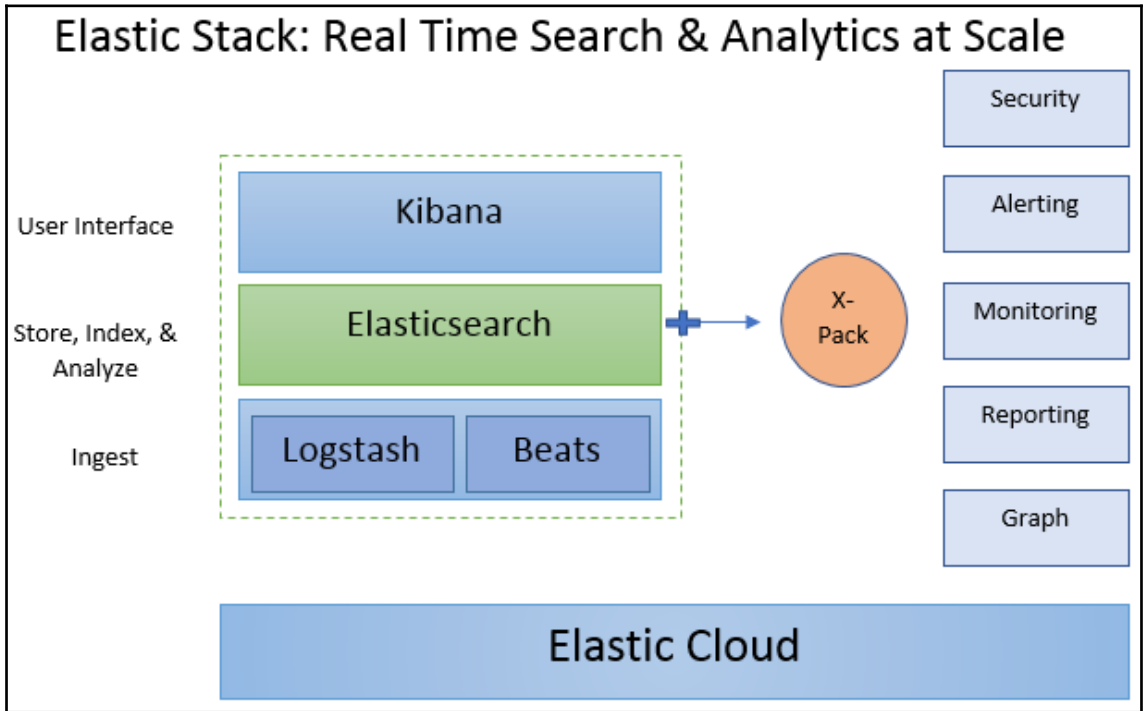
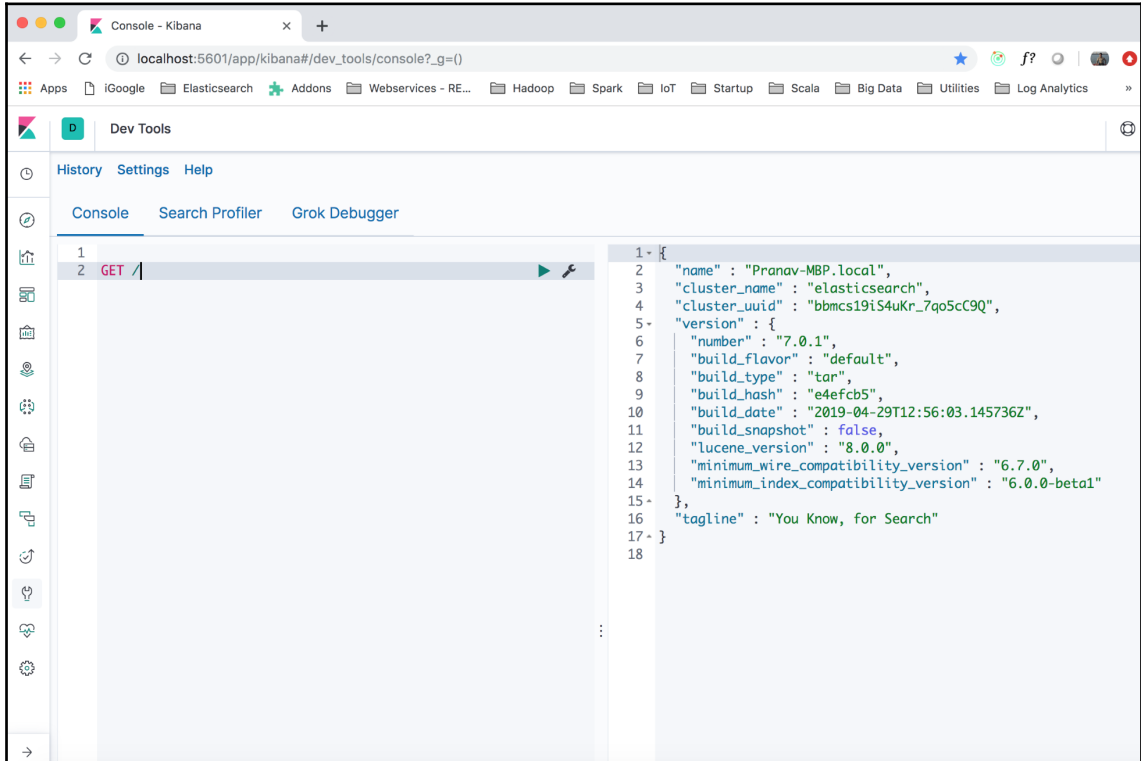


# Chapter 1: Introducing Elastic Stack



```
[ $ curl http://localhost:9200?pretty
{
  "name" : "Pranav-MBP.local",
  "cluster_name" : "elasticsearch",
  "cluster_uuid" : "bbmcs19iS4uKr_7qo5cC9Q",
  "version" : {
    "number" : "7.0.1",
    "build_flavor" : "default",
    "build_type" : "tar",
    "build_hash" : "e4efcb5",
    "build_date" : "2019-04-29T12:56:03.145736Z",
    "build_snapshot" : false,
    "lucene_version" : "8.0.0",
    "minimum_wire_compatibility_version" : "6.7.0",
    "minimum_index_compatibility_version" : "6.0.0-beta1"
  },
  "tagline" : "You Know, for Search"
}
$ █
```

# Chapter 2: Getting Started with Elasticsearch



Console - Kibana

localhost:5601/app/kibana#/dev\_tools/console?\_g=()

Apps | iGoogle | Elasticsearch | Addons | Webservices - RE... | Hadoop | Spark | IoT | Startup | Scala | Big Data | Utilities | Log Analytics

Dev Tools

History Settings Help

Console Search Profiler Grok Debugger

```
1
2 GET /
3
4
5 GET /my_index/_mapping
```

_mapping	endpoint
_mapping/field	endpoint
_mget	endpoint
_migration/deprecations	endpoint
_msearch	endpoint
_msearch/template	endpoint
_mtermvectors	endpoint
_tlm/explain	endpoint

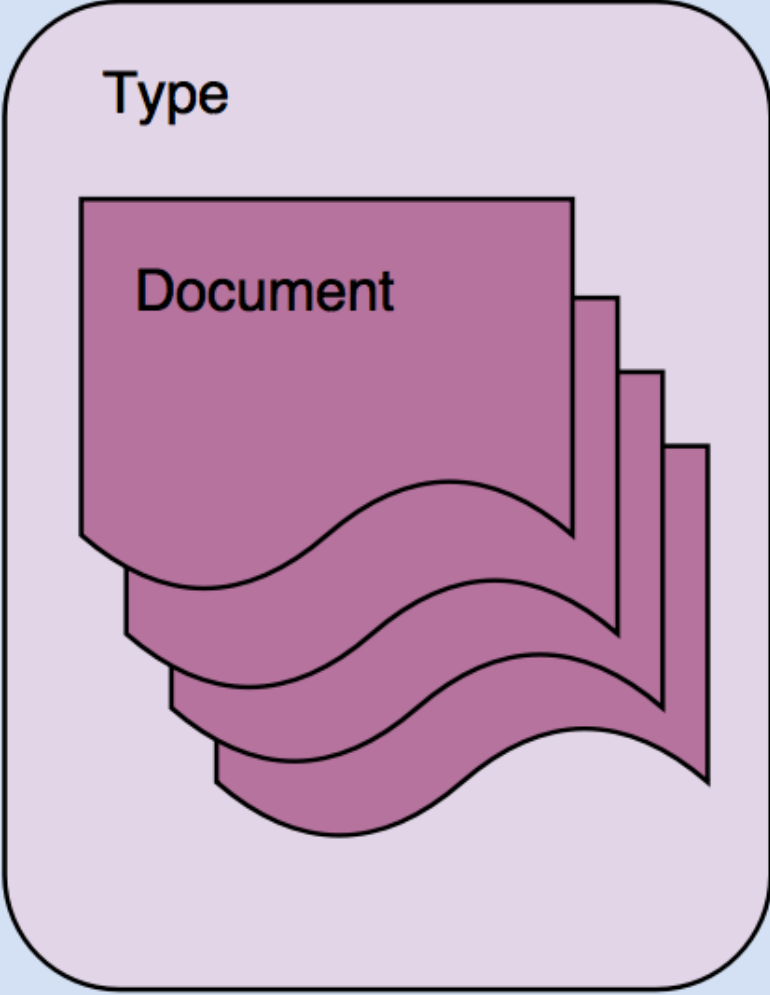
```
1- {
2  "name" : "Pranav-MBP.local",
3  "cluster_name" : "elasticsearch",
4  "cluster_uuid" : "bbmcs19iS4uKr_7qo5cC9Q",
5  "version" : {
6    "number" : "7.0.1",
7    "build_flavor" : "default",
8    "build_type" : "tar",
9    "build_hash" : "e4efcb5",
10   "build_date" : "2019-04-29T12:56:03.145736Z",
11   "build_snapshot" : false,
12   "lucene_version" : "8.0.0",
13   "minimum_wire_compatibility_version" : "6.7.0",
14   "minimum_index_compatibility_version" : "6.0.0-beta1"
15  },
16  "tagline" : "You Know, for Search"
17 }
18
```

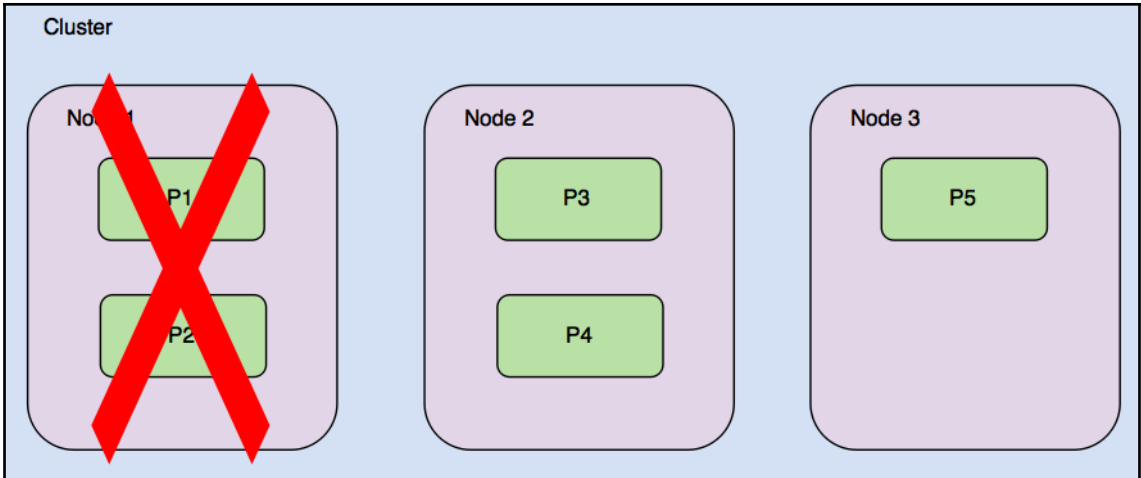

localhost:5601/app/kibana#/dev\_tools/searchprofiler

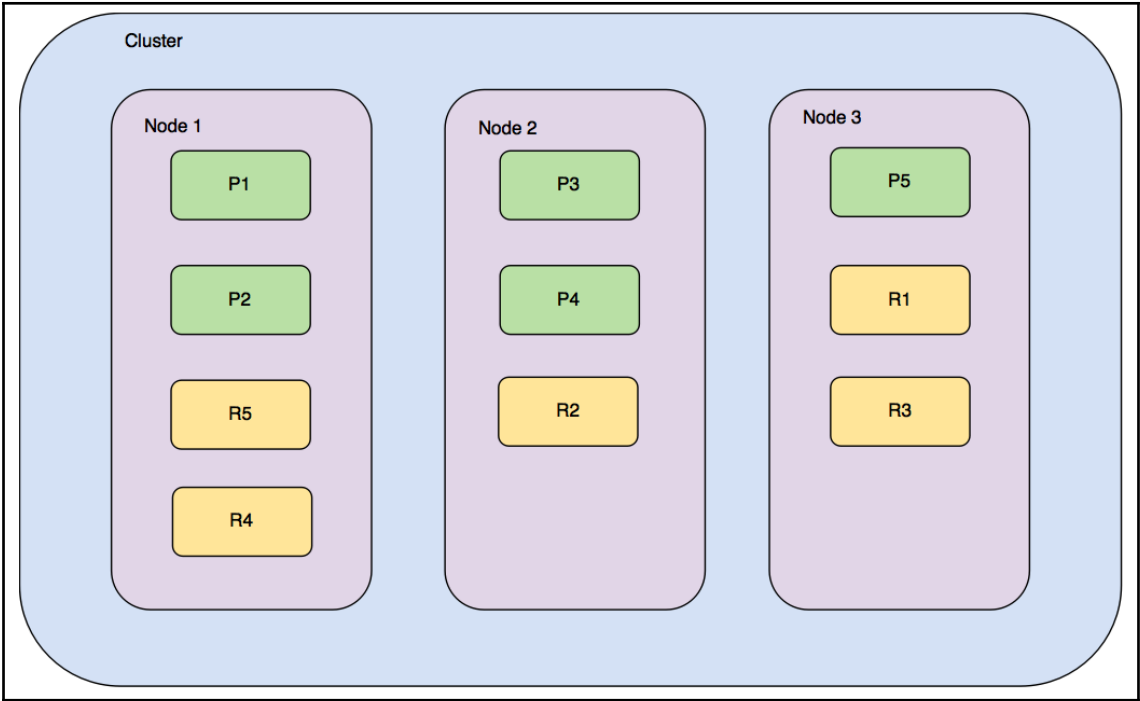
Index

Type

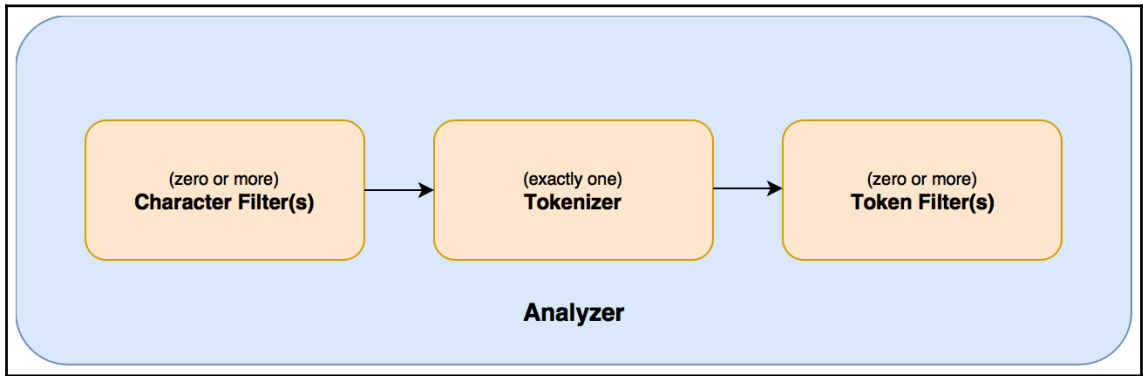
Document







## Chapter 3: Searching - What is Relevant





# Term level query flow

## At indexing time

```
{  
  "manufacturer": "victory multimedia",  
  ...  
}
```

Index time analyzer

Use the analyzer to break into terms  
here it's a noop analyzer

**Terms**  
victory multimedia

Store terms in the  
inverted index

## At query time

```
{  
  "query": {  
    "term": {  
      "manufacturer.raw": "victory multimedia"  
    }  
  }  
}
```

No analysis  
performed on the  
term

Execute the term level query  
directly against the inverted index

Term	Frequency	Documents
victory multimedia	1	1
...	...	...
...	...	...
...	...	...

# High level query flow

## At indexing time

```
{
  "title": "gods & heroes: rome rising",
  ...
}
```

Index time analyzer

Use the analyzer to break into terms

**Terms**  
gods  
heroes  
rome  
rising

Store terms in the inverted index

## At query time

```
{
  "query": {
    "match": {
      "title": "gods heroes"
    }
  }
}
```

Query time analyzer

Use the query/search analyzer to find terms

**Terms**  
gods  
heroes

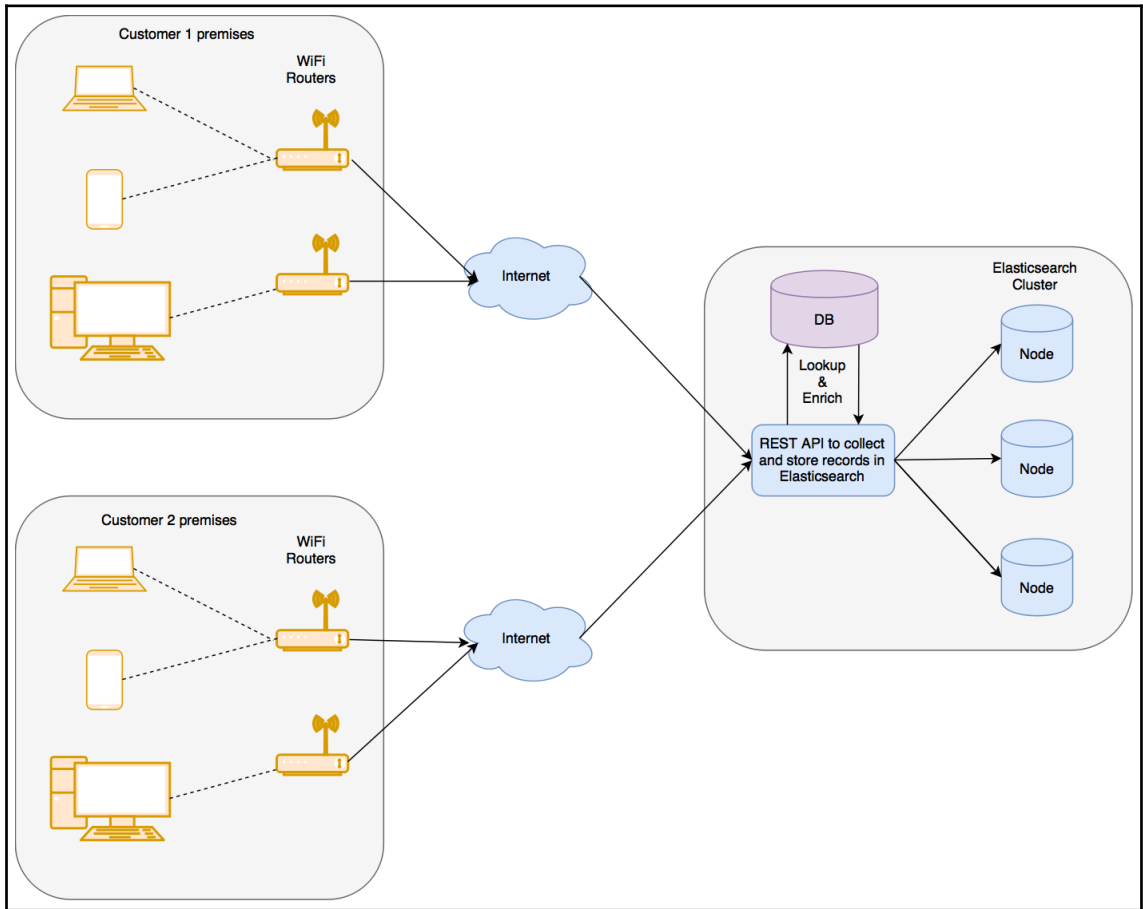
Rewrite in terms of term level queries using individual terms


```
{
  "query": {
    "bool": {
      "should": [
        { "term": { "title": "gods" } },
        { "term": { "title": "heroes" } }
      ]
    }
  }
}
```

Execute the term level query

Term	Frequency	Documents
gods	1	1
heroes	1	1
rome	1	1
rising	1	1

# Chapter 4: Analytics with Elasticsearch





# Chapter 5: Analyzing Log Data

```

2011-05-20 20:06:06.46 Server      This instance of SQL Server last reported using a process ID of 1760 at 5/20/2011 8:03:41 PM (
2011-05-20 20:06:06.46 Server      Registry startup parameters:
-d C:\Program Files\Microsoft SQL Server\MSSQL10_50.MSSQLSERVER\MSSQL\DATA\master.mdf          SQL Server Logs
-e C:\Program Files\Microsoft SQL Server\MSSQL10_50.MSSQLSERVER\MSSQL\Log\ERRORLOG
-l C:\Program Files\Microsoft SQL Server\MSSQL10_50.MSSQLSERVER\MSSQL\DATA\mastlog.ldf
2011-05-20 20:06:06.66 Server      SQL Server is starting at normal priority base (=7). This is an informational message only. No
  
```

```

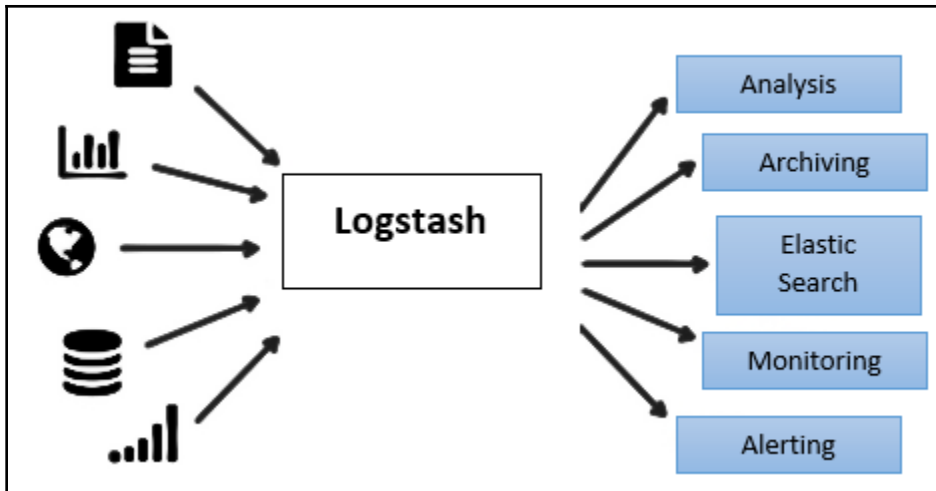
[2011-08-08 10:24:03.550]NBM | [index.store ] [node] [dn_index][3] failed to build store metadata. checking segment info
integrity (with commit: no)
java.nio.file.NoSuchFileException: /u01/pft/pt/es2.3.2/data/ESCluster/nodes/0/index/nr_no_job_data_hot2tp/3/index/segment_1q
at sun.nio.fs.UnixException.rethrowAsIOException(UnixException.java:146)
at sun.nio.fs.UnixException.rethrowAsIOException(UnixException.java:102)
at sun.nio.fs.UnixException.rethrowAsIOException(UnixException.java:127)
at sun.nio.fs.UnixFileSystemProvider.newFileChannel(UnixFileSystemProvider.java:177)
at java.nio.channels.FileChannel.open(FileChannel.java:237)
at java.nio.channels.FileChannel.open(FileChannel.java:235)
at org.apache.lucene.store.NIOFSDirectory.openInput(NIOFSDirectory.java:184)
at org.apache.lucene.store.Filedirectory.openInput(Filedirectory.java:89)
at org.elasticsearch.index.store.StoreMetadataSnapshot.checksumFromLuceneFile(Store.java:930)
at org.elasticsearch.index.store.StoreMetadataSnapshot.loadMetadata(Store.java:940)
at org.elasticsearch.index.store.StoreMetadataSnapshot.close(Store.java:74)
at org.elasticsearch.index.store.Store.getMetadata(Store.java:233)
at org.elasticsearch.index.store.Store.getMetadataEntry(Store.java:192)
at org.elasticsearch.indices.store.TransportNodeListShardStoreMetadata.ListStoreMetadata(TransportNodeListShardStoreMetadata.java:161)
at org.elasticsearch.indices.store.TransportNodeListShardStoreMetadata.nodeOperation(TransportNodeListShardStoreMetadata.java:142)
at org.elasticsearch.indices.store.TransportNodeListShardStoreMetadata.nodeOperation(TransportNodeListShardStoreMetadata.java:67)
at org.elasticsearch.action.support.nodes.TransportNodeAction.nodeOperation(TransportNodeAction.java:92)
at org.elasticsearch.action.support.nodes.TransportNodeActionNodeTransportHandler.messageReceived(TransportNodeAction.java:230)
at org.elasticsearch.action.support.nodes.TransportNodeActionNodeTransportHandler.messageReceived(TransportNodeAction.java:224)
at org.elasticsearch.transport.RequestHandlerRegistry.processMessageReceived(RequestHandlerRegistry.java:75)
at org.elasticsearch.transport.netty.MessageChannelHandler$RequestHandler.doRun(MessageChannelHandler.java:300)
at org.elasticsearch.common.util.concurrent.AbstractRunnable.run(AbstractRunnable.java:37)
  
```

Elasticsearch  
Exceptions

```

93.180.71.3 - - [17/May/2015:08:05:23 +0000] "GET /downloads/product_1 HTTP/1.1" 304 0 "-" "Debian
APT-HTTP/1.3 (0.8.16-exp12ubuntu10.21)"
80.91.33.133 - - [17/May/2015:08:05:24 +0000] "GET /downloads/product_1 HTTP/1.1" 304 0 "-" "Debian
APT-HTTP/1.3 (0.8.16-exp12ubuntu10.17)"
  
```

NGINX logs



← → ↻ <https://www.elastic.co/downloads/logstash-oss>

Downloads

## Download Logstash - OSS Only

Want to upgrade? We'll give you a hand. [Migration Guide](#) »

Version: 7.0.0


Release date: April 10, 2019

License: [Apache 2.0](#)

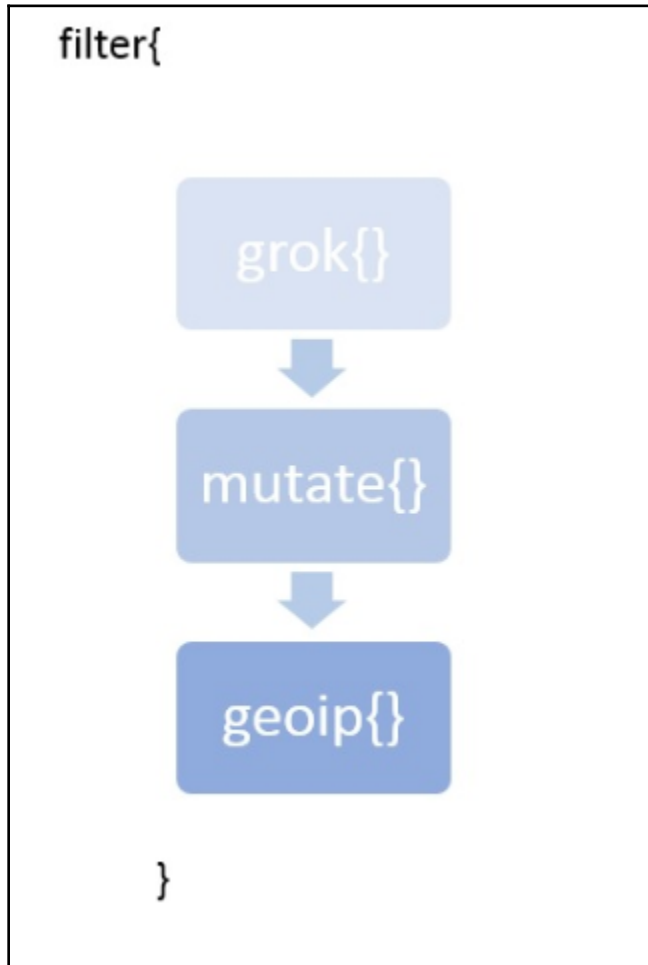
Downloads: [TAR.GZ](#) sha [ZIP](#) sha  
[DEB](#) sha [RPM](#) sha

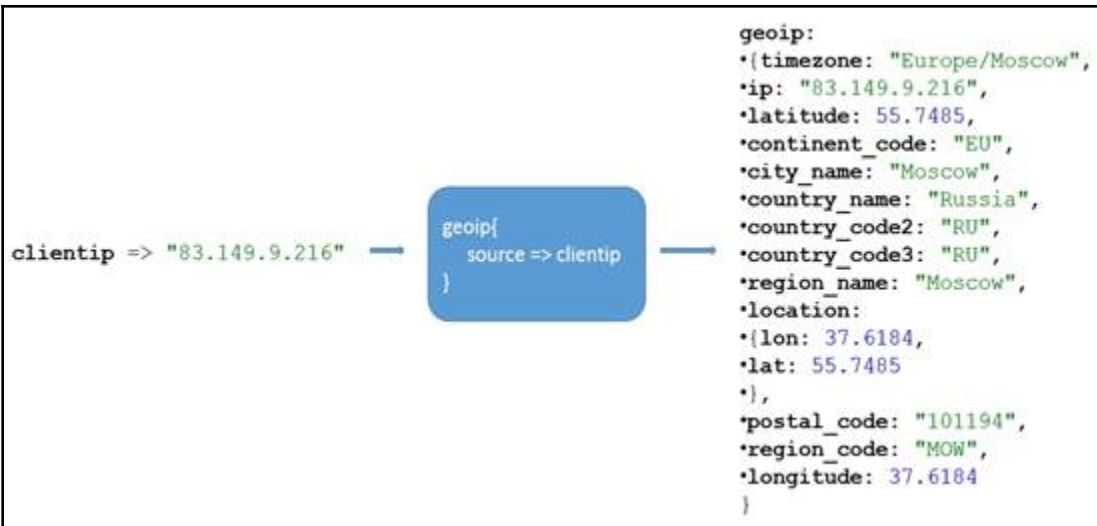
**Notes:** This distribution only includes features licensed under the Apache 2.0 license. To get access to full [set of free features](#), use the [default distribution](#).

View the detailed release notes [here](#).  
Not the version you're looking for? View [past releases](#).  
Java 8 is required for Logstash 6.x and 5.x.

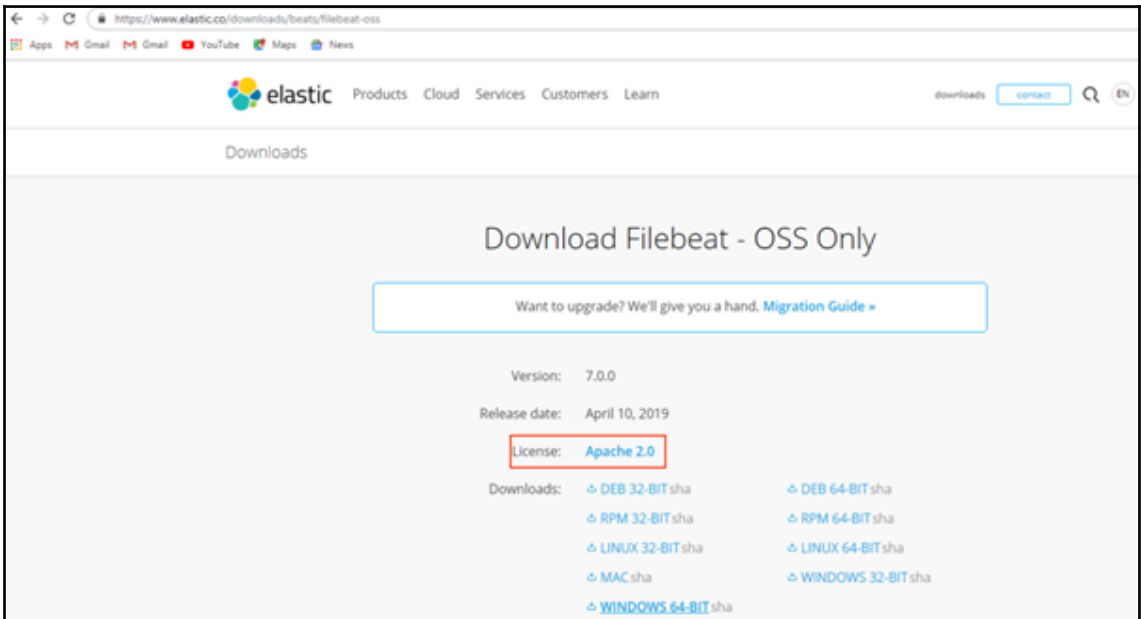
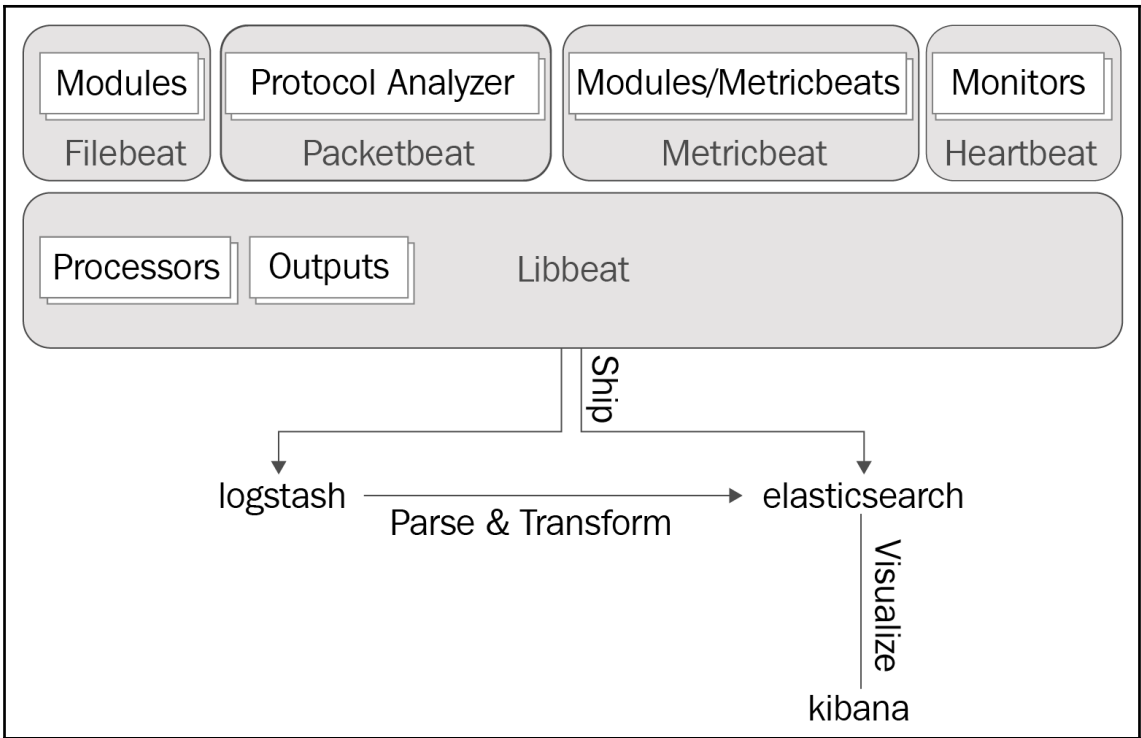


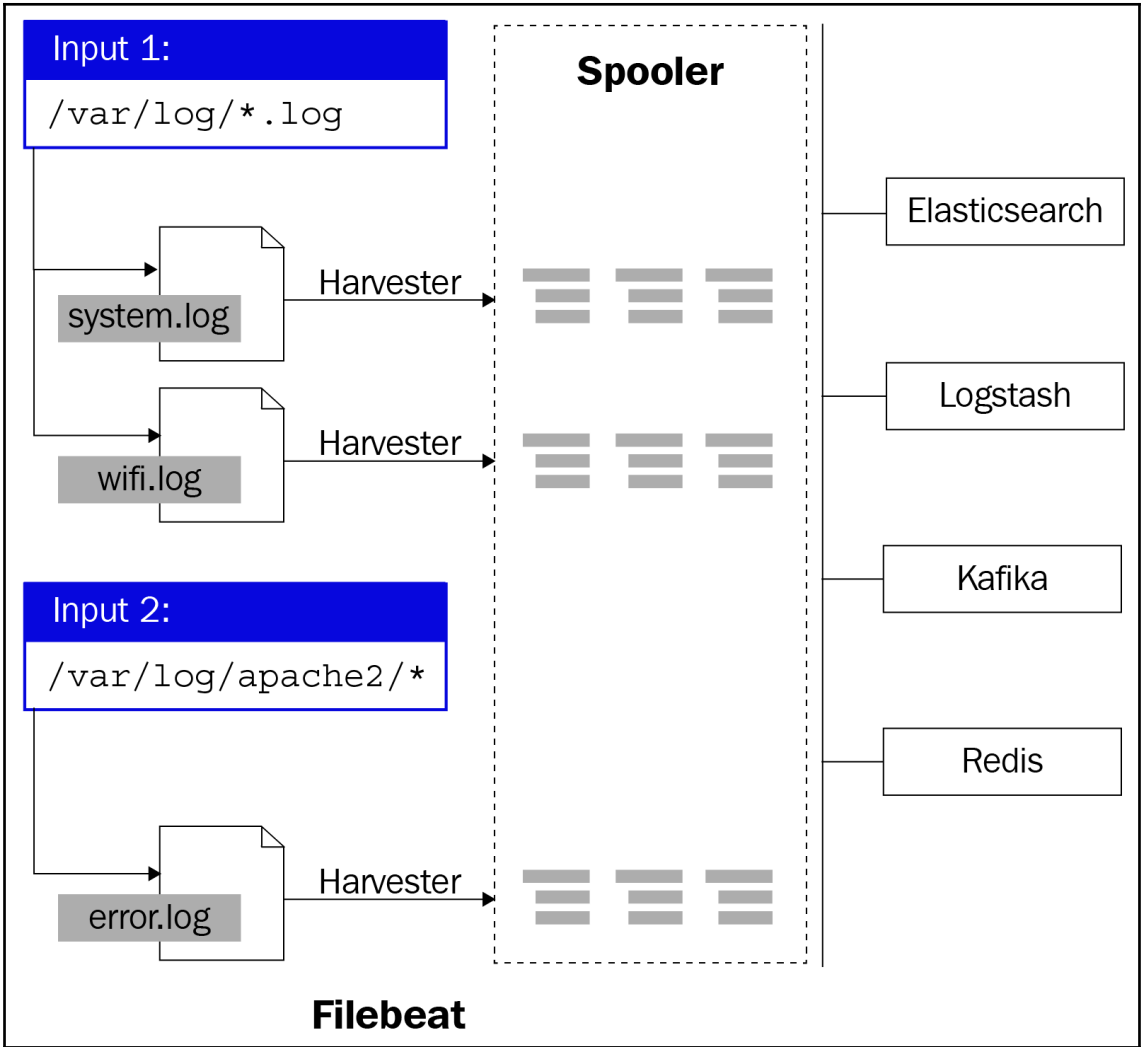
## Chapter 6: Building Data Pipelines with Logstash











```
#===== Filebeat inputs =====  
  
filebeat.inputs:  
  
- type: log  
  # Enable or disable  
  enabled: true  
  paths:  
    - E:\packt\logs\*.log  
    - C:\programdata\elasticsearch\logs\  
  
  exclude_lines: ['^DBG']  
  include_lines: ['^ERR', '^WARN']  
  exclude_files: ['.gz$']  
  fields:  
    level: error_warn_logs  
  tags: ['eslogs']  
  
  multiline.pattern: '^[[:space]]'  
  multiline.negate: false  
  multiline.match: after  
  
- type: docker  
  enabled: true  
  containers.path: "/var/lib/docker/containers"  
  containers.stream: "all"  
  containers.ids: "*"   
  tags: ['dockerlogs']
```

# Chapter 7: Visualizing Data with Kibana

The screenshot shows the Elastic website's download page for Kibana - OSS Only. The page includes a navigation bar with the Elastic logo and links for Products, Cloud, Services, Customers, Learn, downloads, and a contact button. The main heading is "Download Kibana - OSS Only". A callout box asks if the user wants to upgrade and provides a link to the Migration Guide. The page lists the version (7.0.0), release date (April 10, 2019), and license (Apache 2.0, which is highlighted with a red box). It also provides download links for Windows, Linux, and Debian, as well as container instructions for Docker. A notes section explains that this distribution is limited to Apache 2.0 licensed features, and a link is provided to view the detailed release notes.

https://www.elastic.co/downloads/kibana-oss

elastic Products Cloud Services Customers Learn downloads [contact](#) Q EN

## Downloads

### Download Kibana - OSS Only

Want to upgrade? We'll give you a hand. [Migration Guide »](#)

Version: 7.0.0

Release date: April 10, 2019

License: **Apache 2.0**

Downloads: [WINDOWS](#) sha [MAC](#) sha  
[LINUX 64-BIT](#) sha [RPM 64-BIT](#) sha  
[DEB 64-BIT](#) sha

Containers: Run with [Docker](#)

Notes: This distribution only includes features licensed under the Apache 2.0 license. To get access to full [set of free features](#), use the [default distribution](#).

View the detailed release notes [here](#).

localhost:5601/status#?\_g=0

### Server Status

MADSH01-APM01

**Kibana status is Green**

1.42 GB Heap total	71.10 MB Heap used	0.00, 0.00, 0.00 Load
241.33 ms Response time avg	1972.00 ms Response time max	1.80 Requests per second

#### Plugin status

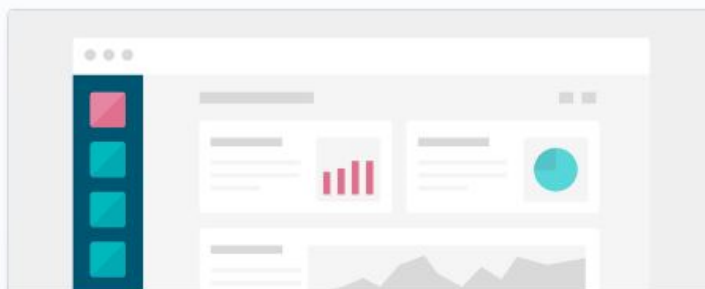
BUILD 23117 COMMIT ee89fda8

ID	Status
plugin:apm_oss@undefined	Ready
plugin:kibana@undefined	Ready
plugin:elasticsearch@undefined	Ready
plugin:interpreter@undefined	Ready
plugin:metrics@undefined	Ready
plugin:console@undefined	Ready
plugin:timelion@undefined	Ready
plugin:tile_map@undefined	Ready



# Welcome to Kibana

Your window into the Elastic Stack

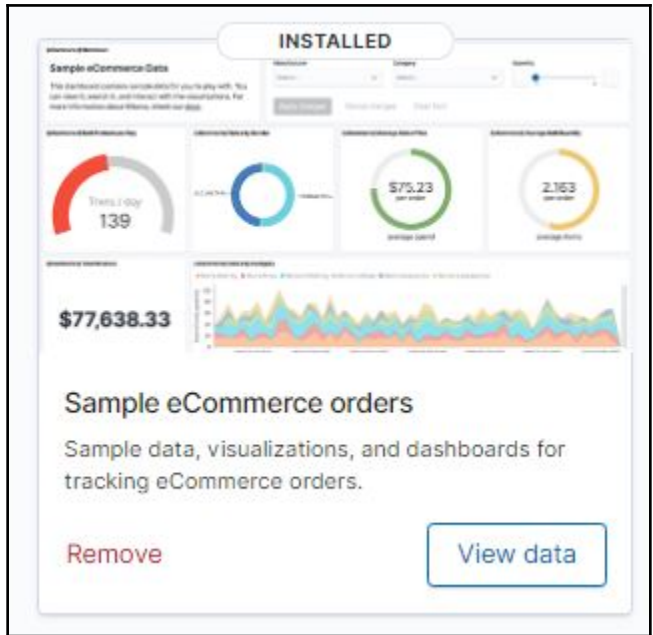
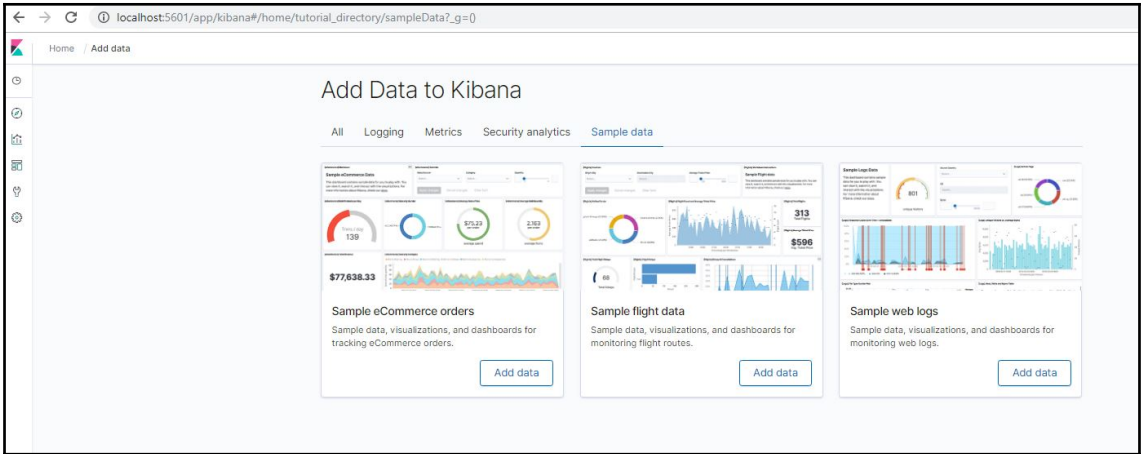


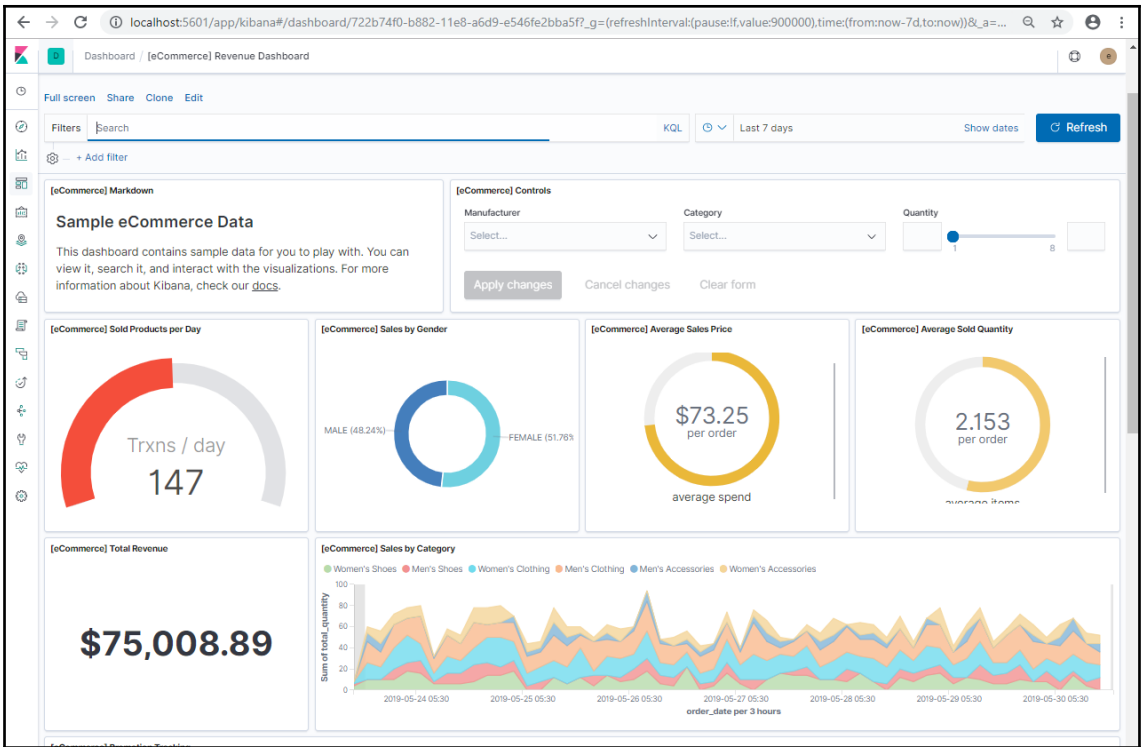
## Let's get started

We noticed that you don't have any data in your cluster. You can try our sample data and dashboards or jump in with your own data.

[Try our sample data](#)

[Explore on my own](#)








Home


## Add Data to Kibana

Use these solutions to quickly turn your data into pre-built dashboards and monitoring systems.




**APM**  
APM automatically collects in-depth performance metrics and errors from inside your applications.

Add APM




**Logging**  
Ingest logs from popular data sources and easily visualize in preconfigured dashboards.

Add log data



**Metrics**  
Collect metrics from the operating system and services running on your servers.

Add metric data



**Security analytics**  
Centralize security events for interactive investigation in ready-to-go visualizations.

Add security events

**Add sample data** **1**  
Load a data set and a Kibana dashboard

**Use Elasticsearch data** **2**  
Connect to your Elasticsearch index

**3**

### Visualize and Explore Data

**Dashboard**  
Display and share a collection of visualizations and saved searches.

**Discover**  
Interactively explore your data by querying and filtering raw documents.

**Visualize**  
Create visualizations and aggregate data stores in your Elasticsearch indices.

### Manage and Administer the Elastic Stack

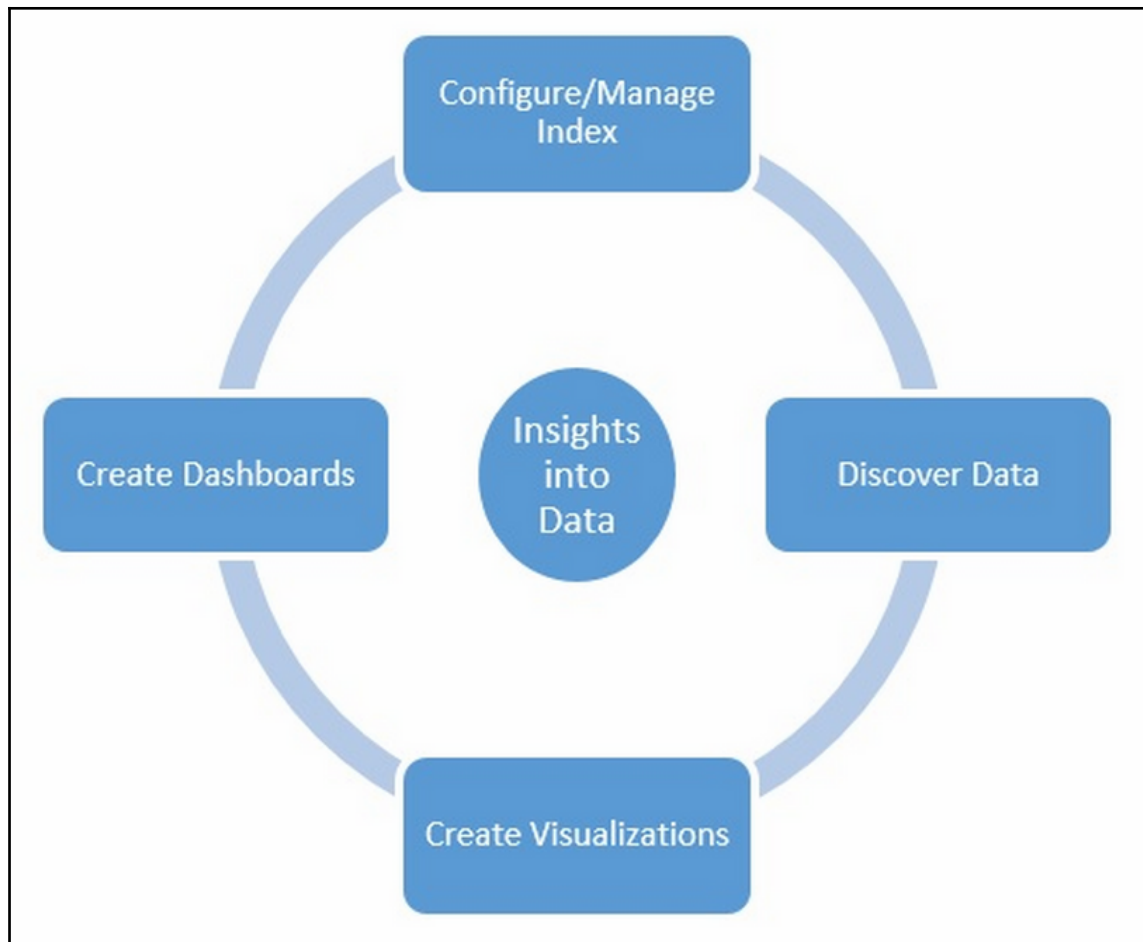
**Console**  
Skip cURL and use this JSON interface to work with your data directly.

**Index Patterns**  
Manage the index patterns that help retrieve your data from Elasticsearch.

**Saved Objects**  
Import, export, and manage your saved searches, visualizations, and dashboards.

Didn't find what you were looking for?

[View full directory of Kibana plugins](#)



### Create index pattern

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

## Create index pattern

Kibana uses index patterns to retrieve data from Elasticsearch indices for things like visualizations.  Include system indices

### Step 1 of 2: Define index pattern

**Index pattern**

logstash-\*

You can use a \* as a wildcard in your index pattern.  
You can't use spaces or the characters \, /, ?, " <, >, |.

[Next step](#)

✓ **Success!** Your index pattern matches **31 indices**.

logstash-2014.05.28
logstash-2014.05.29
logstash-2014.05.30
logstash-2014.05.31
logstash-2014.06.01
logstash-2014.06.02
logstash-2014.06.03
logstash-2014.06.04
logstash-2014.06.05
logstash-2014.06.06

Rows per page: 10

< 1 2 3 4 >

### Create index pattern

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

## Create index pattern

Kibana uses index patterns to retrieve data from Elasticsearch indices for things like visualizations.  Include system indices

### Step 2 of 2: Configure settings

You've defined **logstash-\*** as your index pattern. Now you can specify some settings before we create it.

**Time Filter field name** Refresh

@timestamp

The Time Filter will use this field to filter your data by time.  
You can choose not to have a time field, but you will not be able to narrow down your data by a time range.

[Show advanced options](#)

[Back](#) [Create index pattern](#)

Create index pattern

★ logstash-\*

★ logstash-\*

★ logstash-\*

Time Filter field name: @timestamp

This page lists every field in the **logstash-\*** index and the field's associated core type as recorded by Elasticsearch. To change a field type, use the [Elasticsearch Mapping API](#)

Fields (78) | Scripted fields (0) | Source filters (0)

Filter  All field types ▾

Name	Type	Format	Searchable	Aggregatable	Excluded
@timestamp	date		●	●	
@version	string		●	●	
._id	string		●	●	
._index	string		●	●	
._score	number				
._source	._source				
._type	string		●	●	
agent	string		●		
agent.keyword	string		●	●	
auth	string		●		

Rows per page: 10 ▾

< 1 2 3 4 5 ... 8 >

0 hits

New Save Open Share Inspect

Filters us Lucene  May 28, 2014 @ 00:00:00.0 → Jul 1, 2014 @ 00:00:00.0 Update

+ Add filter

logstash-\*

Selected fields

? \_source

Available fields

No results match your search criteria

Expand your time range

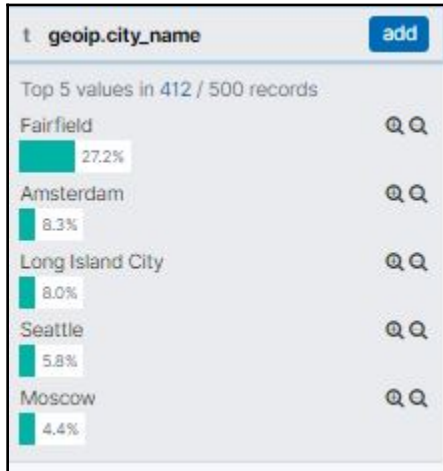
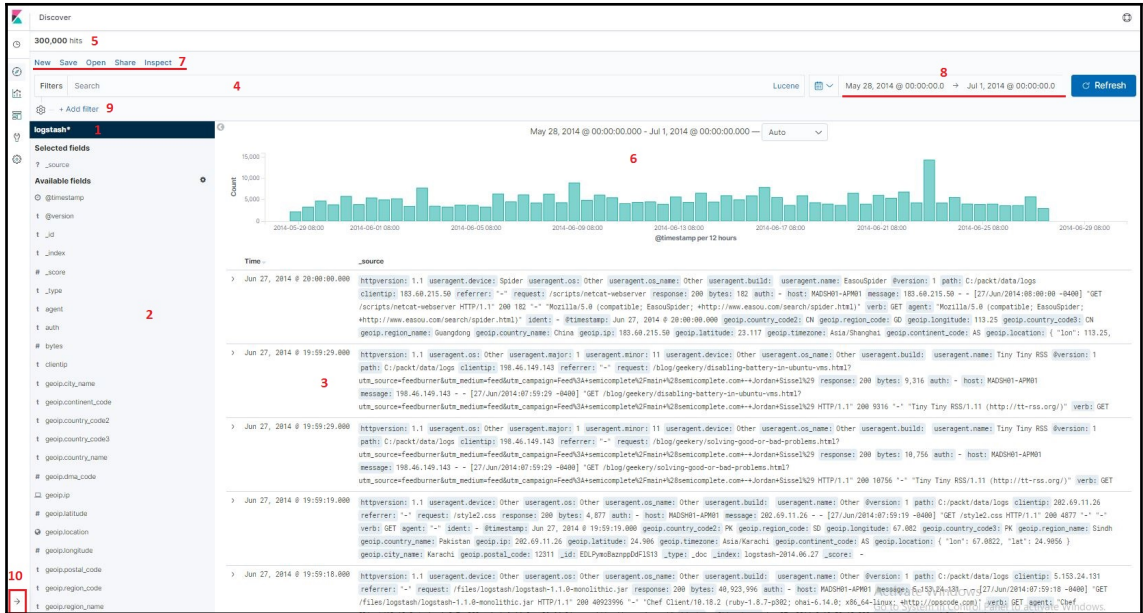
One or more of the indices you're looking at contains a date field. Your query may not find any results because there may not be any data at all in the currently selected time range. You can try changing the time range.

Refine your query

The search bar at the top uses Elasticsearch's support for Lucene Query String syntax.

Calendar: July 2014

2014-07-01 00:00:00.000



Time ▾      \_source

Expand Button

```

June 27th 2014, 17:43:20.000 request: /scripts/netcat-webserver agent: "Mozilla/5.0 (compatible; EasouSpider; +http://www.easou.com/search/spider.html)" geoip.ci
geoip.timezone: Asia/Shanghai geoip.ip: 183.60.215.50 geoip.latitude: 23.117 geoip.country_name: China geoip.country_code2: CN geoip.c
geoip.country_code3: CN geoip.region_name: Guangdong geoip.location: { "lon": 113.25, "lat": 23.1167 } geoip.region_code: 44 geoip.long
ident: - verb: GET useragent.os: Other useragent.build: useragent.name: EasouSpider useragent.os_name: Other useragent.device: Spide
- - [27/Jun/2014:08:00:00 -0400] "GET /scripts/netcat-webserver HTTP/1.1" 200 182 "-" "Mozilla/5.0 (compatible; EasouSpider; +http://

```

Table    JSON    [View surrounding data](#)

@timestamp	June 27th 2014, 17:43:20.000
@version	1
_id	AV4jHLxYxVeTbjX4rA1W
_index	logstash-2014.06.27
_score	-
_type	logs
agent	"Mozilla/5.0 (compatible; EasouSpider; +http://www.easou.com/search/spider.html)"
auth	-
# bytes	182
clientip	183.60.215.50
geoip.city_name	Guangzhou
geoip.continent_code	AS
geoip.country_code2	CN

t geoip.country\_name    🔍 🔍 📄 \*    China

📄 geoip.ip

**Toggle column in table**

Time ▾	geoip.city_name	response	request ✕ ⏪
▶ June 27th 2014, 17:43:20.000	Guangzhou	200	server
▶ June 27th 2014, 17:42:49.000	Buffalo	200	/blog/geekery/solving-good-or-bad-problems.html?utm_source=feedburner&utm_medium=feed&utm_campaign=Fe
▶ June 27th 2014, 17:42:49.000	Buffalo	200	/blog/geekery/disabling-battery-in-ubuntu-vms.html?utm_source=feedburner&utm_medium=feed&utm_campaign=Fe
▶ June 27th 2014, 17:42:39.000	-	200	/style2.css
▶ June 27th 2014, 17:42:38.000	Amsterdam	200	/files/logstash/logstash-1.1.0-monolithic.jar
▶ June 27th 2014, 17:42:37.000	-	200	/images/jordan-80.png
▶ June 27th 2014, 17:42:35.000	-	200	/reset.css
▶ June 27th 2014, 17:42:30.000	-	200	/blog/tags/X11
▶ June 27th 2014, 17:42:12.000	-	200	/images/googledotcom.png

Move column to the left

## SYNTAX OPTIONS

The **Kibana Query Language (KQL)** offers a simplified query syntax and support for scripted fields. KQL also provides autocomplete if you have a Basic license or above. If you turn off KQL, Kibana uses Lucene.

**Kibana Query Language**

Off

88,602 hits

New Save Open Share Inspect

Filters **files logstash** Lucene May 28, 2014 @ 00:00:00.0 → Jul 1, 2014 @ 00:00:00.0 Refresh

logstash-\*

Selected fields

- \_source

Available fields

- @timestamp
- @version
- \_id
- \_index
- \_score
- \_type
- \_type
- agent
- auth
- bytes
- clientip
- clientip\_name
- geosp.country\_code
- geosp.country\_code2
- geosp.country\_name
- geosp.dma\_code
- geospip
- geosp.latitude
- geosp.location
- geosp.longitude

Time-- source

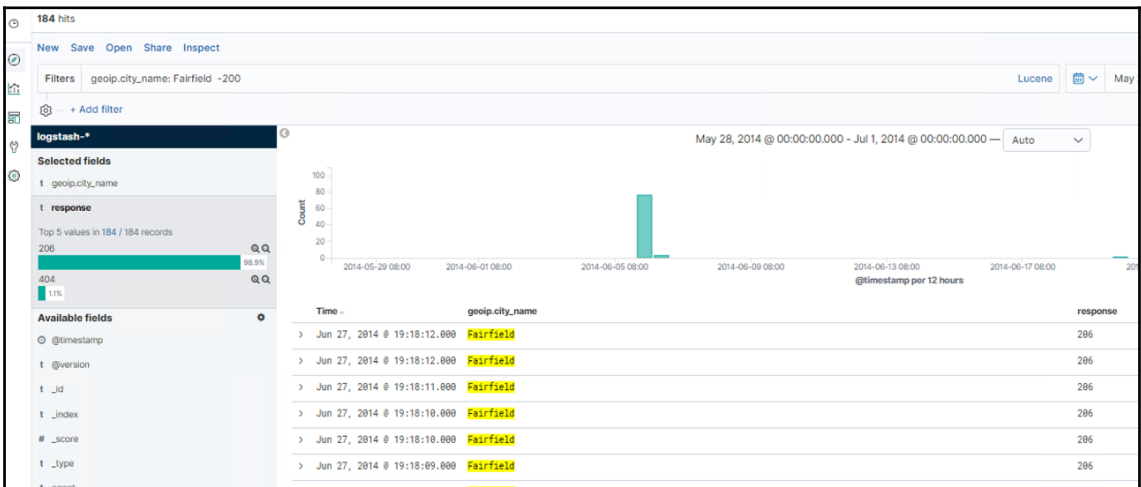
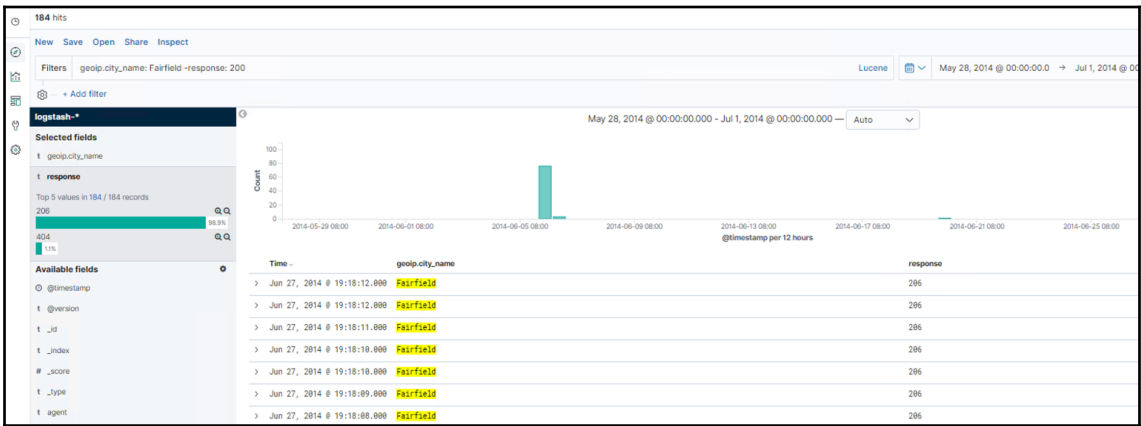
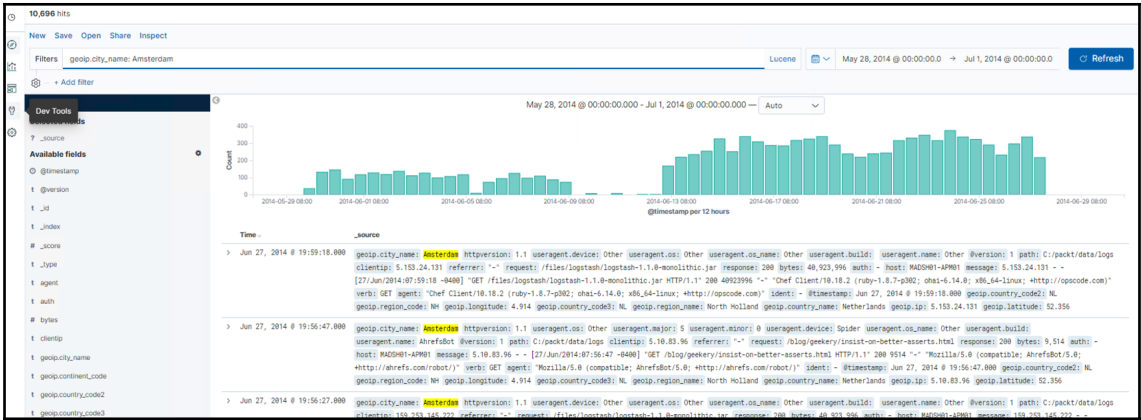
```

> Jun 27, 2014 @ 19:59:18.000 request: /files logstash logstash-1.1.0-monolithic.jar message: 5.153.24.131 - - [27/Jun/2014:07:59:18 -0400] "GET /files logstash logstash-1.1.0-monolithic.jar HTTP/1.1" 200 40923996 "-" "Chef Client/19.18.2 (ruby-1.8.7-p382; ohai-6.14.0; x86_64-linux; http://opscode.com)" httpversion: 1.1 useragent.device: Other useragent.os: Other useragent.os_name: Other useragent.build: useragent.name: Other httpversion: 1 path: C:\pack\data\logs @timestamp: 5.153.24.131 referer: "-" response: 200 bytes: 40.923.996 auth: - host: MAD091-AP01 verb: GET agent: "Chef Client/19.18.2 (ruby-1.8.7-p382; ohai-6.14.0; x86_64-linux; http://opscode.com)" ident: - @timestamp: Jun 27, 2014 @ 19:59:18.000 geosp.country_code2: NL geosp.region_code: NH geospip: 5.153.24.131 geosp.latitude: 52.356 geosp.timezone: Europe/Amsterdam

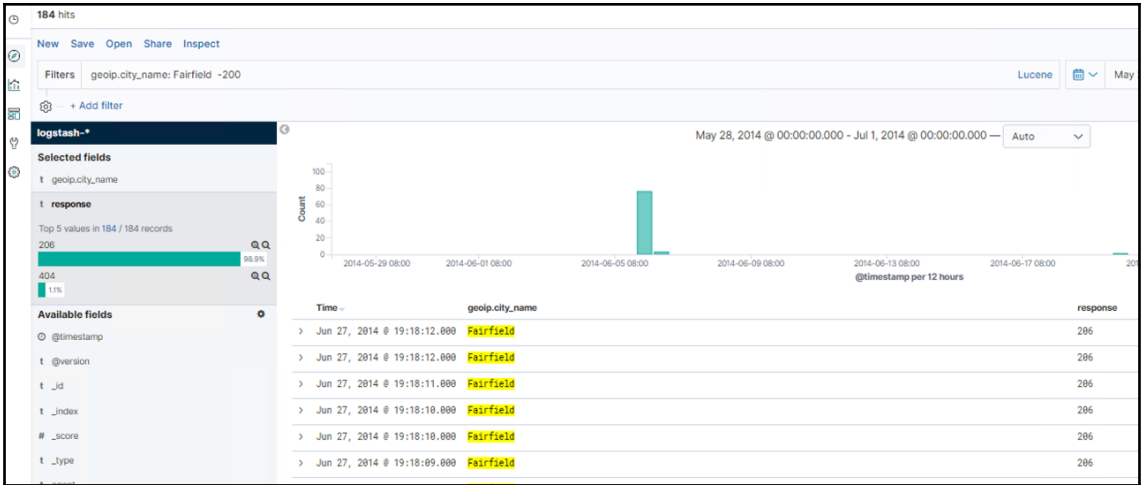
> Jun 27, 2014 @ 19:58:16.000 request: /files logstash logstash-1.1.0-monolithic.jar message: 173.192.221.212 - - [27/Jun/2014:07:58:16 -0400] "GET /files logstash logstash-1.1.0-monolithic.jar HTTP/1.1" 200 40923996 "-" "Chef Client/19.18.2 (ruby-1.8.7-p382; ohai-6.14.0; x86_64-linux; http://opscode.com)" httpversion: 1.1 useragent.device: Other useragent.os: Other useragent.os_name: Other useragent.build: useragent.name: Other httpversion: 1 path: C:\pack\data\logs @timestamp: 173.192.221.212 referer: "-" response: 200 bytes: 40.923.996 auth: - host: MAD091-AP01 verb: GET agent: "Chef Client/19.18.2 (ruby-1.8.7-p382; ohai-6.14.0; x86_64-linux; http://opscode.com)" ident: - @timestamp: Jun 27, 2014 @ 19:58:16.000 geosp.country_code2: US geosp.region_code: VA geospip: 173.192.221.212 geosp.latitude: 38.887 geosp.timezone: America/New_York

> Jun 27, 2014 @ 19:56:27.000 request: /files logstash logstash-1.1.0-monolithic.jar message: 159.253.145.222 - - [27/Jun/2014:07:56:27 -0400] "GET /files logstash logstash-1.1.0-monolithic.jar HTTP/1.1" 200 40923996 "-" "Chef Client/19.18.2 (ruby-1.8.7-p382; ohai-6.14.0; x86_64-linux; http://opscode.com)" httpversion: 1.1 useragent.device: Other useragent.os: Other useragent.os_name: Other useragent.build: useragent.name: Other httpversion: 1 path: C:\pack\data\logs @timestamp: 159.253.145.222 referer: "-" response: 200 bytes: 40.923.996 auth: - host: MAD091-AP01 verb: GET agent: "Chef Client/19.18.2 (ruby-1.8.7-p382; ohai-6.14.0; x86_64-linux; http://opscode.com)" ident: - @timestamp: Jun 27, 2014 @ 19:56:27.000 geosp.country_code2: NL geosp.region_code: NH geospip: 159.253.145.222 geosp.latitude: 52.356

> Jun 27, 2014 @ 19:54:16.000 request: /robots.txt/robots.txt?403-D message: 5.10.83.95 - - [27/Jun/2014:07:54:16 -0400] "GET /robots.txt/robots.txt?403-D HTTP/1.1" 200 955 "-" Mozilla/5.0 (compatible; Ahrefbot/5.0; http://ahrefs.com/robot/) httpversion: 1.1 useragent.device: spider useragent.os_name: Other useragent.build: useragent.name: Ahrefbot httpversion: 1 path: C:\pack\data\logs @timestamp: 5.10.83.95 referer: "-" response: 200 bytes: 955 auth: - host: MAD091-AP01 verb: GET agent: "Mozilla/5.0 (compatible; Ahrefbot/5.0; http://ahrefs.com/robot/)" ident: - @timestamp: Jun 27, 2014 @ 19:54:16.000 geosp.country_code2: NL geosp.region_code: NH geospip: 5.10.83.95 geosp.latitude: 52.356 geosp.timezone: Europe/Amsterdam geosp.continent_code: EU
    
```







Filters response:[301 TO 500]

+ Add filter

logstash-\*

Selected fields

- t geoip.city\_name
- t response

Available fields

Popular

- t geoip.country\_code2

@timestamp

- t @version

- t \_id

- t \_index

- # \_score


- t \_type

- t agent

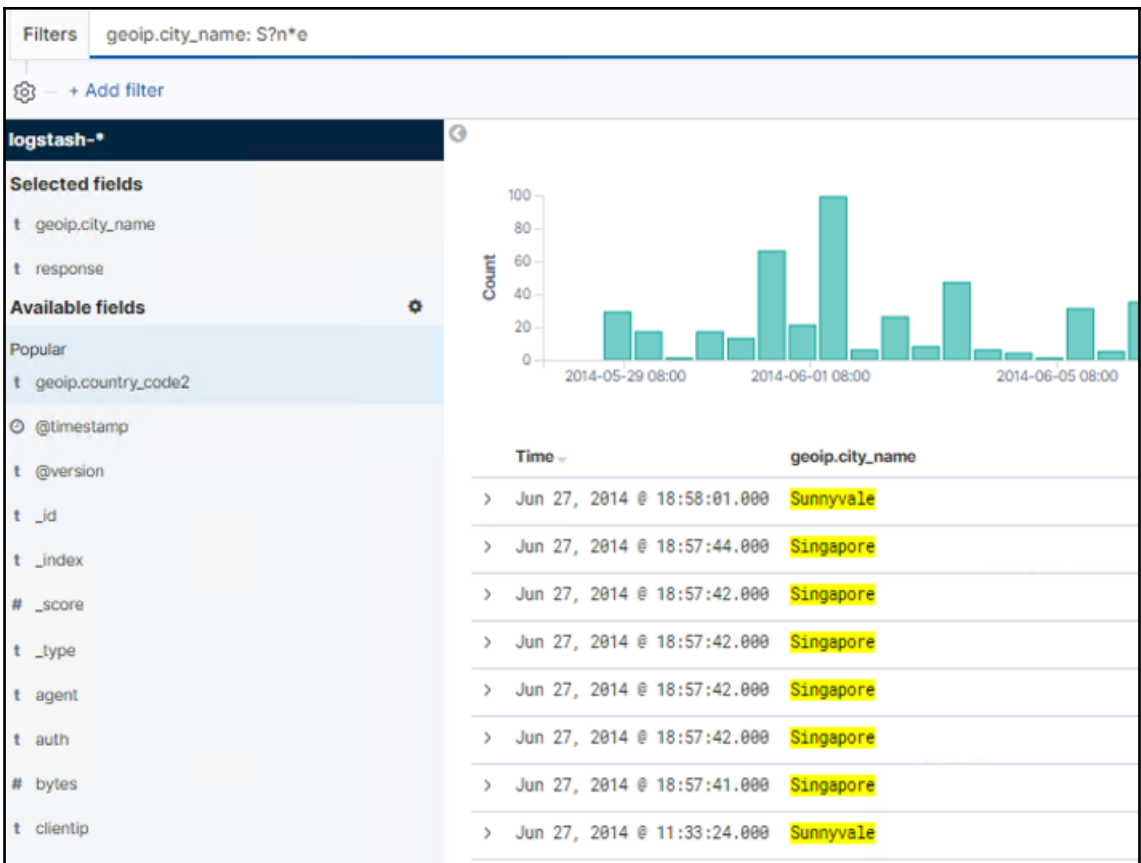
- t auth

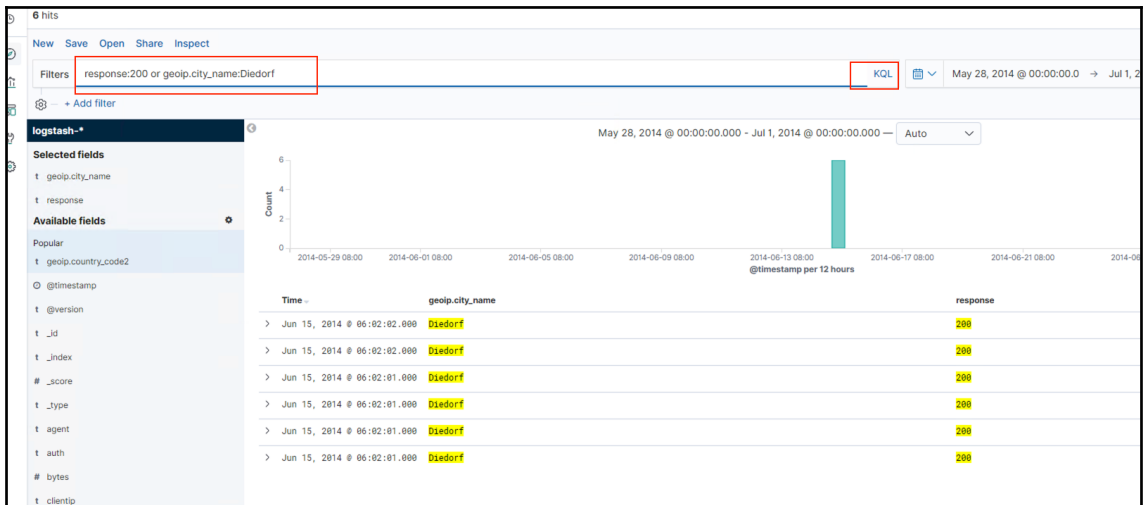
- # bytes


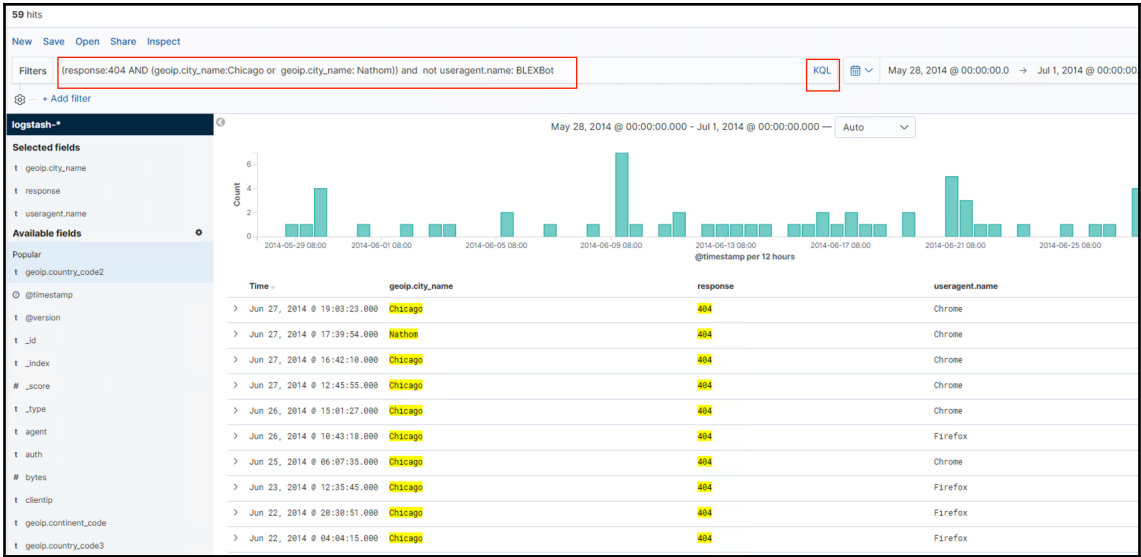
- t clientip



Time	response
> Jun 27, 2014 @ 19:58:29.000	304
> Jun 27, 2014 @ 19:54:26.000	304
> Jun 27, 2014 @ 19:50:27.000	304
> Jun 27, 2014 @ 19:48:45.000	304
> Jun 27, 2014 @ 19:32:50.000	304
> Jun 27, 2014 @ 19:27:28.000	301
> Jun 27, 2014 @ 19:22:30.000	404
> Jun 27, 2014 @ 19:21:38.000	404







## Save search ✕

Save as a new search

Title

custom\_query1

Cancel

Confirm Save

### Open Search ✕

🔍 Search... Manage searches

**Title**

---

custom\_query1

---

responseBy200

---

Rows per page: 15 ▾

**Share** **Inspect**

---

**PERMALINK**

---

Generate the link as

Snapshot ?

Saved object ?

Can't share as saved object until the search has been saved.

✕ Short URL ?

Copy link

## custom\_query1


1 request was made

**Request:** Segment 0

This request queries Elasticsearch to fetch the data for the search.

Statistics [Request](#) Response

```
{
  "version": true,
  "size": 500,
  "sort": [
    {
      "@timestamp": {
        "order": "desc",
        "unmapped_type": "boolean"
      }
    }
  ],
  "_source": {
    "excludes": []
  },
  "aggs": {
    "2": {
      "date_histogram": {
        "field": "@timestamp",
        "interval": "10m",
        "time_zone": "Asia/Singapore",
        "min_doc_count": 1
      }
    }
  },
  "stored_fields": [
    "*"
  ],
  "script_fields": {},
  "docvalue_fields": [
    {
      "field": "@timestamp",
      "format": "date_time"
    }
  ],
  "query": {
    "bool": {
      "must": [
        {
          "bool": {
            "must": [
              {
                "match": {
                  "useragent.name": "IE"
                }
              },
              {
                "match": {
                  "geoip.region_name": "Washington"
                }
              }
            ]
          }
        }
      ]
    }
  }
}
```

 Jun 13, 2014 @ 00:00:00.0 → Jun 13, 2014 @ 12:00:00.0

**Quick select** < >

Last 15 minutes Apply

---

**Commonly used**

<a href="#">Today</a>	<a href="#">This week</a>
<a href="#">This month</a>	<a href="#">This year</a>
<a href="#">Today so far</a>	<a href="#">Week to date</a>
<a href="#">Month to date</a>	<a href="#">Year to date</a>

~ 5 years ago → Jun 13, 2014 @ 12:00:00.000

Absolute Relative Now

5 Years ago

May 19, 2014 @ 10:15:12.134

X Round to the year



May 19, 2014 @ 10:14:20.9 → Jun 13, 2014 @ 12:00:00.0

Absolute

Relative

Now

< May 2014 >							07:30 AM
SU	MO	TU	WE	TH	FR	SA	08:00 AM
27	28	29	30	1	2	3	08:30 AM
4	5	6	7	8	9	10	09:00 AM
11	12	13	14	15	16	17	09:30 AM
18	19	20	21	22	23	24	10:00 AM
25	26	27	28	29	30	31	10:30 AM
							11:00 AM
							11:30 AM
							12:00 PM

2014-05-19 10:14:20.952

May 28, 2014 @ 00:00:00.0 → Jul 1, 2014 @ 00:00:00.0

**Quick select** < >

Last 15 minutes **Apply**

---

**Commonly used**

Today	This week
This month	This year
Today so far	Week to date
Month to date	Year to date

---

**Recently used date ranges**

May 28, 2014 @ 00:00:00.000 to Jul 1, 2014 @ 00:00:00.000

May 28, 2016 @ 00:00:00.000 to Jul 1, 2016 @ 00:00:00.000

Last 15 minutes

Jun 13, 2014 @ 00:00:00.000 to Jun 13, 2014 @ 00:00:00.000

---

**Refresh every**

10 seconds **Start**

The screenshot displays a data analysis interface. At the top, a bar chart shows data over time from 2014-05-29 to 2014-06-17. Below the chart is a log entry for a request on June 27th, 2014, at 17:30:00.000, showing details like IP address (183.60.215.50) and user agent (EasouSpider). The main part of the interface is a table of fields with search and filter icons. Annotations include:

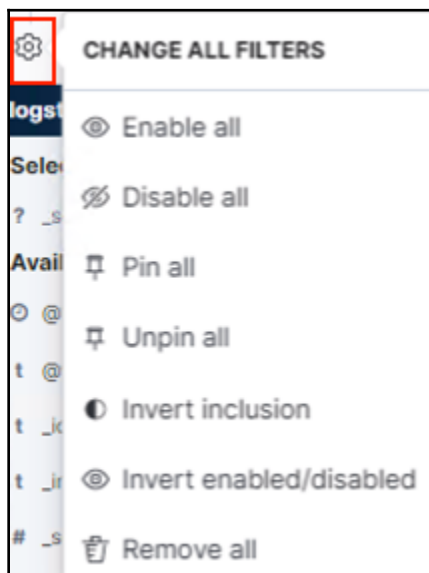
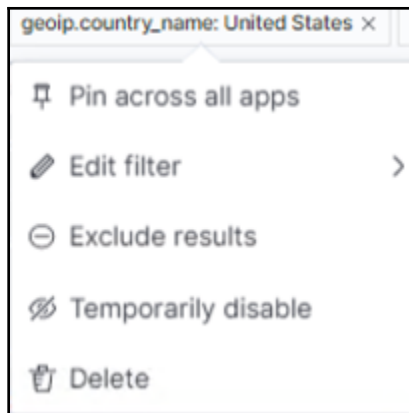
- negative filter:** Points to the search icon for the field `geop.city_name`.
- exists filter:** Points to the search icon for the field `@timestamp`.
- positive filter:** Points to the search icon for the field `geop.continent_code`.
- Filter for field present:** A tooltip that appears over the search icon for `@timestamp`.

On the left, there is a list of available fields and a 'Top 5 values in 393 / 500 records' section showing a bar chart of city names: San Jose (33.3%), Amsterdam (9.2%), Seattle (6.6%), Moscow (4.6%), and Chantilly (4.3%).

The 'EDIT FILTER' dialog box is shown with the following configuration:

- Field:** `geop.country_n...`
- Operator:** `is`
- Value:** `United States`
- Create custom label?:**

Buttons for 'Cancel' and 'Save' are visible at the bottom.



# New Visualization



Filter

## Select a visualization type

Start creating your visualization by selecting a type for that visualization.



Area



Controls



Coordinate Map



Data Table



Gauge



Goal



Heat Map



Horizontal Bar



Line



Markdown



Metric



Pie



Region Map



Tag Cloud



Timelion



Vega

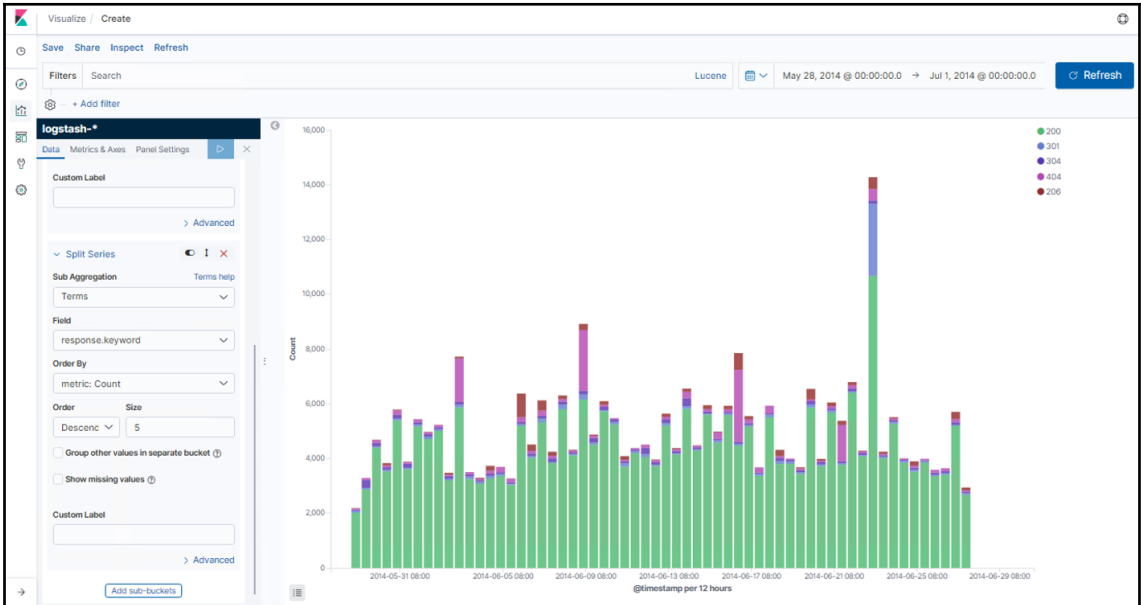
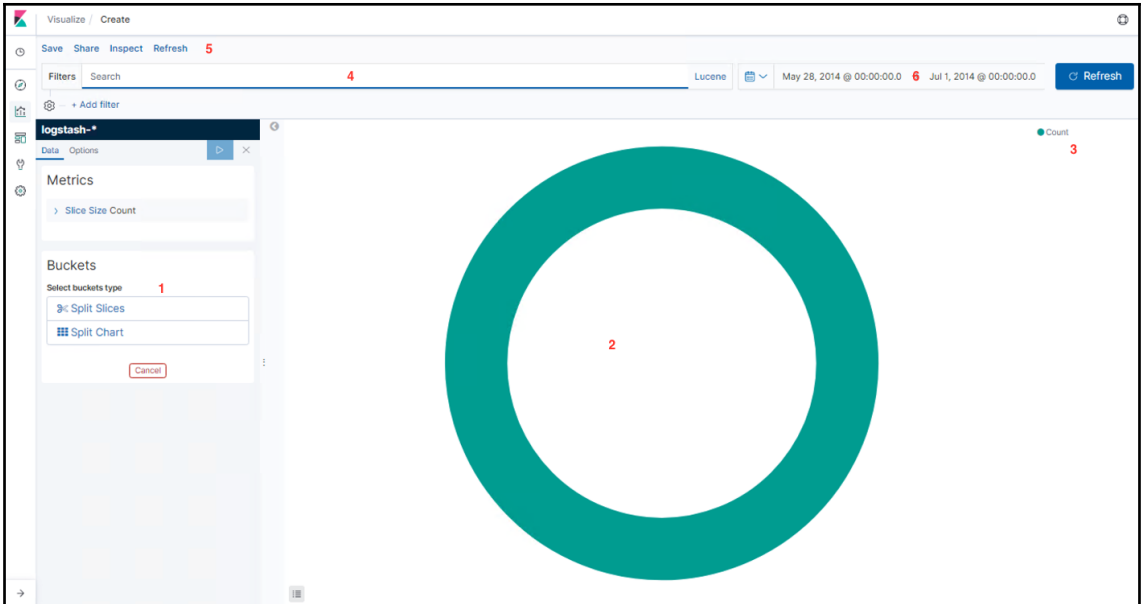


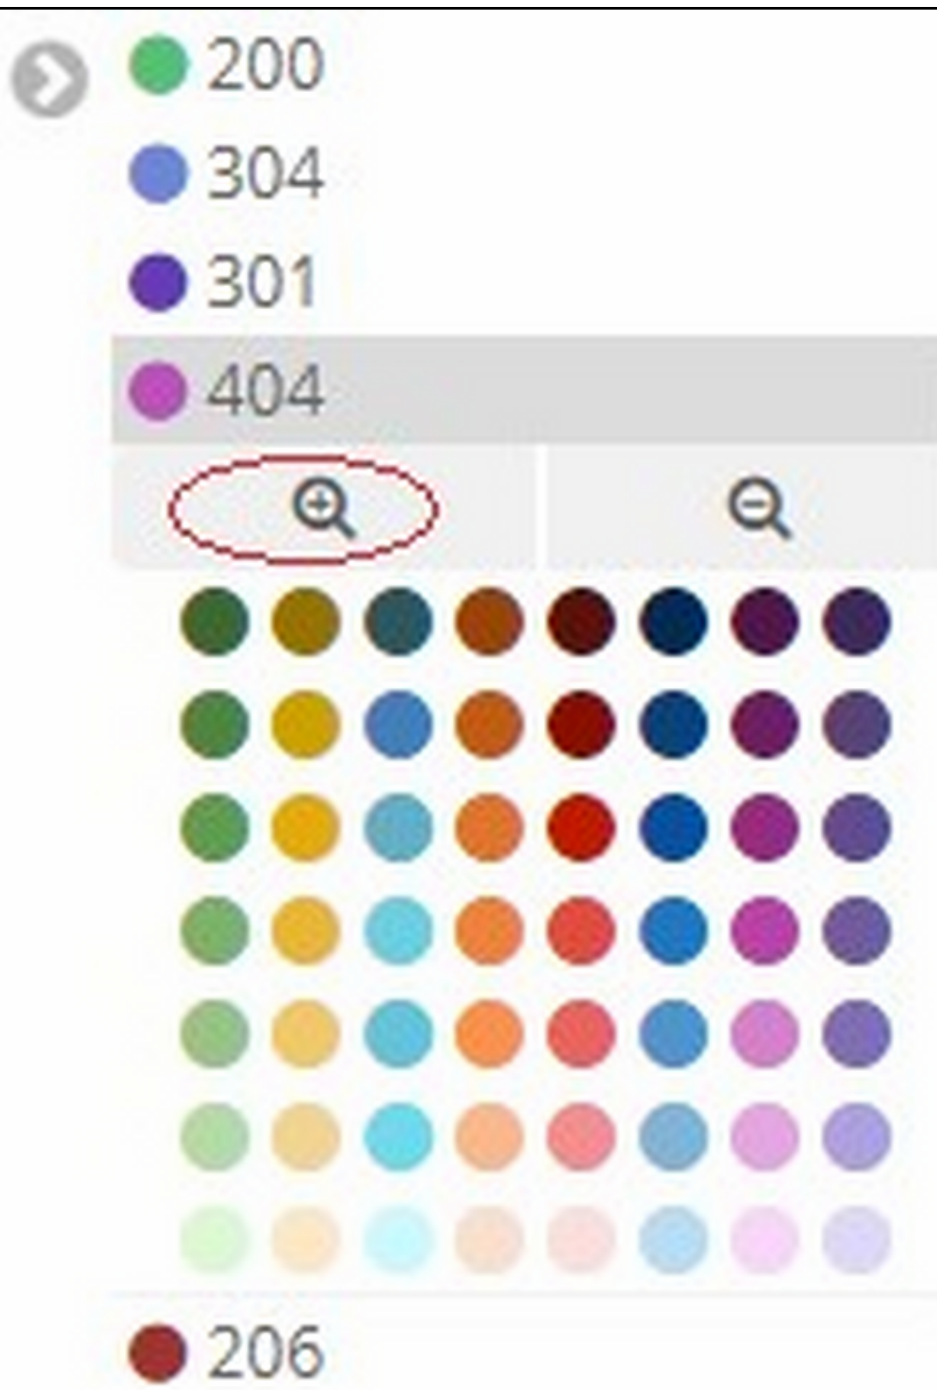
Vertical Bar

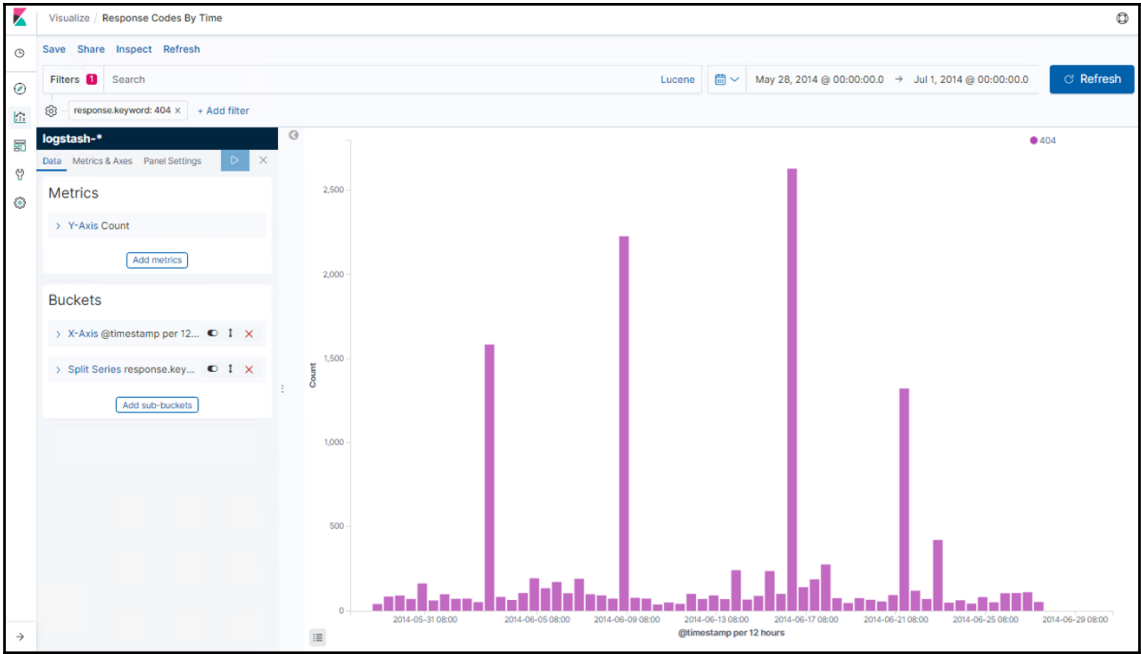


Visual Builder









Visualize / Create

Save Share Inspect Refresh

Filters Search Lucene May 28, 2014 @ 00:00:00.0 → Jul 1, 2014 @ 00:00:00.0 Refresh

+ Add filter

logstash\*

Data Options

Add metrics

Buckets

Split Rows

Aggregation Terms help

Terms

Field request.keyword

Order By metric: Total Requests

Order Descend Size 10

Group other values in separate bucket

Show missing values

Custom Label


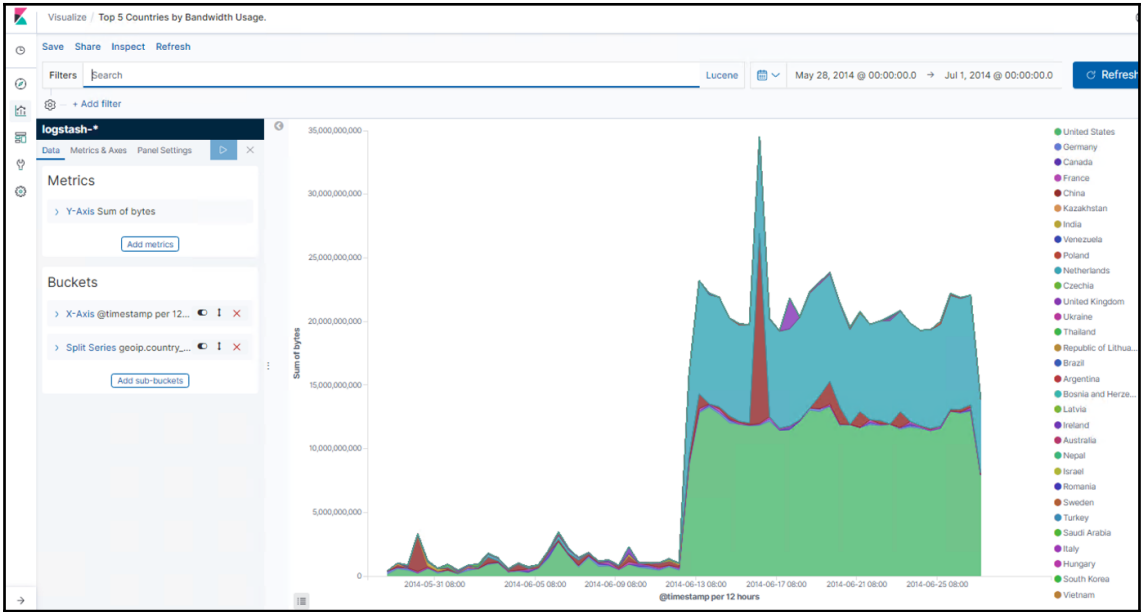
Urls

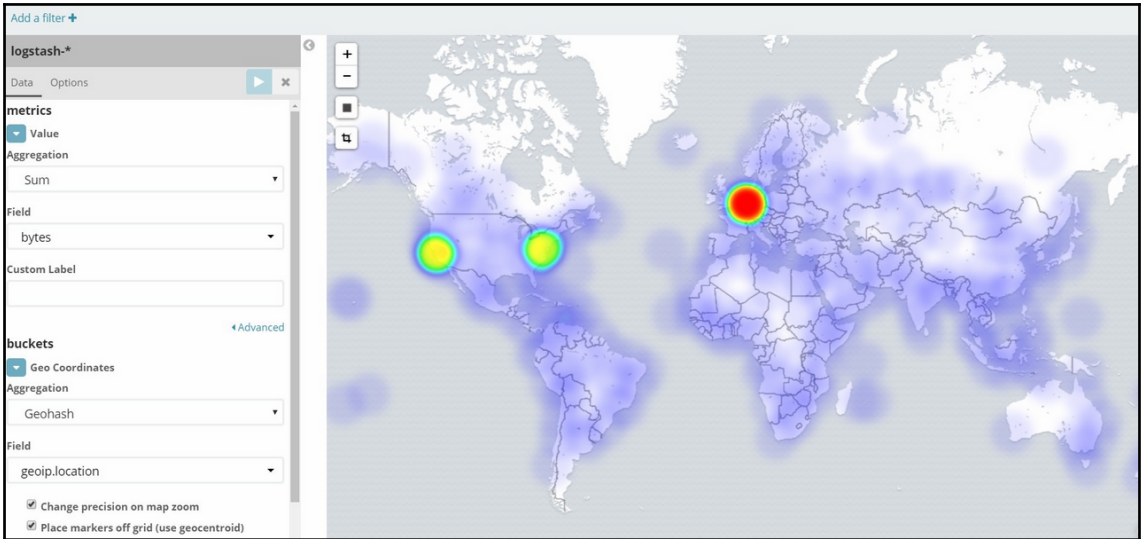
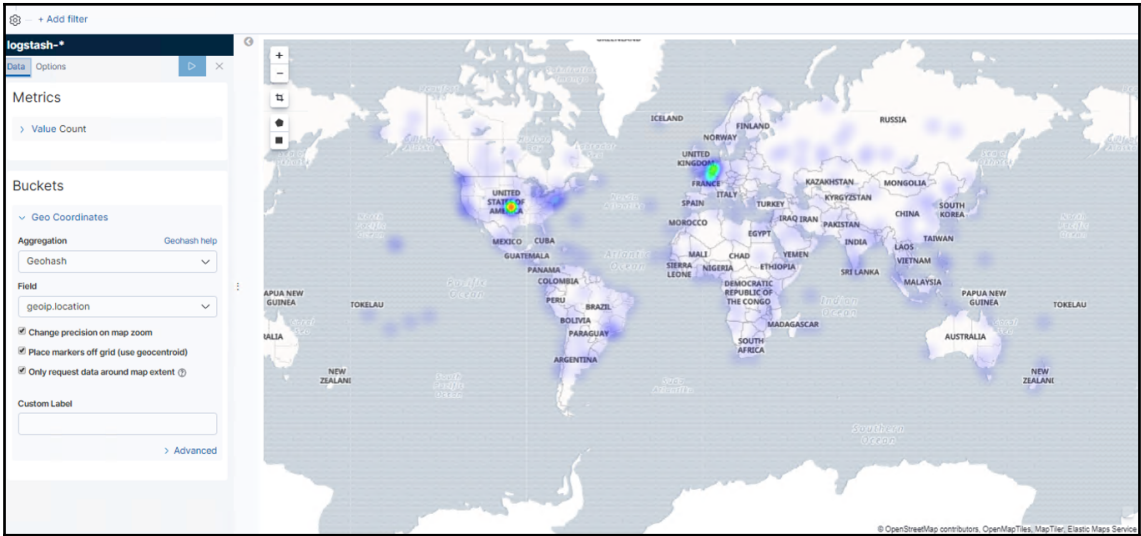
Urls	Total Requests
/favicon.ico	18,893
/files/logstash/logstash-1.1.0-monolithic.jar	14,755
/style2.css	12,925
/reset.css	12,821
/images/jordan-80.png	12,521
/images/web/2009/banner.png	12,236
/blog/tags/puppet?flav=rss20	11,379
/	6,295
/presentations/fpm-scale12x.pdf	5,282
?flav=rss20	5,103

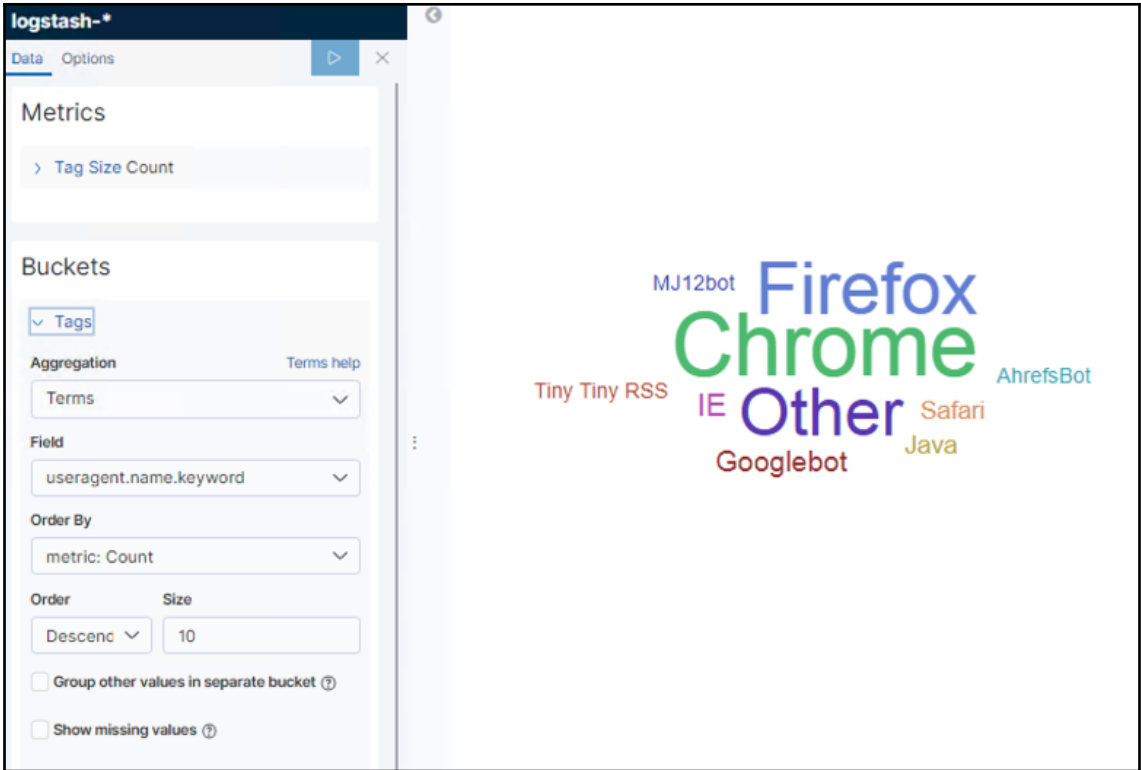
Export: Raw Formatted

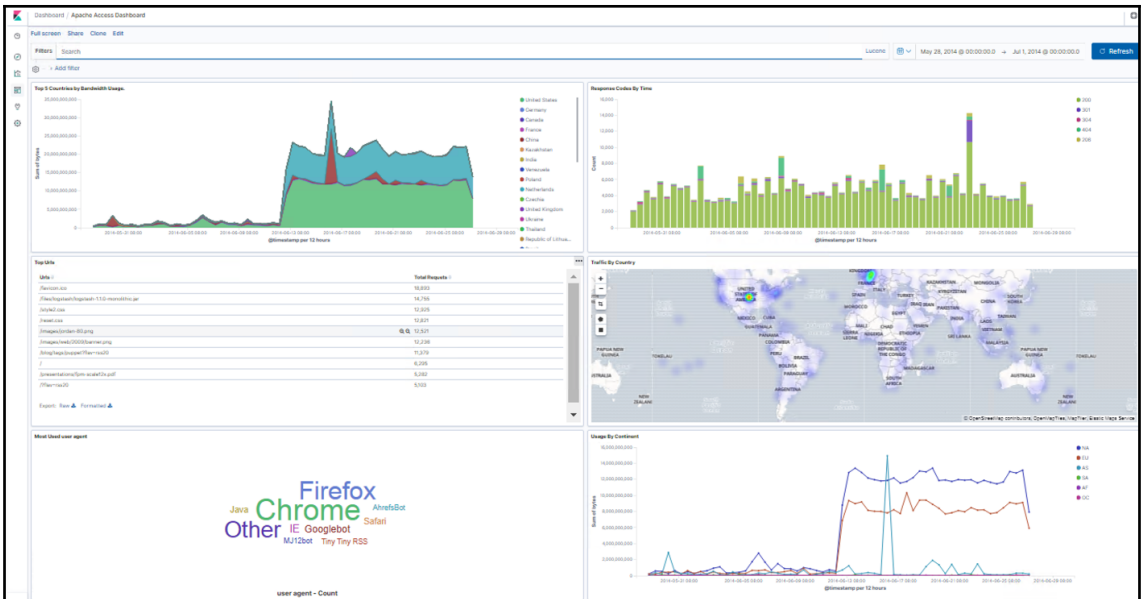
> Advanced











Dashboards

Create your first dashboard

You can combine data views from any Kibana app into one dashboard and see everything in one place.

New to Kibana? Install some sample data to take a test drive.

[+ Create new dashboard](#)

Dashboard - Editing New Dashboard

Save Cancel Add Options Share

Filters Search

+ Add filter

This dashboard is empty. Click the **Add** button in the menu to add a visualization. If you haven't set up any visualizations, create your own.

### Add Panels

Visualization Saved Search

Search... [Add new Visualization](#)

**Title**

- Top 5 Countries by Bandwidth Usage.
- Response Codes By Time
- Top Urls
- Traffic By Country
- Most Used user agent
- Usage By Continent

Rows per page: 15 ▾



## Save dashboard



Save as a new dashboard



**Title**

Apache Access Dashboard

**Description**

**Store time with dashboard**



This changes the time filter to the currently selected time each time this dashboard is loaded.

Cancel

Confirm Save

×

# Clone Dashboard

Please enter a new name for your dashboard.

Cancel Confirm Clone

Share Clone Edit

### SHARE THIS DASHBOARD

>.. Embed code >

[Permalinks](#) >

Share Clone Edit

### < PERMALINK

Generate the link as

Snapshot ⓘ

Saved object ⓘ

X Short URL ⓘ

Copy link

Share Clone Edit

< **EMBED CODE**

Generate the link as

- Snapshot ⓘ
- Saved object ⓘ
- X Short URL ⓘ

Copy iFrame code

# New Visualization

Filter

**Timelion**  
Build time-series using functional expressions

- Area
- Controls ⓘ
- Coordinate Map
- Data Table
- Gauge
- Goal
- Heat Map
- Horizontal Bar
- Line
- Markdown
- Metric
- Pie
- Region Map
- Tag Cloud
- Timelion**
- Vega ⓘ
- Vertical Bar
- Visual Builder



▶ ✕

**Interval**

auto▼

**Timelion Expression**

|

**.abs()** Return the absolute value of each value in the series list (Chainable)

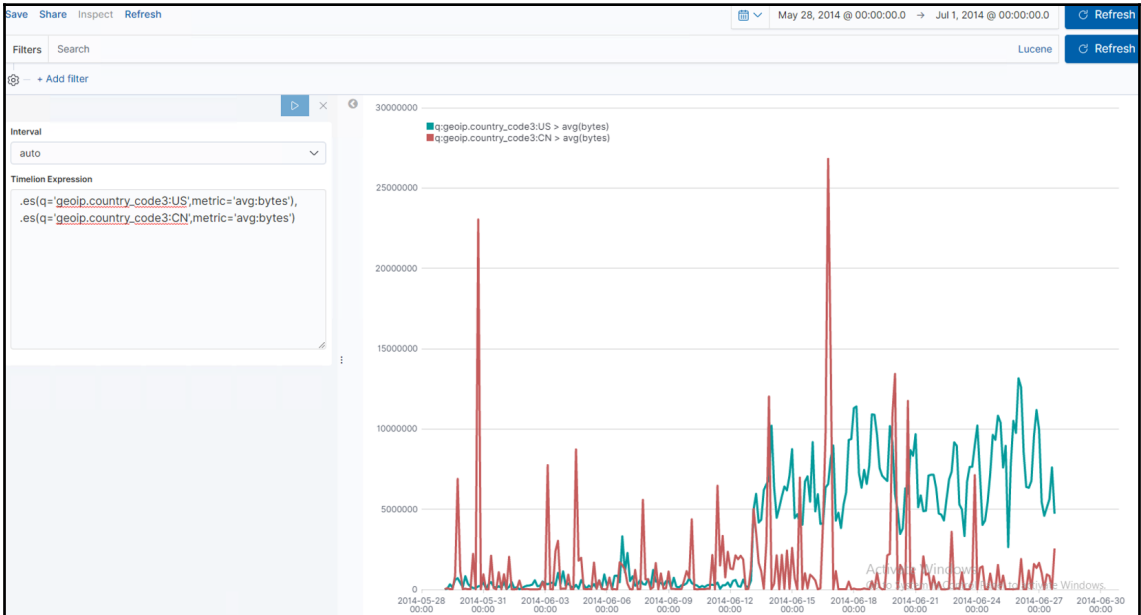
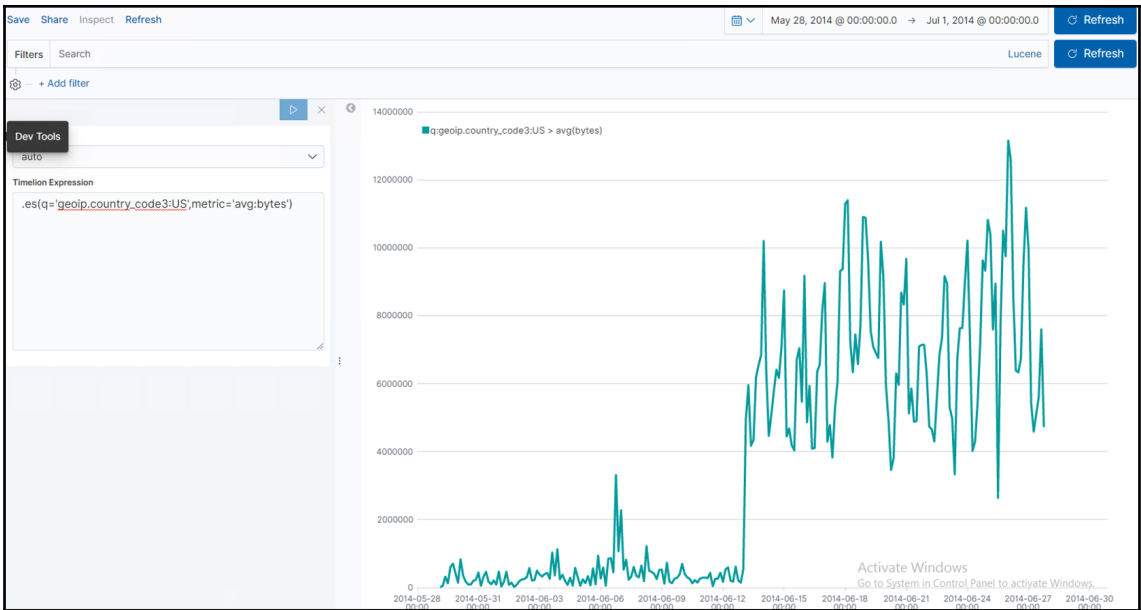
**.add()** Adds the values of one or more series in a seriesList to each position, in each series, of the input seriesList (Chainable)  
**Arguments:** `term=(seriesList | number)`

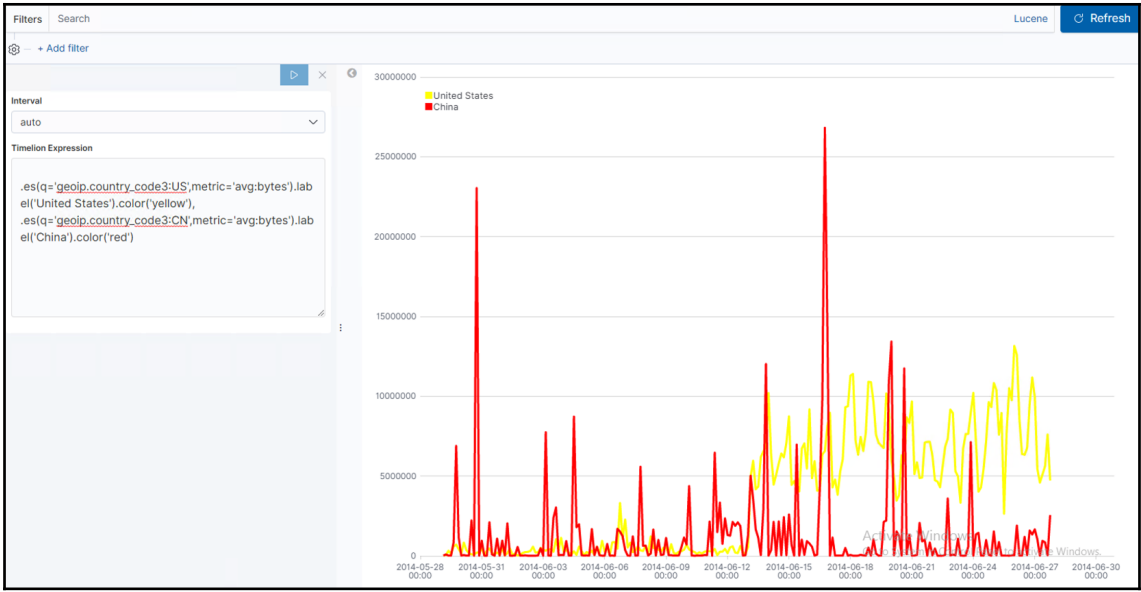
**.aggregate()** Creates a static line based on result of processing all points in the series. Available functions: avg, cardinality, min, max, last, first, sum (Chainable)  
**Arguments:** `function=(string)`

**.bars()** Show the seriesList as bars (Chainable)  
**Arguments:** `width=(number | null) , stack=(boolean | null)`

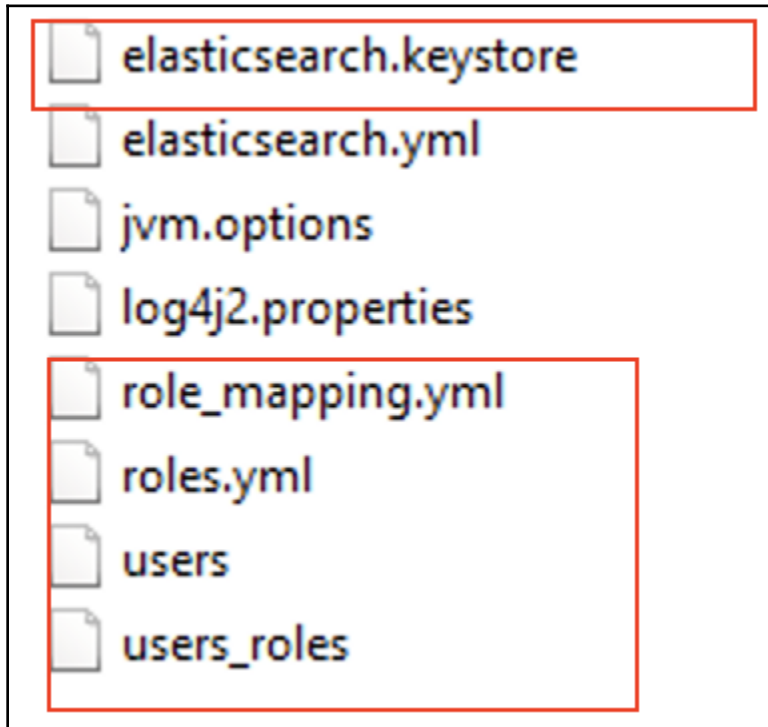
**.color()** Change the color of the series (Chainable)  
**Arguments:** `color=(string)`

**.condition()** Compares each point to a number, or the same point in another series using an operator. then sets its





## Chapter 8: Elastic X-Pack



localhost:5601/app/kibana#/home?\_g=0

Home

Recently viewed

Discover

Visualize

Dashboard

Canvas

Maps

Machine Learning

Infrastructure

Logs

APM

Uptime

Dev Tools

Stack Monitoring

Management

### Add Data to Kibana

Use these solutions to quickly turn your data into pre-built dashboards and monitoring systems.

**APM**  
APM automatically collects in-depth performance metrics and errors from inside your applications.  
[Add APM](#)

**Logging**  
Ingest logs from popular data sources and easily visualize in preconfigured dashboards.  
[Add log data](#)

**Metrics**  
Collect metrics from the operating system and services running on your servers.  
[Add metric data](#)

**Security analytics**  
Centralize security events for interactive investigation in ready-to-go visualizations.  
[Add security events](#)

**Add sample data**  
Load a data set and a Kibana dashboard

**Upload data from log file**  
Import a CSV, NDJSON, or log file

**Use Elasticsearch data**  
Connect to your Elasticsearch index

### Visualize and Explore Data

**APM**  
Automatically collect in-depth performance metrics and errors from inside your applications.

**Canvas**  
Showcase your data in a pixel-perfect way.

**Dashboard**  
Display and share a collection of visualizations and saved searches.

**Discover**  
Interactively explore your data by querying and filtering raw documents.

**Graph**  
Surface and analyze relevant relationships in your Elasticsearch data.

**Infrastructure**  
Explore infrastructure metrics and logs for common servers, containers, and services.

**Logs**  
Stream logs in real time or scroll through historical

**Machine Learning**  
Automatically model the normal behavior of your

### Manage and Administer the Elastic Stack

**Console**  
Skip cURL and use this JSON interface to work with your data directly.

**Index Patterns**  
Manage the index patterns that help retrieve your data from Elasticsearch.

**Monitoring**  
Track the real-time health and performance of your Elastic Stack.

**Saved Objects**  
Import, export, and manage your saved searches, visualizations, and dashboards.

**Spaces**  
Organize your dashboards

**Rollups**  
Summarize and store historical data in a smaller index for future analysis.

**Security Settings**  
Protect your data and easily manage who has access to what with users and roles.

**Watcher**  
Detect changes in your

← Collapse

Management / License management

Elasticsearch

- Index Management
- Index Lifecycle Policies
- Rollup Jobs
- Cross Cluster Replication
- Remote Clusters
- License Management**
- 8.0 Upgrade Assistant

Kibana

- Index Patterns
- Saved Objects
- Spaces
- Reporting
- Advanced Settings

Management

Your Basic license is inactive  
Your license will never expire.

**Update your license**  
If you already have a new license, upload it now.

[Update license](#)

**Start a 30-day trial**  
Experience what security, machine learning, and all our other Platinum features have to offer.

[Start trial](#)

## Start your free 30-day trial

This trial is for the full set of **Platinum features** of the Elastic Stack. You'll get immediate access to:

- Machine learning
- Alerting
- Graph capabilities
- JDBC and ODBC connectivity for SQL

Security features, such as authentication (native, AD/LDAP, SAML, PKI), role-based access control, and auditing, require configuration. See the [documentation](#) for instructions.

By starting this trial, you agree that it is subject to these [terms and conditions](#).

Send basic feature usage statistics to Elastic periodically. [Read more](#) [Cancel](#) [Start my trial](#)

Management / License management

**Elasticsearch**

- Index Management
- Index Lifecycle Policies
- Rollup Jobs
- Cross Cluster Replication
- Remote Clusters
- Watcher
- [License Management](#)
- 8.0 Upgrade Assistant

**Kibana**

- Index Patterns
- Saved Objects

**Your Trial license is active**  
Your license will expire on **June 20, 2019 10:49 AM +08**

**Update your license**  
If you already have a new license, upload it now.

[Update license](#)

**Extend your trial**  
If you'd like to continue using security, machine learning, and our other awesome Platinum features, request an extension now.

[Extend trial](#)

**Revert to Basic license**  
You'll revert to our free features and lose access to security, machine learning and other Platinum features.

[Revert to Basic](#)

← → ↻ ⓘ localhost:9200

**Sign in**  
http://localhost:9200

Username

Password

[Sign in](#) [Cancel](#)



# Welcome to Kibana

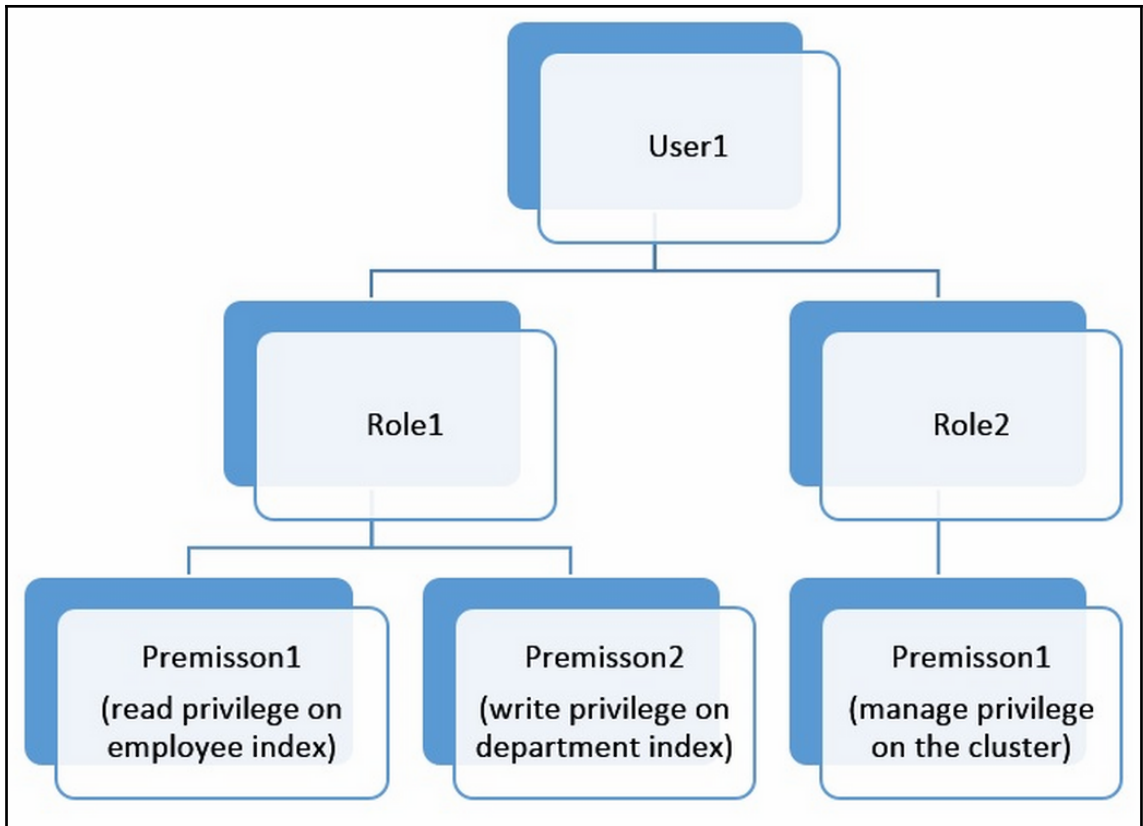
Your window into the Elastic Stack

Username

Password

Log in





localhost:5601/app/kibana#/management/security/users?\_g=0

Management / Users

**Elasticsearch**

- Index Management
- Index Lifecycle Policies
- Rollup Jobs
- Cross Cluster Replication
- Remote Clusters
- Watcher
- License Management
- 8.0 Upgrade Assistant

**Kibana**

- Index Patterns
- Saved Objects
- Spaces
- Reporting
- Advanced Settings

**Logstash**

- Pipelines

**Beats**

- Central Management

**Security**

- Users**
- Roles

### Users

Create new user

Search...

Full Name ↑	User Name	Email Address	Roles	Reserved
<input type="checkbox"/>	elastic		superuser	✓
<input type="checkbox"/>	kibana		kibana_system	✓
<input type="checkbox"/>	logstash_system		logstash_system	✓
<input type="checkbox"/>	beats_system		beats_system	✓
<input type="checkbox"/>	apm_system		apm_system	✓
<input type="checkbox"/>	remote_monitoring_user		remote_monitoring_collector, remote_monitoring_agent	✓

Rows per page: 20

### Users

Create new user

Search...

**Default Users**

Full Name ↑	User Name	Email Address	Roles	Reserved
<input type="checkbox"/>	elastic		superuser	✓
<input type="checkbox"/>	kibana		kibana_system	✓
<input type="checkbox"/>	logstash_system		logstash_system	✓
<input type="checkbox"/>	beats_system		beats_system	✓
<input type="checkbox"/>	apm_system		apm_system	✓
<input type="checkbox"/>	remote_monitoring_user		remote_monitoring_collector, remote_monitoring_agent	✓

Rows per page: 20

## New user

Username

Password

Confirm password

Full name

Email address

Roles



# Users

Create new user

Delete 1 user

Search...

<input type="checkbox"/> Full Name ↑	User Name	Email Address	Roles	Reserved
<input type="checkbox"/>	user1	user1@packt.com		
<input checked="" type="checkbox"/>	user2	user2@packt.com		
<input type="checkbox"/>	elastic		superuser	✓
<input type="checkbox"/>	kibana		kibana_system	✓
<input type="checkbox"/>	logstash_system		logstash_system	✓
<input type="checkbox"/>	beats_system		beats_system	✓
<input type="checkbox"/>	apm_system		apm_system	✓
<input type="checkbox"/>	remote_monitoring_user		remote_monitoring_collector, remote_monitoring_agent	✓

Rows per page: 20

## Edit user1 user

**Username**

Username's cannot be changed after creation.

**Full name**

**Email address**

**Roles**



**Password**

[Change password](#)

---

Management / Roles

**Elasticsearch**

- Index Management
- Index Lifecycle Policies
- Rollup Jobs
- Cross Cluster Replication
- Remote Clusters
- Watcher
- License Management
- 8.0 Upgrade Assistant

**Kibana**

- Index Patterns
- Saved Objects
- Spaces
- Reporting
- Advanced Settings

**Logstash**

- Pipelines

**Beats**

- Central Management

**Security**

- Users
- Roles**

## Roles

Apply roles to groups of users and manage permissions across the stack

Q Search...

[+ Create role](#)

<input type="checkbox"/> Role ↑	Reserved ⓘ
<input type="checkbox"/> apm_system	✓
<input type="checkbox"/> apm_user	✓
<input type="checkbox"/> beats_admin	✓
<input type="checkbox"/> beats_system	✓
<input type="checkbox"/> code_admin	✓
<input type="checkbox"/> code_user	✓
<input type="checkbox"/> ingest_admin	✓
<input type="checkbox"/> kibana_dashboard_only_user	✓
<input type="checkbox"/> kibana_system	✓
<input type="checkbox"/> kibana_user	✓
<input type="checkbox"/> logstash_admin	✓
<input type="checkbox"/> logstash_system	✓
<input type="checkbox"/> machine_learning_admin	✓
<input type="checkbox"/> machine_learning_user	✓
<input type="checkbox"/> monitoring_user	✓
<input type="checkbox"/> remote_monitoring_agent	✓
<input type="checkbox"/> remote_monitoring_collector	✓
<input type="checkbox"/> reporting_user	✓
<input type="checkbox"/> rollup_admin	✓
<input type="checkbox"/> rollup_user	✓

Management / Users / Create

## Create role

Set privileges on your Elasticsearch data and control access to your Kibana spaces.

Role name:

---

**Elasticsearch** hide

### Cluster privileges

Manage the actions this role can perform against your cluster. [Learn more](#)

### Run As privileges

Allow requests to be submitted on the behalf of other users. [Learn more](#)

### Index privileges

Control access to the data in your cluster. [Learn more](#)

Indices:   
Privileges:   
Granted fields (optional):

Grant read privileges to specific documents

[Add index privilege](#)

**Kibana** hide

### Minimum privileges for all spaces

Specify the minimum actions users can perform in your spaces.

No access to spaces

### Higher privileges for individual spaces

Grant more privileges on a per space basis. For example, if the privileges are **read** for all spaces, you can set the privileges to **all** for an individual space.

[Add space privilege](#) [View summary of spaces p](#)

[Create role](#) [Cancel](#)

## Edit user1 user

Username

Username's cannot be changed after creation.

Full name

Email address

Roles

monitor\_role

monitoring\_user

remote\_monitoring\_agent

remote\_monitoring\_collector

Update user

Cancel

Delete user



# Roles

Apply roles to groups of users and manage permissions across the stack

<input checked="" type="checkbox"/> Role ↑	Reserved ⓘ
<input type="checkbox"/> apm_system	✓
<input type="checkbox"/> apm_user	✓
<input type="checkbox"/> beats_admin	✓
<input type="checkbox"/> beats_system	✓
<input type="checkbox"/> code_admin	✓
<input type="checkbox"/> code_user	✓
<input type="checkbox"/> ingest_admin	✓
<input type="checkbox"/> kibana_dashboard_only_user	✓
<input type="checkbox"/> kibana_system	✓
<input type="checkbox"/> kibana_user	✓
<input type="checkbox"/> logstash_admin	✓
<input type="checkbox"/> logstash_system	✓
<input type="checkbox"/> machine_learning_admin	✓
<input type="checkbox"/> machine_learning_user	✓
<input checked="" type="checkbox"/> monitor_role	

Q Search...

Delete

# Edit role

Set privileges on your Elasticsearch data and control access to your Kibana spaces.

**Role name**  
monitor\_role

A role's name cannot be changed once it has been created.

## Elasticsearch hide

### Cluster privileges

Manage the actions this role can perform against your cluster. [Learn more](#)

monitor × manage ×

### Run As privileges

Allow requests to be submitted on the behalf of other users. [Learn more](#)

Add a user...

### Index privileges

Control access to the data in your cluster. [Learn more](#)

Indices:  Prileges:  Granted fields (optional):

Grant read privileges to specific documents

+ Add index privilege

## Kibana hide

### Minimum privileges for all spaces

Specify the minimum actions users can perform in your spaces.

none

No access to spaces

### Higher privileges for individual spaces

Grant more privileges on a per space basis. For example, if the privileges are **read** for all spaces, you can set the privileges to **all** for an individual space.

+ Add space privilege

[View summary of spaces pr](#)

Update role

Cancel

De

# Create role

Set privileges on your Elasticsearch data and control access to your Kibana spaces.

**Role name**  
employee\_read



### Cluster privileges

Manage the actions this role can perform against your cluster. [Learn more](#)

### Run As privileges

Allow requests to be submitted on the behalf of other users. [Learn more](#)

### Index privileges

Control access to the data in your cluster. [Learn more](#)

Indices	Privileges	Granted fields (optional)
employee ×	read ×	gender × state × email ×

Grant read privileges to specific documents

+ Add index privilege

## Edit user2 user

Username

Username's cannot be changed after creation.

Full name

Email address

Roles



Password

[Change password](#)

[Update user](#)

[Cancel](#)

[Delete user](#)

# Create role

Set privileges on your Elasticsearch data and control access to your Kibana spaces.

Role name

department\_IT\_role



## Cluster privileges

Manage the actions this role can perform against your cluster. [Learn more](#)

## Run As privileges

Allow requests to be submitted on the behalf of other users. [Learn more](#)

## Index privileges

Control access to the data in your cluster. [Learn more](#)

Indices

department ×

Privileges

read ×

Granted fields (optional)

\* ×

Grant read privileges to specific documents

Granted documents query

```
("match":{"name": "IT"})
```

## Edit user1 user

Username

Username's cannot be changed after creation.

Full name

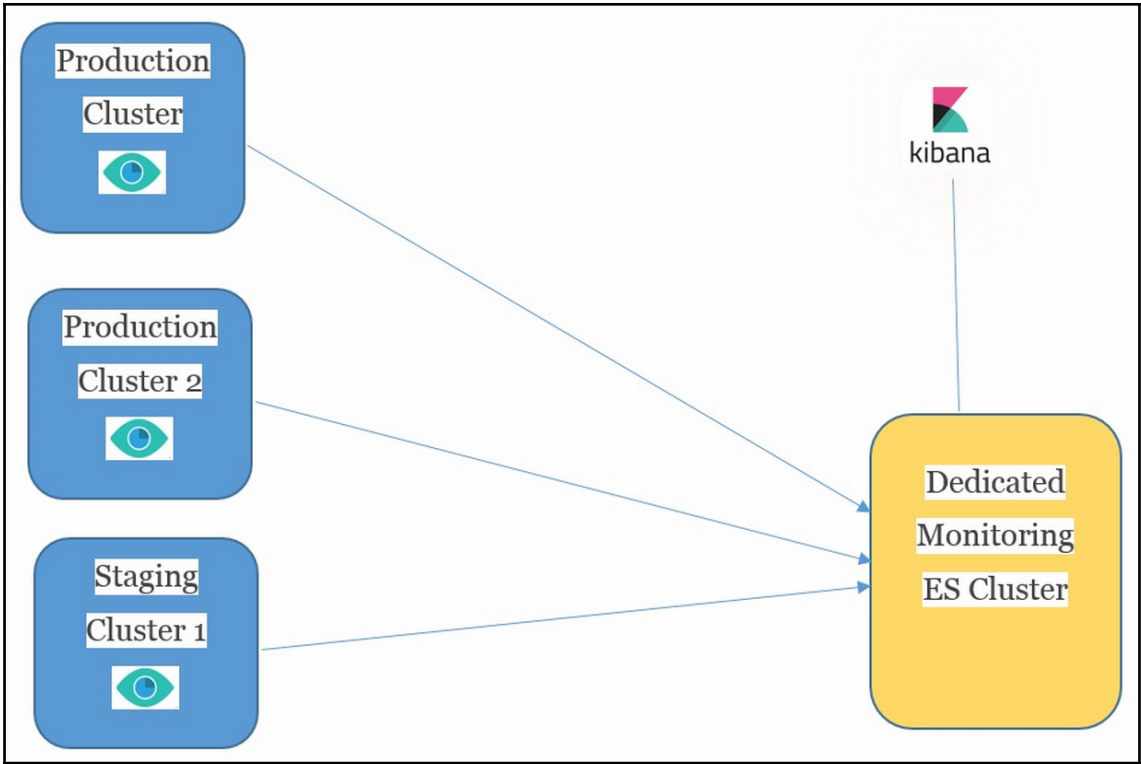
Email address

Roles



Password

[Change password](#)




localhost:5601/app/monitoring#/no-data?\_g=0

Stack Monitoring

Last 1 hour

Clusters




## Monitoring is currently off

Monitoring provides insight to your hardware performance and load.

We checked the cluster defaults settings and found that `xpack.monitoring.collection.enabled` is set to `false`.

Would you like to turn it on?

[Turn on monitoring](#)





Clusters

elasticsearch

Top cluster alerts

Medium severity alert  
Elasticsearch cluster status is yellow. Allocate missing replica shards.  
Last checked May 23, 2019 4:16:36 PM (triggered 14 min ago)

Elasticsearch • Health is yellow Trial license will expire on June 20, 2019

Overview		Nodes: 1		Indices: 11	
Version	7.0.0	Disk Available	87.06%	Documents	1,732
Uptime	2 days		522.1 GB / 599.7 GB	Disk Usage	1.8 MB
Jobs	0	JVM Heap	15.74%	Primary Shards	11
			155.8 MB / 989.9 MB	Replica Shards	0

Kibana • Health is green

Overview		Instances: 1	
Requests	2	Connections	6
Max. Response Time	58 ms	Memory Usage	14.24%
			207.3 MB / 1.4 GB

Clusters / elasticsearch / Elasticsearch

Overview Nodes Indices Jobs CCR

Status	Nodes	Indices	Memory	Total Shards	Unassigned Shards	Documents	Data
Yellow	1	39	663.7 MB / 989.9 MB	69	30	256,743	164.3 MB

Search Rate (/s)

● Total Shards 4.27 /s

Search Latency (ms)

● Search Latency 1.2 ms

Indexing Rate (/s)

● Total Shards 3,666.7 /s

● Primary Shards 3,666.7 /s

Indexing Latency (ms)

● Indexing Latency 0.55 ms

Shard Activity

Completed recoveries

Index	Stage	Total Time	Source / Destination	Files	Bytes	Translog
There are no active shard recoveries for this cluster. Try viewing completed recoveries.						

Clusters / elasticsearch / Elasticsearch

Overview Nodes Indices Jobs CCR

Last 1 hour Show dates Refresh

Status	Nodes	Indices	Memory	Total Shards	Unassigned Shards	Documents	Data
Yellow	1	42	474.6 MB / 989.9 MB	75	33	304,564	176.2 MB

Filter Nodes...

Name ↑	Status	CPU Usage	Load Average	JVM Memory	Disk Free Space	Shards
★ MADSH01-APM01 127.0.0.1:9300	Online	7% ↑ 10% max 0% min	N/A N/A max N/A min	67% ↑ 73% max 22% min	521.5 GB ↓ 522.1 GB max 521.5 GB min	42

Rows per page: 20

Clusters / elasticsearch / Nodes / b5MAMCF5c2B6rTuno3mg

Overview Advanced

Status	Transport Address	JVM Heap	Free Disk Space	Documents	Data	Indices	Shards	Type
Online	1270.0.1:9300	45%	521.5 GB	305k	176.5 MB	42	42	Master Node

JVM Heap (MB) Ⓞ

● Max Heap 889.9 MB  
● Used Heap 575.4 MB

Index Memory (MB) Ⓞ

● Lucene Total 3.9 MB  
● Terms 18.7 kB

CPU Utilization (%) Ⓞ

● CPU Utilization 4.48%

System Load Ⓞ

● No Inpact (Open)

Latency (ms) Ⓞ

● Search 4.35 ms

Segment Count Ⓞ

● Segment Count 232

Shard Legend

Primary  Replica  Relocating  Unassigned Primary  Unassigned Replica

Indices  System Indices

Index: [1] [2] [3] [4] [5] [6] [7] [8] [9] [10] [11] [12] [13] [14] [15] [16] [17] [18] [19] [20] [21] [22] [23] [24] [25] [26] [27] [28] [29] [30] [31] [32] [33] [34] [35] [36] [37] [38] [39] [40] [41] [42]



Clusters / elasticsearch / Elasticsearch

Overview Nodes Indices Jobs OCR

▼ Last 1 hour Show dates Refresh

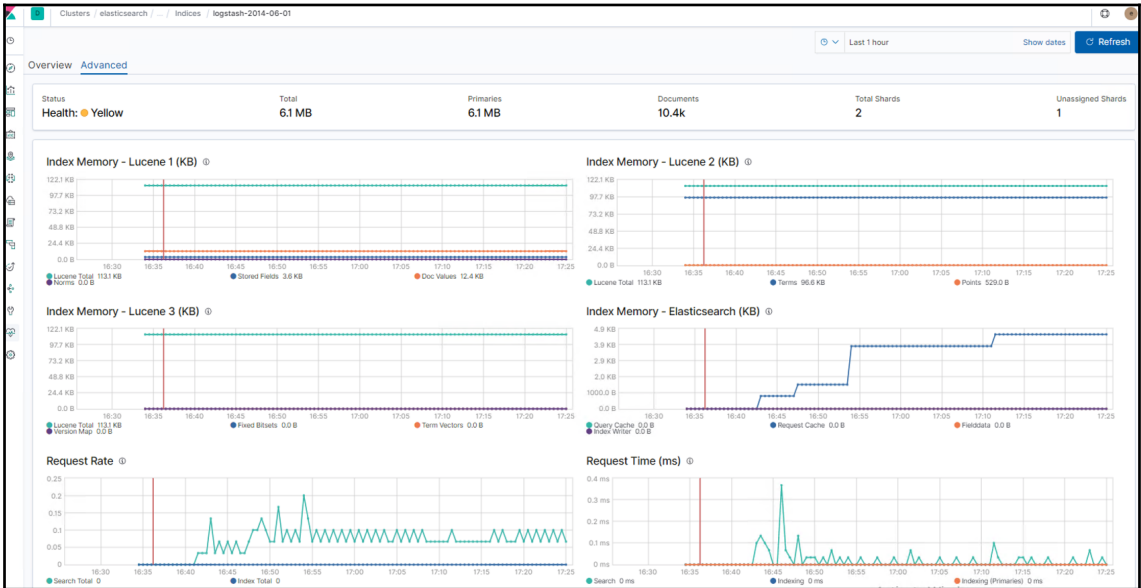
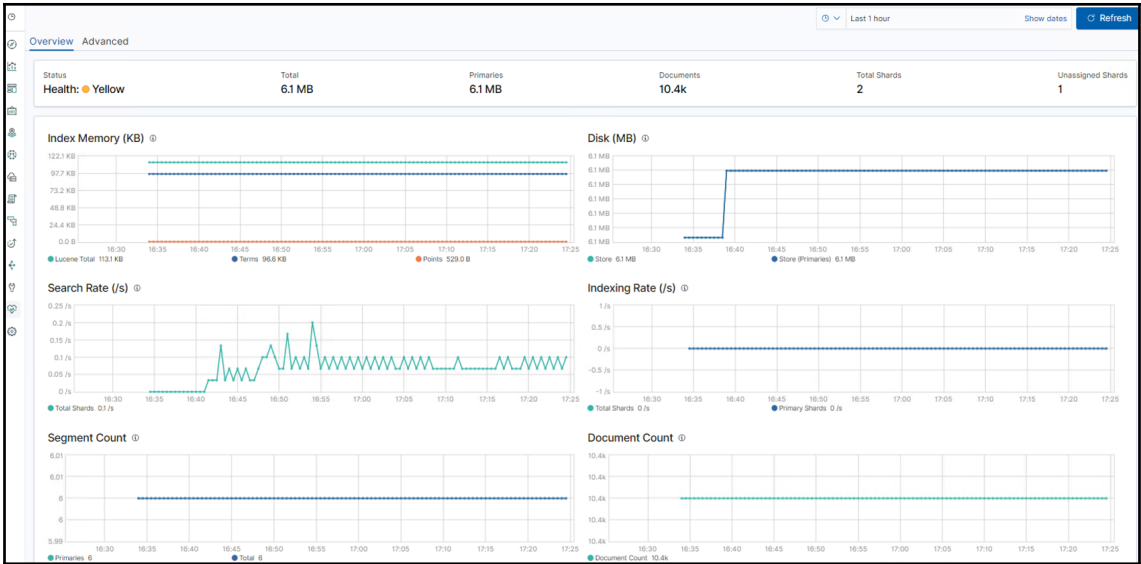
Status: ● Yellow Nodes: 1 Indices: 42 Memory: 471.6 MB / 989.9 MB Total Shards: 75 Unassigned Shards: 33 Documents: 305,452 Date: 177.0 MB

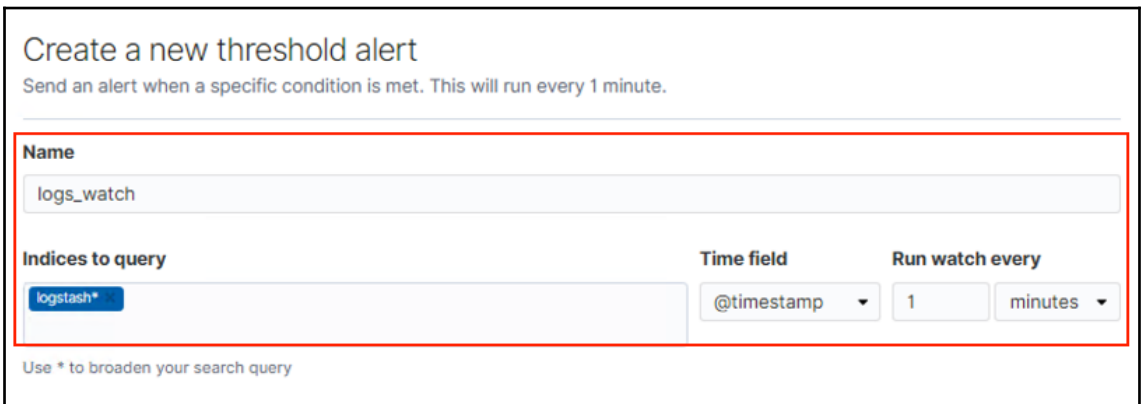
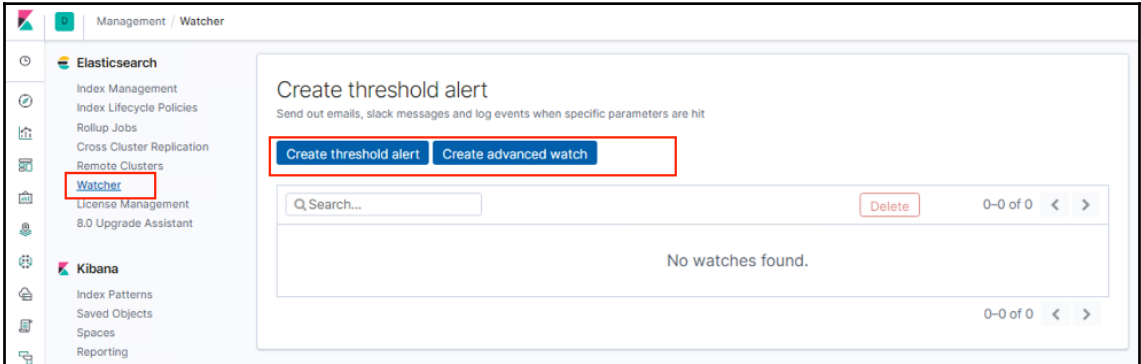
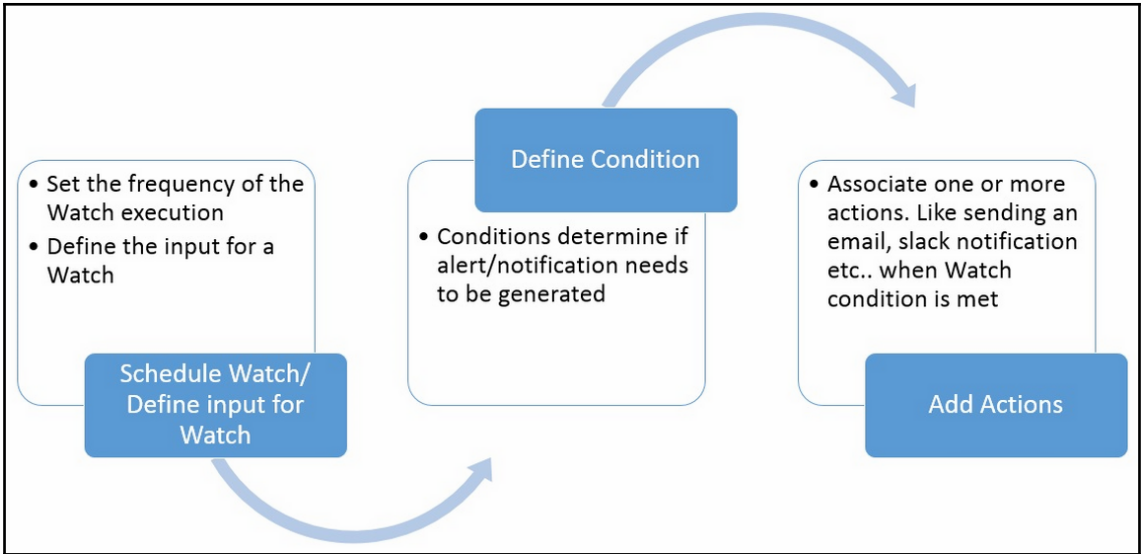
System indices

Filter indices...

Name	Status	Document Count	Data	Index Rate	Search Rate	Unassigned Shards
department	<span style="color: orange;">●</span> Yellow	3	4.0 KB	0/s	0/s	1
employee	<span style="color: orange;">●</span> Yellow	3	7.5 KB	0/s	0/s	1
logstash-2014-05-28	<span style="color: orange;">●</span> Yellow	1k	1.1 MB	0/s	0/s	1
logstash-2014-05-29	<span style="color: orange;">●</span> Yellow	7.9k	5.1 MB	0/s	0/s	1
logstash-2014-05-30	<span style="color: orange;">●</span> Yellow	9.1k	5.6 MB	0.72/s	0/s	1
logstash-2014-05-31	<span style="color: orange;">●</span> Yellow	9.4k	5.9 MB	0/s	0/s	1
logstash-2014-06-01	<span style="color: orange;">●</span> Yellow	10.4k	6.1 MB	0/s	0/s	1
logstash-2014-06-02	<span style="color: orange;">●</span> Yellow	11k	6.3 MB	0/s	0/s	1
logstash-2014-06-03	<span style="color: orange;">●</span> Yellow	7.6k	4.4 MB	0.98/s	0/s	1
logstash-2014-06-04	<span style="color: orange;">●</span> Yellow	6.9k	4.1 MB	0/s	0/s	1
logstash-2014-06-05	<span style="color: orange;">●</span> Yellow	8.5k	5.2 MB	0/s	0/s	1
logstash-2014-06-06	<span style="color: orange;">●</span> Yellow	11.2k	6.6 MB	0/s	0/s	1
logstash-2014-06-07	<span style="color: orange;">●</span> Yellow	10.7k	6.0 MB	1.17/s	0/s	1
logstash-2014-06-08	<span style="color: orange;">●</span> Yellow	13.1k	6.9 MB	0/s	0/s	1
logstash-2014-06-09	<span style="color: orange;">●</span> Yellow	10.7k	5.9 MB	0/s	0/s	1
logstash-2014-06-10	<span style="color: orange;">●</span> Yellow	10.8k	6.6 MB	0.8/s	0/s	1
logstash-2014-06-11	<span style="color: orange;">●</span> Yellow	8.1k	4.4 MB	0/s	0/s	1
logstash-2014-06-12	<span style="color: orange;">●</span> Yellow	10.1k	5.9 MB	0/s	0/s	1
logstash-2014-06-13	<span style="color: orange;">●</span> Yellow	10.9k	6.4 MB	0/s	0/s	1
logstash-2014-06-14	<span style="color: orange;">●</span> Yellow	10.4k	5.9 MB	0.99/s	0/s	1


Rows per page: 20





### Matching the following condition

WHEN count() OVER all documents IS ABOVE 100000 FOR THE LAST 5 minutes



Will perform 1 action once met

Add new action

Logging

Log text

Number of logs : {{{ctx.metadata.name}}} has exceeded the threshold

Remove Logging Action

Log a sample message now

# New watch

Save Cancel

Edit Simulate

ID  
errored\_logs\_watch

Name  
errored\_logs\_watch

## Watch JSON (Syntax)

```
8+   "search": {
9+     "request": {
10+       "body": {
11+         "size": 0,
12+         "query": {
13+           "match": {"message": "error"}
14+         },
15+       },
16+       "indices": [
17+         | "logs"
18+       ]
19+     }
20+   },
21+   "condition": {
22+     "compare": {
23+       "ctx.payload.hits.total": {
24+         "gte": 10
25+       }
26+     },
27+   },
28+   "actions": {
29+     "my-logging-action": {
30+       "logging": {
31+         "text": "There are {{ctx.payload.hits.total}} documents in your index. Threshold is 10."
32+       }
33+     }
34+   }
35+ }
```

```
[2019-05-23T17:43:07.809] INFO Ifo.e.x.v.a.l.ExecutableLoggingAction1 [MADSH01-APP01] The number of errors in logs is 39
[2019-05-23T17:43:37.379] INFO Ifo.e.x.v.a.l.ExecutableLoggingAction1 [MADSH01-APP01] The number of errors in logs is 39
[2019-05-23T17:44:07.343] INFO Ifo.e.x.v.a.l.ExecutableLoggingAction1 [MADSH01-APP01] The number of errors in logs is 39
[2019-05-23T17:44:37.380] INFO Ifo.e.x.v.a.l.ExecutableLoggingAction1 [MADSH01-APP01] The number of errors in logs is 39
[2019-05-23T17:45:07.373] INFO Ifo.e.x.v.a.l.ExecutableLoggingAction1 [MADSH01-APP01] The number of errors in logs is 39
[2019-05-23T17:45:37.293] INFO Ifo.e.x.v.a.l.ExecutableLoggingAction1 [MADSH01-APP01] The number of errors in logs is 39
[2019-05-23T17:46:04.394] INFO Ifo.e.x.v.a.l.ExecutableLoggingAction1 [MADSH01-APP01] There are 39 documents in your index. Threshold is 10.
[2019-05-23T17:46:07.383] INFO Ifo.e.x.v.a.l.ExecutableLoggingAction1 [MADSH01-APP01] The number of errors in logs is 39
[2019-05-23T17:46:37.428] INFO Ifo.e.x.v.a.l.ExecutableLoggingAction1 [MADSH01-APP01] The number of errors in logs is 39
```

## Create threshold alert

Send out emails, slack messages and log events when specific parameters are hit

Create threshold alert

Create advanced watch

Search...

Delete

1-10 of 10 < >

<input type="checkbox"/>	ID ↑	Name	State	Comment	Last Fired	Last Triggered	
<input type="checkbox"/>	252f1c48-98f...	logs_watch	✓ OK				<a href="#">Edit</a>
<input type="checkbox"/>	988aff02-dc8...	logs_error_watch	✓ OK		21 minutes ago	a few seconds...	<a href="#">Edit</a>
<input checked="" type="checkbox"/>	errored_logs_...	errored_logs_watch	✓ OK				<a href="#">Edit</a>
	I2RVLSk2Rr6I...	X-Pack Monitoring: Cluster Status...	🔥 Firing		a minute ago	a minute ago	
	I2RVLSk2Rr6I...	X-Pack Monitoring: Nodes Chang...	✓ OK			a minute ago	
	I2RVLSk2Rr6I...	X-Pack Monitoring: Elasticsearch ...	✓ OK			a minute ago	
	I2RVLSk2Rr6I...	X-Pack Monitoring: Kibana Versio...	✓ OK			a minute ago	
	I2RVLSk2Rr6I...	X-Pack Monitoring: Logstash Vers...	✓ OK			a minute ago	
	I2RVLSk2Rr6I...	X-Pack Monitoring: License Expira...	✓ OK			a minute ago	
<input type="checkbox"/>	logstash_error...		🔥 Firing		a few seconds...	a few seconds...	<a href="#">Edit</a>

1 watch selected

1-10 of 10 < >



## Current Status

[Deactivate](#)[Delete](#)

Action ↑	State
logging_1	✓ OK

## Watch History

Last 1 hour ▾

1-20 of 66



Trigger Time ↓	State	Comment
May 23, 2019 @ 17:19:54.740	✓ OK	
May 23, 2019 @ 17:19:24.729	✓ OK	
May 23, 2019 @ 17:18:54.713	✓ OK	
May 23, 2019 @ 17:18:24.704	✓ OK	
May 23, 2019 @ 17:17:54.694	✓ OK	
May 23, 2019 @ 17:17:24.678	✓ OK	
May 23, 2019 @ 17:16:54.670	✓ OK	
May 23, 2019 @ 17:16:24.659	✓ OK	
May 23, 2019 @ 17:15:54.651	✓ OK	
May 23, 2019 @ 17:15:24.638	✓ OK	
May 23, 2019 @ 17:14:54.630	✓ OK	
May 23, 2019 @ 17:14:24.619	✓ OK	
May 23, 2019 @ 17:13:54.611	✓ OK	
May 23, 2019 @ 17:13:24.602	✓ OK	

# Chapter 9: Running Elastic Stack in Production

pranav.shukla@gmail.com Trial started


Deployments  
Custom plugins  
Account  
Help


## Create deployment

- 1 Name your deployment**

Give your deployment a name
- 2 Select a cloud platform**

Pick your cloud and let us handle the rest. No additional accounts required.

  
Amazon Web Services

  
Google Cloud Platform
- 3 Select a region**
- 4 Set up your deployment**

Stable versions

7.1.0  7.0.1

6.8.0


5.6.16

Select a deployment to restore from its latest snapshot
- 5 Optimize your deployment**

**I/O Optimized**


**Recommended**

Use for search and general all-purpose workloads. Includes a balance of compute, memory, and storage. [Default specs](#)




**Compute Optimized**

Run CPU-intensive workloads or run smaller workloads cost-effectively when you need less memory and storage. [Default specs](#)



**Memory Optimized**


Perform memory-intensive operations efficiently, including workloads with frequent aggregations. [Default specs](#)



**UNAVAILABLE**

**Hot-Warm Architecture**

Use for time-series analytics and logging workloads that benefit from automatic index curation. [Default specs](#)



Elastic Cloud supports many more options to cater to your specific use case such as hot-warm architecture optimized for logging, compute-focused setup optimized for analytics etc. [Learn more ...](#)

**Deployment pricing**

Free! As part of your 14-day trial, you can try it out without a credit card. If you want to unleash the full power of Elasticsearch Service now, you can enter your credit card details or contact [sales@elastic.co](mailto:sales@elastic.co).

pranav.shukla@gmail.com Trial started

test-cluster

# Activity

aws US East (N. Virginia)

Elasticsearch Kibana APM

**Generated user**

You can use the credentials below to login to Elasticsearch or Kibana. Make sure to save the password somewhere as this is the only time we can show it to you.

**Username** elastic

**Password** [REDACTED]

**Cloud ID** test- [REDACTED]

**APM Server secret token** [REDACTED]

Get started with Beats and Logstash quickly. The Cloud ID simplifies sending data to your cluster on Elastic Cloud. [Learn more ...](#)


## Elasticsearch change history

✔ #0 — Applied 4 minutes ago

Show details Reapply

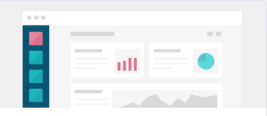
https://0419ac85cf7499d826ba7be273e9399.us-east-1.aws.found.io:9243/app/kibana#/home?\_g=0

Apps Google Elasticsearch Addons Webservices - RE... Hadoop Spark IoT Startup Scala Big Data Utilities Log Analytics DW & Big Data Dist Computing A... Machine Learning Cisco Meraki Clou...



## Welcome to Kibana

Your window into the Elastic Stack



**Let's get started**

We noticed that you don't have any data in your cluster. You can try our sample data and dashboards or jump in with your own data.

[Try our sample data](#) [Explore on my own](#)



### Deployments

test-cluster

Edit

Elasticsearch

Logs

**Snapshots**

API Console

Kibana

APM

Activity

Security

Performance

Custom plugins

Account

Help

test-cluster / Elasticsearch

## Snapshots

aws US East (N. Virginia)

Snapshots are backups of your data that you can restore in the event of an unexpected data loss.

Last successful snapshot **19 minutes ago** Next snapshot **In 11 minutes** Snapshot frequency **30 minutes**

The current settings retain 100 snapshots over a period of 2 days. [Edit settings](#)

[Take snapshot now](#) [Restore from another deployment](#)

Choose an snapshot in another deployment to restore on this deployment

## Snapshots

All snapshots 29 Success 29

Snapshot	Status	Completed	Duration	
scheduled-1558940570-instance-000000001	● Success	19 minutes ago	A few seconds	<a href="#">Restore</a>
scheduled-1558938768-instance-000000001	● Success	An hour ago	A few seconds	<a href="#">Restore</a>
scheduled-1558936966-instance-000000001	● Success	An hour ago	A few seconds	<a href="#">Restore</a>
scheduled-1558935163-instance-000000001	● Success	2 hours ago	A few seconds	<a href="#">Restore</a>
scheduled-1558933361-instance-000000001	● Success	2 hours ago	A few seconds	<a href="#">Restore</a>
scheduled-1558931559-instance-000000001	● Success	3 hours ago	A few seconds	<a href="#">Restore</a>
scheduled-1558929757-instance-000000001	● Success	3 hours ago	A few seconds	<a href="#">Restore</a>
scheduled-1558927955-instance-000000001	● Success	4 hours ago	A few seconds	<a href="#">Restore</a>
scheduled-1558926153-instance-000000001	● Success	4 hours ago	A few seconds	<a href="#">Restore</a>
scheduled-1558924351-instance-000000001	● Success	5 hours ago	A few seconds	<a href="#">Restore</a>



### Deployments

- test-cluster
  - Edit
  - Elasticsearch
    - Logs
    - Snapshots
    - API Console
  - Kibana
    - APM
    - Activity
    - Security
    - Performance

### Custom plugins

### Account

### Help

test-cluster

# scheduled-1558940570- instance-0000000001

aws US East (N. Virginia)

Result	Started	Completed	Duration
Success	21 minutes ago	21 minutes ago	755ms

Total shards	Successful shards	Failed shards
4	4	0

#### UUID

rr0WY62QR7-aggKg1fpgXg

#### Version

7.1.0

#### Indices

- .kibana\_1
- .kibana\_task\_manager
- .security-7
- apm-7.1.0-onboarding-2019.05.26

## Restore Snapshot

### Specify Indices Optional

#### Indices

e.g. index1,index2

To restore all indices in the snapshot and to restore the cluster state, leave this field blank. To limit the indices that will be restored, you can use multi-index syntax. For example, `index1,index2,index3`.

### Rename Indices Optional

#### Match

index\_(.+)

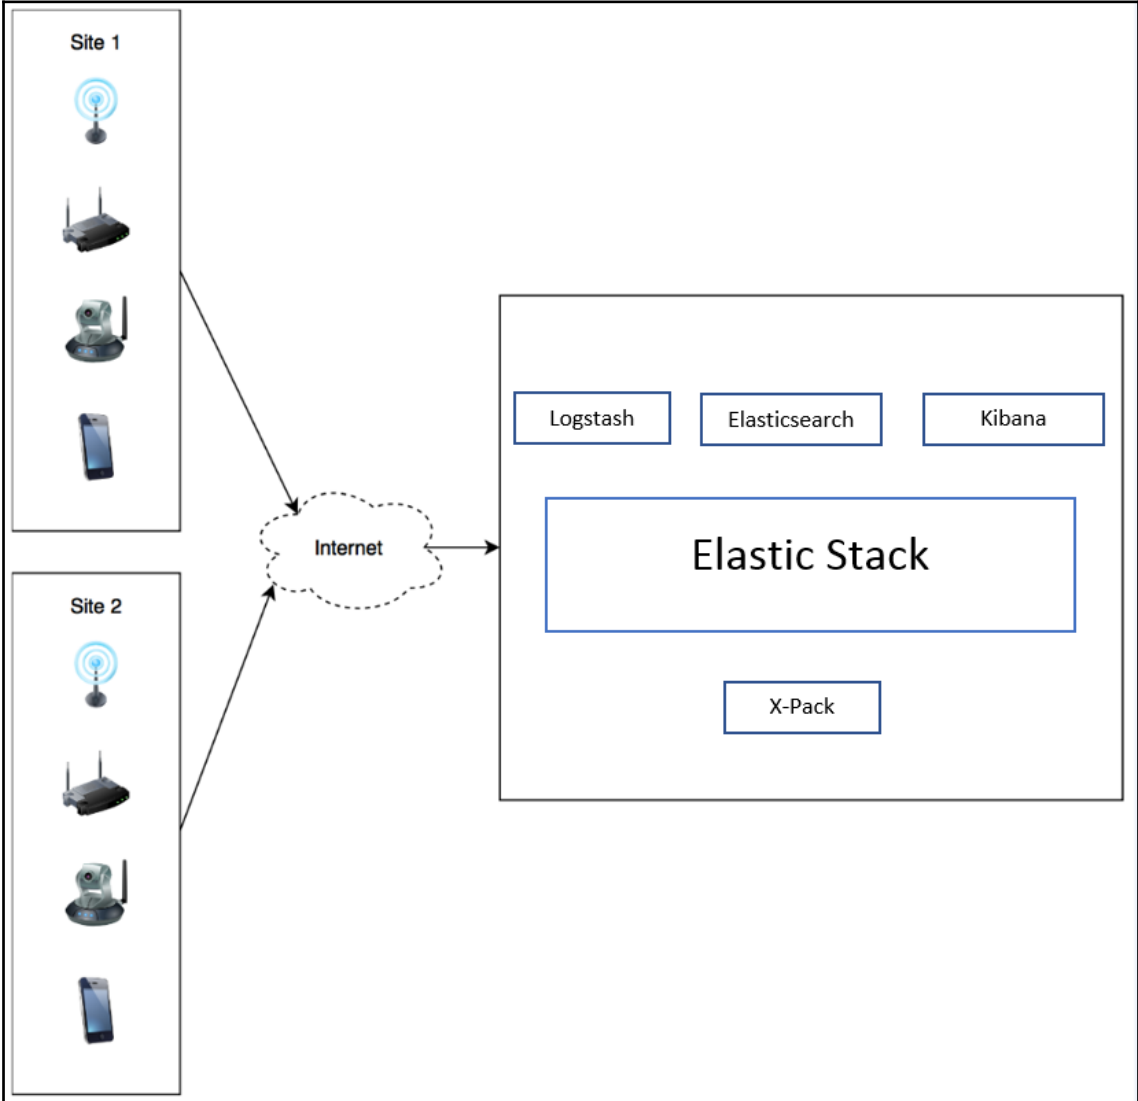
#### Replace with

restored\_index\_\$\$1

Rename indices by matching a pattern with a regular expression and replacing it with a captured expression. For example: To rename the index `index_logging` to `restored_index_logging` when it is restored, match the pattern `index_(.+)` and replace it with `restored_index_$$1`.

Restore snapshot

# Chapter 10: Building a Sensor Data Analytics Application




localhost:5601/app/kibana#/home?\_g=()

Home

## Add Data to Kibana


Use these solutions to quickly turn your data into pre-built dashboards and monitoring systems.



**APM**

APM automatically collects in-depth performance metrics and errors from inside your applications.


[Add APM](#)



**Logging**

Ingest logs from popular data sources and easily visualize in preconfigured dashboards.


[Add log data](#)



**Metrics**

Collect metrics from the operating system and services running on your servers.

[Add metric data](#)



**Security analytics**

Centralize security events for interactive investigation in ready-to-go visualizations.

[Add security events](#)

**Add sample data**

Load a data set and a Kibana dashboard

**Upload data from log file**

Import a CSV, NDJSON, or log file

**Use Elasticsearch data**

Connect to your Elasticsearch index

---

### Visualize and Explore Data

**APM**

Automatically collect in-depth performance metrics and errors from inside your applications.

**Dashboard**

Display and share a collection of visualizations and saved searches.

**Canvas**

Showcase your data in a pixel-perfect way.

**Discover**

Interactively explore your data by querying and filtering raw documents.

### Manage and Administer the Elastic Stack

**Console**

Skip cURL and use this JSON interface to work with your data directly.

**Monitoring**

Track the real-time health and performance of your Elastic Stack.

**Index Patterns**

Manage the index patterns that help retrieve your data from Elasticsearch.

**Rollups**

Summarize and store historical data in a smaller index for future analysis.

Management | Create index pattern

**Elasticsearch**

- Index Management
- Index Lifecycle Policies
- Rollup Jobs
- Cross Cluster Replication
- Remote Clusters
- License Management
- 8.0 Upgrade Assistant

**Kibana**

- Index Patterns**
- Saved Objects
- Spaces
- Reporting
- Advanced Settings

**Create index pattern**

No default Index pattern. You must select or create one to continue.

[Create index pattern](#)

## Create index pattern

Kibana uses index patterns to retrieve data from Elasticsearch indices for things like visualizations.  Include system indices

**Step 1 of 2: Define index pattern**

Index pattern

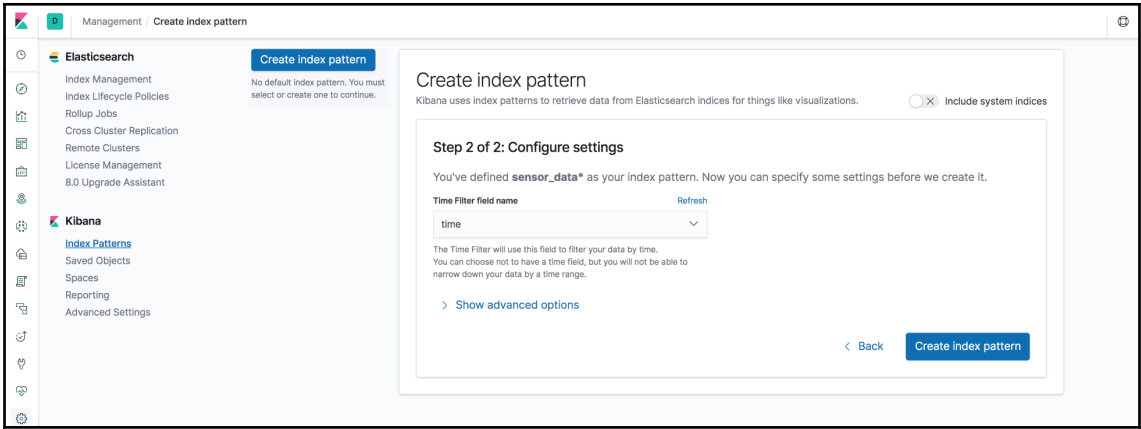
You can use a \* as a wildcard in your index pattern. You can't use spaces or the characters |, /, ? , \* , < , > , |.

[Next step](#)

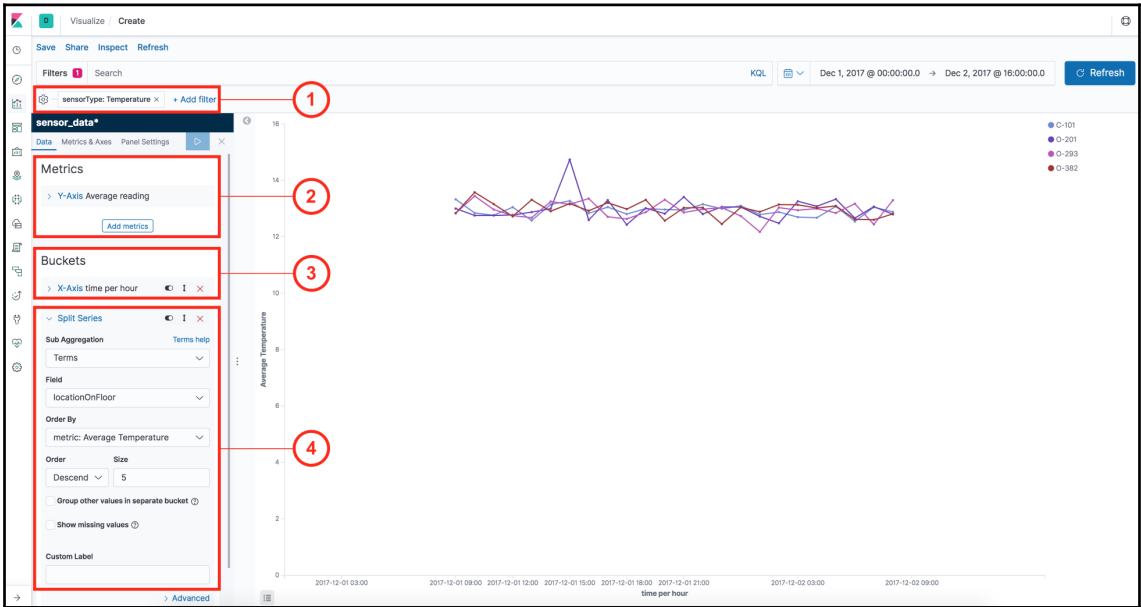
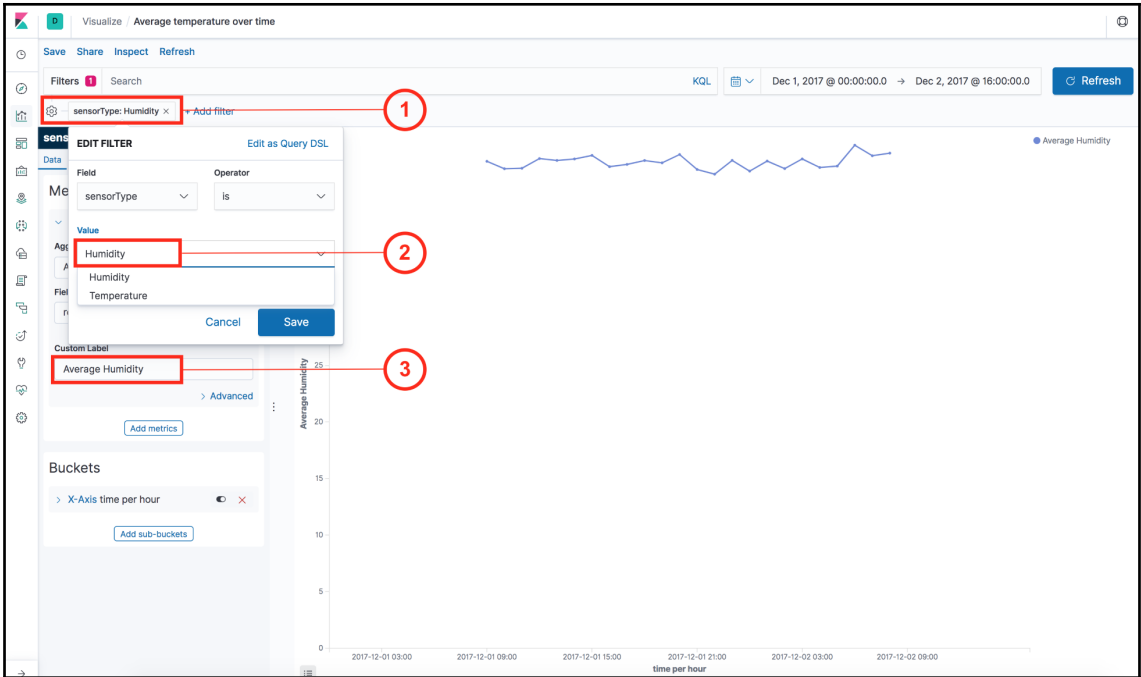
✓ **Success!** Your index pattern matches **1** index.

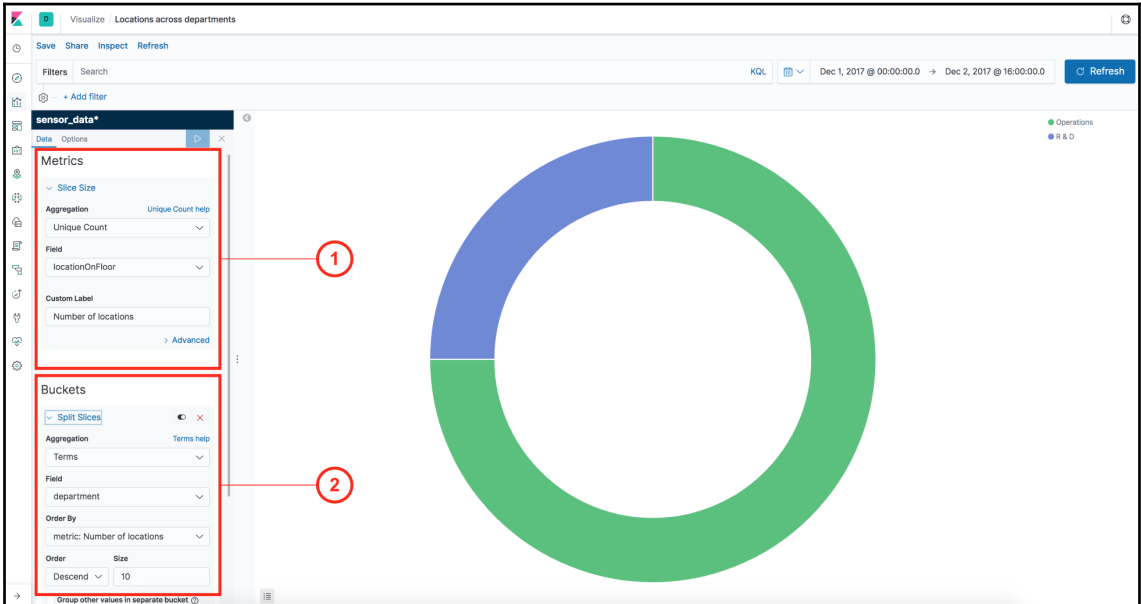
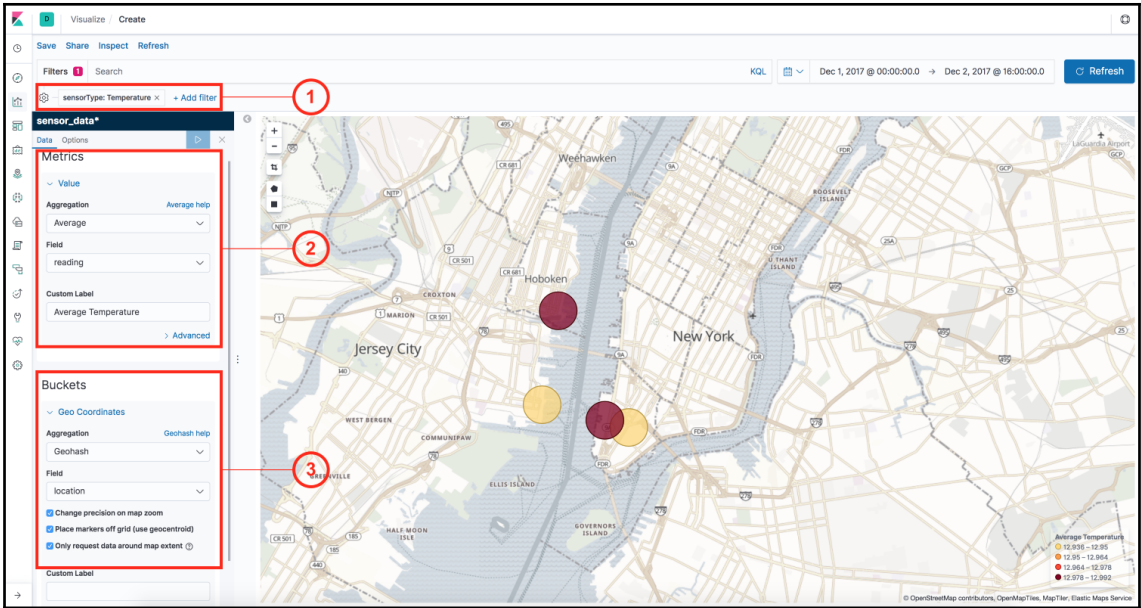
**sensor\_data-2019.05.26**

Rows per page: 10

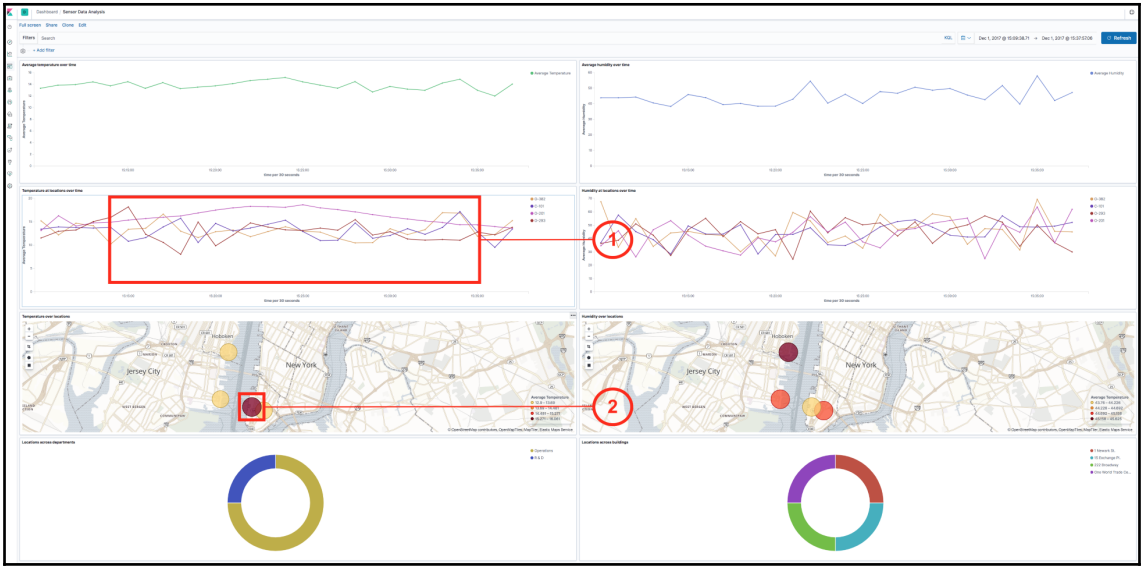












# Chapter 11: Monitoring Server Infrastructure

← → ↻ <https://www.elastic.co/downloads/beats/metricbeat-oss>

elastic [Products](#) [Cloud](#) [Services](#) [Customers](#) [Learn](#) [downloads](#) [contact](#)  [EN](#)

Downloads

## Download Metricbeat - OSS Only

Want to upgrade? We'll give you a hand. [Migration Guide »](#)

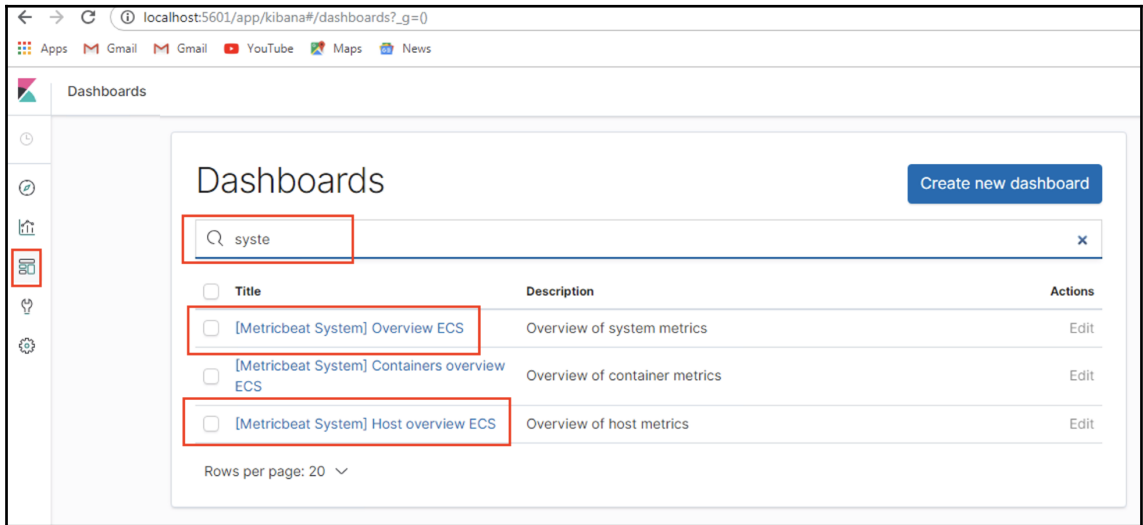
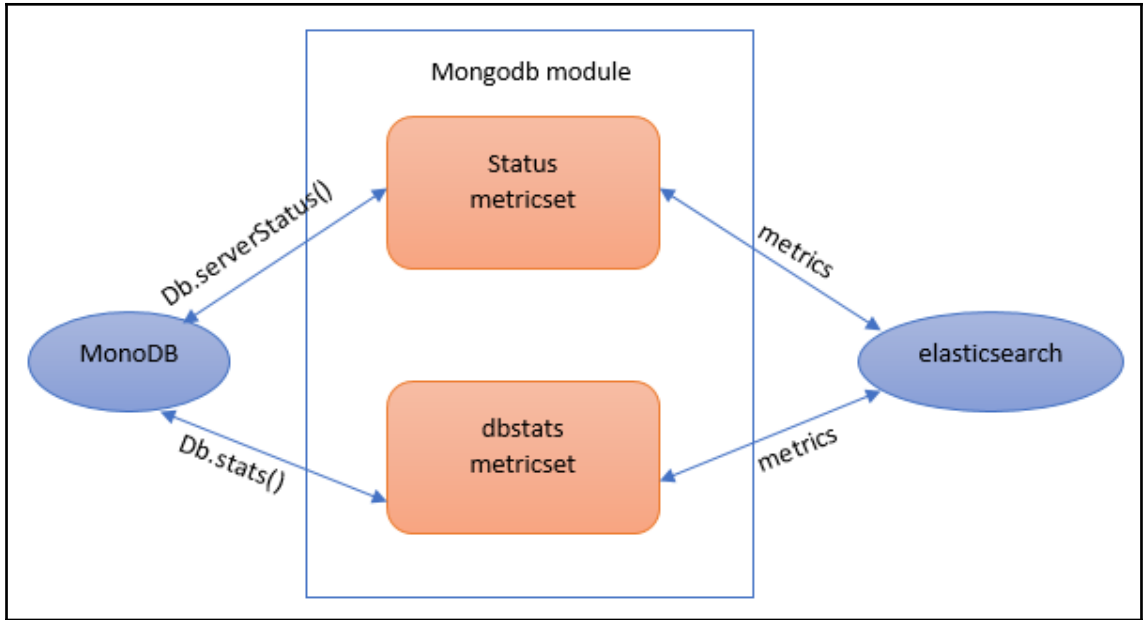
Version: 7.0.0  
Release date: April 10, 2019  
License: **Apache 2.0**

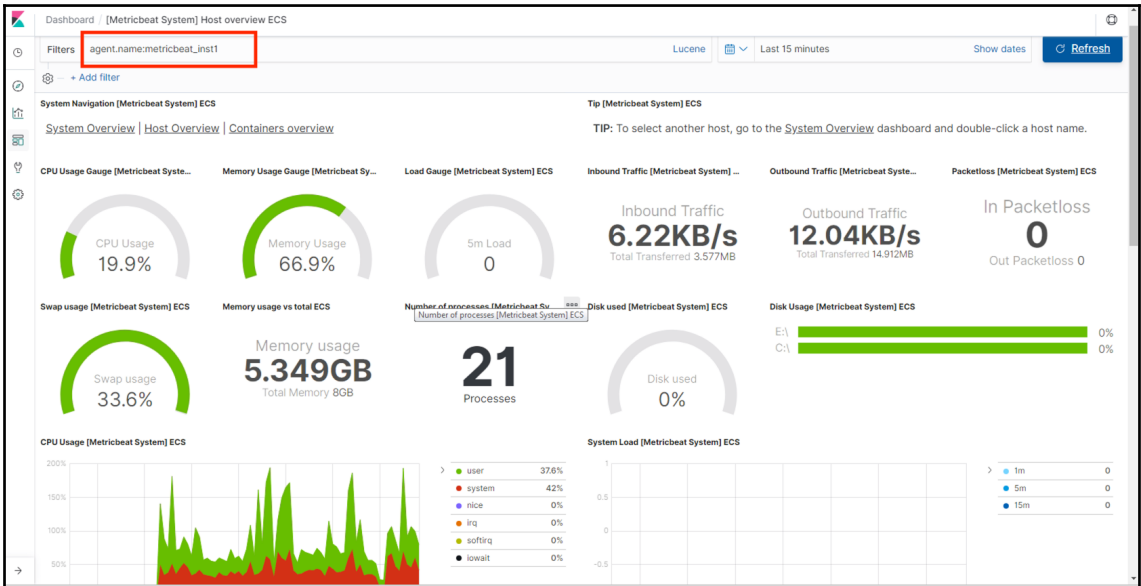
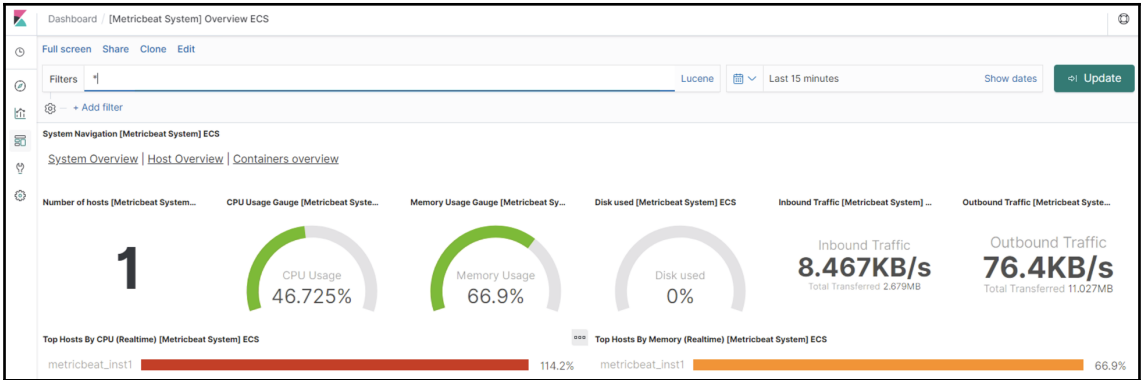
Downloads:

- ↕ [DEB 32-BIT sha](#)
- ↕ [RPM 32-BIT sha](#)
- ↕ [LINUX 32-BIT sha](#)
- ↕ [MAC sha](#)
- ↕ [WINDOWS 64-BIT sha](#)
- ↕ [DEB 64-BIT sha](#)
- ↕ [RPM 64-BIT sha](#)
- ↕ [LINUX 64-BIT sha](#)
- ↕ [WINDOWS 32-BIT sha](#)

Containers: Run with [Docker](#) Run with [Kubernetes](#)

Notes: This distribution only includes features licensed under the Apache 2.0 license. To get access to full [set of free features](#), use











~ 15 minutes ago → now

Quick select



Last



15

minutes



Apply

Commonly used

Today

This week

This month

This year

Today so far

Week to date

Month to date

Year to date

Recently used date ranges

Last 15 minutes

Refresh every

5



seconds



▶ Start

