

# Chapter 1: An Introduction to Kibana



## Prerequisites for installing Kibana 4.1.1

### Installation of Java

#### Installation of Java on Ubuntu 14.04

```
yuvraj@L212:~$ sudo add-apt-repository -y ppa:webupd8team/java
[sudo] password for yuvraj:
gpg: keyring `/tmp/tmpcx59wizk/secring.gpg' created
gpg: keyring `/tmp/tmpcx59wizk/pubring.gpg' created
gpg: requesting key EEA14886 from hkp server keyserver.ubuntu.com
gpg: /tmp/tmpcx59wizk/trustdb.gpg: trustdb created
gpg: key EEA14886: public key "Launchpad VLC" imported
gpg: Total number processed: 1
gpg:      imported: 1 (RSA: 1)
OK
```

```
Hit http://in.archive.ubuntu.com trusty-backports/multiverse Translation-en
Hit http://in.archive.ubuntu.com trusty-backports/restricted Translation-en
Hit http://in.archive.ubuntu.com trusty-backports/universe Translation-en
Ign http://in.archive.ubuntu.com trusty/main Translation-en_IN
Ign http://in.archive.ubuntu.com trusty/multiverse Translation-en_IN
Ign http://in.archive.ubuntu.com trusty/restricted Translation-en_IN
Ign http://in.archive.ubuntu.com trusty/universe Translation-en_IN
Fetched 3,059 kB in 26s (116 kB/s)
Reading package lists... Done
```

```
yuvraj@L212:~$ sudo apt-get -y install oracle-java8-installer
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following packages were automatically installed and are no longer required:
  libunibreak1 libzlibcore-data libzlibcore0.12 libzlibtext-data libzlibtext0.12
  libzlibui-gtk4
Use 'apt-get autoremove' to remove them.
Suggested packages:
  binfmt-support visualvm ttf-baekmuk ttf-unfonts ttf-unfonts-core
  ttf-kochi-gothic ttf-sazanami-gothic ttf-kochi-mincho ttf-sazanami-mincho
  ttf-arphic-uming
The following NEW packages will be installed:
  oracle-java8-installer
0 upgraded, 1 newly installed, 0 to remove and 436 not upgraded.
Need to get 0 B/22.6 kB of archives.
After this operation, 129 kB of additional disk space will be used.
Preconfiguring packages ...
Selecting previously unselected package oracle-java8-installer.
(Reading database ... 179955 files and directories currently installed.)
Preparing to unpack ../oracle-java8-installer_8u45+8u33arm-1-webupd8-1_all.deb ...
oracle-license-v1-1 license has already been accepted
Unpacking oracle-java8-installer (8u45+8u33arm-1-webupd8-1) ...
Processing triggers for gnome-menus (3.10.1-0ubuntu2) ...
Processing triggers for desktop-file-utils (0.22-1ubuntu1) ...
Processing triggers for banfdaemon (0.5.1+14.04.20140409-0ubuntu1) ...
```



```
yuvraj@L212:~$ java -version
java version "1.8.0_45"
Java(TM) SE Runtime Environment (build 1.8.0_45-b14)
Java HotSpot(TM) 64-Bit Server VM (build 25.45-b02, mixed mode)
```

# Installation of Java on Windows

Oracle Technology Network > Java > Java SE > Downloads

Overview Downloads Documentation Community Technologies Training

### Java SE Downloads

[DOWNLOAD +](#)
[DOWNLOAD +](#)

**Java Platform, Standard Edition**

**Java SE 8u60**  
 This releases includes support for ARMv8 processors, Nashorn enhancements, and improvements to Deployment Rule Set functionality.  
 JDK for ARM releases are now available on the same page as the downloads for other platforms.  
[Learn more +](#)

- Installation Instructions
- Release Notes
- Oracle License
- Java SE Products
- Third Party Licenses

[JDK DOWNLOAD +](#)  
[Server JRE DOWNLOAD +](#)

**Java SDKs and Tools**  
[Java SE](#)  
[Java EE and Glassfish](#)  
[Java ME](#)  
[Java Card](#)  
[NetBeans IDE](#)  
[Java Mission Control](#)


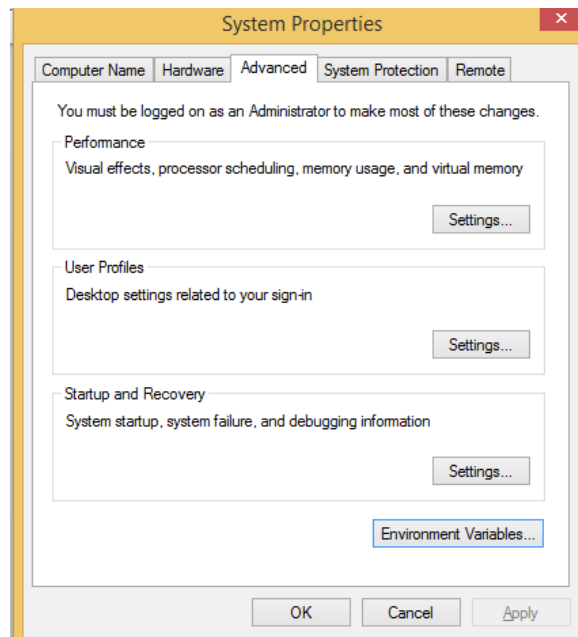
**Java Resources**  
[Java APIs](#)  
[Technical Articles](#)  
[Demos and Videos](#)  
[Forums](#)  
[Java Magazine](#)  
[Java.net](#)  
[Developer Training](#)  
[Tutorials](#)  
[Java.com](#)

## Java SE Development Kit 8u45

**You must accept the [Oracle Binary Code License Agreement for Java SE](#) to download this software.**

Accept License Agreement
  Decline License Agreement

Product / File Description	File Size	Download
Linux x86	146.89 MB	<a href="#">jdk-8u45-linux-i586.rpm</a>
Linux x86	166.88 MB	<a href="#">jdk-8u45-linux-i586.tar.gz</a>
Linux x64	145.19 MB	<a href="#">jdk-8u45-linux-x64.rpm</a>
Linux x64	165.24 MB	<a href="#">jdk-8u45-linux-x64.tar.gz</a>
Mac OS X x64	221.98 MB	<a href="#">jdk-8u45-macosx-x64.dmg</a>
Solaris SPARC 64-bit (SVR4 package)	131.73 MB	<a href="#">jdk-8u45-solaris-sparcv9.tar.Z</a>
Solaris SPARC 64-bit	92.9 MB	<a href="#">jdk-8u45-solaris-sparcv9.tar.gz</a>
Solaris x64 (SVR4 package)	139.51 MB	<a href="#">jdk-8u45-solaris-x64.tar.Z</a>
Solaris x64	95.88 MB	<a href="#">jdk-8u45-solaris-x64.tar.gz</a>
Windows x86	175.98 MB	<a href="#">jdk-8u45-windows-i586.exe</a>
Windows x64	180.44 MB	<a href="#">jdk-8u45-windows-x64.exe</a>



```
C:\Windows\system32\cmd.exe
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.

C:\Users\ygupta>java -version
java version "1.8.0_45"
Java(TM) SE Runtime Environment (build 1.8.0_45-b15)
Java HotSpot(TM) 64-Bit Server VM (build 25.45-b02, mixed mode)
```

# Installation of Elasticsearch

## Installation of Elasticsearch on Ubuntu 14.04


```
yuvraj@L212:~$ curl -L -O https://download.elastic.co/elasticsearch/elasticsearch/elasticsearch-1.5.2.tar.gz
% Total % Received % Xferd Average Speed Time Time Time Current
Dload Upload Total Spent Left Speed
100 26.8M 100 26.8M 0 0 432k 0 0:01:03 0:01:03 --:--:-- 1043k
```

```
yuvraj@L212:~/elasticsearch-1.5.2/bin$ ./elasticsearch
[2015-06-11 15:16:56,261][INFO ][node ] [Fan Boy] version[1.5.2], pid[26804], build[62ff986/2015-04-27T09:21:06Z]
[2015-06-11 15:16:56,262][INFO ][node ] [Fan Boy] initializing ...
[2015-06-11 15:16:56,267][INFO ][plugins] [Fan Boy] loaded [], sites []
[2015-06-11 15:16:59,417][INFO ][node ] [Fan Boy] initialized
[2015-06-11 15:16:59,418][INFO ][node ] [Fan Boy] starting ...
[2015-06-11 15:16:59,551][INFO ][transport] [Fan Boy] bound_address {inet[/0:0:0:0:0:0:9300]}, publish_address {inet[/172.16.8.93:9300]}
[2015-06-11 15:16:59,692][INFO ][discovery] [Fan Boy] elasticsearch/JUC1UHVQSS6PBlesY3ENaA
[2015-06-11 15:17:03,507][INFO ][cluster.service] [Fan Boy] new_master [Fan Boy][JUC1UHVQSS6PBlesY3ENaA][L212][inet[/172.16.8.93:9300]]
, reason: zen-disco-join (elected_as_master)
[2015-06-11 15:17:03,554][INFO ][http] [Fan Boy] bound_address {inet[/0:0:0:0:0:0:9200]}, publish_address {inet[/172.16.8.93:9200]}
[2015-06-11 15:17:03,555][INFO ][node ] [Fan Boy] started
[2015-06-11 15:17:03,636][INFO ][gateway] [Fan Boy] recovered [0] indices into cluster_state
```

```
{
  "status" : 200,
  "name" : "Fan Boy",
  "cluster_name" : "elasticsearch",
  "version" : {
    "number" : "1.5.2",
    "build_hash" : "62ff9868b4c8a0c45860bebb259e21980778ab1c",
    "build_timestamp" : "2015-04-27T09:21:06Z",
    "build_snapshot" : false,
    "lucene_version" : "4.10.4"
  },
  "tagline" : "You Know, for Search"
}
```

## Installation of Elasticsearch on Windows

### Installation of GIT



```
ygupta@XE-T-YGUPTA /C:/Users/ygupta/Desktop/Personal
$ curl -L -O https://download.elastic.co/elasticsearch/elasticsearch/elasticsearch-1.5.2.zip
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left   Speed
  1 29.9M    1 319k    0     0    6252      0  1:23:47  0:00:52  1:22:55  9648
```

```
Elasticsearch 1.5.2
[2015-06-27 11:43:26.708][INFO ][node ] [Astronomer] version
[1.5.2], pid[4700], build[62ff986/2015-04-27T09:21:06Z]
[2015-06-27 11:43:26.716][INFO ][node ] [Astronomer] initial
izing ...
[2015-06-27 11:43:26.813][INFO ][plugins ] [Astronomer] loaded
[river-twitter], sites [head]
[2015-06-27 11:43:30.192][INFO ][node ] [Astronomer] initial
ized
[2015-06-27 11:43:39.290][INFO ][node ] [Astronomer] startin
g ...
[2015-06-27 11:43:39.661][INFO ][transport ] [Astronomer] bound_a
ddress <inet[/0:0:0:0:0:0:9300]>, publish_address <inet[/192.168.1.111:9300]
>
[2015-06-27 11:43:39.887][INFO ][discovery ] [Astronomer] elastic
search/Qg9rCrAySDaKQp0m2sAg7A
[2015-06-27 11:43:43.692][INFO ][cluster.service ] [Astronomer] new_mas
ter [Astronomer] [Qg9rCrAySDaKQp0m2sAg7A] [XE-T-YGUPTA] [inet[/192.168.1.111:9300]]
, reason: zen-disco-join (elected_as_master)
[2015-06-27 11:43:43.838][INFO ][http ] [Astronomer] bound_a
ddress <inet[/0:0:0:0:0:0:9200]>, publish_address <inet[/192.168.1.111:9200]
>
[2015-06-27 11:43:43.838][INFO ][node ] [Astronomer] started
```

```
{
  "status" : 200,
  "name" : "Astronomer",
  "cluster_name" : "elasticsearch",
  "version" : {
    "number" : "1.5.2",
    "build_hash" : "62ff9868b4c8a0c45860bebb259e21980778ab1c",
    "build_timestamp" : "2015-04-27T09:21:06Z",
    "build_snapshot" : false,
    "lucene_version" : "4.10.4"
  },
  "tagline" : "You Know, for Search"
}
```

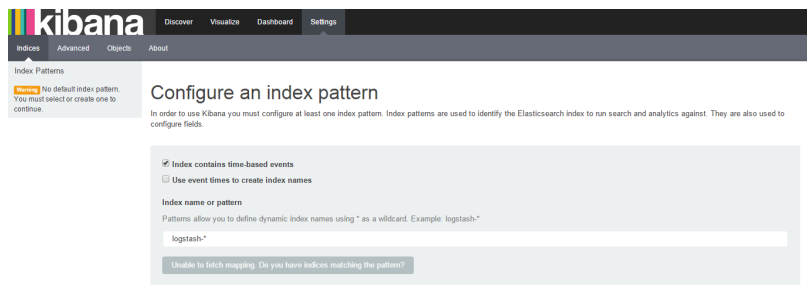
# Installation of Kibana

## Installation of Kibana on Ubuntu 14.04

```
yuvraj@L212:~$ curl -L -O https://download.elasticsearch.org/kibana/kibana/kibana-4.1.1-linux-x64.tar.gz
% Total % Received % Xferd Average Speed Time Time Time Current
         Dload Upload Total Spent Left Speed
100 11.1M 100 11.1M 0 0 520k 0 0:00:21 0:00:21 --:--:-- 622k
```

```
yuvraj@L212:~$ cd kibana-4.1.1-linux-x64/bin/
yuvraj@L212:~/kibana-4.1.1-linux-x64/bin$ ./kibana
{"name": "Kibana", "hostname": "L212", "pid": 25104, "level": 30, "msg": "Found kibana index", "time": "2015-09-26T16:00:05.889Z", "v": 0}
{"name": "Kibana", "hostname": "L212", "pid": 25104, "level": 30, "msg": "Listening on 0.0.0.0:5601", "time": "2015-09-26T16:00:06.123Z", "v": 0}
```

```
yuvraj@L212:~/kibana-4.1.1-linux-x64/bin$ ./kibana
{"name": "Kibana", "hostname": "L212", "pid": 25145, "level": 50, "err": "Request error, retrying -- connect ECONNREFUSED", "msg": "", "time": "2015-09-26T16:00:23.105Z", "v": 0}
{"name": "Kibana", "hostname": "L212", "pid": 25145, "level": 40, "msg": "Unable to revive connection: http://localhost:9200/", "time": "2015-09-26T16:00:23.110Z", "v": 0}
{"name": "Kibana", "hostname": "L212", "pid": 25145, "level": 40, "msg": "No living connections", "time": "2015-09-26T16:00:23.110Z", "v": 0}
{"name": "Kibana", "hostname": "L212", "pid": 25145, "level": 30, "msg": "Unable to connect to elasticsearch at http://localhost:9200. Retrying in 2.5 seconds.", "time": "2015-09-26T16:00:23.113Z", "v": 0}
```



## Installation of Kibana on Windows

```
ygupta@XE-T-YGUPTA /C:/Users/ygupta/Desktop
$ curl -L -O https://download.elasticsearch.org/kibana/kibana/kibana-4.1.1-wind
ows.zip
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload    Total   Spent    Left   Speed
 16 10.0M    16 1663k    0     0   110k    0  0:01:33  0:00:15  0:01:18  121k
```

```
Kibana Server 4.1.1
{"name":"Kibana","hostname":"XE-T-YGUPTA","pid":2136,"level":30,"msg":"Found kib
ana index","time":"2015-09-26T08:06:11.853Z","v":0}
{"name":"Kibana","hostname":"XE-T-YGUPTA","pid":2136,"level":30,"msg":"Listening
on 0.0.0.0:5601","time":"2015-09-26T08:06:11.867Z","v":0}
```

```
Kibana Server 4.1.1
{"name":"Kibana","hostname":"XE-T-YGUPTA","pid":6556,"level":50,"err":"Request e
rror, retrying -- connect ECONNREFUSED","msg":"","time":"2015-09-26T08:04:00.439
Z","v":0}
{"name":"Kibana","hostname":"XE-T-YGUPTA","pid":6556,"level":40,"msg":"Unable to
revive connection: http://localhost:9200/","time":"2015-09-26T08:04:00.554Z","v
":0}
{"name":"Kibana","hostname":"XE-T-YGUPTA","pid":6556,"level":40,"msg":"No living
connections","time":"2015-09-26T08:04:00.555Z","v":0}
{"name":"Kibana","hostname":"XE-T-YGUPTA","pid":6556,"level":30,"msg":"Unable to
connect to elasticsearch at http://localhost:9200. Retrying in 2.5 seconds.","t
ime":"2015-09-26T08:04:00.556Z","v":0}
```

The screenshot shows the Kibana web interface. At the top, there is a navigation bar with the Kibana logo and tabs for 'Discover', 'Visualize', 'Dashboard', and 'Settings'. Below the navigation bar, there are links for 'Indices', 'Advanced', 'Objects', and 'About'. The main content area is titled 'Configure an index pattern'. It contains a warning message: 'Warning: No default index pattern. You must select or create one to continue.' Below this, there is a section for 'Index name or pattern' with a text input field containing 'logstash-'. There are also two checkboxes: 'Index contains time-based events' (checked) and 'Use event times to create index names' (unchecked). At the bottom, there is a message: 'Unable to fetch mapping. Do you have indices matching the pattern?'.



## Additional information

### Changing the Elasticsearch configuration

```
##### Elasticsearch Configuration Example #####
# This file contains an overview of various configuration settings,
# targeted at operations staff. Application developers should
# consult the guide at <http://elasticsearch.org/guide>.
#
# The installation procedure is covered at
# <http://elasticsearch.org/guide/en/elasticsearch/reference/current/setup.html>.
#
# Elasticsearch comes with reasonable defaults for most settings,
# so you can try it out without bothering with configuration.
#
# Most of the time, these defaults are just fine for running a production
# cluster. If you're fine-tuning your cluster, or wondering about the
# effect of certain configuration option, please do ask on the
# mailing list or IRC channel [http://elasticsearch.org/community].
#
# Any element in the configuration can be replaced with environment variables
# by placing them in ${...} notation. For example:
#
#node.rack: ${RACK_ENV_VAR}
#
# For information on supported formats and syntax for the config file, see
# <http://elasticsearch.org/guide/en/elasticsearch/reference/current/setup-configuration.html>
#
##### Cluster #####
#
# Cluster name identifies your cluster for auto-discovery. If you're running
# multiple clusters on the same network, make sure you're using unique names.
#
#cluster.name: elasticsearch
```

```
##### Elasticsearch Configuration Example #####
# This file contains an overview of various configuration settings,
# targeted at operations staff. Application developers should
# consult the guide at <http://elasticsearch.org/guide>.
#
# The installation procedure is covered at
# <http://elasticsearch.org/guide/en/elasticsearch/reference/current/setup.html>.
#
# Elasticsearch comes with reasonable defaults for most settings,
# so you can try it out without bothering with configuration.
#
# Most of the time, these defaults are just fine for running a production
# cluster. If you're fine-tuning your cluster, or wondering about the
# effect of certain configuration option, please do ask on the
# mailing list or IRC channel [http://elasticsearch.org/community].
#
# Any element in the configuration can be replaced with environment variables
# by placing them in ${...} notation. For example:
#
#node.rack: ${RACK_ENV_VAR}
#
# For information on supported formats and syntax for the config file, see
# <http://elasticsearch.org/guide/en/elasticsearch/reference/current/setup-configuration.html>
#
##### Cluster #####
#
# Cluster name identifies your cluster for auto-discovery. If you're running
# multiple clusters on the same network, make sure you're using unique names.
#
#cluster.name: test
```

```

{
  "status" : 200,
  "name" : "Slick",
  "cluster_name" : "test",
  "version" : {
    "number" : "1.5.2",
    "build_hash" : "62ff9868b4c8a0c45860bebb259e21980778ab1c",
    "build_timestamp" : "2015-04-27T09:21:06Z",
    "build_snapshot" : false,
    "lucene_version" : "4.10.4"
  },
  "tagline" : "You Know, for Search"
}

```

## Changing Kibana configuration

```

# Kibana is served by a back end server. This controls which port to use.
port: 5601

# The host to bind the server to.
host: "0.0.0.0"

# The Elasticsearch instance to use for all your queries.
elasticsearch_url: "http://localhost:9200"

# preserve_elasticsearch_host true will send the hostname specified in 'elasticsearch'. If you set it to false,
# then the host you use to connect to *this* Kibana instance will be sent.
elasticsearch_preserve_host: true

```

```

# Kibana is served by a back end server. This controls which port to use.
port: 5604

# The host to bind the server to.
host: "0.0.0.0"

# The Elasticsearch instance to use for all your queries.
elasticsearch_url: "http://localhost:9200"

# preserve_elasticsearch_host true will send the hostname specified in 'elasticsearch'. If you set it to false,
# then the host you use to connect to *this* Kibana instance will be sent.
elasticsearch_preserve_host: true

```

localhost:5604/#/settings/indices/?g=0

Discover Visualize Dashboard Settings

Indices Advanced Objects About

**Warning** No default index pattern. You must select or create one to continue.

### 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

Use event times to create index names

Index name or pattern

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

logstash-\*

Unable to fetch mapping. Do you have indices matching the pattern?

# Importing a JSON file into Elasticsearch

## Installation of npm

### Installation of npm on Ubuntu 14.04

```
yuvraj@L212:~$ curl --silent --location https://deb.nodesource.com/setup_0.12 | sudo bash -  
  
## Installing the NodeSource Node.js 0.12 repo...  
  
## Populating apt-get cache...  
  
+ apt-get update  
Ign http://dl.google.com stable InRelease  
Hit http://dl.google.com stable Release.gpg  
Hit http://dl.google.com stable Release  
Hit http://dl.google.com stable/main amd64 Packages  
Hit http://dl.google.com stable/main i386 Packages  
Ign http://extras.ubuntu.com trusty InRelease  
Ign http://security.ubuntu.com trusty-security InRelease  
Ign http://archive.canonical.com trusty InRelease  
Hit http://security.ubuntu.com trusty-security Release.gpg  
Hit http://archive.canonical.com trusty Release.gpg  
Hit http://security.ubuntu.com trusty-security Release  
Hit http://archive.canonical.com trusty Release  
Ign http://dl.google.com stable/main Translation-en_IN  
Ign http://dl.google.com stable/main Translation-en  
Hit http://security.ubuntu.com trusty-security/main Sources  
Hit http://security.ubuntu.com trusty-security/restricted Sources  
Hit http://security.ubuntu.com trusty-security/universe Sources
```

```
Hit http://in.archive.ubuntu.com trusty-backports/multiverse i386 Packages  
Hit http://in.archive.ubuntu.com trusty-backports/main Translation-en  
Hit http://in.archive.ubuntu.com trusty-backports/multiverse Translation-en  
Hit http://in.archive.ubuntu.com trusty-backports/restricted Translation-en  
Hit http://in.archive.ubuntu.com trusty-backports/universe Translation-en  
Ign http://in.archive.ubuntu.com trusty/main Translation-en_IN  
Ign http://in.archive.ubuntu.com trusty/multiverse Translation-en_IN  
Ign http://in.archive.ubuntu.com trusty/restricted Translation-en_IN  
Ign http://in.archive.ubuntu.com trusty/universe Translation-en_IN  
Fetched 6,627 B in 31s (210 B/s)  
Reading package lists... Done  
  
## Run `apt-get install nodejs` (as root) to install Node.js 0.12 and npm
```

```
yuvraj@L212:~$ sudo apt-get install --yes nodejs
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following packages were automatically installed and are no longer required:
  gyp libamd2.3.1 libbabl-0.1-0 libc-ares-dev libc-ares2 libcamd2.3.1
  libccolamd2.8.0 libcholmod2.1.2 libgegl-0.2-0 libgfortran3
  libjavascriptcoregtk-1.0-0 libjs-node-uuid liblapack3 libmng2 libssl-dev
  libssl-doc libumfpack5.6.2 libunibreak1 libv8-3.14-dev libv8-3.14.5
  libwebkitgtk-1.0-0 libwebkitgtk-1.0-common libzlib-dev libzlib1-dev
  libzlib1-gtk-1.0-0 libzlib1-gtk-1.0-common libzlib1-gtk-1.0-data libzlib1-gtk-1.0-dev
Use 'apt-get autoremove' to remove them.
The following NEW packages will be installed:
  nodejs
0 upgraded, 1 newly installed, 0 to remove and 107 not upgraded.
Need to get 5,416 kB of archives.
After this operation, 26.8 MB of additional disk space will be used.
Get:1 https://deb.nodesource.com/node_0.12/ trusty/main nodejs amd64 0.12.7-1nodesource1-trusty1 [5,416 kB]
Fetched 5,416 kB in 10min 16s (8,780 B/s)
Selecting previously unselected package nodejs.
(Reading database ... 225502 files and directories currently installed.)
Preparing to unpack .../nodejs_0.12.7-1nodesource1-trusty1_amd64.deb ...
Unpacking nodejs (0.12.7-1nodesource1-trusty1) ...
Processing triggers for man-db (2.6.7.1-1ubuntu1) ...
Setting up nodejs (0.12.7-1nodesource1-trusty1) ...
```

## Installing elasticsearch

### Installing elasticsearch on Ubuntu 14.04

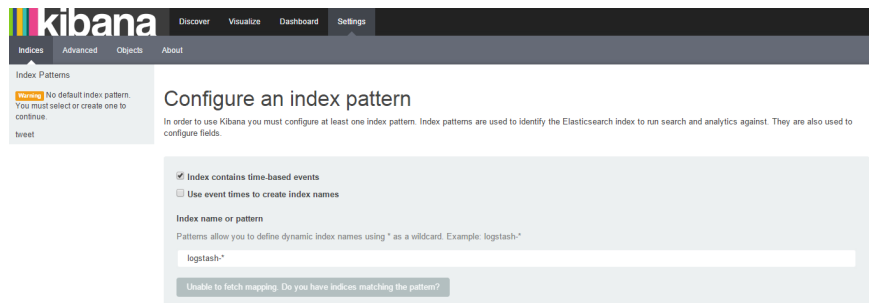
```
yuvraj@L212:~$ sudo npm install elasticsearch -g
[sudo] password for yuvraj:
/usr/bin/elasticsearch -> /usr/lib/node_modules/elasticsearch/bin/elasticsearch
/usr/bin/multielasticdump -> /usr/lib/node_modules/elasticsearch/bin/multielasticdump
elasticsearch@0.14.1 /usr/lib/node_modules/elasticsearch
├── line-reader@0.2.4
├── JSONStream@0.9.0 (through@2.3.7, jsonparse@0.0.5)
├── optimist@0.6.1 (wordwrap@0.0.3, minimist@0.0.10)
├── request@2.58.0 (caseless@0.10.0, aws-sign2@0.5.0, forever-agent@0.6.1, stringstream@0.0.4, tunnel-agent@0.4.0, oauth-sign@0.8.0, isstream@0.1.2, extend@2.0.1, json-stringify-safe@5.0.1, node-uuid@1.4.3, qs@3.1.0, tough-cookie@2.0.0, mime-types@2.0.14, combined-stream@1.0.5, http-signature@0.11.0, form-data@1.0.0-rc1, hawk@2.3.1, bl@0.9.4, har-validator@1.8.0)
```

```
yuvraj@L212:~$ elasticsearch \
> --bulk=true \
> --input="/home/yuvraj/Desktop/tweet.json" \
> --output=http://localhost:9200/
Wed, 01 Jul 2015 10:24:44 GMT | starting dump
Wed, 01 Jul 2015 10:24:44 GMT | got 100 objects from source file (offset: 0)
Wed, 01 Jul 2015 10:24:46 GMT | sent 100 objects to destination elasticsearch, wrote 100
Wed, 01 Jul 2015 10:24:46 GMT | got 100 objects from source file (offset: 100)
Wed, 01 Jul 2015 10:24:46 GMT | sent 100 objects to destination elasticsearch, wrote 100
Wed, 01 Jul 2015 10:24:46 GMT | got 100 objects from source file (offset: 200)
```

## Installing elasticsearch-dump on Windows

```
ygupta@XE-T-YGUPTA ~/Desktop/Personal/Kibana/elasticsearch-1.5.2
$ npm install elasticsearch-dump -g
C:\Users\ygupta\AppData\Roaming\npm\multielasticsearch-dump -> C:\Users\ygupta\AppData\Roaming\npm\node_modules\elasticsearch-dump\bin\multielasticsearch-dump
C:\Users\ygupta\AppData\Roaming\npm\elasticsearch-dump -> C:\Users\ygupta\AppData\Roaming\npm\node_modules\elasticsearch-dump\bin\elasticsearch-dump
elasticsearch-dump@0.14.1 C:\Users\ygupta\AppData\Roaming\npm\node_modules\elasticsearch-dump
├── line-reader@0.2.4
├── optimist@0.6.1 (wordwrap@0.0.3, minimist@0.0.10)
├── JSONStream@0.9.0 (through@2.3.7, jsonparse@0.0.5)
├── request@2.58.0 (caseless@0.10.0, forever-agent@0.6.1, aws-sign2@0.5.0, stringstream@0.0.4, tunnel-agent@0.4.0, oauth-sign@0.8.0, isstream@0.1.2, extend@2.0.1, json-stringify-safe@5.0.1, node-uuid@1.4.3, qs@3.1.0, combined-stream@1.0.5, mime-types@2.0.14, tough-cookie@2.0.0, http-signature@0.11.0, bl@0.9.4, hawk@2.3.1, form-data@1.0.0-rc1, har-validator@1.8.0)
```

```
MINGW32:/C:/Users/ygupta/Desktop
ygupta@XE-T-YGUPTA /C:/Users/ygupta/Desktop
$ elasticsearch-dump \
> --bulk=true \
> --input="C:\Users\ygupta\Desktop\tweet.json" \
> --output=http://localhost:9200/
Sat, 27 Jun 2015 10:08:25 GMT : starting dump
Sat, 27 Jun 2015 10:08:25 GMT : got 100 objects from source file (offset: 0)
Sat, 27 Jun 2015 10:08:25 GMT : sent 100 objects to destination elasticsearch, w
rote 100
Sat, 27 Jun 2015 10:08:25 GMT : got 100 objects from source file (offset: 100)
Sat, 27 Jun 2015 10:08:26 GMT : sent 100 objects to destination elasticsearch, w
rote 100
Sat, 27 Jun 2015 10:08:26 GMT : got 100 objects from source file (offset: 200)
Sat, 27 Jun 2015 10:08:26 GMT : sent 100 objects to destination elasticsearch, w
rote 100
```



Index Patterns

**Warning** No default index pattern. You must select or create one to continue.

tweet

## 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
- Use event times to create index names

**Index name or pattern**

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

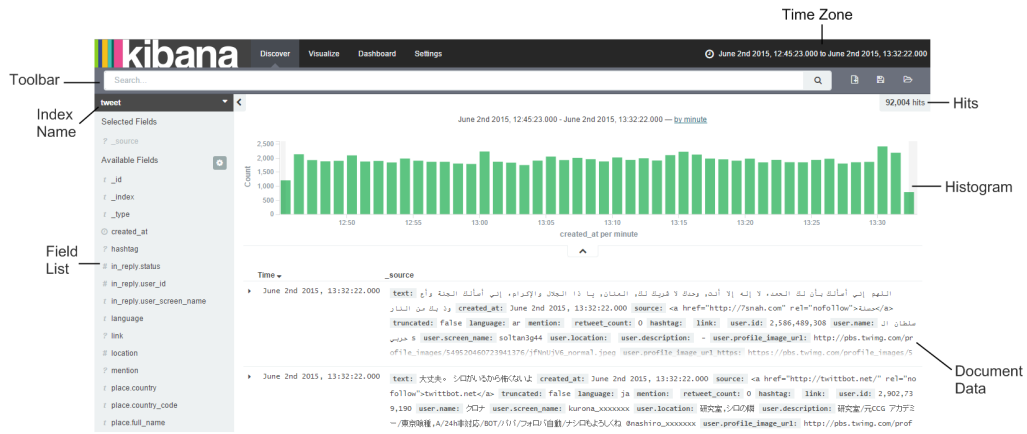
tweet

**Time-field name** refresh fields

created\_at



Create

# Chapter 2: Exploring the Discover Page



## Understanding the time filter

### Setting the time filter



kibana Discover Visualize Dashboard Settings Auto-refresh June 2nd 2015, 12:45:23.000 to June 2nd 2015, 13:32:22.000

Quick

From: June 3rd 2015, 06:07:41.833 To: Now

Relative  Minutes ago  Go

round to the minute

Absolute

kibana Discover Visualize Dashboard Settings Auto-refresh June 2nd 2015, 12:45:23.000 to June 2nd 2015, 13:32:22.000

Quick

From: June 3rd 2015, 06:08:00.000 To: Now

Relative  Minutes ago  Go

round to the minute

Absolute

YYYY defines the year in 4 digits (Ex: 2015)

MM defines the month in 2 digits (From 01-12)

DD defines the day in 2 digits (From 01-31)

HH defines the hours in 2 digits (From 00-23)

mm defines the minutes in 2 digits (From 00-59)

ss defines the seconds in 2 digits (From 00-59)

SSS defines milliseconds in 3 digits (From 000-999)

} Specify Date

} Specify Time

kibana Discover Visualize Dashboard Settings Auto-refresh June 2nd 2015, 12:45:23.000 to June 2nd 2015, 13:32:22.000

Quick

Relative

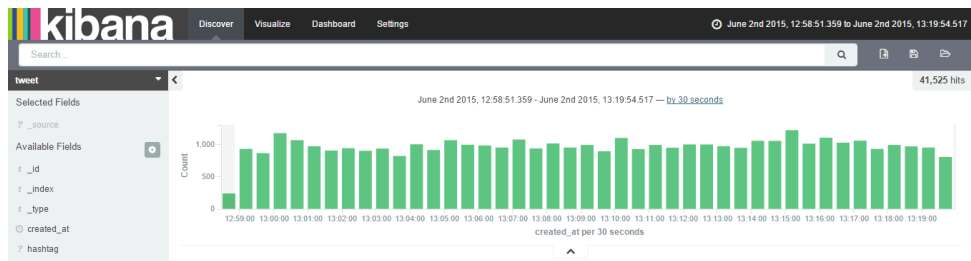
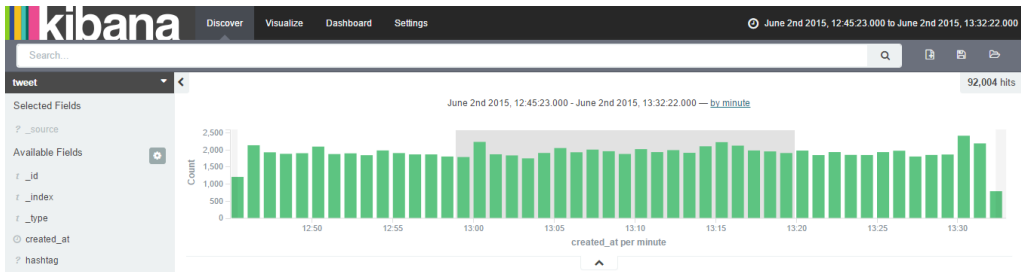
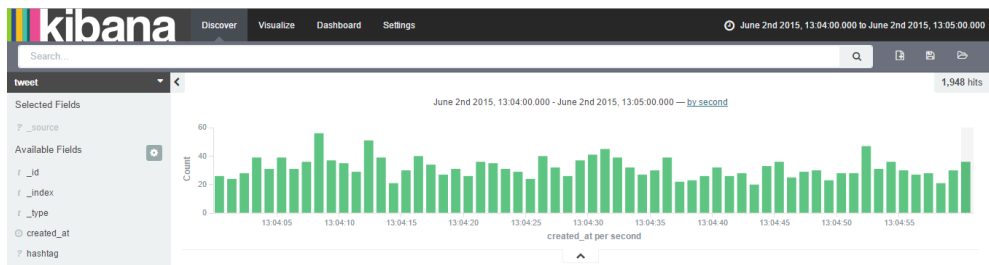
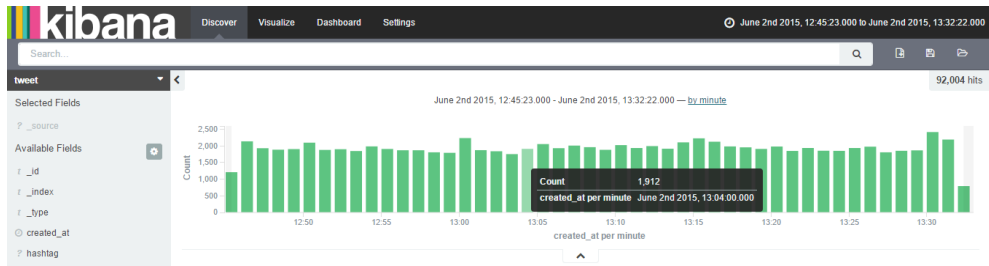
Absolute

From:  To:  Go

YYYY-MM-DD HH:mm:ss.SSS

June 2015							June 2015						
Sun	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Mon	Tue	Wed	Thu	Fri	Sat
31	01	02	03	04	05	06	31	01	02	03	04	05	06
07	08	09	10	11	12	13	07	08	09	10	11	12	13
14	15	16	17	18	19	20	14	15	16	17	18	19	20
21	22	23	24	25	26	27	21	22	23	24	25	26	27
28	29	30	01	02	03	04	28	29	30	01	02	03	04





kibana Discover Visualize Dashboard Settings Auto-refresh June 2nd 2015, 12:45:23.000 to June 2nd 2015, 13:32:22.000

Off	5 seconds	1 minute	1 hour
	10 seconds	5 minutes	2 hour
	30 seconds	15 minutes	12 hour
	45 seconds	30 minutes	1 day

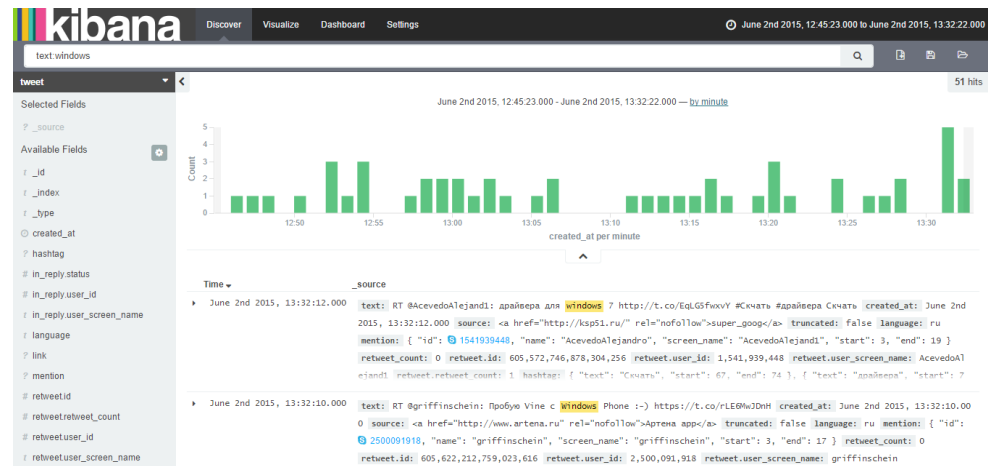
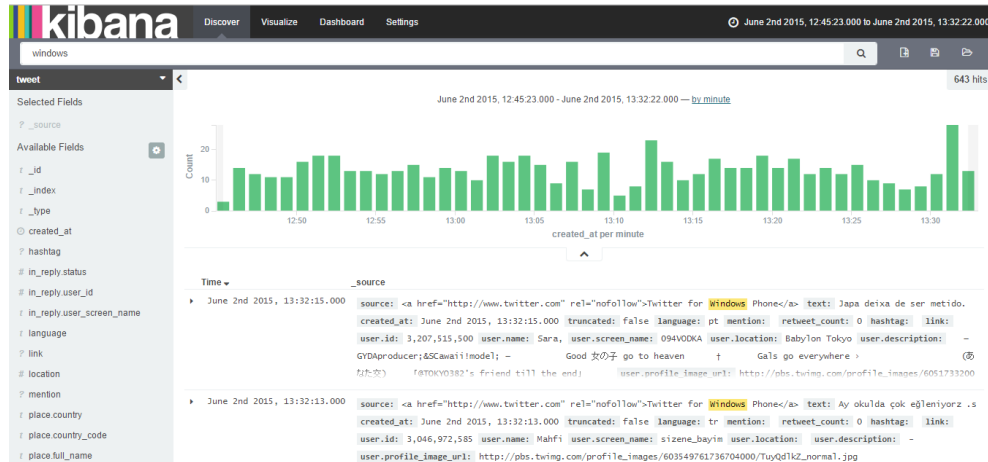
kibana Discover Visualize Dashboard Settings 10 seconds June 2nd 2015, 12:45:23.000 to June 2nd 2015, 13:32:22.000

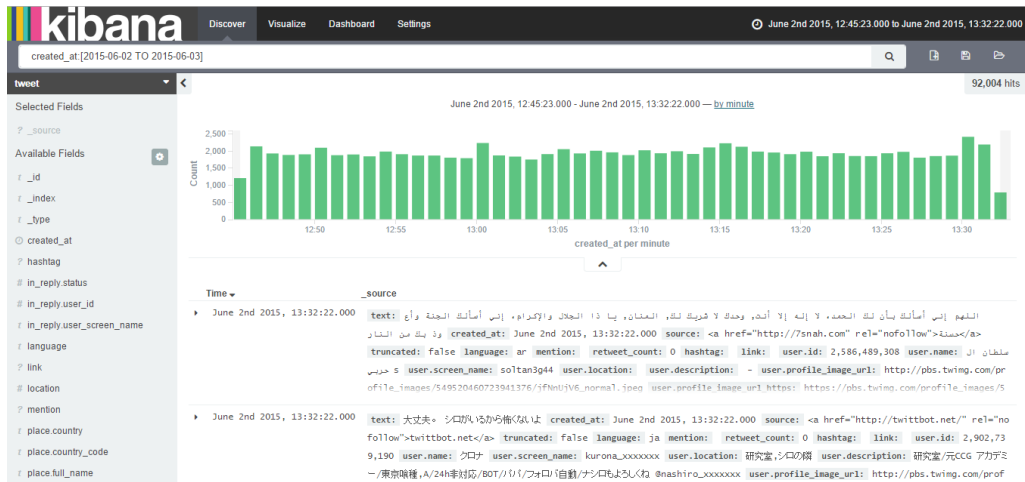
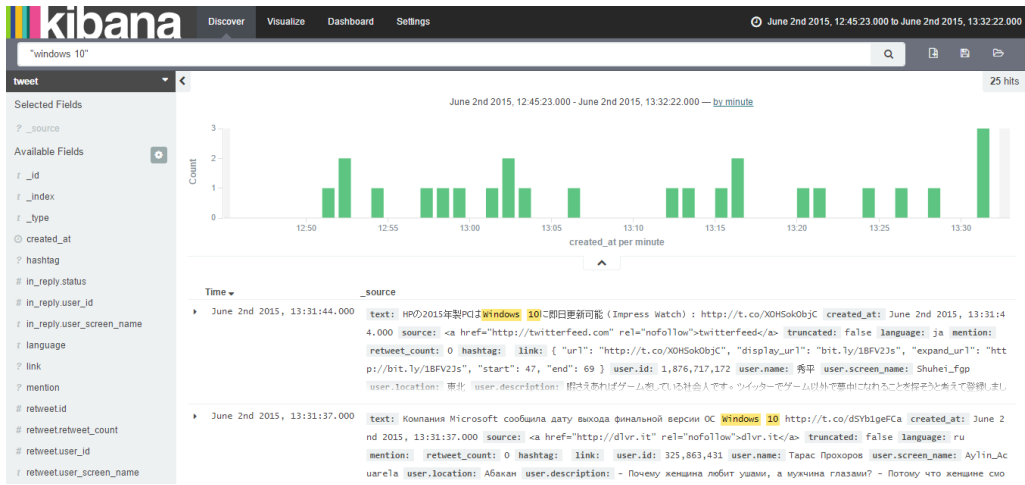
Off	5 seconds	1 minute	1 hour
	<b>10 seconds</b>	5 minutes	2 hour
	30 seconds	15 minutes	12 hour
	45 seconds	30 minutes	1 day

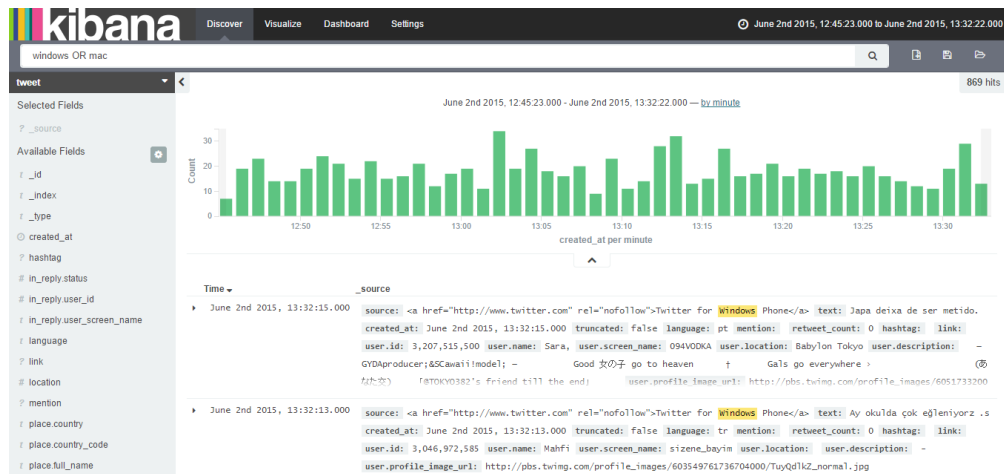
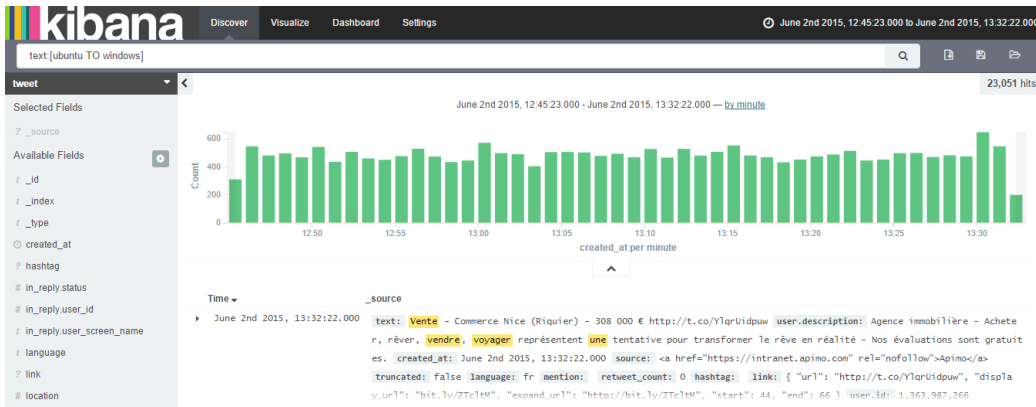
Refresh Interval Selected  
Option Displayed

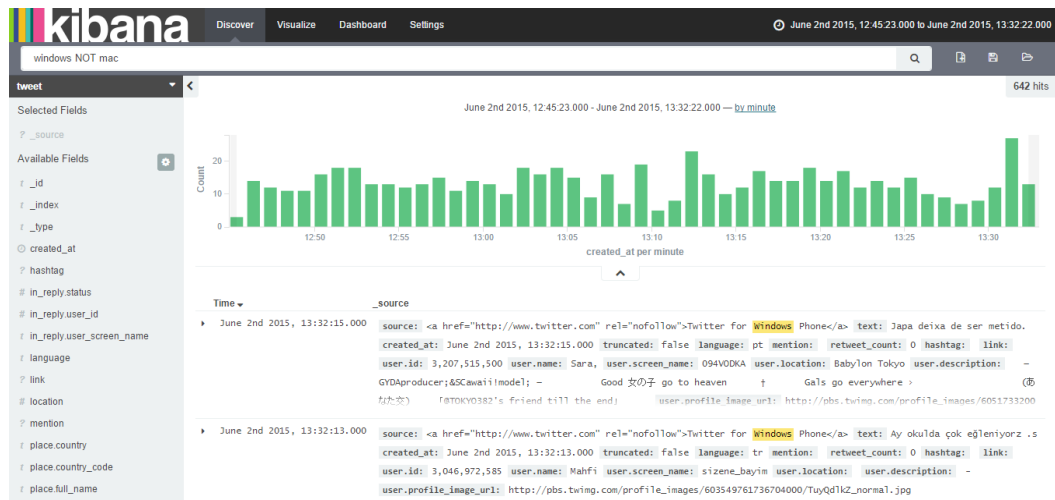
# Understanding the toolbar

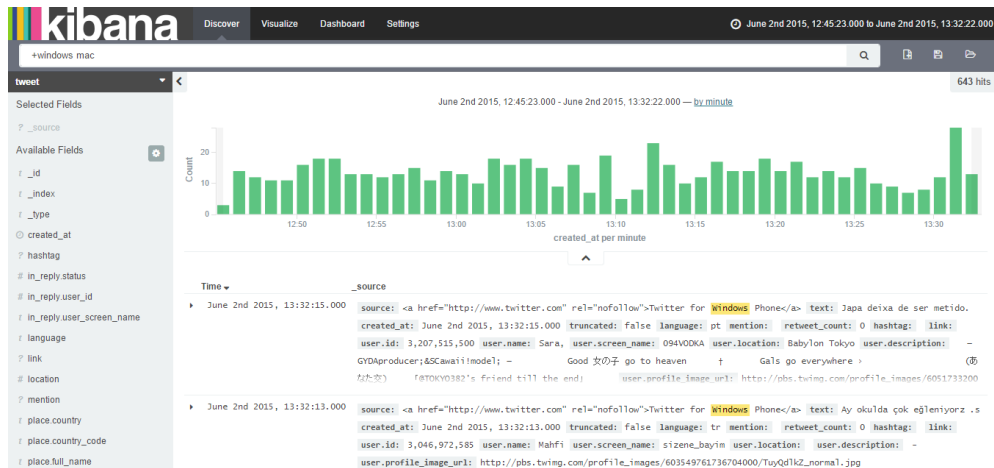
## Using the search bar







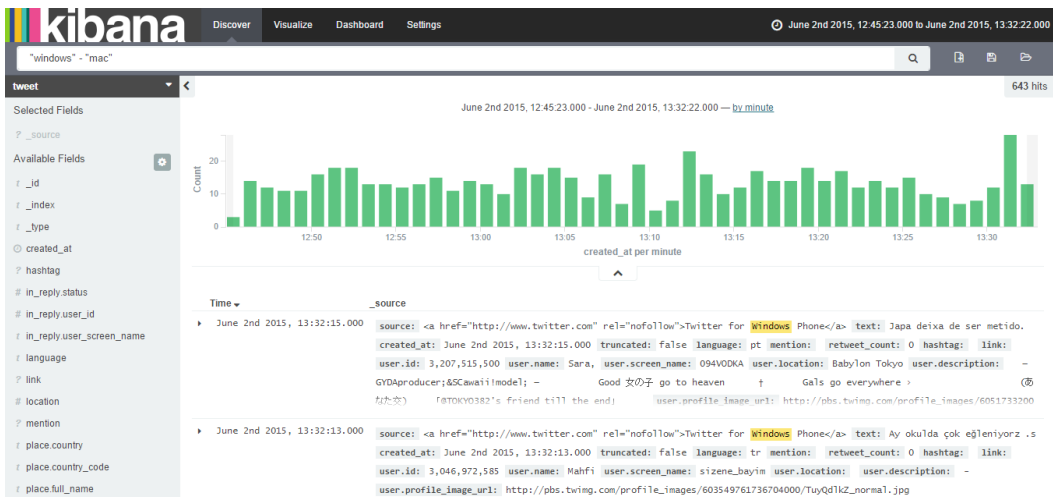


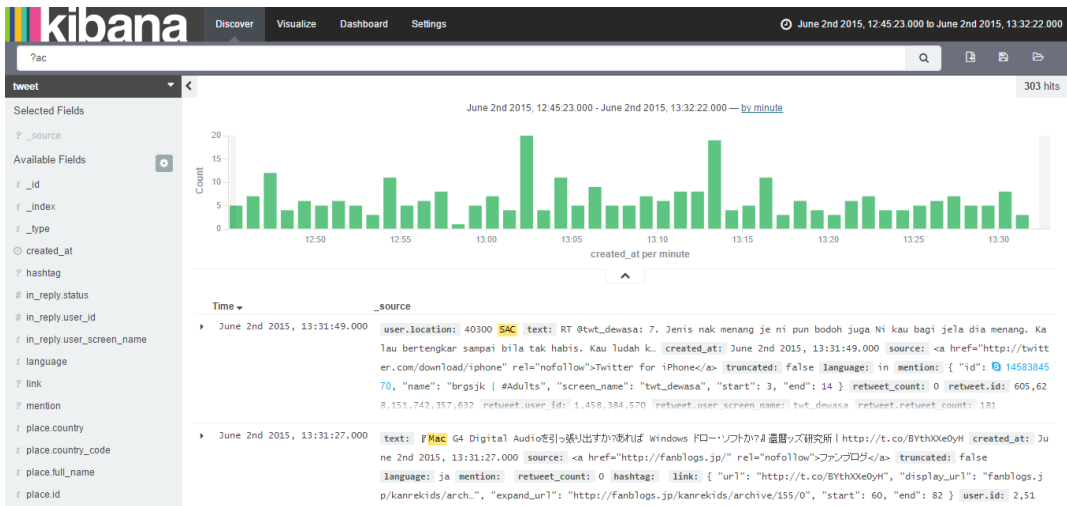
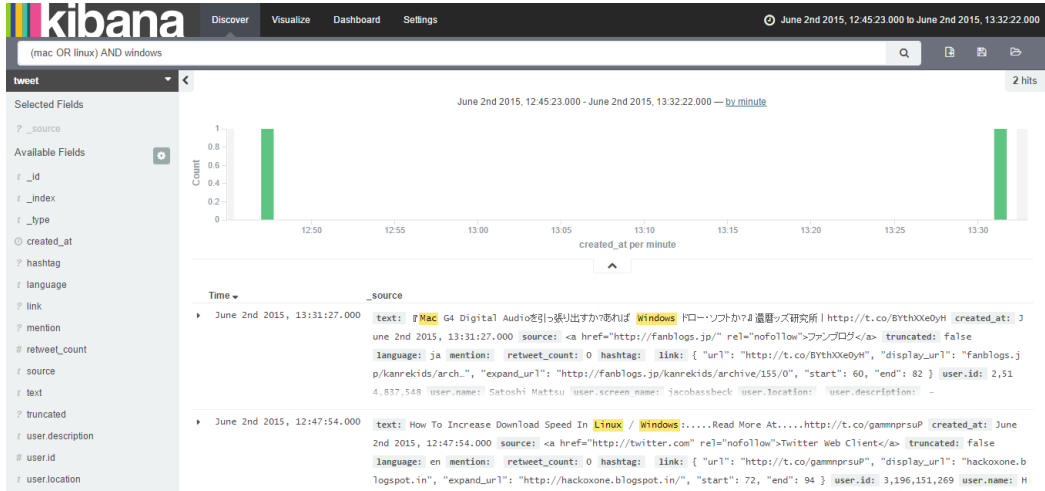


- June 2nd 2015, 13:31:28.000

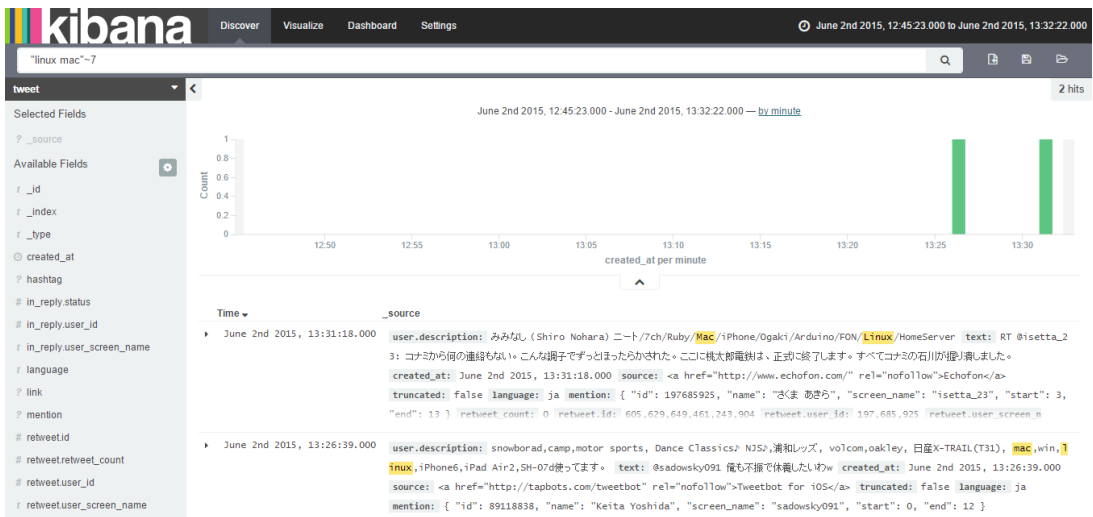
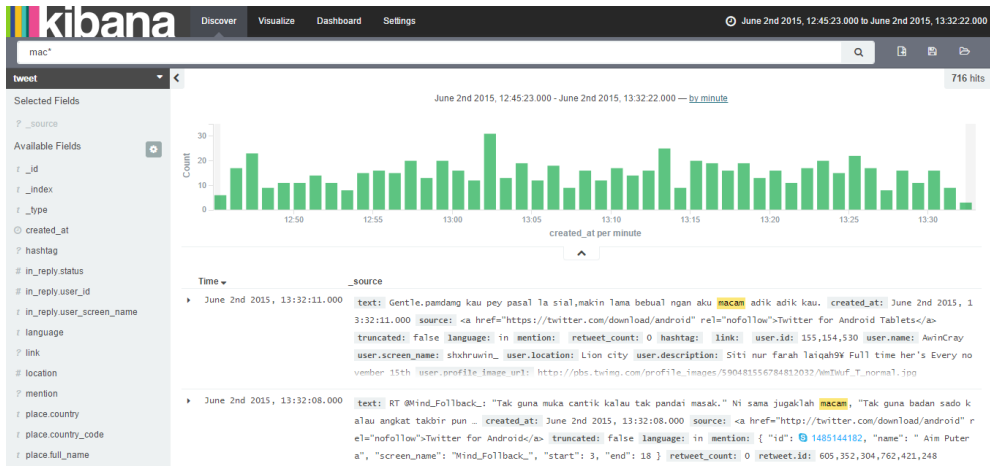
text: Software Update - including Windows Update. All optional updates will need manual update <http://t.co/tajSFQg0tx> <http://t.co/9Y8a0WskGP> created\_at: June 2nd 2015, 13:31:28.000 source: <a href="http://www.hootsuite.com" rel="nofollow">Hootsuite</a> truncated: false language: en mention: retweet\_count: 0 hashtag: link: { "url": "http://t.co/tajSFQg0tx", "display\_url": "ow.ly/NKzpu", "expand\_url": "http://ow.ly/NKzpu", "start": 89, "end": 111 } user.id: 18,796,389 user.name: Geoff Williamson user.screen\_name: ituk user.location: Crawley, England user.description: IT Engineer since 1
- June 2nd 2015, 13:31:27.000

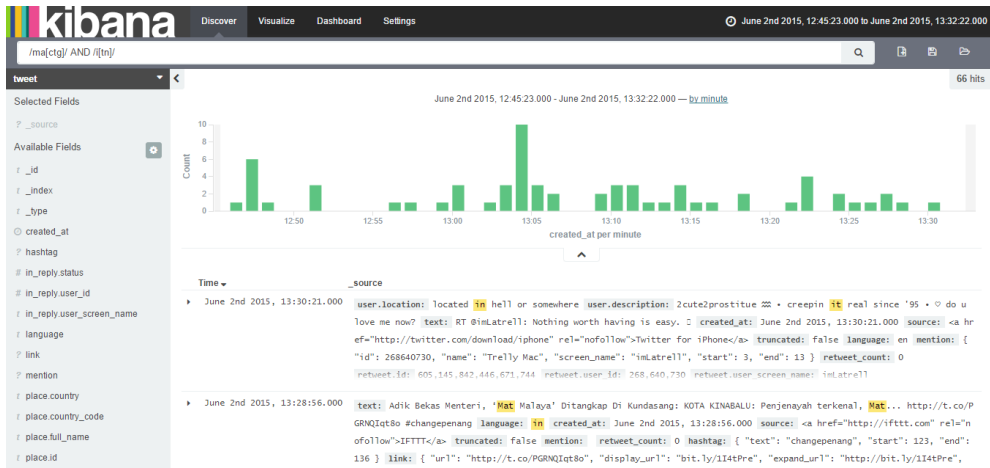
text: 『Mac G4 Digital Audioを引継ぎ出すか?』あれば Windows ドローソフトか?』 遺囑ツ研究所 | <http://t.co/BythXe0yH> created\_at: June 2nd 2015, 13:31:27.000 source: <a href="http://fanblogs.jp/" rel="nofollow">ファンブログ</a> truncated: false language: ja mention: retweet\_count: 0 hashtag: link: { "url": "http://t.co/BythXe0yH", "display\_url": "fanblogs.jp/kanrekids/arch...", "expand\_url": "http://fanblogs.jp/kanrekids/archive/155/0", "start": 60, "end": 82 } user.id: 2,514,837,548 user.name: Satoshi Matsu user.screen\_name: jacobassbeck user.location: user.description: -











The search bar contains the query `windows`. The interface shows the search type as `tweet` and a `New Search` button. The results count is `643 hits`.

The search bar contains the query `windows`. The interface shows the search type as `tweet` and a `Save Search` button. The results count is `643 hits`.

The "Save Search" dialog box is displayed. It contains a text input field with the value `Search_Twitter` and a `Save` button.

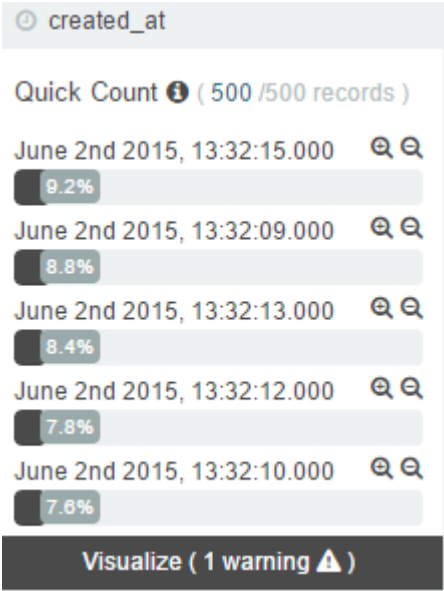
The search bar contains the query `windows`. The interface shows the search type as `tweet` and a `Load Saved Search` button.

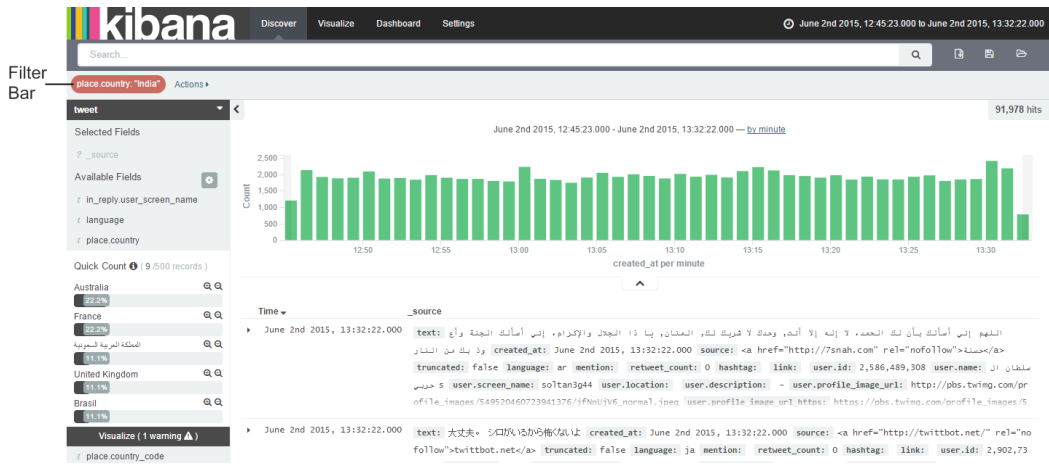
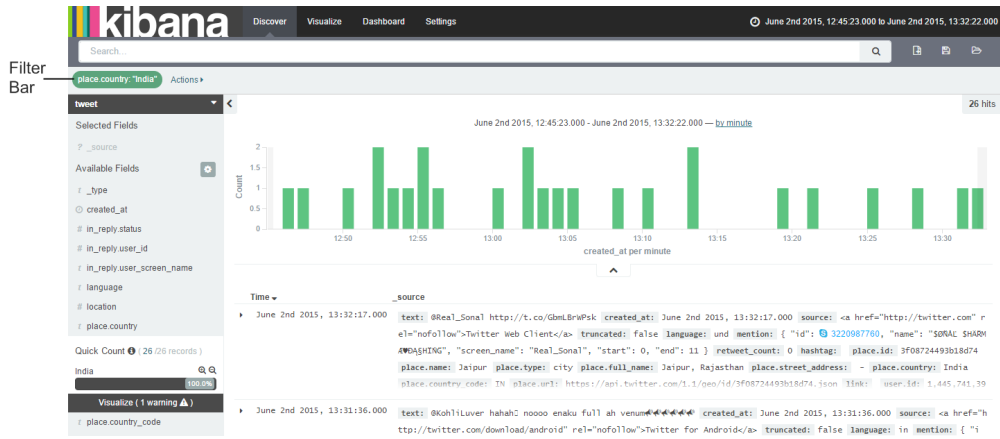
**kibana** Discover Visualize Dashboard Settings June 2nd 2015, 12:45:23.000 to June 2nd 2015, 13:32:22.000

windows

Saved Search Filter manage searches

Search\_Twitter 1 searches





Time

source

June 2nd 2015, 13:32:22.000

```

text: الفهم إلى أماتك بأن لك الحمد، لا إله إلا أنت، وحده لا شريك لك، العنان، يا ذا الجلال والإكرام، إلى أماتك الجنة وأج
المنار <a href="http://tsnah.com" rel="nofollow">حسنة</a> created_at: June 2nd 2015, 13:32:22.000 source: <a href="http://tsnah.com" rel="nofollow">حسنة</a>
truncated: false language: ar mention: retweet_count: 0 hashtag: link: user.id: 2,586,489,308 user.name: سلطان ال
سلطان ال user.screen_name: soltan3944 user.location: user.description: - user.profile_image_url: http://pbs.twimg.com/pr
ofile_images/549520460723941376/5fNnUjV6_normal.jpg user.profile_image_url_https: https://pbs.twimg.com/profile_images/5

```

[Link to /tweet/tweet/605645619340795904](http://twitter.com/tweet/tweet/605645619340795904)

Table	JSON
_id	605645619340795904
_index	tweet
_type	tweet
created_at	June 2nd 2015, 13:32:22.000


place.country: "India" Actions ▾

All filters: Enable Disable Pin Unpin Invert Toggle Remove

### The Disable filter

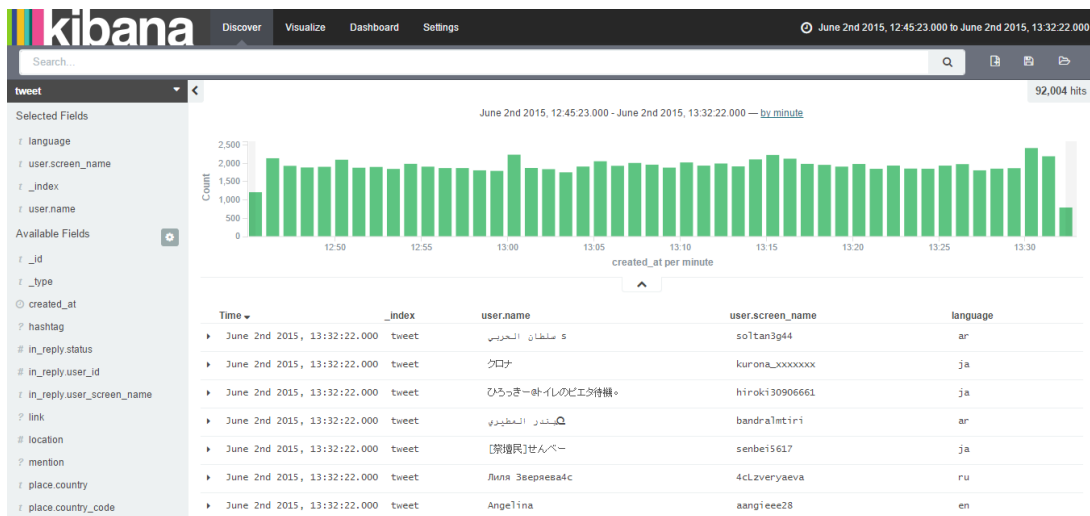
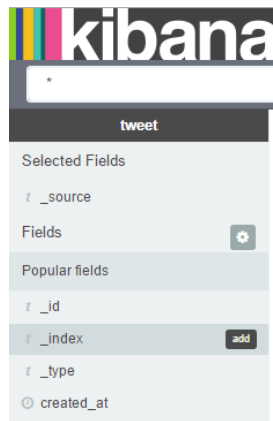


### The Invert filter

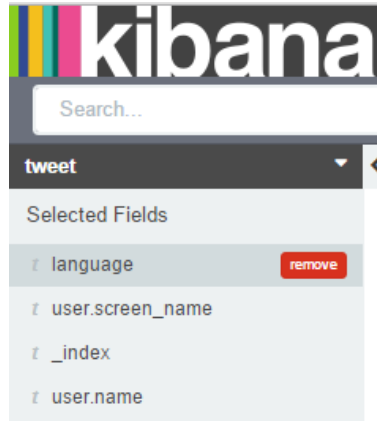


# Understanding document data

## Add field to document data



## Remove field from document data



## View data

	Time	_index	user.name	user.screen_name	language
	June 2nd 2015, 13:32:22.000	tweet	سلطان الحريسي	so1tan3g44	ar
					<a href="#">Link to /tweet/tweet/605645619340795904</a>
					<a href="#">Table</a> <a href="#">JSON</a>
# _id			605645619340795904		
# _index		tweet			
# _type		tweet			
created_at	June 2nd 2015, 13:32:22.000				
# hashtag					
# language				ar	
# link					
# mention					
# retweet_count				0	
# source				<a href="http://7snah.com" rel="nofollow">حسنة</a>	
# text				السلام إن شاء الله الحمد، لا إله إلا أنت، وحده لا شريك له، المنان، يا ذا الجلال والإكرام. إن شاء الله الجنة وأعوذ بك من النار	
# truncated				false	
# user.description				-	
# user.id				2,586,489,308	

Time	_index	user.name	user.screen_name	language
June 2nd 2015, 13:32:22.000	tweet	س سلطان الحريسي	so1tan3g44	ar

Table	JSON
1	{
2	"index": "tweet",
3	"type": "tweet",
4	"id": "605645619340795904",
5	"score": null,
6	"source": {
7	"text": "الهدم إنتي أمأنتك بآن لك الحمد، لا إله إلا أنت وحده لا شريك لك، المشانق بنا ذا الصلح والفرح، إنتي أمأنتك الحمد وأمر بك من الشكر",
8	"created_at": "2015-06-02T13:32:22.000Z",
9	"source": { "ce_href": "http://5snah.com" "rel": "nofollow">خمسة</ce_href>"},
10	"truncated": false,
11	"language": "ar",
12	"mentions": [],
13	"retweet_count": 0,
14	"hashtags": [],
15	"links": [],
16	"user": {
17	"id": "3598489308",
18	"name": "س سلطان الحريسي",
19	"screen_name": "so1tan3g44",
20	"location": "",
21	"description": null,
22	"profile_image_url": "http://pbs.twimg.com/profile_images/549520460723041376/2fhu03W_normal.jpeg",
23	"profile_image_url_https": "https://pbs.twimg.com/profile_images/549520460723041376/2fhu03W_normal.jpeg",
24	}
25	}
26	},
27	"fields": {
28	"created_at": {
29	"type": "date",
30	"value": "2015-06-02T13:32:22.000Z",
31	}
32	"sort": {
33	"type": "number",
34	"value": "143332143000",
35	}

[Link to /tweet/tweet/605645619340795904](#)

Table	JSON
1	{
2	"id": "605645619340795904",
3	"score": null,
4	"source": {
5	"text": "الهدم إنتي أمأنتك بآن لك الحمد، لا إله إلا أنت وحده لا شريك لك، المشانق بنا ذا الصلح والفرح، إنتي أمأنتك الحمد وأمر بك من الشكر",
6	"created_at": "2015-06-02T13:32:22.000Z",
7	"source": { "ce_href": "http://5snah.com" "rel": "nofollow">خمسة</ce_href>"},
8	"truncated": false,
9	"language": "ar",
10	"mentions": [],
11	"retweet_count": 0,
12	"hashtags": [],
13	"links": [],
14	"user": {
15	"id": "3598489308",
16	"name": "س سلطان الحريسي",
17	"screen_name": "so1tan3g44",
18	"location": "",
19	"description": null,
20	"profile_image_url": "http://pbs.twimg.com/profile_images/549520460723041376/2fhu03W_normal.jpeg",
21	"profile_image_url_https": "https://pbs.twimg.com/profile_images/549520460723041376/2fhu03W_normal.jpeg",
22	}
23	}
24	},
25	"fields": {
26	"created_at": {
27	"type": "date",
28	"value": "2015-06-02T13:32:22.000Z",
29	}
30	"sort": {
31	"type": "number",
32	"value": "143332143000",
33	}
34	},

localhost:5601/#/doc/tweet/tweet/tweet?id=605645619340795904&\_g=()

## Sorting document

created_at	Time	_index	user.name	user.screen_name	language
	June 2nd 2015, 13:32:22.000	tweet	س سلطان الحريسي	so1tan3g44	ar

Sort by user name

## Moving fields in document data

user.description	Time	_index	user.name	user.screen_name	language
	June 2nd 2015, 13:32:22.000	tweet	س سلطان الحريسي	so1tan3g44	ar

Move column to the right

user.description	Time	_index	user.name	user.screen_name	language
	June 2nd 2015, 13:32:22.000	tweet	س سلطان الحريسي	so1tan3g44	ar

Move column to the left

user.description	Time	_index	user.name	user.screen_name	language
	June 2nd 2015, 13:32:22.000	tweet	س سلطان الحريسي	so1tan3g44	ar

Move column to the left

user.description	Time	_index	user.name	user.screen_name	language
	June 2nd 2015, 13:32:22.000	tweet	س سلطان الحريسي	so1tan3g44	ar

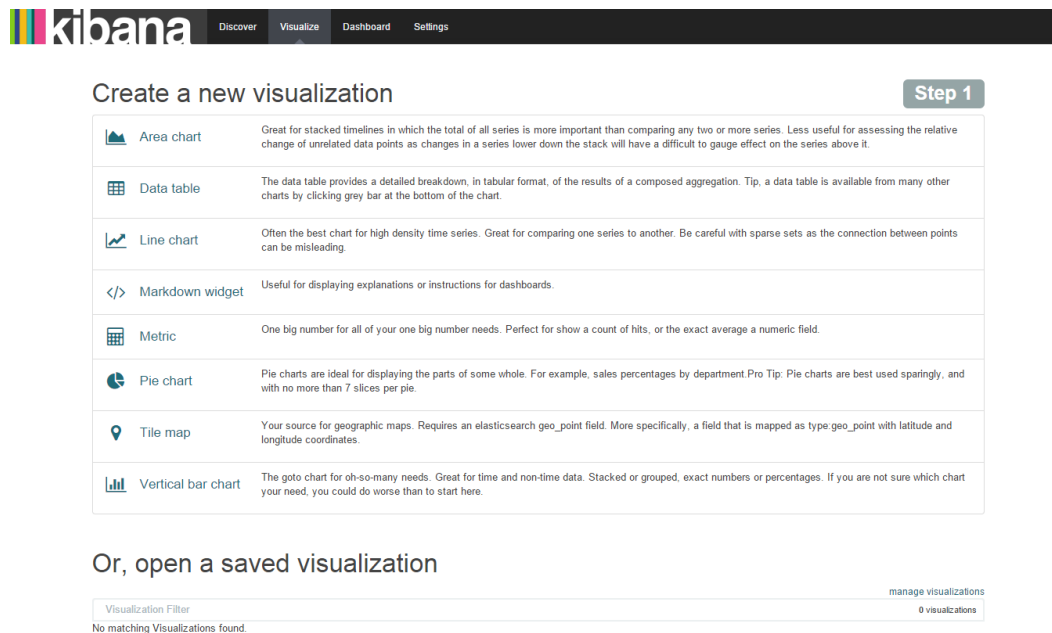
Move column to the right



# Chapter 3: Exploring the Visualize Page









## Steps for designing visualization

### Step 1 – selecting a visualization type



**Step 1**

Create a new visualization

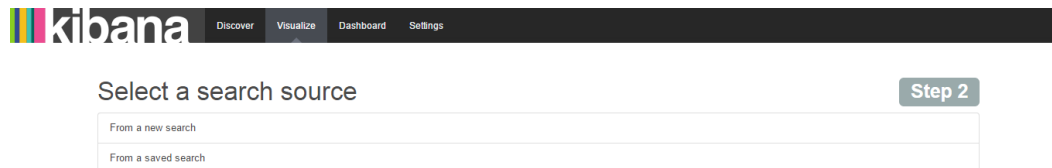
 Area chart	Great for stacked timelines in which the total of all series is more important than comparing any two or more series. Less useful for assessing the relative change of unrelated data points as changes in a series lower down the stack will have a difficult to gauge effect on the series above 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 show 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.
 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 your need, you could do worse than to start here.

Or, open a saved visualization

Visualization Filter manage visualizations  
0 visualizations

No matching Visualizations found.

### Step 2 – selecting search data source



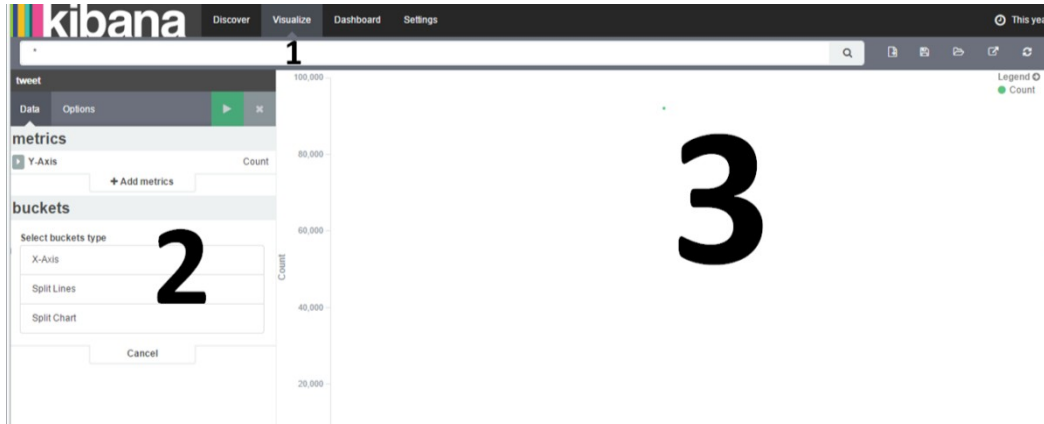
**Step 2**

Select a search source

From a new search

From a saved search

## Step 3 – the visualization canvas




### Toolbar

#### New Visualization



#### Save Visualization



## Load Saved Visualization



## Share Visualization



Embed this visualization. Add to your html source. Note all clients must still be able to access kibana

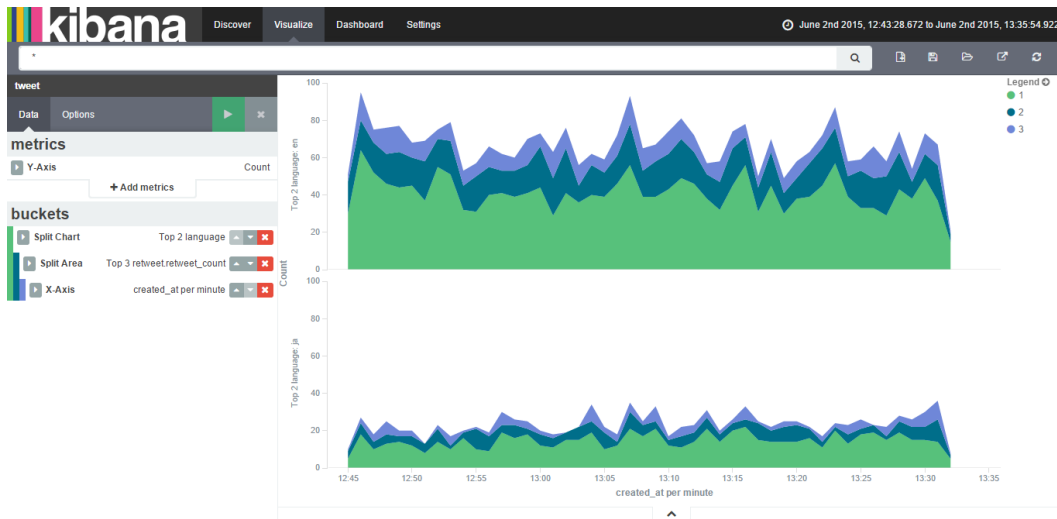
```
<iframe src="http://localhost:5601/#/visualize/create?type=area&indexPattern=tweet&_g=(refreshInterval:(display:Off,pause:1f,section:0,value:0),time:(from:now%2Fy,mode:quick,to:now%2Fy))&_a=(filters:1(),linked:1f,query:(query_string:(analyze_wildcard:1t,query:"))),vis:(aggs:1!(id:1,params:(),schema:metric,type:count)),listeners:(),params:(addLegend:1t,addTimeMarker:1f,addTooltip:1t,defaultYExtents:1f,interpolate:linear,mode:stacked,scale:linear,setYExtents:1f,shareYAxis:1t,smoothLines:1f,times:1!(,yAxis:()),type:area))" height="600" width="800"></iframe>
```

Share a link


```
http://localhost:5601/#/visualize/create?type=area&indexPattern=tweet&_g=(refreshInterval:(display:Off,pause:1f,section:0,value:0),time:(from:now%2Fy,mode:quick,to:now%2Fy))&_a=(filters:1!(,linked:1f,query:(query_string:(analyze_wildcard:1t,query:"))),vis:(aggs:1!(id:1,params:(),schema:metric,type:count)),listeners:(),params:(addLegend:1t,addTimeMarker:1f,addTooltip:1t,defaultYExtents:1f,interpolate:linear,mode:stacked,scale:linear,setYExtents:1f,shareYAxis:1t,smoothLines:1f,times:1!(,yAxis:()),type:area))
```

# Explanation of visualization types

## Area Chart




## Overlap



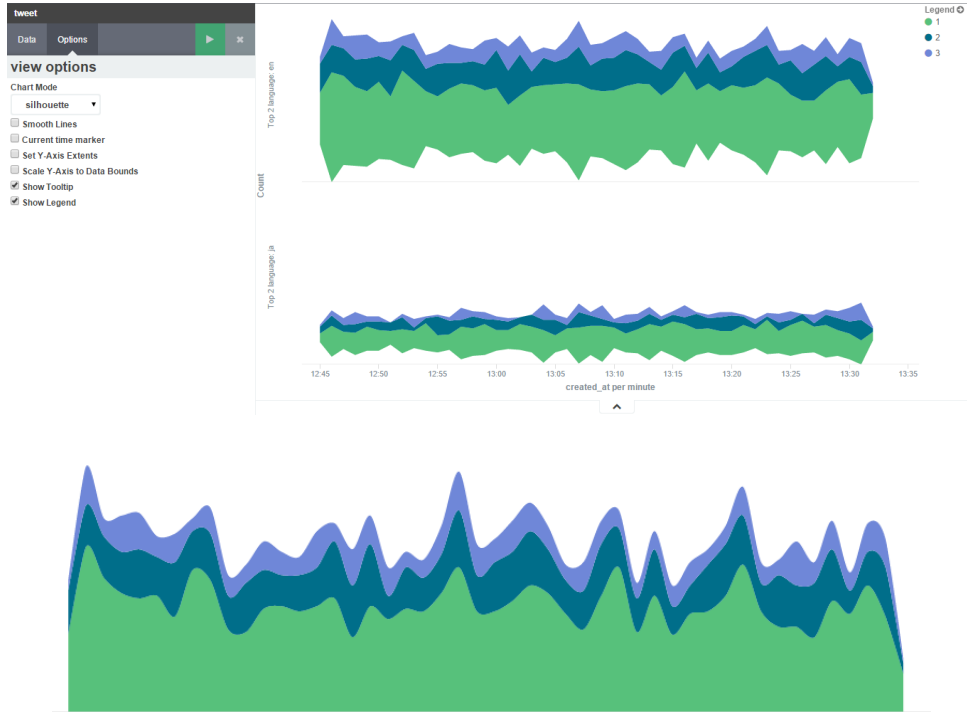
# Percentage



# Wiggle



# Silhouette



# Data Table

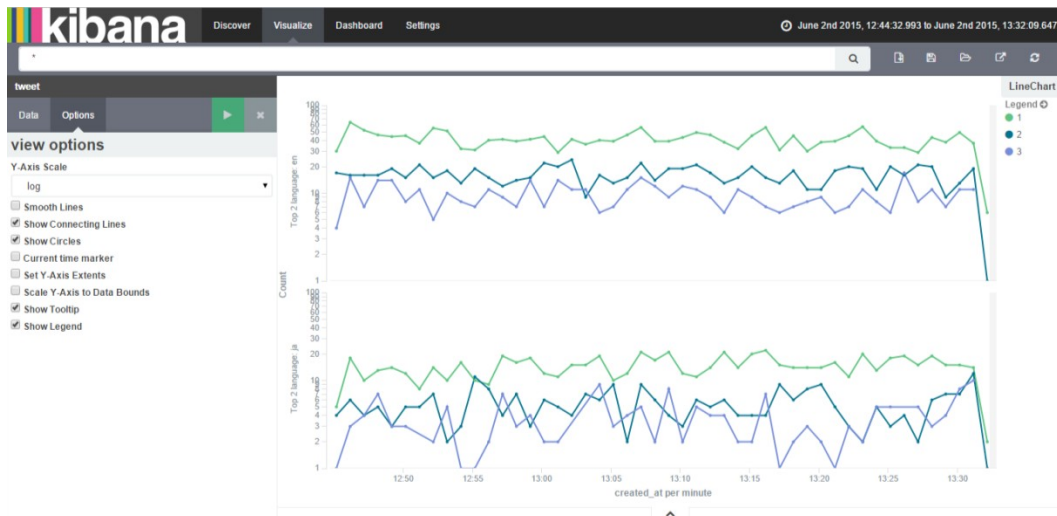
Export: [Raw](#) [Formatted](#)

Top 5 language	retweet.retweet_count ranges	Count
en	0	4,354
en	10	754
ja	0	1,575
ja	10	395
ar	0	1,160
ar	10	230
in	0	554
in	10	125
es	0	727
es	10	118

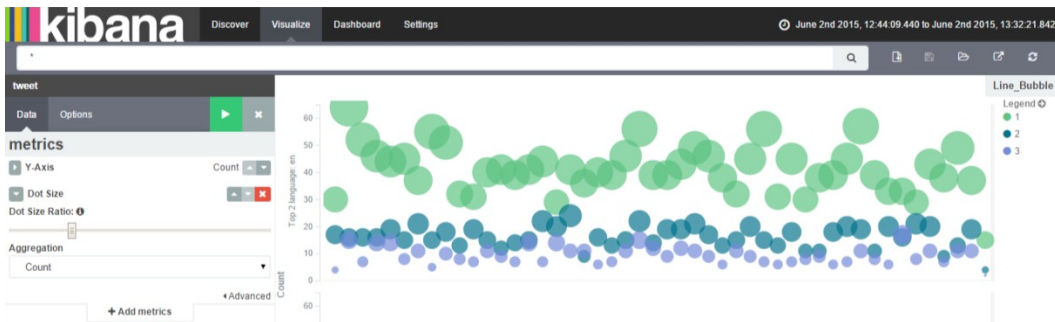

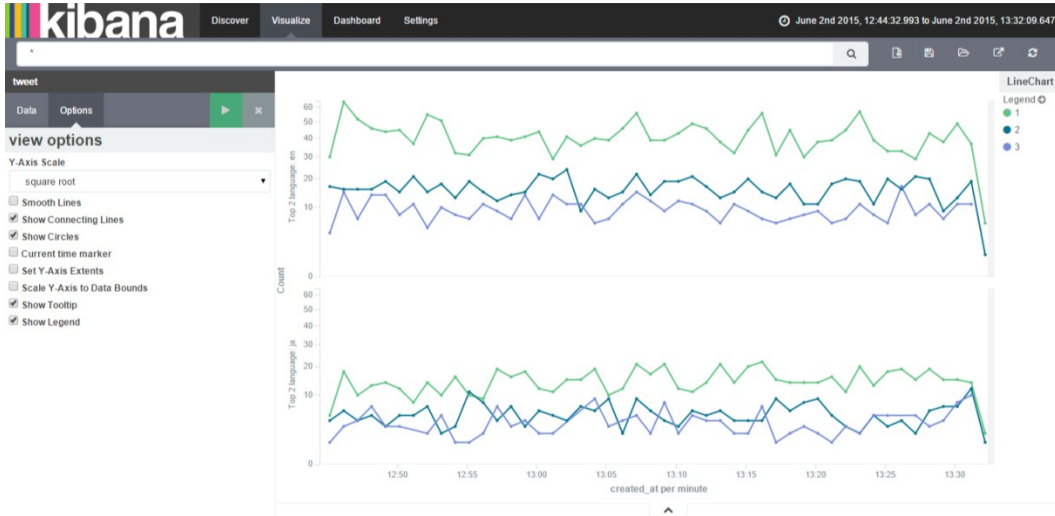
# Line Chart



# Log

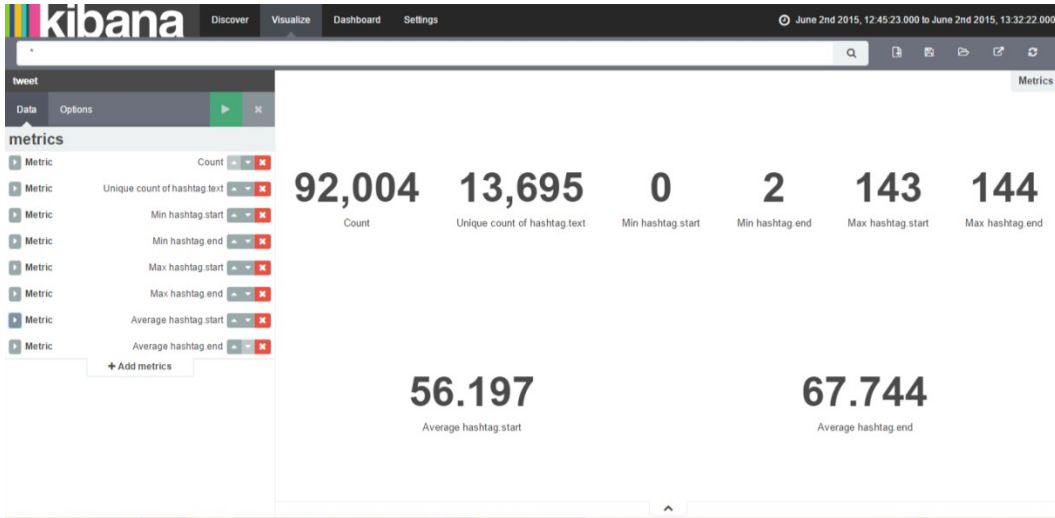


# Square root

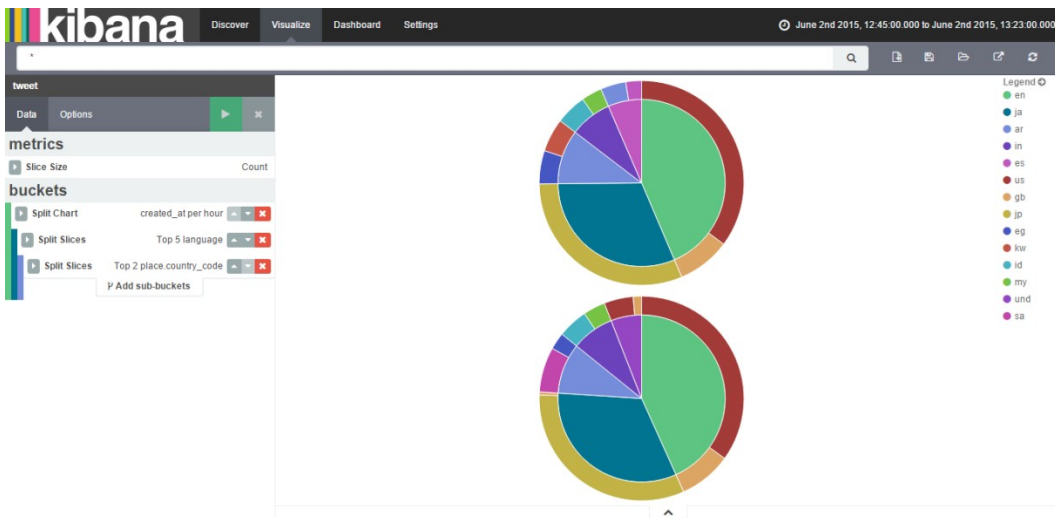






# Metric



# Pie Chart






tweet

Data Options

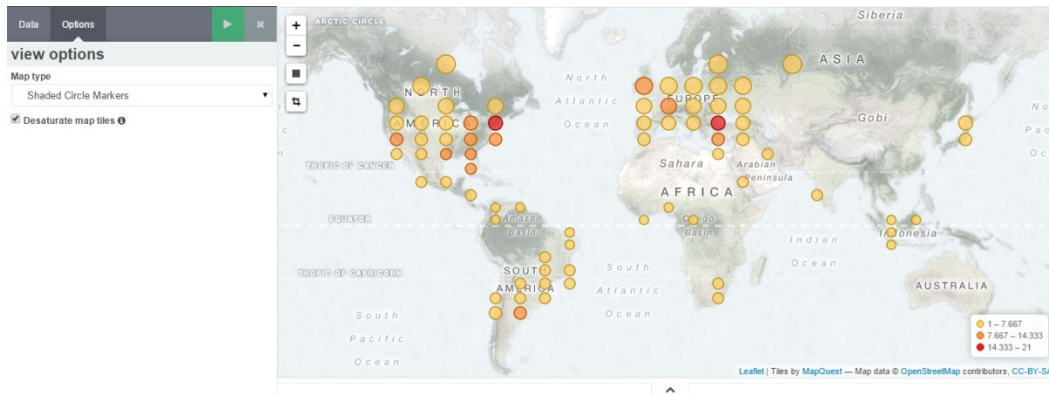
view options

- Donut
- Show Tooltip
- Show Legend

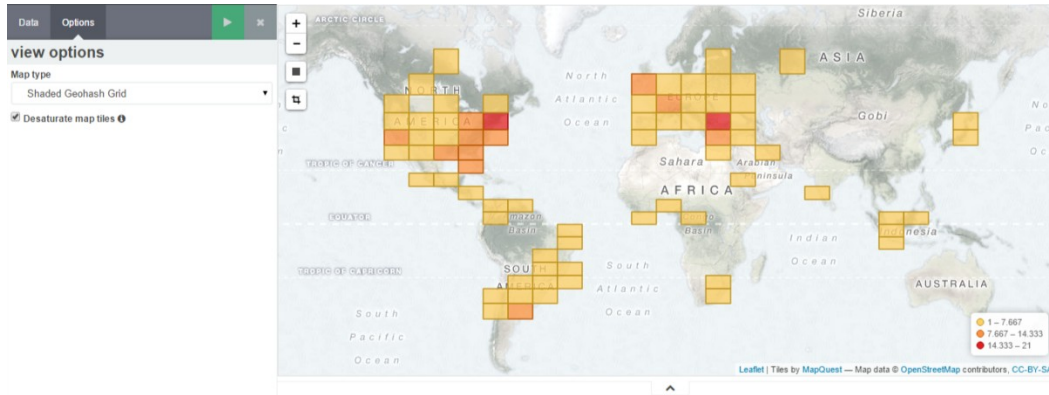


- Legend
- en
  - ja
  - ar
  - in
  - es
  - us
  - gb
  - jp
  - sa
  - eg
  - id
  - my

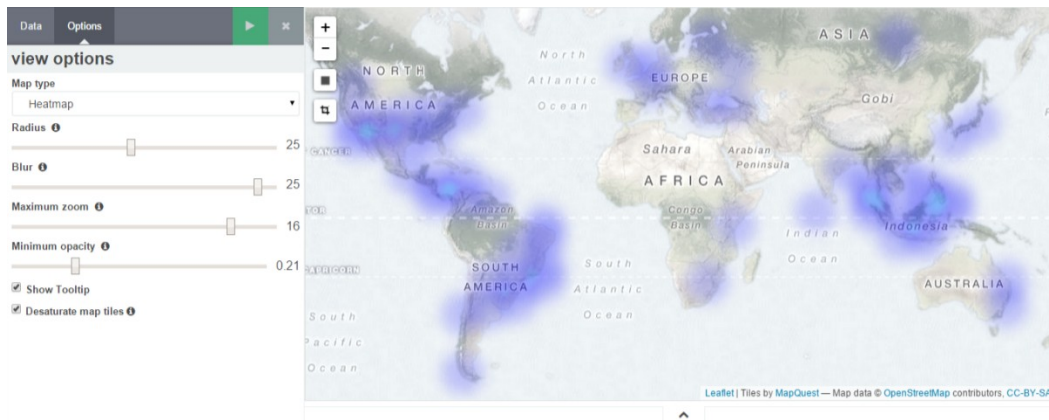
# Tile Map



## Shaded GeoHash Grid




## Desaturate map tiles




# Vertical Bar Chart



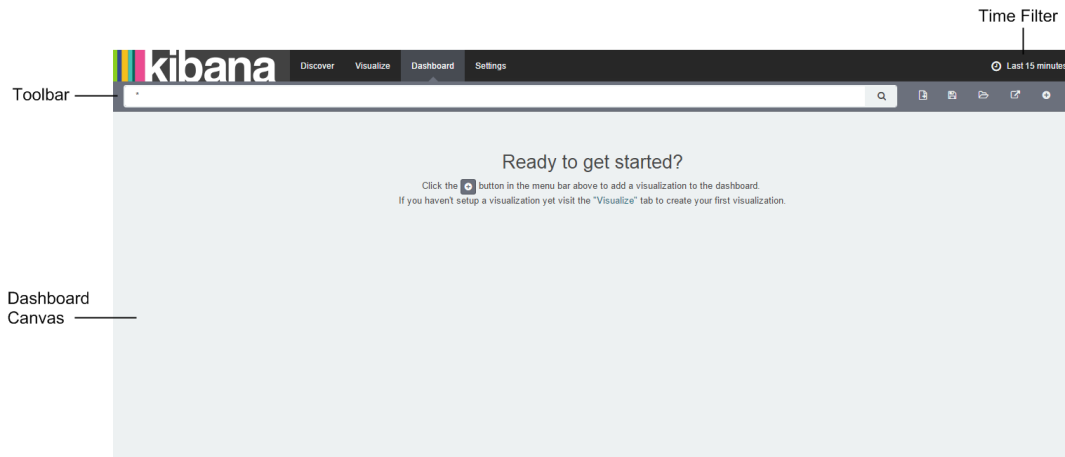
# Percentage



# Grouped




## Chapter 4: Exploring the Dashboard Page



## Understanding the toolbar

### The New Dashboard option



### Adding visualizations



**kibana** Discover Visualize Dashboard Settings June 2nd 2015, 12:45:00.000 to June 2nd 2015, 13:33:00.000

Visualizations Searches

Visualization Filter manage visualizations  
8 visualizations

- AreaChart
- BarChart
- DataTable
- Line\_Bubble
- LineChart

1 2 >

**kibana** Discover Visualize Dashboard Settings June 2nd 2015, 12:45:00.000 to June 2nd 2015, 13:33:00.000

AreaChart Legend

Count Top... June 2nd... created\_at per minute

PieChart Legend

Count Top... June 2nd... Legend

LineChart Legend

Count Top... June 2nd... created\_at per minute

BarChart Legend



Count Top... June 2nd... Top 5 language

Line\_Bubble Legend


Count Top... June 2nd... created\_at per minute




## The Save Dashboard option



## The Load Saved Dashboard option



## Sharing the saved dashboard

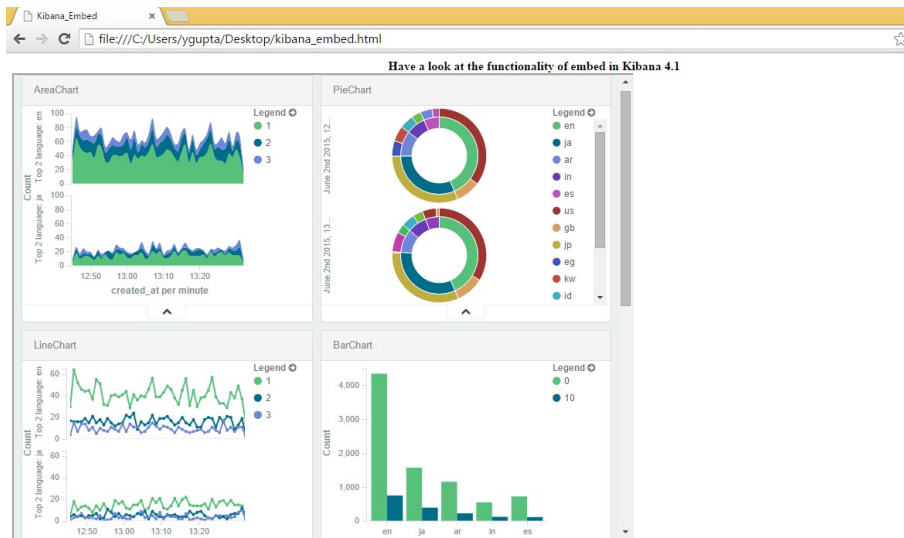


# Understanding the dashboard canvas

## Removing visualizations

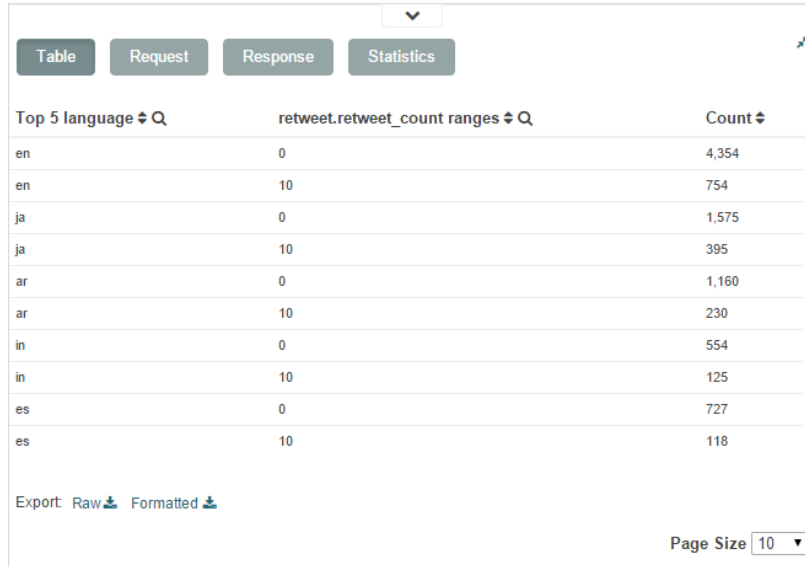


# Embedding dashboard in a web page



# Understanding the debug panel

## Table

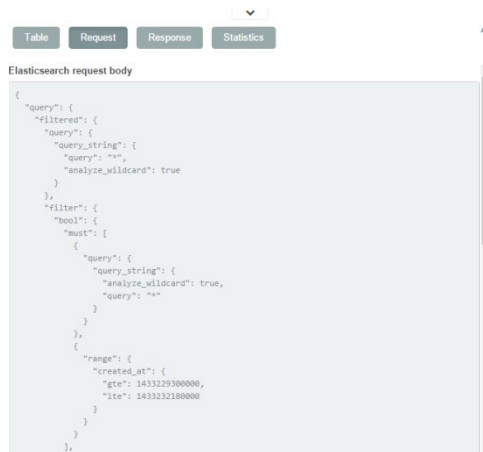


The screenshot shows a debug panel with a table of results. At the top, there are tabs for 'Table', 'Request', 'Response', and 'Statistics'. The 'Table' tab is selected. The table has three columns: 'Top 5 language', 'retweet.retweet\_count ranges', and 'Count'. The data is as follows:

Top 5 language	retweet.retweet_count ranges	Count
en	0	4,354
en	10	754
ja	0	1,575
ja	10	395
ar	0	1,160
ar	10	230
in	0	554
in	10	125
es	0	727
es	10	118

At the bottom of the table, there are options to 'Export Raw' and 'Formatted'. On the right side, there is a 'Page Size' dropdown menu set to '10'.

## Request



The screenshot shows a debug panel with tabs for 'Table', 'Request', 'Response', and 'Statistics'. The 'Request' tab is selected. The panel displays the 'Elasticsearch request body' as a JSON object:

```
{
  "query": {
    "filtered": {
      "query": {
        "query_string": {
          "query": "*",
          "analyze_wildcard": true
        }
      },
      "filter": {
        "bool": {
          "must": [
            {
              "query": {
                "query_string": {
                  "analyze_wildcard": true,
                  "query": "*"
                }
              }
            }
          ]
        }
      },
      "range": {
        "created_at": {
          "gte": 1433229000000,
          "lte": 1433221800000
        }
      }
    }
  }
}
```

## Response

Table Request Response Statistics

Elasticsearch response body

```
{
  "took": 26,
  "timed_out": false,
  "_shards": {
    "total": 5,
    "successful": 5,
    "failed": 0
  },
  "hits": {
    "total": 92004,
    "max_score": 0,
    "hits": []
  },
  "aggregations": {
    "2": {
      "doc_count_error_upper_bound": 113,
      "sum_other_doc_count": 23423,
      "buckets": [
        {
          "3": {
            "buckets": {
              "0.0-10.0": {
                "from": 0,
                "from_as_string": "0.0",
                "to": 10,
                "to_as_string": "10.0",
                "doc_count": 4354
              },
              "10.0-20.0": {
```

## Statistics

Table Request Response Statistics

Query Duration	26ms
Request Duration	327ms
Hits	92004
Index	"tweet"

# Chapter 5: Exploring the Settings Page

## Indices

### Configuring an index pattern

The screenshot shows the Kibana Settings page for configuring an index pattern. The page is titled "Configure an index pattern" and includes a sidebar with "Indices" and "Index Patterns" (with a star next to "tweet"). The main content area contains the following elements:

- Index Patterns:** A list showing "★ tweet" and "kibana-\*".
- Configuration Form:**
  - Index contains time-based events
  - Use event times to create index names
  - Index name or pattern:** A text input field containing "logstash-\*". Below it, a message states: "Unable to fetch mapping. Do you have indices matching the pattern?"
- Advanced Configuration Form:**
  - Index contains time-based events
  - Use event times to create index names
  - Index pattern interval:** A dropdown menu set to "Daily".
  - Index name or pattern:** A text input field containing "[logstash-]YYYY.MM.DD". Below it, a message states: "Pattern does not match any existing indices".

Index contains time-based events

Use event times to create index names

**Index pattern interval** ⓘ

Daily

**Index name or pattern**

Patterns allow you to define dynamic index names. Static text in an index name is denoted using brackets. Example: [logstash-]YYYY.MM.DD. Please note that weeks are setup to use ISO weeks which start on Monday. — [Date Format Documentation](#)

[kibana1-]YYYY-MM-DD


Pattern matches 100% of existing indices and aliases

- kibana1-2015-06-02
- kibana1-2015-07-27
- kibana1-2015-08-12

**Time-field name** ⓘ refresh fields

created\_at

Create



**kibana** Discover Visualize Dashboard Settings

Indices Advanced Objects About

Index Patterns +Add New

★ tweet

[kibana1-]YYYY-MM-DD

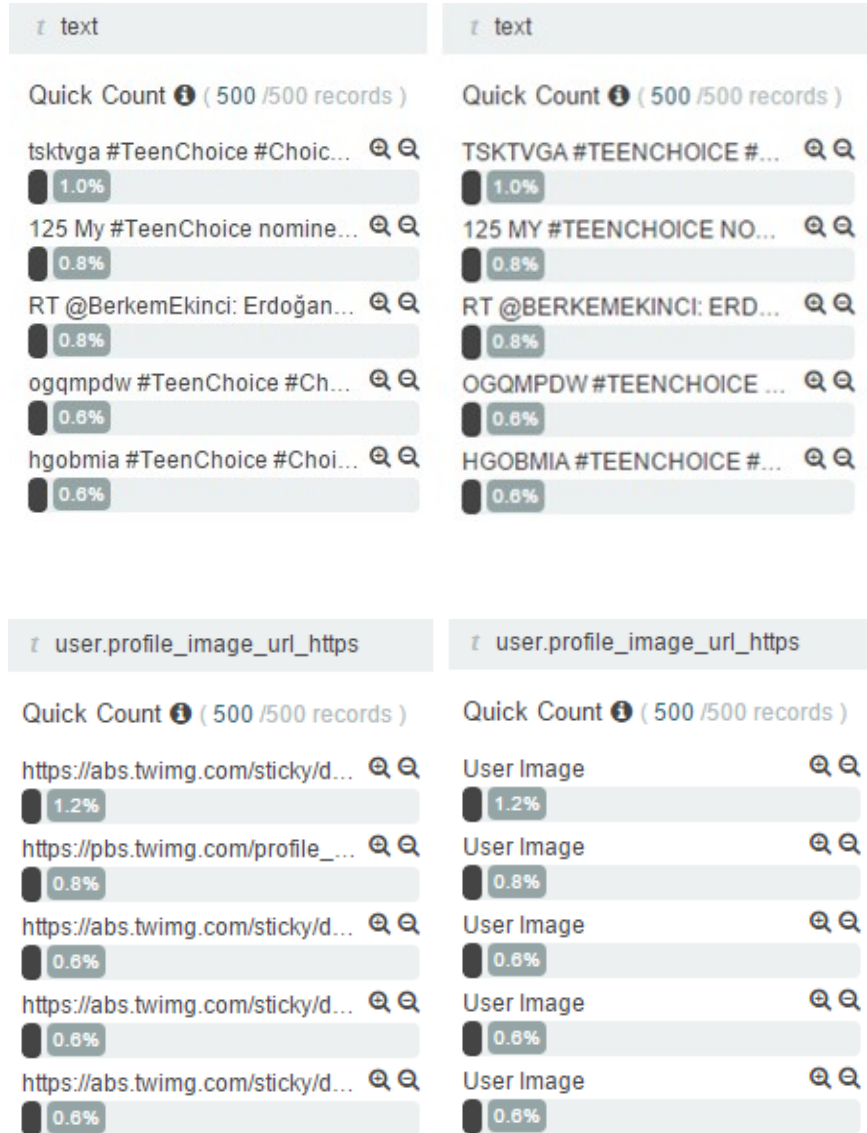
kibana-\*

This page lists every field in the **tweet** index and the field's associated core type as recorded by Elasticsearch. While this list allows you to view the core type of each field, changing field types must be done using Elasticsearch's Mapping API %

name	type	format	analyzed	indexed	controls
mention.name	string		✓	✓	🔍
mention.screen_name	string		✓	✓	🔍
language	string		✓	✓	🔍
retweet.id	number			✓	🔍
in_reply.status	number			✓	🔍
source	string		✓	✓	🔍
retweet.count	number			✓	🔍
mention.end	number			✓	🔍
place.country_code	string		✓	✓	🔍

# Managing the field properties

## The field type format





created\_at

Quick Count ( 500 /500 records )

June 2nd 2015, 13:32:15.000 🔍 🔍

9.2%

June 2nd 2015, 13:32:09.000 🔍 🔍

8.8%

June 2nd 2015, 13:32:13.000 🔍 🔍

8.4%

June 2nd 2015, 13:32:12.000 🔍 🔍

7.8%

June 2nd 2015, 13:32:10.000 🔍 🔍

7.6%

created\_at

Quick Count ( 500 /500 records )

02-Jun-2015 13:32:15.000 PM 🔍 🔍

9.2%

02-Jun-2015 13:32:09.000 PM 🔍 🔍

8.8%

02-Jun-2015 13:32:13.000 PM 🔍 🔍

8.4%

02-Jun-2015 13:32:12.000 PM 🔍 🔍

7.8%

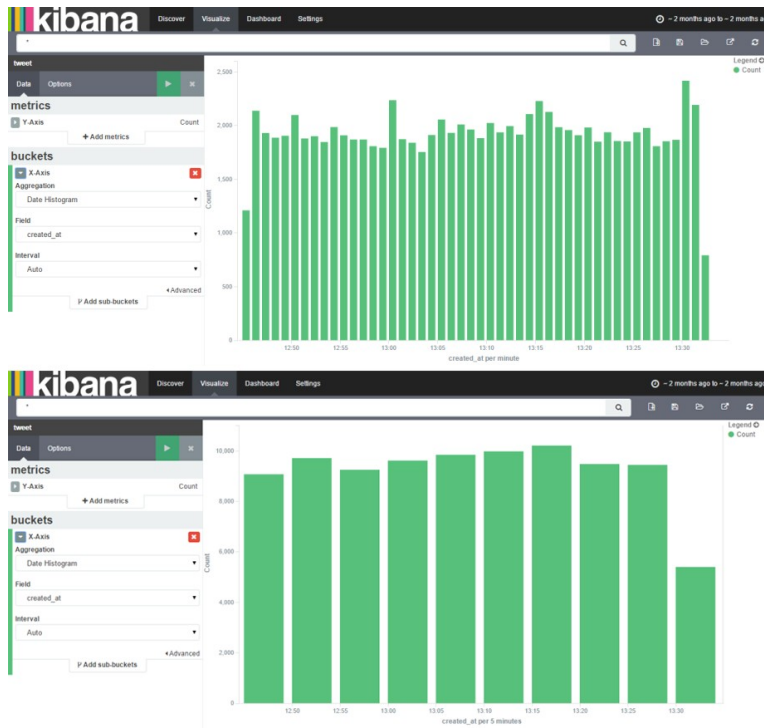
02-Jun-2015 13:32:10.000 PM 🔍 🔍

7.6%

# Advanced

These are the first 500 documents matching your search, refine your search to see others. [Back to top.](#)

These are the first 100 documents matching your search, refine your search to see others. [Back to top.](#)

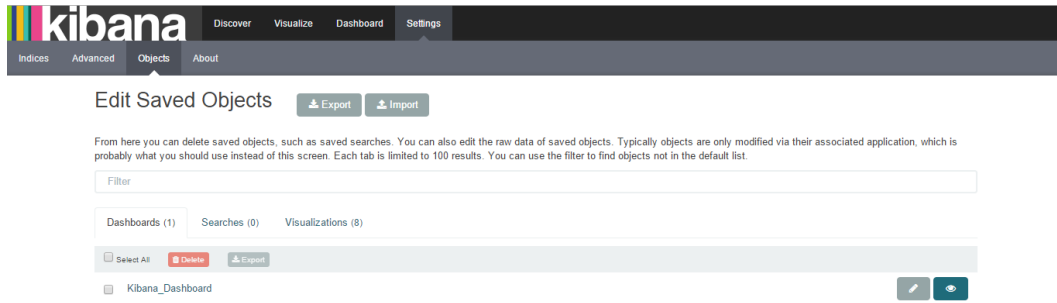


```
"Top 5 language","retweet.retweet_count ranges",Count
en,"0.0-10.0",4354
en,"10.0-20.0",754
ja,"0.0-10.0",1575
ja,"10.0-20.0",395
ar,"0.0-10.0",1160
ar,"10.0-20.0",230
in,"0.0-10.0",554
in,"10.0-20.0",125
es,"0.0-10.0",727
es,"10.0-20.0",118
```

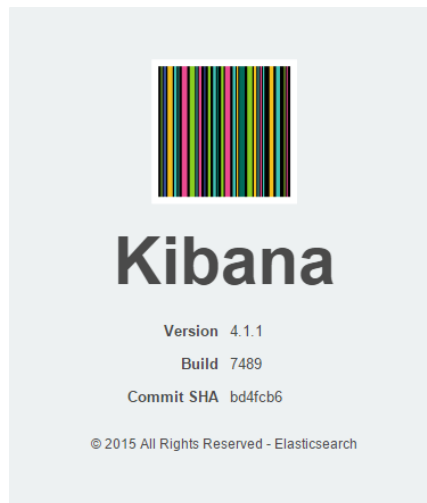
```
Top 5 language:retweet.retweet_count ranges:Count
en:0.0-10.0:4354
en:10.0-20.0:754
ja:0.0-10.0:1575
ja:10.0-20.0:395
ar:0.0-10.0:1160
ar:10.0-20.0:230
in:0.0-10.0:554
in:10.0-20.0:125
es:0.0-10.0:727
es:10.0-20.0:118
```

# Objects

## Managing saved Searches, Visualizations, and Dashboards



# About





# Creating a Twitter developer account

## Mobile

Expand your experience, get closer, and stay current.

### Add your phone number.

Enter your phone number in the box below. We'll send you a text message with a confirmation code. Text message fees may apply.

Country/region

Phone number

Continue

## Create an application

### Application Details

#### Name \*

Your application name. This is used to attribute the source of a tweet and in user-facing authorization screens. 32 characters max.

#### Description \*

Your application description, which will be shown in user-facing authorization screens. Between 10 and 200 characters max.

#### Website \*

Your application's publicly accessible home page, where users can go to download, make use of, or find out more information about your application. This fully-qualified URL is used in the source attribution for tweets created by your application and will be shown in user-facing authorization screens.  
(If you don't have a URL yet, just put a placeholder here but remember to change it later.)

#### Callback URL

Where should we return after successfully authenticating? [OAuth 1.0a](#) applications should explicitly specify their `oauth_callback` URL on the request token step, regardless of the value given here. To restrict your application from using callbacks, leave this field blank.

### Application Settings

Keep the "Consumer Secret" a secret. This key should never be human-readable in your application.

Consumer Key (API Key)	
Consumer Secret (API Secret)	
Access Level	Read and write ( <a href="#">modify app permissions</a> )
Owner	YuvrajGupta14
Owner ID	1031926392

#### Application Actions

[Regenerate Consumer Key and Secret](#) [Change App Permissions](#)

### Your Access Token

This access token can be used to make API requests on your own account's behalf. Do not share your access token secret with anyone.

Access Token	
Access Token Secret	
Access Level	Read and write
Owner	YuvrajGupta14
Owner ID	1031926392

#### Token Actions


[Regenerate My Access Token and Token Secret](#) [Revoke Token Access](#)

## Creating a Logstash configuration file


```
Logstash startup completed
```

## Creating visualizations for scenarios


### Number of tweets over a period of time



## Number of tweets in different languages




## Number of tweets from different geographical locations






## Number of tweets from Android, iPhone, iPad, and Web devices




## Number of tweets in various languages using different devices


:




## Number of tweets from various countries using different devices



## The most retweeted user screen name tweeting using different devices



## The most tweeted user's screen name



## Popular hashtags



# Twitter metrics

