# Chapter 1: Introduction to Data Driven Architecture

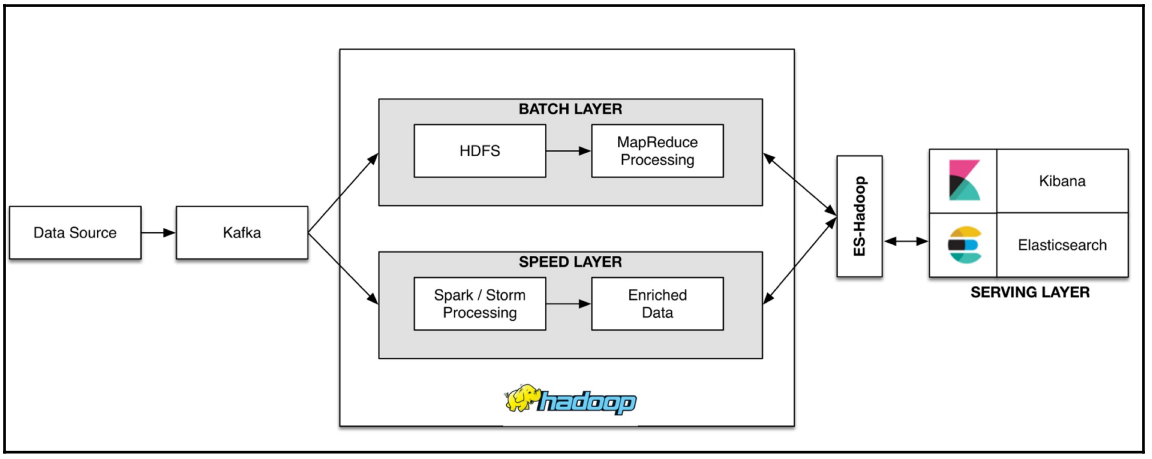

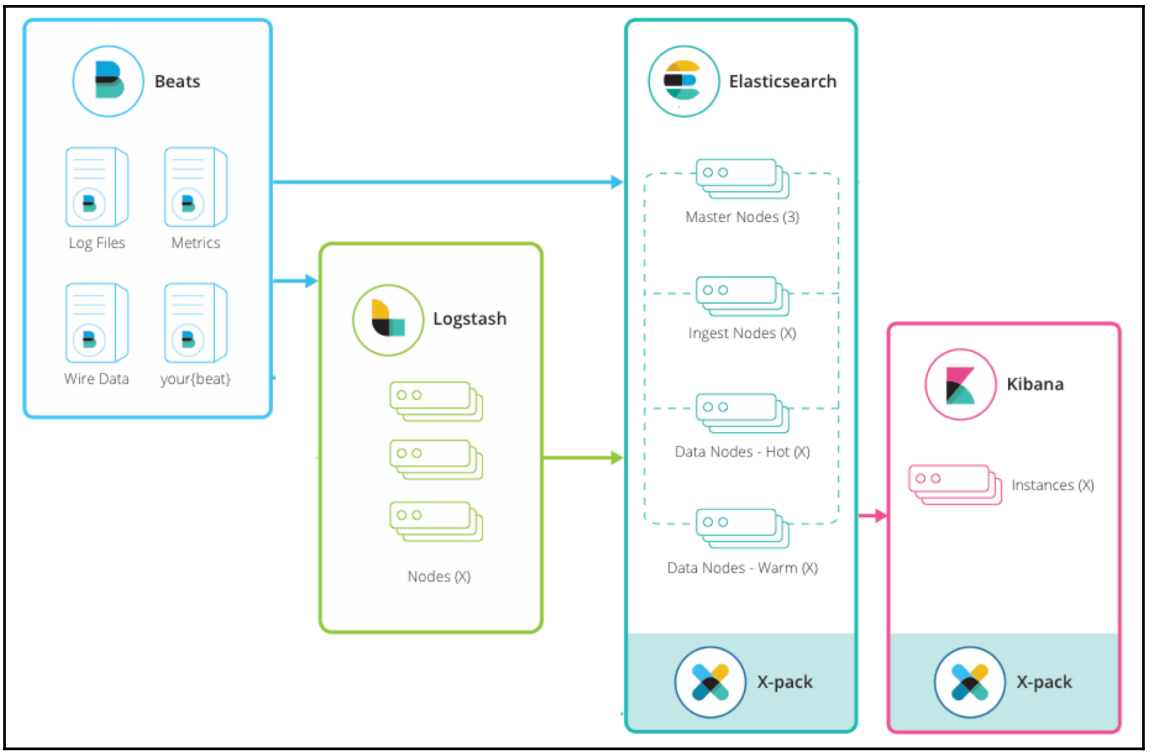
Beats
Log Files
Metrics
Wire Data
your{beat}
Logstash
Nodes (X)
Elasticsearch
Master Nodes (3)
Ingest Nodes (X)
Data Nodes - Hot (X)
Data Nodes - Warm (X)
X-pack
Kibana
Instances (X)
X-pack

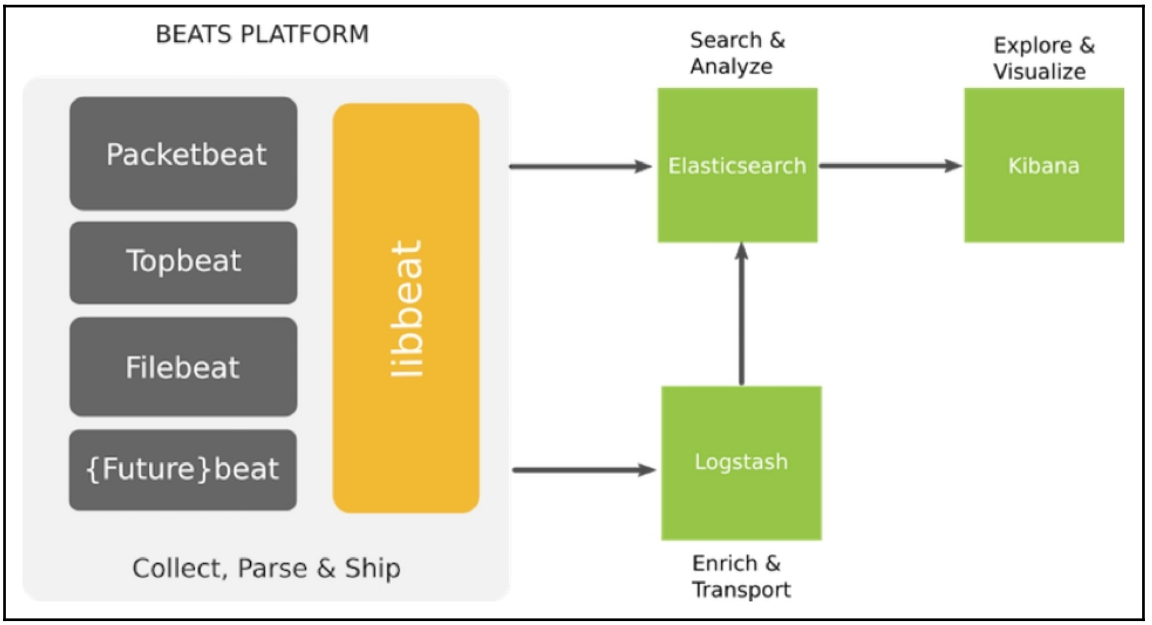
BEATS PLATFORM
Packetbeat
Topbeat
Filebeat
{Future}beat
libbeat
Collect, Parse & Ship
Search &
Analyze
Elasticsearch
Explore &
Visualize
Kibana
Logstash
Enrich &
Transport

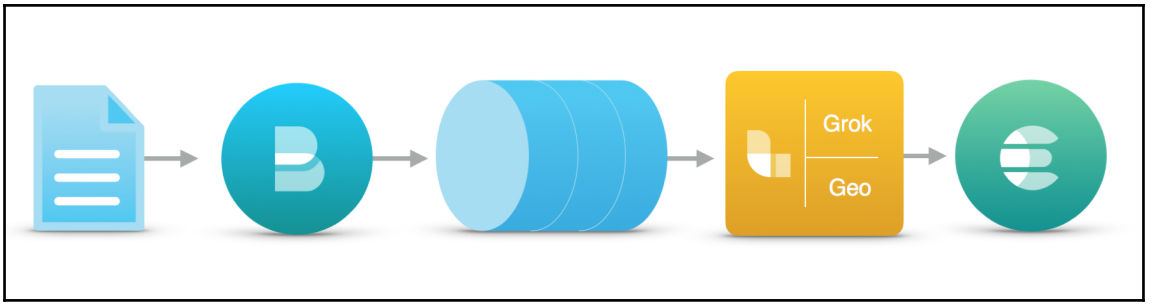
Grok
Geo

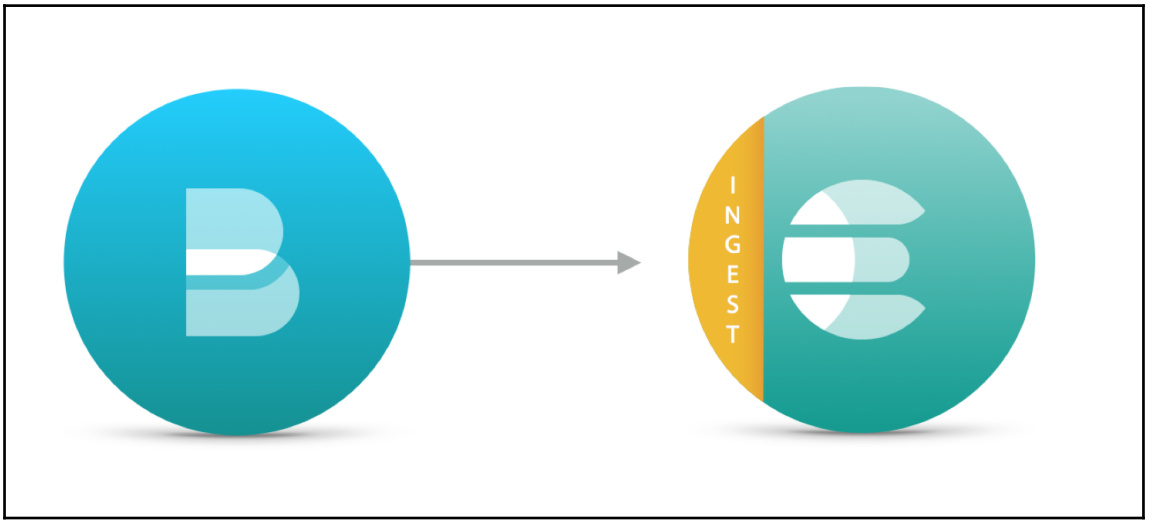
INGEST

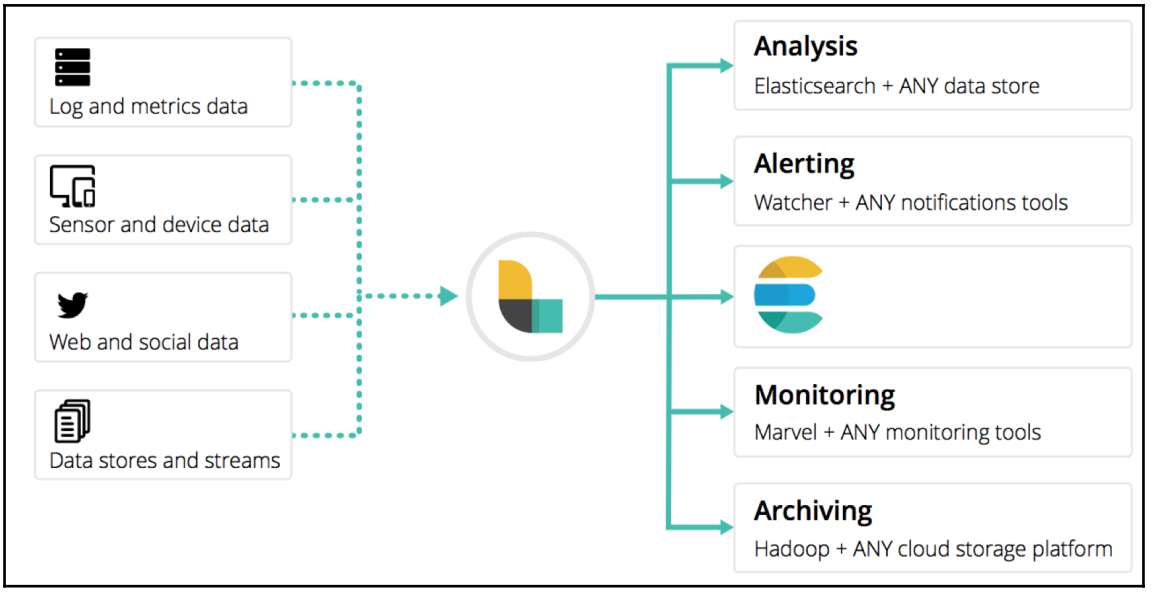
Log and metrics data
Sensor and device data
Web and social data
Data stores and streams
Analysis
Elasticsearch + ANY data store
Alerting
Watcher + ANY notifications tools
Monitoring
Marvel + ANY monitoring tools
Archiving
Hadoop + ANY cloud storage platform

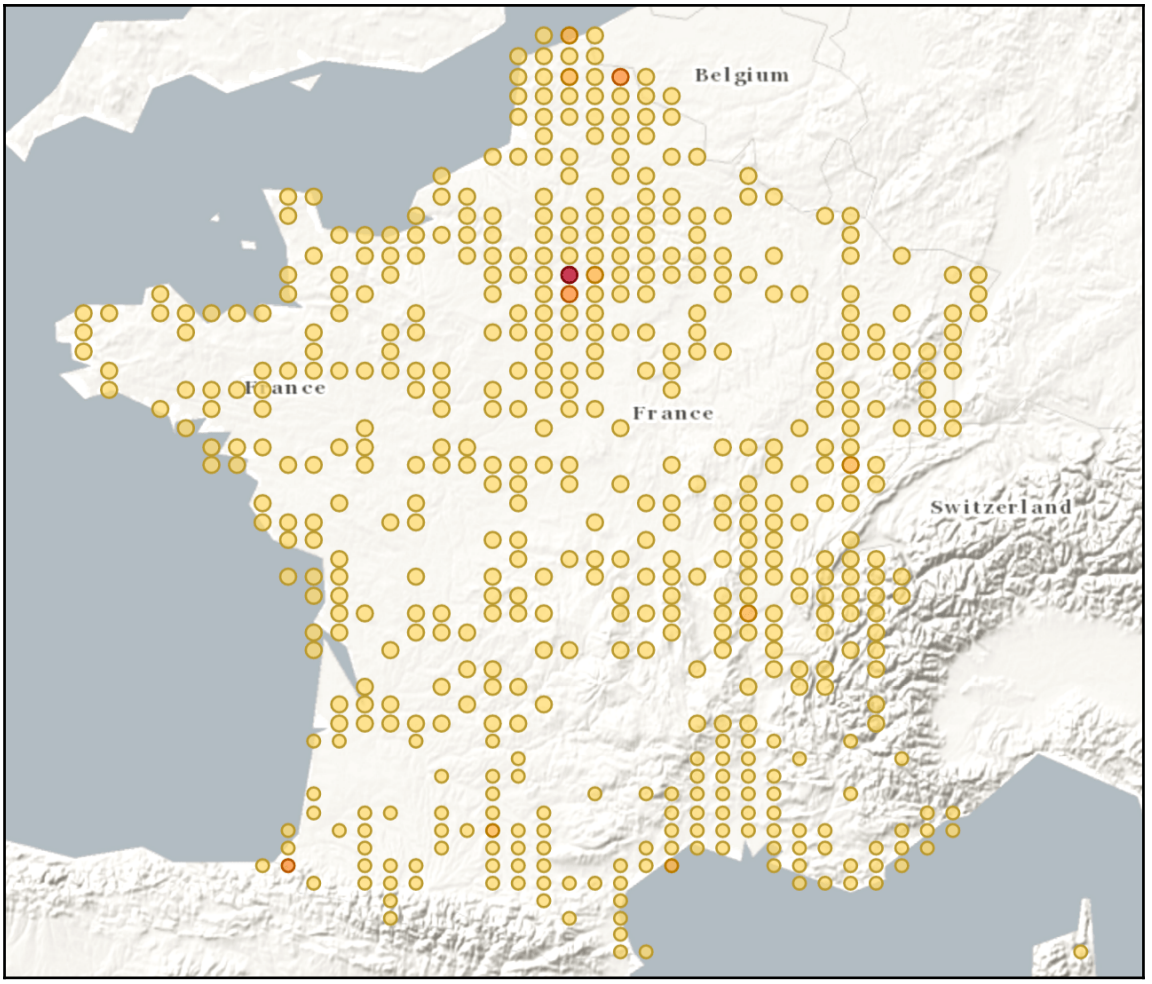
Belgium
France
France
Switzerland

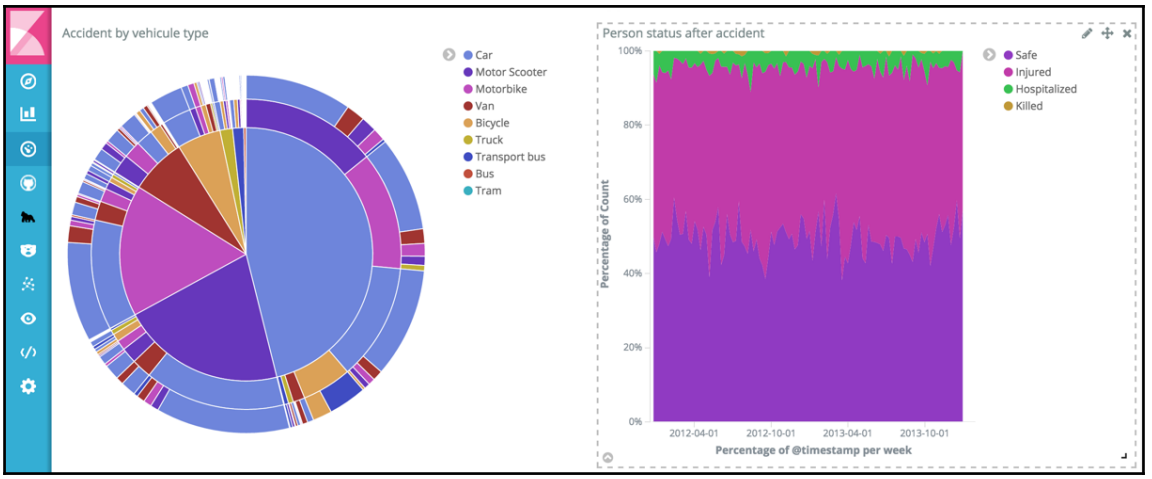
Accident by vehicule type
Car
Motor Scooter
Motorbike
Van
Bicycle
Truck
Transport bus
Bus
Tram
Person status after accident
Safe
Injured
Hospitalized
Killed
100%
80%
60%
40%
20%
0%
Percentage of Count
2012-04-01
2012-10-01
2013-04-01
2013-10-01
Percentage of @timestamp per week

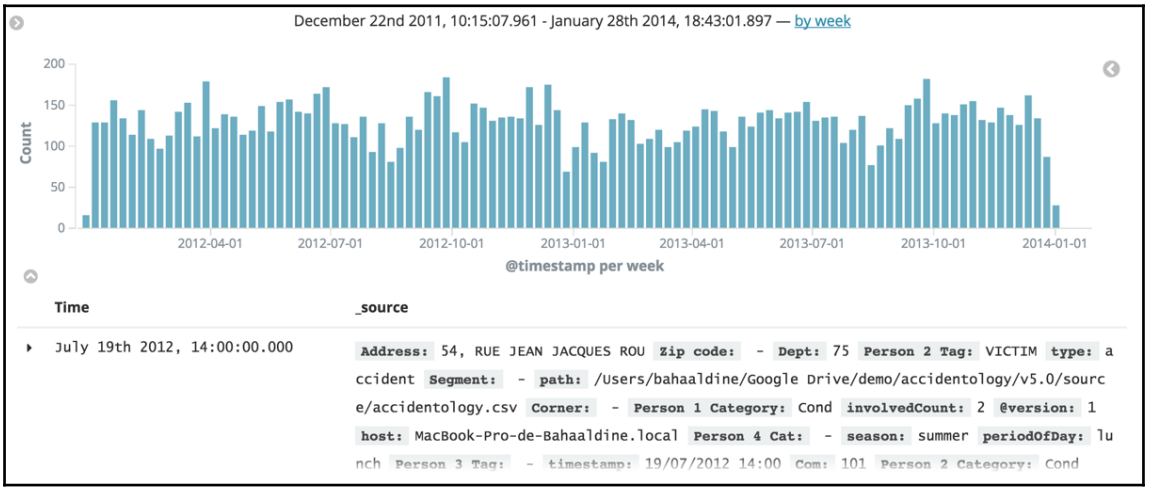
December 22nd 2011, 10:15:07.961 - January 28th 2014, 18:43:01.897 — by week
200
150
100
50
0
Count
2012-04-01
2012-07-01
2012-10-01
2013-01-01
2013-04-01
2013-07-01
2013-10-01
2014-01-01
@timestamp per week
Time
_source
July 19th 2012, 14:00:00.000
Address: 54, RUE JEAN JACQUES ROU Zip code: - Dept: 75 Person 2 Tag: VICTIM type: accident Segment: - path: /Users/bahaaldine/Google Drive/demo/accidentology/v5.0/source/accidentology.csv Corner: - Person 1 Category: Cond involvedCount: 2 @version: 1 host: MacBook-Pro-de-Bahaaldine.local Person 4 Cat: - season: summer periodOfDay: lunch Person 3 Tag: - timestamp: 19/07/2012 14:00 Com: 101 Person 2 Category: Cond

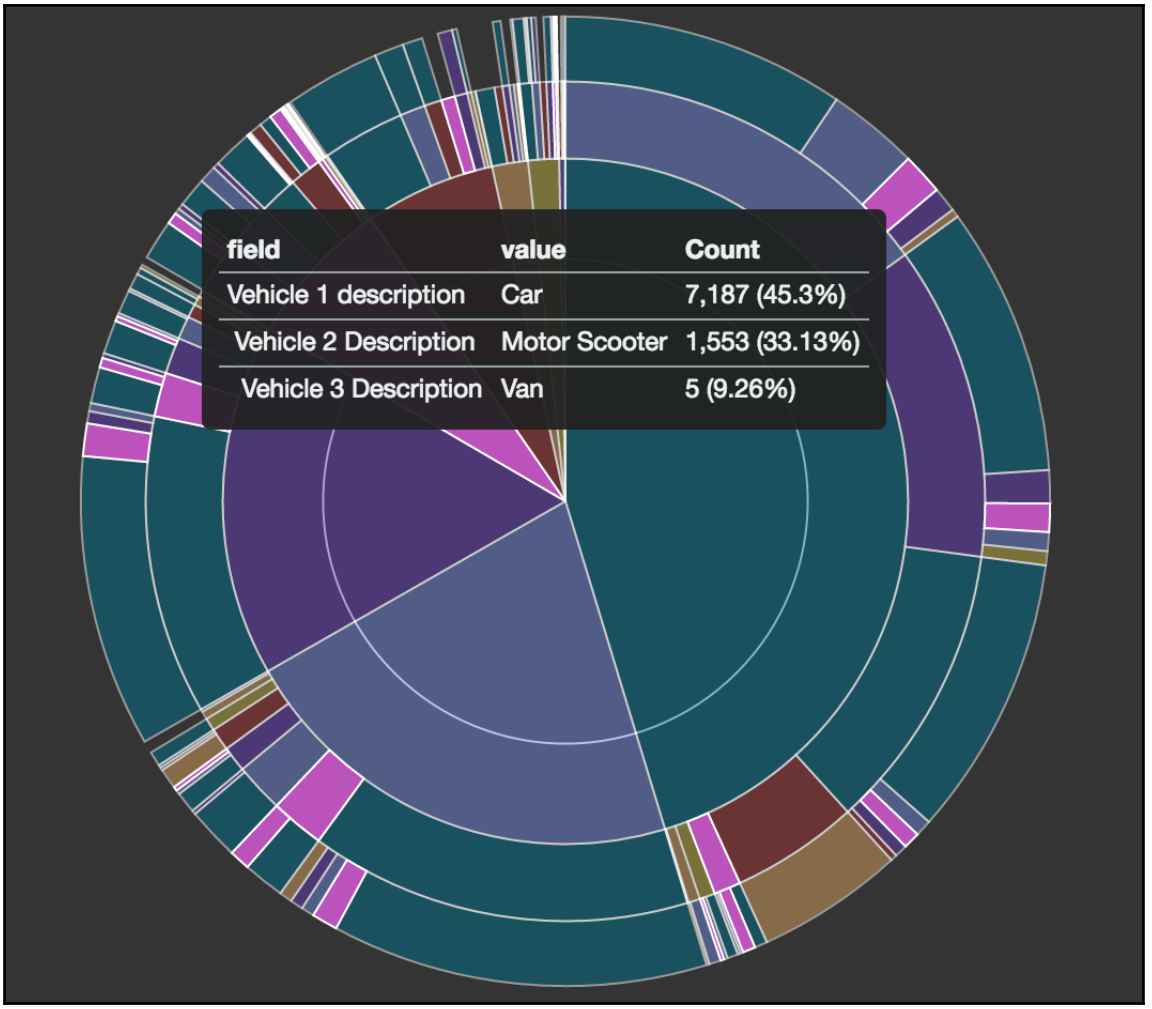
field
value
Count
Vehicle 1 description
Car
7,187 (45.3%)
Vehicle 2 Description
Motor Scooter
1,553 (33.13%)
Vehicle 3 Description
Van
5 (9.26%)

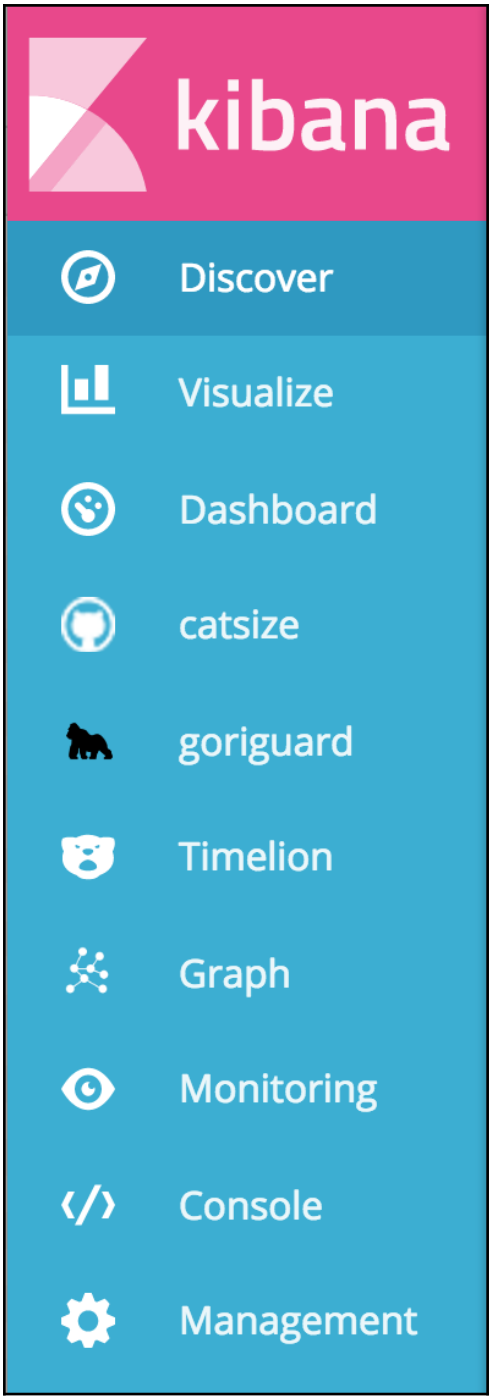
kibana
Discover
Visualize
Dashboard
catsize
goriguard
Timelion
Graph
Monitoring
Console
Management

Clusters / elasticsearch / Kibana

macbook-pro-de-bahaaldine.home

Status: Green ✔ **localhost:5603** OS Free Memory: **3.5 GB**

## OS Load

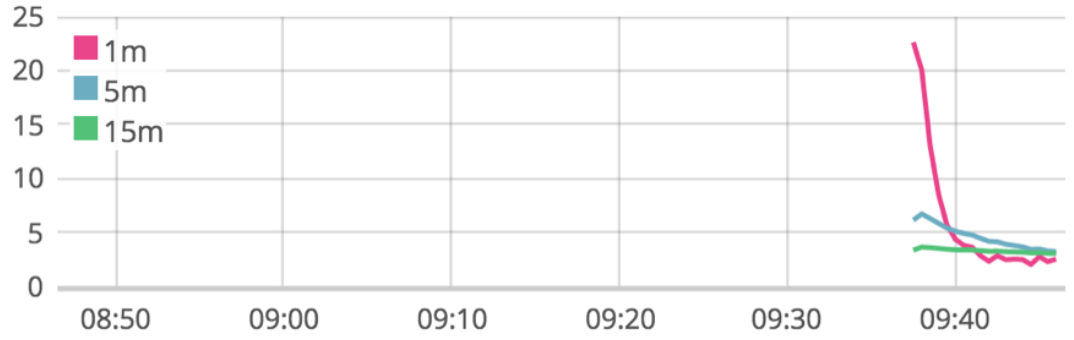


## Event Loop Delay (ms)

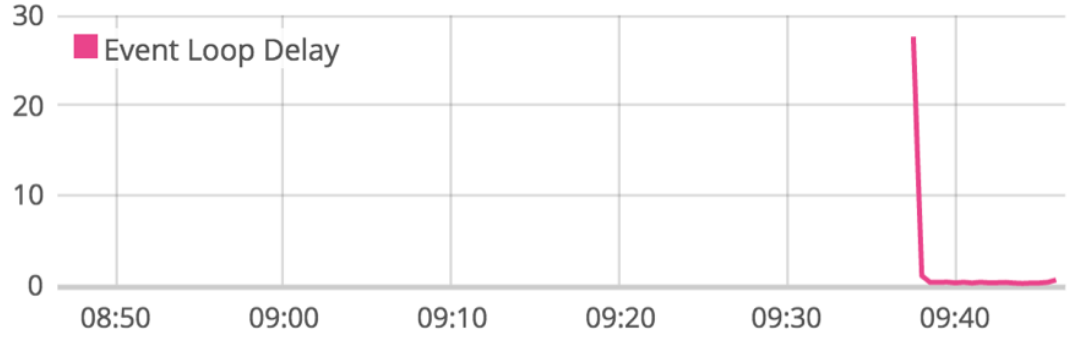


10:00
12:00
14:00
16:00
18:00
20:00
22:00

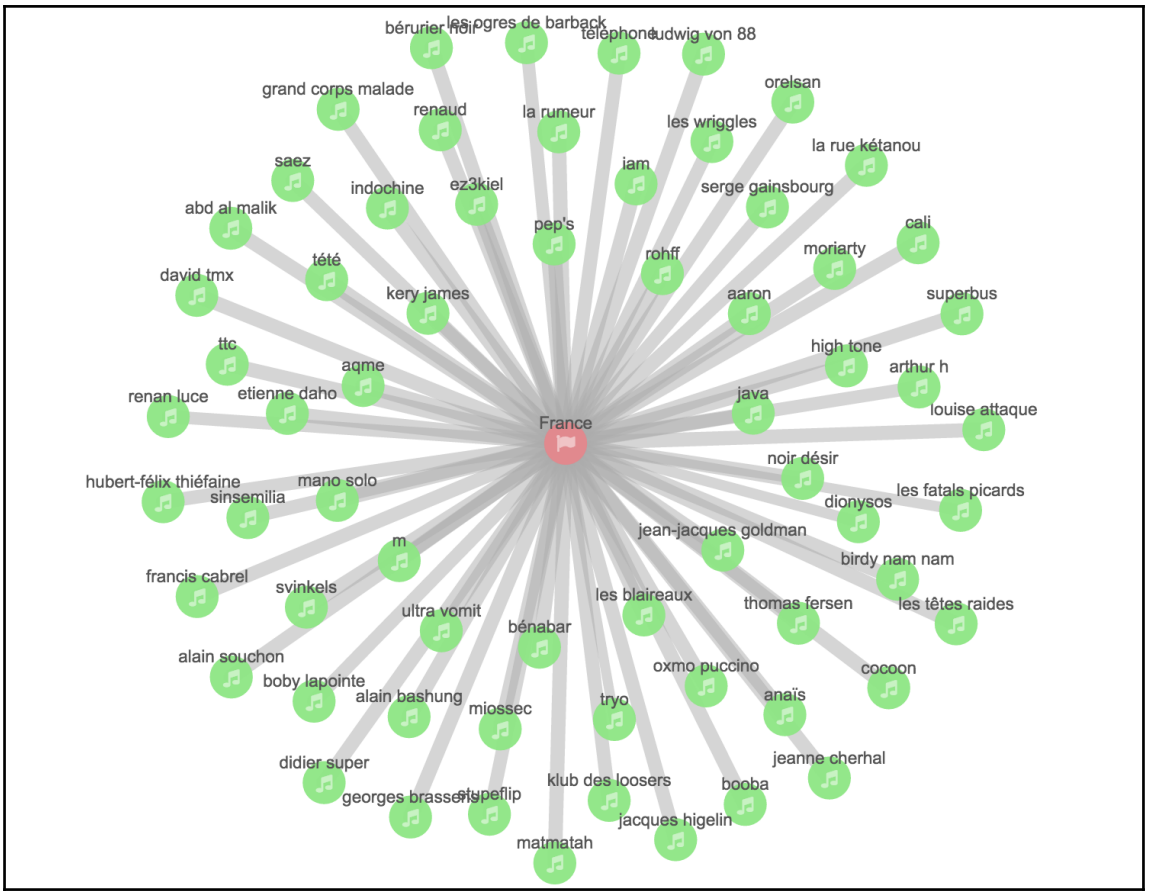
France
bérurier noir
les ogres de barback
téléphone
ludwig von 88
orelsan
grand corps malade
renaud
la rumeur
les wriggles
la rue kétanou
saez
iam
serge gainsbourg
indochine
ez3kiel
abd al malik
pep's
cali
moriarty
rohff
david tmx
tété
kery james
aaron
superbus
ttc
high tone
aqme
arthur h
renan luce
etienne daho
java
louise attaque
noir désir
hubert-félix thiéfaine
mano solo
sinsemilia
dionysos
les fatals picards
jean-jacques goldman
m
birdy nam nam
francis cabrel
svinkels
les blaireaux
thomas fersen
les têtes raides
ultra vomit
bénabar
alain souchon
oxmo puccino
cocoon
boby lapointe
alain bashung
miossec
tryo
anaïs
jeanne cherhal
didier super
klub des loosers
booba
georges brassens
stupeflip
jacques higelin
matmatah

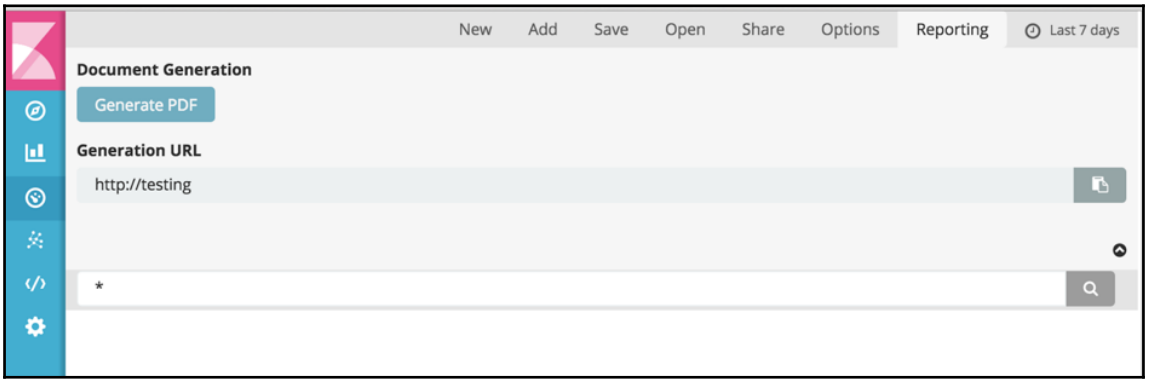
New
Add
Save
Open
Share
Options
Reporting
Last 7 days
Document Generation
Generate PDF
Generation URL
http://testing
*

# Chapter 2: Installing and Setting up Kibana 5.0

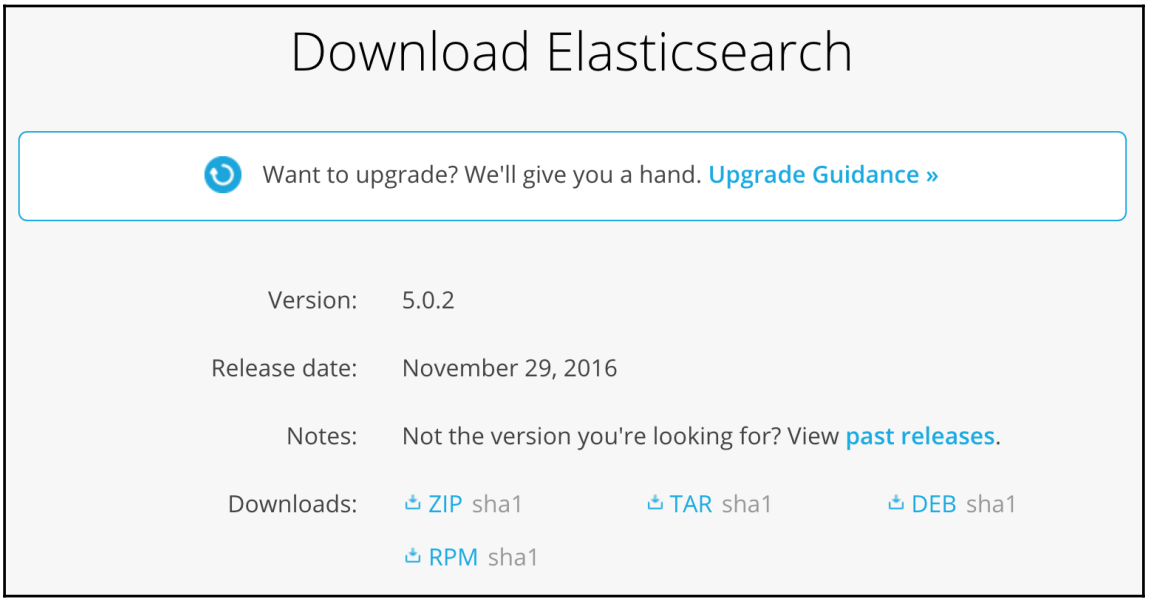

Downloads

# Kibana 5.0.0-alpha4

Not for production use! Requires Elasticsearch 5.0.0-alpha4.

| | | | |
|---|---|---|---|
| WINDOWS sha1 | MAC sha1 | LINUX 64-BIT sha1 | LINUX 32-BIT sha1 |
| RPM 64-BIT sha1 | RPM 32-BIT sha1 | DEB 64-BIT sha1 | DEB 32-BIT sha1 |

## Installation Steps

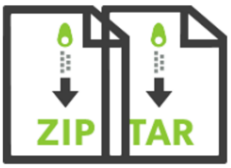


Download and unzip Kibana


- Extract your archive
- Open *config/kibana.yml* in


- Point your browser at *http://yourhost.com:5601*

localhost:5601/status

# Status: Green

pc55.home

| | | |
|---|---|---|
| Heap Total (MB) 84.27 | Heap Used (MB) 63.12 | Load 1.55, 1.73, 1.84 |
| Response Time Avg (ms) 0.00 | Response Time Max (ms) 0.00 | Requests Per Second 0.00 |

## Status Breakdown

| ID | Status |
|---|---|
| ui settings | Ready |
| plugin:kibana@1.0.0 | Ready |
| plugin:elasticsearch@1.0.0 | Kibana index ready |
| plugin:console@1.0.0 | Ready |
| plugin:kbn_doc_views@1.0.0 | Ready |
| plugin:kbn_vislib_vis_types@1.0.0 | Ready |
| plugin:markdown_vis@1.0.0 | Ready |
| plugin:metric_vis@1.0.0 | Ready |
| plugin:spy_modes@1.0.0 | Ready |
| plugin:status_page@1.0.0 | Ready |
| plugin:table_vis@1.0.0 | Ready |

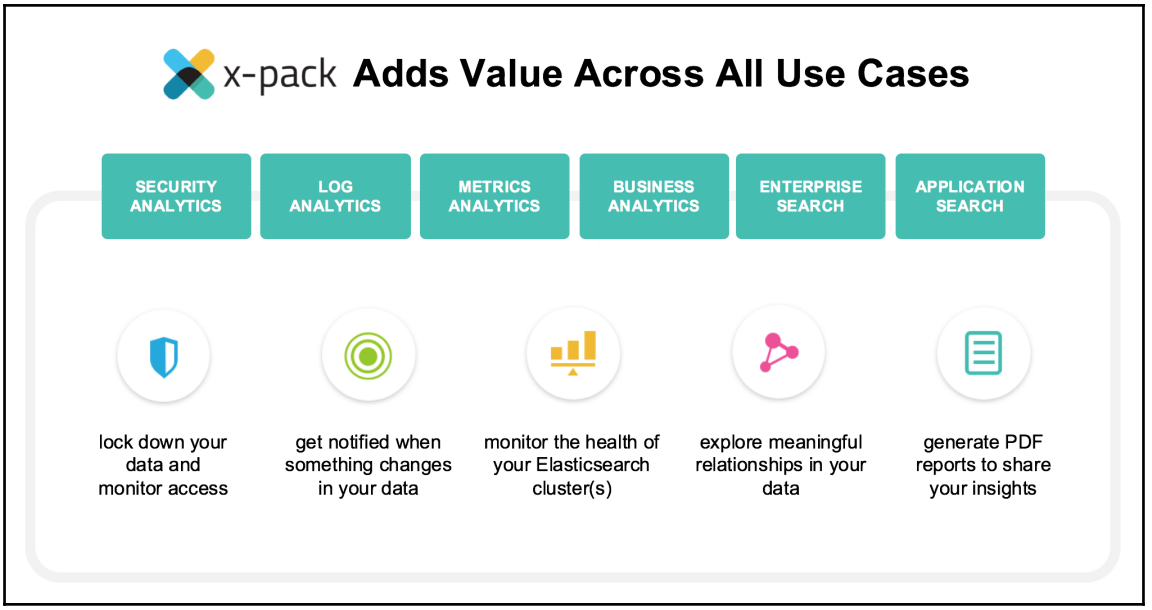
x-pack Adds Value Across All Use Cases
SECURITY ANALYTICS
LOG ANALYTICS
METRICS ANALYTICS
BUSINESS ANALYTICS
ENTERPRISE SEARCH
APPLICATION SEARCH
lock down your data and monitor access
get notified when something changes in your data
monitor the health of your Elasticsearch cluster(s)
explore meaningful relationships in your data
generate PDF reports to share your insights


localhost:9200
localhost:9200
Authentification requise
http://localhost:9200 nécessite un nom d'utilisateur et un mot de passe.
Nom d'utilisateur :
elastic
Mot de passe :
Annuler
Se connecter

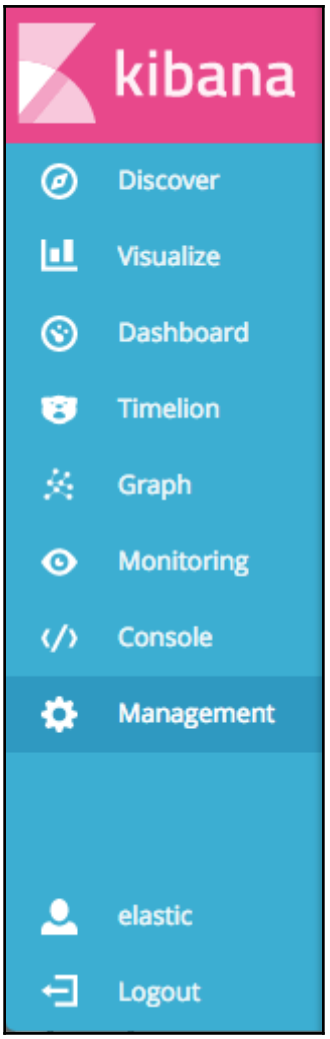
kibana
Discover
Visualize
Dashboard
Timelion
Graph
Monitoring
Console
Management
elastic
Logout


Import data
Create Index pattern
Discover data
Create Viz
Create Dashboards
More Viz

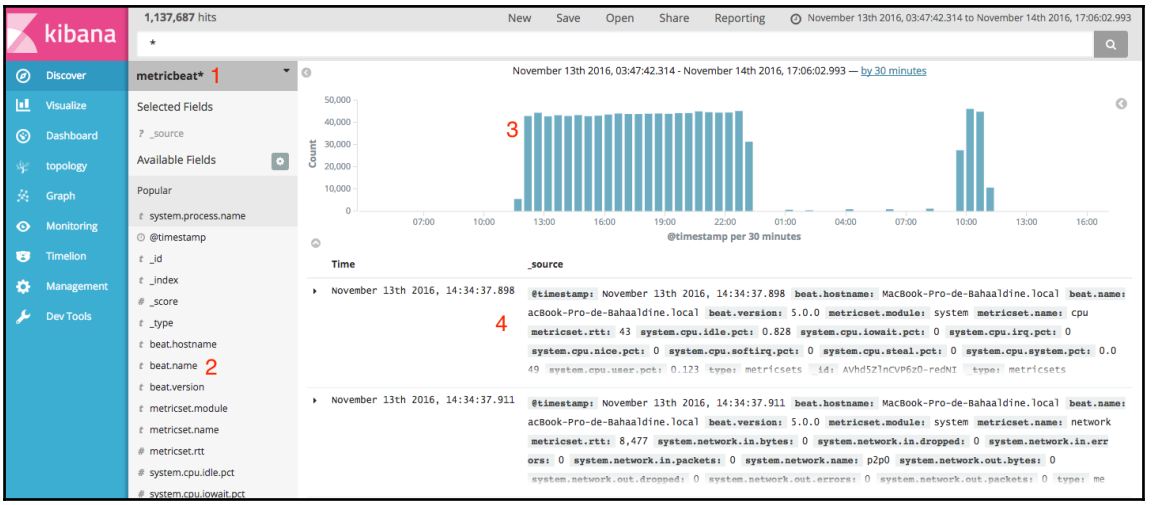
kibana
1,137,687 hits
New
Save
Open
Share
Reporting
November 13th 2016, 03:47:42.314 to November 14th 2016, 17:06:02.993
*
Discover
Visualize
Dashboard
topology
Graph
Monitoring
Timelion
Management
Dev Tools
metricbeat* 1
Selected Fields
? _source
Available Fields
Popular
t system.process.name
@timestamp
t _id
t _index
# _score
t _type
t beat.hostname
t beat.name 2
t beat.version
t metricset.module
t metricset.name
# metricset.rtt
# system.cpu.idle.pct
# system.cpu.iowait.pct
November 13th 2016, 03:47:42.314 - November 14th 2016, 17:06:02.993 — by 30 minutes
3
Count
50,000
40,000
30,000
20,000
10,000
0
07:00
10:00
13:00
16:00
19:00
22:00
01:00
04:00
07:00
10:00
13:00
16:00
@timestamp per 30 minutes
Time
_source
November 13th 2016, 14:34:37.898
4
@timestamp: November 13th 2016, 14:34:37.898 beat.hostname: MacBook-Pro-de-Bahaaldine.local beat.name: acBook-Pro-de-Bahaaldine.local beat.version: 5.0.0 metricset.module: system metricset.name: cpu metricset.rtt: 43 system.cpu.idle.pct: 0.828 system.cpu.iowait.pct: 0 system.cpu.irq.pct: 0 system.cpu.nice.pct: 0 system.cpu.softirq.pct: 0 system.cpu.steal.pct: 0 system.cpu.system.pct: 0.0
November 13th 2016, 14:34:37.911
@timestamp: November 13th 2016, 14:34:37.911 beat.hostname: MacBook-Pro-de-Bahaaldine.local beat.name: acBook-Pro-de-Bahaaldine.local beat.version: 5.0.0 metricset.module: system metricset.name: network metricset.rtt: 8,477 system.network.in.bytes: 0 system.network.in.dropped: 0 system.network.in.errors: 0 system.network.in.packets: 0 system.network.name: p2p0 system.network.out.bytes: 0

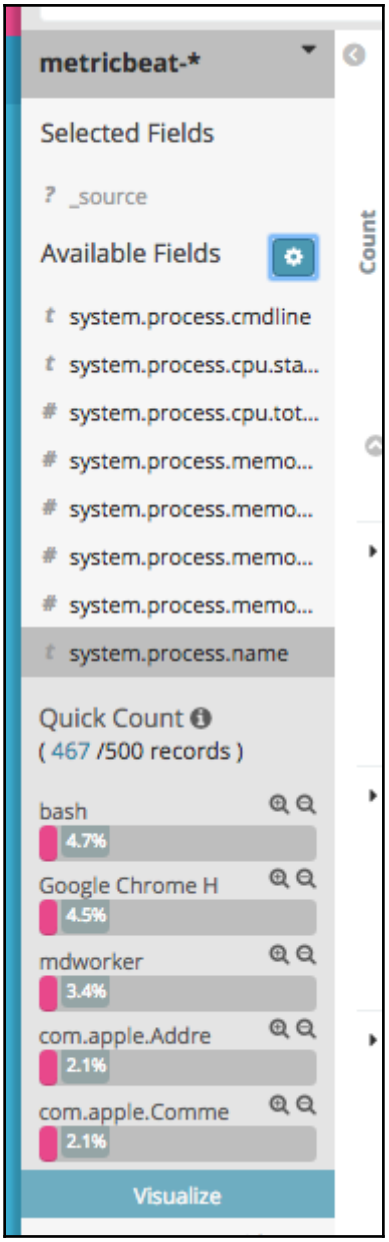
metricbeat-*
Selected Fields
? _source
Available Fields
t system.process.cmdline
t system.process.cpu.sta...
# system.process.cpu.tot...
# system.process.memo...
# system.process.memo...
# system.process.memo...
# system.process.memo...
t system.process.name
Quick Count
( 467 /500 records )
bash
4.7%
Google Chrome H
4.5%
mdworker
3.4%
com.apple.Addre
2.1%
com.apple.Comme
2.1%
Visualize
Count

| Time | _source |
| --- | --- |
| July 13th 2016, 10:50:07.492 | **@timestamp:** July 13th 2016, 10:50:07.492 **beat.hostname:** MacBook-Pro-de-Bahaaldine.local **beat.name:** MacBook-Pro-de-Bahaaldine.local **metricset.module:** system **metricset.name:** memory **metricset.rtt:** 2,051 **system.memory.actual.free:** 5,575,241,728 **system.memory.actual.used.bytes:** 10.808GB **system.memory.actual.used.pct:** 67.55% **system.memory.free:** 1.241GB **system.memory.swap.free:** 1,112, |

Table JSON

Link to /metricbeat-2016.07.13/metricsets/AVXjcuxCzpFq4rY7i6V3

| | | |
| --- | --- | --- |
| @timestamp | | July 13th 2016, 10:50:07.492 |
| t | _id | AVXjcuxCzpFq4rY7i6V3 |
| t | _index | metricbeat-2016.07.13 |
| # | _score | 2 |
| t | _type | metricsets |
| t | beat.hostname | MacBook-Pro-de-Bahaaldine.local |
| t | beat.name | MacBook-Pro-de-Bahaaldine.local |
| t | metricset.module | system |
| t | metricset.name | memory |
| # | metricset.rtt | 2,051 |
| # | system.memory.actual.free | 5,575,241,728 |
| # | system.memory.actual.used.bytes | 10.808GB |
| # | system.memory.actual.used.pct | 67.55% |
| # | system.memory.free | 1.241GB |
| # | system.memory.swap.free | 1,112,014,848 |
| # | system.memory.swap.total | 4,294,967,296 |
| # | system.memory.swap.used.bytes | 2.964GB |
| # | system.memory.swap.used.pct | 74.11% |
| # | system.memory.total | 16GB |
| # | system.memory.used.bytes | 14.759GB |
| # | system.memory.used.pct | 92.24% |
| t | type | metricsets |

## Create New Visualization

### Area chart

Great for stacked timelines in which the total of all series is more important than comparing any two or more series. Less useful for assessing the relative change of unrelated data points as changes in a series lower down the stack will have a difficult to gauge effect on the series above it.

### Data table

The data table provides a detailed breakdown, in tabular format, of the results of a composed aggregation. Tip, a data table is available from many other charts by clicking grey bar at the bottom of the chart.

### Line chart

Often the best chart for high density time series. Great for comparing one series to another. Be careful with sparse sets as the connection between points can be misleading.

### Markdown widget

Useful for displaying explanations or instructions for dashboards.

### Metric

One big number for all of your one big number needs. Perfect for showing a count of hits, or the exact average a numeric field.

### Pie chart

Pie charts are ideal for displaying the parts of some whole. For example, sales percentages by department.Pro Tip: Pie charts are best used sparingly, and with no more than 7 slices per pie.

### Tile map

Your source for geographic maps. Requires an elasticsearch geo_point field. More specifically, a field that is mapped as type:geo_point with latitude and longitude coordinates.

### Timeseries

Create timeseries charts using the timelion expression language. Perfect for computing and combining timeseries set with functions suchs as derivatives and moving averages

### Vertical bar chart

The goto chart for oh-so-many needs. Great for time and non-time data. Stacked or grouped, exact numbers or percentages. If you are not sure which chart you need, you could do worse than to start here.

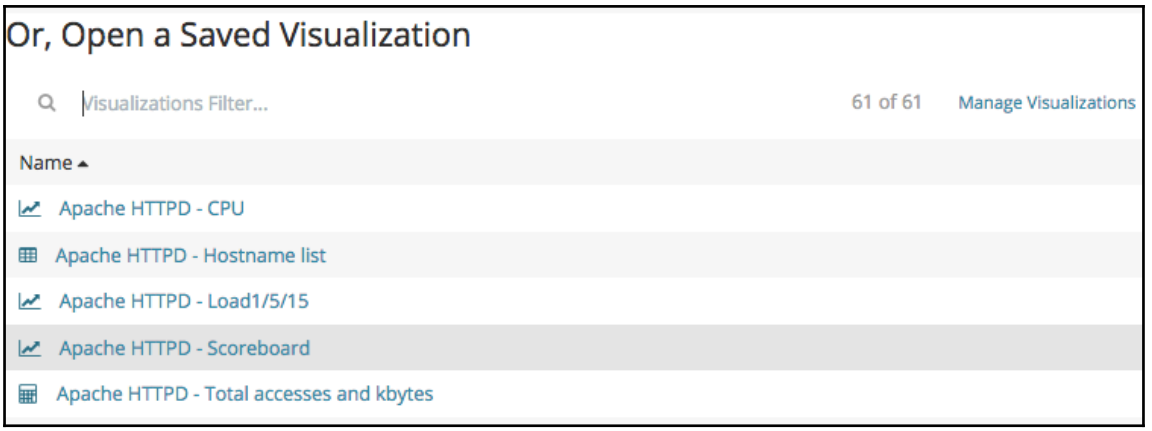
Or, Open a Saved Visualization
Visualizations Filter...
61 of 61
Manage Visualizations
Name
Apache HTTPD - CPU
Apache HTTPD - Hostname list
Apache HTTPD - Load1/5/15
Apache HTTPD - Scoreboard
Apache HTTPD - Total accesses and kbytes

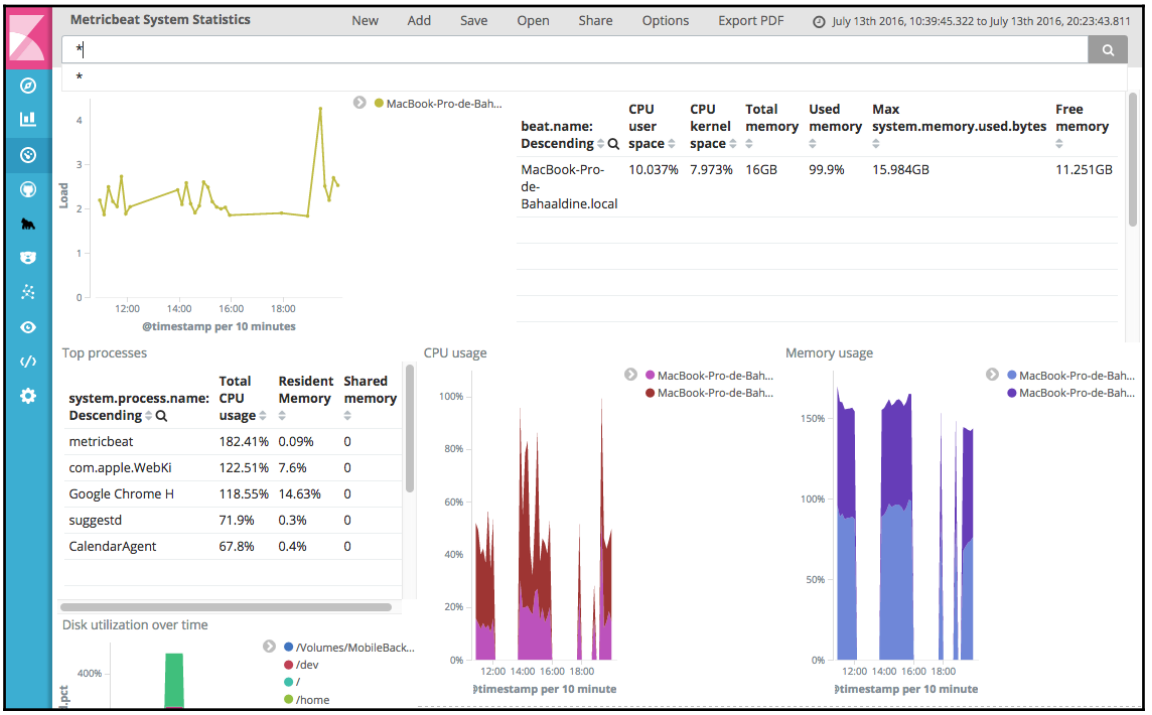
Metricbeat System Statistics
New
Add
Save
Open
Share
Options
Export PDF
July 13th 2016, 10:39:45.322 to July 13th 2016, 20:23:43.811
MacBook-Pro-de-Bah...
Load
@timestamp per 10 minutes
beat.name: Descending
CPU user space
CPU kernel space
Total memory
Used memory
Max system.memory.used.bytes
Free memory
MacBook-Pro-de-Bahaaldine.local
10.037%
7.973%
16GB
99.9%
15.984GB
11.251GB
Top processes
system.process.name: Descending
Total CPU usage
Resident Memory
Shared memory
metricbeat 182.41% 0.09% 0
com.apple.WebKi 122.51% 7.6% 0
Google Chrome H 118.55% 14.63% 0
suggestd 71.9% 0.3% 0
CalendarAgent 67.8% 0.4% 0
CPU usage
Memory usage
Disk utilization over time
/Volumes/MobileBack...
/dev
/
/home
@timestamp per 10 minute

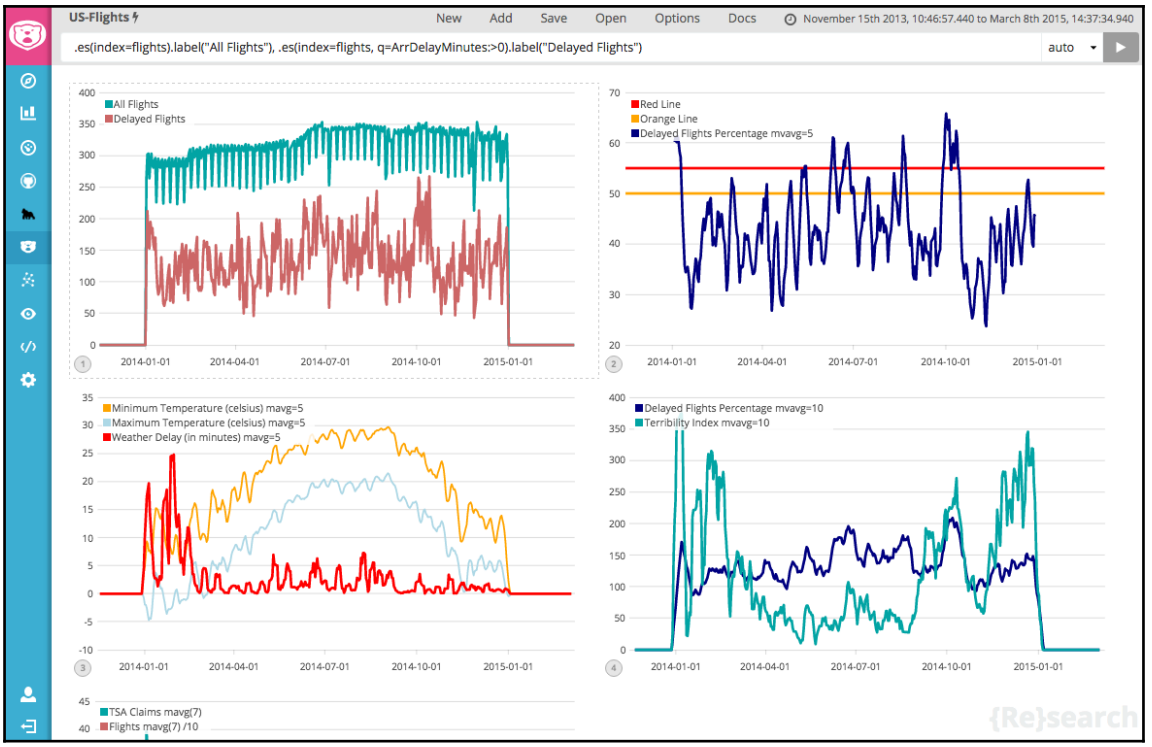
US-Flights
New
Add
Save
Open
Options
Docs
November 15th 2013, 10:46:57.440 to March 8th 2015, 14:37:34.940
.es(index=flights).label("All Flights"), .es(index=flights, q=ArrDelayMinutes:>0).label("Delayed Flights")
auto
All Flights
Delayed Flights
Red Line
Orange Line
Delayed Flights Percentage mvavg=5
Minimum Temperature (celsius) mavg=5
Maximum Temperature (celsius) mavg=5
Weather Delay (in minutes) mavg=5
Delayed Flights Percentage mvavg=10
Terribility Index mvavg=10
TSA Claims mavg(7)
Flights mavg(7) /10
2014-01-01
2014-04-01
2014-07-01
2014-10-01
2015-01-01
{Re}search

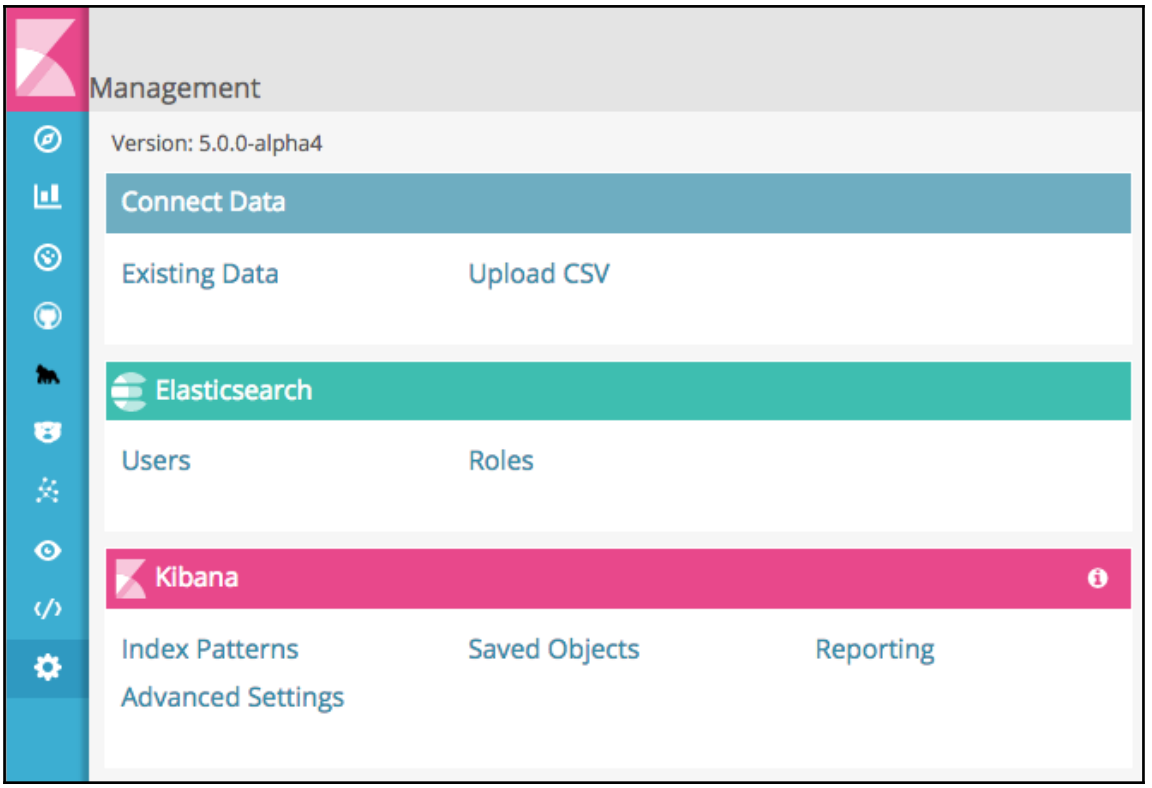
Management
Version: 5.0.0-alpha4
Connect Data
Existing Data
Upload CSV
Elasticsearch
Users
Roles
Kibana
Index Patterns
Saved Objects
Reporting
Advanced Settings

Dev Tools

Console

History Settings Help

```
GET metric*/_search
{
  "query": {
    "bool": {
      "must": [
        {
          "constant_score": {
            "filter": {
              "exists": {
                "field": "system.process.cpu.total.pct"
              }
            }
          }
        }
      ]
    }
  }
}
```


```
{
  "took": 42,
  "timed_out": false,
  "_shards": {
    "total": 10,
    "successful": 10,
    "failed": 0
  },
  "hits": {
    "total": 1038728,
    "max_score": 1,
    "hits": [
      {
        "_index": "metricbeat-2016.11.13",
        "_type": "metricsets",
        "_id": "AVhd5ZlnCVP6z0-redNY",
        "_score": 1,
        "_source": {
          "@timestamp": "2016-11-13T13:34:37.911Z",
          "beat": {
            "hostname": "MacBook-Pro-de-Bahaaldine.local",
            "name": "MacBook-Pro-de-Bahaaldine.local",
            "version": "5.0.0"
          },
          "metricset": {
            "module": "system",
            "name": "process",
            "rtt": 12557
          },
          "system": {
            "process": {
              "cmdline": "/System/Library/Frameworks/QuickLook
.framework/Resources/quicklookd.app/Contents/MacOS/quicklookd",
              "cpu": {
                "start_time": "2016-11-13T13:32:45.210Z",
                "total": {
                  "pct": 0
                }
              },
              "memory": {
                "rss": {
                  "bytes": 29442048,
                  "pct": 0.0017
                },
                "share": 0,
                "size": 3087671296
              },
              "name": "quicklookd",
              "pgid": 3313,
              "pid": 3313,
              "ppid": 1,
              "state": "running",
              "username": "bahaaldine"
```


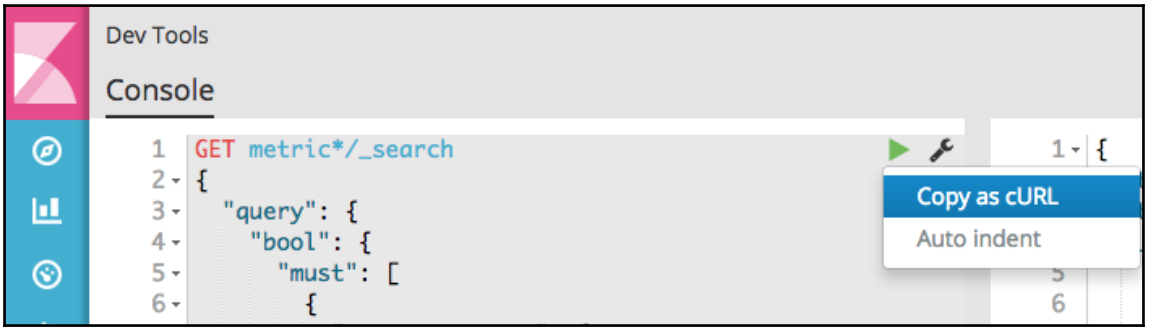

Dev Tools

Console Profiler

_all

Index Type

```
1 {
2   "query":{
3     "match_all" : {}
4   }
5 }
```

**Index:** .monitoring-kibana-2-2016.12.12

> **[gSXZ7sxgSLCWhytmHKIcJA][0]** 8.217ms

Cumulative Time: 24.686ms

**Index:** .monitoring-kibana-2-2016.12.06

> **[gSXZ7sxgSLCWhytmHKIcJA][0]** 24.686ms

Cumulative Time: 0.001ms

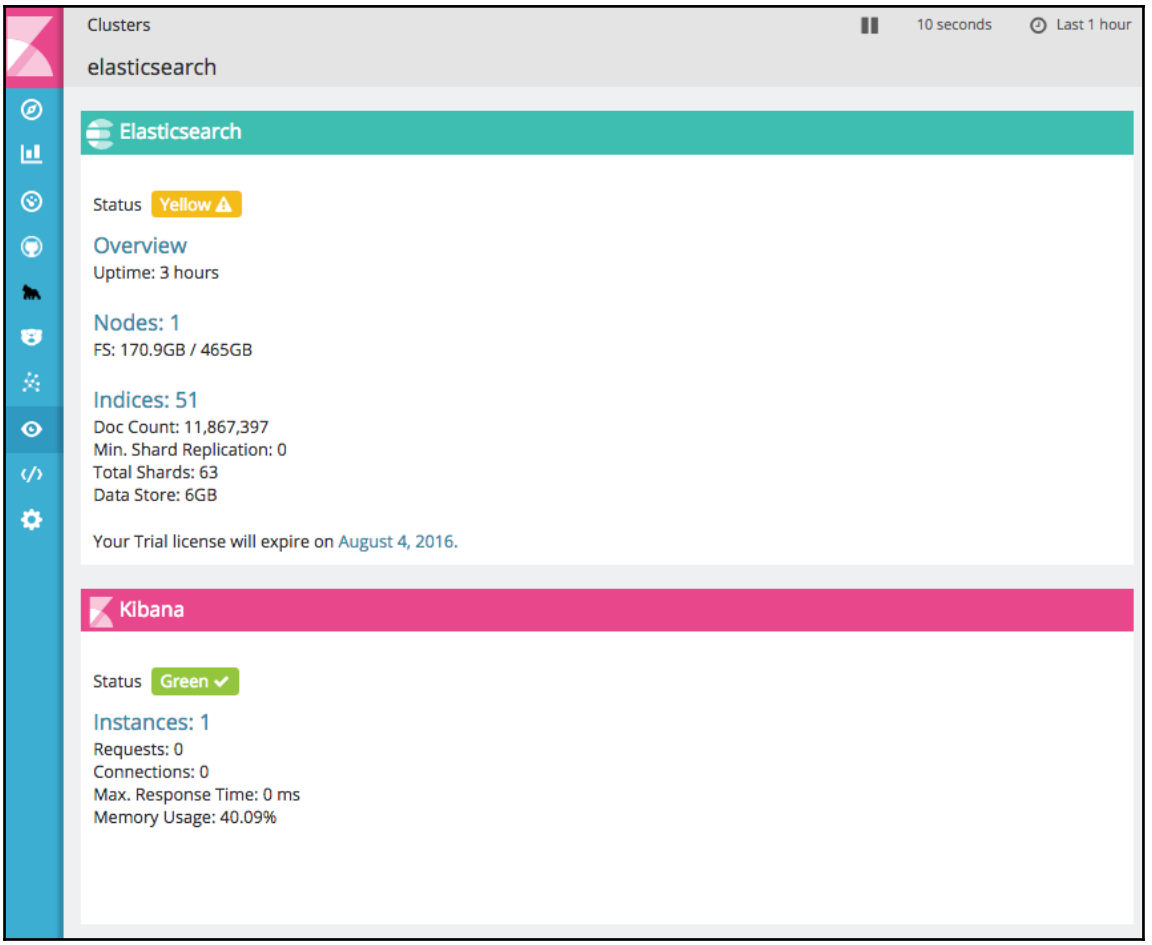
Clusters
10 seconds
Last 1 hour
elasticsearch
Elasticsearch
Status Yellow
Overview
Uptime: 3 hours
Nodes: 1
FS: 170.9GB / 465GB
Indices: 51
Doc Count: 11,867,397
Min. Shard Replication: 0
Total Shards: 63
Data Store: 6GB
Your Trial license will expire on August 4, 2016.
Kibana
Status Green
Instances: 1
Requests: 0
Connections: 0
Max. Response Time: 0 ms
Memory Usage: 40.09%

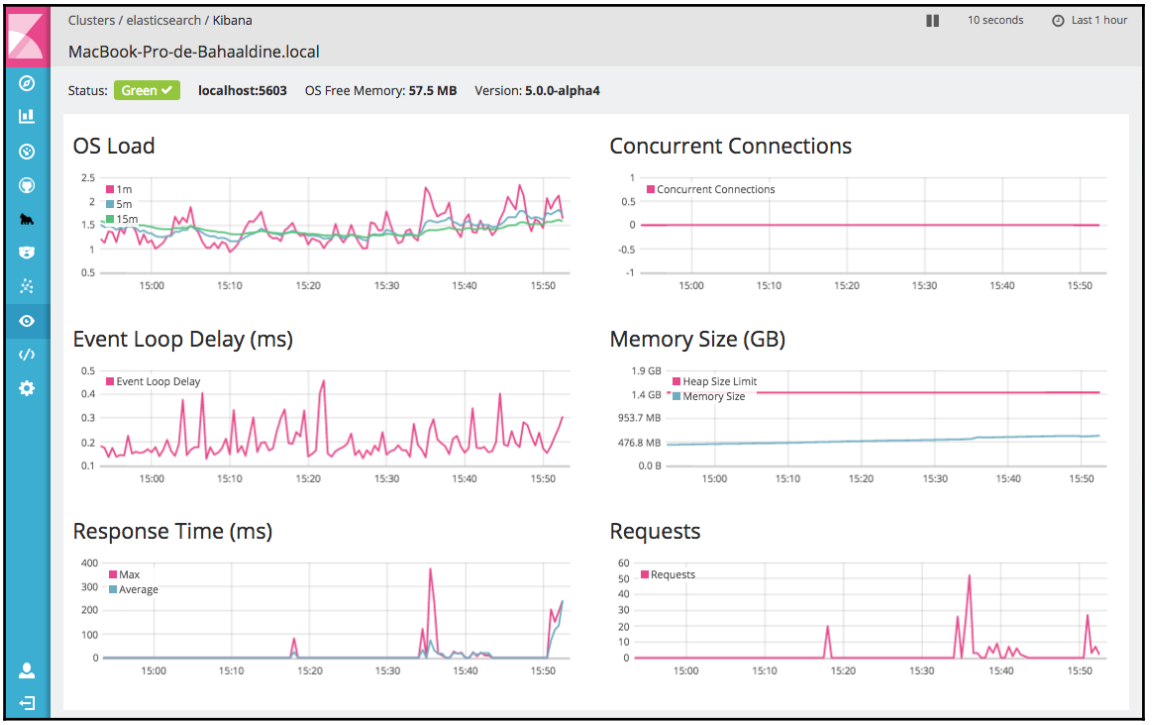
Clusters / elasticsearch / Kibana
10 seconds
Last 1 hour
MacBook-Pro-de-Bahaaldine.local
Status: Green
localhost:5603
OS Free Memory: 57.5 MB
Version: 5.0.0-alpha4
OS Load
1m
5m
15m
Concurrent Connections
Concurrent Connections
Event Loop Delay (ms)
Event Loop Delay
Memory Size (GB)
Heap Size Limit
Memory Size
Response Time (ms)
Max
Average
Requests
Requests
15:00
15:10
15:20
15:30
15:40
15:50

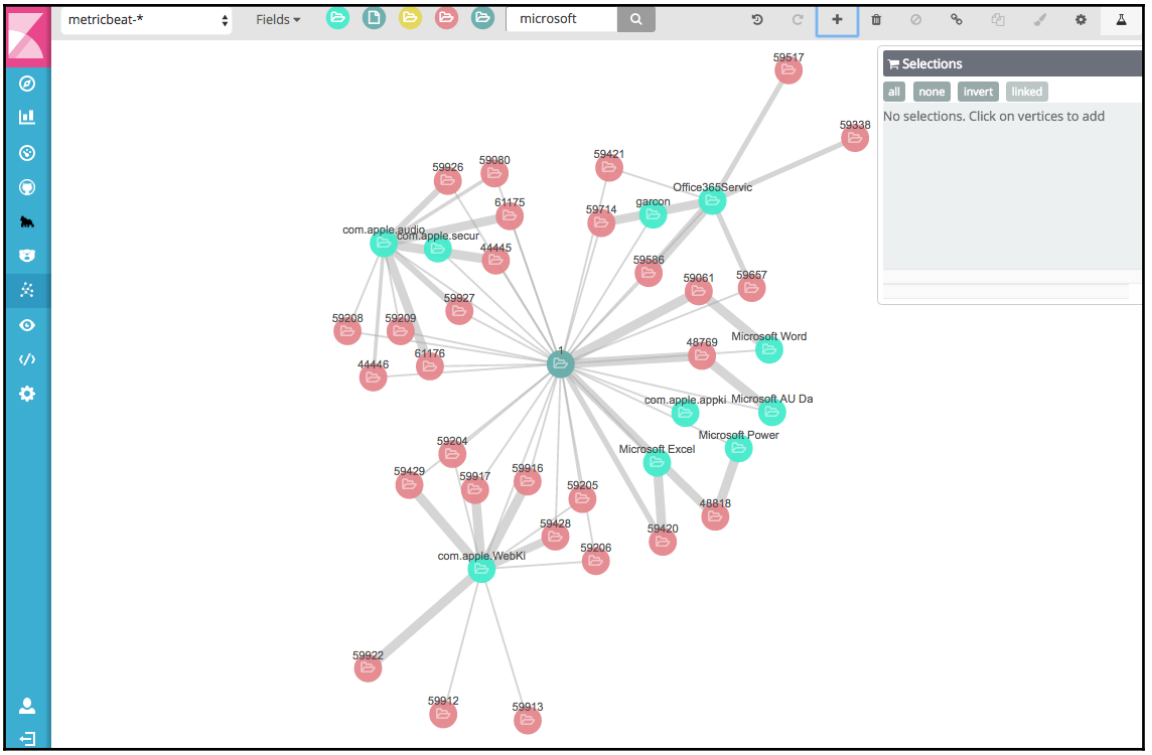
metricbeat-*
Fields
microsoft
Selections
all
none
invert
linked
No selections. Click on vertices to add
59517
59338
59421
59926
59080
Office365Servic
61175
59714
garcon
com.apple.audio
com.apple.secur
44445
59586
59061
59657
59927
59208
59209
48769
Microsoft Word
61176
44446
1
com.apple.appki
Microsoft AU Da
Microsoft Power
Microsoft Excel
59204
59429
59917
59916
59205
48818
59428
59420
59206
com.apple.WebKi
59922
59912
59913

# Chapter 3: Business Analytics with Kibana 5.0

Visualize / Step / 1

## Create New Visualization

### Area chart

Great for stacked timelines in which the total of all series is more important than comparing any two or more series. Less useful for assessing the relative change of unrelated data points as changes in a series lower down the stack will have a difficult to gauge effect on the series above it.

### Data table

The data table provides a detailed breakdown, in tabular format, of the results of a composed aggregation. Tip, a data table is available from many other charts by clicking grey bar at the bottom of the chart.

### Line chart

Often the best chart for high density time series. Great for comparing one series to another. Be careful with sparse sets as the connection between points can be misleading.

Visualize / Step / 2

## From a New Search, Select Index

Filter...

10 of 10

Name

accident*

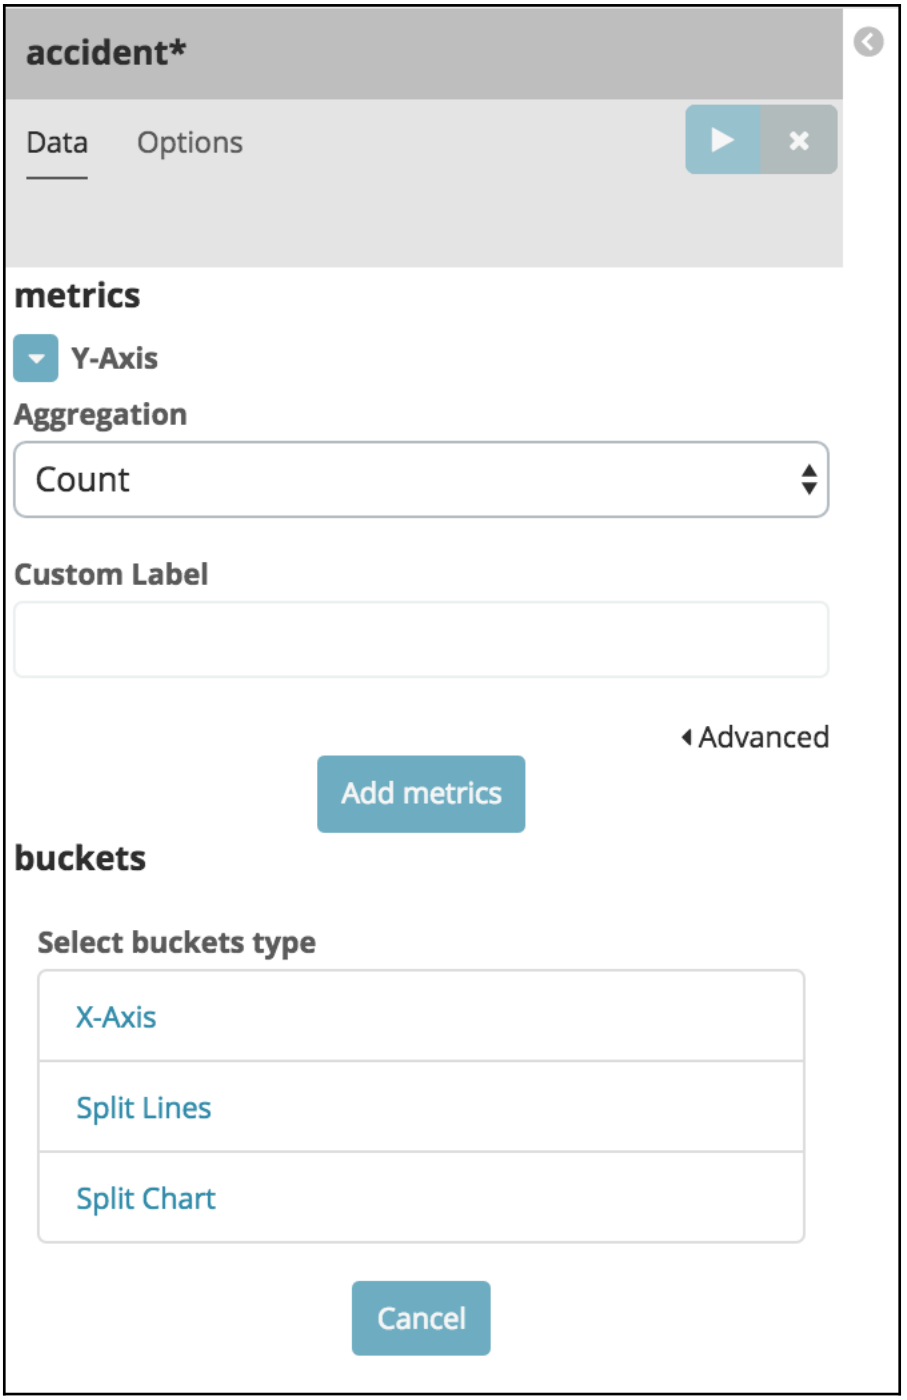
accident*
Data
Options
metrics
Y-Axis
Aggregation
Count
Custom Label
Advanced
Add metrics
buckets
Select buckets type
X-Axis
Split Lines
Split Chart
Cancel

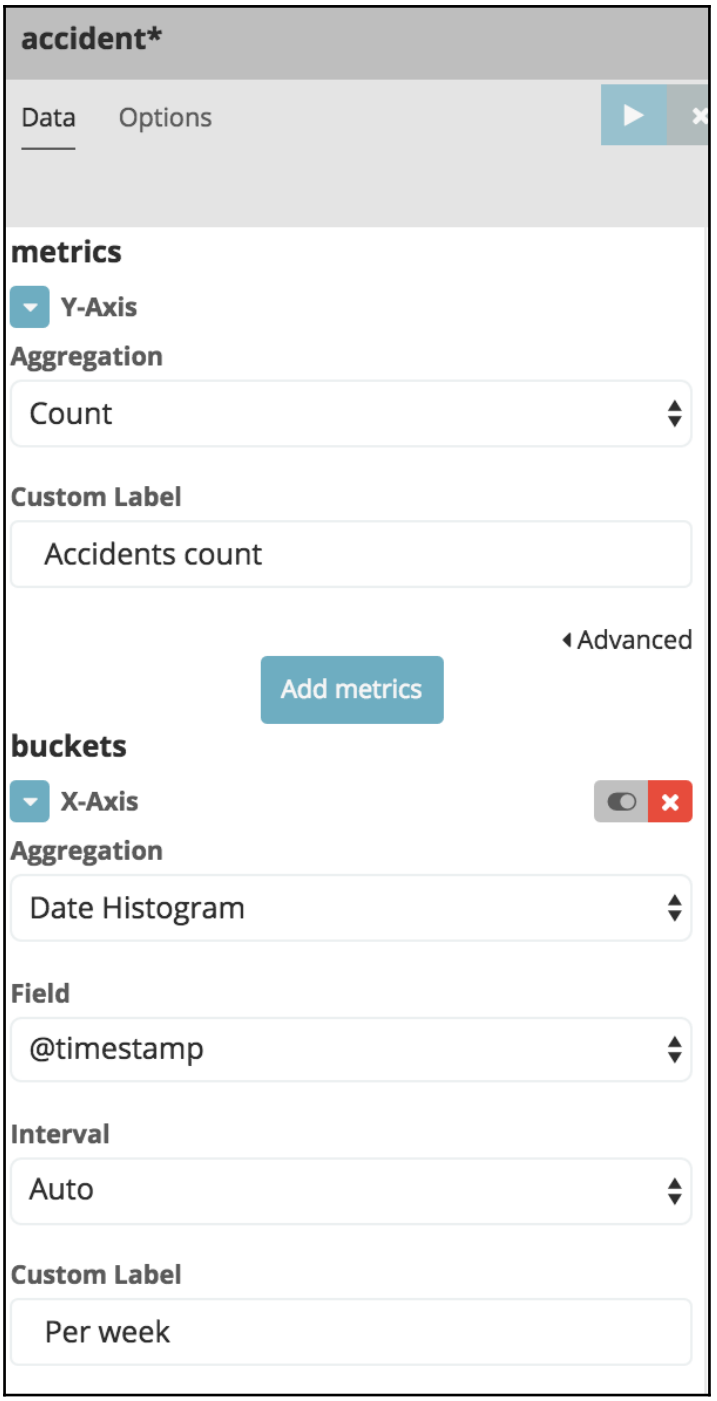
accident*
Data
Options
metrics
Y-Axis
Aggregation
Count
Custom Label
Accidents count
Advanced
Add metrics
buckets
X-Axis
Aggregation
Date Histogram
Field
@timestamp
Interval
Auto
Custom Label
Per week

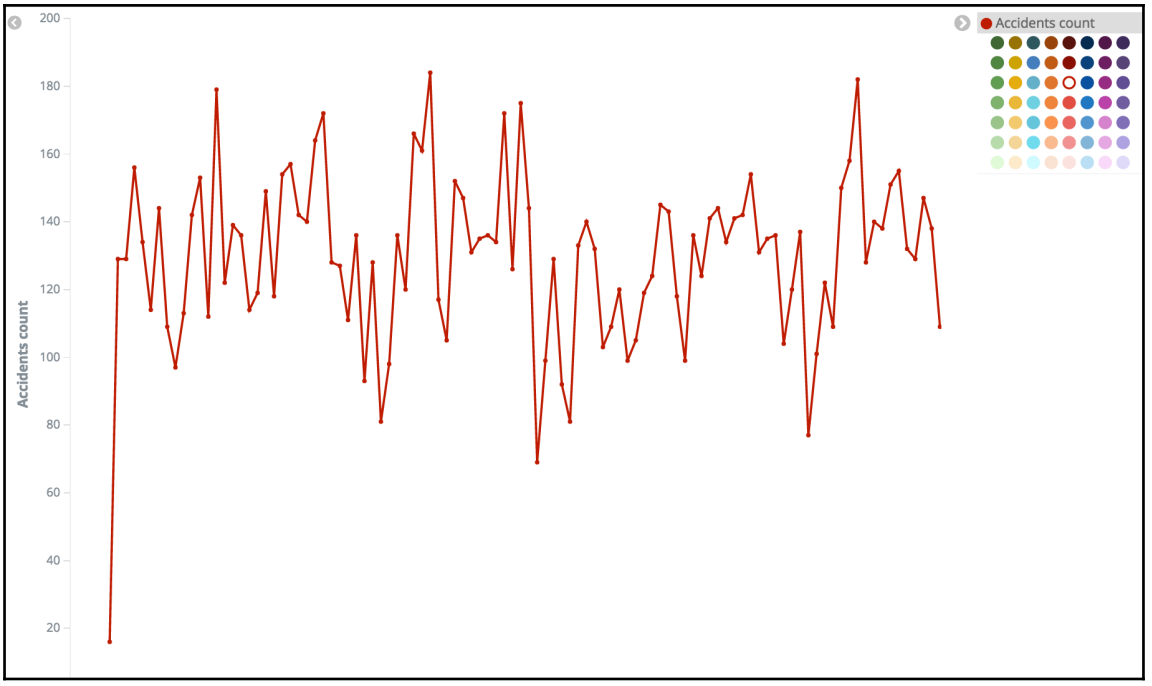
Accidents count
200
180
160
140
120
100
80
60
40
20
Accidents count

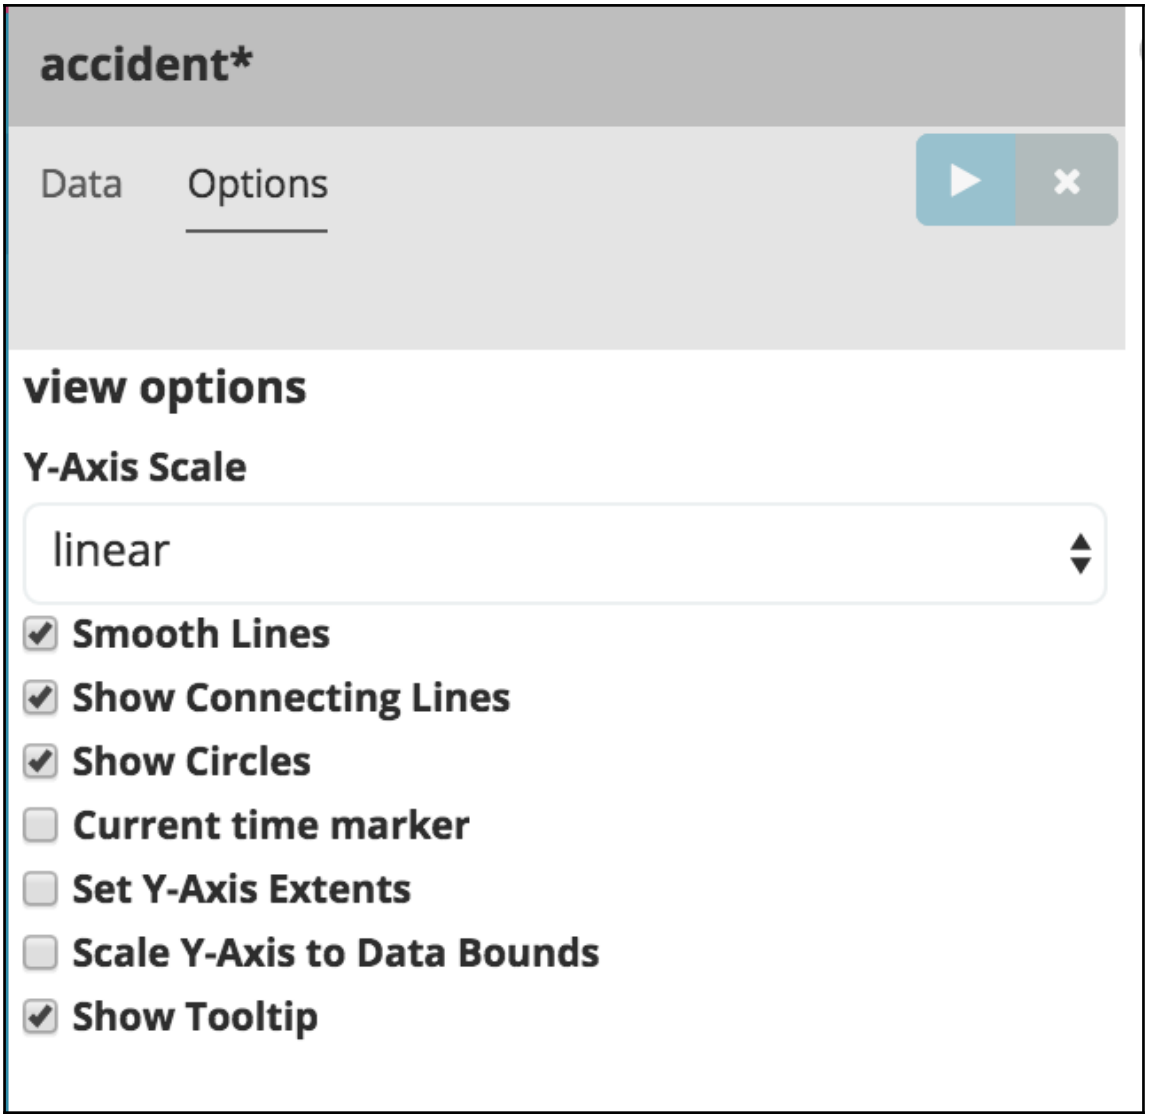
accident*
Data
Options
view options
Y-Axis Scale
linear
Smooth Lines
Show Connecting Lines
Show Circles
Current time marker
Set Y-Axis Extents
Scale Y-Axis to Data Bounds
Show Tooltip

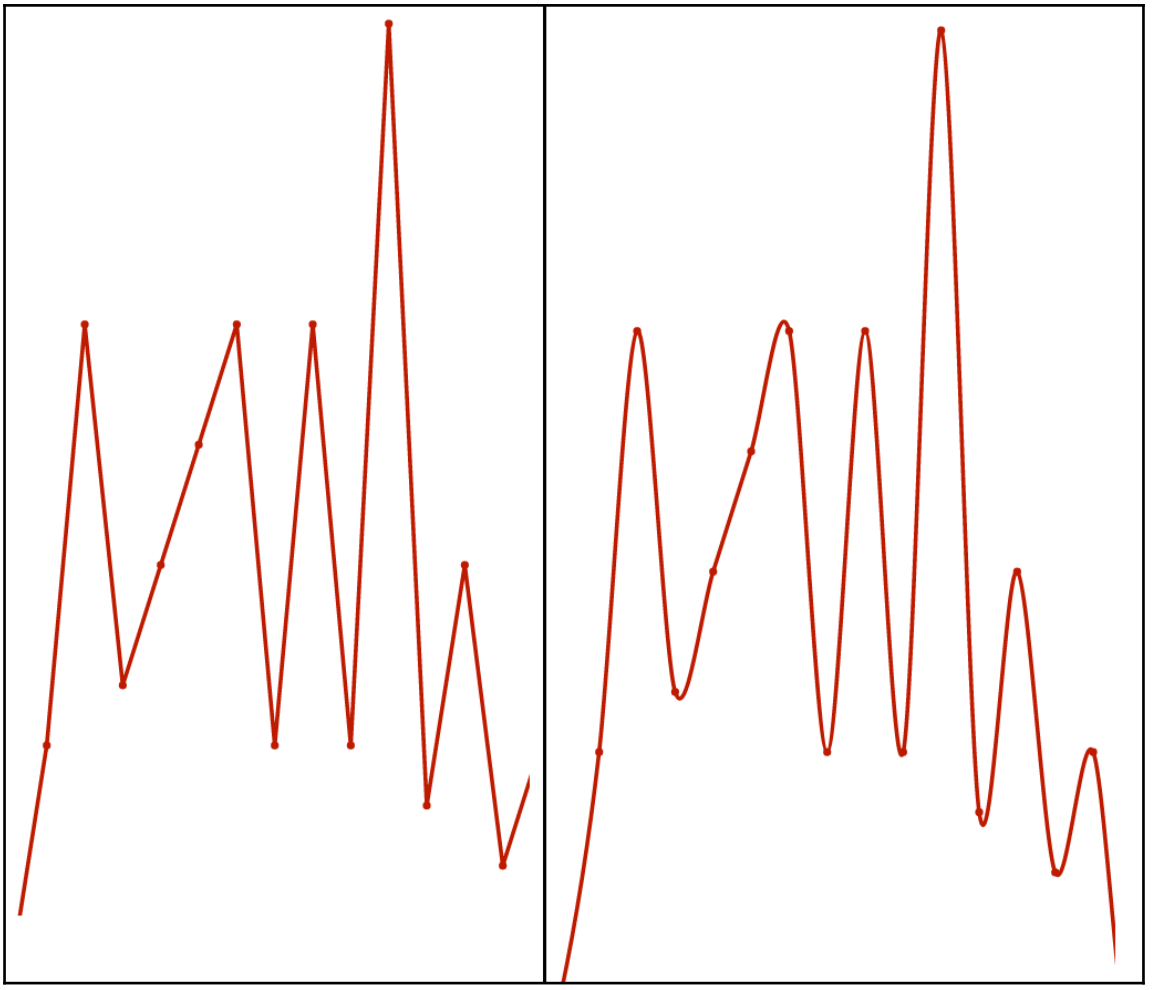


20
10
0
2012-04-01
2012-07-01

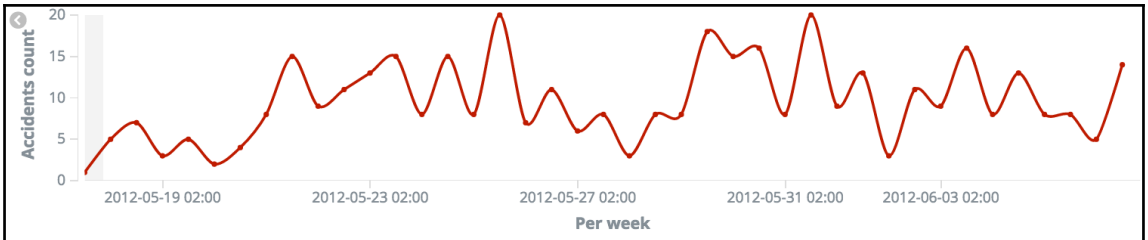


Table | Request | Response | Statistics | Debug

| Per week | Accidents count |
|---|---|
| May 17th 2012, 12:00:00.000 | 1 |
| May 18th 2012, 00:00:00.000 | 5 |
| May 18th 2012, 12:00:00.000 | 7 |
| May 19th 2012, 00:00:00.000 | 3 |
| May 19th 2012, 12:00:00.000 | 5 |
| May 20th 2012, 00:00:00.000 | 2 |
| May 20th 2012, 12:00:00.000 | 4 |
| May 21st 2012, 00:00:00.000 | 8 |
| May 21st 2012, 12:00:00.000 | 15 |
| May 22nd 2012, 00:00:00.000 | 9 |

Export: Raw Formatted

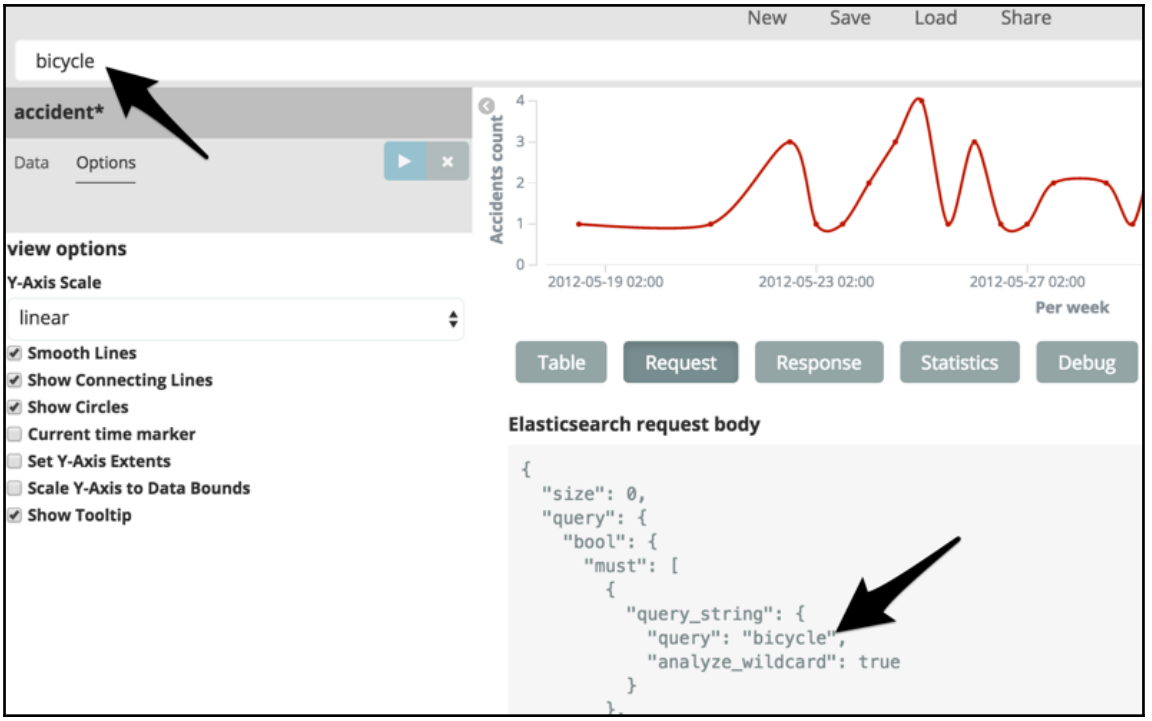
New
Save
Load
Share
bicycle
accident*
Data
Options
view options
Y-Axis Scale
linear
Smooth Lines
Show Connecting Lines
Show Circles
Current time marker
Set Y-Axis Extents
Scale Y-Axis to Data Bounds
Show Tooltip
Accidents count
4
3
2
1
0
2012-05-19 02:00
2012-05-23 02:00
2012-05-27 02:00
Per week
Table
Request
Response
Statistics
Debug
Elasticsearch request body
{
"size": 0,
"query": {
"bool": {
"must": [
{
"query_string": {
"query": "bicycle",
"analyze_wildcard": true
}

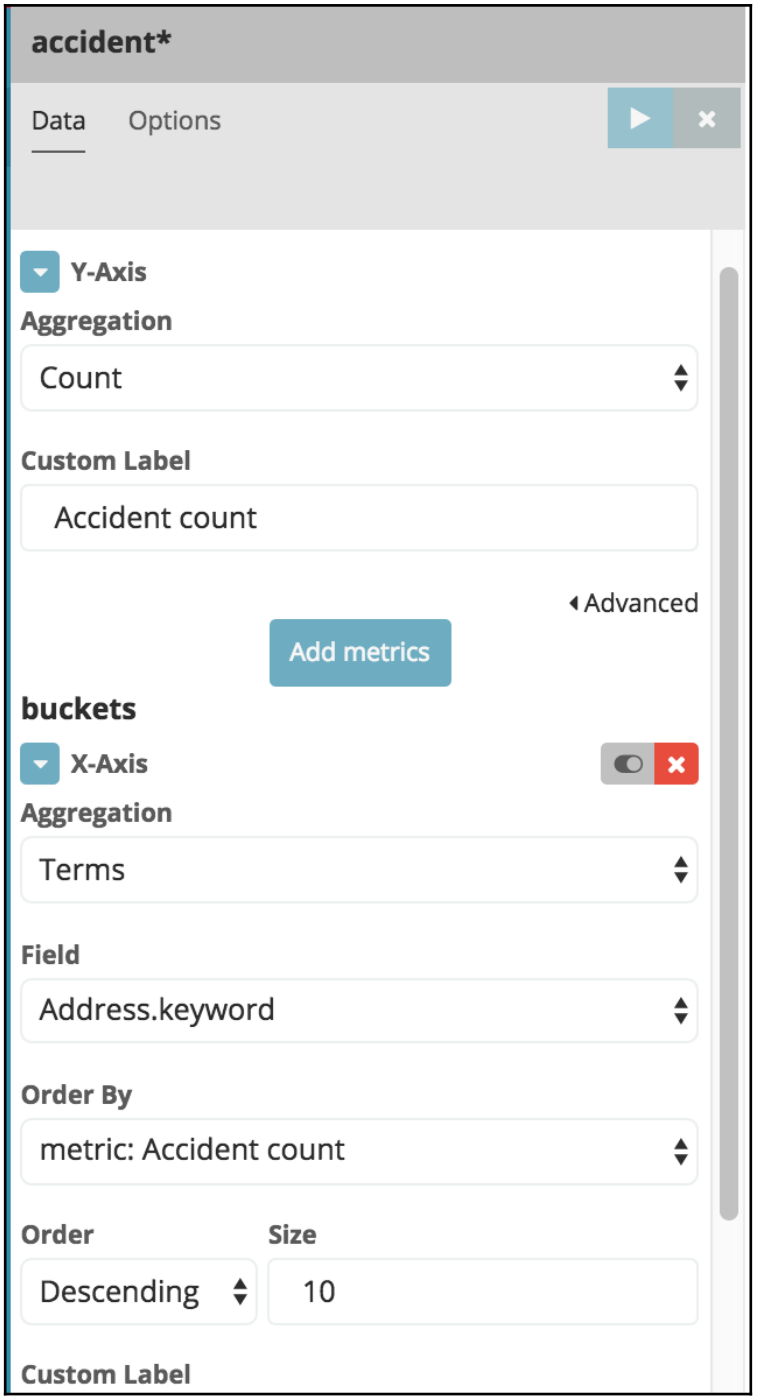
accident*
Data
Options
Y-Axis
Aggregation
Count
Custom Label
Accident count
Advanced
Add metrics
buckets
X-Axis
Aggregation
Terms
Field
Address.keyword
Order By
metric: Accident count
Order
Descending
Size
10
Custom Label

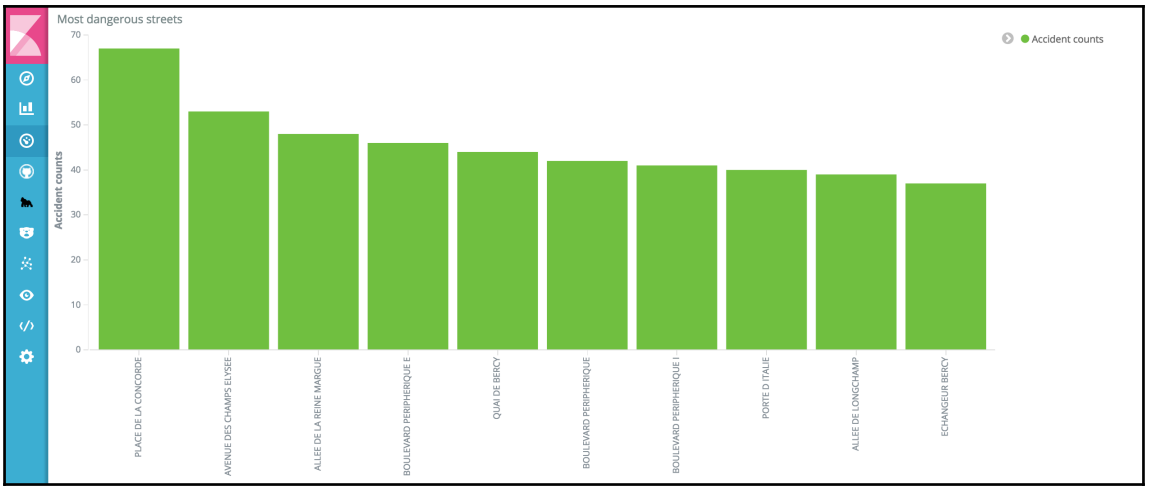
Most dangerous streets
Accident counts
Accident counts
70
60
50
40
30
20
10
0
PLACE DE LA CONCORDE
AVENUE DES CHAMPS ELYSEE
ALLEE DE LA REINE MARGUE
BOULEVARD PERIPHERIQUE E
QUAI DE BERCY
BOULEVARD PERIPHERIQUE
BOULEVARD PERIPHERIQUE I
PORTE D ITALIE
ALLEE DE LONGCHAMP
ECHANGEUR BERCY

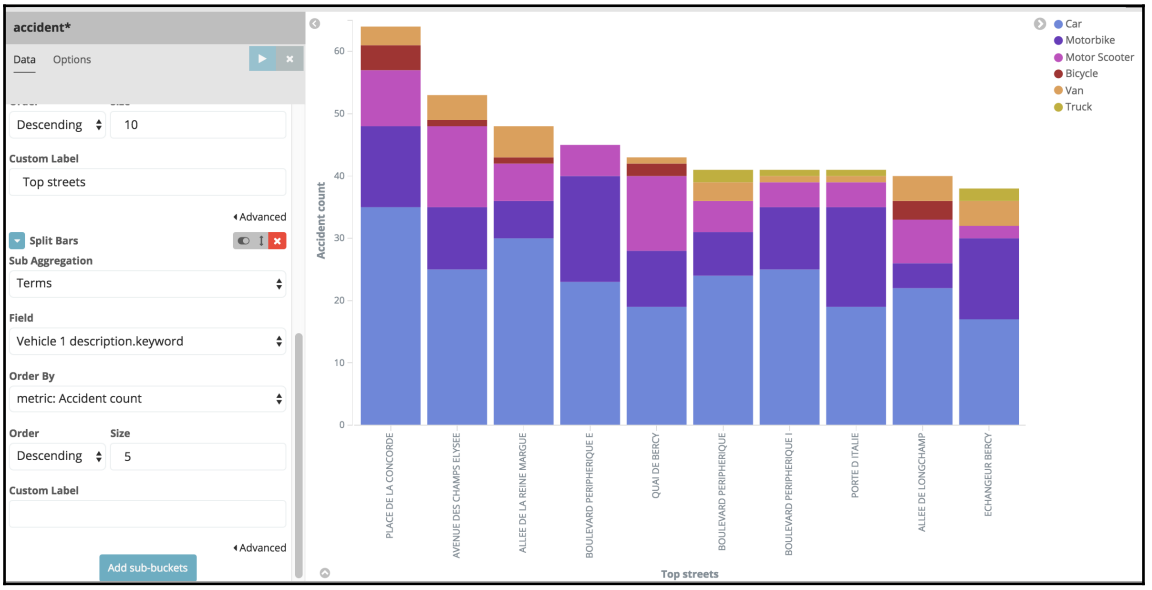
accident*
Data
Options
Descending
10
Custom Label
Top streets
Advanced
Split Bars
Sub Aggregation
Terms
Field
Vehicle 1 description.keyword
Order By
metric: Accident count
Order
Size
Descending
5
Custom Label
Advanced
Add sub-buckets
Car
Motorbike
Motor Scooter
Bicycle
Van
Truck
Accident count
60
50
40
30
20
10
0
PLACE DE LA CONCORDE
AVENUE DES CHAMPS ELYSEE
ALLEE DE LA REINE MARGUE
BOULEVARD PERIPHERIQUE E
QUAI DE BERCY
BOULEVARD PERIPHERIQUE
BOULEVARD PERIPHERIQUE I
PORTE D ITALIE
ALLEE DE LONGCHAMP
ECHANGEUR BERCY
Top streets

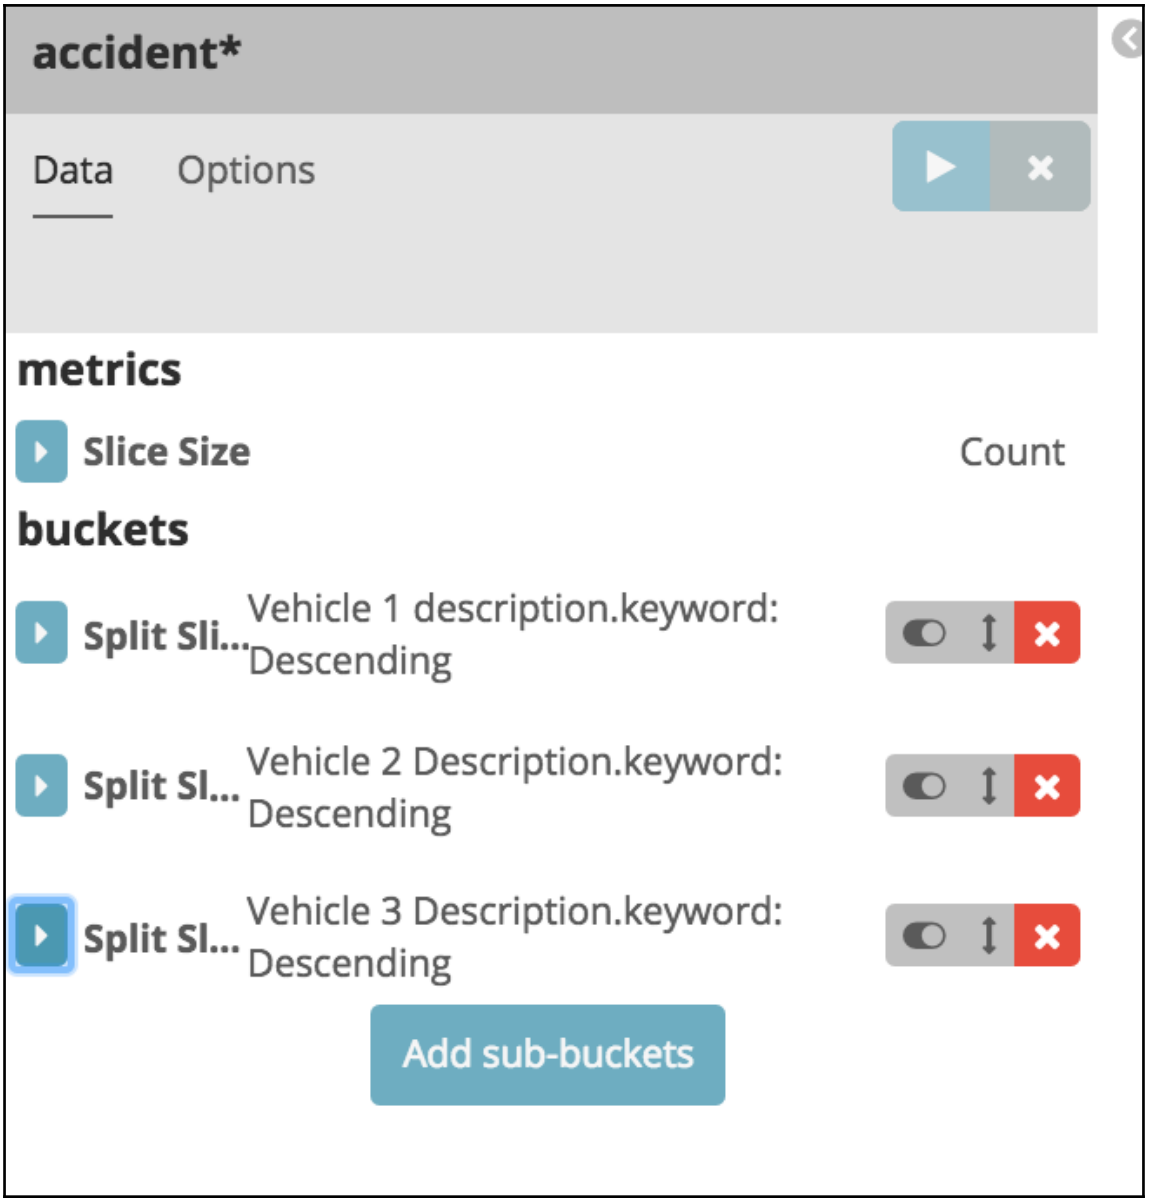
accident*
Data
Options
metrics
Slice Size
Count
buckets
Split Sli...
Vehicle 1 description.keyword: Descending
Split Sl...
Vehicle 2 Description.keyword: Descending
Split Sl...
Vehicle 3 Description.keyword: Descending
Add sub-buckets

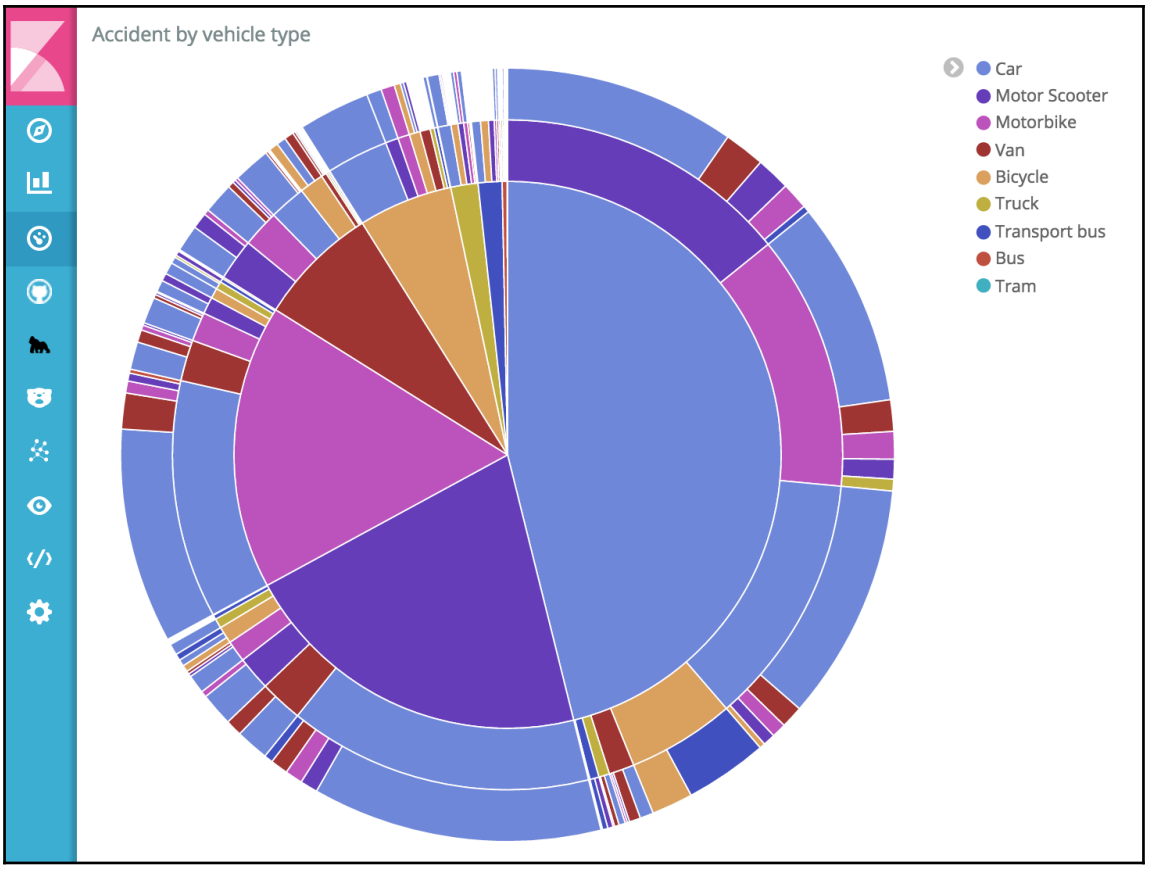
Accident by vehicle type
Car
Motor Scooter
Motorbike
Van
Bicycle
Truck
Transport bus
Bus
Tram

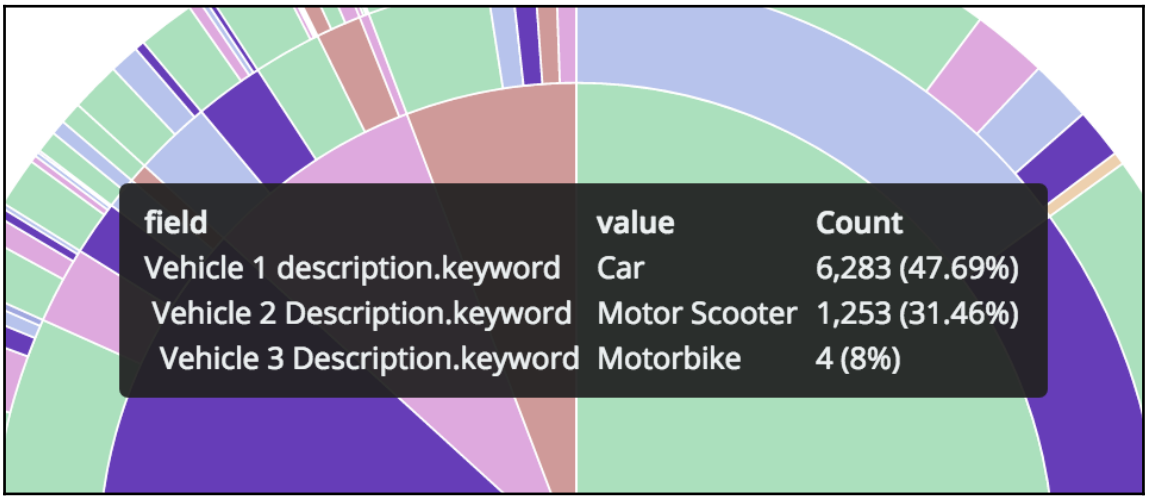
field
value
Count
Vehicle 1 description.keyword
Car
6,283 (47.69%)
Vehicle 2 Description.keyword
Motor Scooter
1,253 (31.46%)
Vehicle 3 Description.keyword
Motorbike
4 (8%)

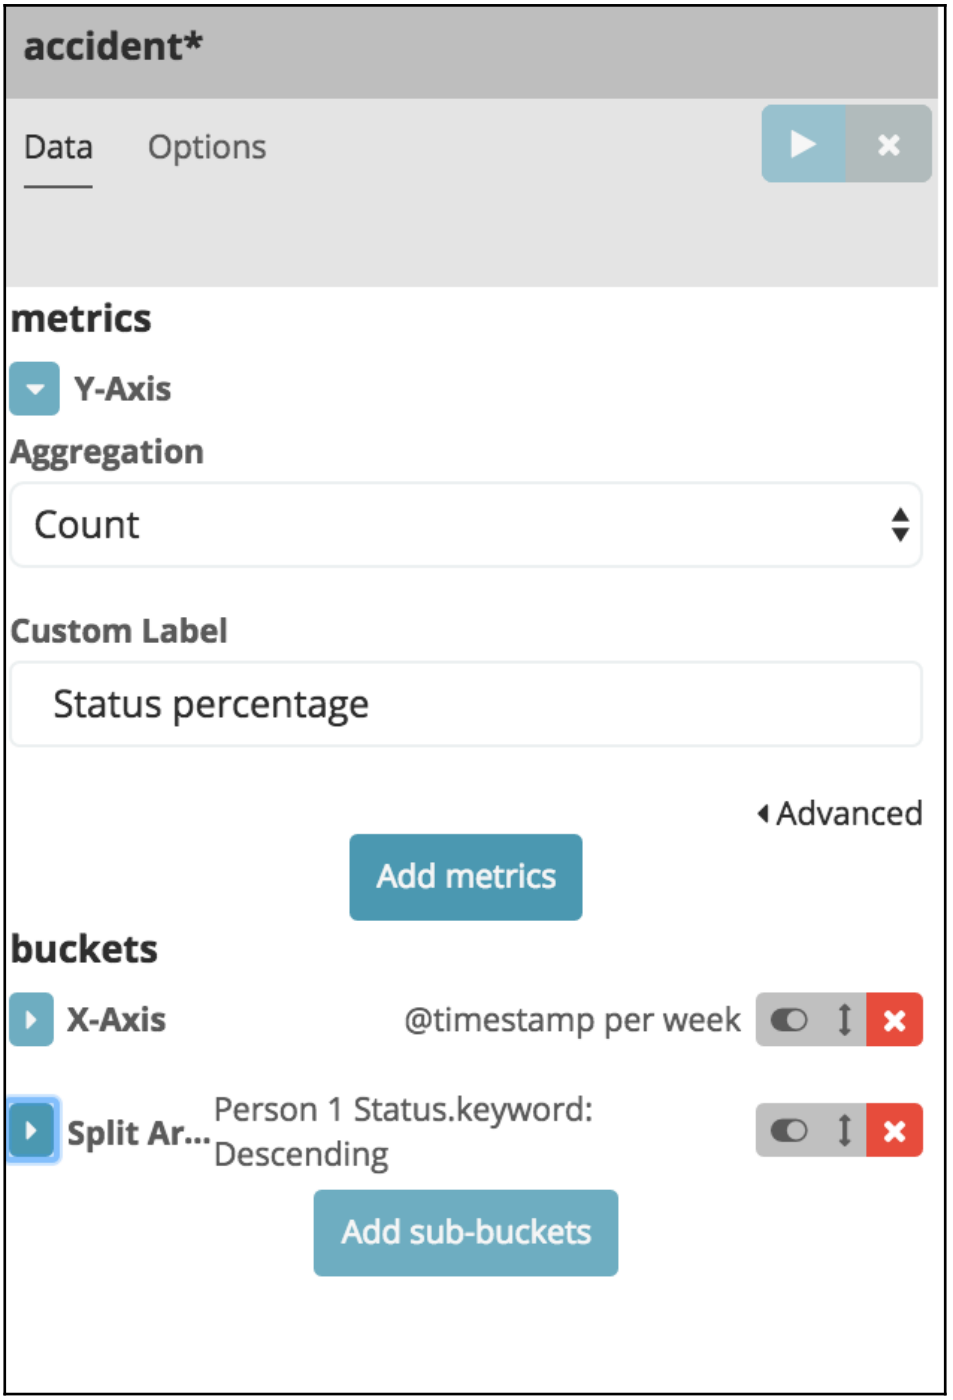
accident*
Data
Options
metrics
Y-Axis
Aggregation
Count
Custom Label
Status percentage
Advanced
Add metrics
buckets
X-Axis
@timestamp per week
Split Ar...
Person 1 Status.keyword: Descending
Add sub-buckets

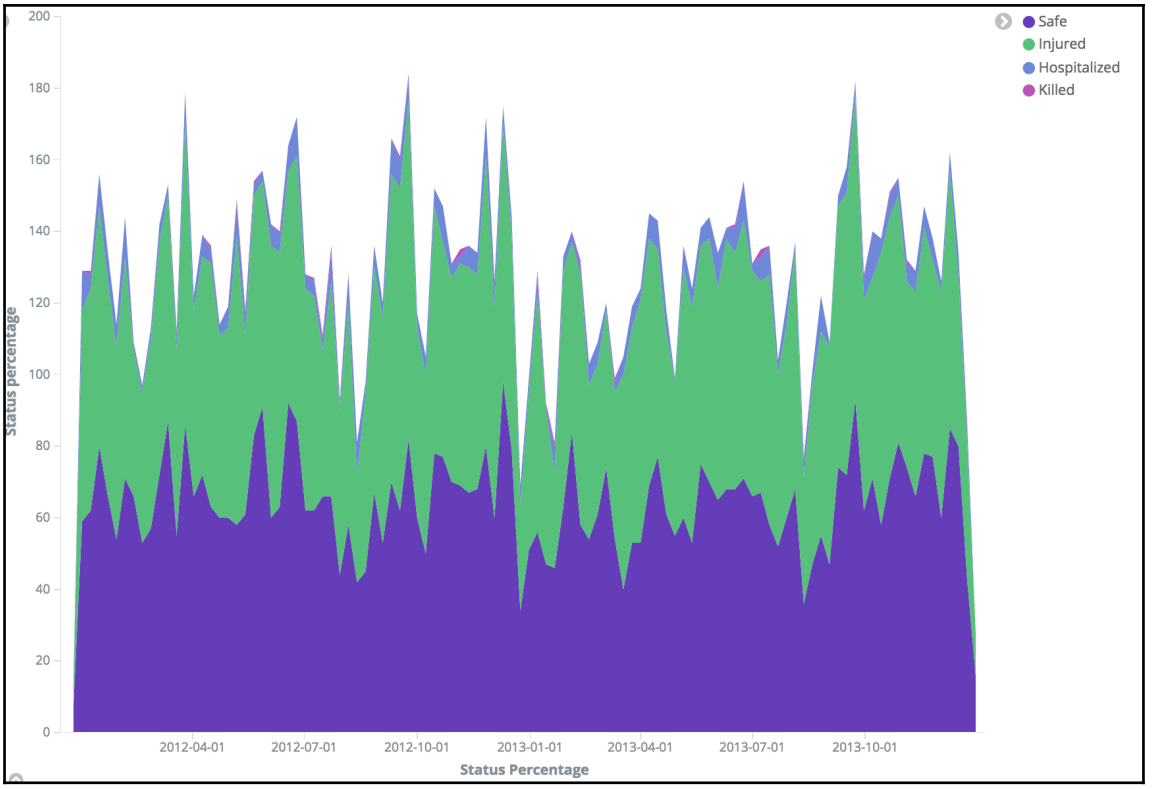
Safe
Injured
Hospitalized
Killed
Status percentage
200
180
160
140
120
100
80
60
40
20
0
2012-04-01
2012-07-01
2012-10-01
2013-01-01
2013-04-01
2013-07-01
2013-10-01
Status Percentage

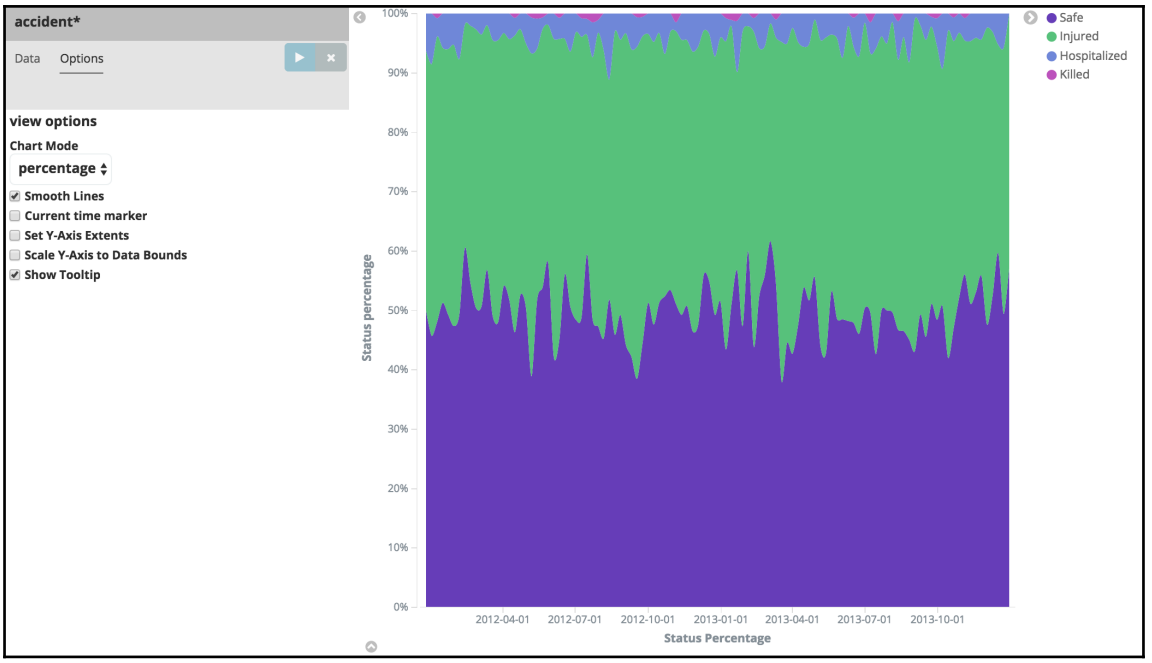
accident*
Data
Options
view options
Chart Mode
percentage
Smooth Lines
Current time marker
Set Y-Axis Extents
Scale Y-Axis to Data Bounds
Show Tooltip
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
Status percentage
2012-04-01
2012-07-01
2012-10-01
2013-01-01
2013-04-01
2013-07-01
2013-10-01
Status Percentage
Safe
Injured
Hospitalized
Killed

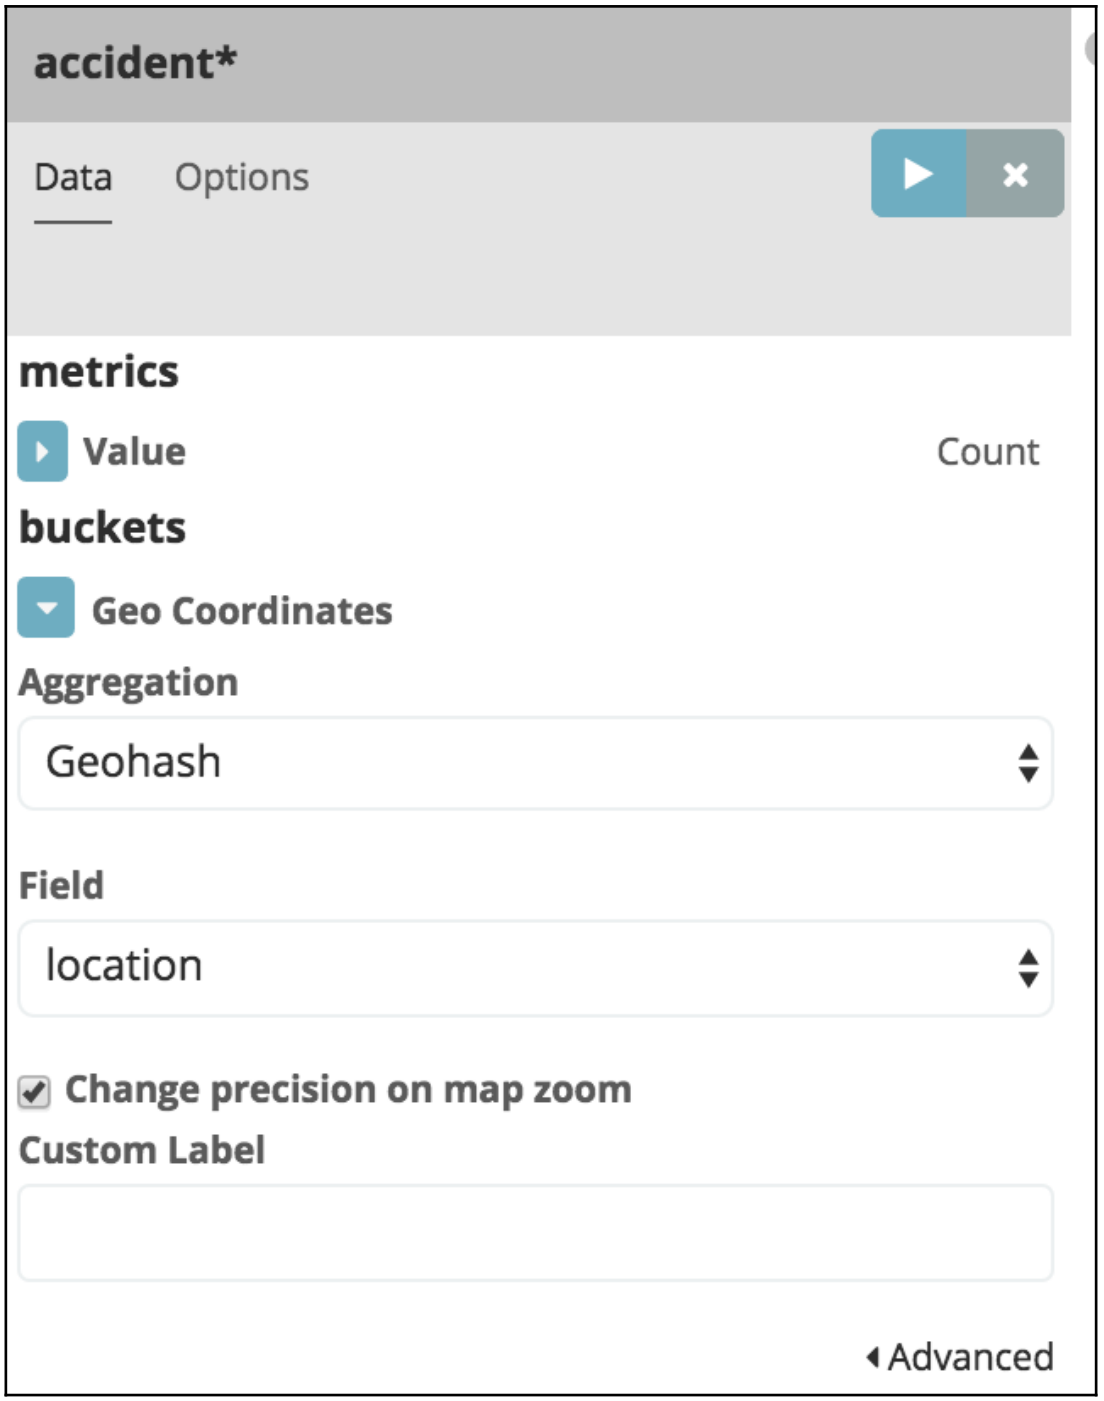
accident*
Data
Options
metrics
Value
Count
buckets
Geo Coordinates
Aggregation
Geohash
Field
location
Change precision on map zoom
Custom Label
Advanced

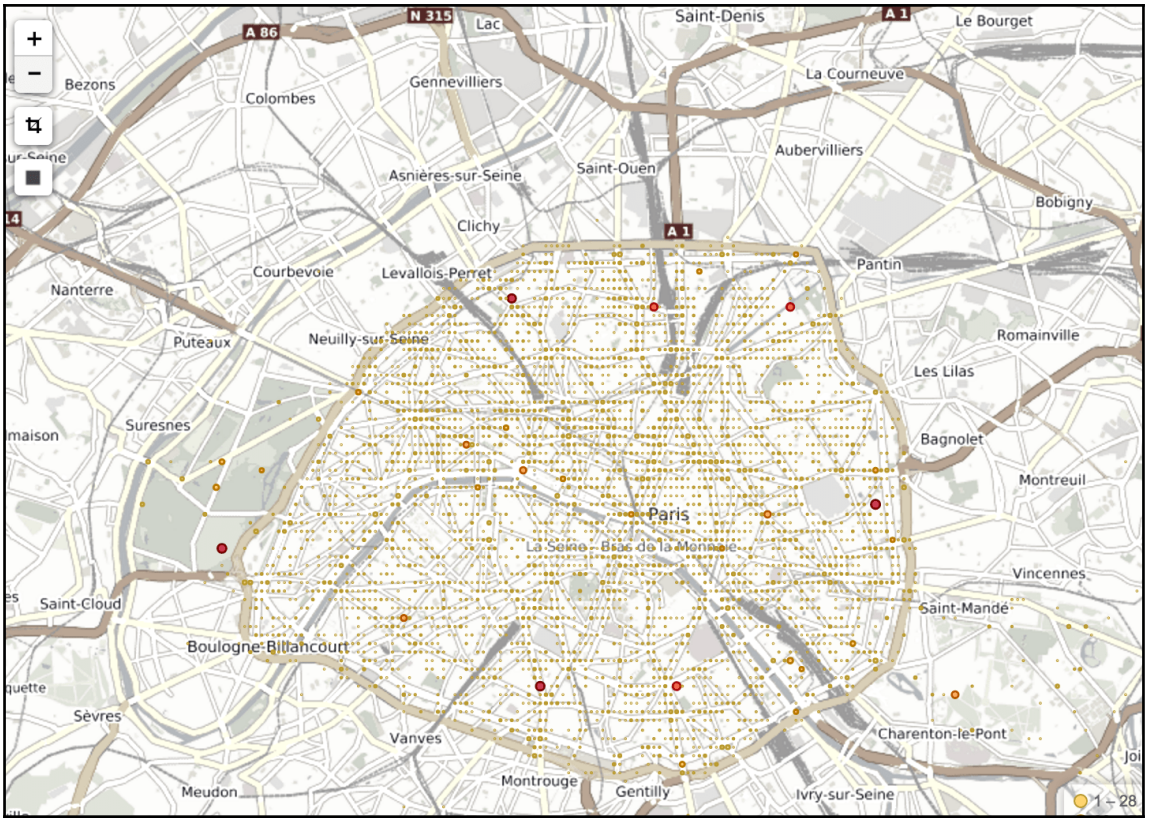
A 86
N 315
Lac
Saint-Denis
A 1
Le Bourget
Bezons
Colombes
Gennevilliers
La Courneuve
Aubervilliers
Asnières-sur-Seine
Saint-Ouen
Bobigny
Clichy
A 1
Courbevoie
Levallois-Perret
Pantin
Nanterre
Puteaux
Neuilly-sur-Seine
Romainville
Les Lilas
Suresnes
Bagnolet
Montreuil
Paris
La Seine - Bras de la Monnaie
Vincennes
Saint-Cloud
Saint-Mandé
Boulogne-Billancourt
Sèvres
Vanves
Charenton-le-Pont
Meudon
Montrouge
Gentilly
Ivry-sur-Seine
1 – 28

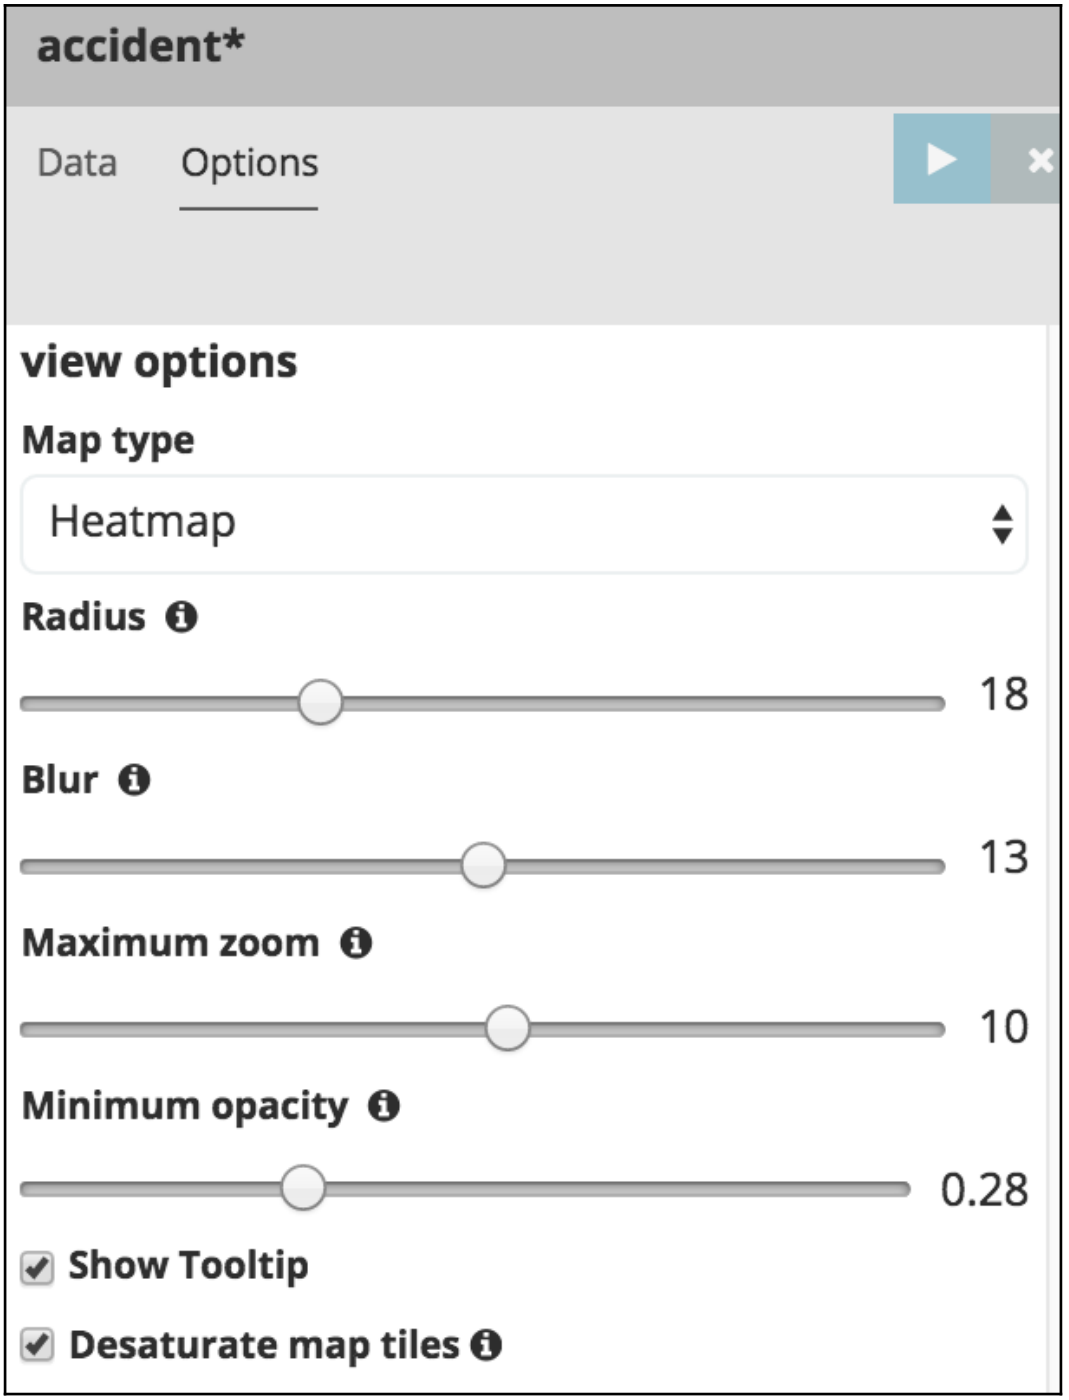
accident*
Data
Options
view options
Map type
Heatmap
Radius
18
Blur
13
Maximum zoom
10
Minimum opacity
0.28
Show Tooltip
Desaturate map tiles

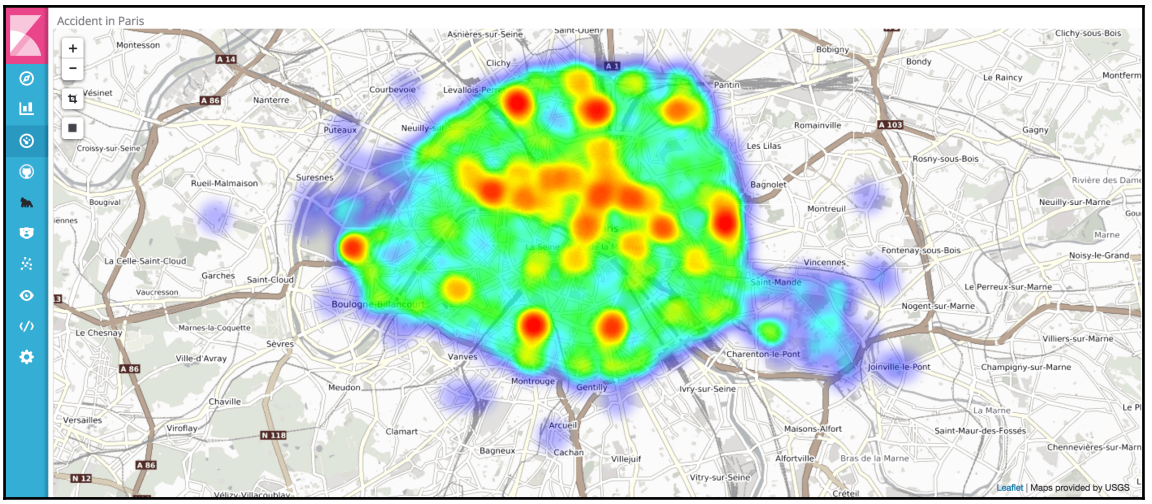
Accident in Paris
Montesson
Vésinet
Nanterre
Courbevoie
Levallois-Perret
Clichy
Asnières-sur-Seine
Saint-Ouen
Bobigny
Bondy
Le Raincy
Clichy-sous-Bois
Pantin
Puteaux
Neuilly-sur-Seine
Romainville
Gagny
Les Lilas
Rosny-sous-Bois
Croissy-sur-Seine
Rueil-Malmaison
Suresnes
Bagnolet
Montreuil
Bougival
Neuilly-sur-Marne
Marne
Noisy-le-Grand
La Celle-Saint-Cloud
Garches
Saint-Cloud
Vaucresson
Vincennes
Fontenay-sous-Bois
Saint-Mandé
Le Perreux-sur-Marne
Boulogne-Billancourt
Nogent-sur-Marne
Le Chesnay
Marnes-la-Coquette
Sèvres
Villiers-sur-Marne
Ville-d'Avray
Vanves
Charenton-le-Pont
Joinville-le-Pont
Champigny-sur-Marne
Meudon
Montrouge
Gentilly
Ivry-sur-Seine
Chaville
Versailles
Viroflay
Clamart
Arcueil
Maisons-Alfort
La Marne
Saint-Maur-des-Fossés
Bagneux
Cachan
Villejuif
Alfortville
Bras de la Marne
Chennevières-sur-Marn
Vitry-sur-Seine
Créteil
Leaflet | Maps provided by USGS


Accident over time
Count
@timestamp per week
2012-04-01
2012-07-01
2012-10-01
2013-01-01
2013-04-01
2013-07-01
Most dangerous streets
Count
ALLEE DE LA REIN...
ALLEE DE LONGCHAMP
QUAI DE BERCY
QUAI DES TUILERIES
ALLEE DU BORD DE...
QUAI DE LA RAPEE
VOIE GEORGES P...
PLACE DE LA CON...
PORTE D ITALIE
AVENUE DE L HIP...
Address: Descending
Accident by vehicule type
Car
Motor Scooter
Motorbike
Van
Bicycle
Transport bus
Truck
Bus
Tram
Person status after accident
Percentage of Count
Safe
Injured
Hospitalized
Killed
Percentage of @timestamp per week

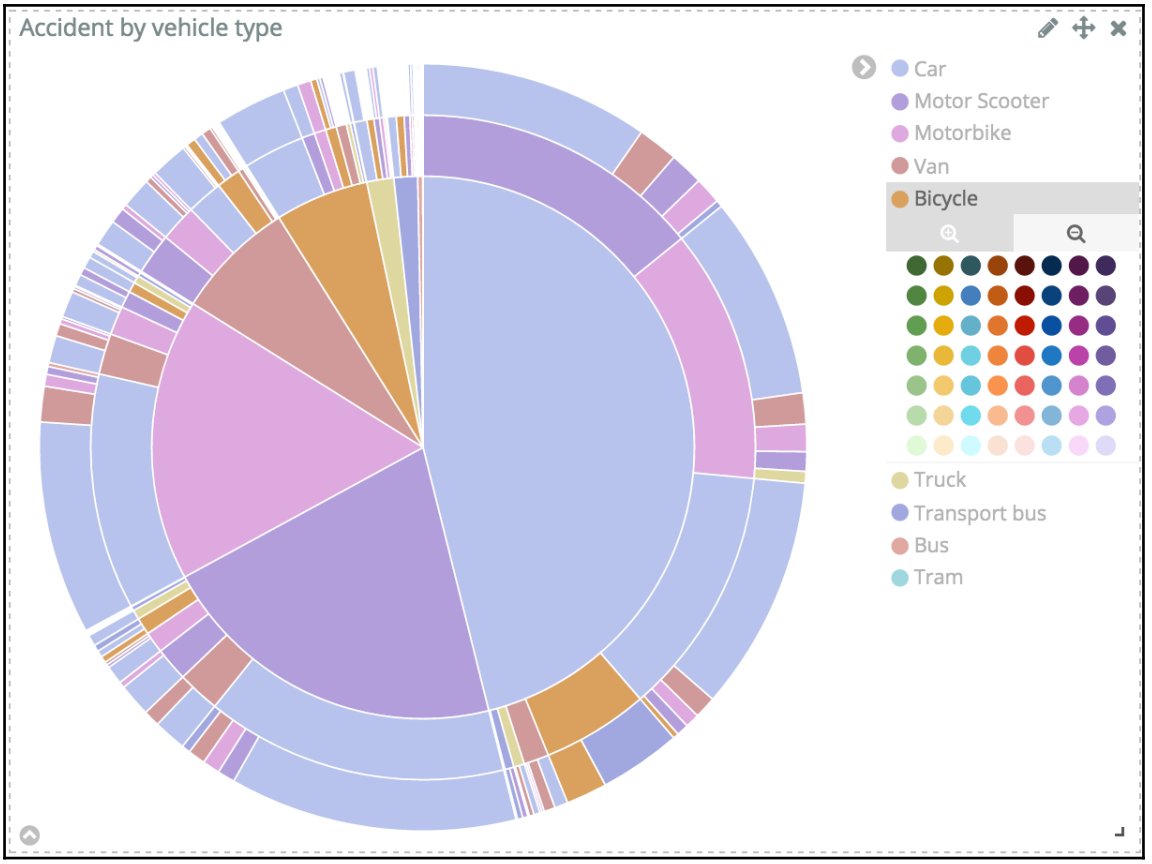
Accident by vehicle type
Car
Motor Scooter
Motorbike
Van
Bicycle
Truck
Transport bus
Bus
Tram


Vehicle 1 description.keyword: "Bicycle"

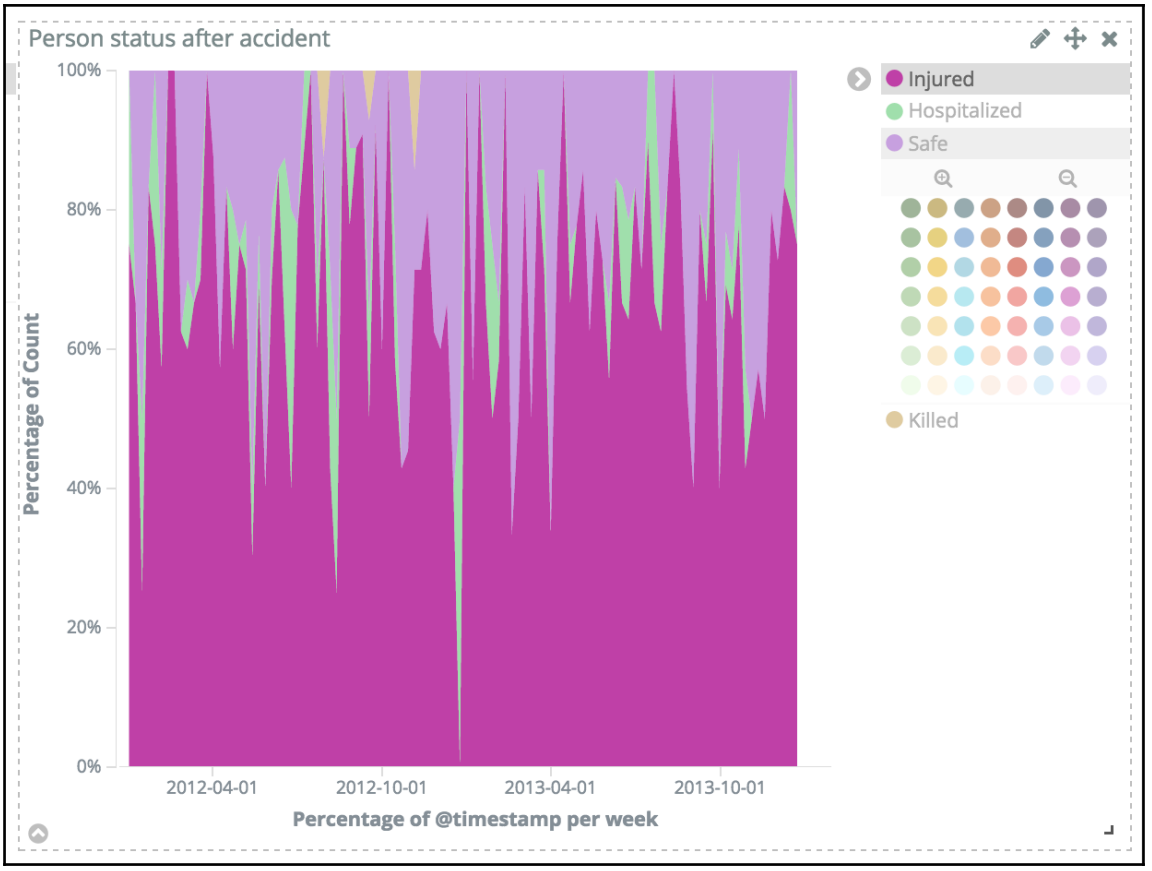
Person status after accident
100%
80%
60%
40%
20%
0%
Percentage of Count
2012-04-01
2012-10-01
2013-04-01
2013-10-01
Percentage of @timestamp per week
Injured
Hospitalized
Safe
Killed

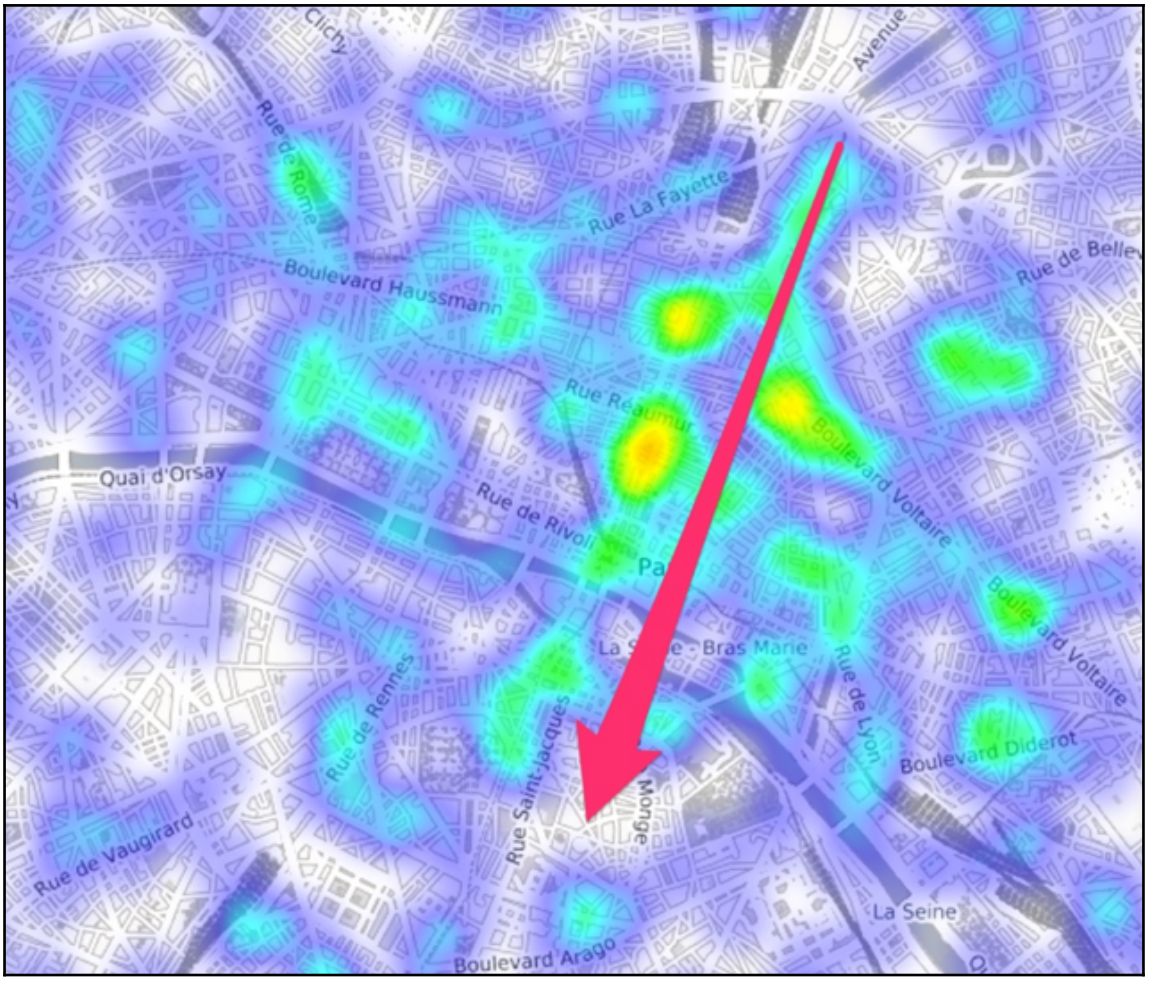
Clichy
Avenue
Rue de Rome
Rue La Fayette
Boulevard Haussmann
Rue de Belle
Rue Réaumur
Boulevard Voltaire
Quai d'Orsay
Rue de Rivoli
Boulevard Voltaire
Bras Marie
Rue de Rennes
Rue Saint-Jacques
Rue de Lyon
Boulevard Diderot
Monge
Rue de Vaugirard
La Seine
Boulevard Arago

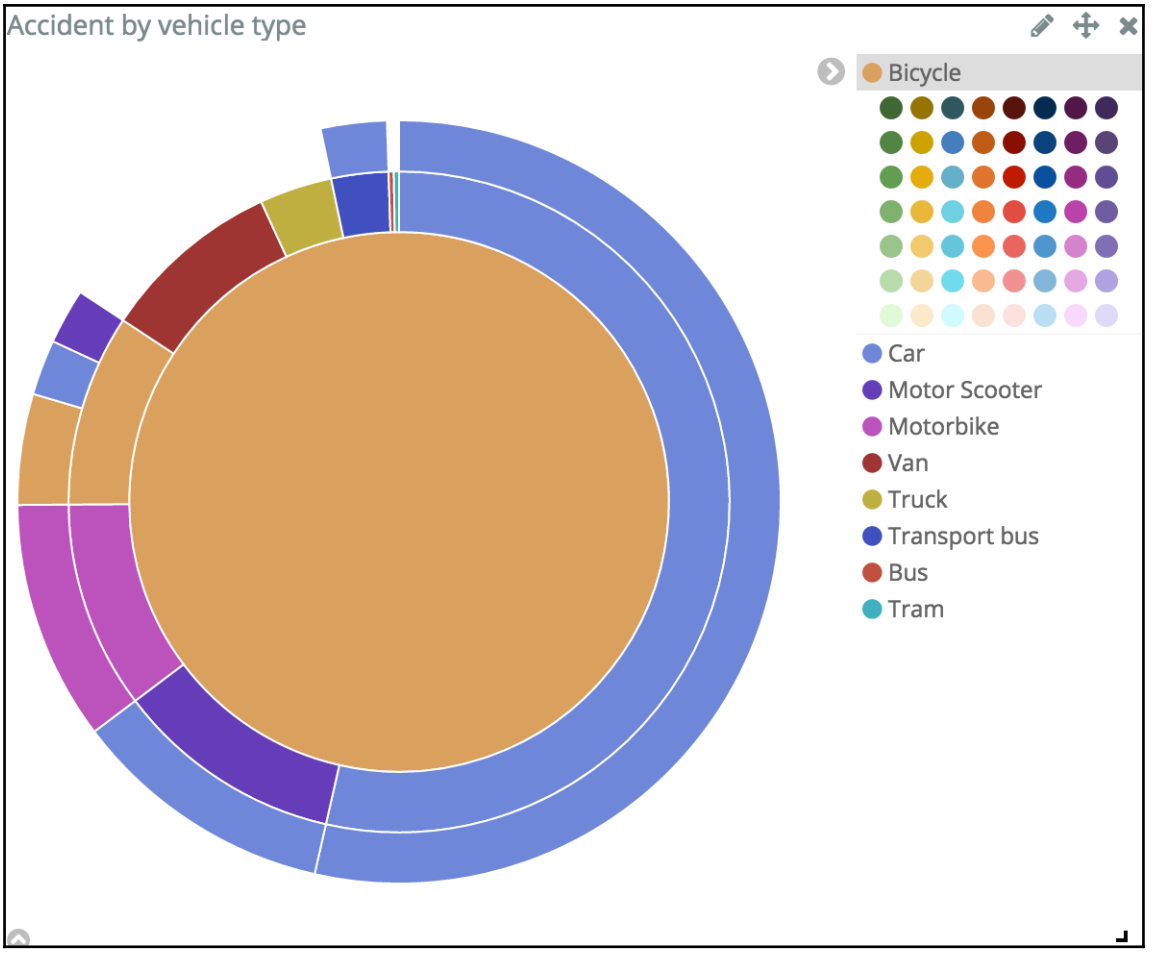
Accident by vehicle type
Bicycle
Car
Motor Scooter
Motorbike
Van
Truck
Transport bus
Bus
Tram


A 1
Courbevoie
Levallois-Perret
Pantin
Puteaux
Neuilly-sur-Seine
Romainville
A 103
Les Lilas
Suresnes
Bagnolet
Montreuil
Vincennes
Saint-Mandé
Boulogne-Billancourt
Sèvres
Vanves
Charenton-le-Pont
Meudon
Montrouge
Gentilly
Ivry-sur-Seine

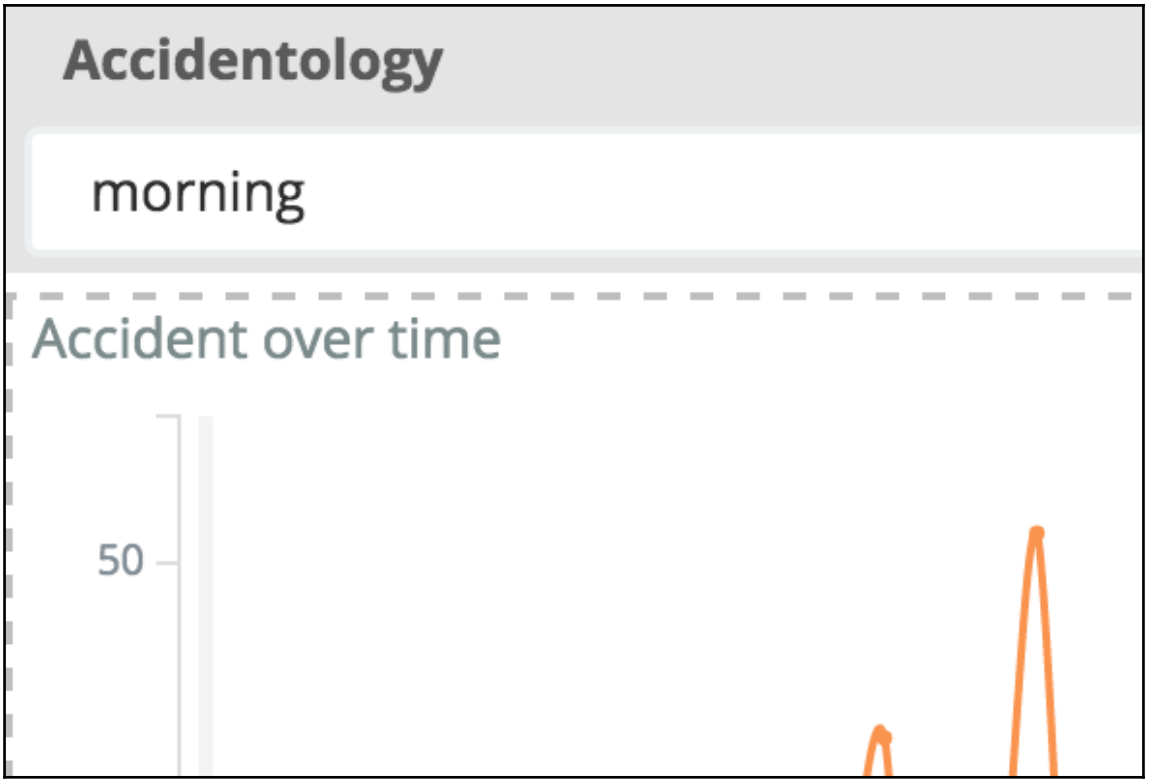
Accidentology
morning
Accident over time
50

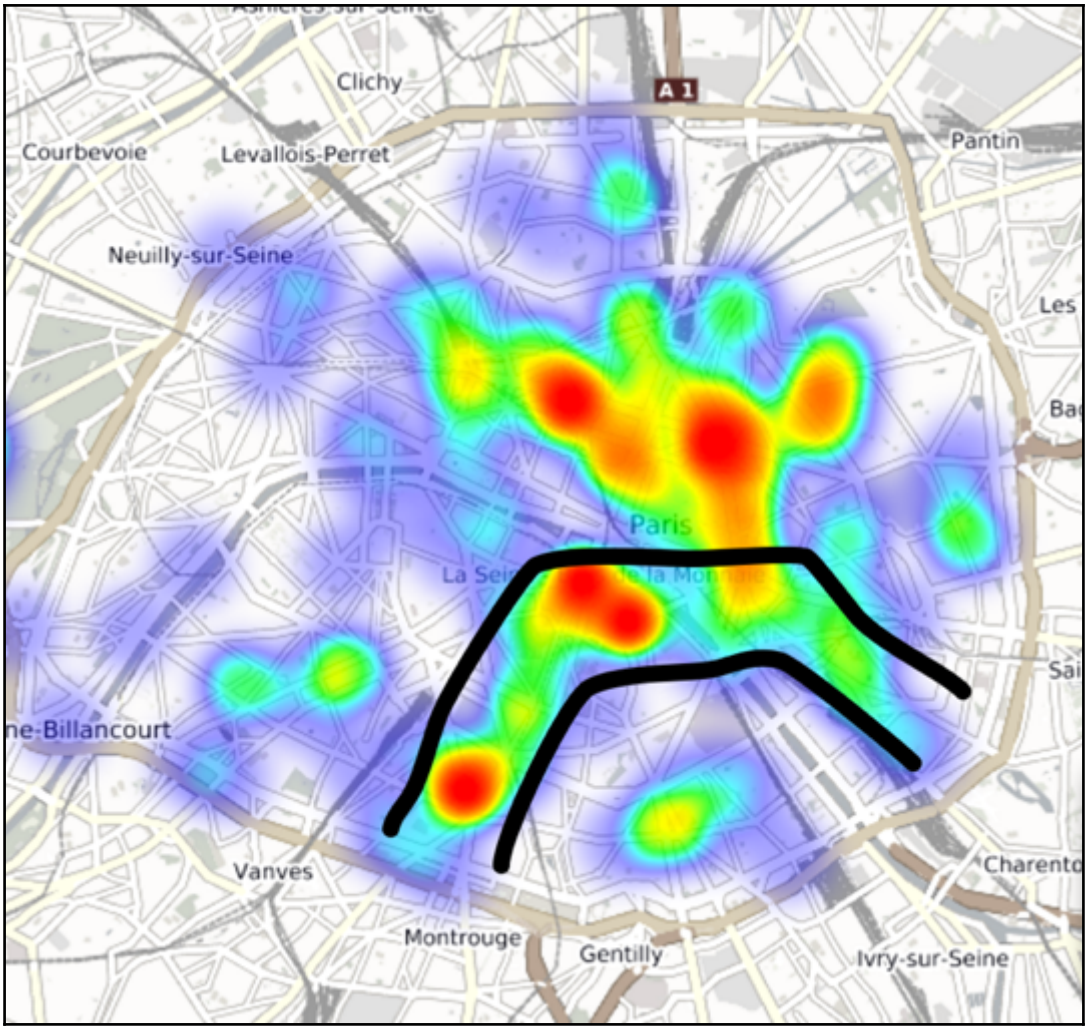
Clichy
A 1
Courbevoie
Levallois-Perret
Pantin
Neuilly-sur-Seine
Les
Paris
Vanves
Montrouge
Gentilly
Ivry-sur-Seine

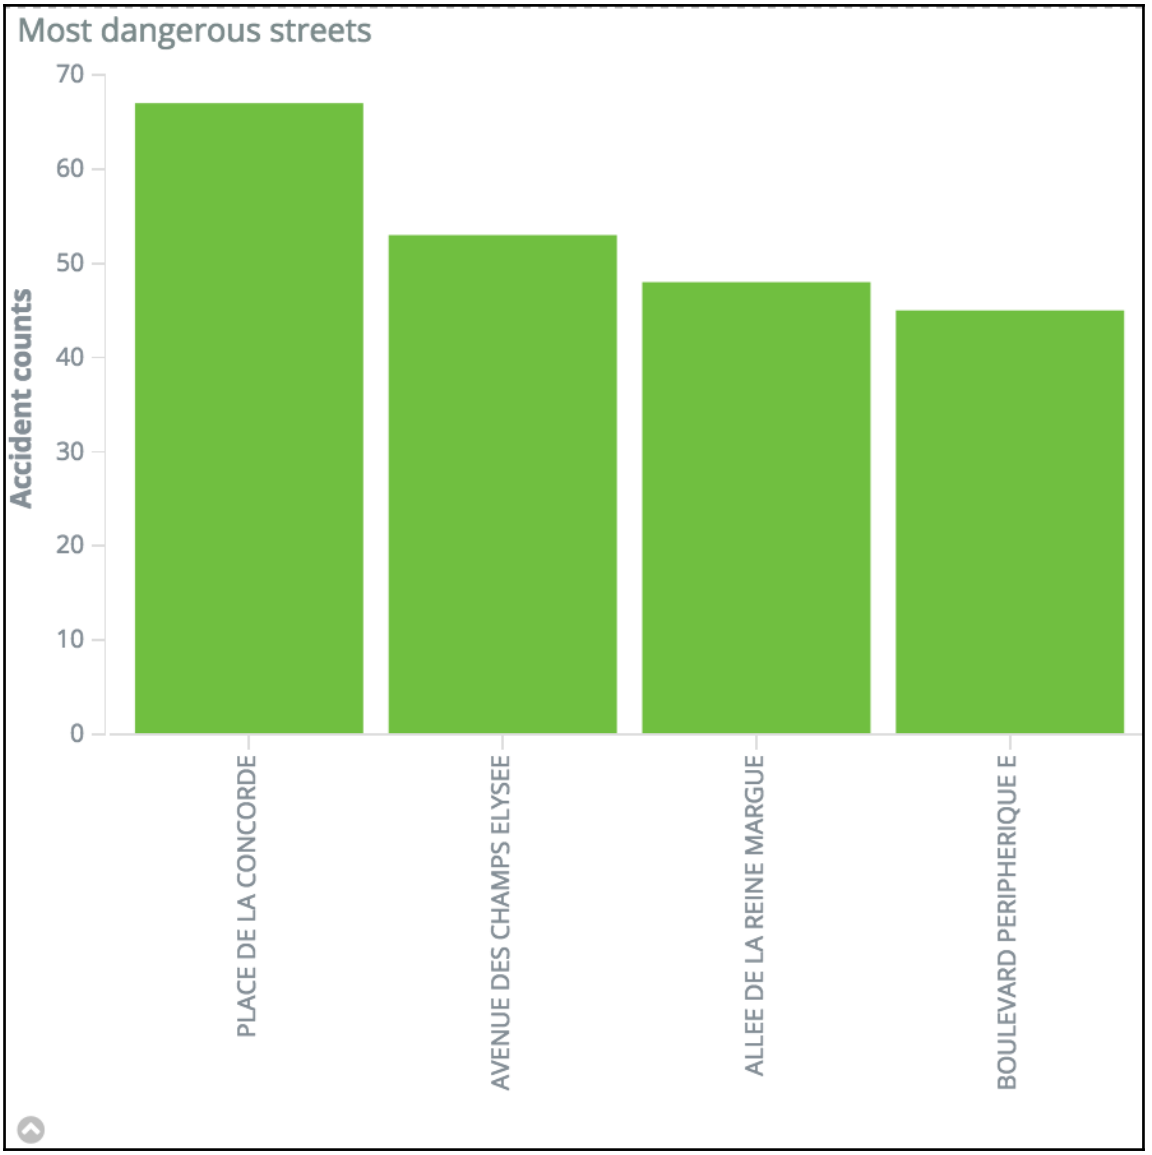
Most dangerous streets
Accident counts
70
60
50
40
30
20
10
0
PLACE DE LA CONCORDE
AVENUE DES CHAMPS ELYSEE
ALLEE DE LA REINE MARGUE
BOULEVARD PERIPHERIQUE E

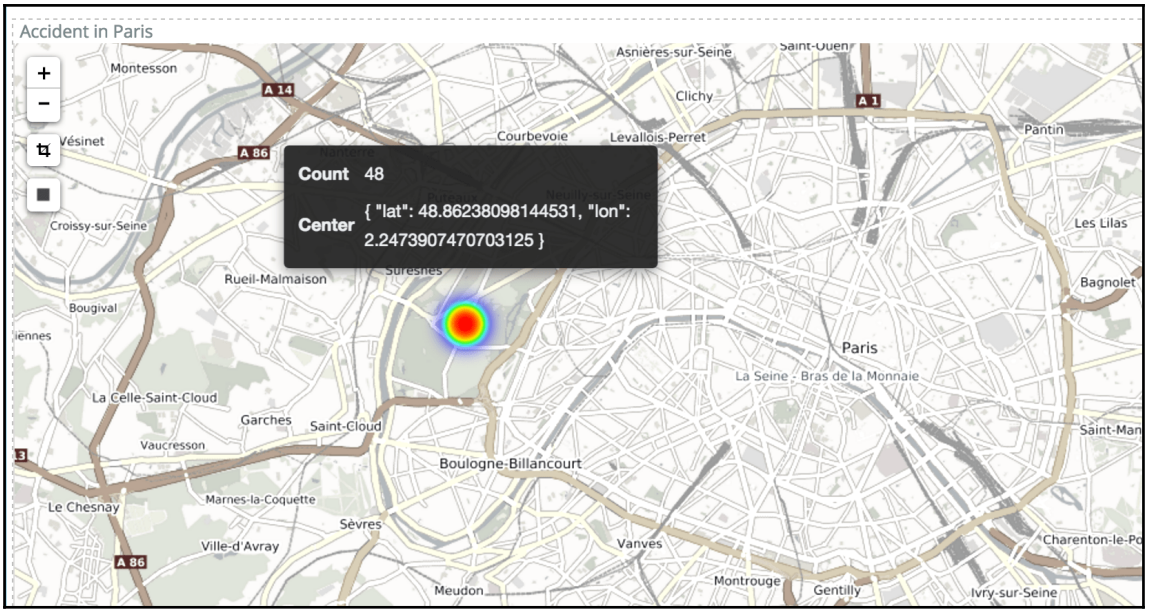
Accident in Paris
Count 48
Center { "lat": 48.86238098144531, "lon": 2.2473907470703125 }
Montesson
Vésinet
Croissy-sur-Seine
Bougival
Rueil-Malmaison
Suresnes
Courbevoie
Levallois-Perret
Asnières-sur-Seine
Clichy
Pantin
Les Lilas
Bagnolet
Paris
La Seine - Bras de la Monnaie
La Celle-Saint-Cloud
Garches
Saint-Cloud
Vaucresson
Boulogne-Billancourt
Le Chesnay
Marnes-la-Coquette
Sèvres
Ville-d'Avray
Vanves
Meudon
Montrouge
Gentilly
Ivry-sur-Seine
A 14
A 86
A 1

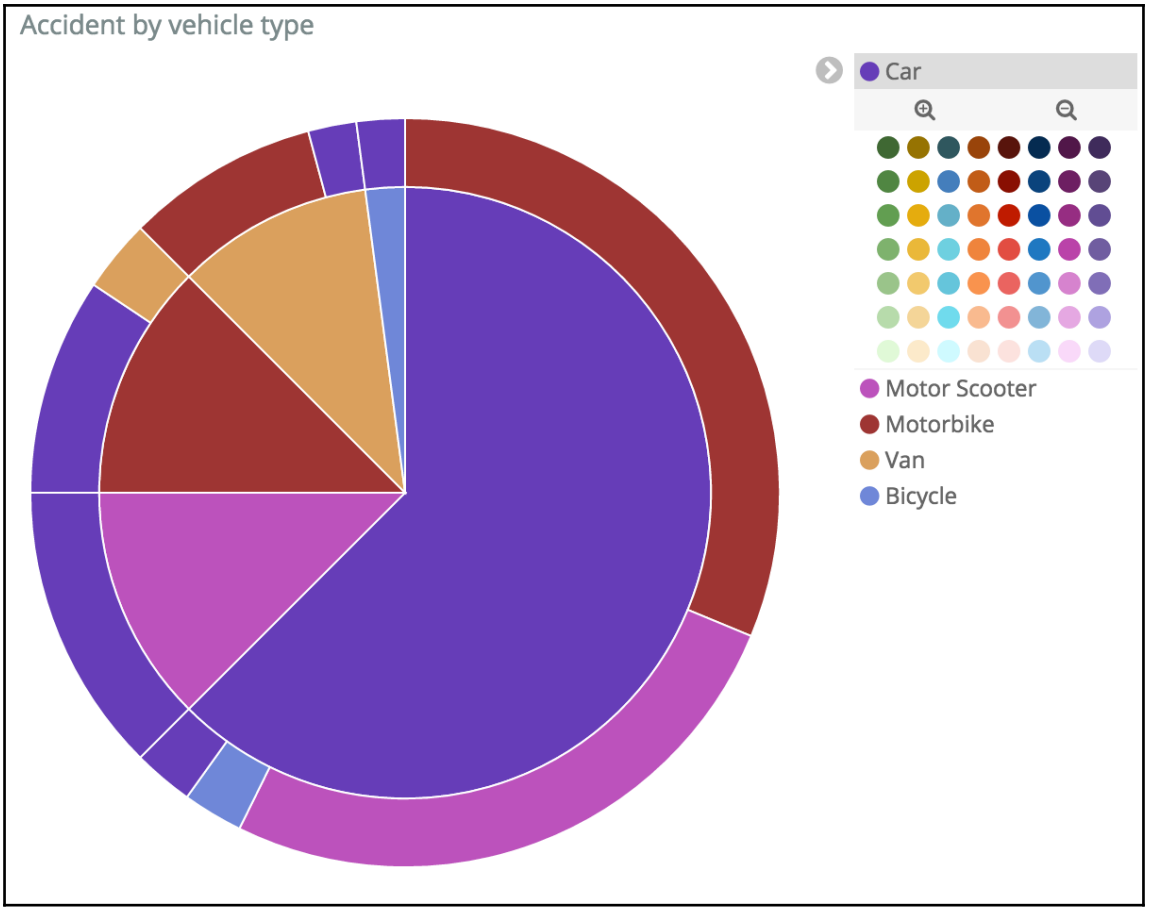
Accident by vehicle type
Car
Motor Scooter
Motorbike
Van
Bicycle

# Chapter 4: Logging Analytics with Kibana 5.0

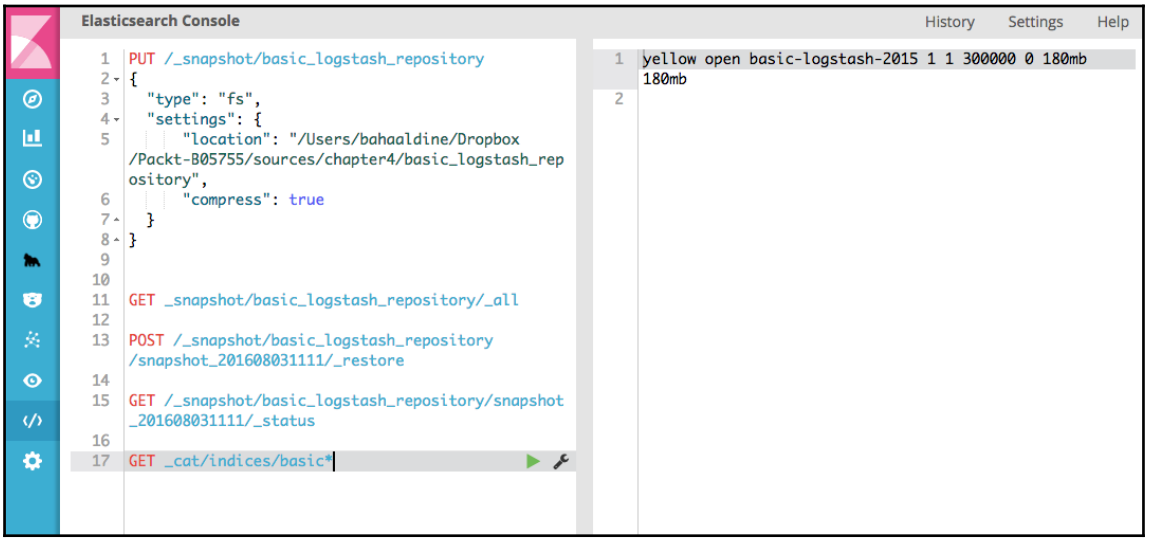

Management / Kibana

Index Patterns | Saved Objects | Reporting | Advanced Settings

# Edit Saved Objects

Export Everything | Import

From here you can delete saved objects, such as saved searches. You can also edit the raw data of saved objects. Typically objects are only modified via their associated application, which is probably what you should use instead of this screen. Each tab is limited to 100 results. You can use the filter to find objects not in the default list.

Basic

Dashboards (0) | Searches (0) | Visualizations (13)

Select All | Delete | Export

- Basic Elastic Summary
- Basic ELK - Summary
- Basic Logstash - Summary
- Basic Logstash - US vs CA
- Basic Logstash 404 vs 404 mavg(5)
- BasicLogstash - Bandwidth by Country
- BasicLogstash - HeatMap
- BasicLogstash - Map
- BasicLogstash - Metrics
- BasicLogstash - Requests by Agent
- BasicLogstash - Response Codes over Time
- BasicLogstash - Significant Countries by Response Code
- BasicLogstash - Top Requested Resources

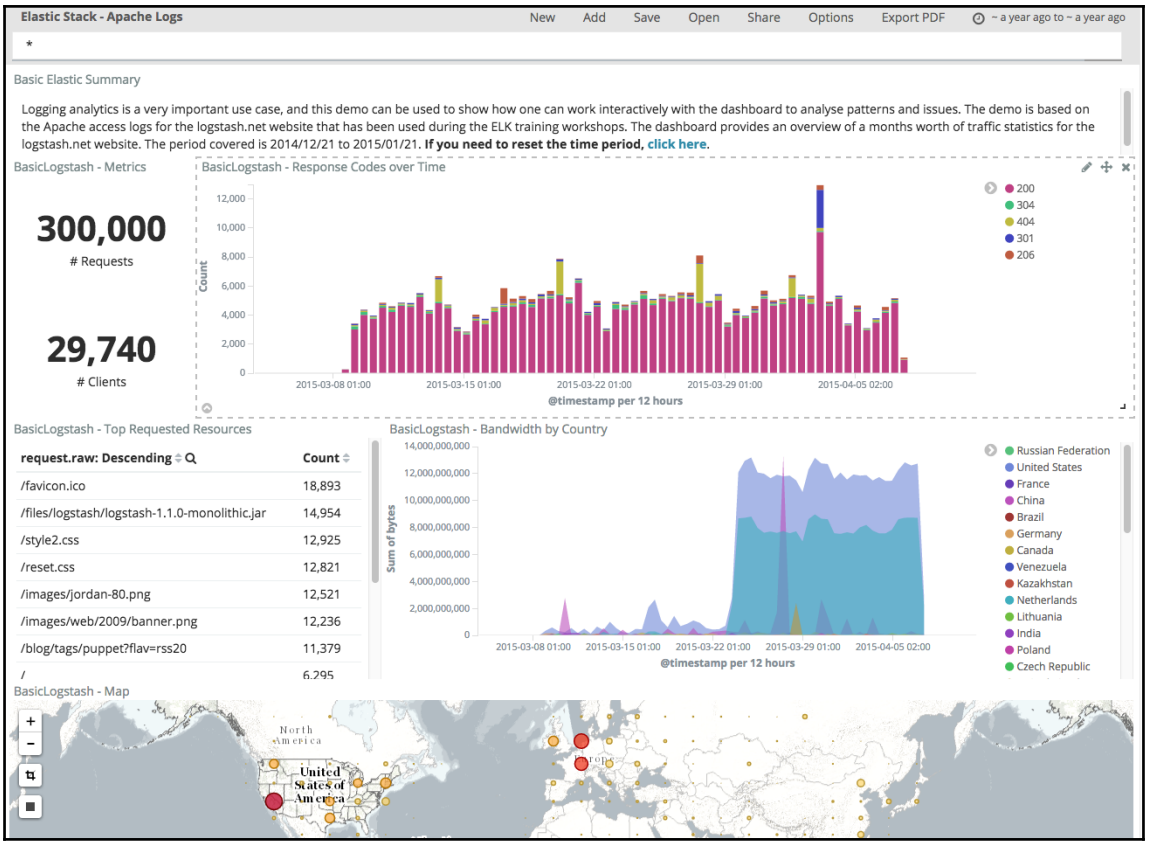


Basic Elastic Summary

Logging analytics is a very important use case, and this demo can be used to show how one can work interactively with the dashboard to analyse patterns and issues. The demo is based on the Apache access logs for the logstash.net website. The dashboard provides an overview of a months worth of traffic statistics for the logstash.net website. The period covered is 2014/12/21 to 2015/01/21. **If you need to reset the time period, click here**.

BasicLogstash - Metrics

**300,000**

# Requests

**29,740**

# Clients

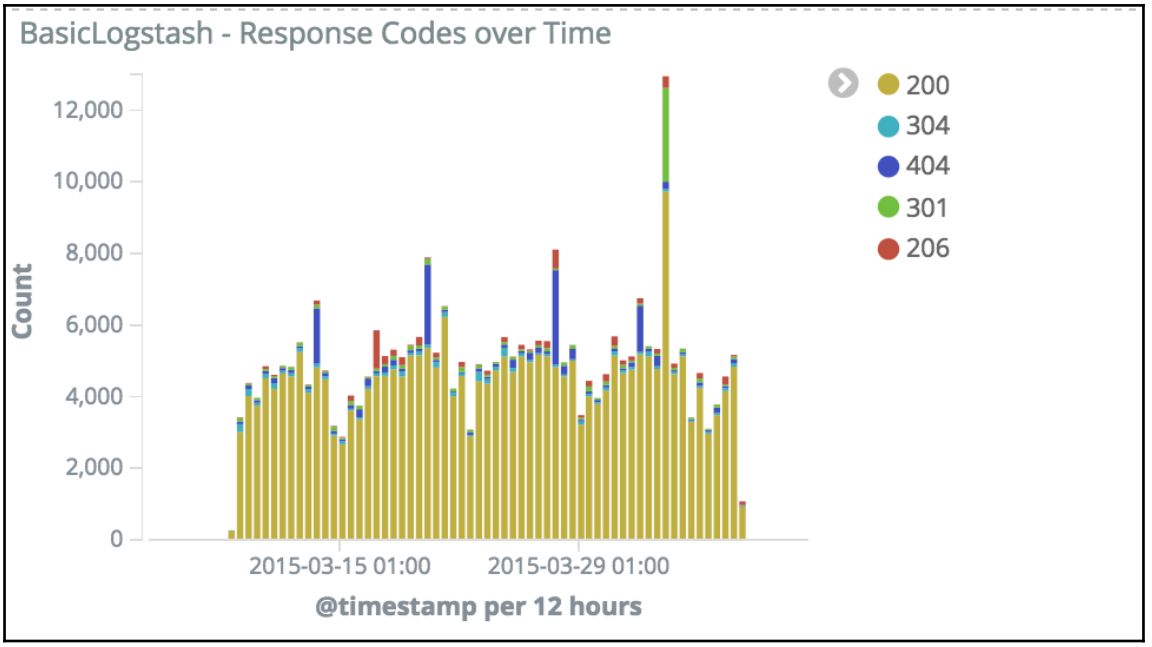
BasicLogstash - Response Codes over Time
Count
12,000
10,000
8,000
6,000
4,000
2,000
0
2015-03-15 01:00
2015-03-29 01:00
@timestamp per 12 hours
200
304
404
301
206

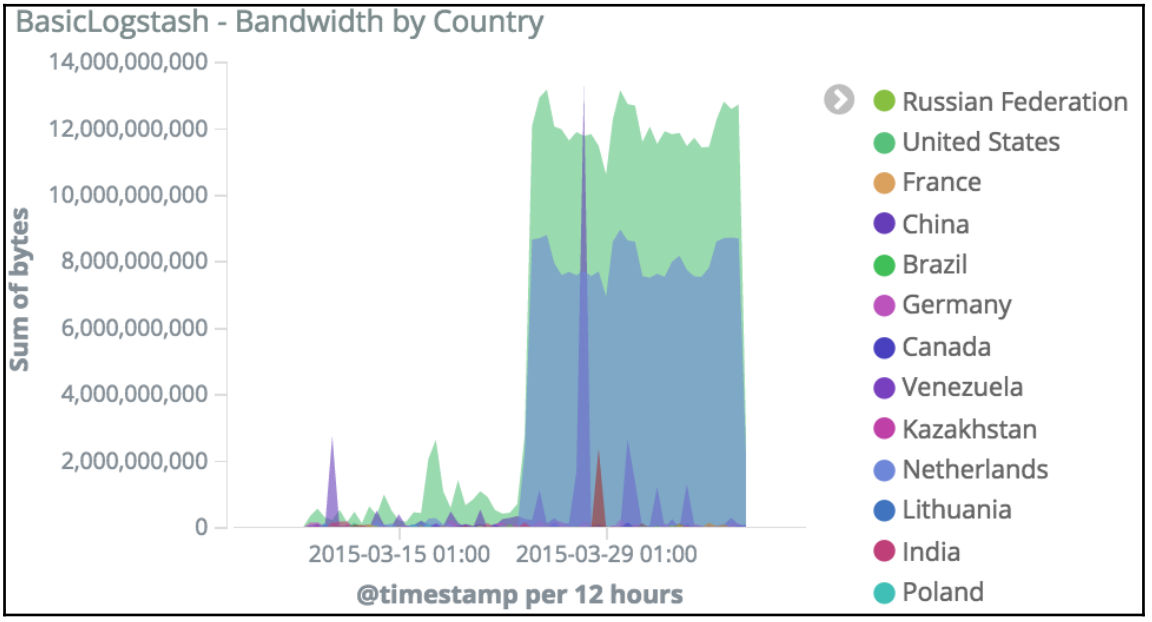
BasicLogstash - Bandwidth by Country
Sum of bytes
14,000,000,000
12,000,000,000
10,000,000,000
8,000,000,000
6,000,000,000
4,000,000,000
2,000,000,000
0
2015-03-15 01:00
2015-03-29 01:00
@timestamp per 12 hours
Russian Federation
United States
France
China
Brazil
Germany
Canada
Venezuela
Kazakhstan
Netherlands
Lithuania
India
Poland

BasicLogstash - Requests by Agent

| agent.raw: Descending | Count |
|---|---|
| "Chef Client/10.18.2 (ruby-1.8.7-p302; ohai-6.14.0; x86_64-linux; +http://opscode.com)" | 14,072 |
| "-" | 11,092 |
| "Mozilla/5.0 (Windows NT 6.1; WOW64; rv:27.0) Gecko/20100101 Firefox/27.0" | 10,151 |
| "Mozilla/5.0 (X11; Ubuntu; Linux x86_64; rv:27.0) Gecko/20100101 Firefox/27.0" | 8,773 |
| "UniversalFeedParser/4.2-pre-314-svn +http://feedparser.org/" | 8,529 |

## BasicLogstash - Top Requested Resources

| | |
|---|---|
| /images/jordan-80.png | 12,521 |
| /images/web/2009/banner.png | 12,236 |
| /blog/tags/puppet?flav=rss20 | 11,379 |
| / | 6,295 |
| /presentations/fpm-scale12x.pdf | 5,327 |
| /?flav=rss20 | 5,103 |

Export: Raw Formatted

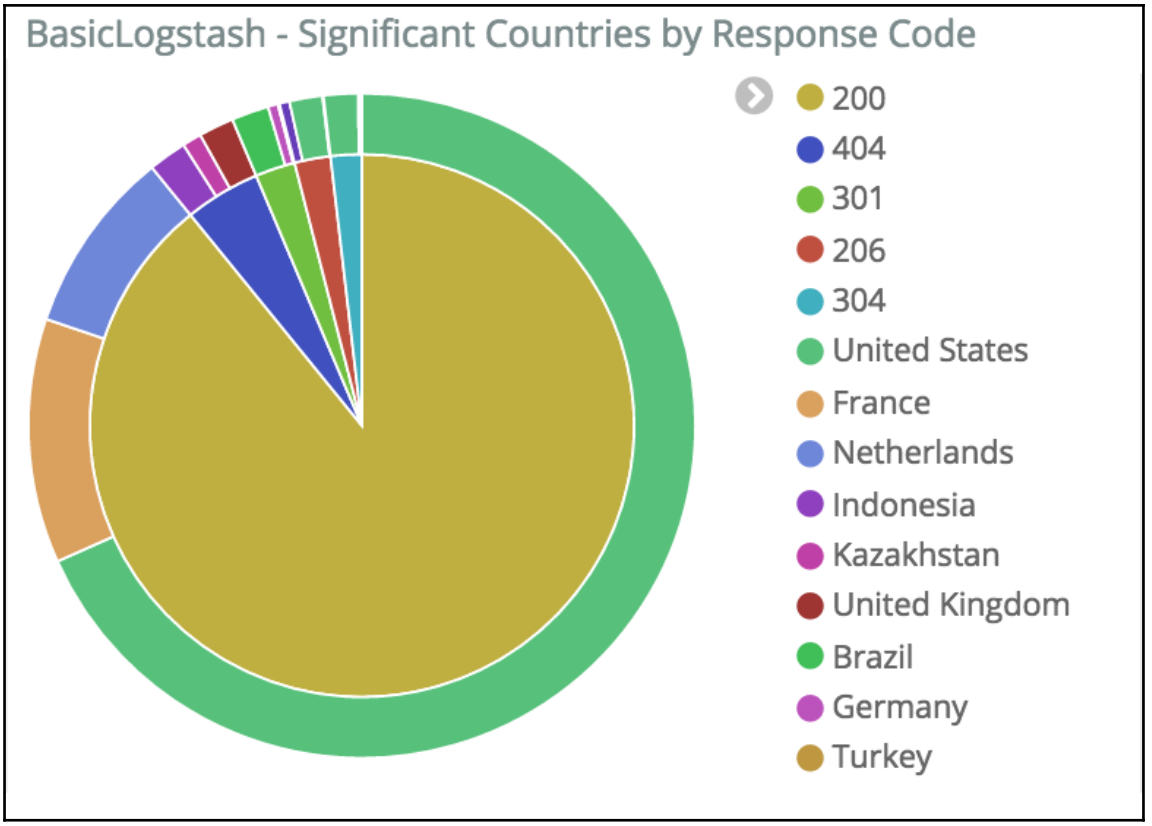
BasicLogstash - Significant Countries by Response Code
200
404
301
206
304
United States
France
Netherlands
Indonesia
Kazakhstan
United Kingdom
Brazil
Germany
Turkey

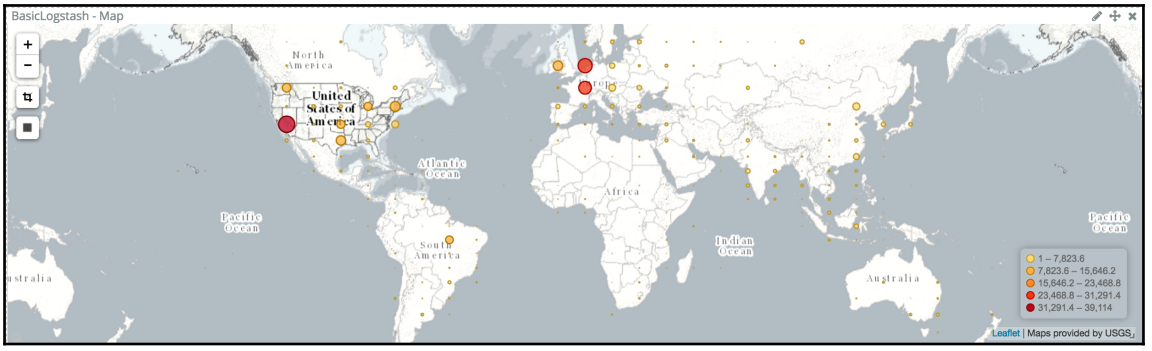
BasicLogstash - Map
North America
United States of America
Atlantic Ocean
Pacific Ocean
South America
Africa
Indian Ocean
Australia
1 – 7,823.6
7,823.6 – 15,646.2
15,646.2 – 23,468.8
23,468.8 – 31,291.4
31,291.4 – 39,114
Leaflet | Maps provided by USGS

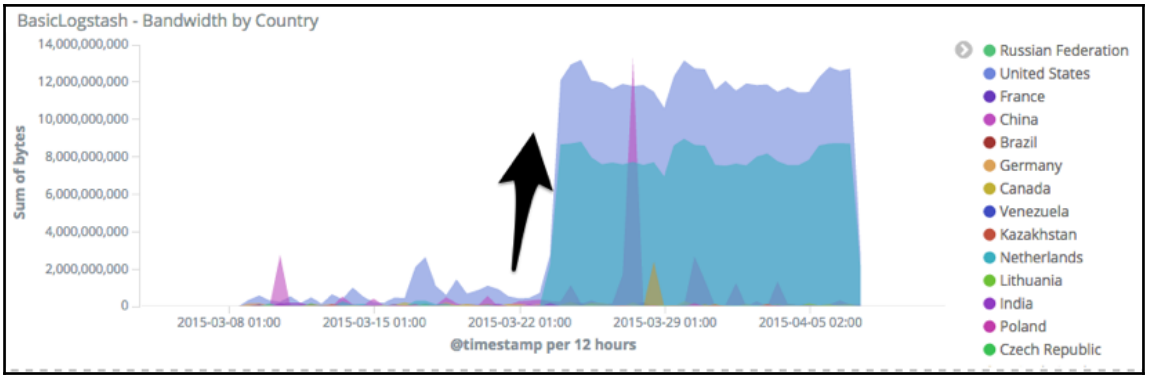
BasicLogstash - Bandwidth by Country
Sum of bytes
14,000,000,000
12,000,000,000
10,000,000,000
8,000,000,000
6,000,000,000
4,000,000,000
2,000,000,000
0
2015-03-08 01:00
2015-03-15 01:00
2015-03-22 01:00
2015-03-29 01:00
2015-04-05 02:00
@timestamp per 12 hours
Russian Federation
United States
France
China
Brazil
Germany
Canada
Venezuela
Kazakhstan
Netherlands
Lithuania
India
Poland
Czech Republic

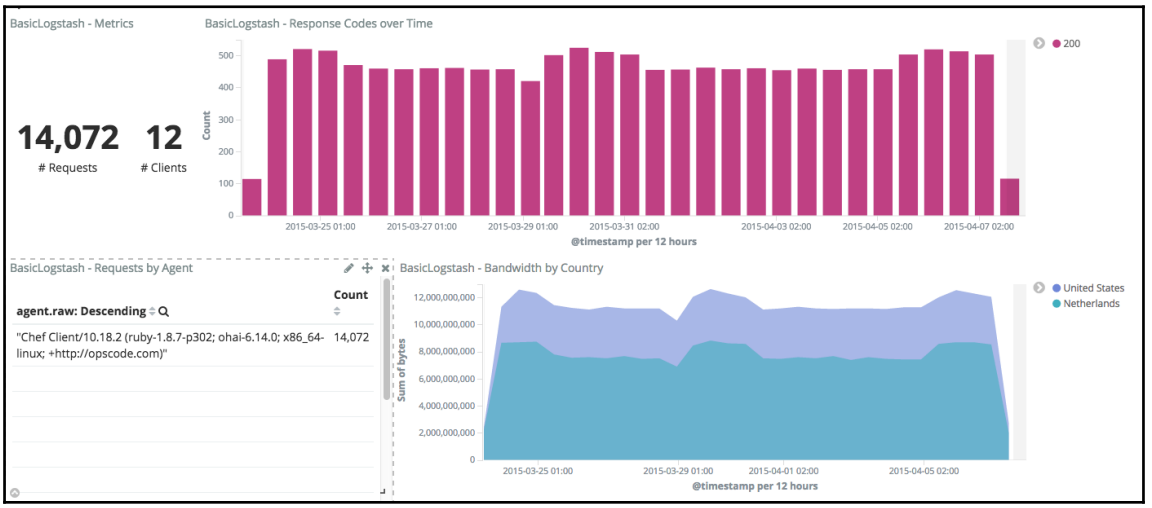
BasicLogstash - Metrics
14,072
# Requests
12
# Clients
BasicLogstash - Response Codes over Time
Count
500
400
300
200
100
0
2015-03-25 01:00
2015-03-27 01:00
2015-03-29 01:00
2015-03-31 02:00
2015-04-03 02:00
2015-04-05 02:00
2015-04-07 02:00
@timestamp per 12 hours
200
BasicLogstash - Requests by Agent
agent.raw: Descending
Count
"Chef Client/10.18.2 (ruby-1.8.7-p302; ohai-6.14.0; x86_64-linux; +http://opscode.com)"
14,072
BasicLogstash - Bandwidth by Country
Sum of bytes
12,000,000,000
10,000,000,000
8,000,000,000
6,000,000,000
4,000,000,000
2,000,000,000
0
2015-03-25 01:00
2015-03-29 01:00
2015-04-01 02:00
2015-04-05 02:00
@timestamp per 12 hours
United States
Netherlands

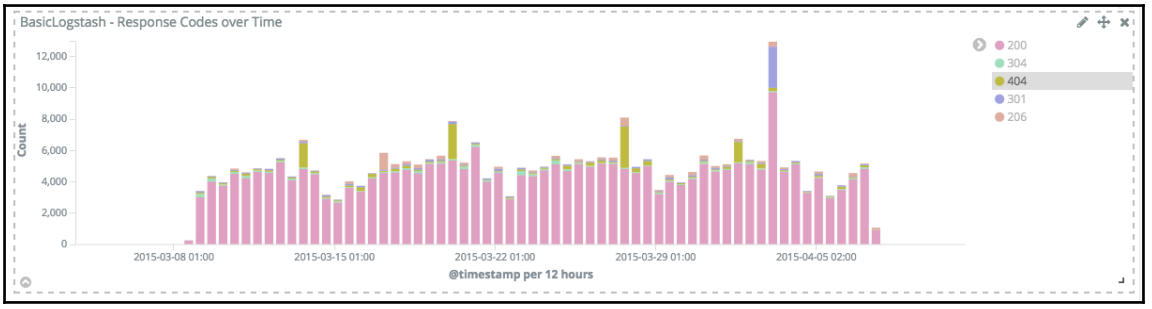
BasicLogstash - Response Codes over Time
Count
12,000
10,000
8,000
6,000
4,000
2,000
0
2015-03-08 01:00
2015-03-15 01:00
2015-03-22 01:00
2015-03-29 01:00
2015-04-05 02:00
@timestamp per 12 hours
200
304
404
301
206

BasicLogstash - Top Requested Resources

| request.raw: Descending | Count |
| --- | --- |
| /wp/wp-admin/ | 79 |
| /blog/wp-admin/ | 76 |
| /wordpress/wp-admin/ | 73 |
| /wp-admin/ | 71 |

# Chapter 5: Metric Analytics with Metricbeat and Kibana 5.0

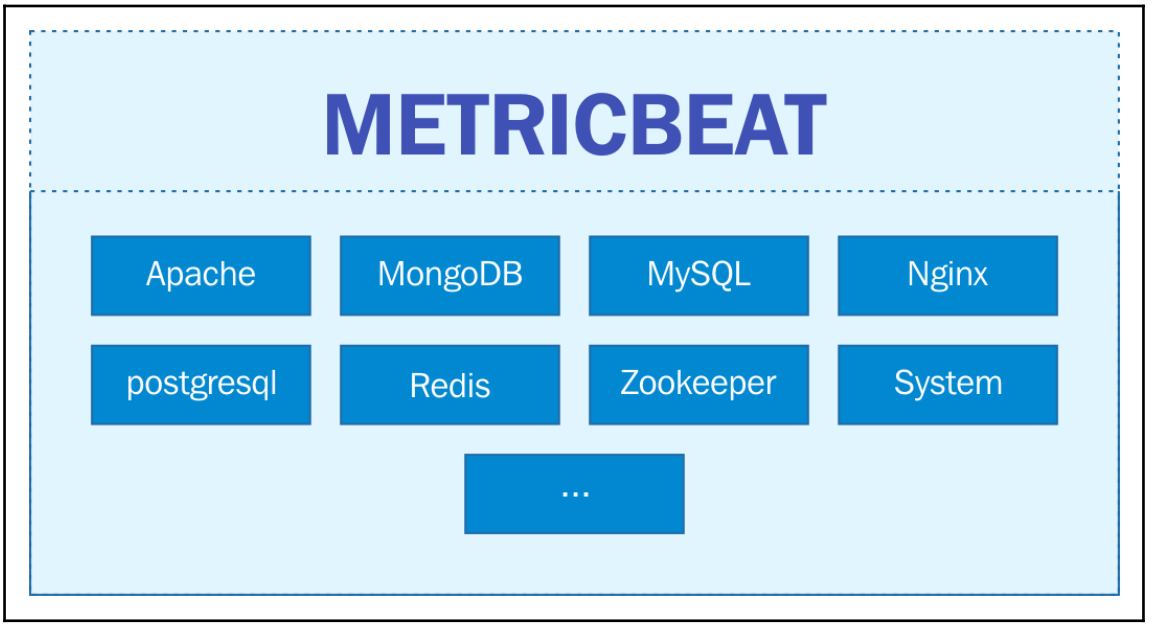


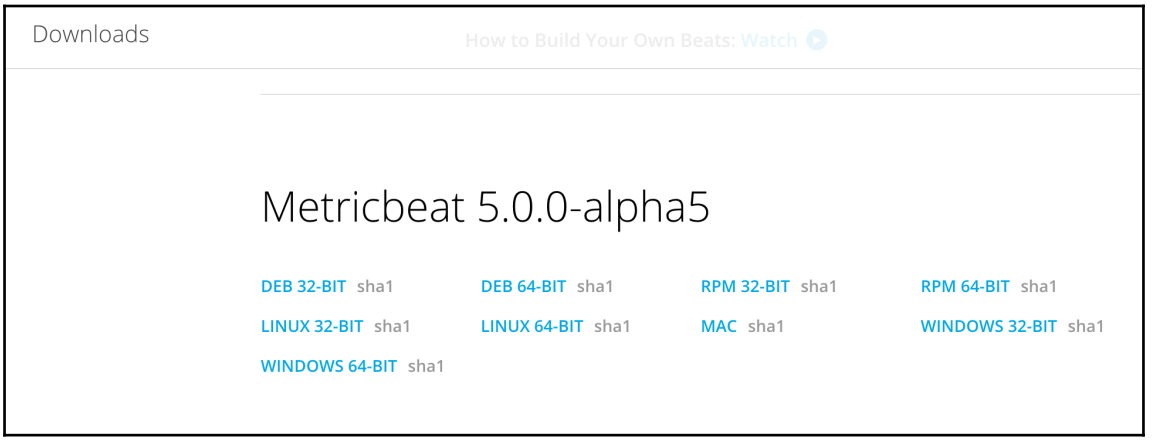

Management

Version: 5.0.0

Elasticsearch

Users Roles

Kibana

Index Patterns Saved Objects Reporting Advanced Settings

Management / Elasticsearch

Users Roles

Filter...

Delete New Role

| Role | Description | |
|---|---|---|
| kibana_user | | Reserved |
| superuser | | Reserved |
| transport_client | | Reserved |

Management / Elasticsearch / Roles / Edit

Users Roles

Edit Role

Return to All Roles Delete Save

**Name**

metricbeat_user

Cluster Privileges

☑ **all** ☐ **monitor** ☐ **manage** ☐ **manage_security** ☐ **manage_index_templates**

Run As Privileges

Add a user...

Indices Privileges

metricbeat-* ✕ | all ✕ | +

Add a query... | * ✕ | 🗑

Management / Elasticsearch / Users

Users Roles

Return to All Users Save

New User

**Username**

metricbeat

**Password**

••••••

••••••

Change Password

**Full Name**

Metricbeat

**Email**

metricbeat@beat.go

Roles

metricbeat_user

3,001 hits New Save Open Share Export PDF Last 15 minutes

*

metricbeat-*

Selected Fields

? _source

Available Fields

September 6th 2016, 18:18:57.043 - September 6th 2016, 18:33:57.043 — by 30 seconds

Count

800 600 400 200 0

18:20:00 18:21:00 18:22:00 18:23:00 18:24:00 18:25:00 18:26:00 18:27:00 18:28:00 18:29:00 18:30:00 18:31:00 18:32:00 18:33:00

@timestamp per 30 seconds

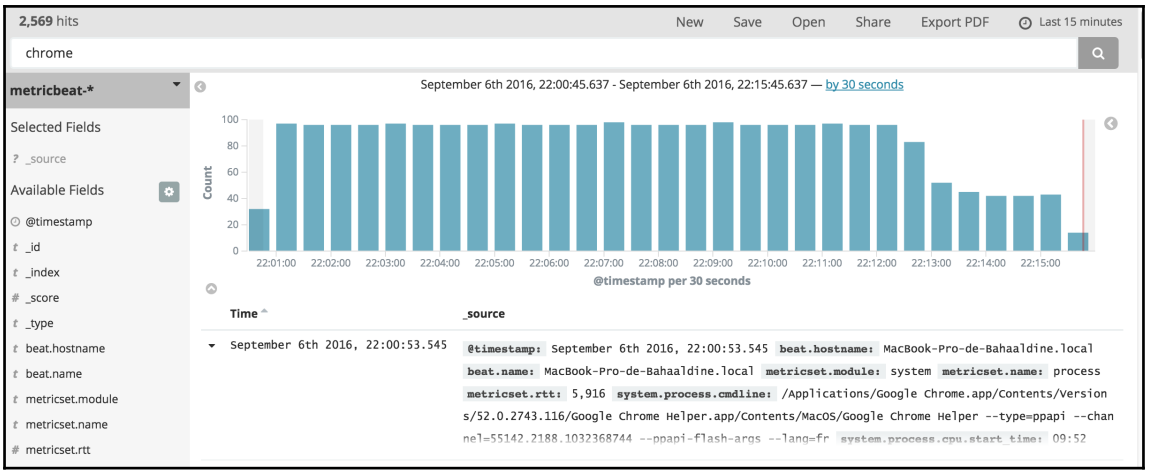
2,569 hits
New
Save
Open
Share
Export PDF
Last 15 minutes
chrome
metricbeat-*
September 6th 2016, 22:00:45.637 - September 6th 2016, 22:15:45.637 — by 30 seconds
Selected Fields
? _source
Available Fields
@timestamp
t _id
t _index
# _score
t _type
t beat.hostname
t beat.name
t metricset.module
t metricset.name
# metricset.rtt
Count
100
80
60
40
20
0
22:01:00
22:02:00
22:03:00
22:04:00
22:05:00
22:06:00
22:07:00
22:08:00
22:09:00
22:10:00
22:11:00
22:12:00
22:13:00
22:14:00
22:15:00
@timestamp per 30 seconds
Time
_source
September 6th 2016, 22:00:53.545
@timestamp: September 6th 2016, 22:00:53.545 beat.hostname: MacBook-Pro-de-Bahaaldine.local
beat.name: MacBook-Pro-de-Bahaaldine.local metricset.module: system metricset.name: process
metricset.rtt: 5,916 system.process.cmdline: /Applications/Google Chrome.app/Contents/Version
s/52.0.2743.116/Google Chrome Helper.app/Contents/MacOS/Google Chrome Helper --type=ppapi --chan

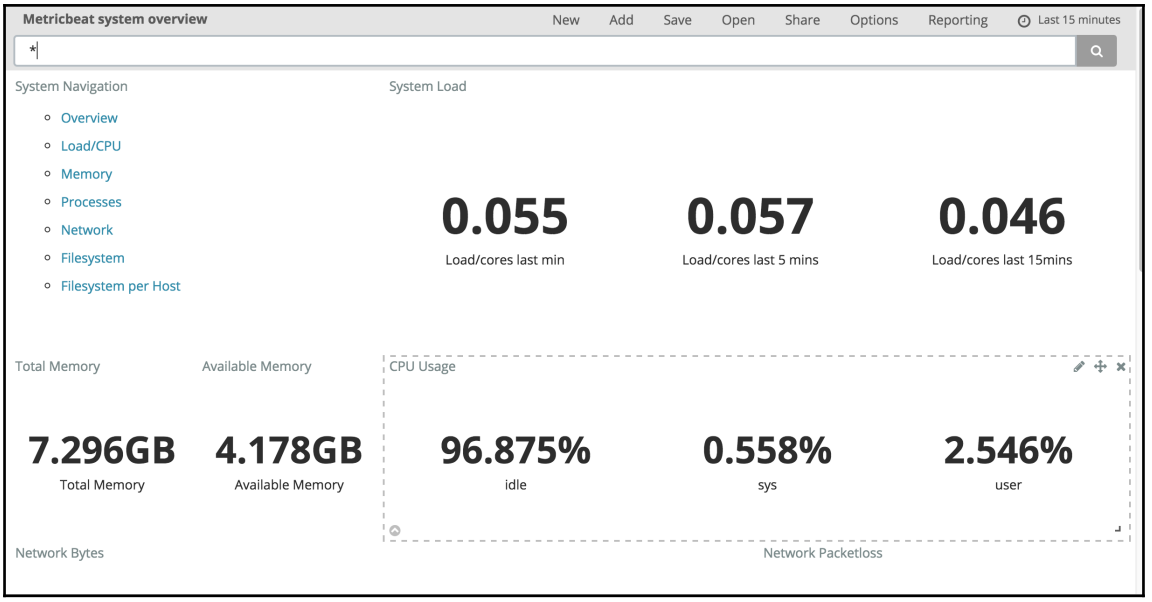
Metricbeat system overview
New
Add
Save
Open
Share
Options
Reporting
Last 15 minutes
*
System Navigation
Overview
Load/CPU
Memory
Processes
Network
Filesystem
Filesystem per Host
System Load
0.055
Load/cores last min
0.057
Load/cores last 5 mins
0.046
Load/cores last 15mins
Total Memory
7.296GB
Total Memory
Available Memory
4.178GB
Available Memory
CPU Usage
96.875%
idle
0.558%
sys
2.546%
user
Network Bytes
Network Packetloss

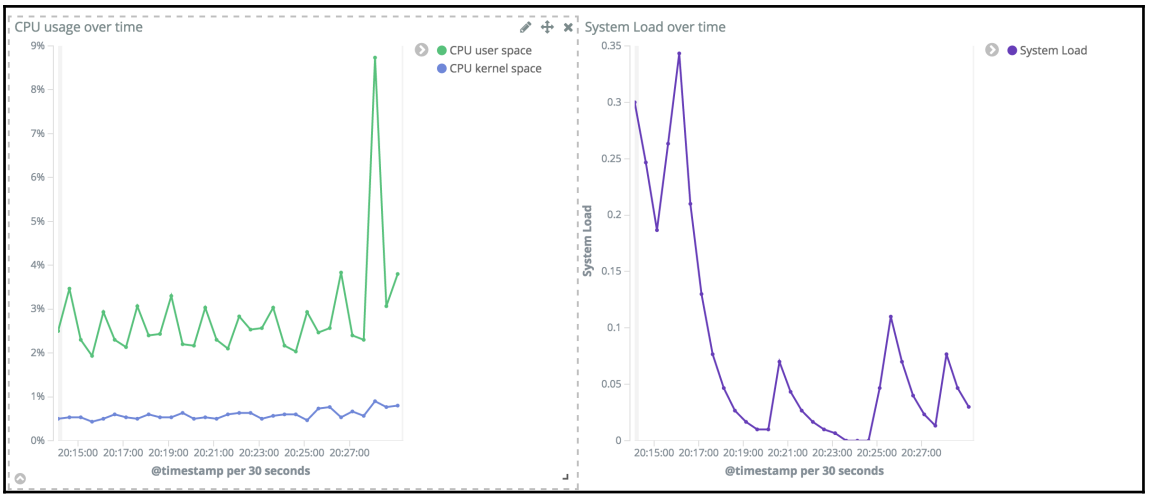
CPU usage over time
CPU user space
CPU kernel space
9%
8%
7%
6%
5%
4%
3%
2%
1%
0%
20:15:00 20:17:00 20:19:00 20:21:00 20:23:00 20:25:00 20:27:00
@timestamp per 30 seconds
System Load over time
System Load
System Load
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
20:15:00 20:17:00 20:19:00 20:21:00 20:23:00 20:25:00 20:27:00
@timestamp per 30 seconds

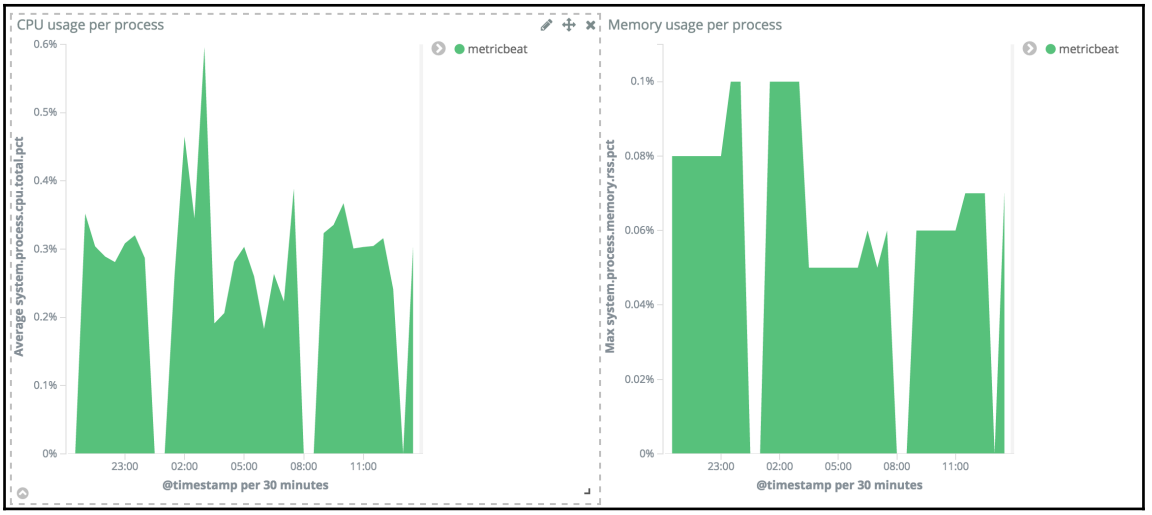
CPU usage per process
metricbeat
Average system.process.cpu.total.pct
0.6%
0.5%
0.4%
0.3%
0.2%
0.1%
0%
23:00 02:00 05:00 08:00 11:00
@timestamp per 30 minutes
Memory usage per process
metricbeat
Max system.process.memory.rss.pct
0.1%
0.08%
0.06%
0.04%
0.02%
0%
23:00 02:00 05:00 08:00 11:00
@timestamp per 30 minutes

Top processes

| system.process.name: Descending | Total CPU usage | Resident Memory | Shared memory |
|---|---|---|---|
| zoom.us | 147.98% | 1.59% | 0 |
| Google Chrome H | 110.54% | 6.94% | 0 |
| CalendarAgent | 107.03% | 0.31% | 0 |
| bsdtar | 91.77% | 0.03% | 0 |
| suggestd | 77.2% | 0.2% | 0 |

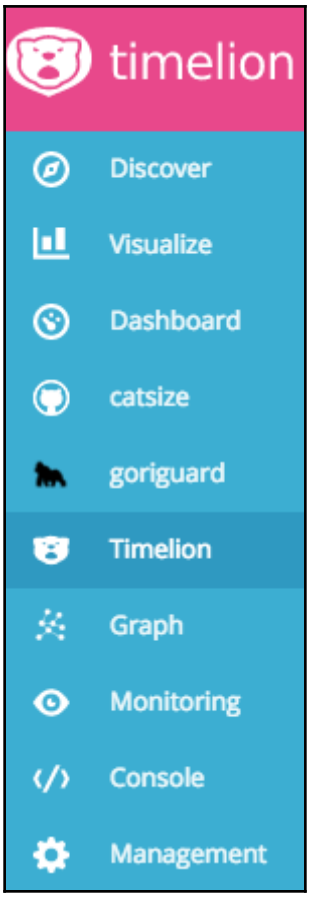
timelion
Discover
Visualize
Dashboard
catsize
goriguard
Timelion
Graph
Monitoring
Console
Management

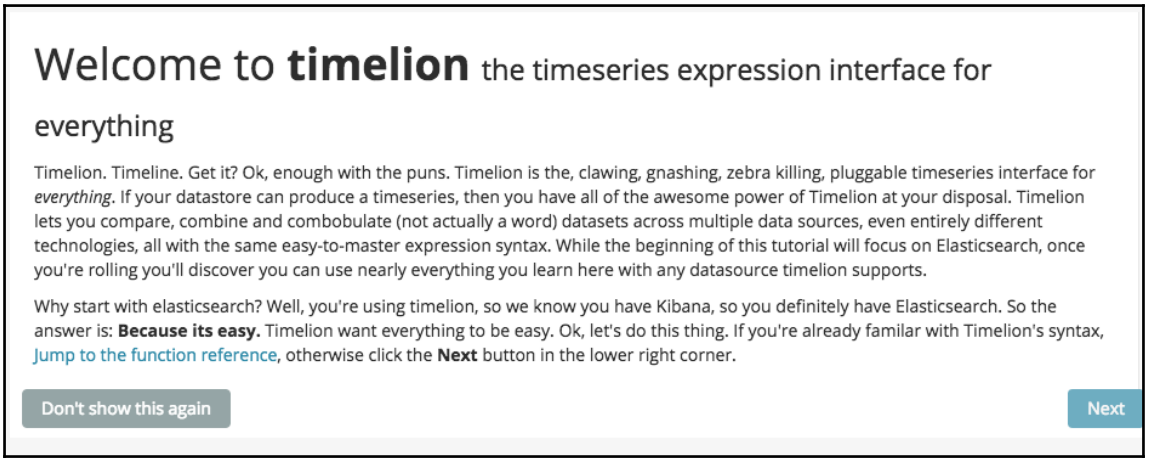
Welcome to timelion the timeseries expression interface for everything
Timelion. Timeline. Get it? Ok, enough with the puns. Timelion is the, clawing, gnashing, zebra killing, pluggable timeseries interface for *everything*. If your datastore can produce a timeseries, then you have all of the awesome power of Timelion at your disposal. Timelion lets you compare, combine and combobulate (not actually a word) datasets across multiple data sources, even entirely different technologies, all with the same easy-to-master expression syntax. While the beginning of this tutorial will focus on Elasticsearch, once you're rolling you'll discover you can use nearly everything you learn here with any datasource timelion supports.
Why start with elasticsearch? Well, you're using timelion, so we know you have Kibana, so you definitely have Elasticsearch. So the answer is: **Because its easy.** Timelion want everything to be easy. Ok, let's do this thing. If you're already familar with Timelion's syntax, Jump to the function reference, otherwise click the **Next** button in the lower right corner.
Don't show this again
Next

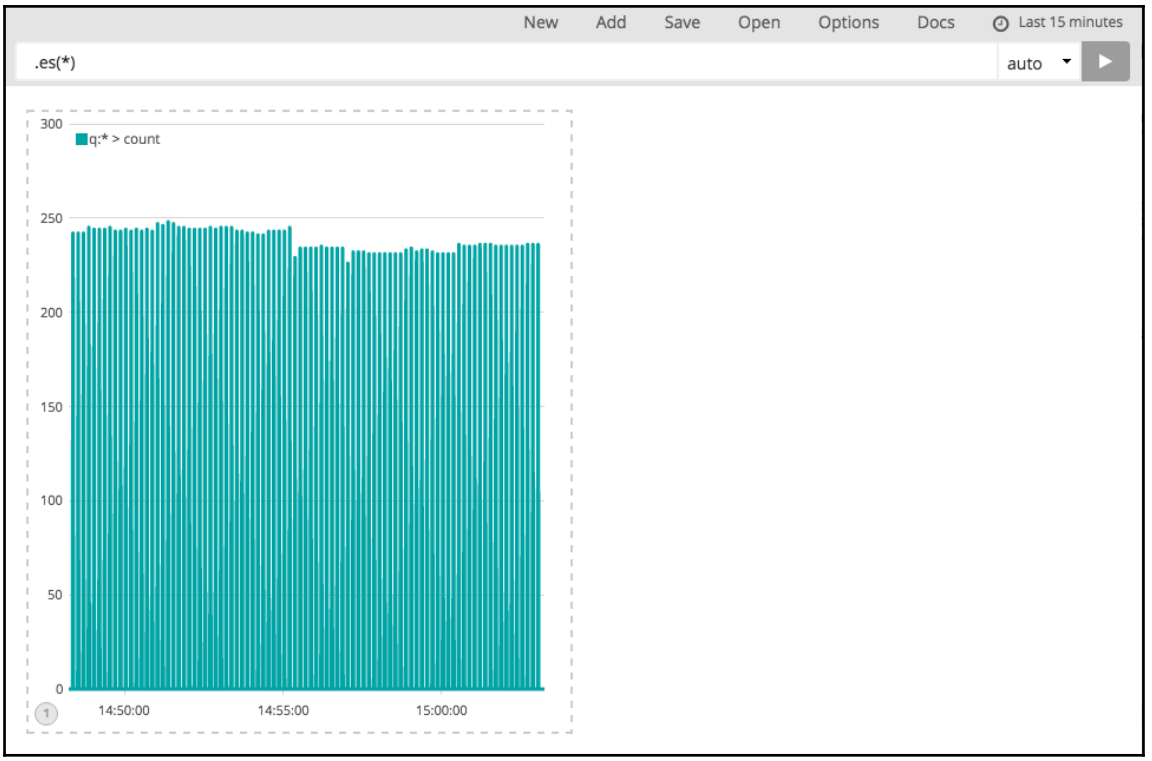
New
Add
Save
Open
Options
Docs
Last 15 minutes
.es(*)
auto
q:* > count
300
250
200
150
100
50
0
14:50:00
14:55:00
15:00:00


Timelion: Error: Max buckets exceeded: 31622400 of 2000 allowed. Choose a larger interval or a shorter time span
276s
More Info
OK

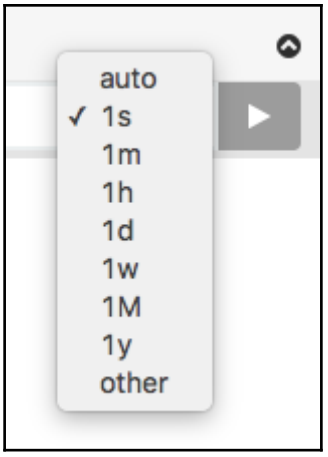
auto
1s
1m
1h
1d
1w
1M
1y
other

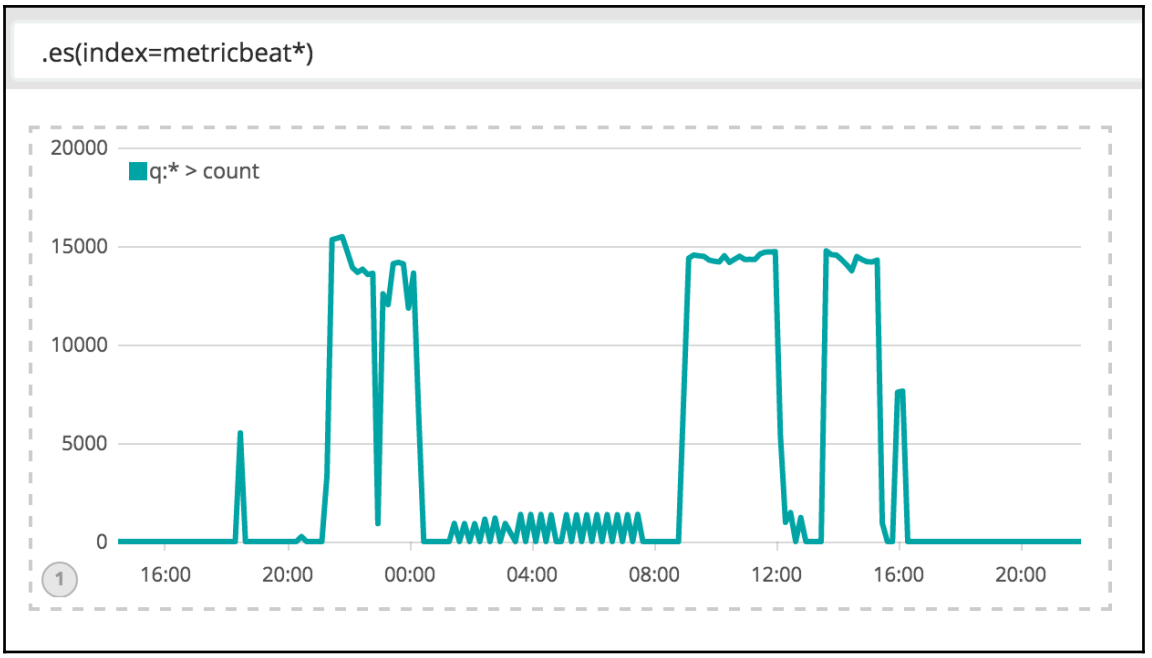
.es(index=metricbeat*)
q:* > count
20000
15000
10000
5000
0
1
16:00
20:00
00:00
04:00
08:00
12:00
16:00
20:00

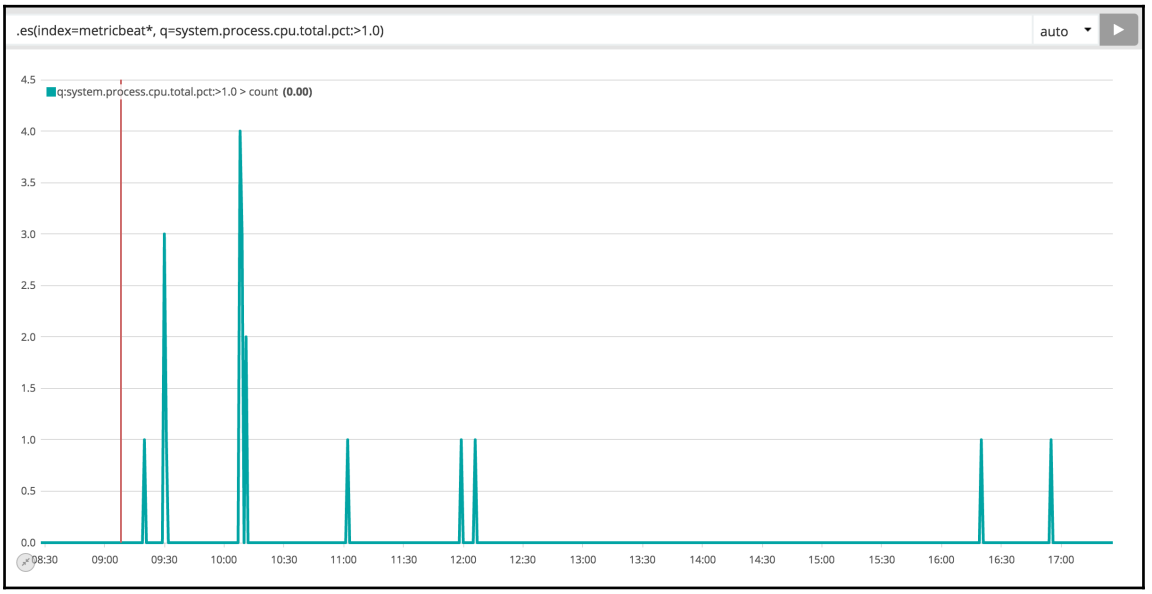
.es(index=metricbeat*, q=system.process.cpu.total.pct:>1.0)
auto
q:system.process.cpu.total.pct:>1.0 > count (0.00)
4.5
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
08:30
09:00
09:30
10:00
10:30
11:00
11:30
12:00
12:30
13:00
13:30
14:00
14:30
15:00
15:30
16:00
16:30
17:00

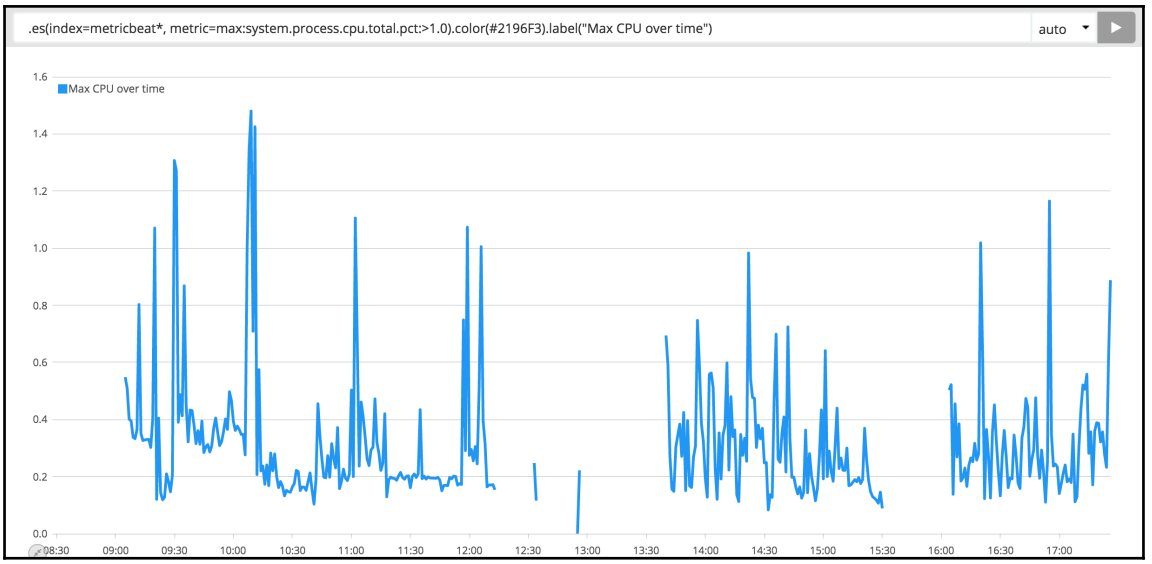
.es(index=metricbeat*, metric=max:system.process.cpu.total.pct:>1.0).color(#2196F3).label("Max CPU over time")
auto
Max CPU over time
1.6
1.4
1.2
1.0
0.8
0.6
0.4
0.2
0.0
08:30
09:00
09:30
10:00
10:30
11:00
11:30
12:00
12:30
13:00
13:30
14:00
14:30
15:00
15:30
16:00
16:30
17:00

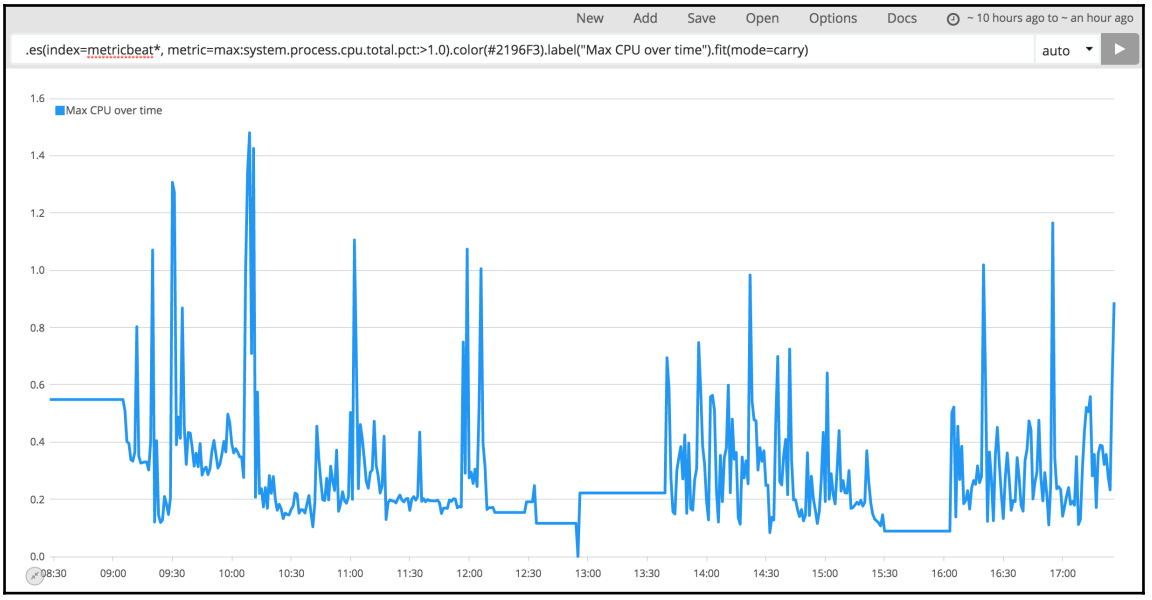
New
Add
Save
Open
Options
Docs
~ 10 hours ago to ~ an hour ago
.es(index=metricbeat*, metric=max:system.process.cpu.total.pct:>1.0).color(#2196F3).label("Max CPU over time").fit(mode=carry)
auto
Max CPU over time
1.6
1.4
1.2
1.0
0.8
0.6
0.4
0.2
0.0
08:30
09:00
09:30
10:00
10:30
11:00
11:30
12:00
12:30
13:00
13:30
14:00
14:30
15:00
15:30
16:00
16:30
17:00

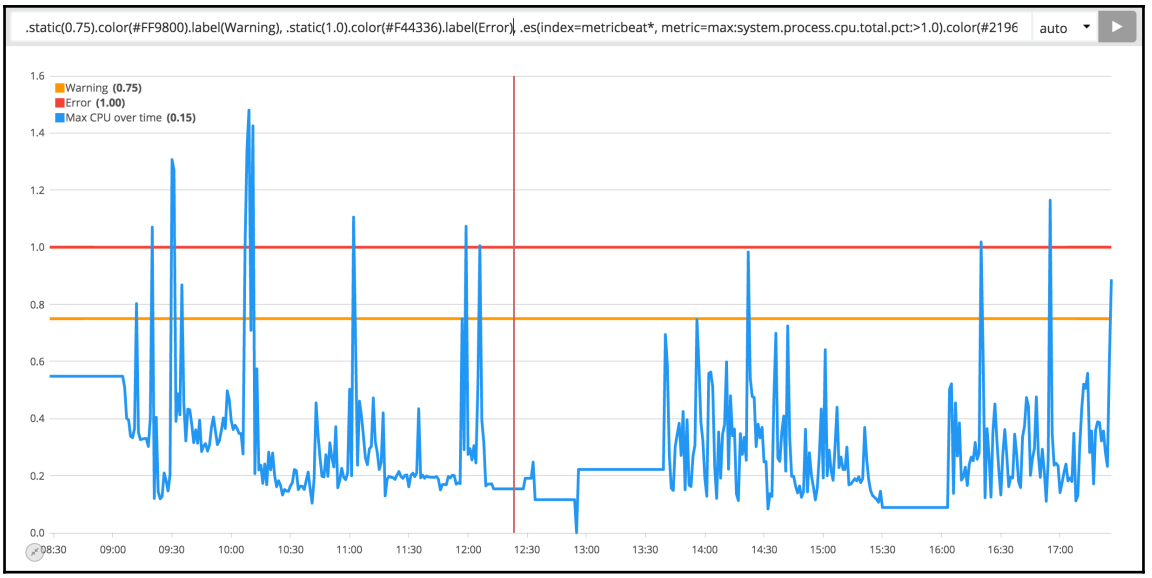
.static(0.75).color(#FF9800).label(Warning), .static(1.0).color(#F44336).label(Error), .es(index=metricbeat*, metric=max:system.process.cpu.total.pct:>1.0).color(#2196
auto
Warning (0.75)
Error (1.00)
Max CPU over time (0.15)
1.6
1.4
1.2
1.0
0.8
0.6
0.4
0.2
0.0
08:30
09:00
09:30
10:00
10:30
11:00
11:30
12:00
12:30
13:00
13:30
14:00
14:30
15:00
15:30
16:00
16:30
17:00

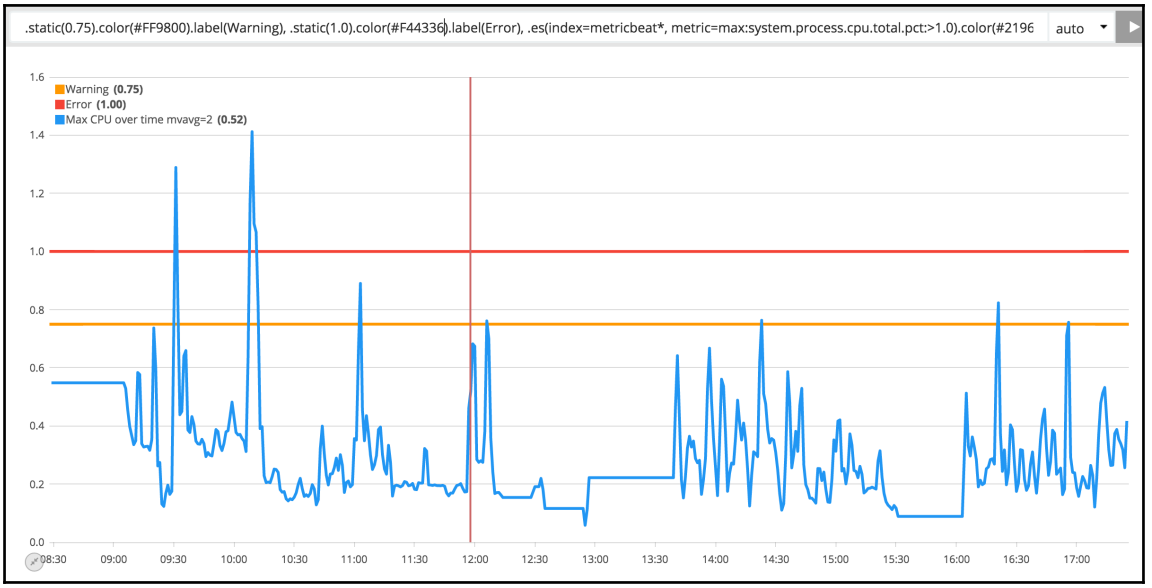
.static(0.75).color(#FF9800).label(Warning), .static(1.0).color(#F44336).label(Error), .es(index=metricbeat*, metric=max:system.process.cpu.total.pct:>1.0).color(#2196
auto
Warning (0.75)
Error (1.00)
Max CPU over time mvavg=2 (0.52)
1.6
1.4
1.2
1.0
0.8
0.6
0.4
0.2
0.0
08:30
09:00
09:30
10:00
10:30
11:00
11:30
12:00
12:30
13:00
13:30
14:00
14:30
15:00
15:30
16:00
16:30
17:00

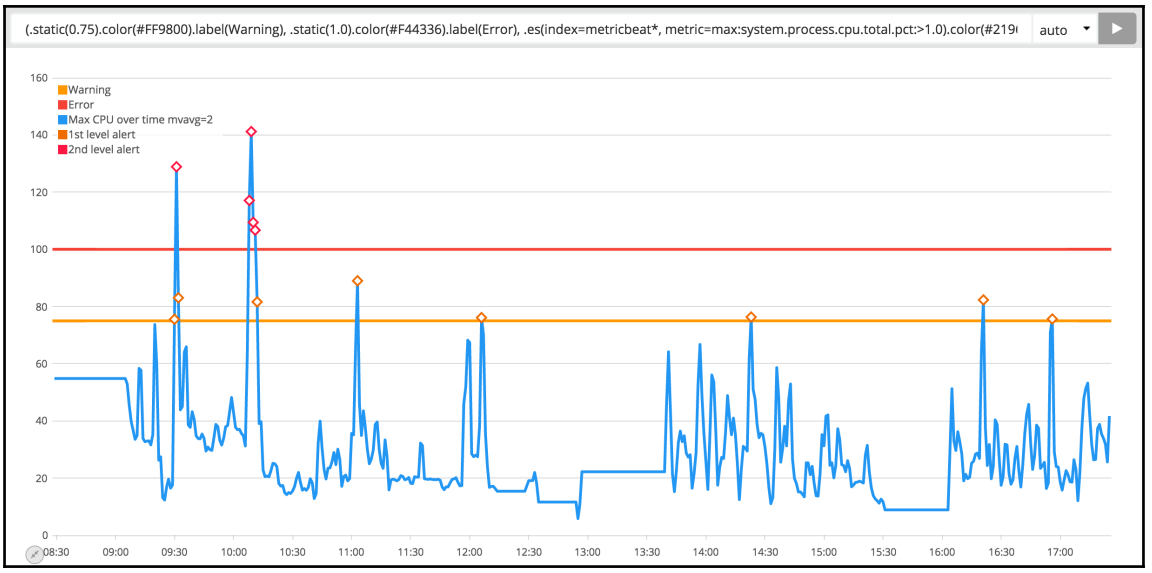
(.static(0.75).color(#FF9800).label(Warning), .static(1.0).color(#F44336).label(Error), .es(index=metricbeat*, metric=max:system.process.cpu.total.pct:>1.0).color(#219
auto
Warning
Error
Max CPU over time mvavg=2
1st level alert
2nd level alert
160
140
120
100
80
60
40
20
0
08:30
09:00
09:30
10:00
10:30
11:00
11:30
12:00
12:30
13:00
13:30
14:00
14:30
15:00
15:30
16:00
16:30
17:00

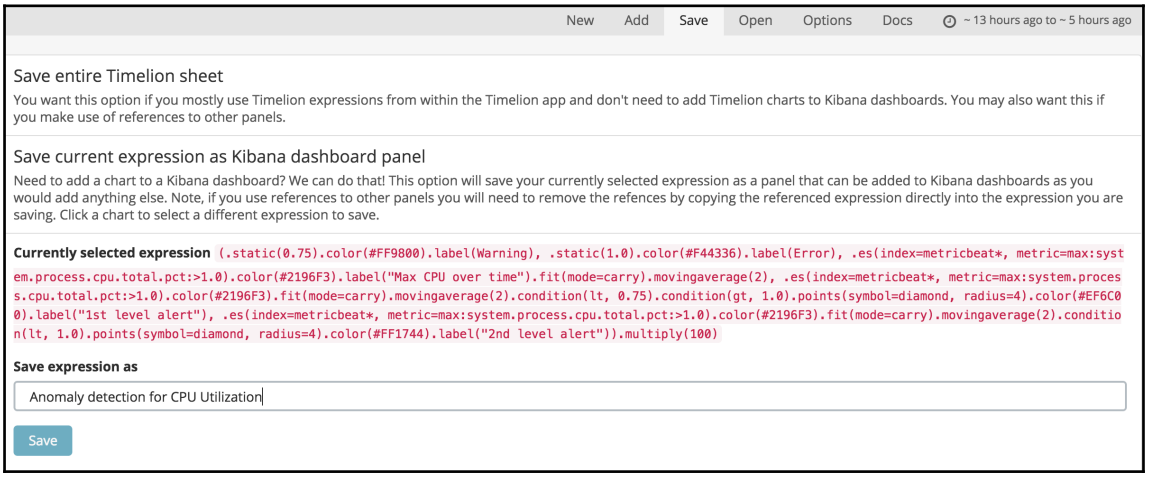
New
Add
Save
Open
Options
Docs
~ 13 hours ago to ~ 5 hours ago
Save entire Timelion sheet
You want this option if you mostly use Timelion expressions from within the Timelion app and don't need to add Timelion charts to Kibana dashboards. You may also want this if you make use of references to other panels.
Save current expression as Kibana dashboard panel
Need to add a chart to a Kibana dashboard? We can do that! This option will save your currently selected expression as a panel that can be added to Kibana dashboards as you would add anything else. Note, if you use references to other panels you will need to remove the refences by copying the referenced expression directly into the expression you are saving. Click a chart to select a different expression to save.
Currently selected expression (.static(0.75).color(#FF9800).label(Warning), .static(1.0).color(#F44336).label(Error), .es(index=metricbeat*, metric=max:system.process.cpu.total.pct:>1.0).color(#2196F3).label("Max CPU over time").fit(mode=carry).movingaverage(2), .es(index=metricbeat*, metric=max:system.process.cpu.total.pct:>1.0).color(#2196F3).fit(mode=carry).movingaverage(2).condition(lt, 0.75).condition(gt, 1.0).points(symbol=diamond, radius=4).color(#EF6C00).label("1st level alert"), .es(index=metricbeat*, metric=max:system.process.cpu.total.pct:>1.0).color(#2196F3).fit(mode=carry).movingaverage(2).condition(lt, 1.0).points(symbol=diamond, radius=4).color(#FF1744).label("2nd level alert")).multiply(100)
Save expression as
Anomaly detection for CPU Utilization
Save

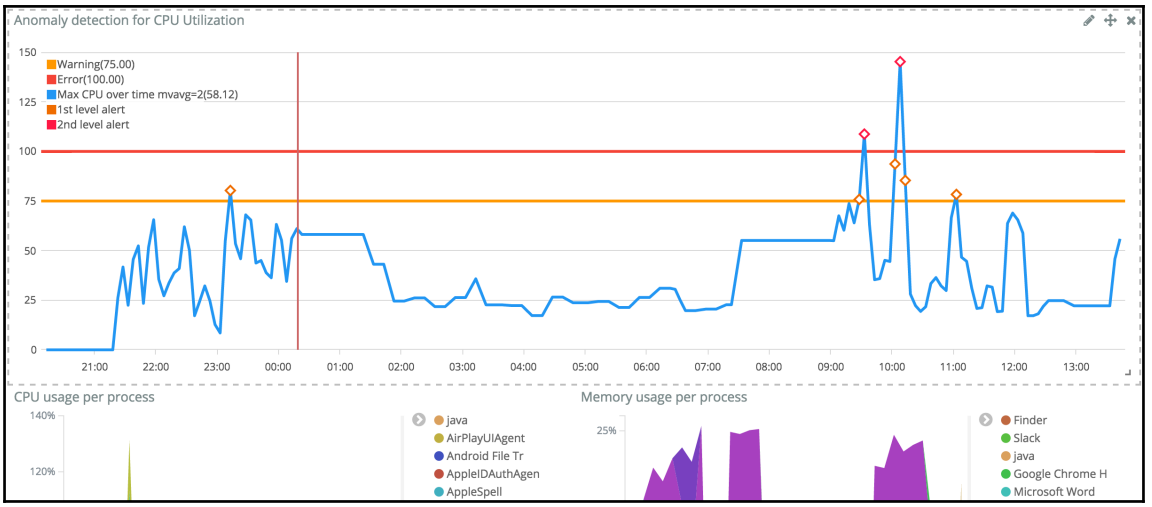
Anomaly detection for CPU Utilization
Warning(75.00)
Error(100.00)
Max CPU over time mvavg=2(58.12)
1st level alert
2nd level alert
150
125
100
75
50
25
0
21:00
22:00
23:00
00:00
01:00
02:00
03:00
04:00
05:00
06:00
07:00
08:00
09:00
10:00
11:00
12:00
13:00
CPU usage per process
140%
120%
java
AirPlayUIAgent
Android File Tr
AppleIDAuthAgen
AppleSpell
Memory usage per process
25%
Finder
Slack
java
Google Chrome H
Microsoft Word

Elasticsearch Console

```
84  }
85
86
87
88  PUT _xpack/watcher/watch/cpu_watch
89  {
90    "trigger": {
91      "schedule": {
92        "interval": "10s"
93      }
94    },
95    "input": {
96      "search": {
97        "request": {
98          "indices": [
99            "metricbeat*"
100         ],
101         "body": {
102           "size": 0,
103           "aggs": {
104             "max_cpu": {
105               "max": {
106                 "field": "system.process.cpu.total.pct"
107               }
108             }
109           },
110           "query": {
111             "bool": {
112               "must": [
113                 {
114                   "range": {
115                   "@timestamp": {
116                     "gte": "now-10s"
```

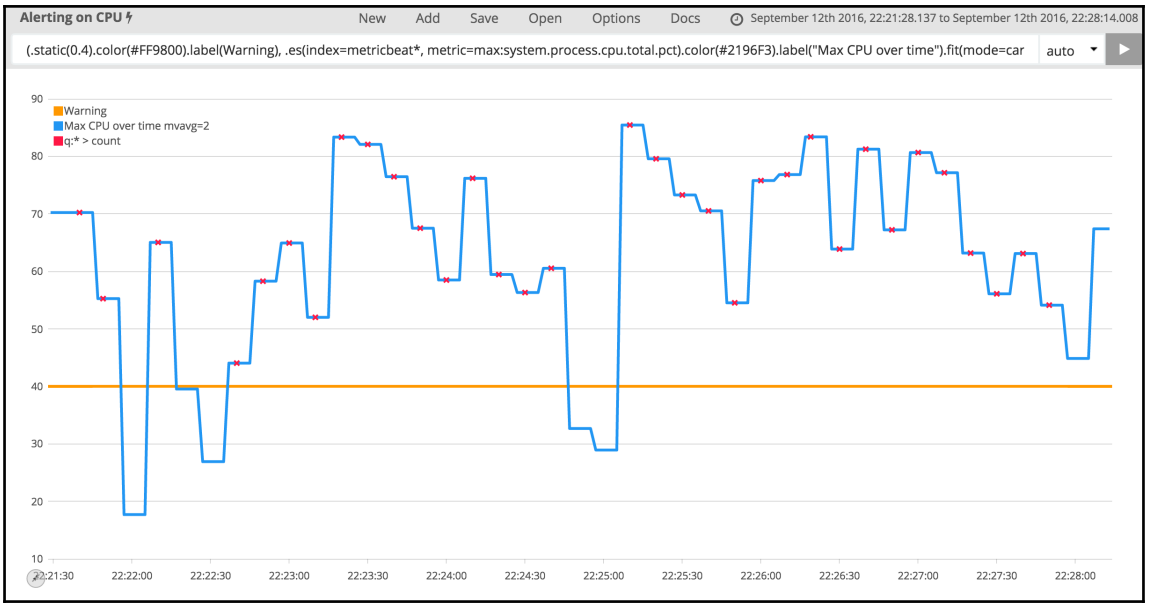
Alerting on CPU
New
Add
Save
Open
Options
Docs
September 12th 2016, 22:21:28.137 to September 12th 2016, 22:28:14.008
(.static(0.4).color(#FF9800).label(Warning), .es(index=metricbeat*, metric=max:system.process.cpu.total.pct).color(#2196F3).label("Max CPU over time").fit(mode=car
auto
Warning
Max CPU over time mvavg=2
q:* > count
90
80
70
60
50
40
30
20
10
22:21:30
22:22:00
22:22:30
22:23:00
22:23:30
22:24:00
22:24:30
22:25:00
22:25:30
22:26:00
22:26:30
22:27:00
22:27:30
22:28:00

# Chapter 6: Graph Exploration in Kibana

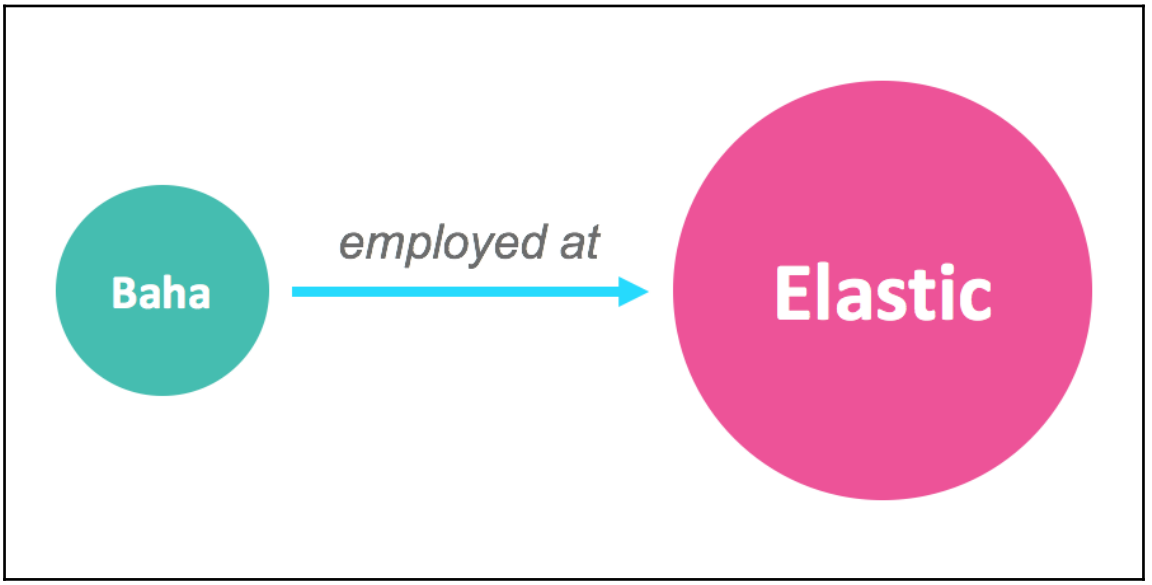


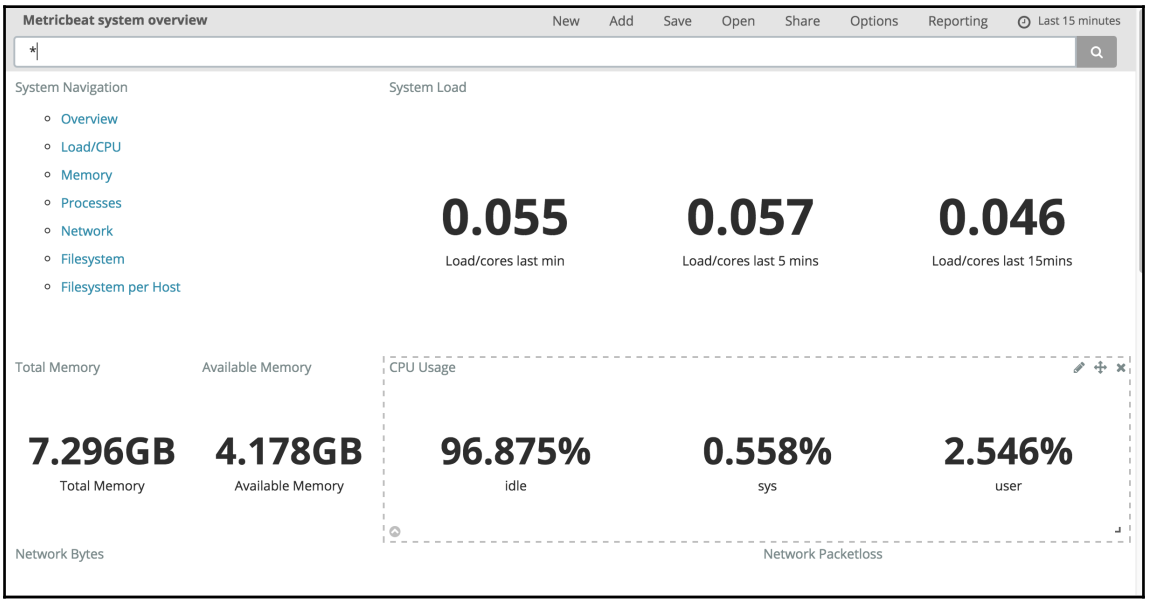



**Most frequent connections**

Emphasis the ***popularity*** on the record ! Poor Bach …

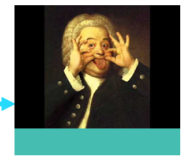

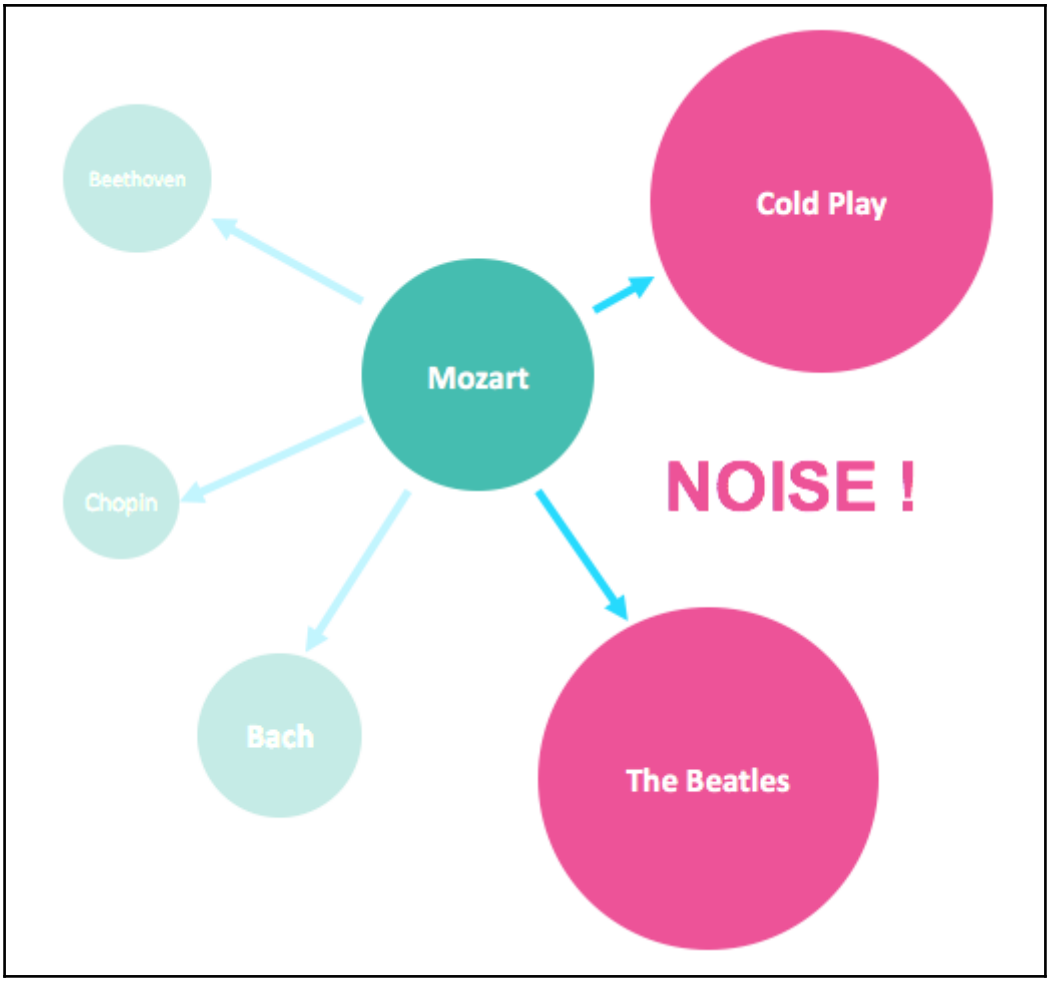
Beethoven
Cold Play
Mozart
Chopin
NOISE !
Bach
The Beatles

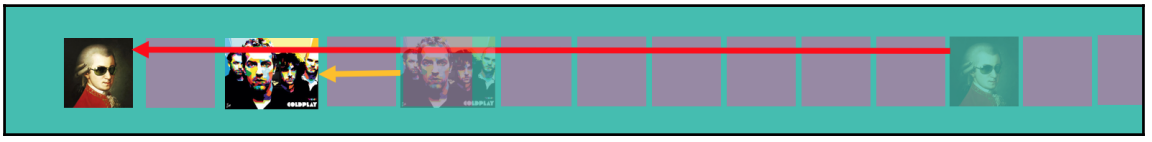

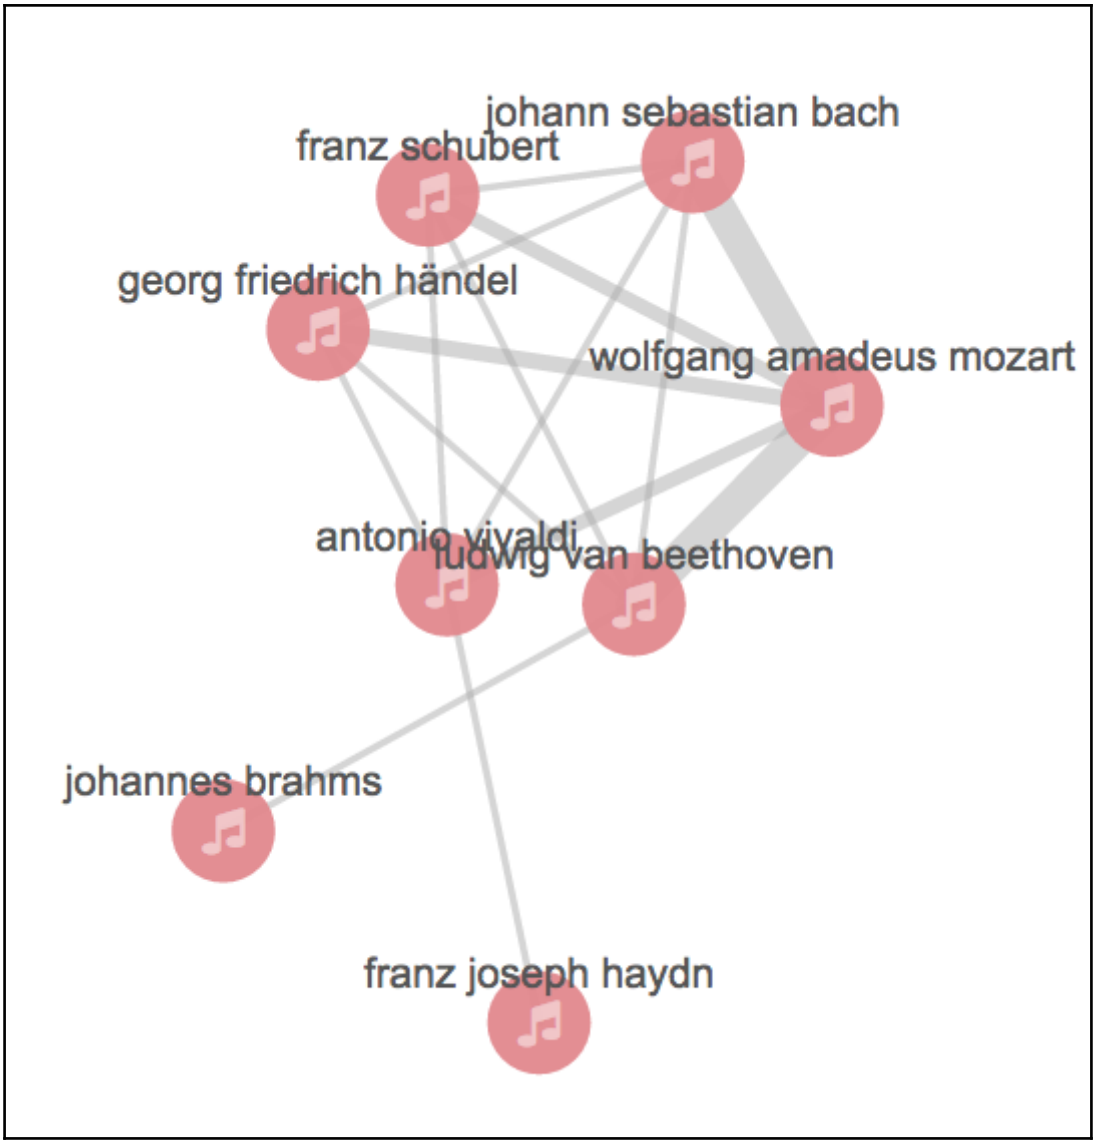
johann sebastian bach
franz schubert
georg friedrich händel
wolfgang amadeus mozart
antonio vivaldi
ludwig van beethoven
johannes brahms
franz joseph haydn

## Significance-led connections

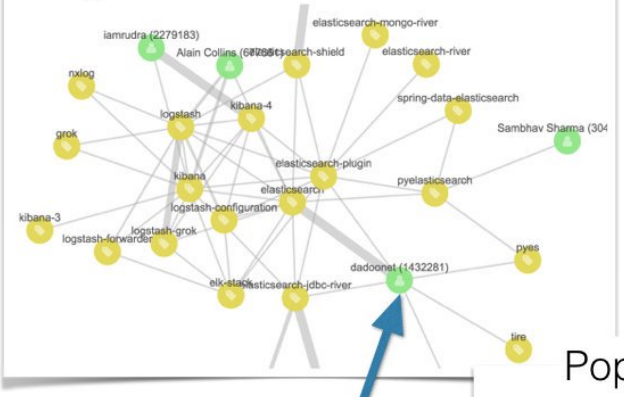


## Popularity-led connections

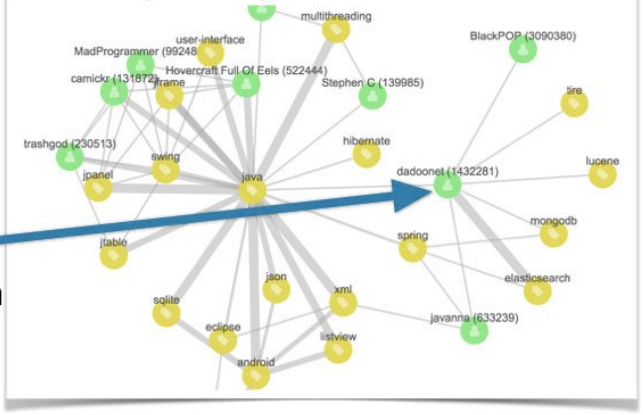


- Same StackOverflow data
- Same start point
- Different graph exploration algos

**Mark Harwood** @elasticmark · 13 juin
Graph exploration needs relevance eg the StackOverflow links of our very own @dadoonet with and without significance

12 8 •••

Dev Tools

Console History Settings Help

```
1  GET _cat/indices/stack*
```

```
1  green open stackoverflow GPR8PhlwQwO9CLzf0pBSkA 1 0 11192635 0 2.1gb 2.1gb
2
```

# Configure an index pattern

In order to use Kibana you must configure at least one index pattern. Index patterns are used to identify the Elasticsearch index to run search and analytics against. They are also used to configure fields.

☐ **Index contains time-based events**

**Index name or pattern**

Patterns allow you to define dynamic index names using * as a wildcard. Example: logstash-*

stackoverflow

Create

**tag:** javascript, jquery **user:** learnmore (1742289), ckoverflow **_score:** 1

Table JSON


```
1  {
2    "_index": "stackoverflow",
3    "_type": "qna",
4    "_id": "AVgSYpTT1vXMTQ2oWJCN",
5    "_score": 1,
6    "_source": {
7      "tag": [
8        "javascript",
9        "jquery"
10     ],
11     "user": [
12       "learnmore (1742289)",
13       "Vohuman (848164)"
14     ]
15   }
16 }
```


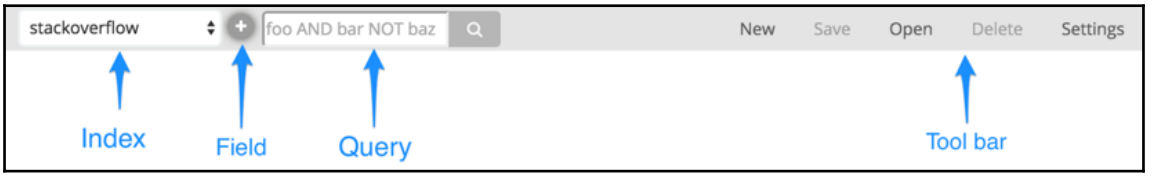

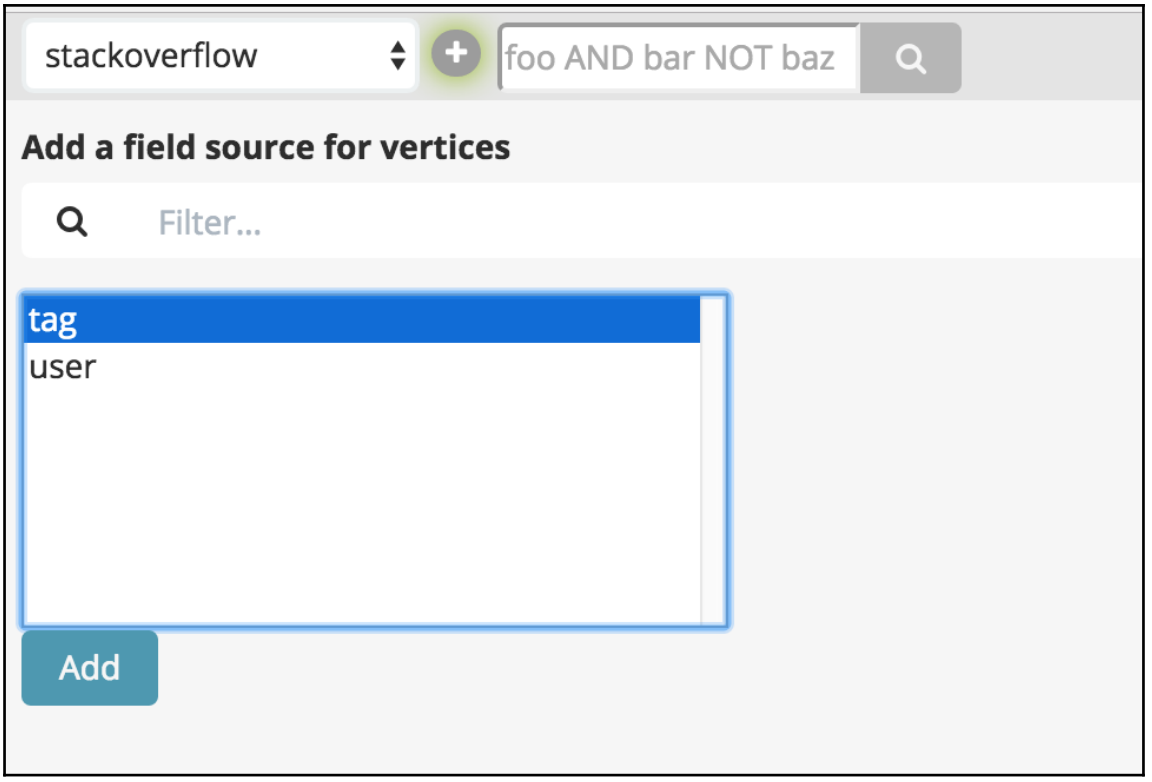
stackoverflow
foo AND bar NOT baz
Add a field source for vertices
Filter...
tag
user
Add

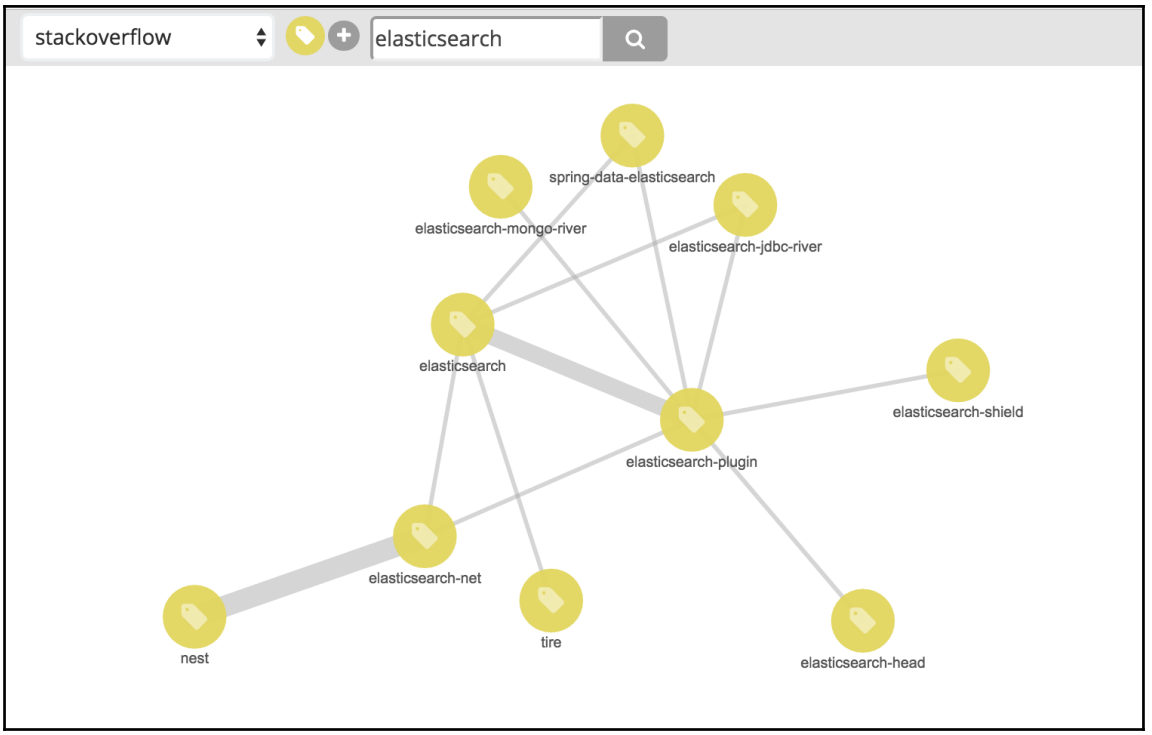
stackoverflow
elasticsearch
spring-data-elasticsearch
elasticsearch-mongo-river
elasticsearch-jdbc-river
elasticsearch
elasticsearch-shield
elasticsearch-plugin
elasticsearch-net
nest
tire
elasticsearch-head

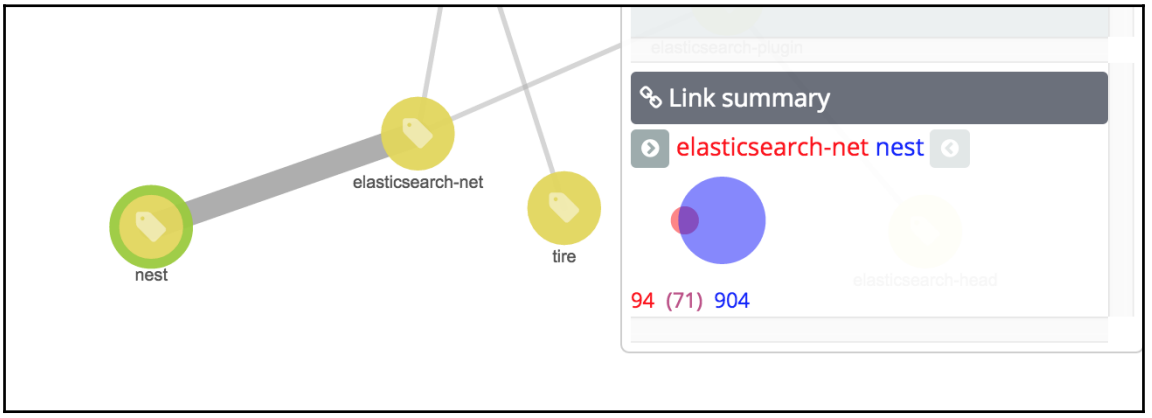
elasticsearch-net
nest
tire
Link summary
elasticsearch-net nest
94 (71) 904

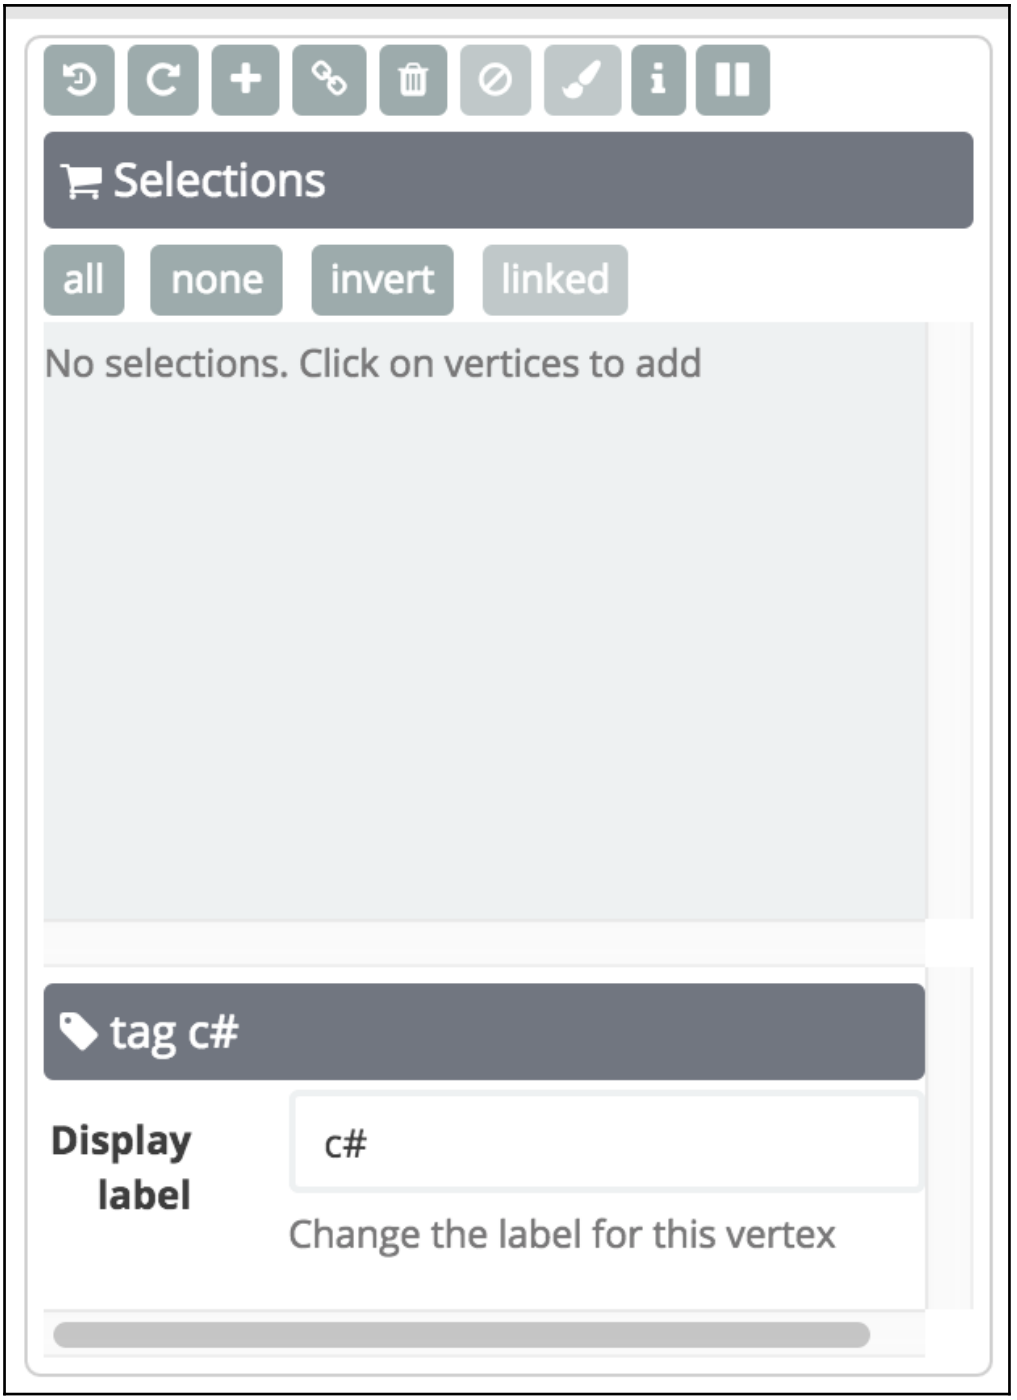
Selections
all
none
invert
linked
No selections. Click on vertices to add
tag c#
Display label
c#
Change the label for this vertex

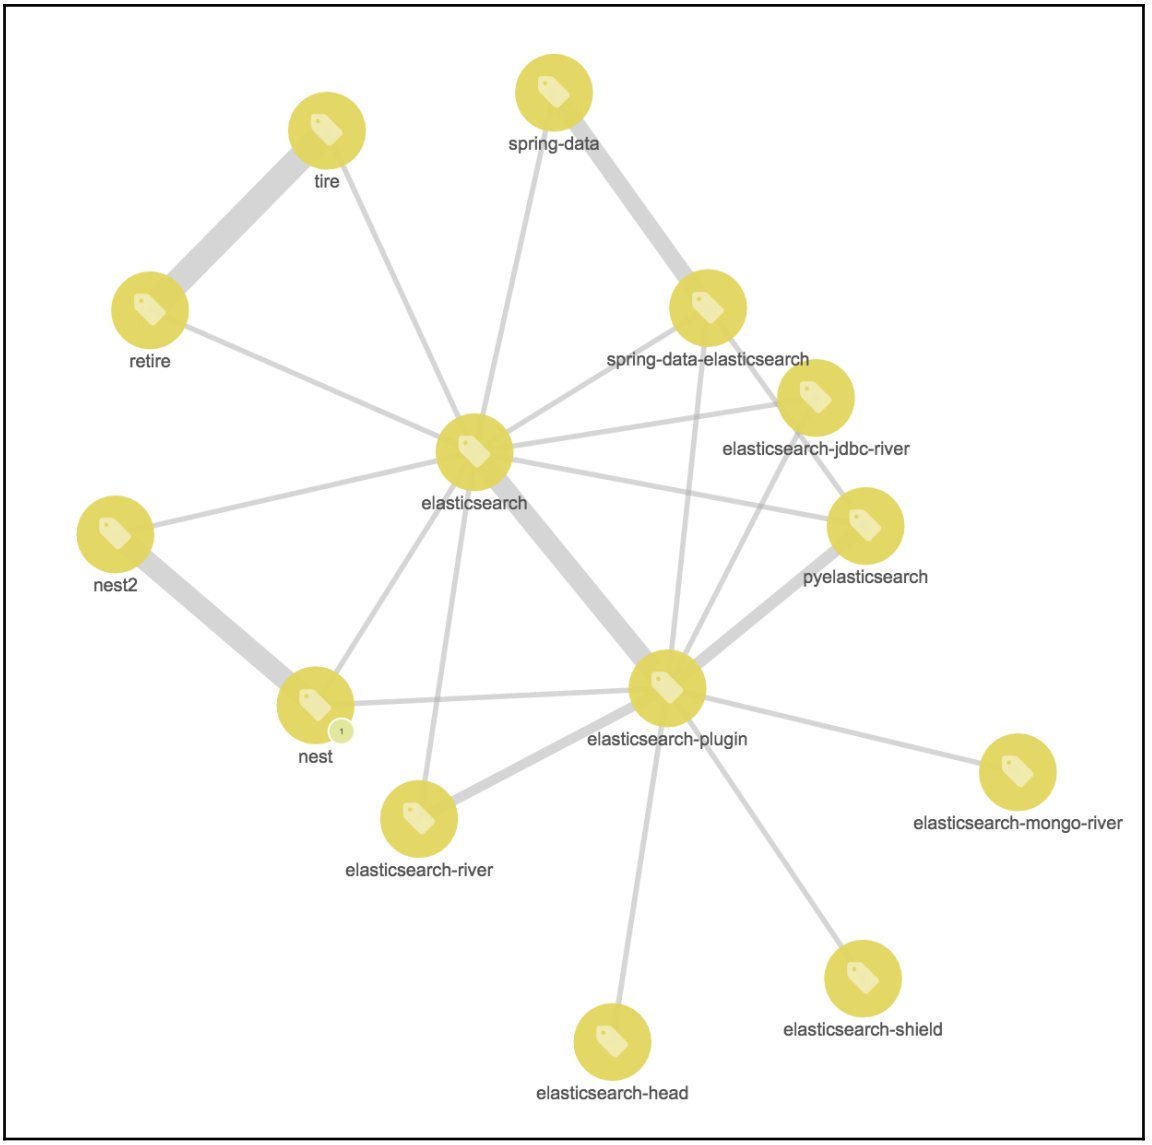
spring-data
tire
retire
spring-data-elasticsearch
elasticsearch-jdbc-river
elasticsearch
nest2
pyelasticsearch
nest
elasticsearch-plugin
elasticsearch-mongo-river
elasticsearch-river
elasticsearch-shield
elasticsearch-head

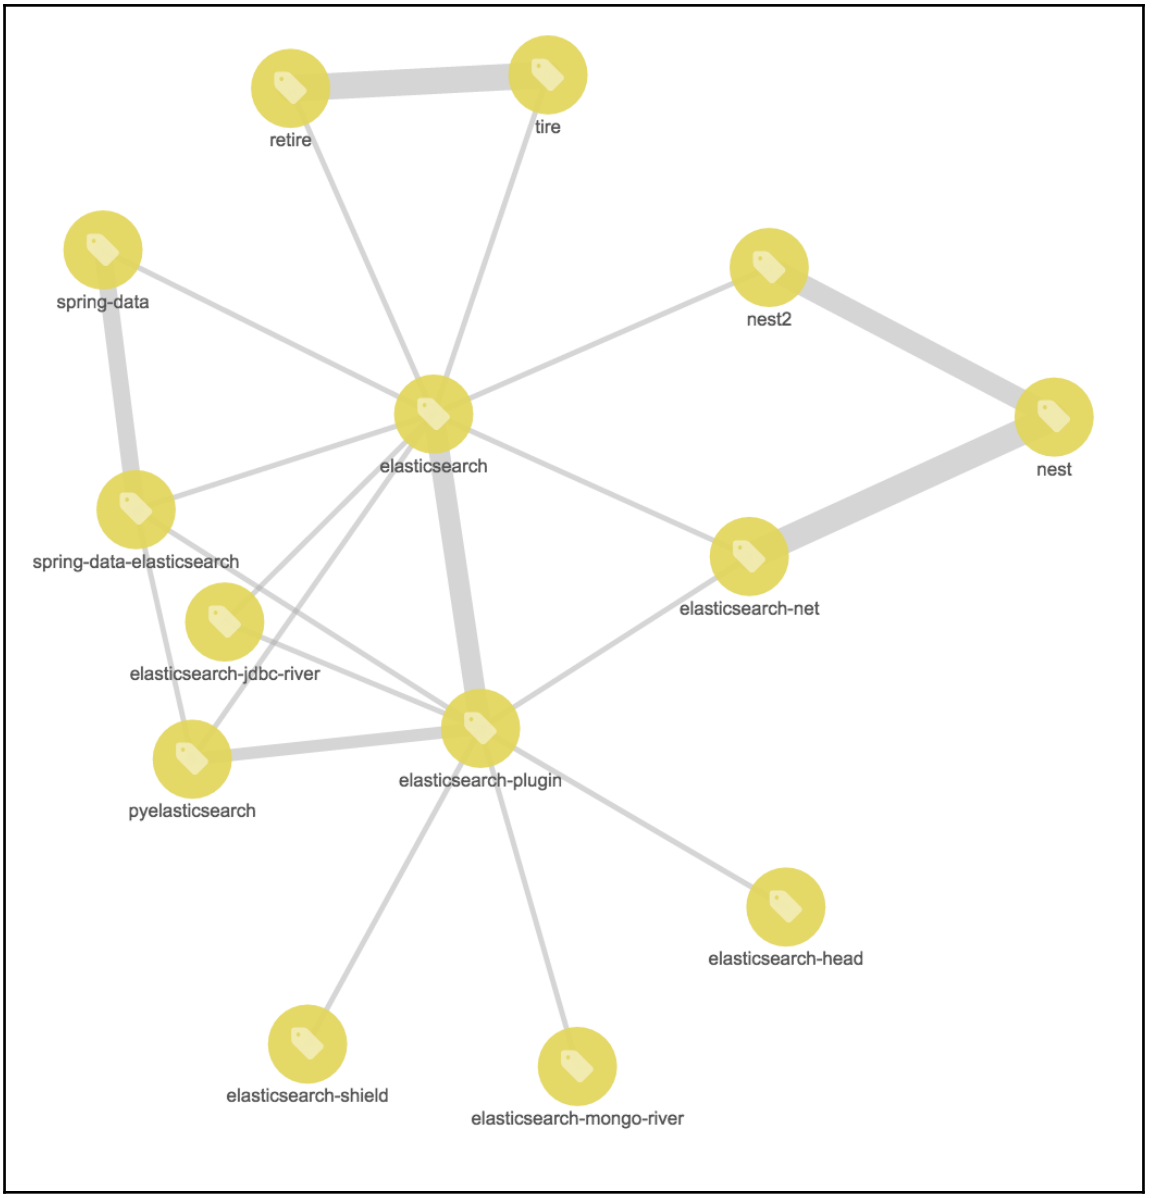
retire
tire
spring-data
nest2
elasticsearch
nest
spring-data-elasticsearch
elasticsearch-net
elasticsearch-jdbc-river
elasticsearch-plugin
pyelasticsearch
elasticsearch-head
elasticsearch-shield
elasticsearch-mongo-river

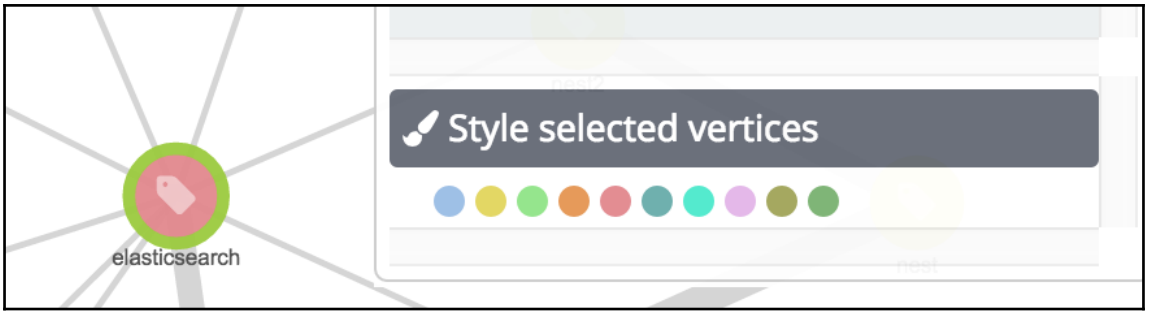
elasticsearch
Style selected vertices

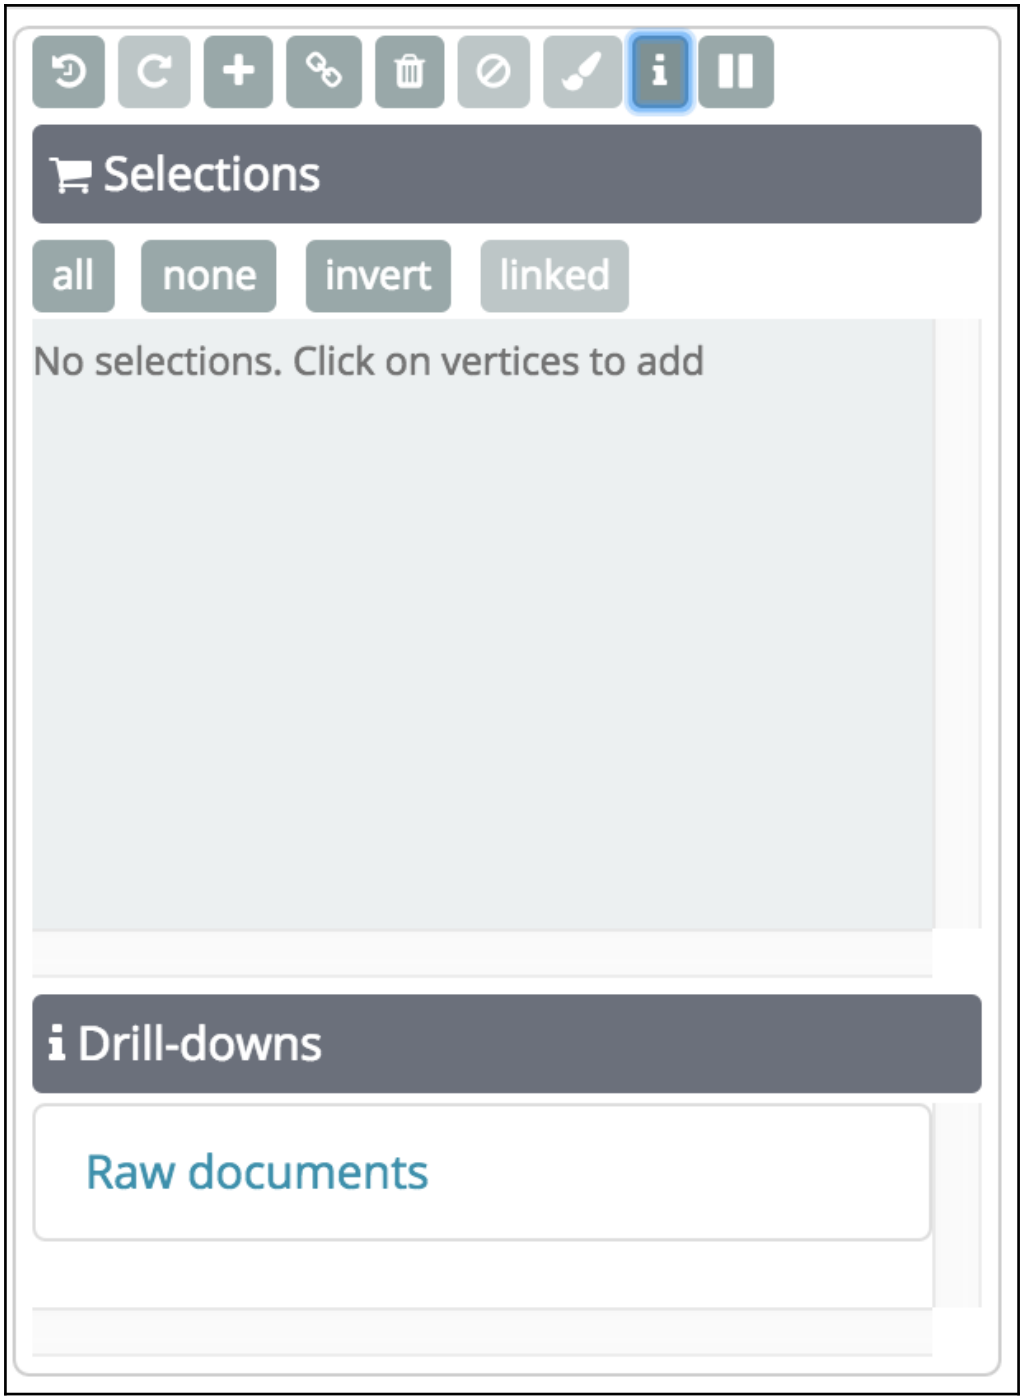
Selections
all
none
invert
linked
No selections. Click on vertices to add
Drill-downs
Raw documents

15,599 hits

New Save Open Share Reporting

{"bool":{"minimum_should_match":1,"should":[{"term":{"tag":"elasticsearch"}}]}}

**stackoverflow**

Selected Fields

? _source

Available Fields

t _id

t _index

# _score

t _type

t tag

t user

**_source**

tag: elasticsearch, tire user: karmi (95696), Zack Fernandes (144903) _id: AVgSYpTT1vXMTQ2oWJDd _type: qna _index: stackoverflow _score: 6.565

tag: lucene, elasticsearch user: Rajni (1746198), dadoonet (1432281), imotov (783043) _id: AVgSYpcW1vXMTQ2oWJ1B _type: qna _index: stackoverflow _score: 6.565

tag: java, lucene, elasticsearch, apache-tika user: rstuart85 (1713202), farid (371236) _id: AVgSYpgp1vXMTQ2oWKUI _type: qna _index: stackoverflow _score: 6.565

tag: php, symfony2, elasticsearch user: CedricD (1713283), dbrumann (1166880), Robert Harvey (102937) _id: AVgSYpgp1vXMTQ2oWKYM _type: qna _index: stackoverflow _score: 6.565

Last request Blacklist Advanced settings Drill-downs

**URL** http://stackoverflow.com/questions/tagged/{{gquery}}

Define template URLs using {{gquery}} where the selected vertex terms are inserted

**Title** Display questions

**URL parameter type** plain text

Text of selected vertex labels as a plain url-encoded string

**Toolbar icon**

Reset Save

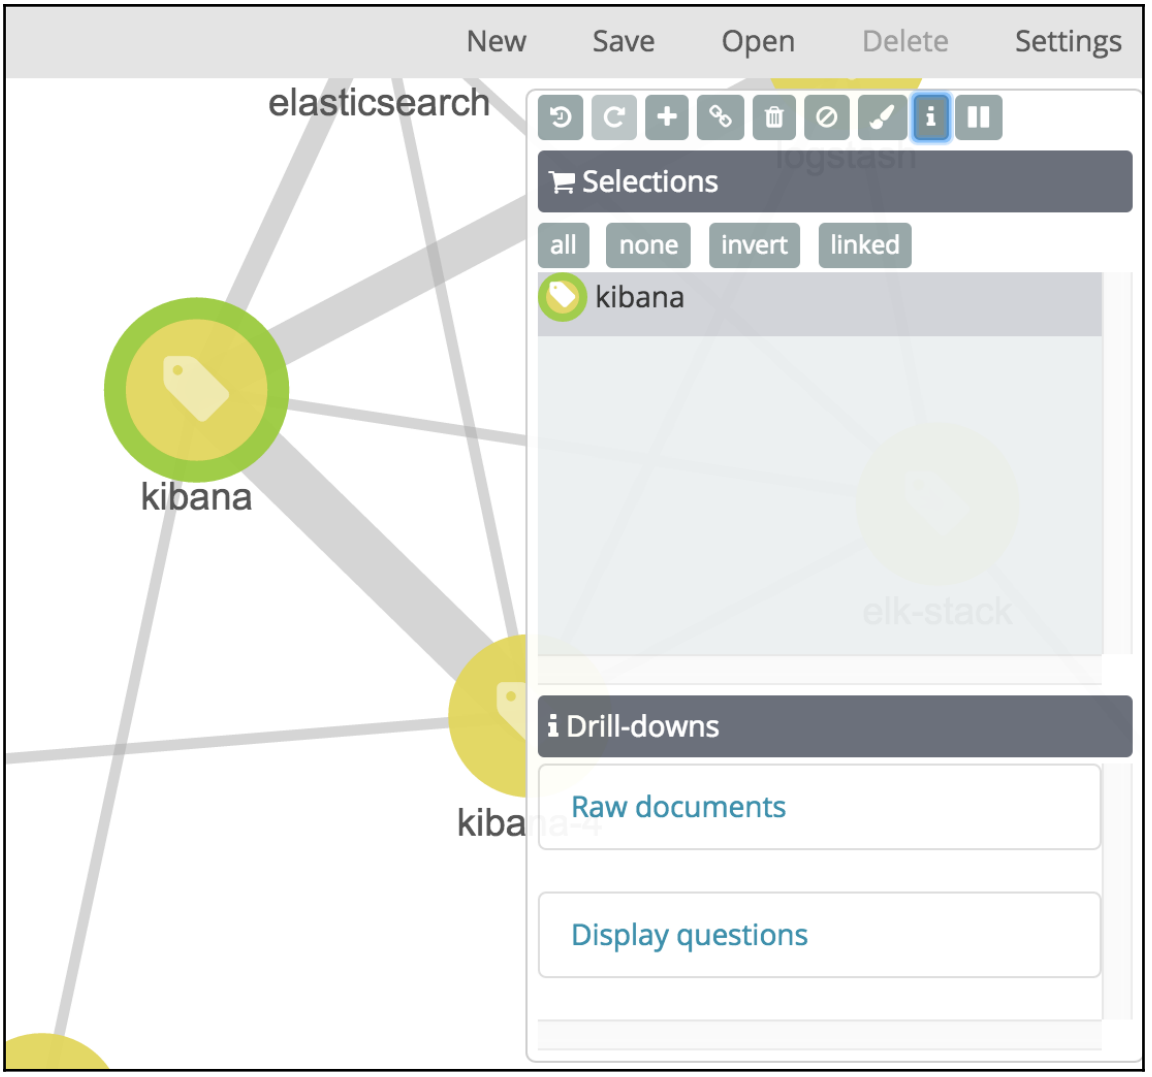
New
Save
Open
Delete
Settings
elasticsearch
Selections
all
none
invert
linked
kibana
kibana
elk-stack
Drill-downs
Raw documents
Display questions

StackExchange ▾ sign up log in

stackoverflow

Questions Jobs Documentation

## Tagged Questions

info newest frequent votes active unanswered

Kibana is an application for exploring and visualizing your data. It helps you find and demonstrate trends in your data with tools for searching, creating visualizations, and combining those visualizations to build dashboards.

learn more… top users synonyms

0 votes
0 answers
4 views

### Can I add visualizations related to mulitple indexes to single kibana dashboard?

I have two indexes which configured in kibana. I want to display graphs from two indexes in one dashboard is kibana support this? In kibana4 and above has this but I would like to know is this ...

kibana kibana-4 kibana-5

asked 46 mins ago
BEJGAM SHIVA PRASAD
119 • 6

0 votes
2 answers
14 views

### How to draw multi-lines from multiple queries in Kibana

I am new to Kibana and need some help. I can draw this line chart for a single query (java): Now I would like to another line for another query (for example python) in the same chart. I am not so ...

kibana kibana-4 elasticsearch-2.0

asked yesterday
kee
2,024 • 8 • 31 • 63

**Sample size** 2000

Terms are identified from samples of the most relevant documents. Bigger is not necessarily better - can be slower and less relevant.

☐ Significant links

Identify terms that are "significant" rather than simply popular

**Certainty** 3

The min number of documents that are required as evidence before introducing a related term

**Diversity field** [No diversification]

To avoid document samples being dominated by a single voice, pick the field that helps identify the source of bias. *This must be a single-term field or searches will be rejected with an error*

**Timeout (ms)** 5000

Max time in milliseconds a request can run

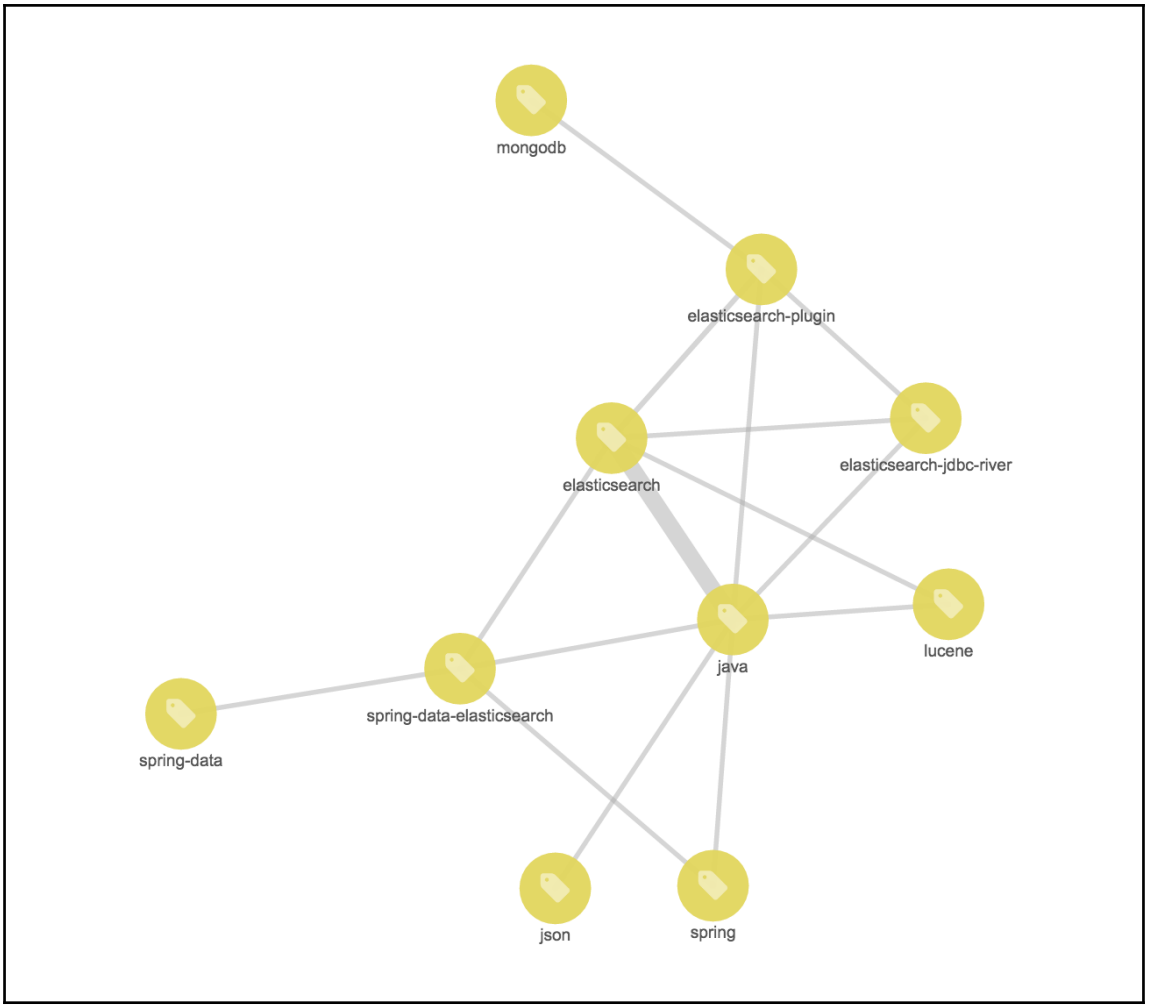
mongodb
elasticsearch-plugin
elasticsearch
elasticsearch-jdbc-river
lucene
java
spring-data-elasticsearch
spring-data
json
spring

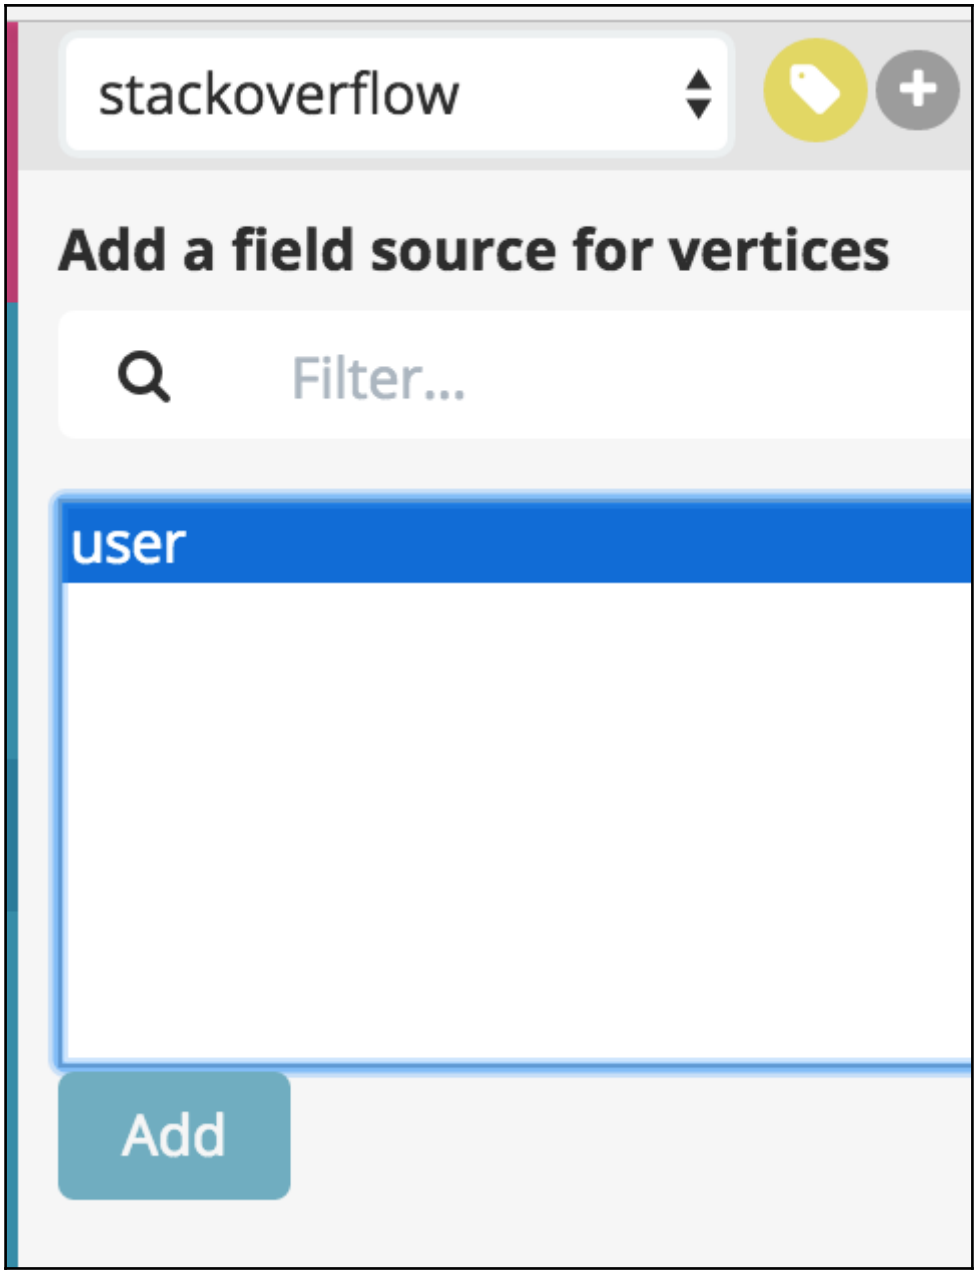
stackoverflow
Add a field source for vertices
Filter...
user
Add

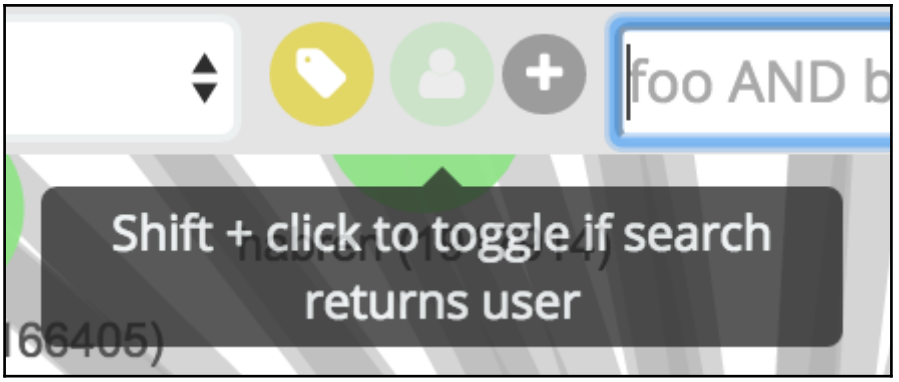
foo AND b
Shift + click to toggle if search returns user

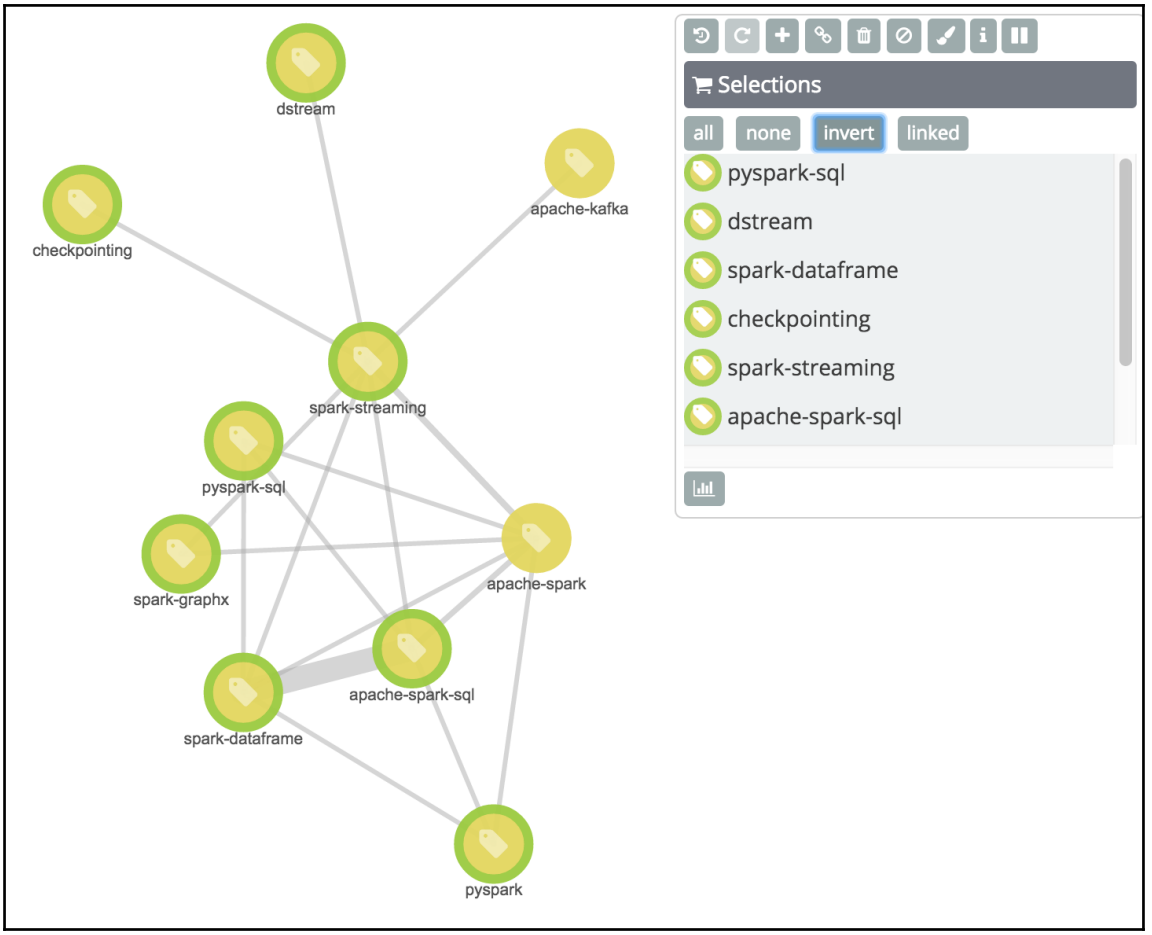
dstream
apache-kafka
checkpointing
spark-streaming
pyspark-sql
spark-graphx
apache-spark
apache-spark-sql
spark-dataframe
pyspark
Selections
all
none
invert
linked
pyspark-sql
dstream
spark-dataframe
checkpointing
spark-streaming
apache-spark-sql

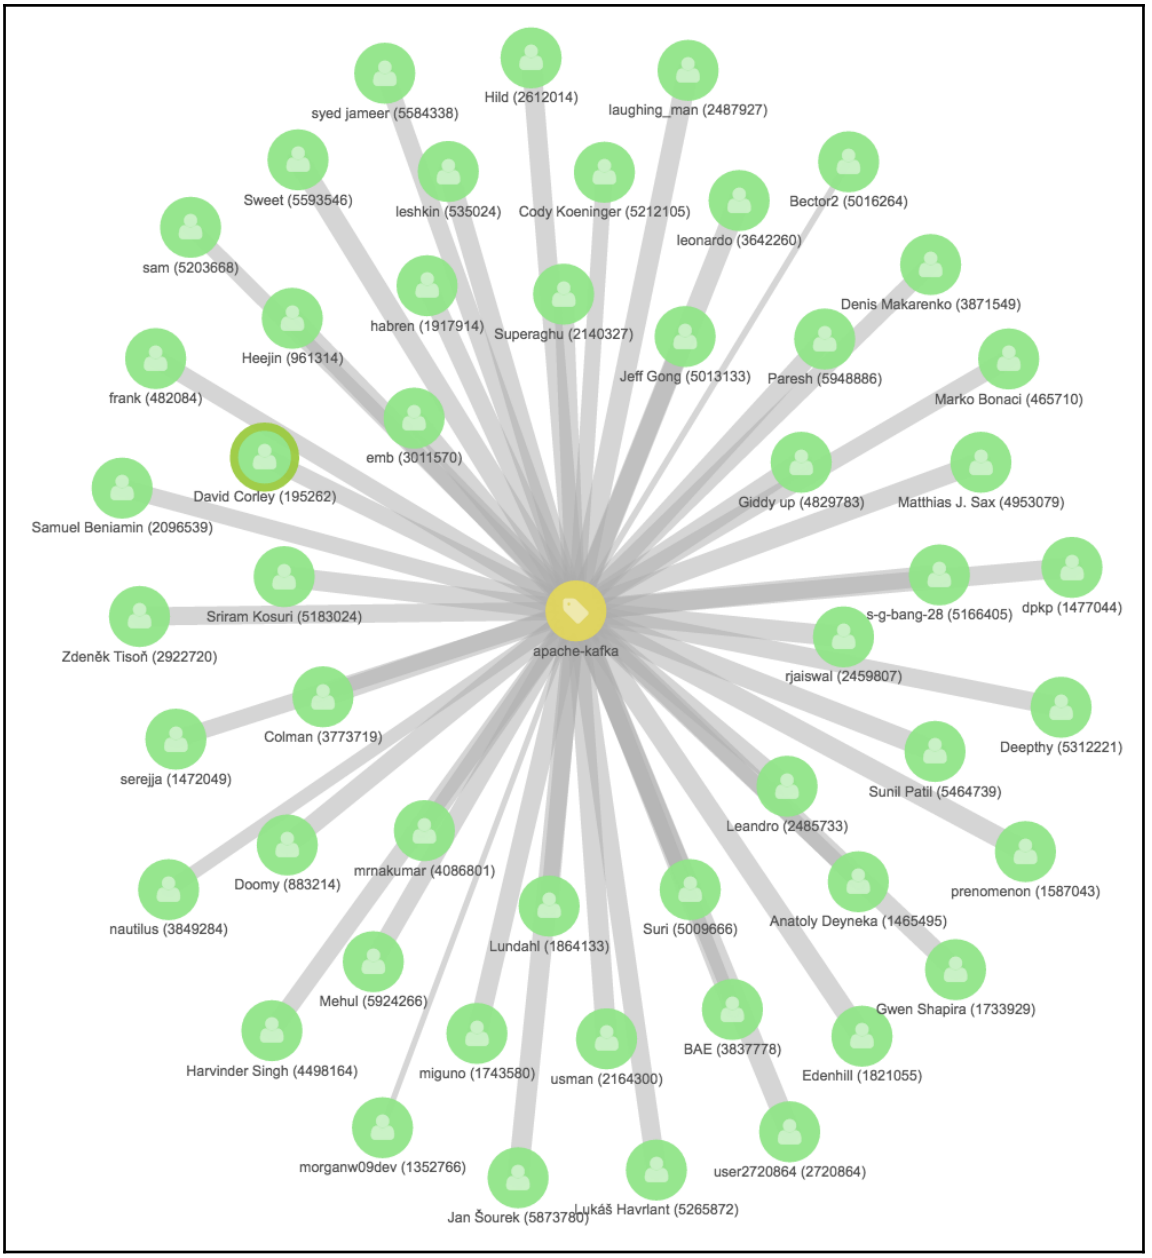
Hild (2612014)
syed jameer (5584338)
laughing_man (2487927)
Sweet (5593546)
leshkin (535024)
Cody Koeninger (5212105)
Bector2 (5016264)
leonardo (3642260)
sam (5203668)
Denis Makarenko (3871549)
habren (1917914)
Superaghu (2140327)
Heejin (961314)
Jeff Gong (5013133)
Paresh (5948886)
frank (482084)
Marko Bonaci (465710)
emb (3011570)
David Corley (195262)
Giddy up (4829783)
Matthias J. Sax (4953079)
Samuel Beniamin (2096539)
s-g-bang-28 (5166405)
dpkp (1477044)
Sriram Kosuri (5183024)
Zdeněk Tisoň (2922720)
apache-kafka
rjaiswal (2459807)
Colman (3773719)
Deepthy (5312221)
serejja (1472049)
Sunil Patil (5464739)
Leandro (2485733)
mrnakumar (4086801)
Doomy (883214)
prenomenon (1587043)
nautilus (3849284)
Anatoly Deyneka (1465495)
Suri (5009666)
Lundahl (1864133)
Mehul (5924266)
Gwen Shapira (1733929)
BAE (3837778)
Harvinder Singh (4498164)
miguno (1743580)
usman (2164300)
Edenhill (1821055)
morganw09dev (1352766)
user2720864 (2720864)
Jan Šourek (5873780)
Lukáš Havrlant (5265872)

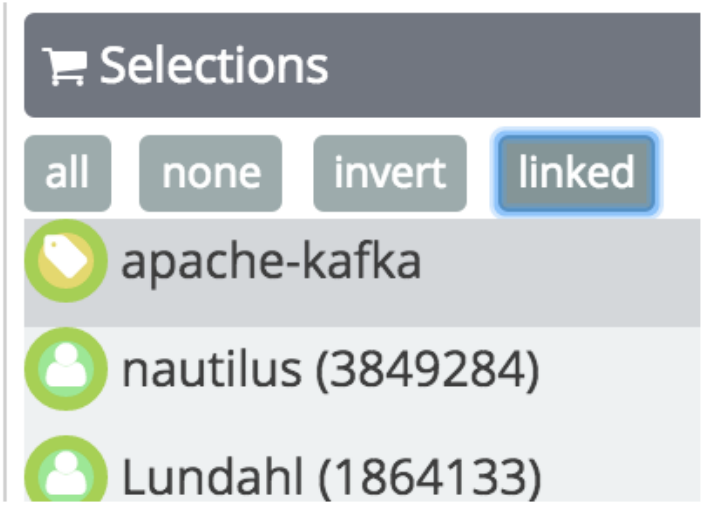
Selections
all
none
invert
linked
apache-kafka
nautilus (3849284)
Lundahl (1864133)

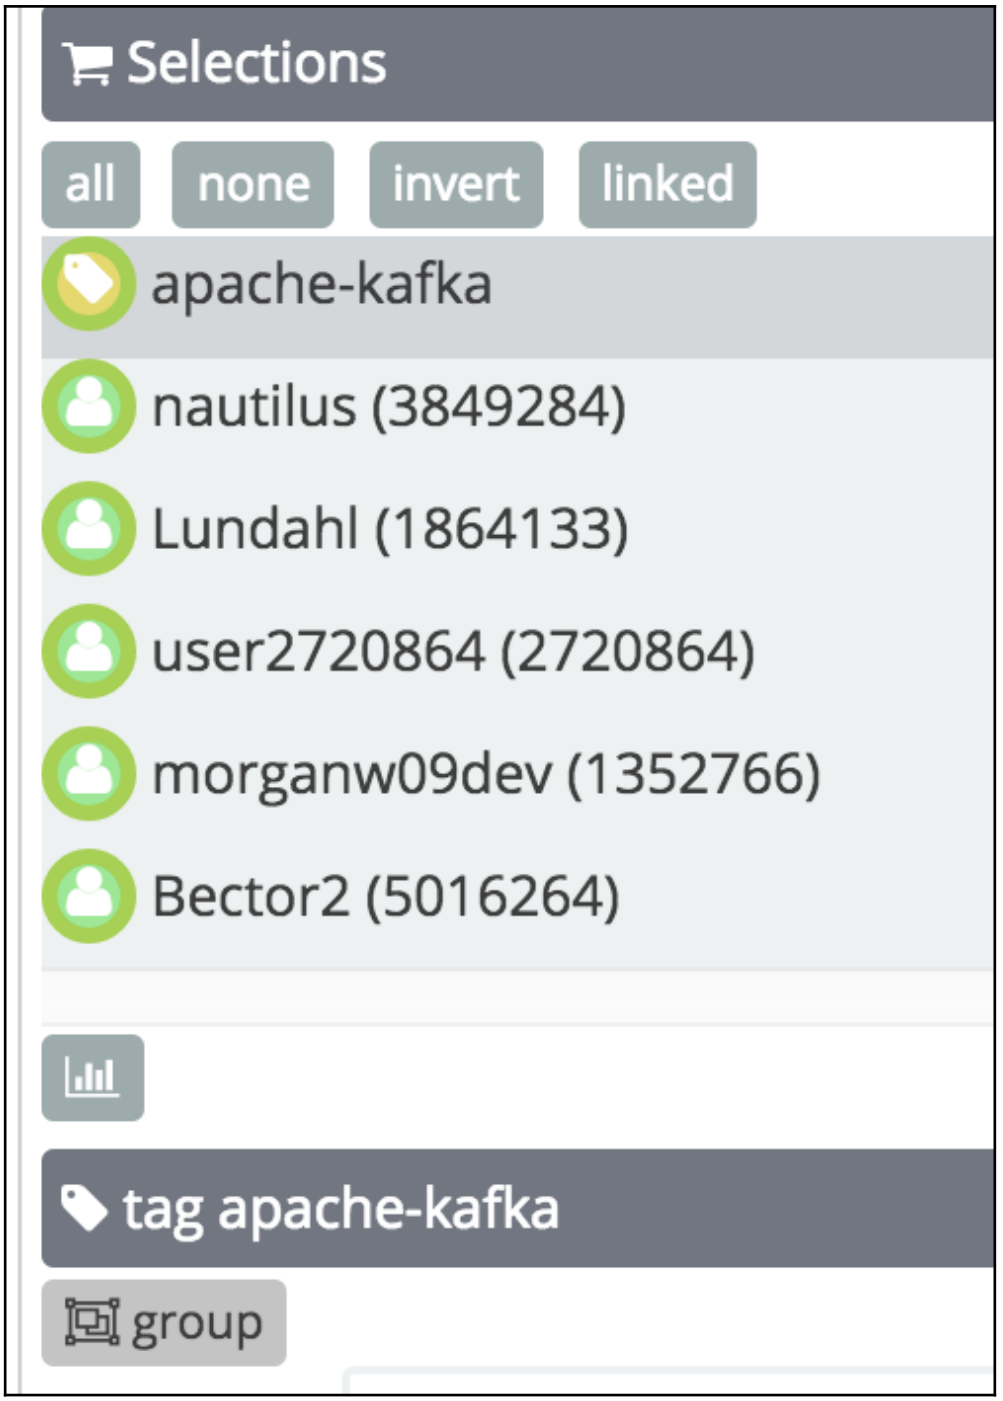
Selections
all
none
invert
linked
apache-kafka
nautilus (3849284)
Lundahl (1864133)
user2720864 (2720864)
morganw09dev (1352766)
Bector2 (5016264)
tag apache-kafka
group

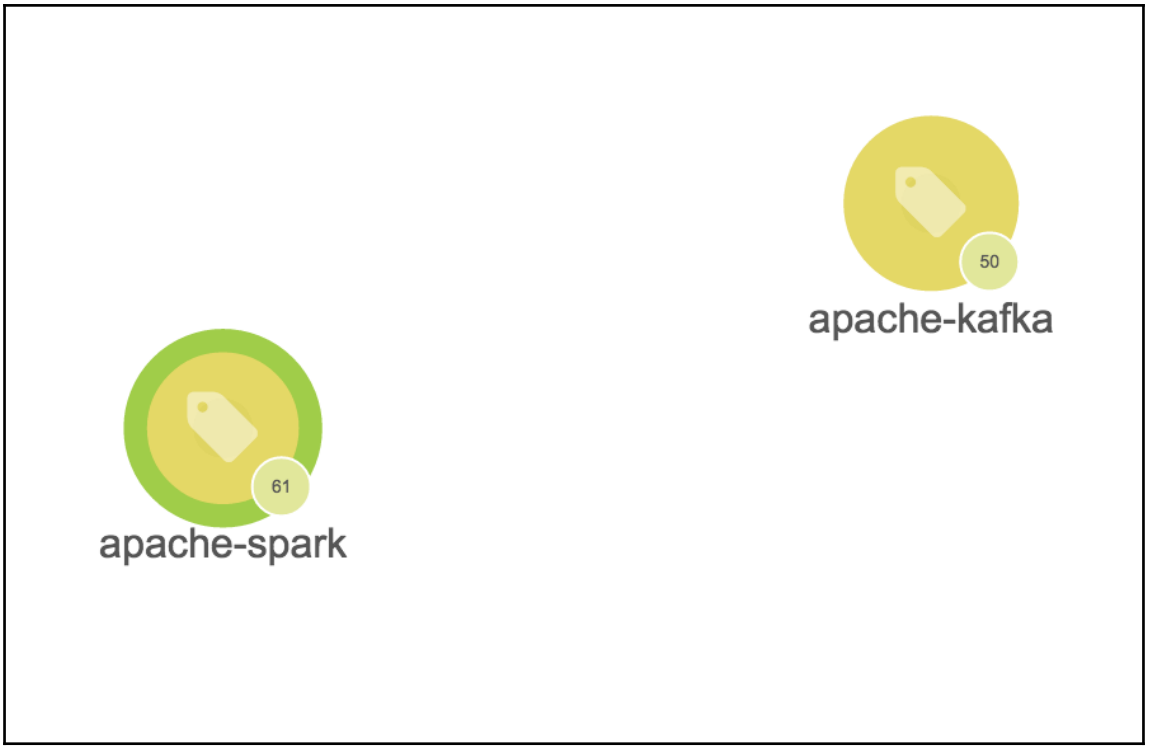
50
apache-kafka
61
apache-spark

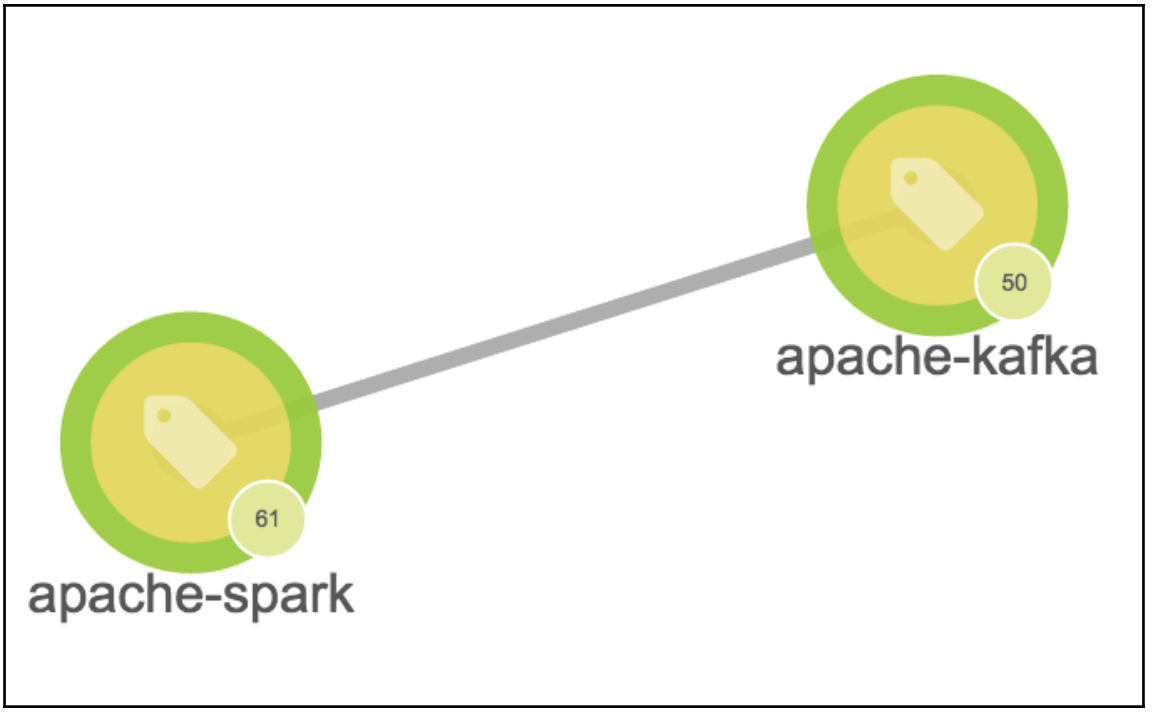
50
apache-kafka
61
apache-spark

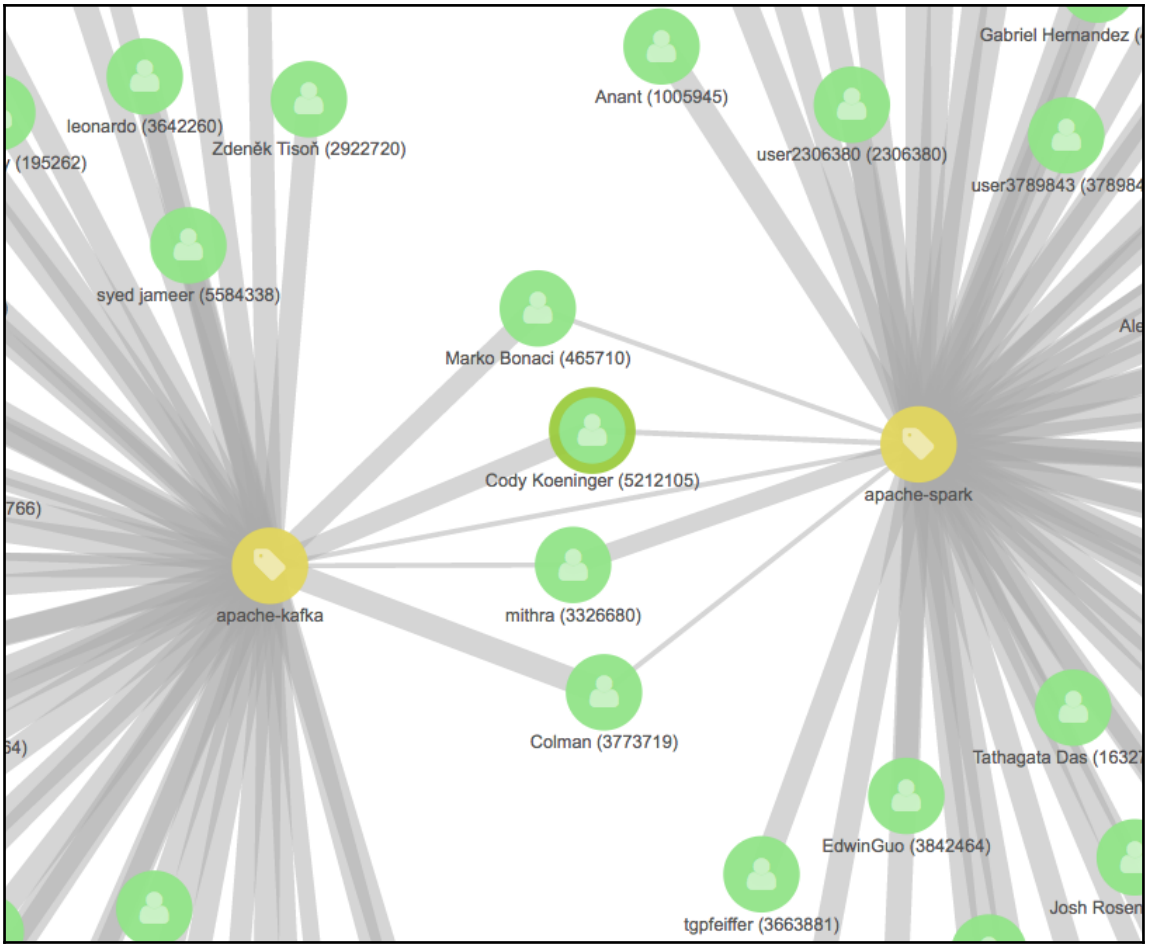
Gabriel Hernandez (
Anant (1005945)
leonardo (3642260)
Zdeněk Tisoň (2922720)
(195262)
user2306380 (2306380)
user3789843 (378984
syed jameer (5584338)
Ale
Marko Bonaci (465710)
Cody Koeninger (5212105)
apache-spark
766)
apache-kafka
mithra (3326680)
Colman (3773719)
Tathagata Das (1632
EdwinGuo (3842464)
tgpfeiffer (3663881)
Josh Rosen

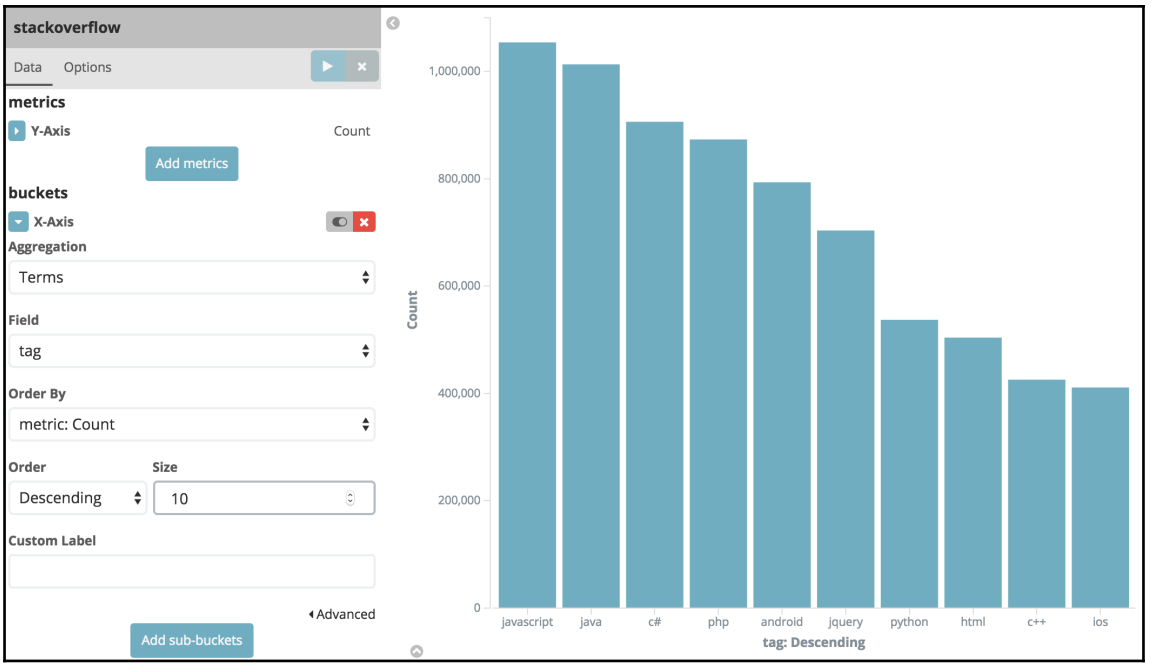
stackoverflow
Data
Options
metrics
Y-Axis
Count
Add metrics
buckets
X-Axis
Aggregation
Terms
Field
tag
Order By
metric: Count
Order
Size
Descending
10
Custom Label
Advanced
Add sub-buckets
1,000,000
800,000
600,000
400,000
200,000
0
Count
javascript
java
c#
php
android
jquery
python
html
c++
ios
tag: Descending

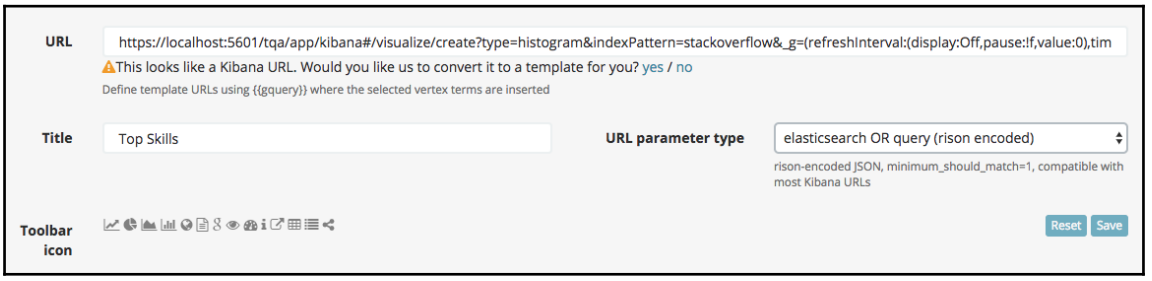
URL
https://localhost:5601/tqa/app/kibana#/visualize/create?type=histogram&indexPattern=stackoverflow&_g=(refreshInterval:(display:Off,pause:!f,value:0),tim
This looks like a Kibana URL. Would you like us to convert it to a template for you? yes / no
Define template URLs using {{gquery}} where the selected vertex terms are inserted
Title
Top Skills
URL parameter type
elasticsearch OR query (rison encoded)
rison-encoded JSON, minimum_should_match=1, compatible with most Kibana URLs
Toolbar icon
Reset
Save

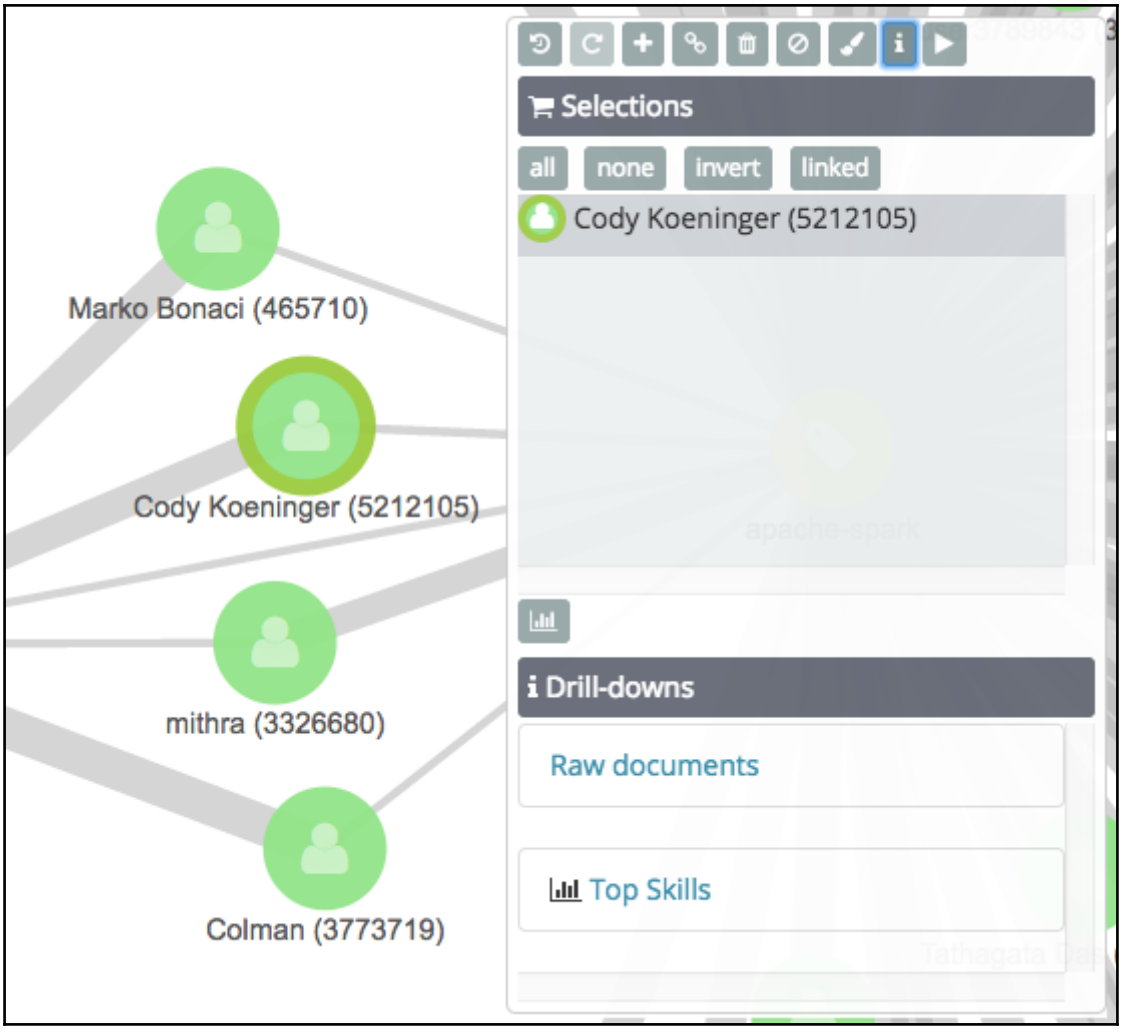
Selections
all
none
invert
linked
Cody Koeninger (5212105)
Drill-downs
Raw documents
Top Skills
Marko Bonaci (465710)
Cody Koeninger (5212105)
mithra (3326680)
Colman (3773719)

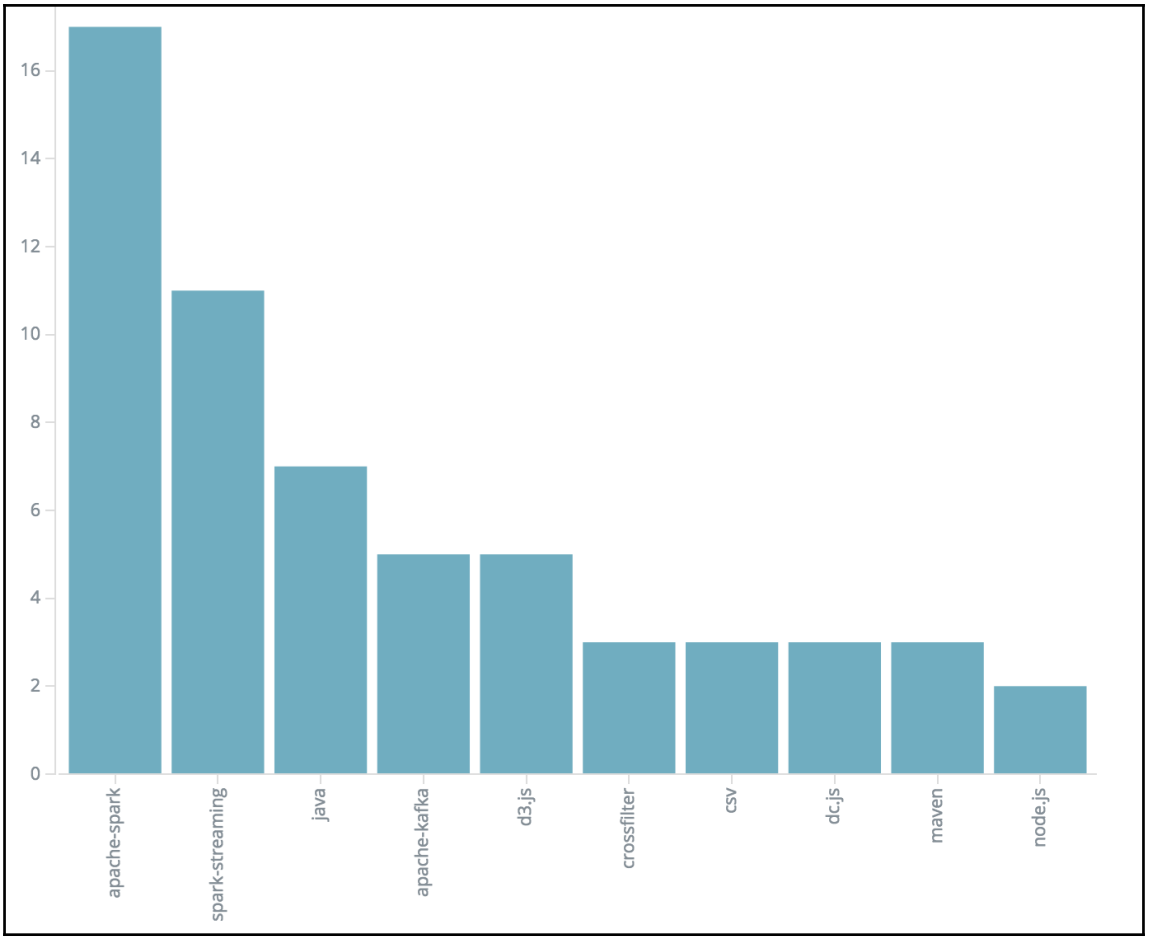
16
14
12
10
8
6
4
2
0
apache-spark
spark-streaming
java
apache-kafka
d3.js
crossfilter
csv
dc.js
maven
node.js

# Chapter 7: Customizing Kibana 5.0 Timelion

This repository Search Pull requests Issues Gist

elastic / timelion

Code Issues 26 Pull requests 0 Projects 0 Wiki Pulse

Branch: master timelion / server /

rashidkpc Fix precision in worldbank api

| | |
|---|---|
| .. | |
| fit_functions | Remove .only |
| handlers | Fix handling of multi-valued args |
| lib | Use new function API instead of interacting directly with files |
| routes | Use function API, fix issue with sorting of function reference |
| series_functions | Fix precision in worldbank api |

kibana

.es(*).multiply(2)

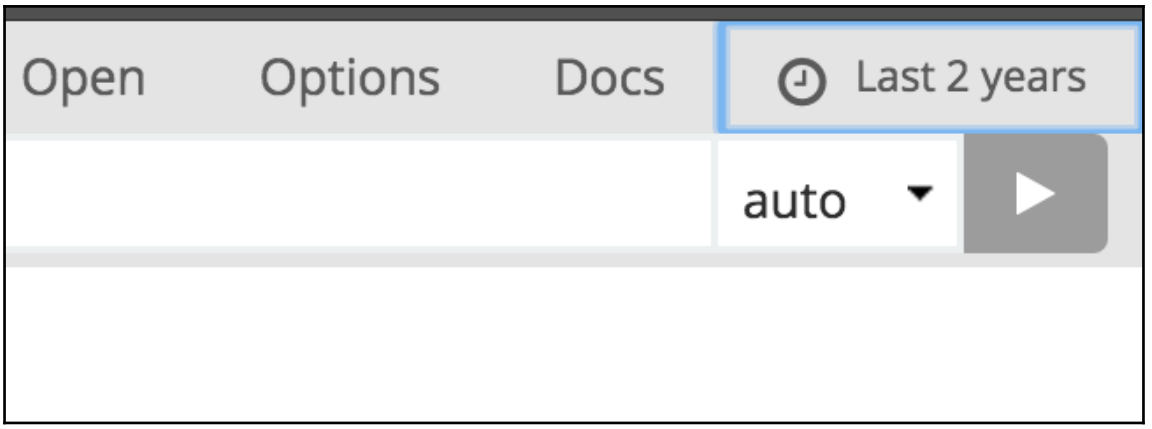


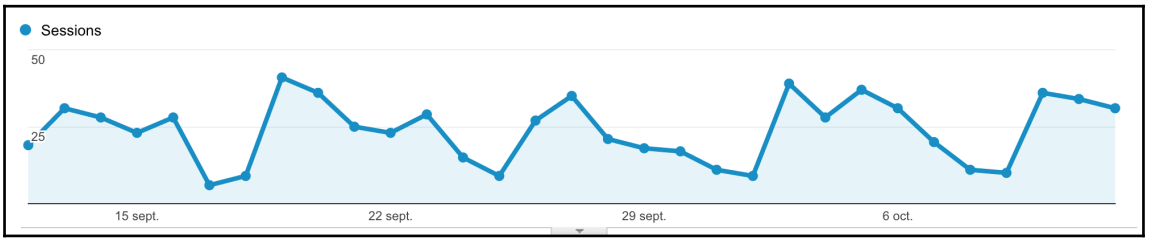


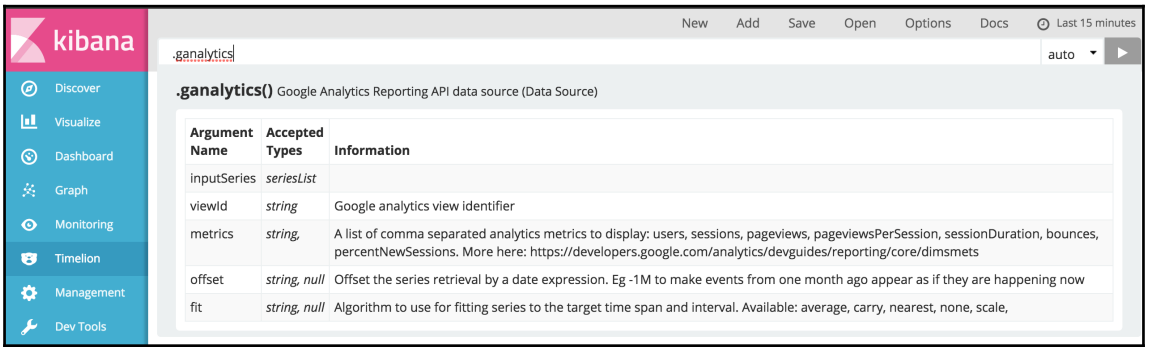


| Argument Name | Accepted Types | Information |
|---|---|---|
| inputSeries | *seriesList* | |
| viewId | *string* | Google analytics view identifier |
| metrics | *string,* | A list of comma separated analytics metrics to display: users, sessions, pageviews, pageviewsPerSession, sessionDuration, bounces, percentNewSessions. More here: https://developers.google.com/analytics/devguides/reporting/core/dimsmets |
| offset | *string, null* | Offset the series retrieval by a date expression. Eg -1M to make events from one month ago appear as if they are happening now |
| fit | *string, null* | Algorithm to use for fitting series to the target time span and interval. Available: average, carry, nearest, none, scale, |

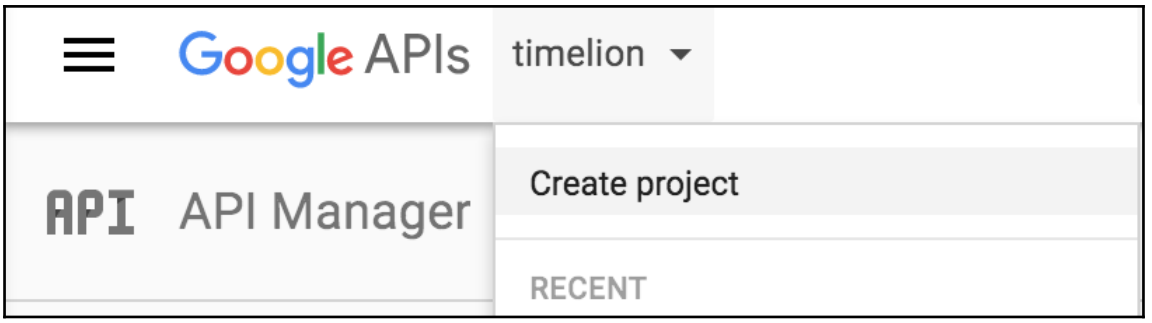
Google APIs
timelion
API
API Manager
Create project
RECENT

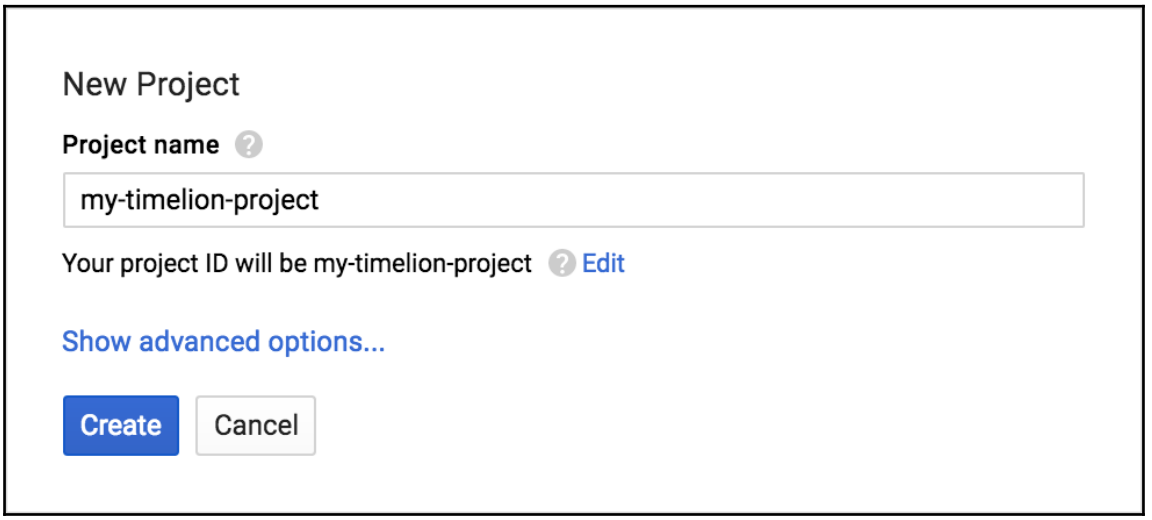
New Project
Project name
my-timelion-project
Your project ID will be my-timelion-project
Edit
Show advanced options...
Create
Cancel

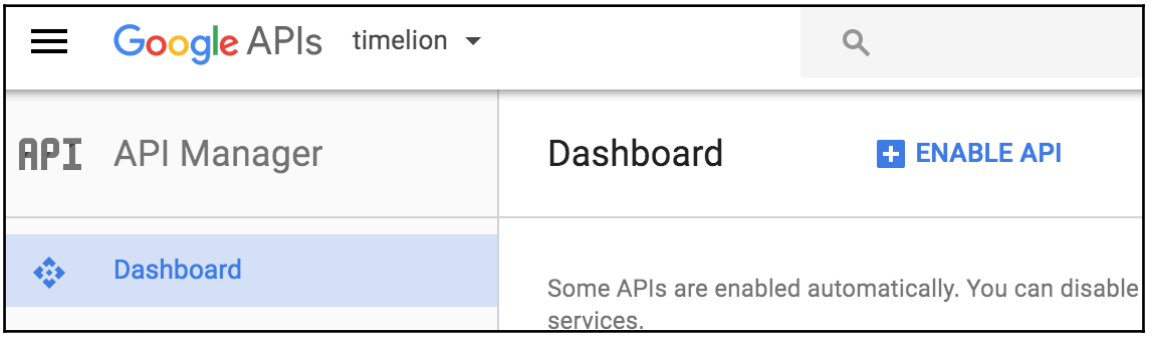
Google APIs
timelion
API
API Manager
Dashboard
ENABLE API
Dashboard
Some APIs are enabled automatically. You can disable
services.

API Manager

Library

Dashboard

Library

Credentials

Google APIs

google analytics reporting api

Back to popular APIs

| Name | Description |
| --- | --- |
| Google Analytics Reporting API | Accesses **Analytics** report data. |

API Manager

Google Analytics Reporting API

ENABLE

IAM & Admin

Service Accounts

CREATE SERVICE ACCOUNT

Create service account

**Service account name**

my-service-account

**Role**

Owner

**Service account ID**

my-service-account

☑ Furnish a new private key

Downloads a file that contains the private key. Store the file securely because this key can't be recovered if lost.

**Key type**

◉ JSON

Recommended

○ P12

For backward compatibility with code using the P12 format

☐ Enable G Suite Domain-wide Delegation

Grants a client access to all users' data on a G Suite domain without manual authorization on their part. Learn more

Create Cancel

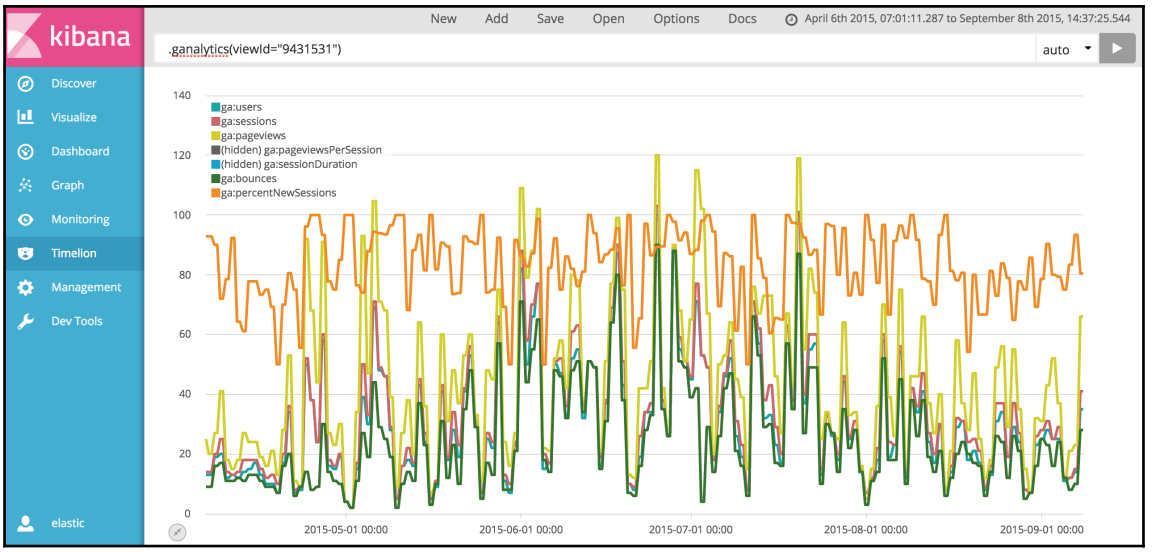
kibana
New
Add
Save
Open
Options
Docs
April 6th 2015, 07:01:11.287 to September 8th 2015, 14:37:25.544
.ganalytics(viewId="9431531")
auto
Discover
Visualize
Dashboard
Graph
Monitoring
Timelion
Management
Dev Tools
elastic
ga:users
ga:sessions
ga:pageviews
(hidden) ga:pageviewsPerSession
(hidden) ga:sessionDuration
ga:bounces
ga:percentNewSessions
140
120
100
80
60
40
20
0
2015-05-01 00:00
2015-06-01 00:00
2015-07-01 00:00
2015-08-01 00:00
2015-09-01 00:00

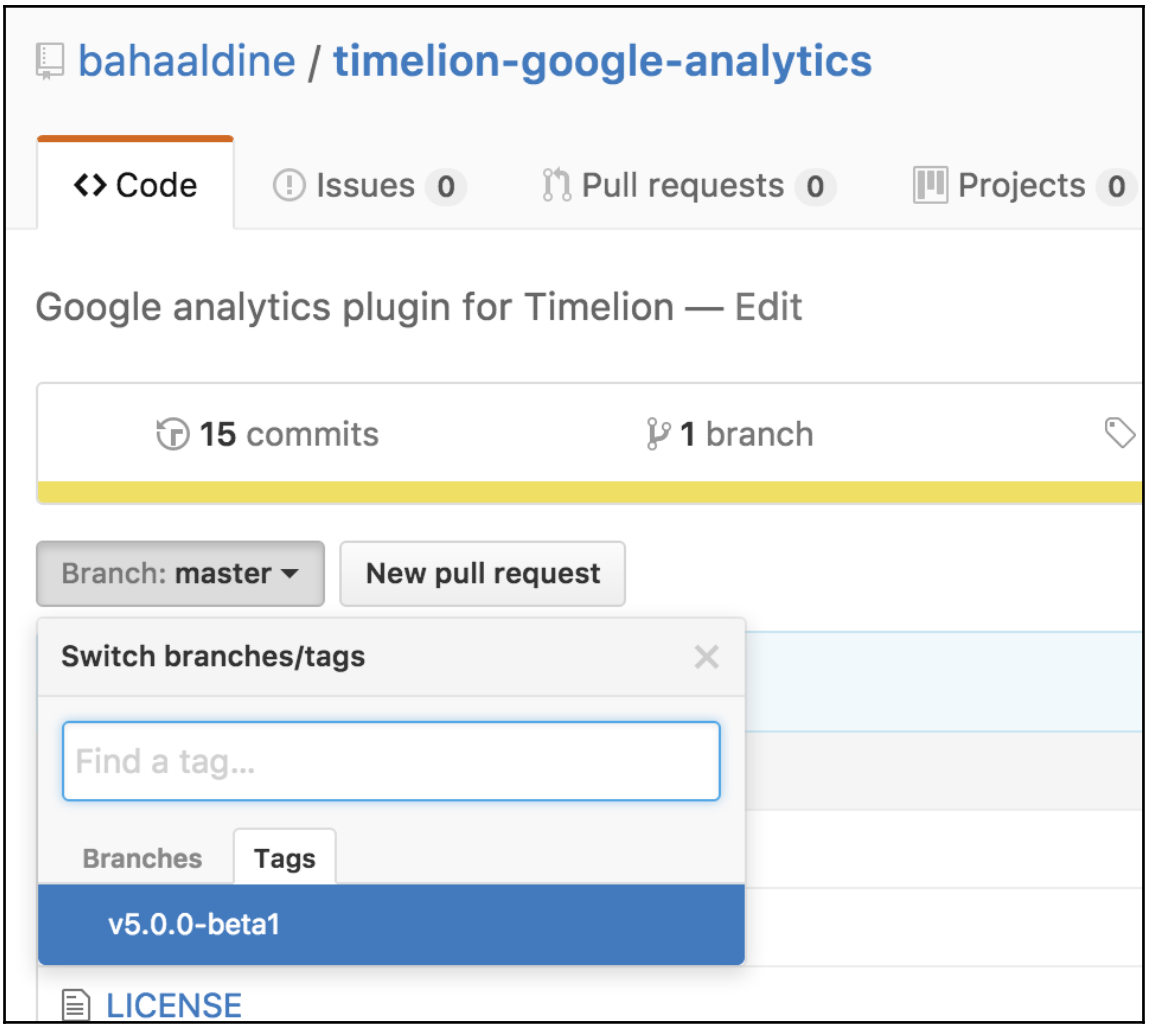
bahaaldine / timelion-google-analytics
Code
Issues 0
Pull requests 0
Projects 0
Google analytics plugin for Timelion — Edit
15 commits
1 branch
Branch: master
New pull request
Switch branches/tags
Find a tag...
Branches
Tags
v5.0.0-beta1
LICENSE

bahaaldine / timelion-google-analytics

Unwatch 1 | Star 1 | Fork 0

Code | Issues 0 | Pull requests 0 | Projects 0 | Wiki | Pulse | Graphs | Settings

Releases | Tags

Draft a new release

Latest release

v5.0.0-beta1

0b0f3af

# v5.0.0-beta1

Edit

**bahaaldine** released this 2 days ago · **11 commits** to master since this release

```
TGA for KB v5.0.0-beta1
```

## Downloads

Source code (zip)

Source code (tar.gz)

# Chapter 8: Anomaly Detection in Kibana 5.0


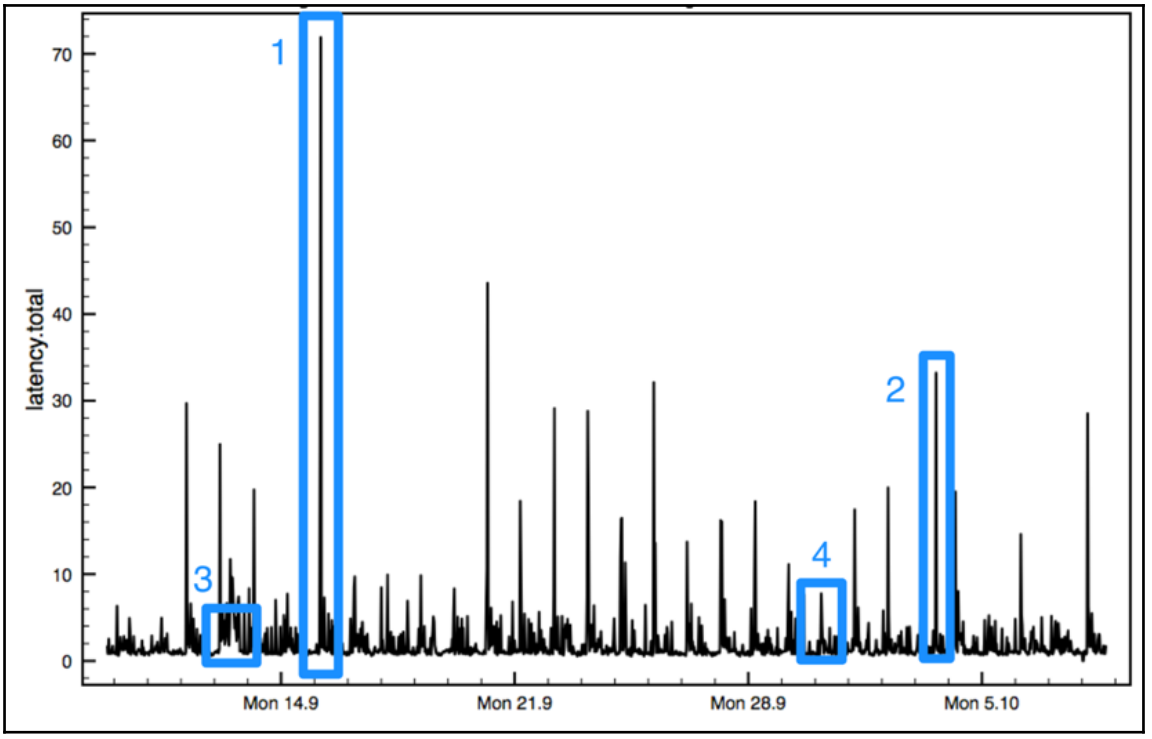

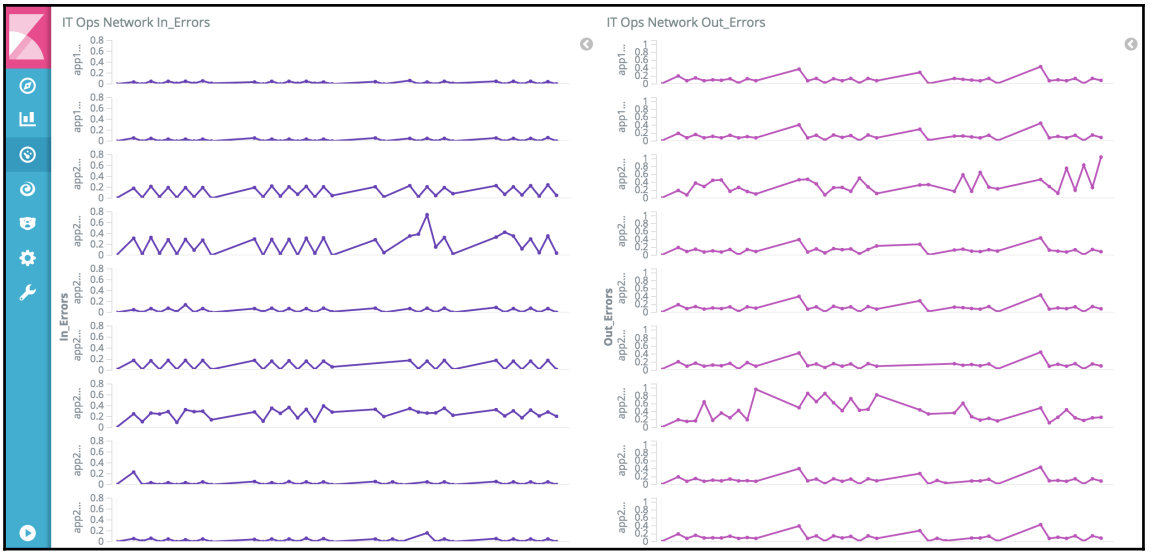
IT Ops Network In_Errors
IT Ops Network Out_Errors
In_Errors
Out_Errors

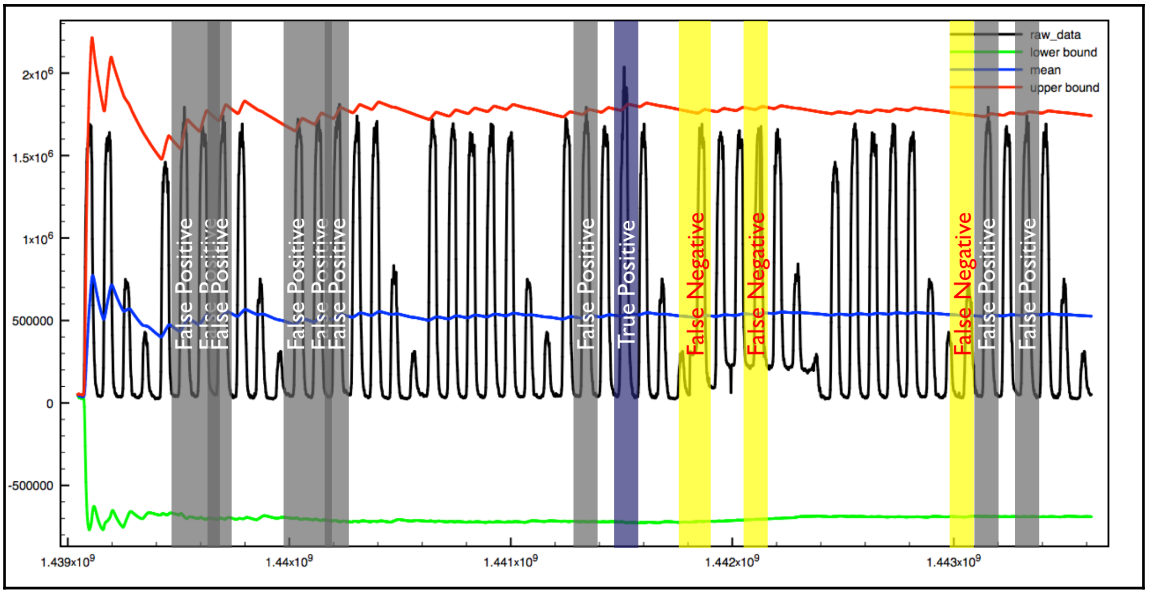
raw_data
lower bound
mean
upper bound
2x10^6
1.5x10^6
1x10^6
500000
0
-500000
1.439x10^9
1.44x10^9
1.441x10^9
1.442x10^9
1.443x10^9
False Positive
True Positive
False Negative

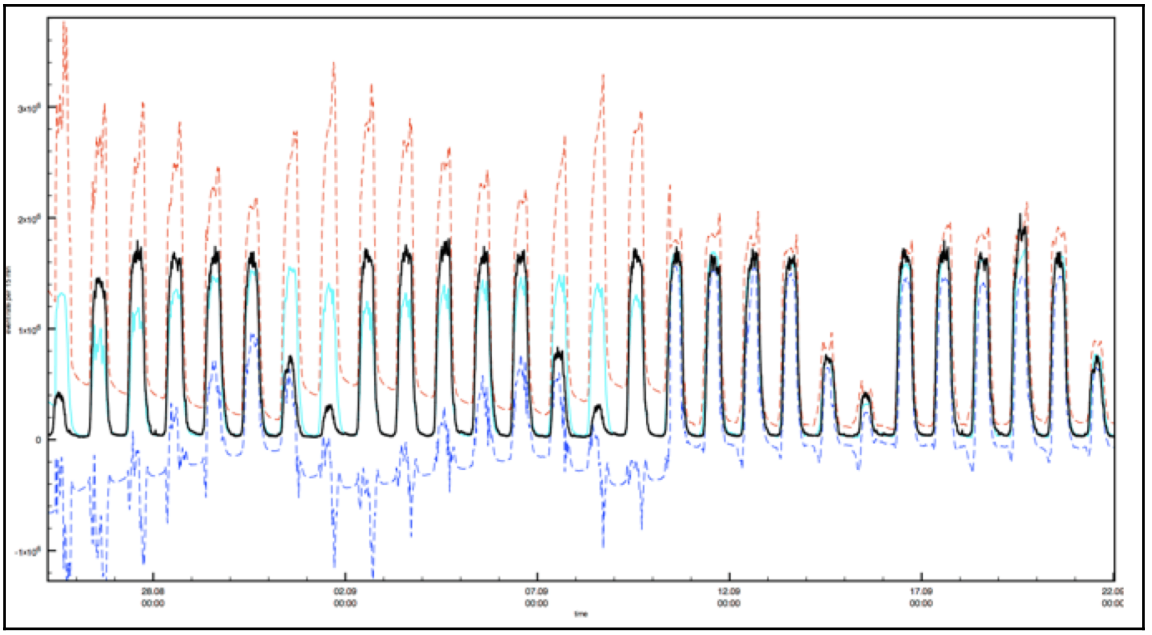
3x10^6
2x10^6
1x10^6
0
-1x10^6
28.08
00:00
02.09
00:00
07.09
00:00
12.09
00:00
17.09
00:00
time

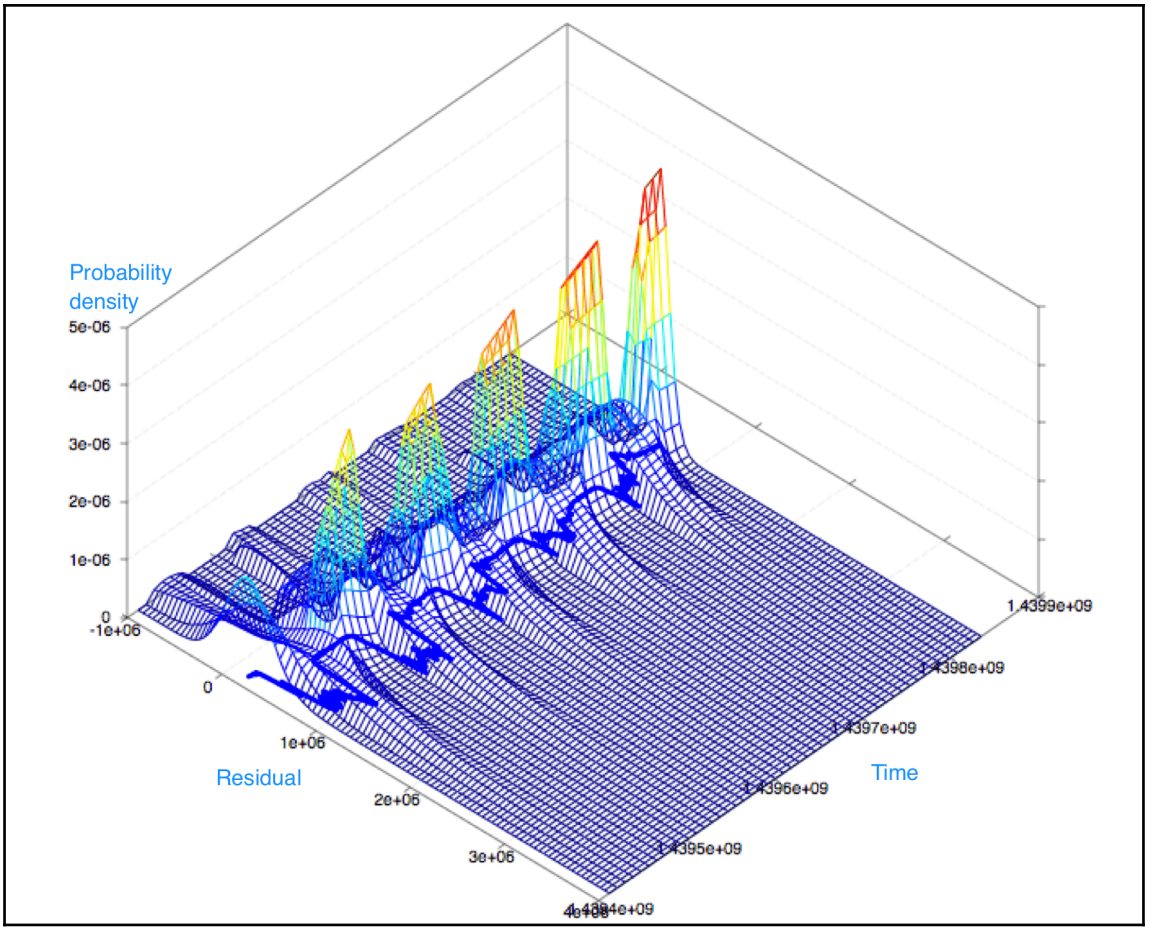
Probability density
5e-06
4e-06
3e-06
2e-06
1e-06
0
-1e+06
0
1e+06
2e+06
3e+06
4e+06
Residual
1.4399e+09
1.4398e+09
1.4397e+09
1.4396e+09
1.4395e+09
1.4394e+09
Time

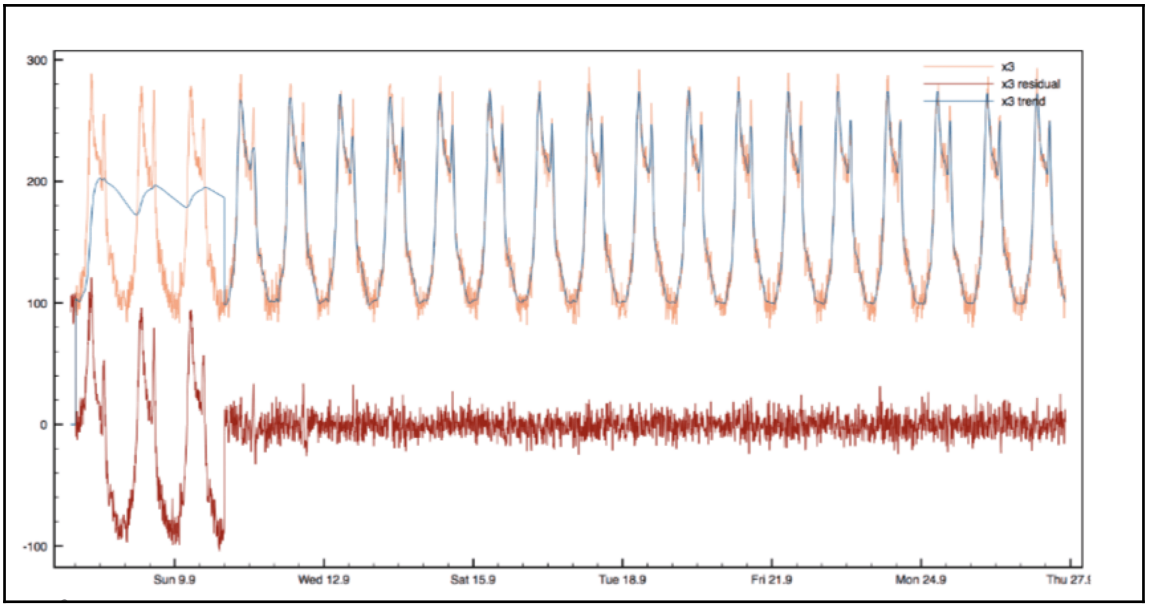
300
200
100
0
-100
x3
x3 residual
x3 trend
Sun 9.9
Wed 12.9
Sat 15.9
Tue 18.9
Fri 21.9
Mon 24.9

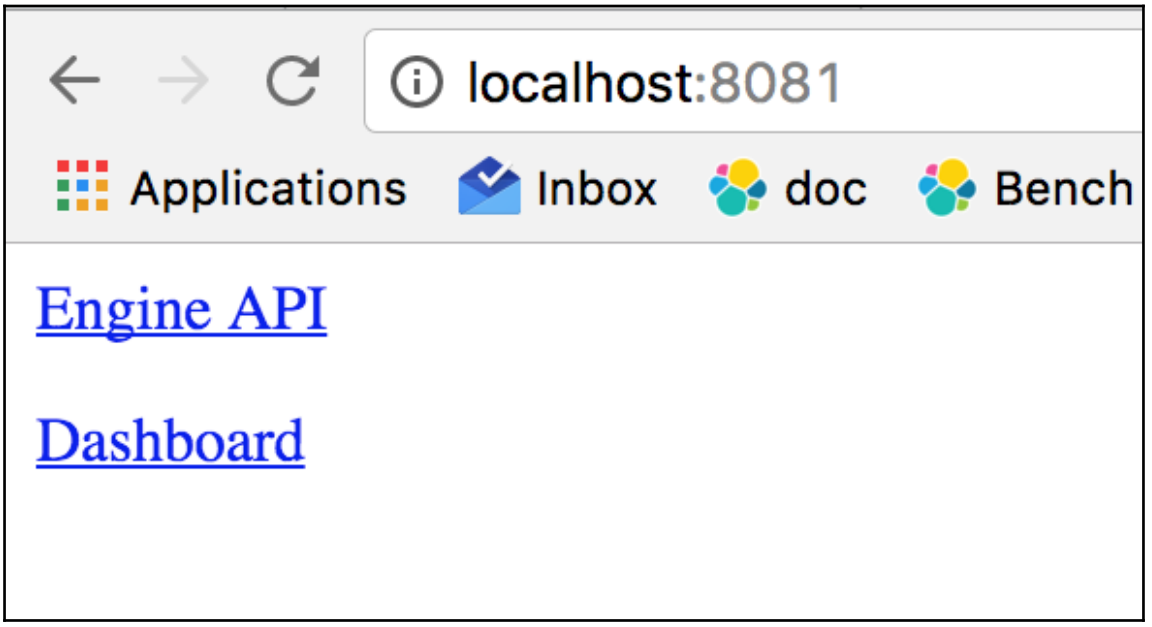
localhost:8081
Applications
Inbox
doc
Bench
Engine API
Dashboard

# Prelert Engine API 2.1.2

## Analytics Version:

Model State Version 31
Quantile State Version 3
prelert_autodetect_api (64 bit): Version 6.3.1 (Build 77be586fe61a46) Copyright (c) Prelert Ltd 2006-2016

| | |
|---|---|
| Hostname | MacBook-Pro-de-Bahaaldine.local |
| OS Name | Mac OS X |
| OS Version | 10.12 |
| Total Memory Size MB | 16384 |
| Total Disk MB | 475928 |
| Available Disk MB | 100333 |

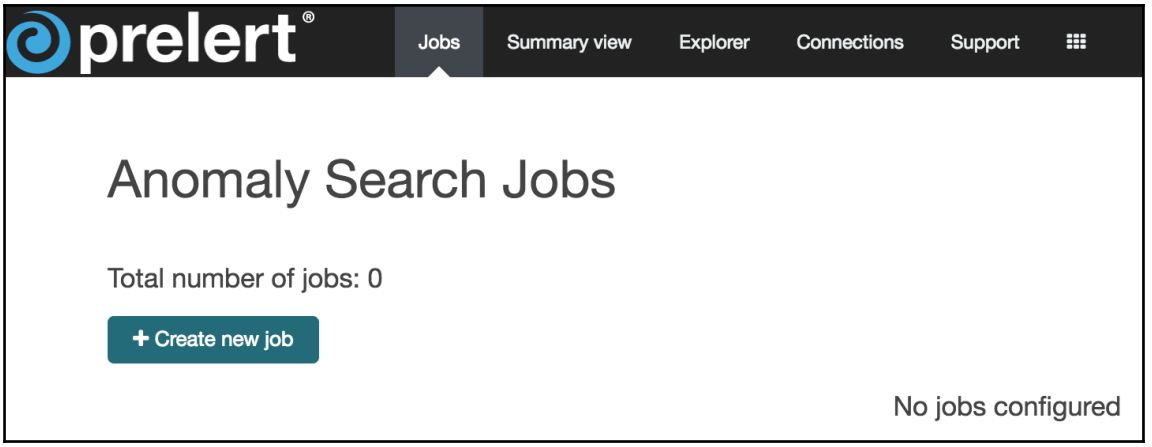

# Create a new job

**Choose a data source**

## Elasticsearch server

*Specify the address of an Elasticsearch server.*



## File upload

*Upload a file containing a data set. Maximum size is 100MB.*



## Other data source

*Create a job without reference to source data. Specify fields manually and upload data later using the API.*

# Create a new job

**Elasticsearch server address**

http://localhost:9200

☐ **Authenticated**

◉ **Input index**

○ **Choose index from list**

**Index**

metricbeat*

**Types**

☐ All types

☐ _default_

☑ metricsets

**Time-field name**

@timestamp

**Time format**

yyyy-MM-dd'T'HH:mm:ss.SSSX

*e.g. 2016-03-24'T'16:20:24.611+0100*

Back Next

# Create a new job

Job Details | Transforms | Analysis Configuration | Data Description | Scheduler | Edit JSON | Data Preview

**Name**

cpu-anomaly-detection

**Description**

Using Prelert to analyze CPU data generated by Metricbeat

**Custom URLs**

+ Add Custom URL

Save | Cancel

## Add new detector

**Description**

Process CPU detector

| function | fieldName | byFieldName |
|---|---|---|
| metric | system.process.cpu.total.pct | system.process.name |

| overFieldName | partitionFieldName | excludeFrequent |
|---|---|---|
| | | |

*Help for metric*

Add | Cancel

# Create a new job

Job Details | Transforms | Analysis Configuration | Data Description | Scheduler | Edit JSON | Data Preview

☑ **Scheduled job**

**Data source**

Elasticsearch 2.x, 5.0

**Query**

{"match_all":{}}

**Query delay (seconds)**

60

**Frequency (seconds)**

150

**scrollSize**

1000

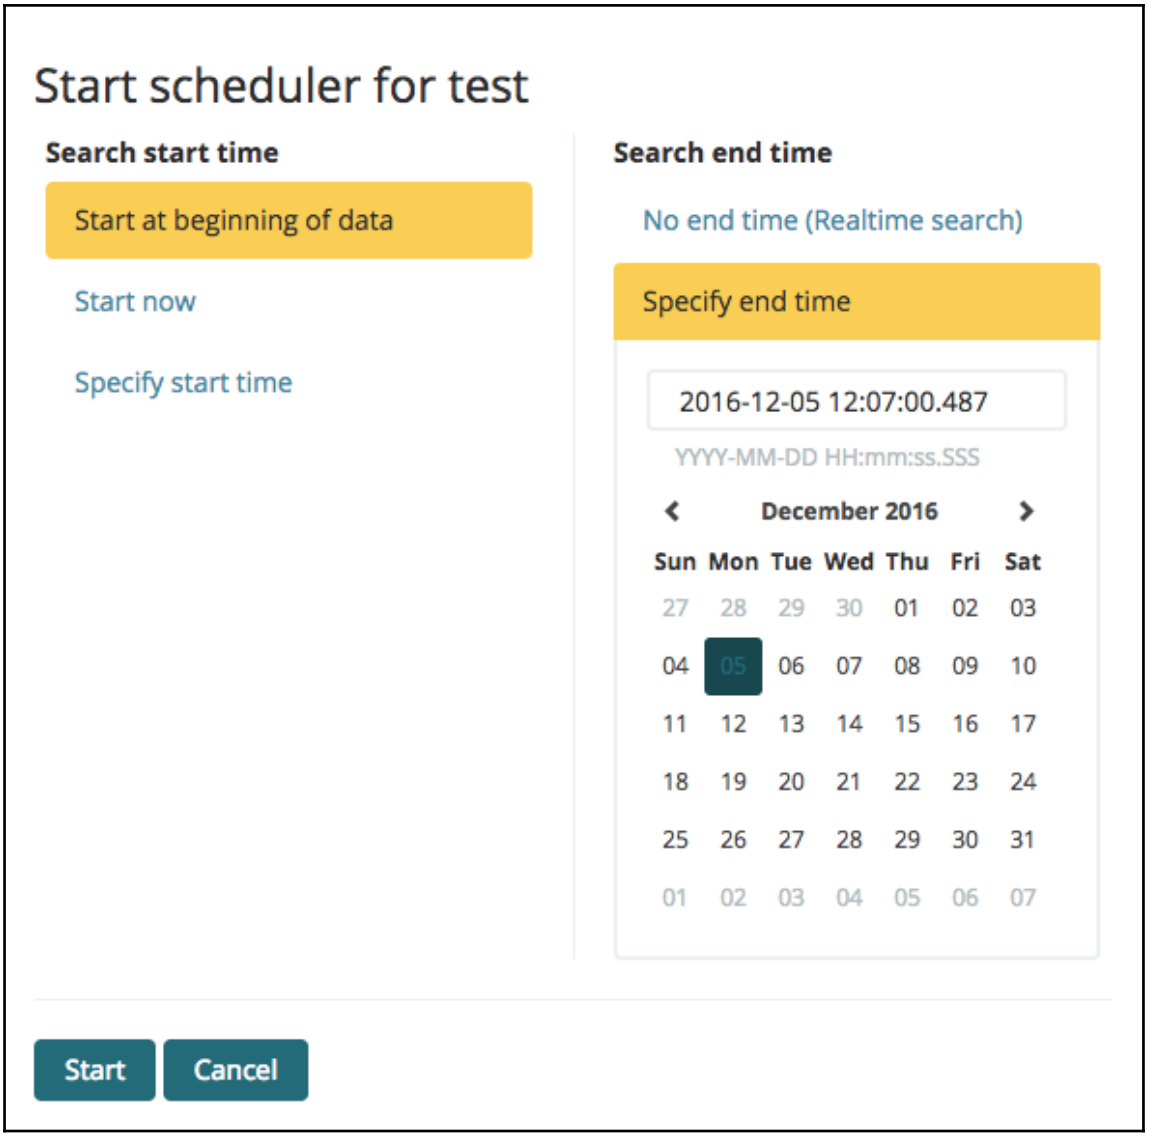


| Search name | Description | Processed records | Memory status | Job status | Scheduler status | Latest timestamp | Actions |
|---|---|---|---|---|---|---|---|
| cpu-anomaly-detection | Using Prelert to analyze CPU data generated by Metricbeat | 117,991 | OK | RUNNING | STARTED | 2016-11-13T12:17:27+0000 | |

timestamp Actions

Open results for cpu-anomaly-detection

Page Size 10

Nov 13 19:30 Nov 13 21:00 Nov 13 22:30

Nov 13 22:30

Highest anomaly per detector

50 Using Prelert to analyze CPU data generated by

November 13th 2016, 22:30:00

Process CPU detector

93 **critical anomaly in *Process CPU detector***
**for system.process.name com.apple.WebKi**
metric(system.process.cpu.total.pct)
actual: 0.194
typical: 0.000760183

88 **critical anomaly in *Process CPU detector***
**for system.process.name EvernoteHelper**
metric(system.process.cpu.total.pct)
actual: 0.0038
typical: 0.000140552

63 **major anomaly in *Process CPU detector***
**for system.process.name loginwindow**
metric(system.process.cpu.total.pct)
actual: 0.0193
typical: 0.0022705

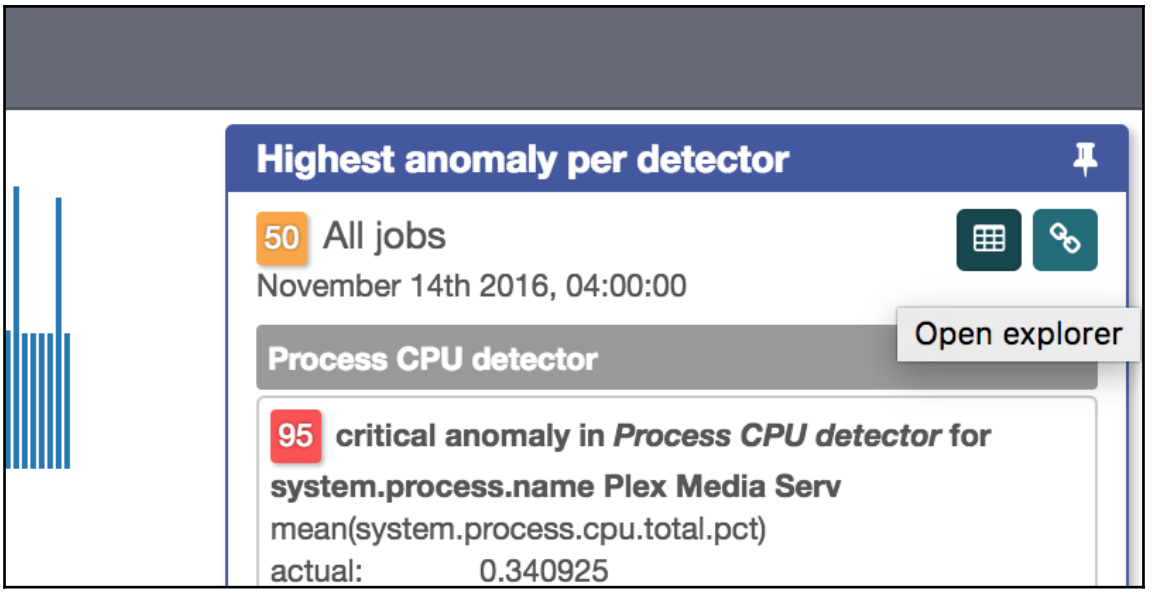


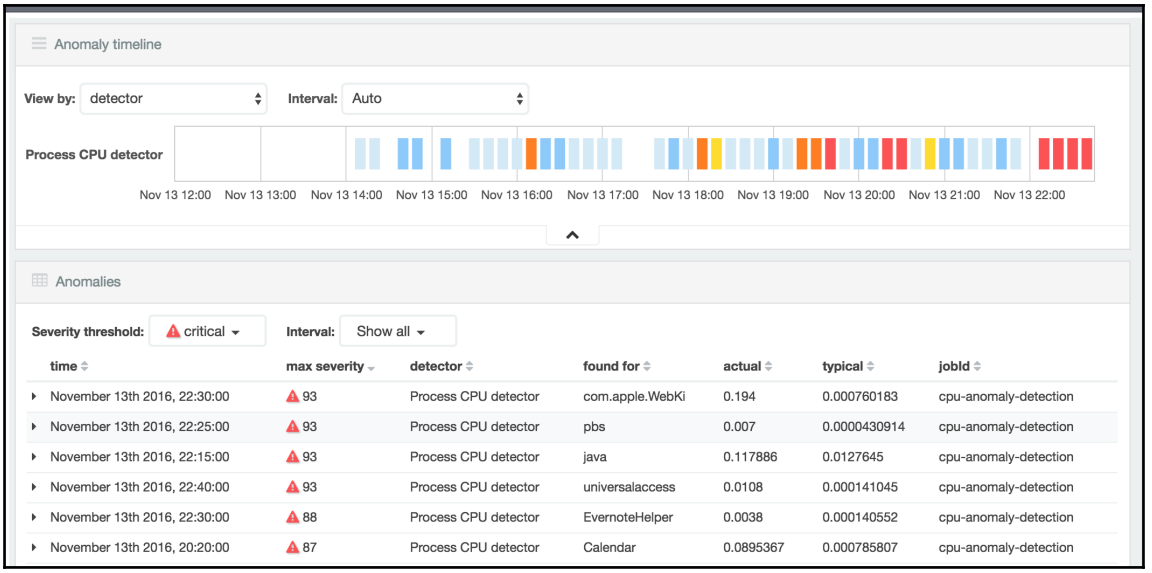


| time | max severity | detector | found for | actual | typical | jobId |
|---|---|---|---|---|---|---|
| November 13th 2016, 22:30:00 | 93 | Process CPU detector | com.apple.WebKi | 0.194 | 0.000760183 | cpu-anomaly-detection |
| November 13th 2016, 22:25:00 | 93 | Process CPU detector | pbs | 0.007 | 0.0000430914 | cpu-anomaly-detection |
| November 13th 2016, 22:15:00 | 93 | Process CPU detector | java | 0.117886 | 0.0127645 | cpu-anomaly-detection |
| November 13th 2016, 22:40:00 | 93 | Process CPU detector | universalaccess | 0.0108 | 0.000141045 | cpu-anomaly-detection |
| November 13th 2016, 22:30:00 | 88 | Process CPU detector | EvernoteHelper | 0.0038 | 0.000140552 | cpu-anomaly-detection |
| November 13th 2016, 20:20:00 | 87 | Process CPU detector | Calendar | 0.0895367 | 0.000785807 | cpu-anomaly-detection |

Anomalies

**Severity threshold:** critical **Interval:** Show all

| time | max severity | detector |
|---|---|---|
| November 13th 2016, 22:30:00 | 93 | Process CPU detector |

**Description:**
critical anomaly in Process CPU detector found for system.process.name com.apple.WebKi

**Anomaly Details:**

| | |
|---|---|
| system.process.name: | com.apple.WebKi |
| time: | November 13th 2016, 22:30:00 to November 13th 2016, 22:35:00 |
| function: | metric |
| fieldName: | system.process.cpu.total.pct |
| actual: | 0.194 |
| typical: | 0.000760183 |
| job ID: | cpu-anomaly-detection |
| probability: | 1.47031e-21 |

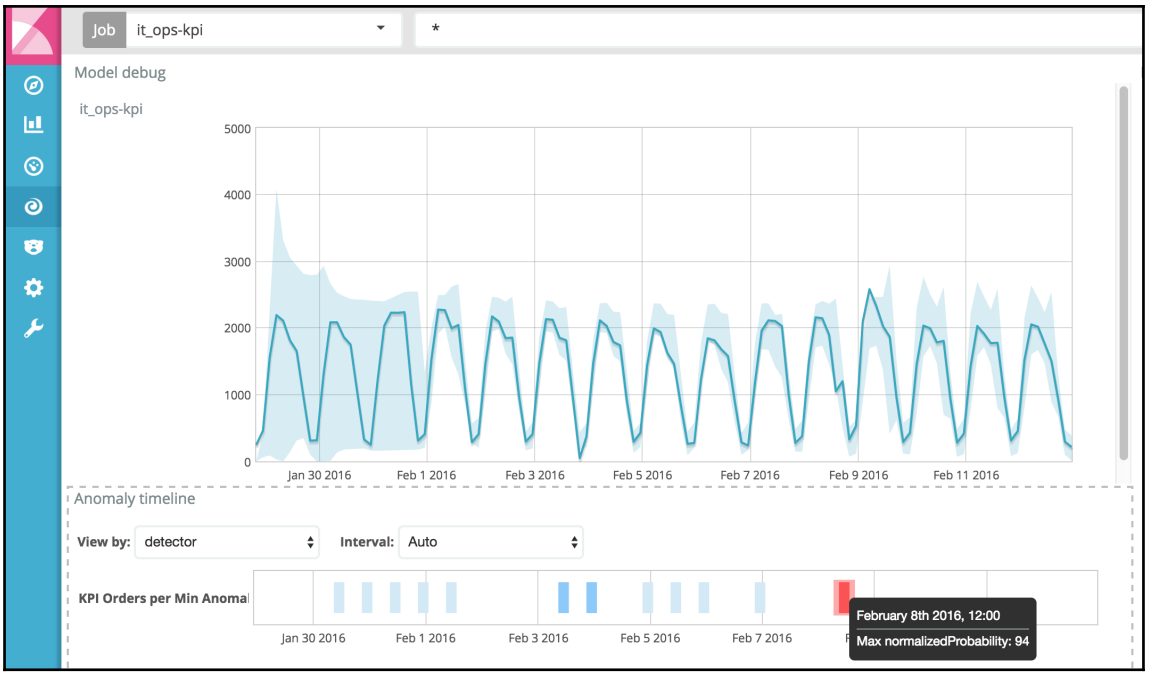
Job
it_ops-kpi
*
Model debug
it_ops-kpi
5000
4000
3000
2000
1000
0
Jan 30 2016
Feb 1 2016
Feb 3 2016
Feb 5 2016
Feb 7 2016
Feb 9 2016
Feb 11 2016
Anomaly timeline
View by:
detector
Interval:
Auto
KPI Orders per Min Anomal
February 8th 2016, 12:00
Max normalizedProbability: 94

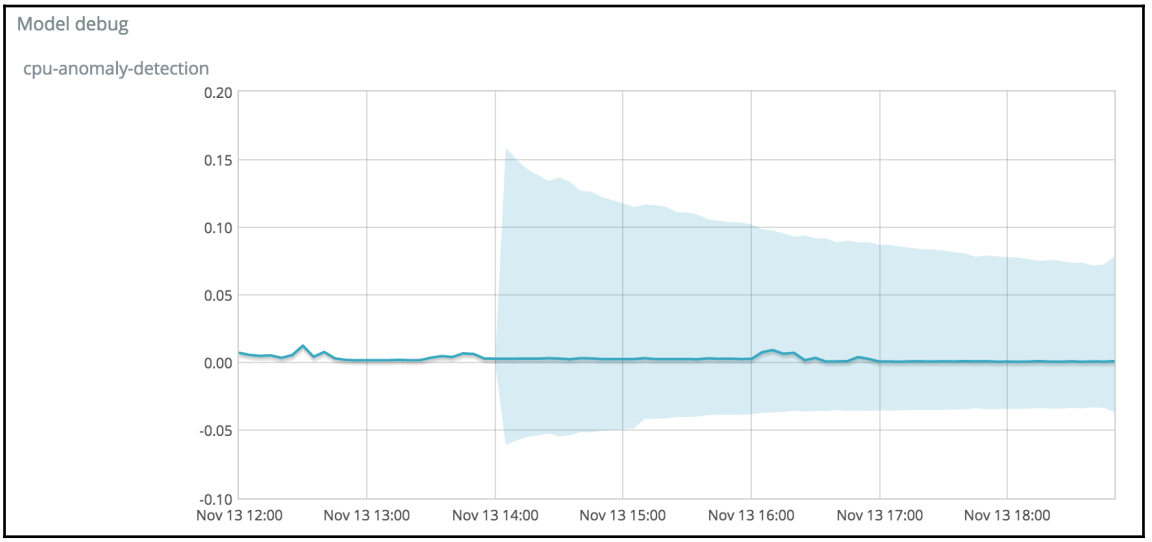
Model debug
cpu-anomaly-detection
0.20
0.15
0.10
0.05
0.00
-0.05
-0.10
Nov 13 12:00
Nov 13 13:00
Nov 13 14:00
Nov 13 15:00
Nov 13 16:00
Nov 13 17:00
Nov 13 18:00

| Search name | Description | Processed records | Memory status | Job status | Scheduler status | Latest timestamp | Actions |
|---|---|---|---|---|---|---|---|
| cpu-anomaly-detection | Using Prelert to analyze CPU data generated by Metricbeat | 950,308 | OK | CLOSED | STOPPED | 2016-11-13T21:45:27+0000 | |

Job settings | Job config | Scheduler | Counts | JSON | Job Messages

**General**

| | |
|---|---|
| location | http://localhost:8081/engine/v2/jobs/cpu-anomaly-detection |
| id | cpu-anomaly-detection |
| status | CLOSED |
| timeout | 0 |
| description | Using Prelert to analyze CPU data generated by Metricbeat |
| averageBucketPr... | 39ms |
| schedulerStatus | STOPPED |
| createTime | 2016-11-13T21:45:26.862+0000 |
| finishedTime | 2016-11-13T21:46:14.813+0000 |
| lastDataTime | 2016-11-13T21:46:13.562+0000 |

**Endpoint links**

| | |
|---|---|
| data | http://localhost:8081/engine/v2/data/cpu-anomaly-detection |
| alertsLongPoll | http://localhost:8081/engine/v2/alerts_longpoll/cpu-anomaly-detection |
| records | http://localhost:8081/engine/v2/results/cpu-anomaly-detection/records |
| buckets | http://localhost:8081/engine/v2/results/cpu-anomaly-detection/buckets |
| categoryDefinitions | http://localhost:8081/engine/v2/results/cpu-anomaly-detection/categorydefinitions |
| logs | http://localhost:8081/engine/v2/logs/cpu-anomaly-detection |
| modelSnapshots | http://localhost:8081/engine/v2/modelsnapshots/cpu-anomaly-detection |

Dev Tools

Console

```
1 GET _cat/indices/prelert
2
```

```
1 yellow open prelert ego1iHrQQe2FSpWgUKSosA 5 1 1100 0 428.3kb 428.3kb
2
```

$warning=.es(index=cpu-anomaly-detection-results, timefield=timestamp).divide(.es(index=cpu-anomaly-detection-results, timefield=timestamp)).multiply(.es(inde: auto

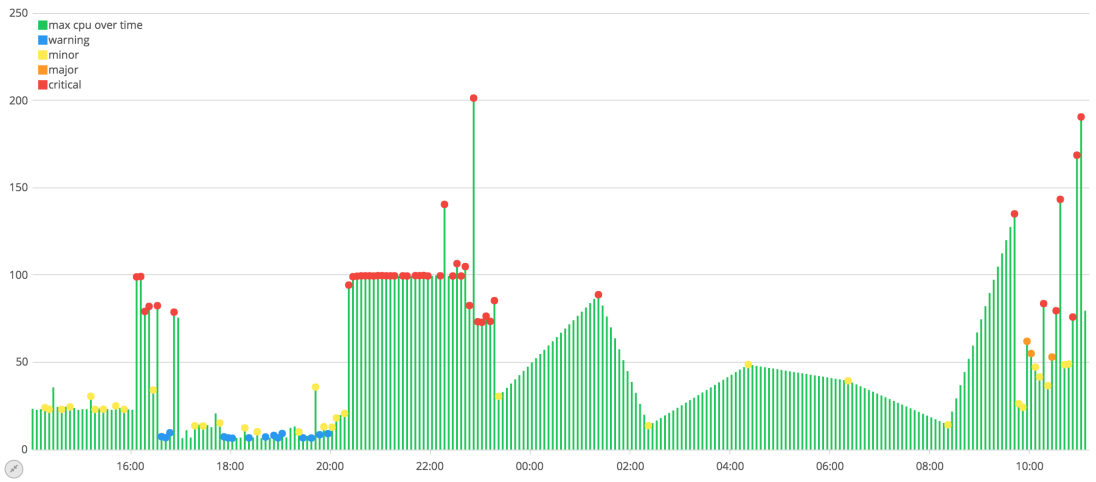

## Save entire Timelion sheet

You want this option if you mostly use Timelion expressions from within the Timelion app and don't need to add Timelion charts to Kibana dashboards. You may also want this if you make use of references to other panels.

## Save current expression as Kibana dashboard panel

Need to add a chart to a Kibana dashboard? We can do that! This option will save your currently selected expression as a panel that can be added to Kibana dashboards as you would add anything else. Note, if you use references to other panels you will need to remove the refences by copying the referenced expression directly into the expression you are saving. Click a chart to select a different expression to save.

**metricbeat***

Data | Options

Max

**Field**

system.process.cpu.total.pct

**Custom Label**

Max CPU

Advanced

Add metrics

**buckets**

**Split Rows**

**Aggregation**

Terms

**Field**

system.process.name

**Order By**

metric: Max CPU

**Order** | **Size**

Descending | 20

**Custom Label**

Process name

| Process name | Max CPU |
|---|---|
| metricbeat | 2.239 |
| java | 2.013 |
| pbzip2 | 1.672 |
| Google Chrome H | 1.433 |
| CalendarAgent | 1.35 |
| Sublime Text 2 | 1.161 |
| node | 1.012 |
| Microsoft Word | 1.012 |
| yes | 0.997 |
| Google Chrome | 0.688 |

Export: Raw Formatted

system.process.cpu.total

Fields (2) Scripted fields (0)

| name | type | format | searchable | aggregatable | analyzed | controls |
|---|---|---|---|---|---|---|
| system.process.cpu.total.ticks | number | | | | | |
| system.process.cpu.total.pct | number | | ✔ | ✔ | | |

# metricbeat*

## system.process.cpu.total.pct

**Type**

number

**Format** (Default: *Number*)

- ✓ - default -
- Url
- Bytes
- Duration
- Number
- Percentage
- String
- Color
- Boolean

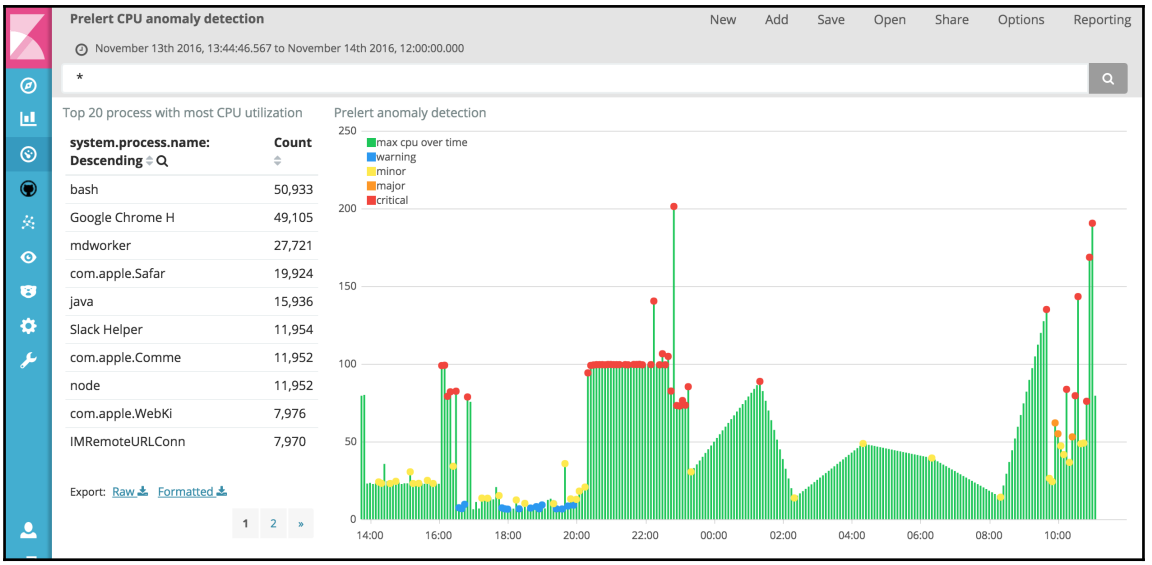
Prelert CPU anomaly detection
New
Add
Save
Open
Share
Options
Reporting
November 13th 2016, 13:44:46.567 to November 14th 2016, 12:00:00.000
*
Top 20 process with most CPU utilization
system.process.name: Descending
Count
bash 50,933
Google Chrome H 49,105
mdworker 27,721
com.apple.Safar 19,924
java 15,936
Slack Helper 11,954
com.apple.Comme 11,952
node 11,952
com.apple.WebKi 7,976
IMRemoteURLConn 7,970
Export: Raw Formatted
1 2 »
Prelert anomaly detection
max cpu over time
warning
minor
major
critical
250
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

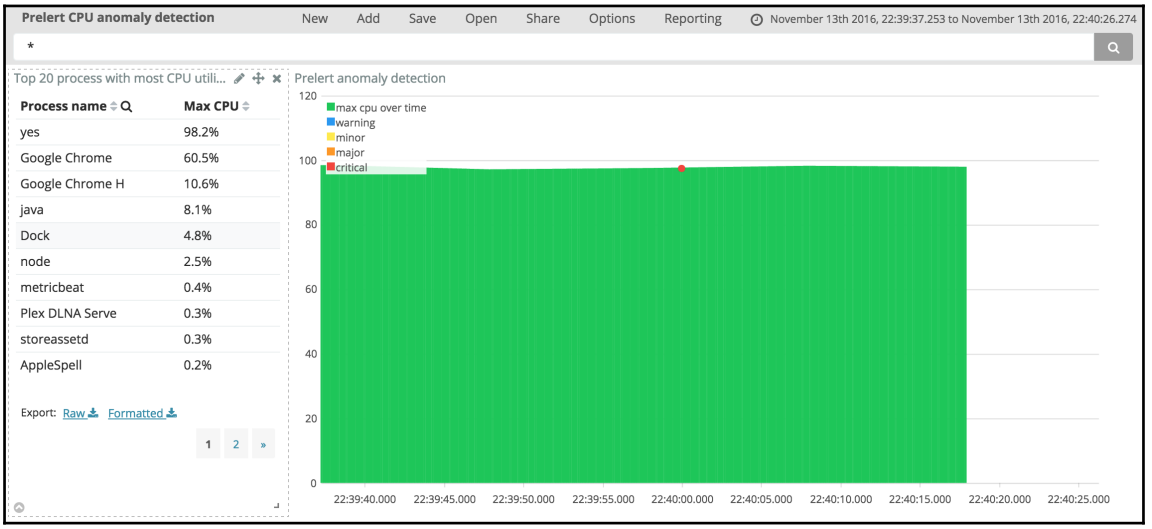
Prelert CPU anomaly detection
New
Add
Save
Open
Share
Options
Reporting
November 13th 2016, 22:39:37.253 to November 13th 2016, 22:40:26.274
*
Top 20 process with most CPU utili...
Process name
Max CPU
yes 98.2%
Google Chrome 60.5%
Google Chrome H 10.6%
java 8.1%
Dock 4.8%
node 2.5%
metricbeat 0.4%
Plex DLNA Serve 0.3%
storeassetd 0.3%
AppleSpell 0.2%
Export: Raw Formatted
1 2 »
Prelert anomaly detection
max cpu over time
warning
minor
major
critical
120
100
80
60
40
20
0
22:39:40.000
22:39:45.000
22:39:50.000
22:39:55.000
22:40:00.000
22:40:05.000
22:40:10.000
22:40:15.000
22:40:20.000
22:40:25.000

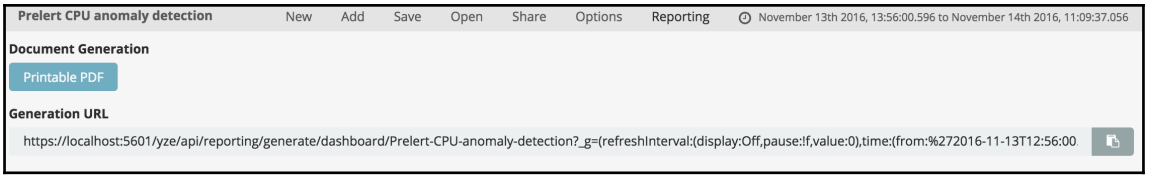
Prelert CPU anomaly detection
New
Add
Save
Open
Share
Options
Reporting
November 13th 2016, 13:56:00.596 to November 14th 2016, 11:09:37.056
Document Generation
Printable PDF
Generation URL
https://localhost:5601/yze/api/reporting/generate/dashboard/Prelert-CPU-anomaly-detection?_g=(refreshInterval:(display:Off,pause:!f,value:0),time:(from:%272016-11-13T12:56:00

Management / Kibana

Index Patterns Saved Objects Reporting Advanced Settings

# Generated Reports

Filter Reports: ☑ **Only show my reports**

| Document | Added |
| --- | --- |
| Prelert CPU anomaly detection<br>dashboard | 2016-11-14 @ 5:09 PM<br>elastic |
| Prelert CPU anomaly detection<br>dashboard | 2016-11-14 @ 3:05 PM<br>elastic |

# Chapter 9: Creating a Custom Plugin for Kibana 5.0

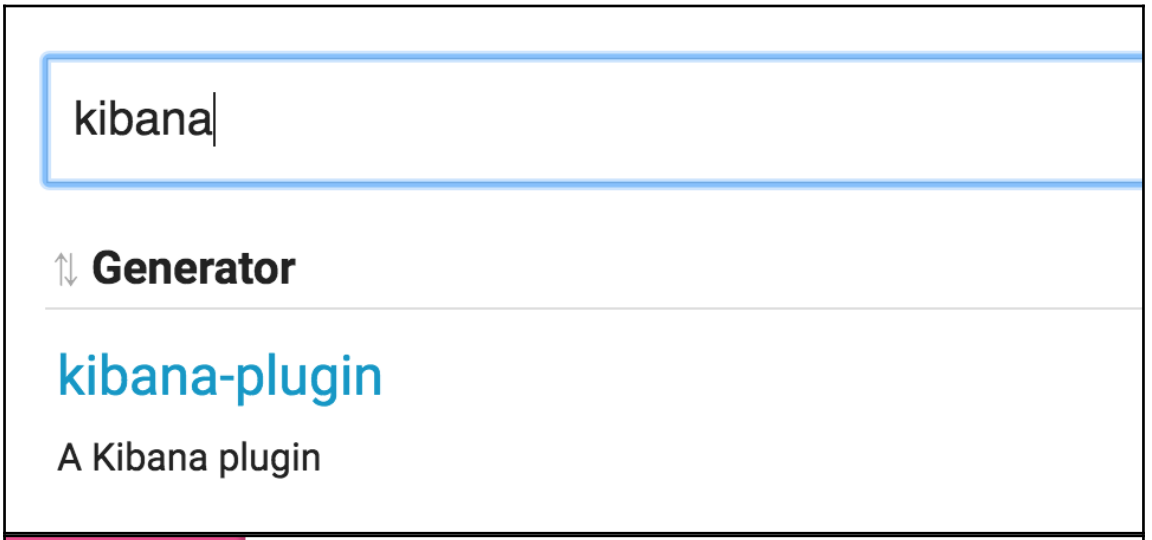


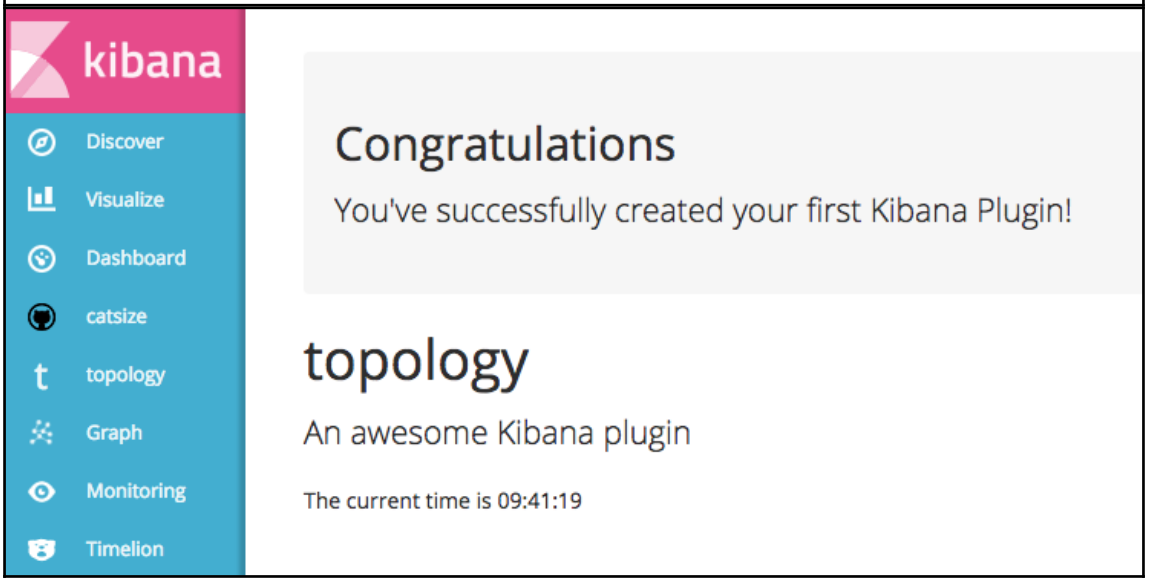

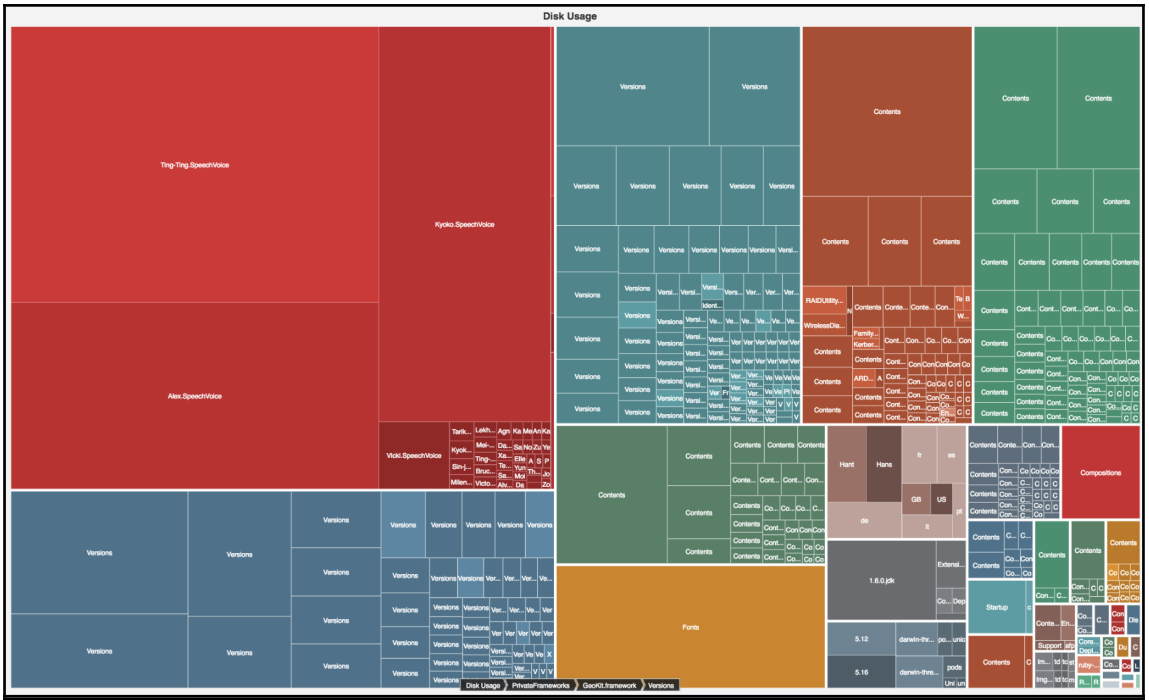
Disk Usage
Ting-Ting.SpeechVoice
Kyoko.SpeechVoice
Alex.SpeechVoice
Vicki.SpeechVoice
Versions
Contents
Fonts
Compositions
1.6.0.jdk
Disk Usage
PrivateFrameworks
GeoKit.framework
Versions

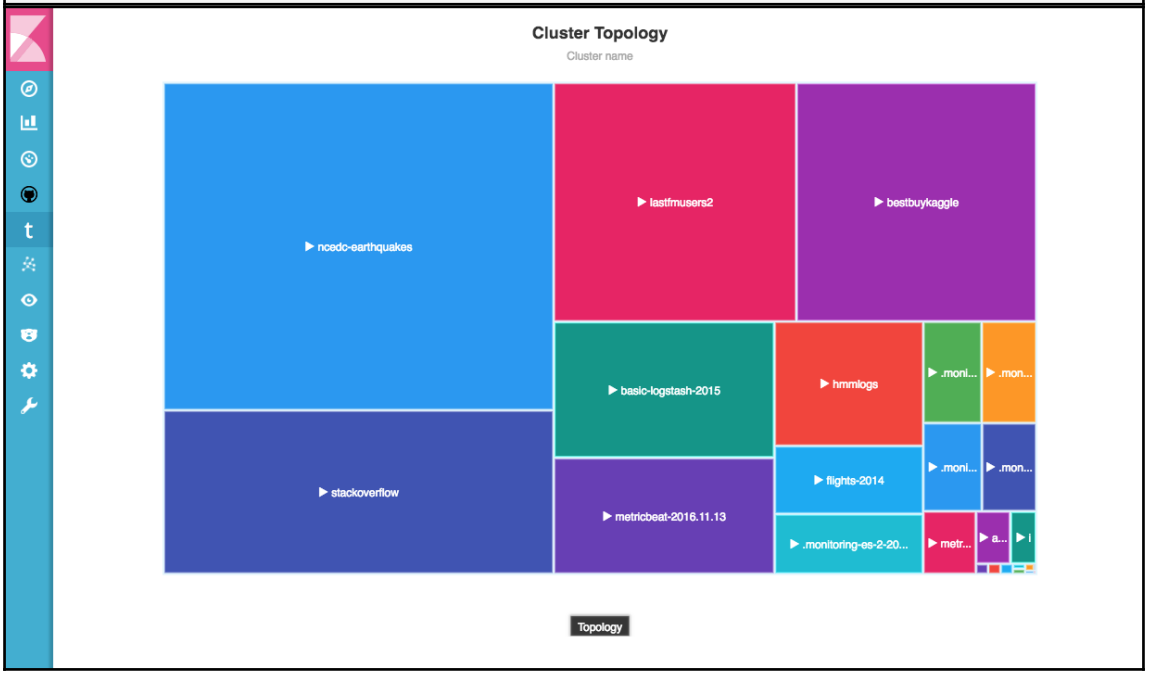
Cluster Topology
Cluster name
nceqc-earthquakes
lastfmusers2
bestbuykaggle
basic-logstash-2015
hmmlogs
stackoverflow
metricbeat-2016.11.13
flights-2014
.monitoring-es-2-20...
Topology

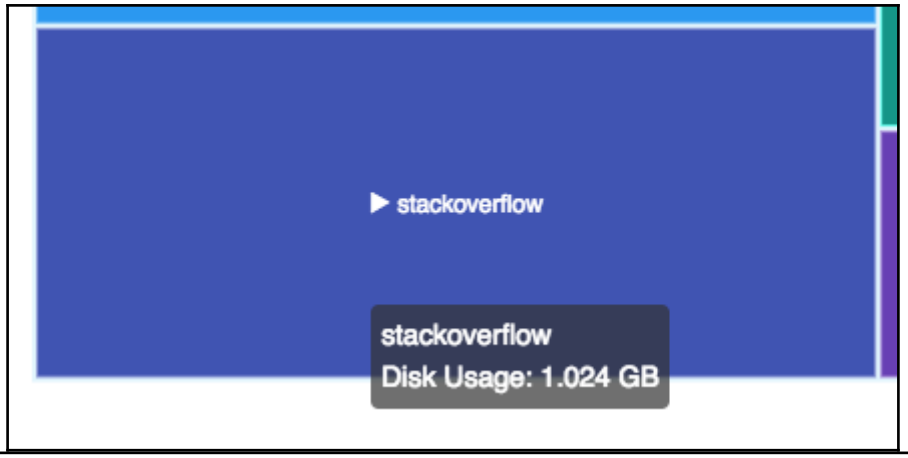
stackoverflow
stackoverflow
Disk Usage: 1.024 GB


Cluster Topology
Cluster name
_pr
_mh
_ae
_gj
_d7
stackoverflow/primary/0/_d7
Disk Usage: 215 MB
_49
_qo
_rs
_qe
_r8
_q3
_ri
_pk
_d
_jc
_7c
_qy
_s2
_n2
_qd
_s1
Topology
stackoverflow
primary
0