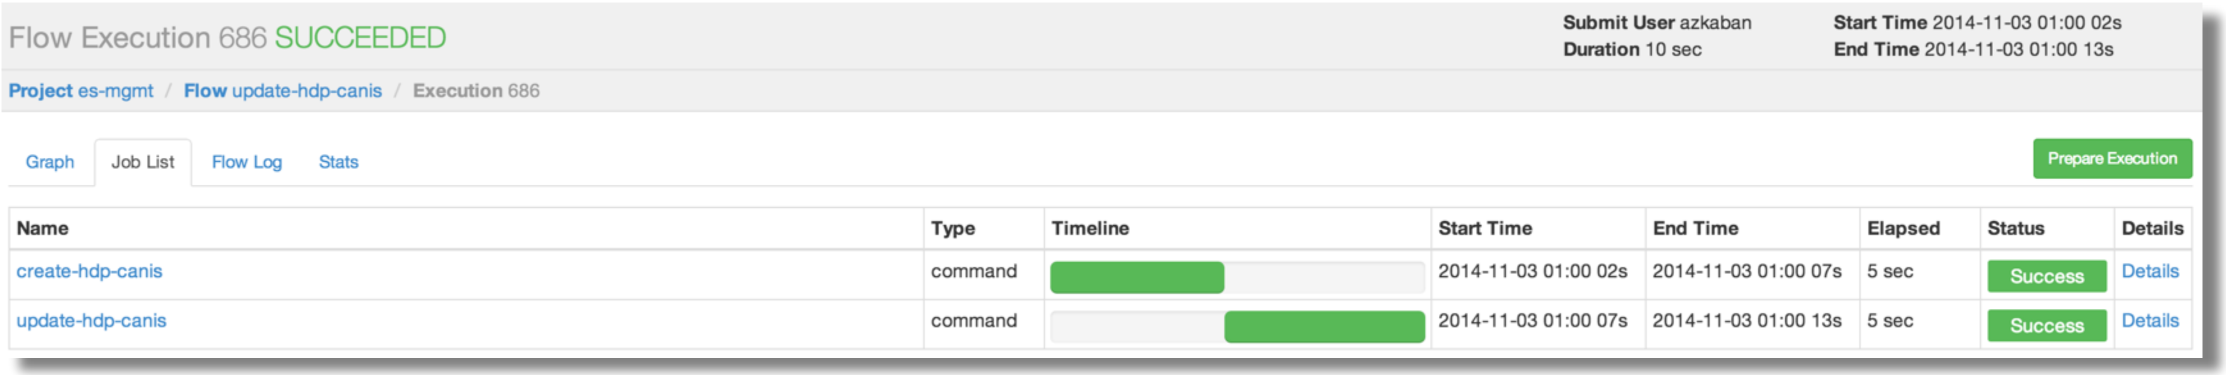

Flow Execution 686 SUCCEEDED

Submit User azkaban
Duration 10 sec

Start Time 2014-11-03 01:00 02s
End Time 2014-11-03 01:00 13s

Project es-mgmt / **Flow** update-hdp-canis / **Execution** 686

Name	Type	Timeline	Start Time	End Time	Elapsed	Status	Details
create-hdp-canis	command		2014-11-03 01:00 02s	2014-11-03 01:00 07s	5 sec	Success	Details
update-hdp-canis	command		2014-11-03 01:00 07s	2014-11-03 01:00 13s	5 sec	Success	Details

Kibana

- Webapp for analytics and visualisation
- Intuitive for most
- Easy sharing capabilities
- Pretty graphs!
    - Which *may* melt your cluster :-)

Kibana

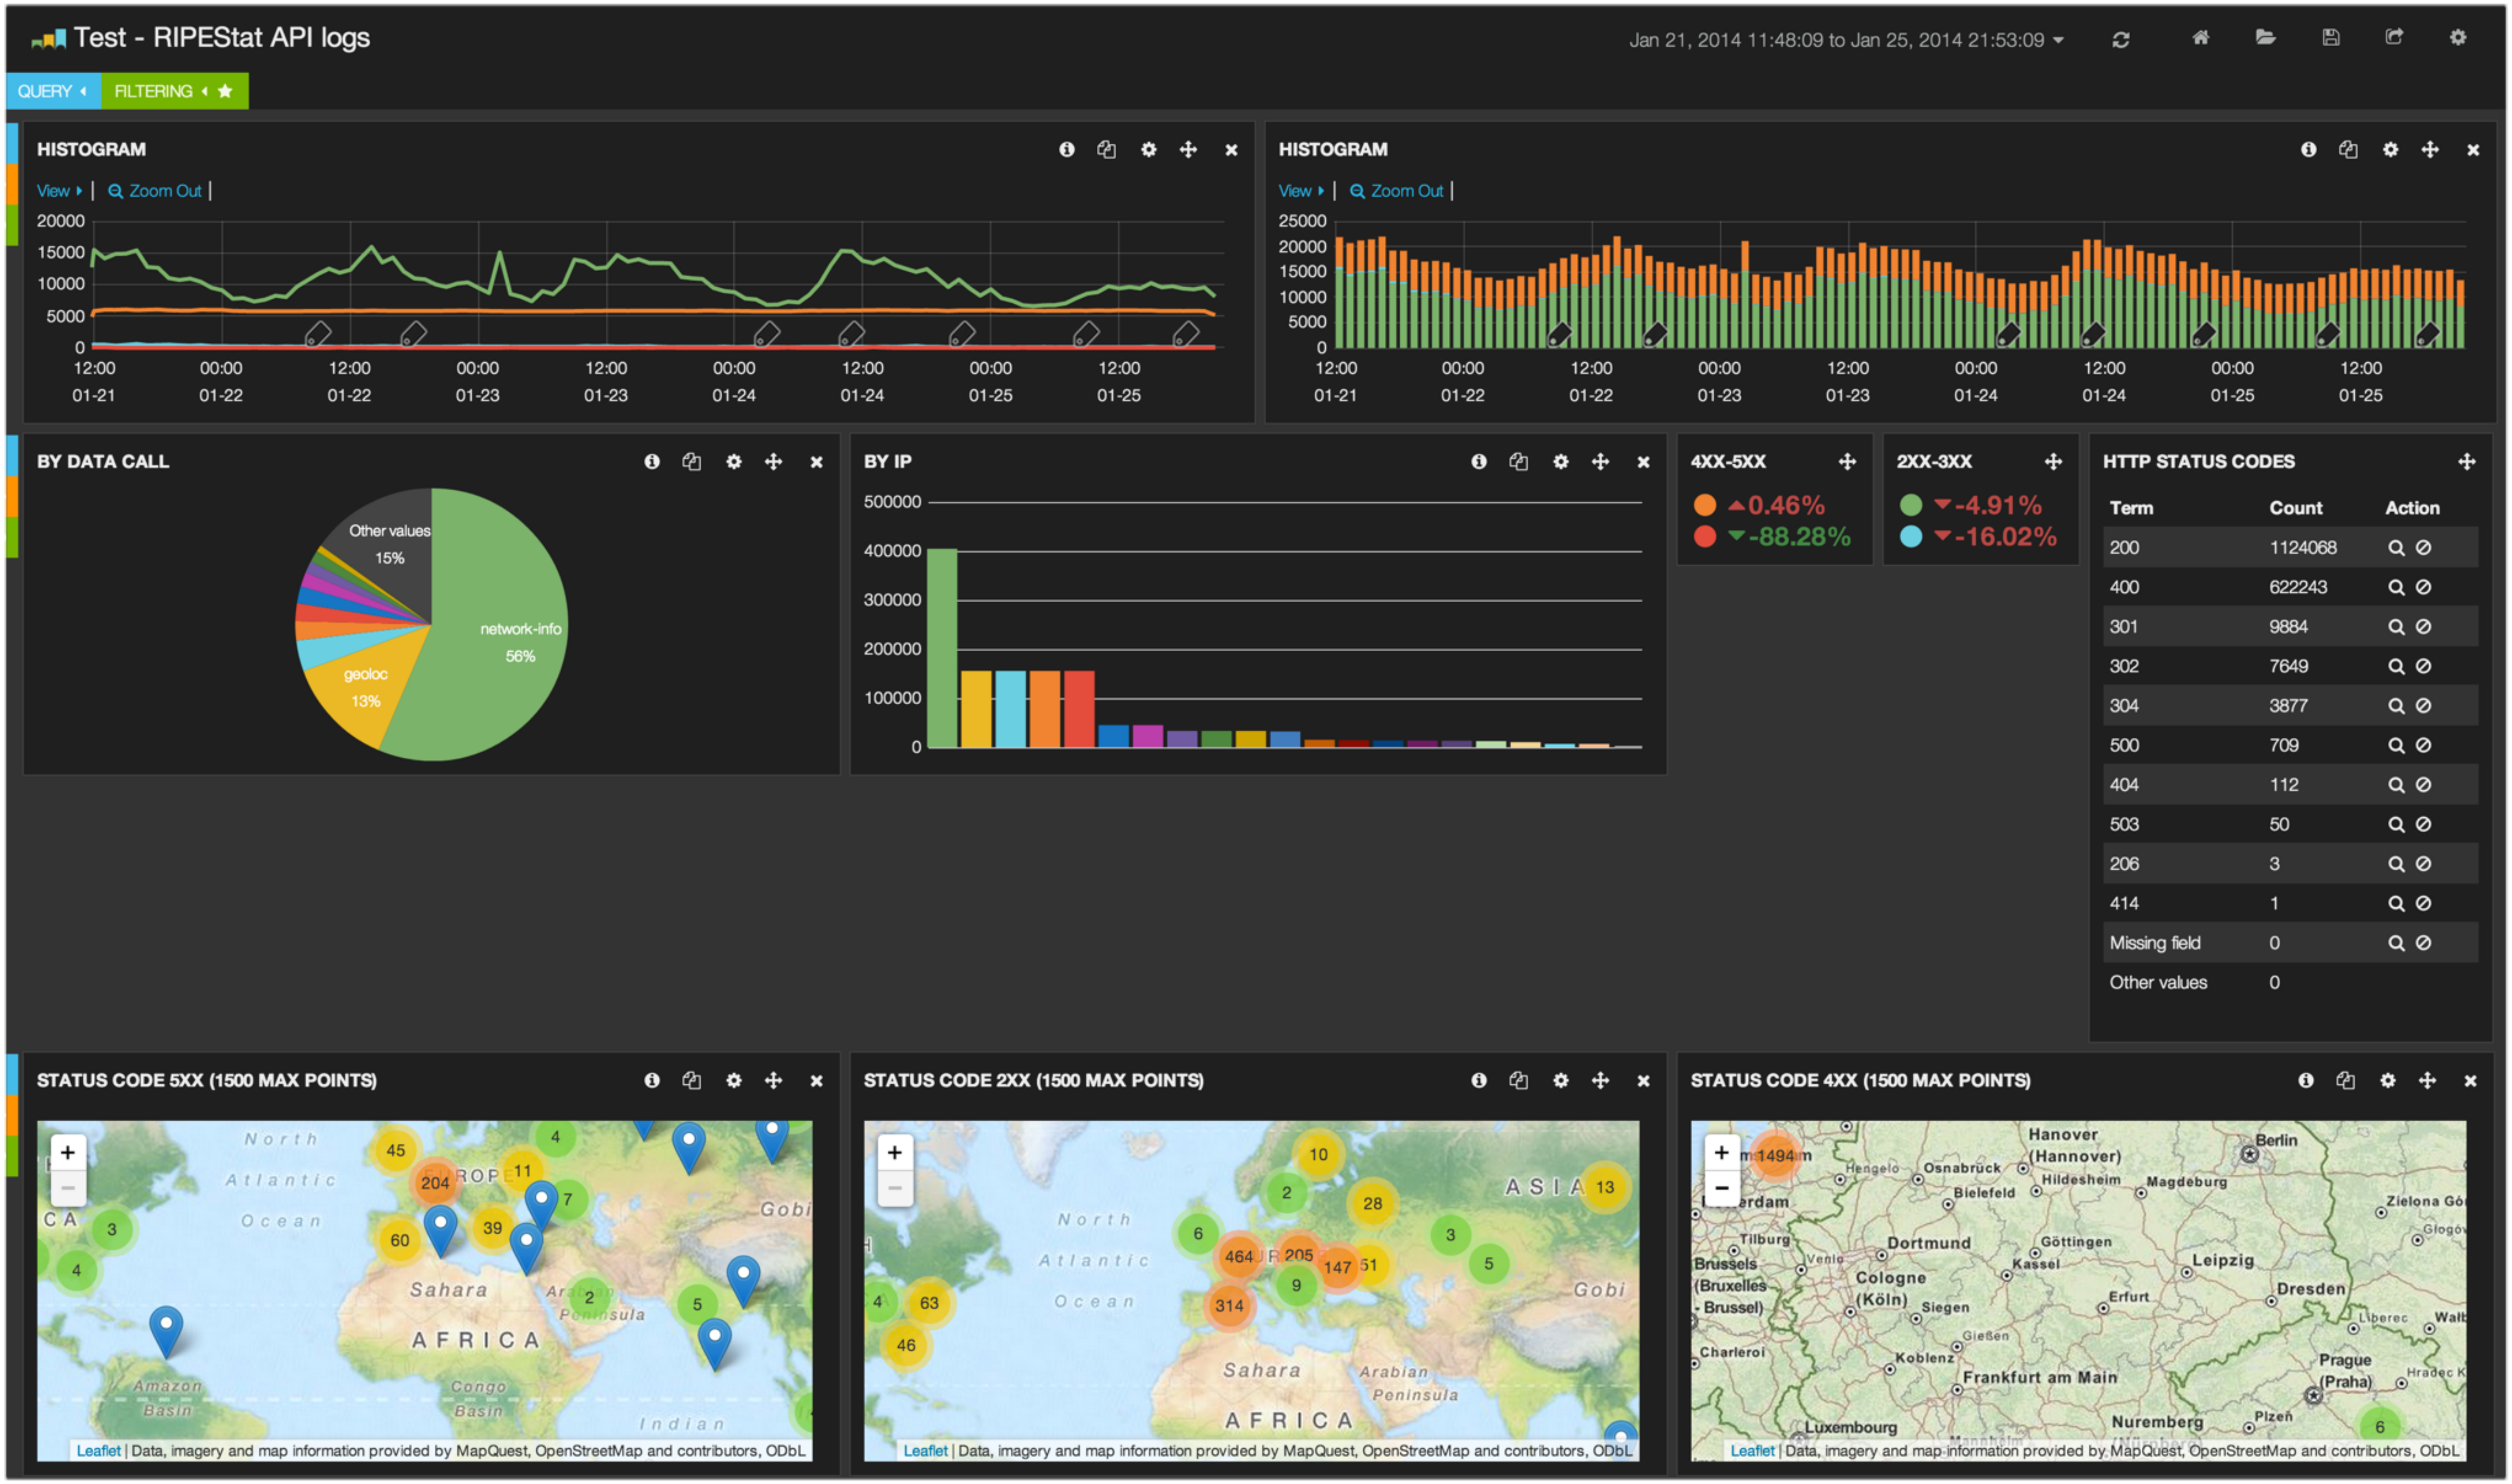

Test - RIPEStat API logs

Jan 21, 2014 11:48:09 to Jan 25, 2014 21:53:09

HTTP STATUS CODES

Term	Count	Action
200	1124068	
400	622243	
301	9884	
302	7649	
304	3877	
500	709	
404	112	
503	50	
206	3	
414	1	
Missing field	0	
Other values	0	

Kibana

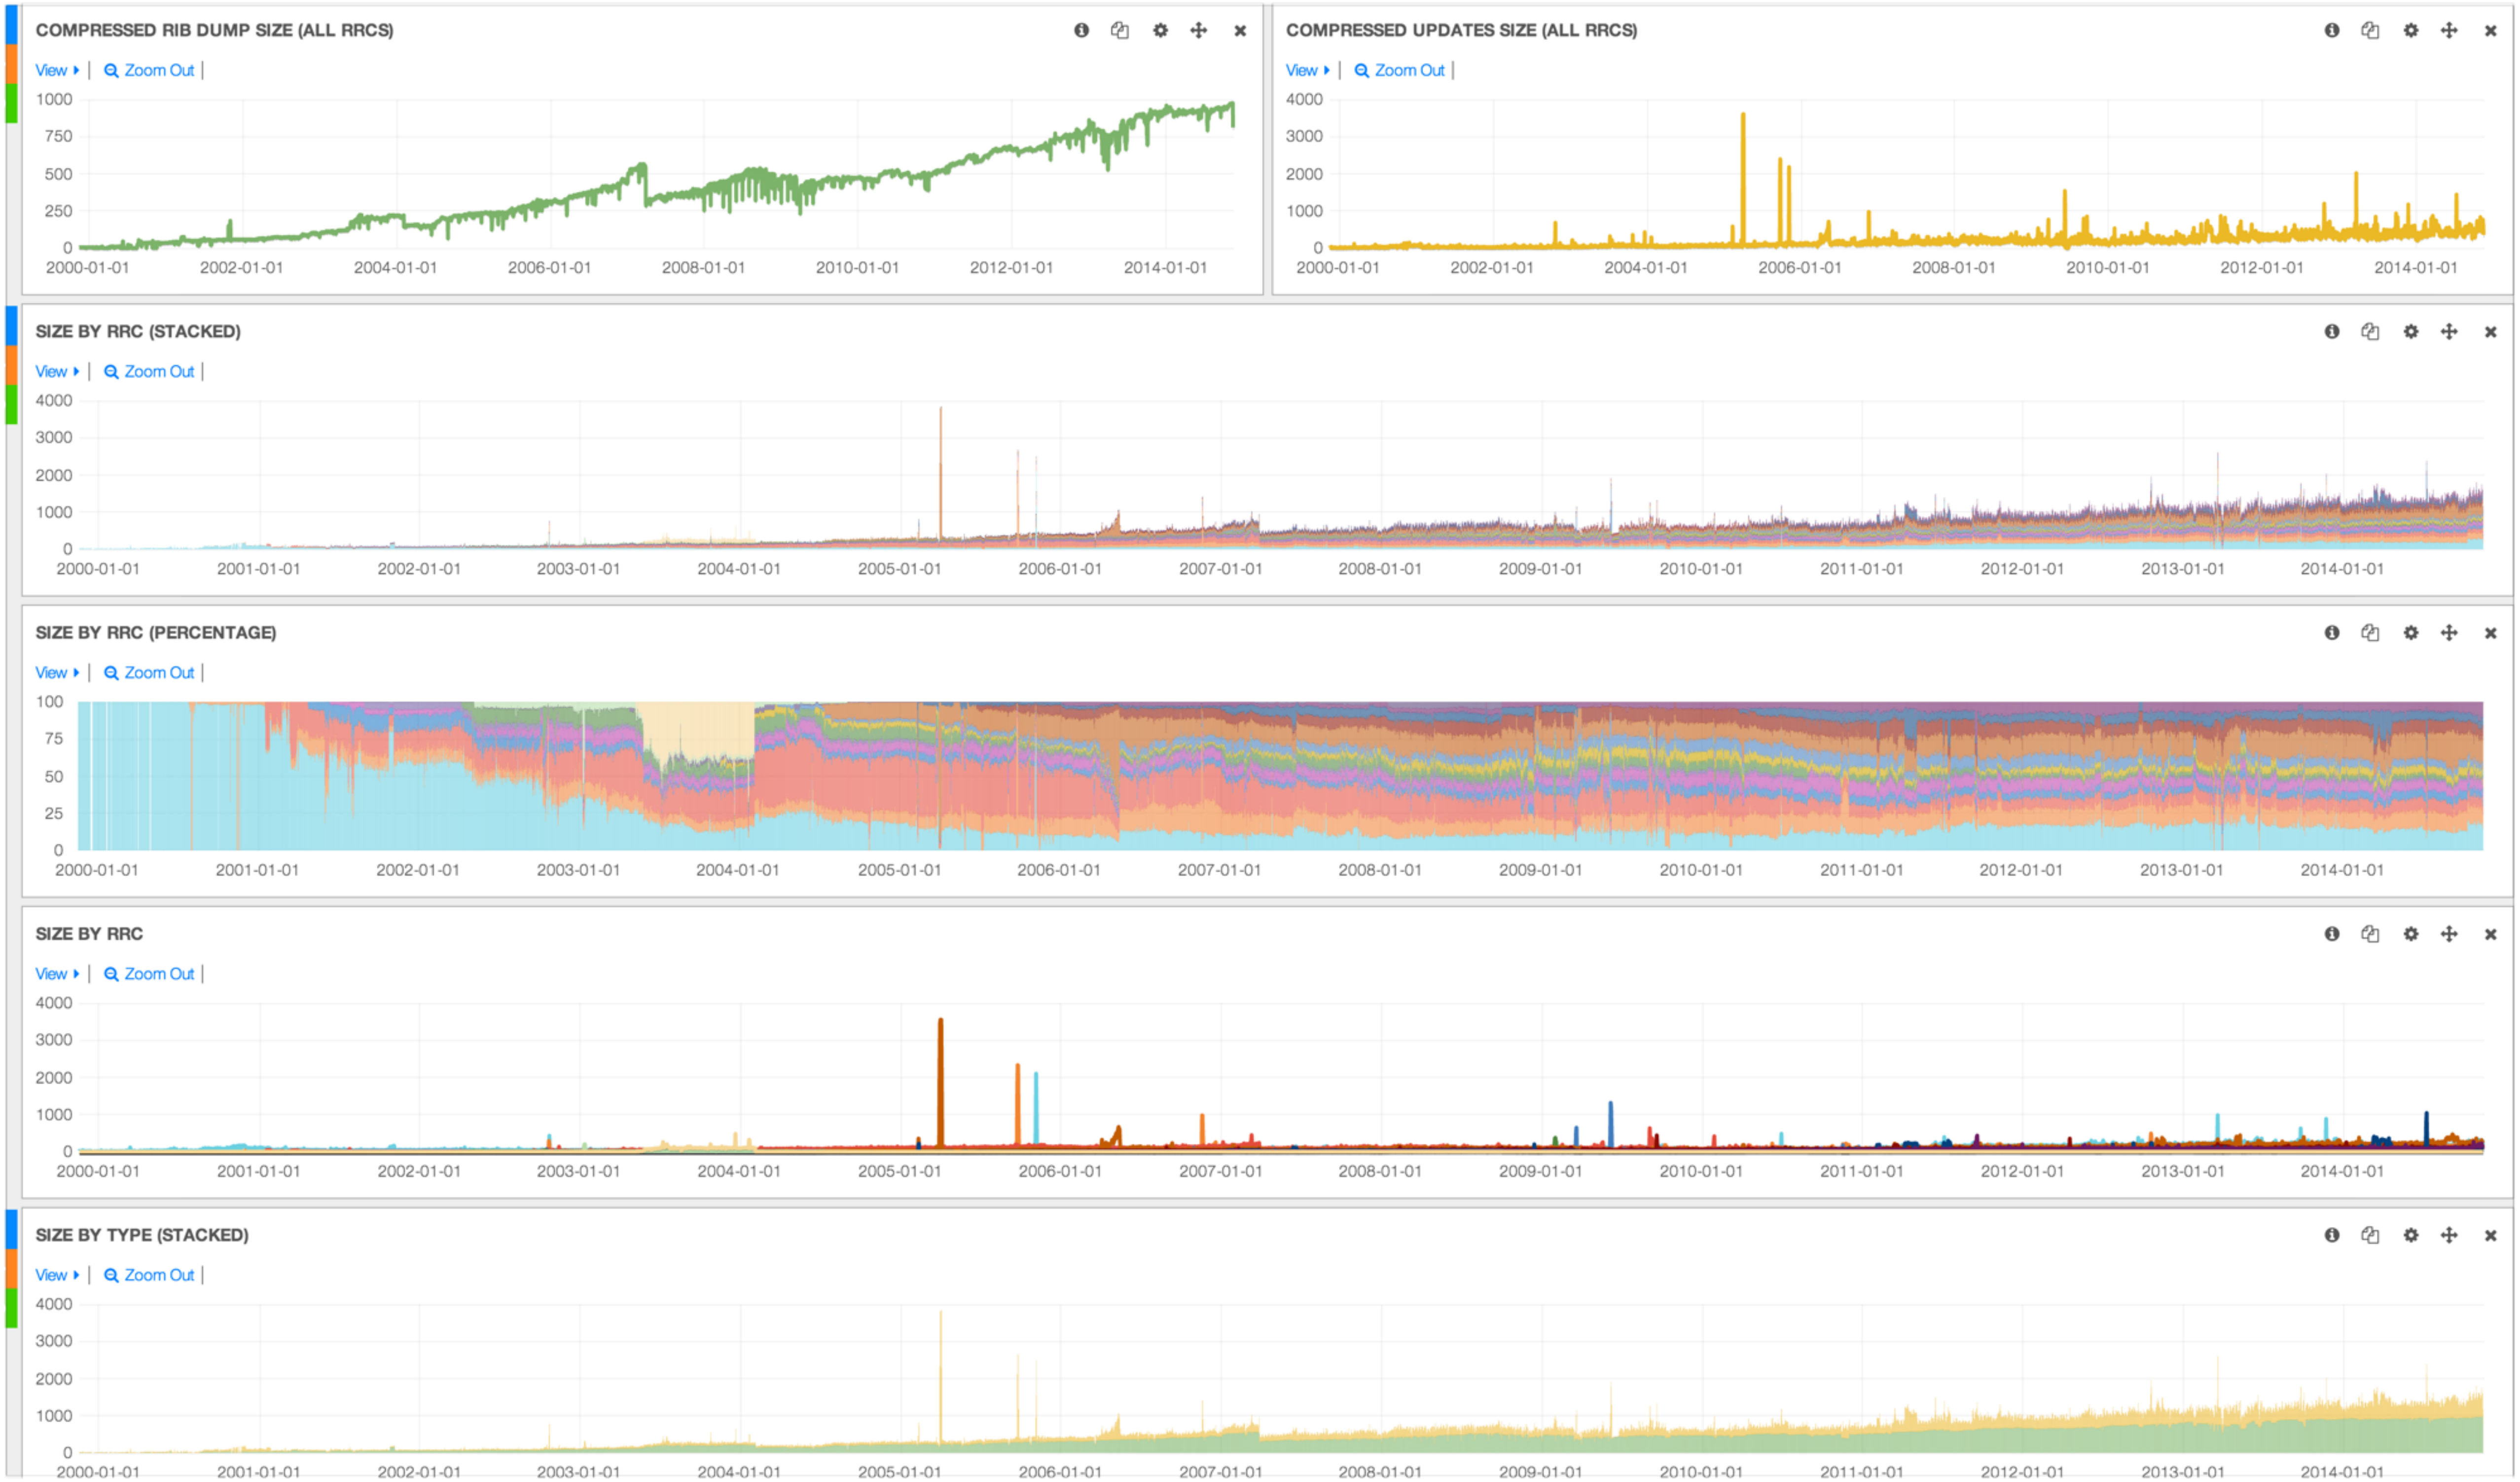
COMPRESSED RIB DUMP SIZE (ALL RRCS)
View ▸ | Zoom Out |
1000
750
500
250
0
2000-01-01
2002-01-01
2004-01-01
2006-01-01
2008-01-01
2010-01-01
2012-01-01
2014-01-01
COMPRESSED UPDATES SIZE (ALL RRCS)
View ▸ | Zoom Out |
4000
3000
2000
1000
0
2000-01-01
2002-01-01
2004-01-01
2006-01-01
2008-01-01
2010-01-01
2012-01-01
2014-01-01
SIZE BY RRC (STACKED)
View ▸ | Zoom Out |
4000
3000
2000
1000
0
2000-01-01
2001-01-01
2002-01-01
2003-01-01
2004-01-01
2005-01-01
2006-01-01
2007-01-01
2008-01-01
2009-01-01
2010-01-01
2011-01-01
2012-01-01
2013-01-01
2014-01-01
SIZE BY RRC (PERCENTAGE)
View ▸ | Zoom Out |
100
75
50
25
0
2000-01-01
2001-01-01
2002-01-01
2003-01-01
2004-01-01
2005-01-01
2006-01-01
2007-01-01
2008-01-01
2009-01-01
2010-01-01
2011-01-01
2012-01-01
2013-01-01
2014-01-01
SIZE BY RRC
View ▸ | Zoom Out |
4000
3000
2000
1000
0
2000-01-01
2001-01-01
2002-01-01
2003-01-01
2004-01-01
2005-01-01
2006-01-01
2007-01-01
2008-01-01
2009-01-01
2010-01-01
2011-01-01
2012-01-01
2013-01-01
2014-01-01
SIZE BY TYPE (STACKED)
View ▸ | Zoom Out |
4000
3000
2000
1000
0
2000-01-01
2001-01-01
2002-01-01
2003-01-01
2004-01-01
2005-01-01
2006-01-01
2007-01-01
2008-01-01
2009-01-01
2010-01-01
2011-01-01
2012-01-01
2013-01-01
2014-01-01

Kibana

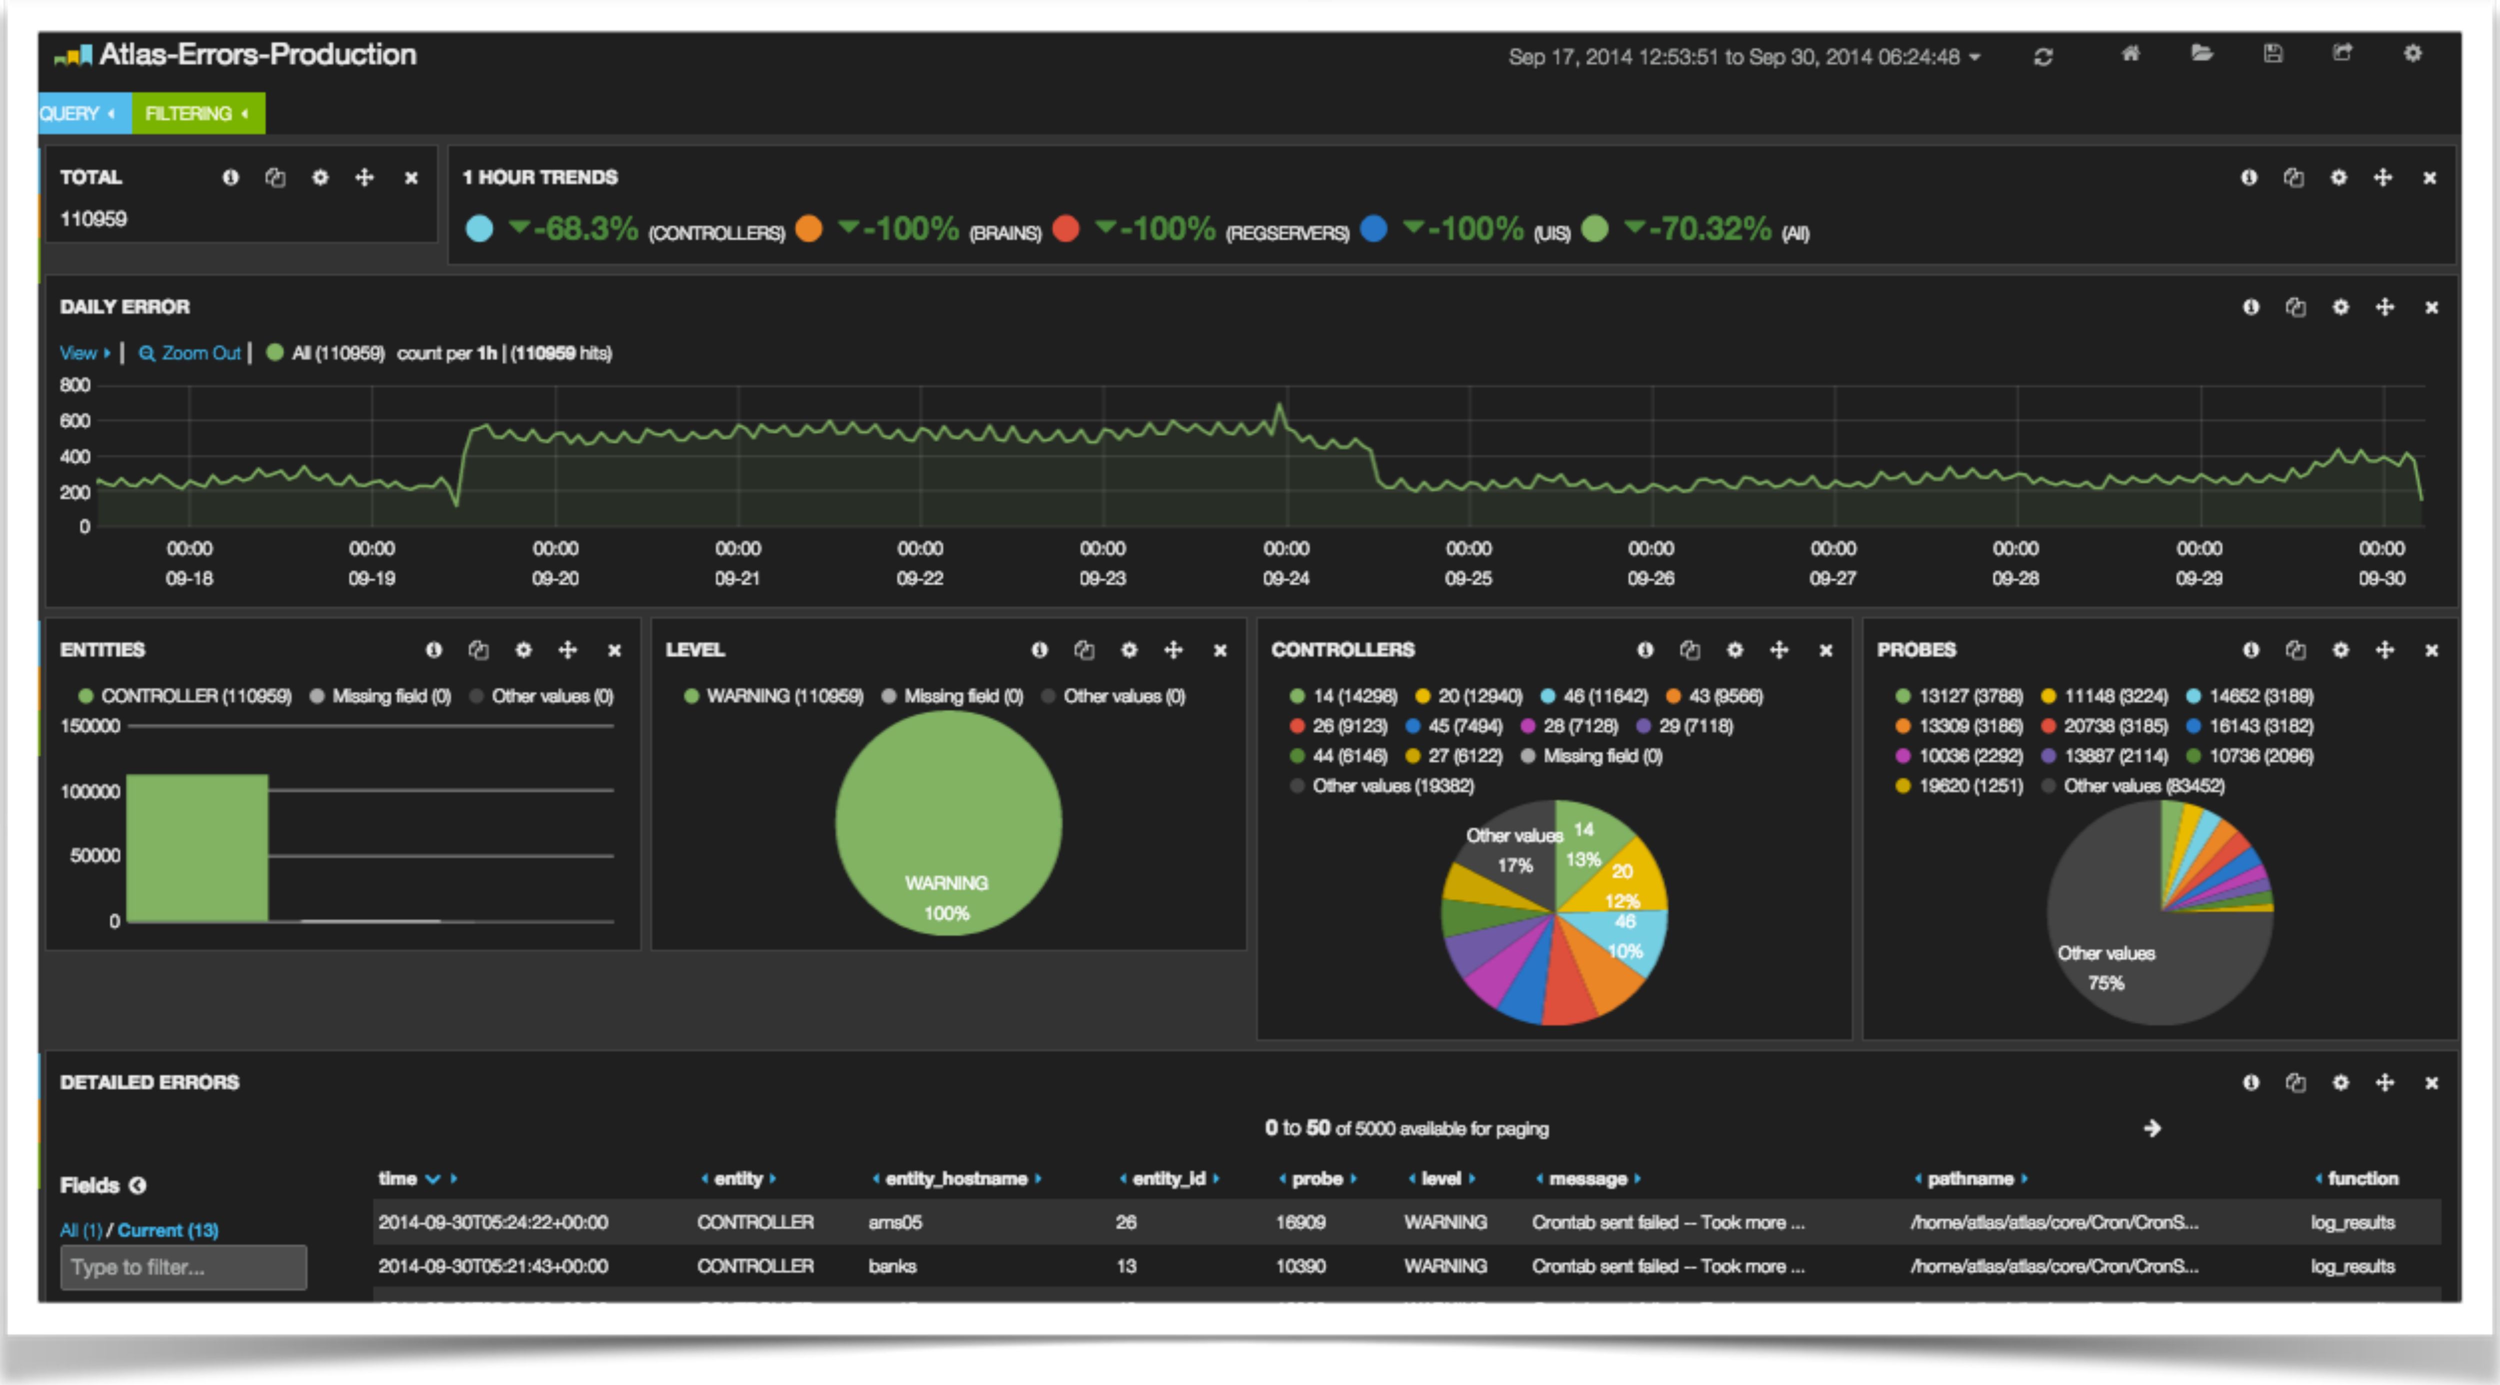

Atlas-Errors-Production

Sep 17, 2014 12:53:51 to Sep 30, 2014 06:24:48

QUERY FILTERING

TOTAL

110959

1 HOUR TRENDS

● -68.3% (CONTROLLERS) ● -100% (BRAINS) ● -100% (REGSERVERS) ● -100% (UIS) ● -70.32% (All)

DAILY ERROR

View | Zoom Out | ● All (110959) count per 1h | (110959 hits)

ENTITIES

● CONTROLLER (110959) ● Missing field (0) ● Other values (0)

LEVEL

● WARNING (110959) ● Missing field (0) ● Other values (0)

CONTROLLERS

● 14 (14298) ● 20 (12940) ● 46 (11642) ● 43 (9566)
● 26 (9123) ● 45 (7494) ● 28 (7128) ● 29 (7118)
● 44 (6146) ● 27 (6122) ● Missing field (0)
● Other values (19382)

PROBES

● 13127 (3788) ● 11148 (3224) ● 14652 (3189)
● 13309 (3186) ● 20738 (3185) ● 16143 (3182)
● 10036 (2292) ● 13887 (2114) ● 10736 (2096)
● 19620 (1251) ● Other values (83452)

DETAILED ERRORS

0 to **50** of 5000 available for paging

Fields

All (1) / **Current (13)**

Type to filter...

time	entity	entity_hostname	entity_id	probe	level	message	pathname	function
2014-09-30T05:24:22+00:00	CONTROLLER	ams05	26	16909	WARNING	Crontab sent failed -- Took more ...	/home/atlas/atlas/core/Cron/CronS...	log_results
2014-09-30T05:21:43+00:00	CONTROLLER	banks	13	10390	WARNING	Crontab sent failed -- Took more ...	/home/atlas/atlas/core/Cron/CronS...	log_results

Recap

- Production
- Collection
- Transport
- Queueing, buffering
- Massaging
- Storage
- Search and analysis

Future plans

- Kibana4
  - Operational simplicity
  - Superior capabilities and UX
  - Migrate all dashboards
  - Encourage data exploration
- Kafka
  - “*Stateless*” broker
  - Superior performance (backlog ingestion)
  - Unified pipeline
- Data quality checks

Thanks!

- Mail, XMPP: iortiz@ripe.net
- Twitter: @ioc32

Questions?