Introduction to Elasticsearch and Kibana geospatial capabilities

Jorge Sanz , Kibana, Elastic

xurxosanz jsanz ✉ jorge.sanz@elastic.co

2022-08-22

https://ela.st/foss4g22-workshop

elastic

Training objectives

- ✓ To be able to **run** the Elastic stack
- ✓ What is **Kibana** and its different components
- ✓ How to use **Elastic Maps** to visualize geospatial data
- ✓ Understand Elasticsearch **geospatial** capabilities
- ✓ How the Elastic stack fits in a geospatial application **architecture**

elastic

Agenda

1. Introduction to the Elastic stack
2. Lab setup
3. Kibana introduction
4. Elastics Maps
5. Elasticsearch Geo
6. Web mapping and Elasticsearch

The Elastic Search Platform

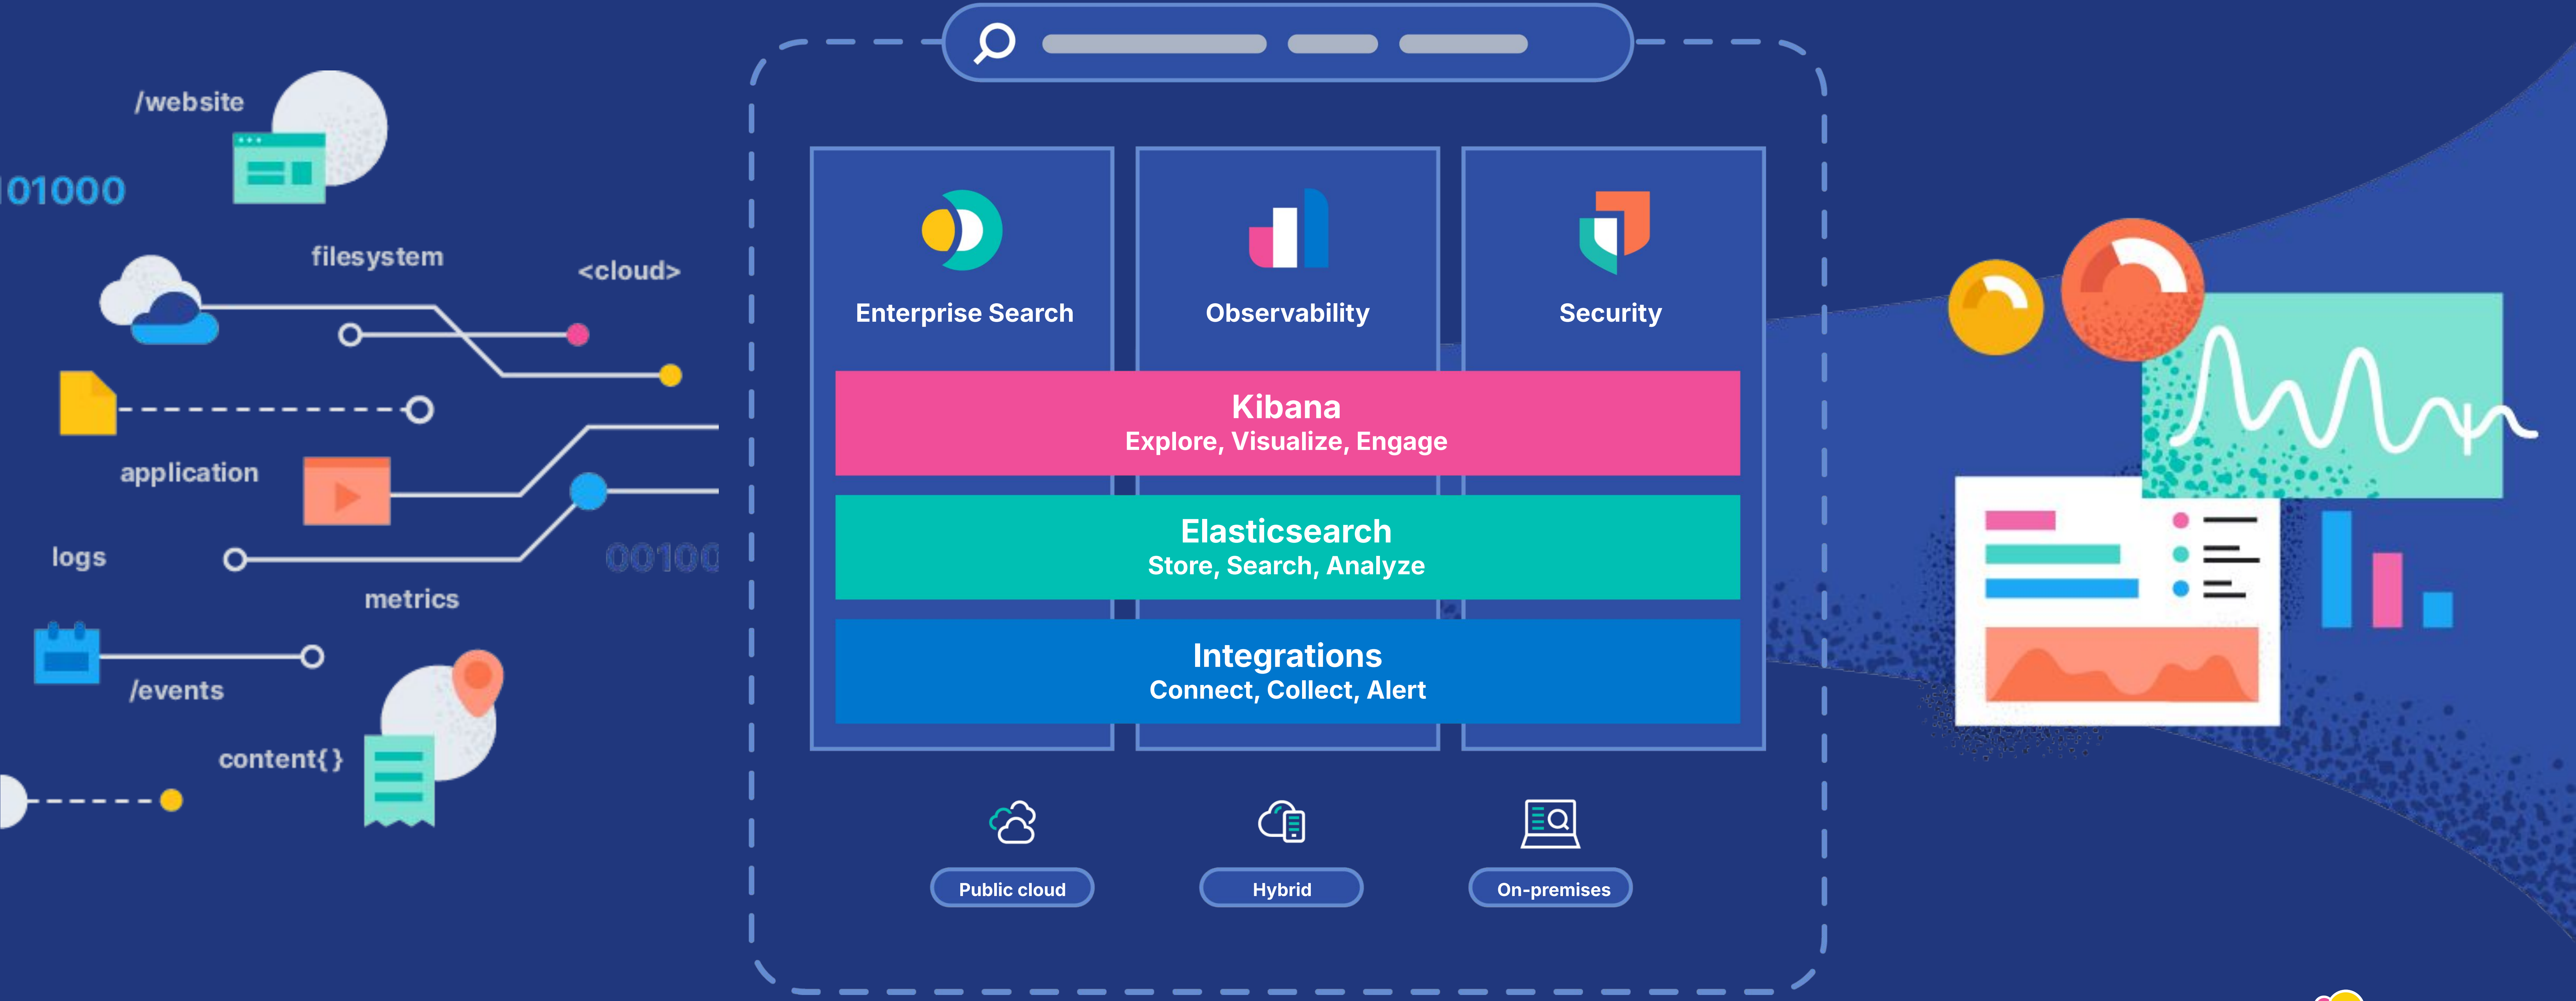

elastic

Communit y

https://github.com/elastic

https://ela.st/slack

https://discuss.elastic.co

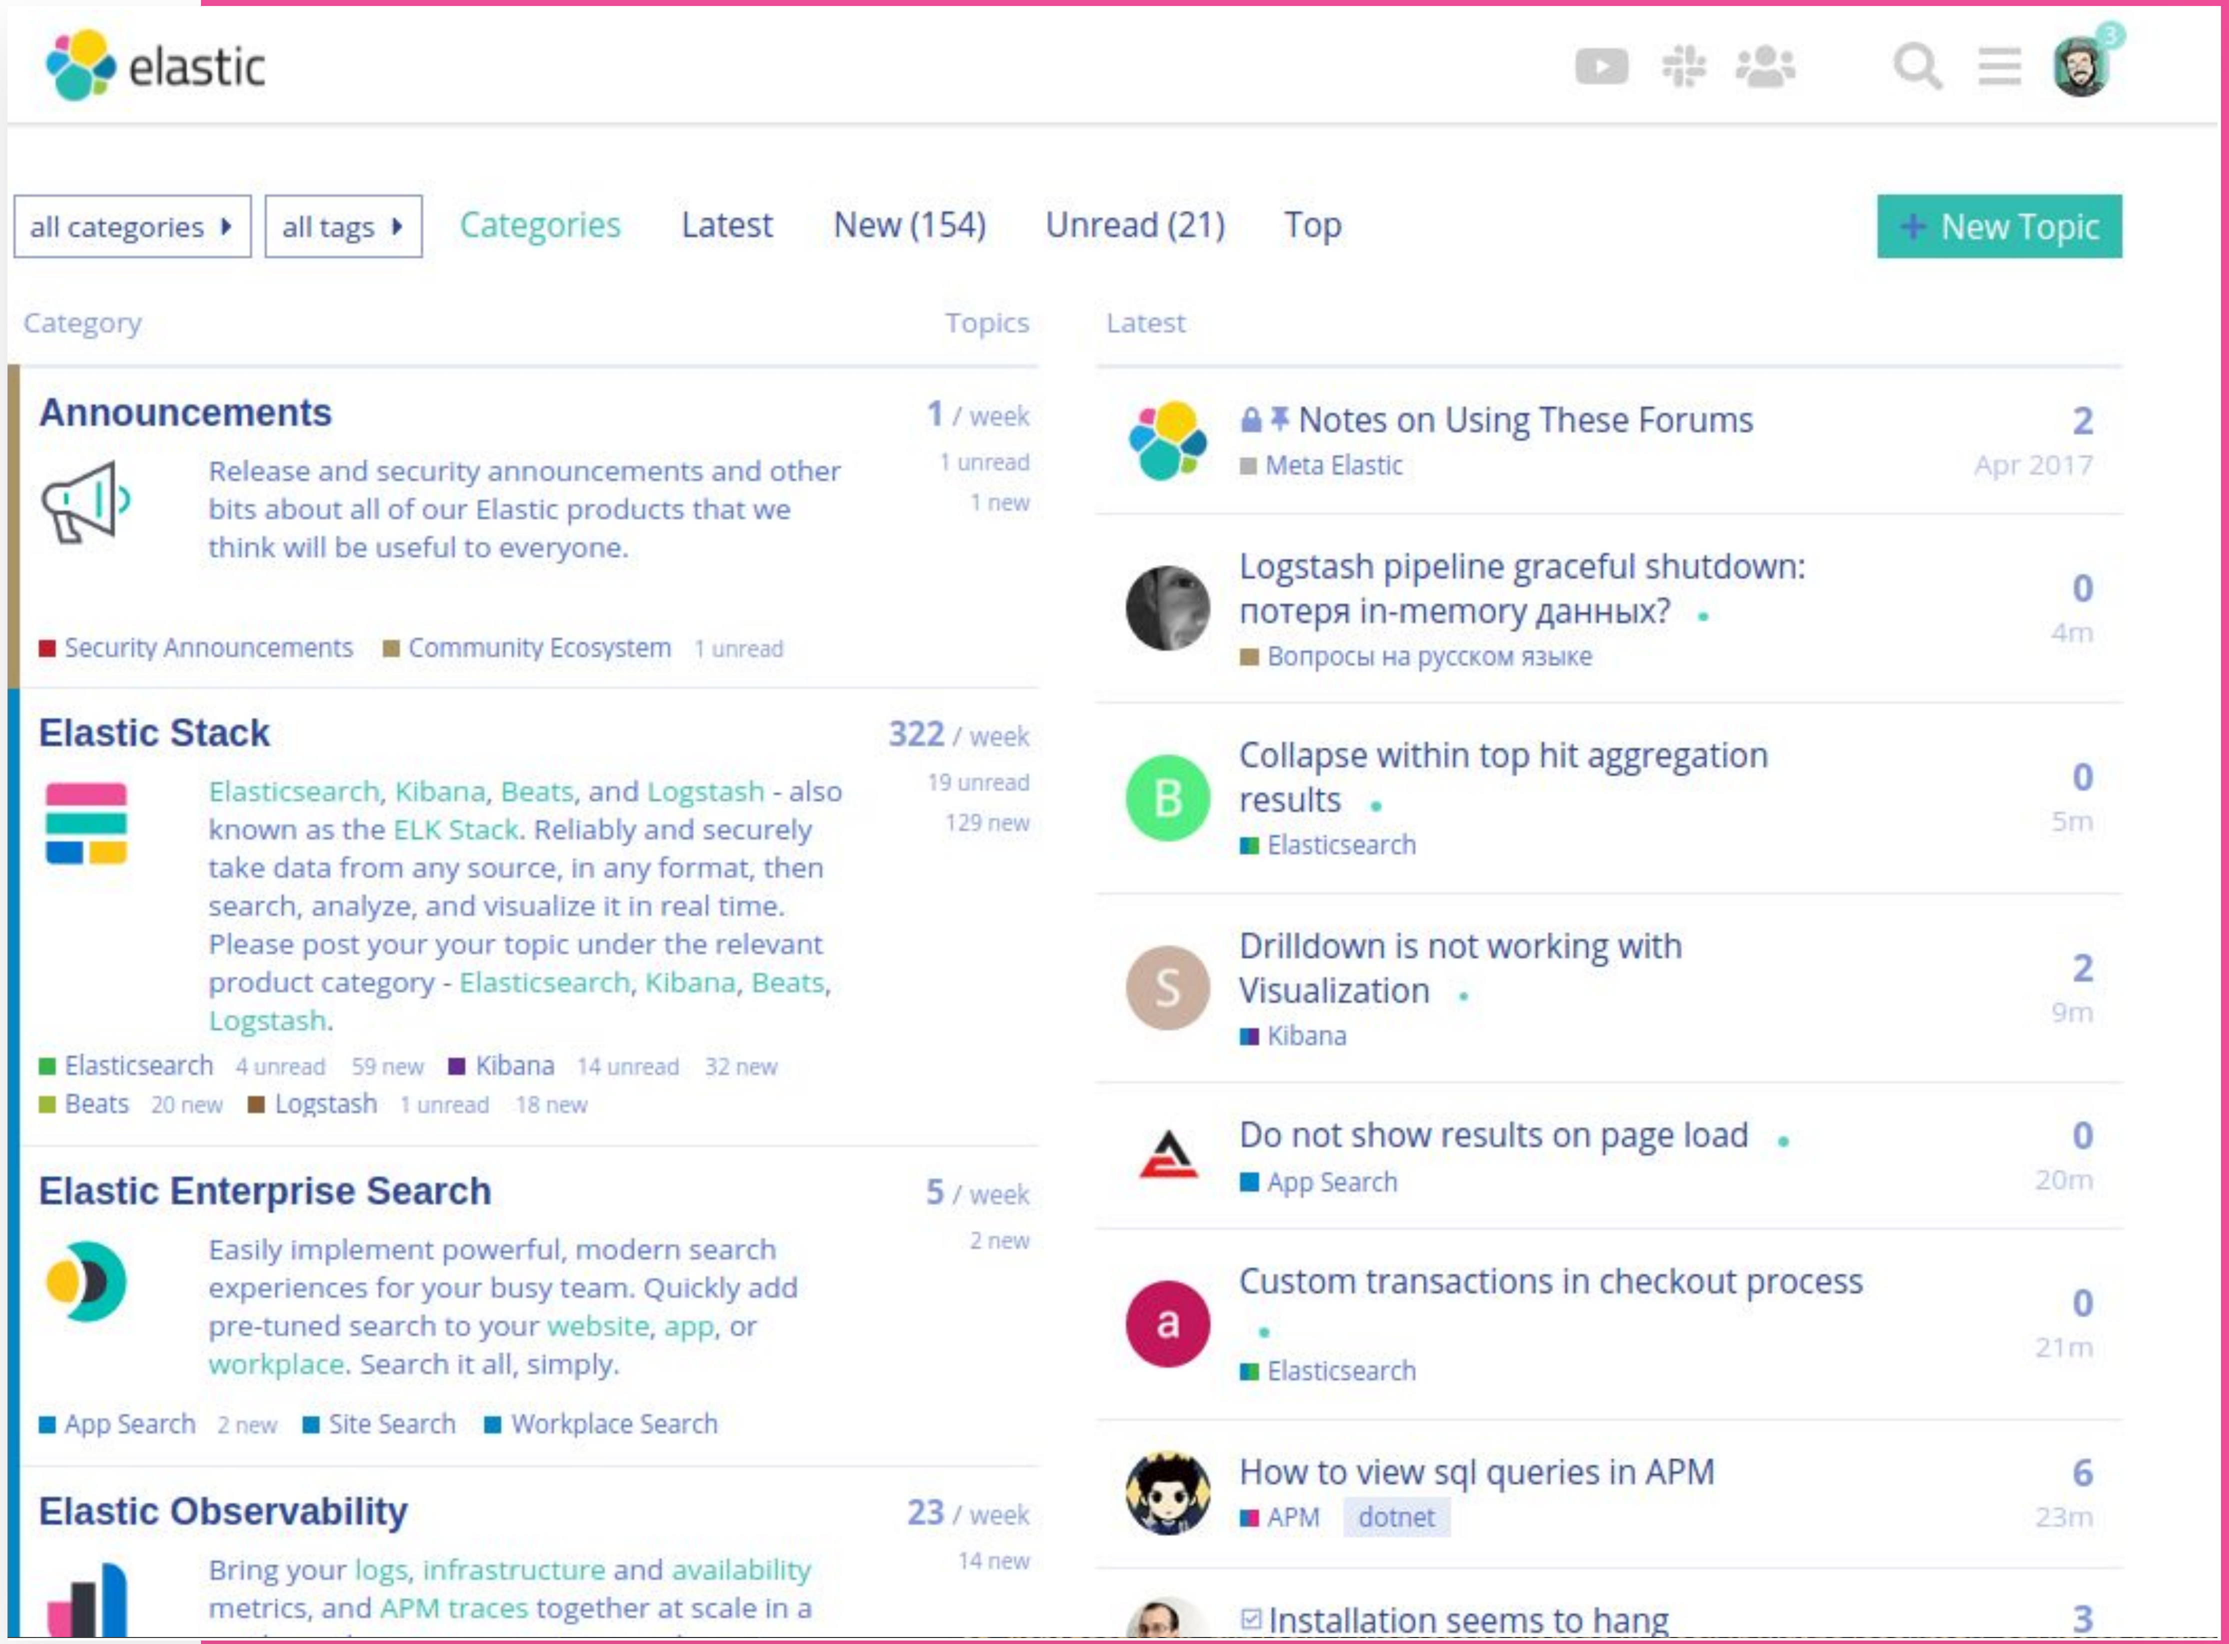

Category	Topics
Announcements Release and security announcements and other bits about all of our Elastic products that we think will be useful to everyone. Security Announcements · Community Ecosystem 1 unread	1 / week 1 unread 1 new
Elastic Stack Elasticsearch, Kibana, Beats, and Logstash - also known as the ELK Stack. Reliably and securely take data from any source, in any format, then search, analyze, and visualize it in real time. Please post your your topic under the relevant product category - Elasticsearch, Kibana, Beats, Logstash. Elasticsearch 4 unread 59 new · Kibana 14 unread 32 new · Beats 20 new · Logstash 1 unread 18 new	322 / week 19 unread 129 new
Elastic Enterprise Search Easily implement powerful, modern search experiences for your busy team. Quickly add pre-tuned search to your website, app, or workplace. Search it all, simply. App Search 2 new · Site Search · Workplace Search	5 / week 2 new
Elastic Observability Bring your logs, infrastructure and availability metrics, and APM traces together at scale in a	23 / week 14 new

Latest		
Notes on Using These Forums Meta Elastic	2	Apr 2017
Logstash pipeline graceful shutdown: потеря in-memory данных? Вопросы на русском языке	0	4m
Collapse within top hit aggregation results Elasticsearch	0	5m
Drilldown is not working with Visualization Kibana	2	9m
Do not show results on page load App Search	0	20m
Custom transactions in checkout process Elasticsearch	0	21m
How to view sql queries in APM APM dotnet	6	23m
☑ Installation seems to hang	3	

Elasticsearch

All data is welcome

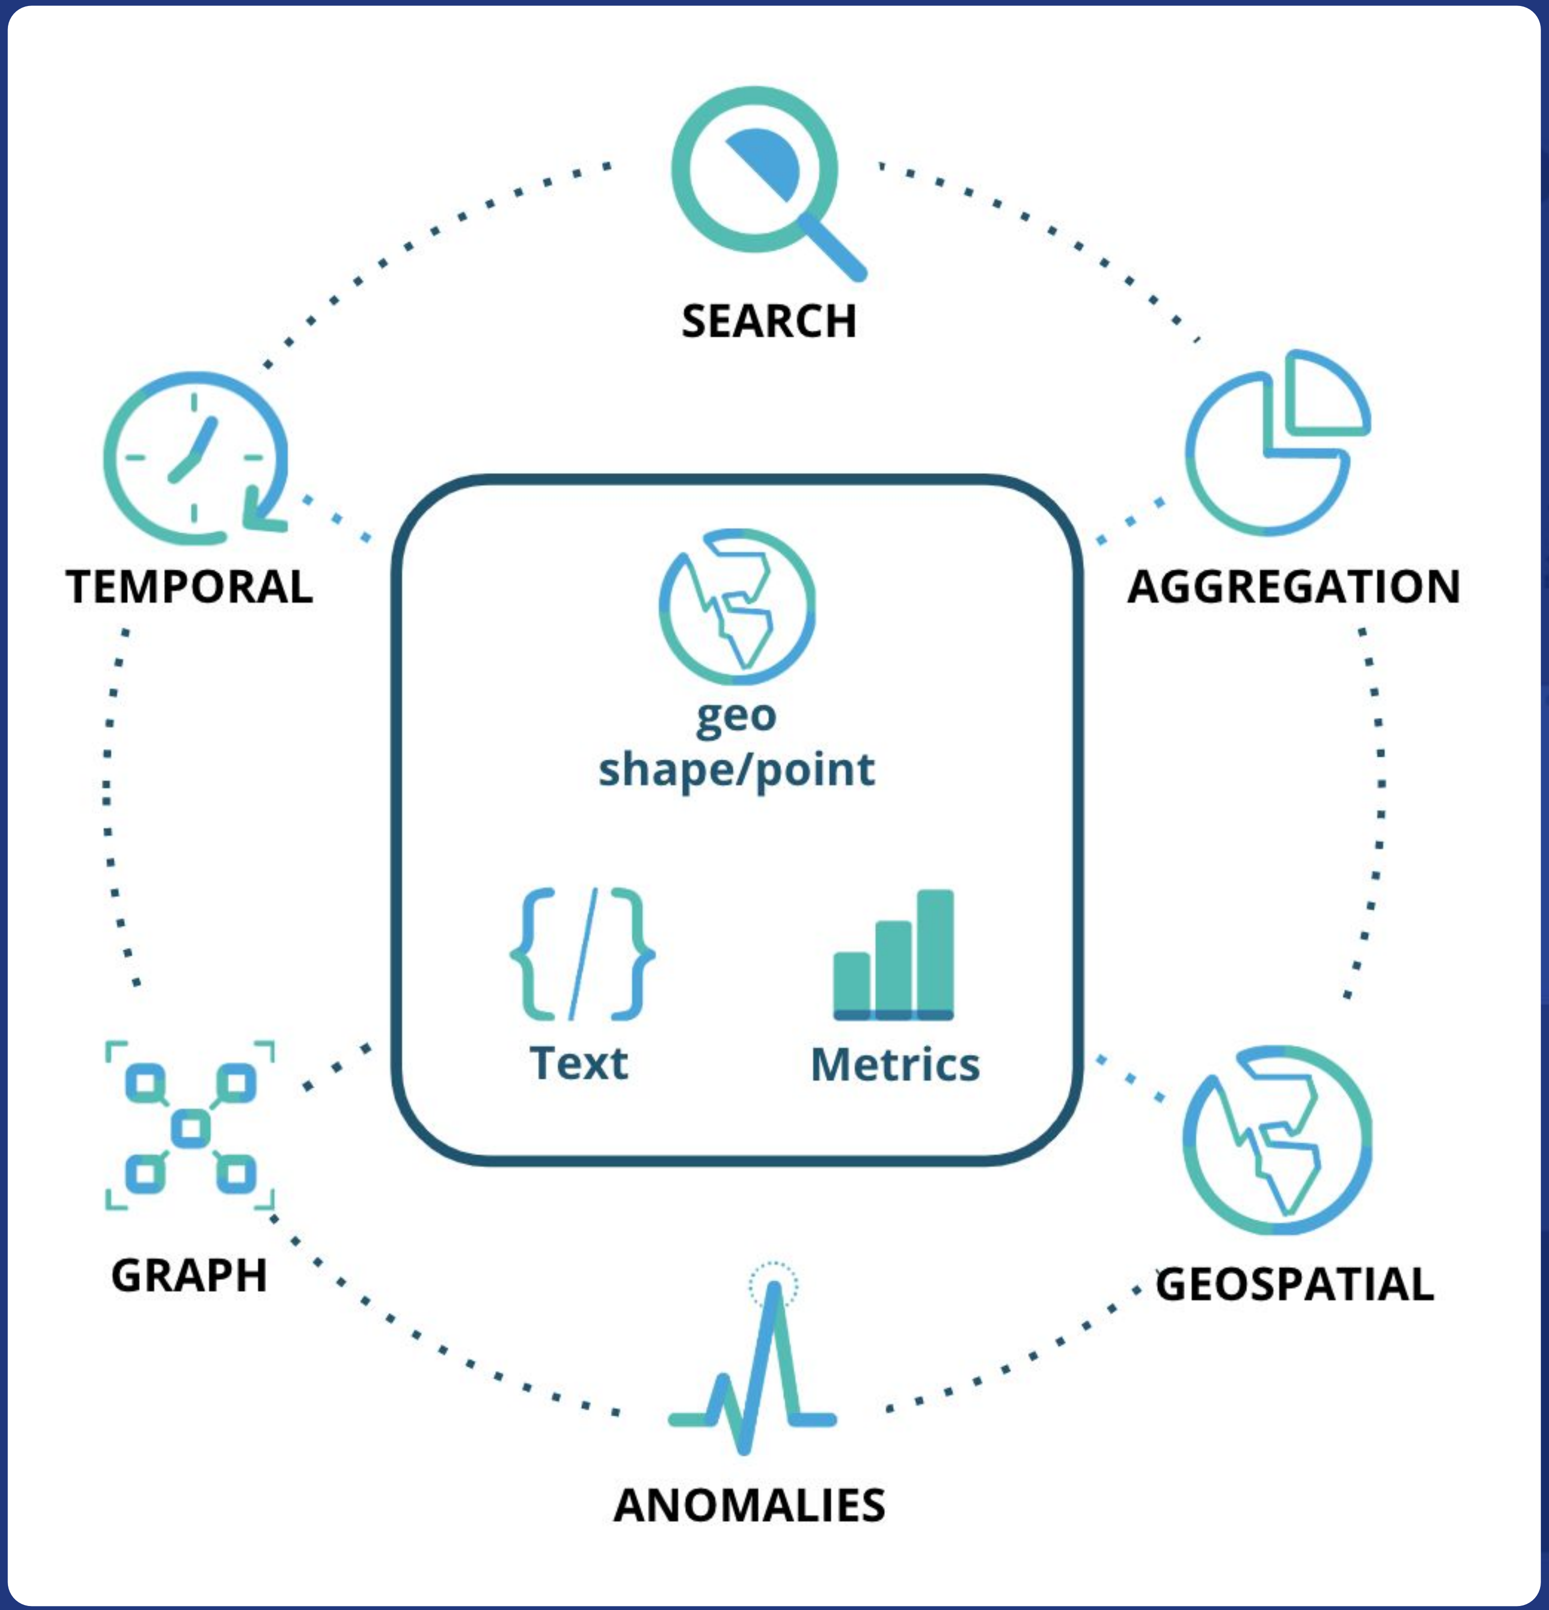

Elasticsearch components

- Cluster
- Node
- Shard
- Index
- Mapping
- Document
- Field

Communicating with Elasticsearch

- All communication through HTTP endpoints 📖

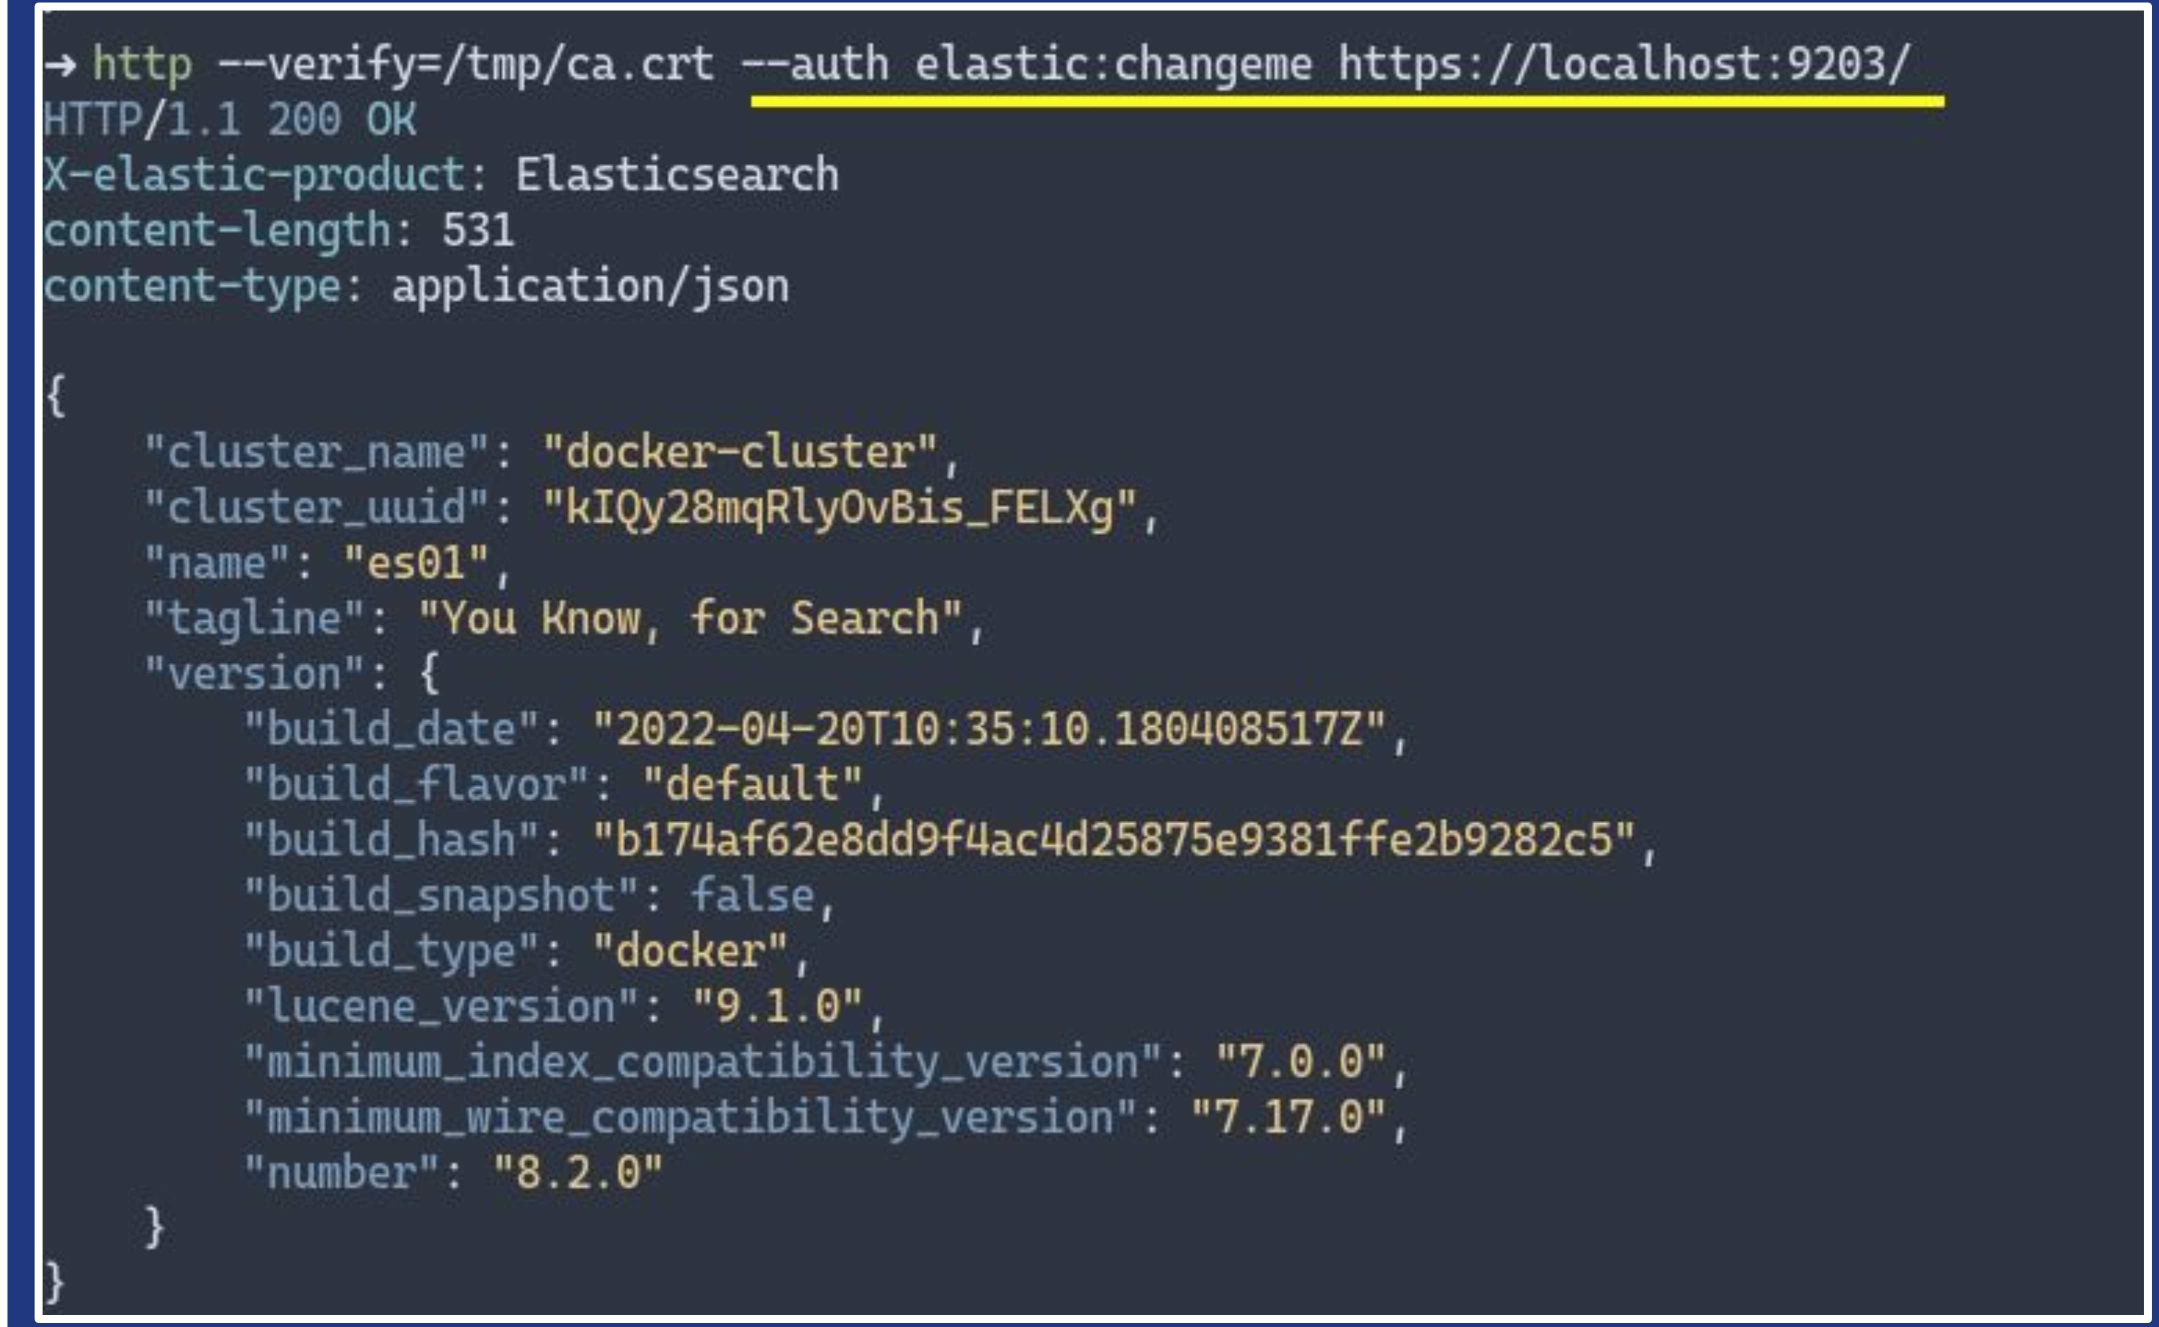

```
→ http --verify=/tmp/ca.crt --auth elastic:changeme https://localhost:9203/
HTTP/1.1 200 OK
X-elastic-product: Elasticsearch
content-length: 531
content-type: application/json

{
    "cluster_name": "docker-cluster",
    "cluster_uuid": "kIQy28mqRlyOvBis_FELXg",
    "name": "es01",
    "tagline": "You Know, for Search",
    "version": {
        "build_date": "2022-04-20T10:35:10.180408517Z",
        "build_flavor": "default",
        "build_hash": "b174af62e8dd9f4ac4d25875e9381ffe2b9282c5",
        "build_snapshot": false,
        "build_type": "docker",
        "lucene_version": "9.1.0",
        "minimum_index_compatibility_version": "7.0.0",
        "minimum_wire_compatibility_version": "7.17.0",
        "number": "8.2.0"
    }
}
```

Communicating with Elasticsearch

- All communication through HTTP endpoints 📖
- JSON

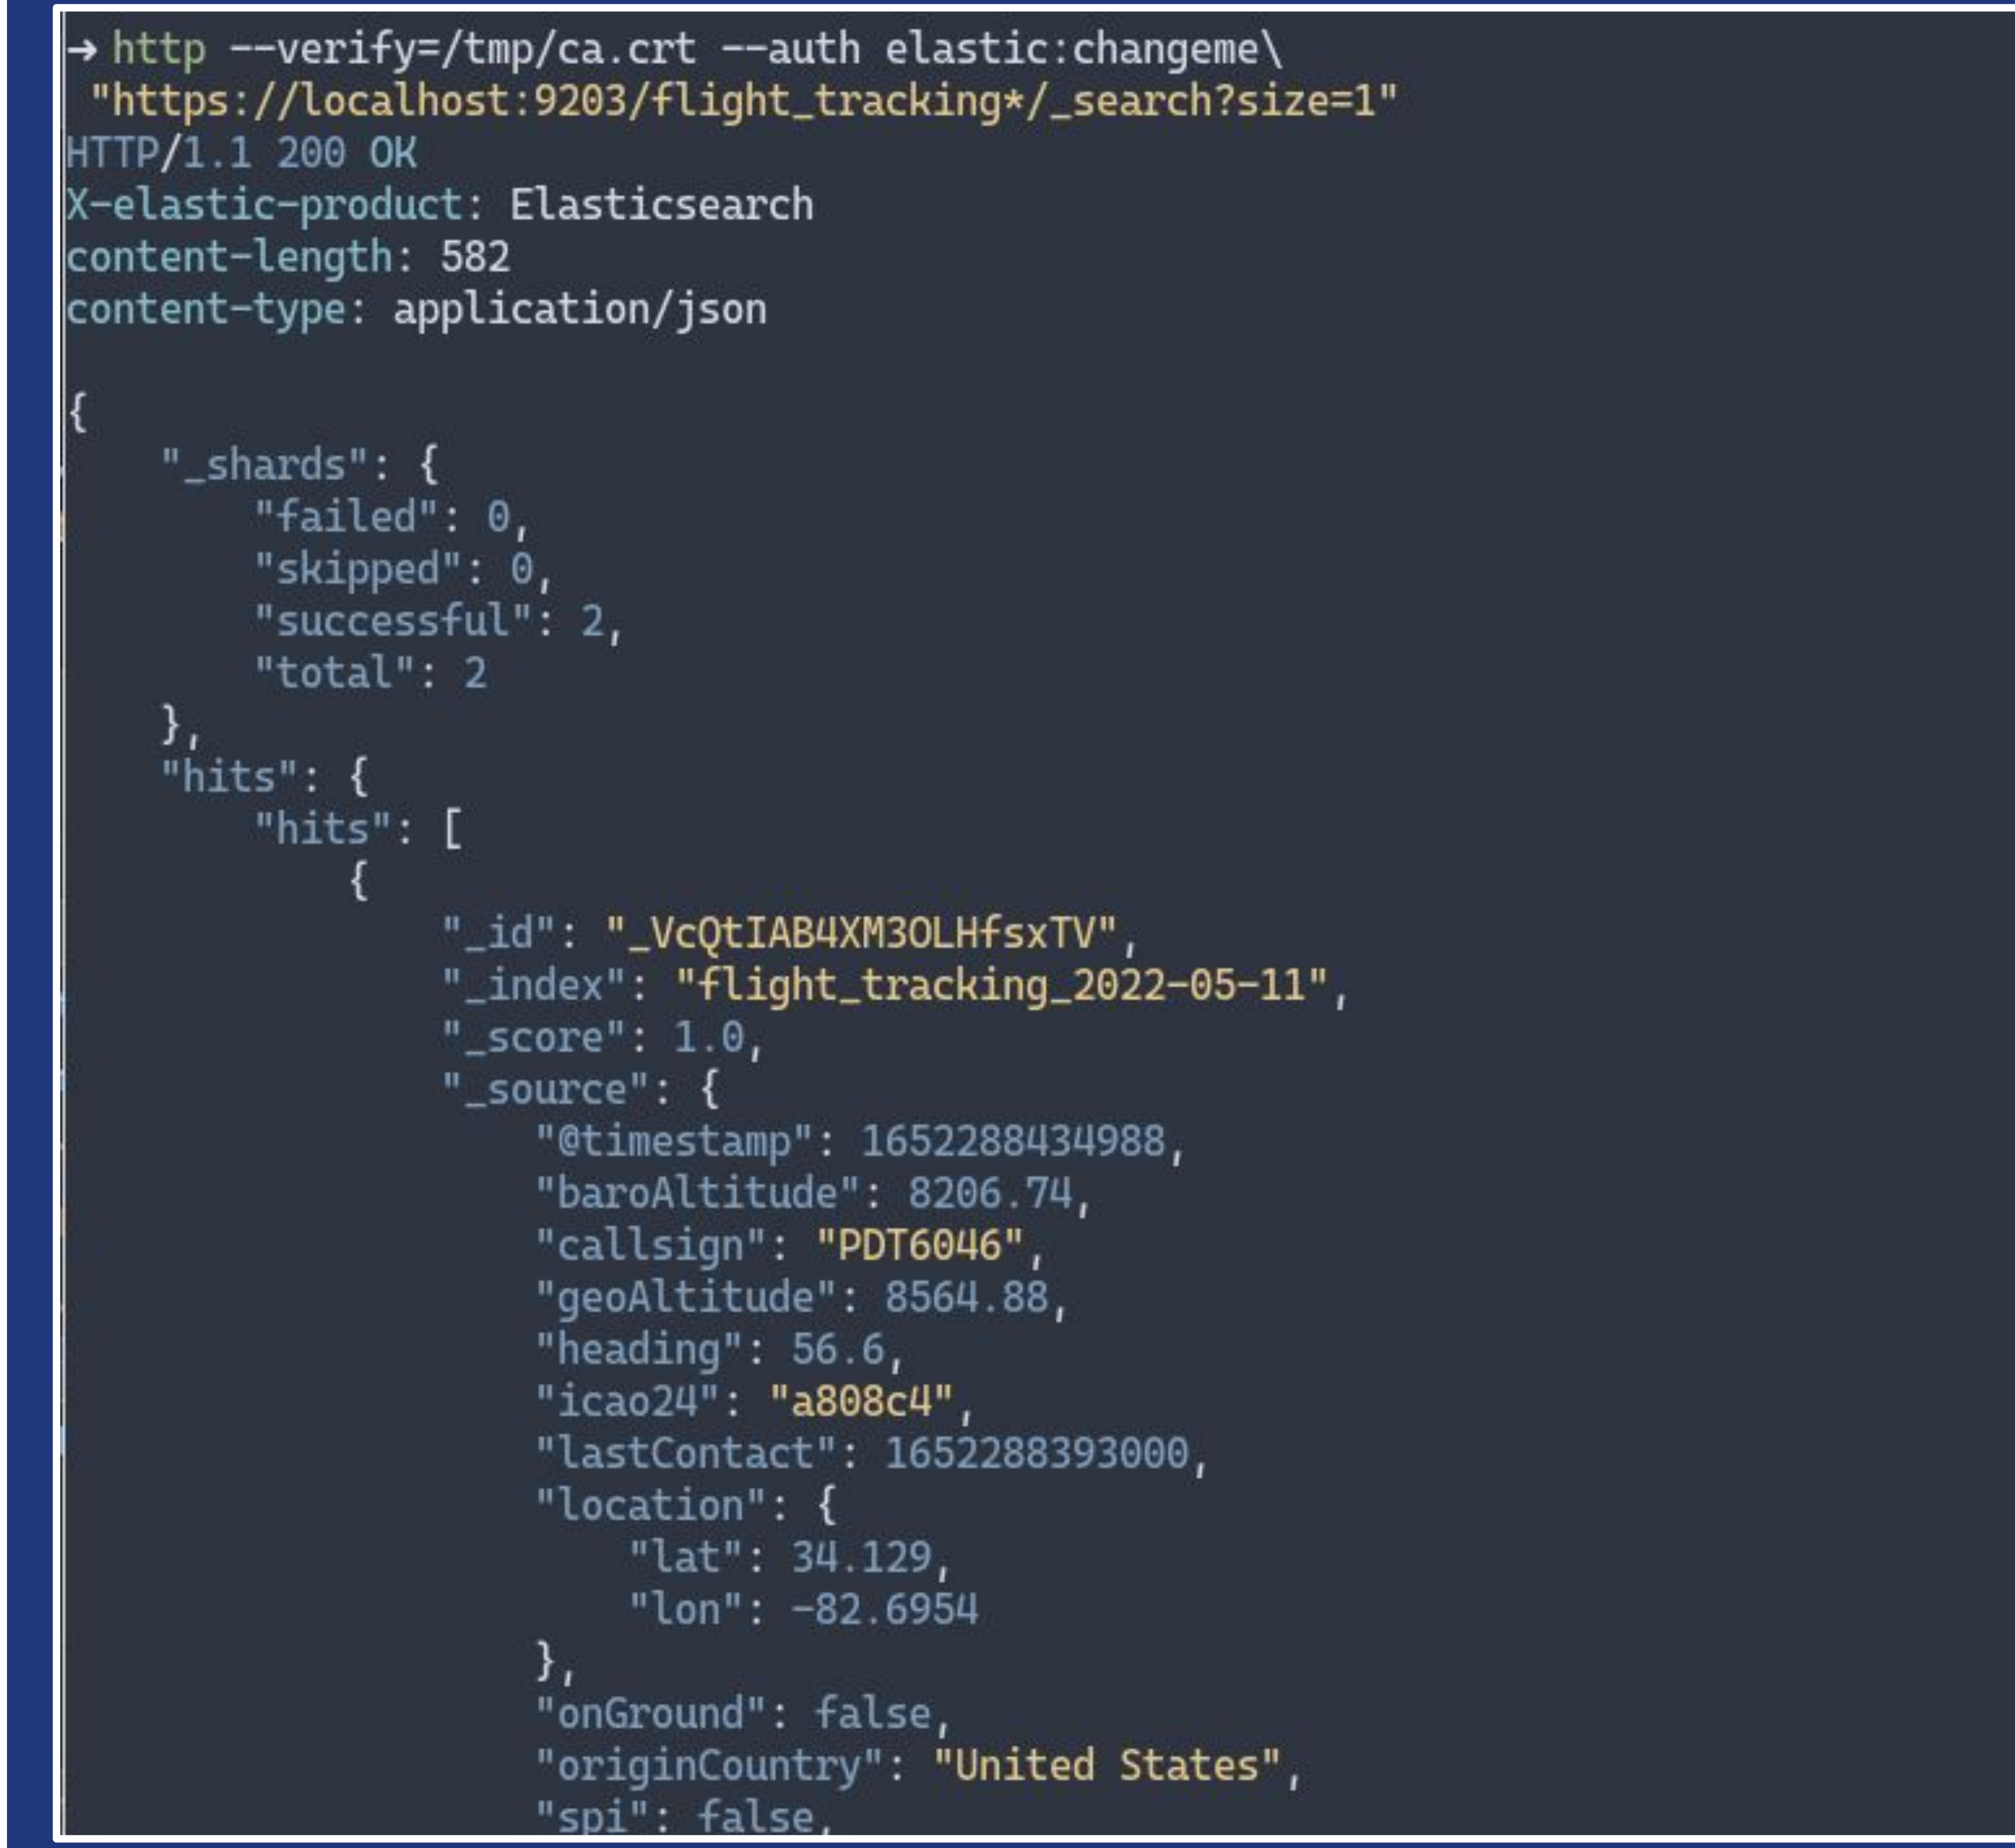

```
→ http --verify=/tmp/ca.crt --auth elastic:changeme\
 "https://localhost:9203/flight_tracking*/_search?size=1"
HTTP/1.1 200 OK
X-elastic-product: Elasticsearch
content-length: 582
content-type: application/json

{
    "_shards": {
        "failed": 0,
        "skipped": 0,
        "successful": 2,
        "total": 2
    },
    "hits": {
        "hits": [
            {
                "_id": "_VcQtIAB4XM3OLHfsxTV",
                "_index": "flight_tracking_2022-05-11",
                "_score": 1.0,
                "_source": {
                    "@timestamp": 1652288434988,
                    "baroAltitude": 8206.74,
                    "callsign": "PDT6046",
                    "geoAltitude": 8564.88,
                    "heading": 56.6,
                    "icao24": "a808c4",
                    "lastContact": 1652288393000,
                    "location": {
                        "lat": 34.129,
                        "lon": -82.6954
                    },
                    "onGround": false,
                    "originCountry": "United States",
                    "spi": false,
```

Communicating with Elasticsearch

- All communication through HTTP endpoints 📖
- JSON
- REST methods: GET, PUT, POST, DELETE

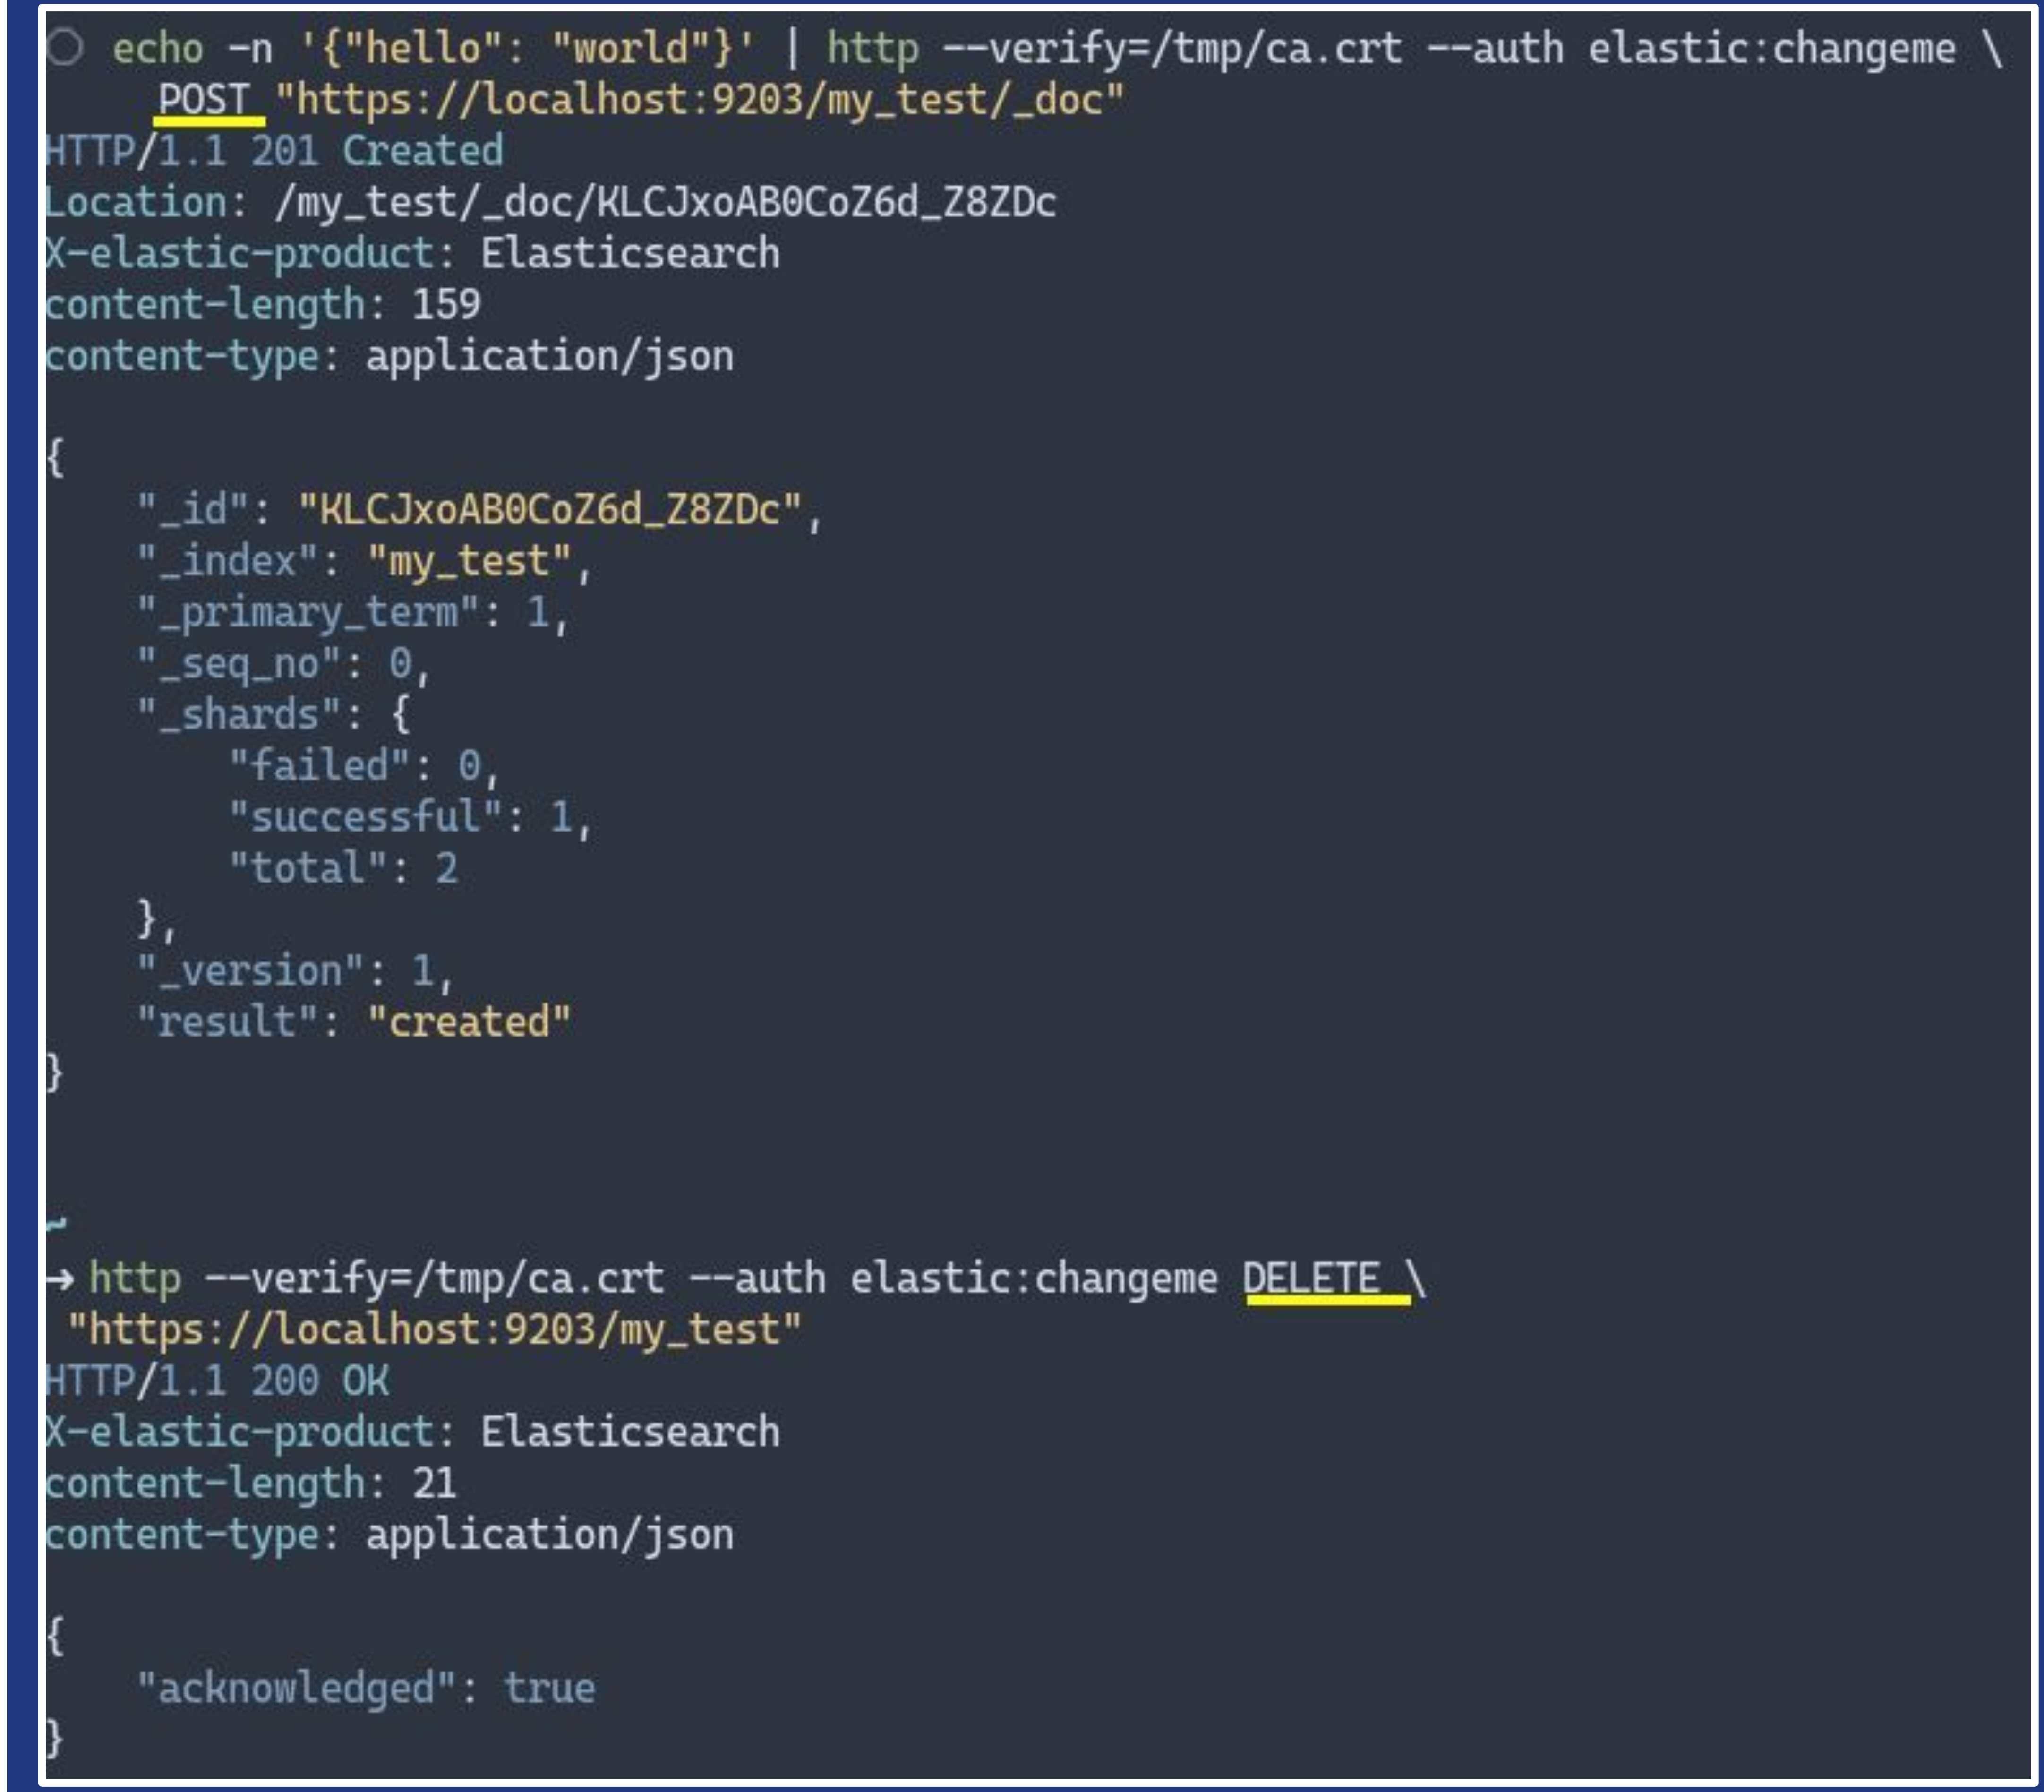

```
○ echo -n '{"hello": "world"}' | http --verify=/tmp/ca.crt --auth elastic:changeme \
    POST "https://localhost:9203/my_test/_doc"
HTTP/1.1 201 Created
Location: /my_test/_doc/KLCJxoAB0CoZ6d_Z8ZDc
X-elastic-product: Elasticsearch
content-length: 159
content-type: application/json

{
    "_id": "KLCJxoAB0CoZ6d_Z8ZDc",
    "_index": "my_test",
    "_primary_term": 1,
    "_seq_no": 0,
    "_shards": {
        "failed": 0,
        "successful": 1,
        "total": 2
    },
    "_version": 1,
    "result": "created"
}



~
→ http --verify=/tmp/ca.crt --auth elastic:changeme DELETE \
 "https://localhost:9203/my_test"
HTTP/1.1 200 OK
X-elastic-product: Elasticsearch
content-length: 21
content-type: application/json

{
    "acknowledged": true
}
```

Communicating with Elasticsearch

- All communication through HTTP endpoints 📖
- JSON
- REST methods: GET, PUT, POST, DELETE
- _cat API for human readable display 📖

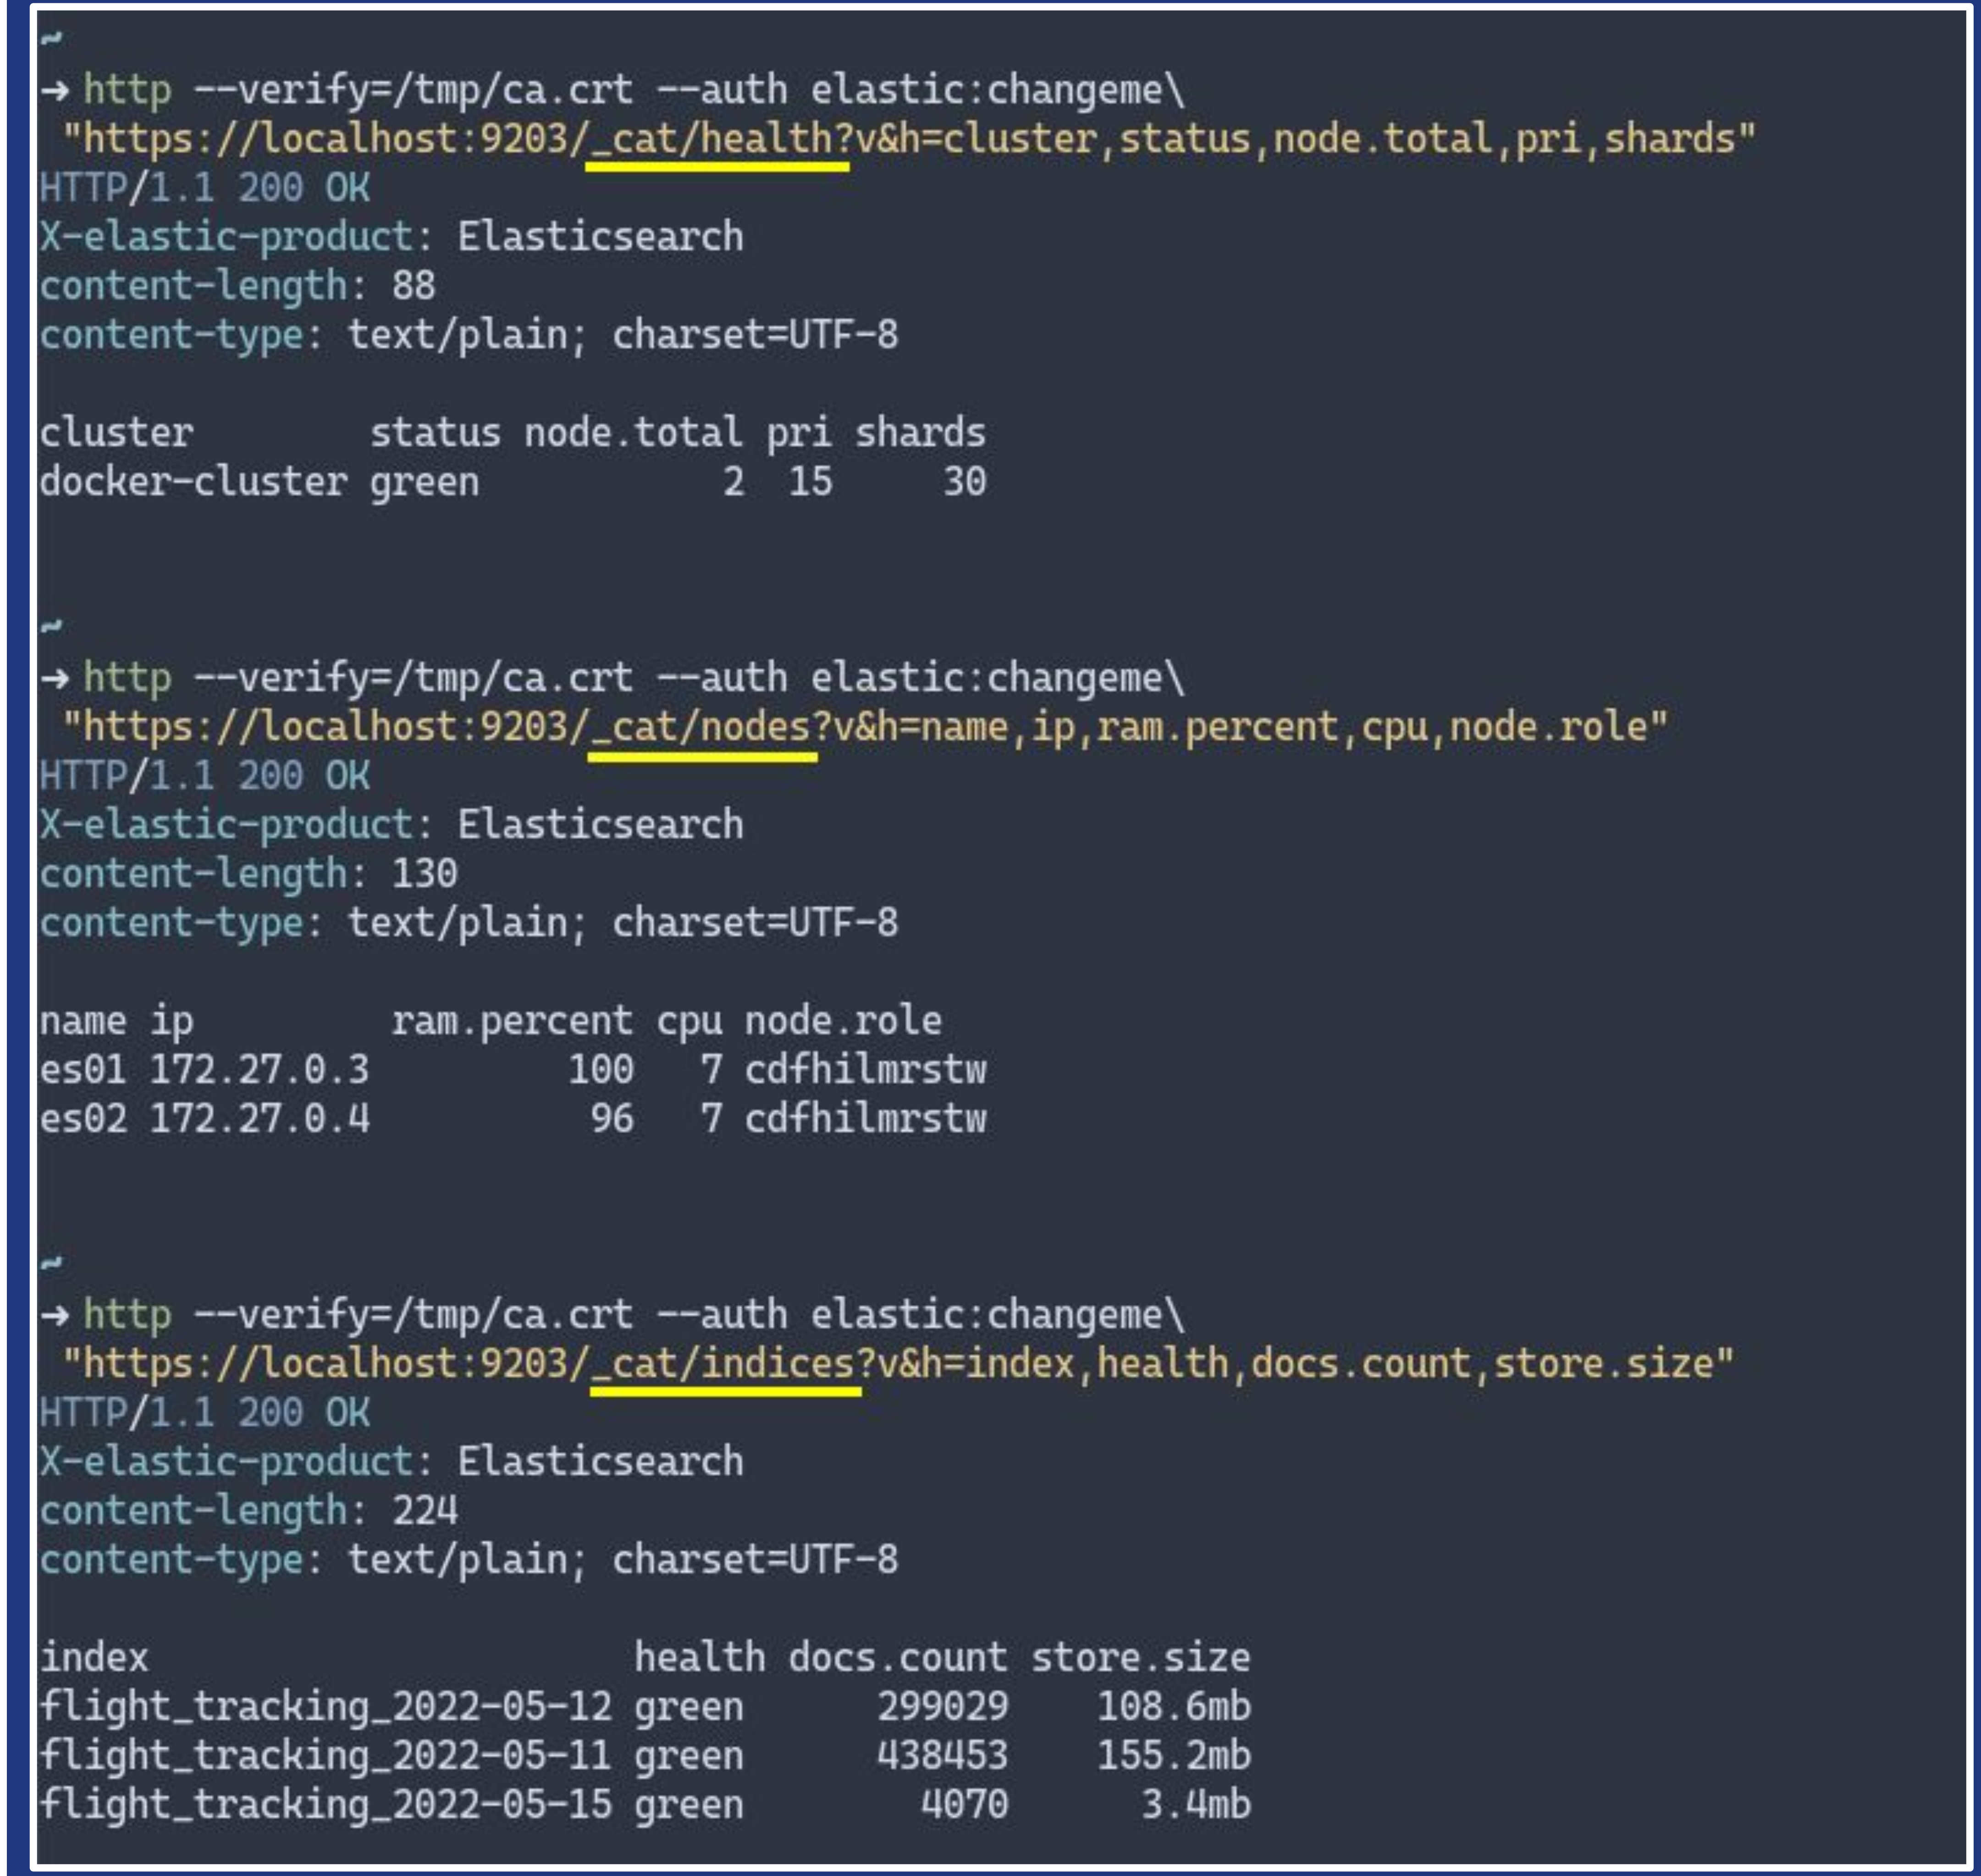

```
~
→ http --verify=/tmp/ca.crt --auth elastic:changeme\
 "https://localhost:9203/_cat/health?v&h=cluster,status,node.total,pri,shards"
HTTP/1.1 200 OK
X-elastic-product: Elasticsearch
content-length: 88
content-type: text/plain; charset=UTF-8

cluster        status node.total pri shards
docker-cluster green           2  15     30


~
→ http --verify=/tmp/ca.crt --auth elastic:changeme\
 "https://localhost:9203/_cat/nodes?v&h=name,ip,ram.percent,cpu,node.role"
HTTP/1.1 200 OK
X-elastic-product: Elasticsearch
content-length: 130
content-type: text/plain; charset=UTF-8

name ip         ram.percent cpu node.role
es01 172.27.0.3         100   7 cdfhilmrstw
es02 172.27.0.4          96   7 cdfhilmrstw


~
→ http --verify=/tmp/ca.crt --auth elastic:changeme\
 "https://localhost:9203/_cat/indices?v&h=index,health,docs.count,store.size"
HTTP/1.1 200 OK
X-elastic-product: Elasticsearch
content-length: 224
content-type: text/plain; charset=UTF-8

index                      health docs.count store.size
flight_tracking_2022-05-12 green      299029    108.6mb
flight_tracking_2022-05-11 green      438453    155.2mb
flight_tracking_2022-05-15 green        4070      3.4mb
```

Beats

Lightweight data shippers

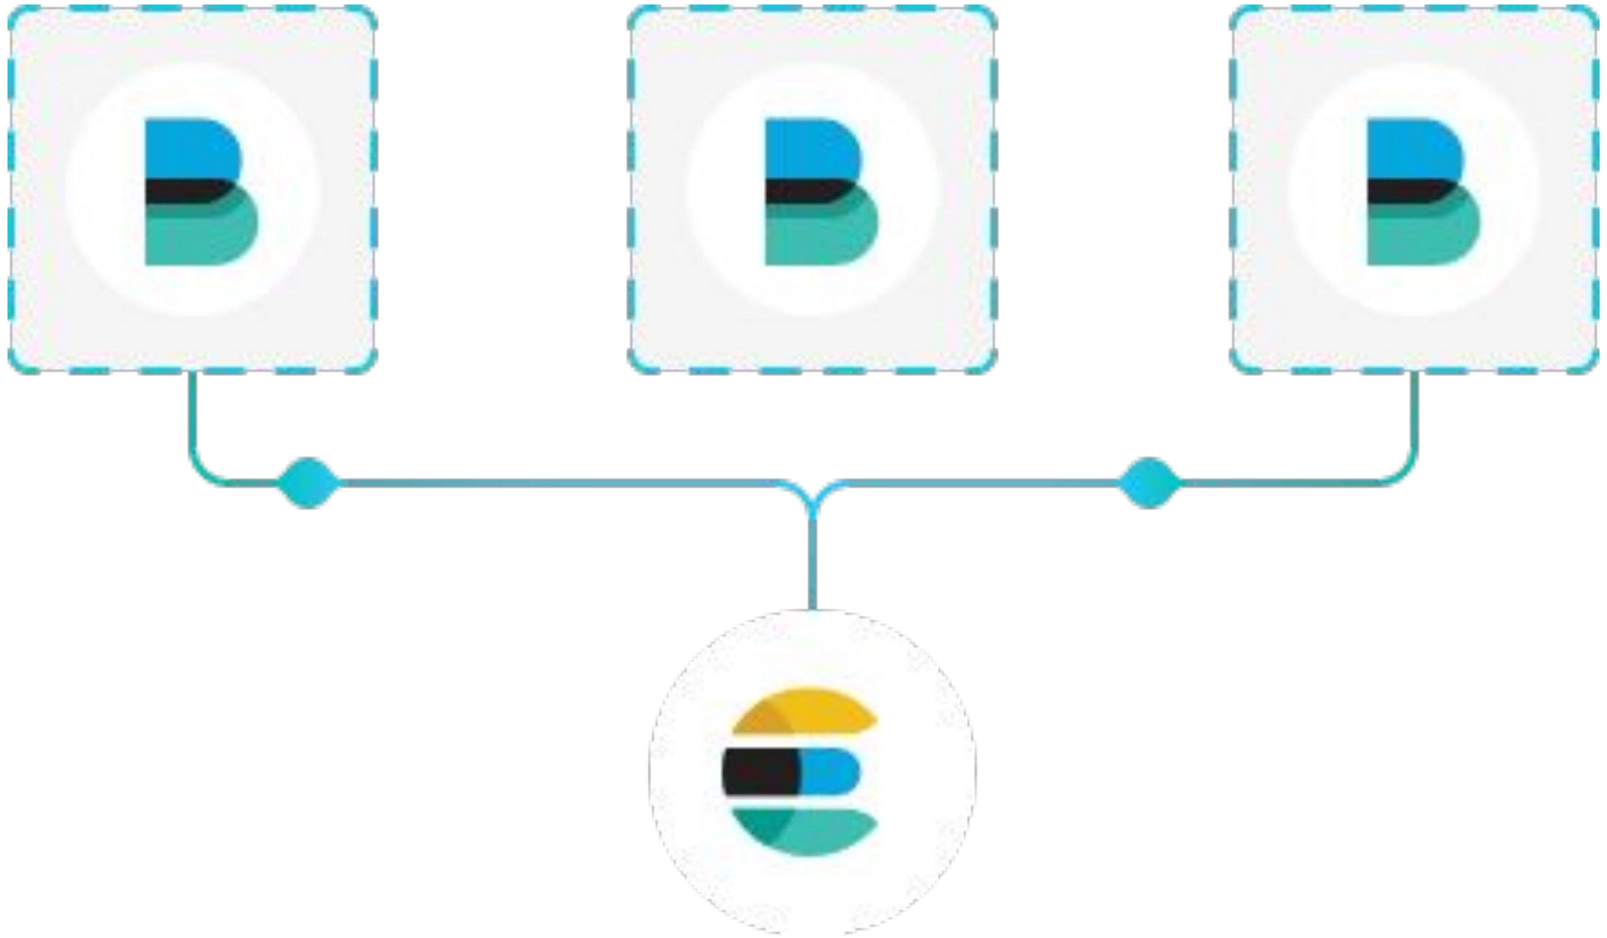

Ship data from the source	Ship and centralize in Elasticsearch	Ship to Logstash for transformation and parsing
Ship to Elastic Cloud	Libbeat: API framework to build custom beats	70+ community Beats

Beats

Lightweight data shippers

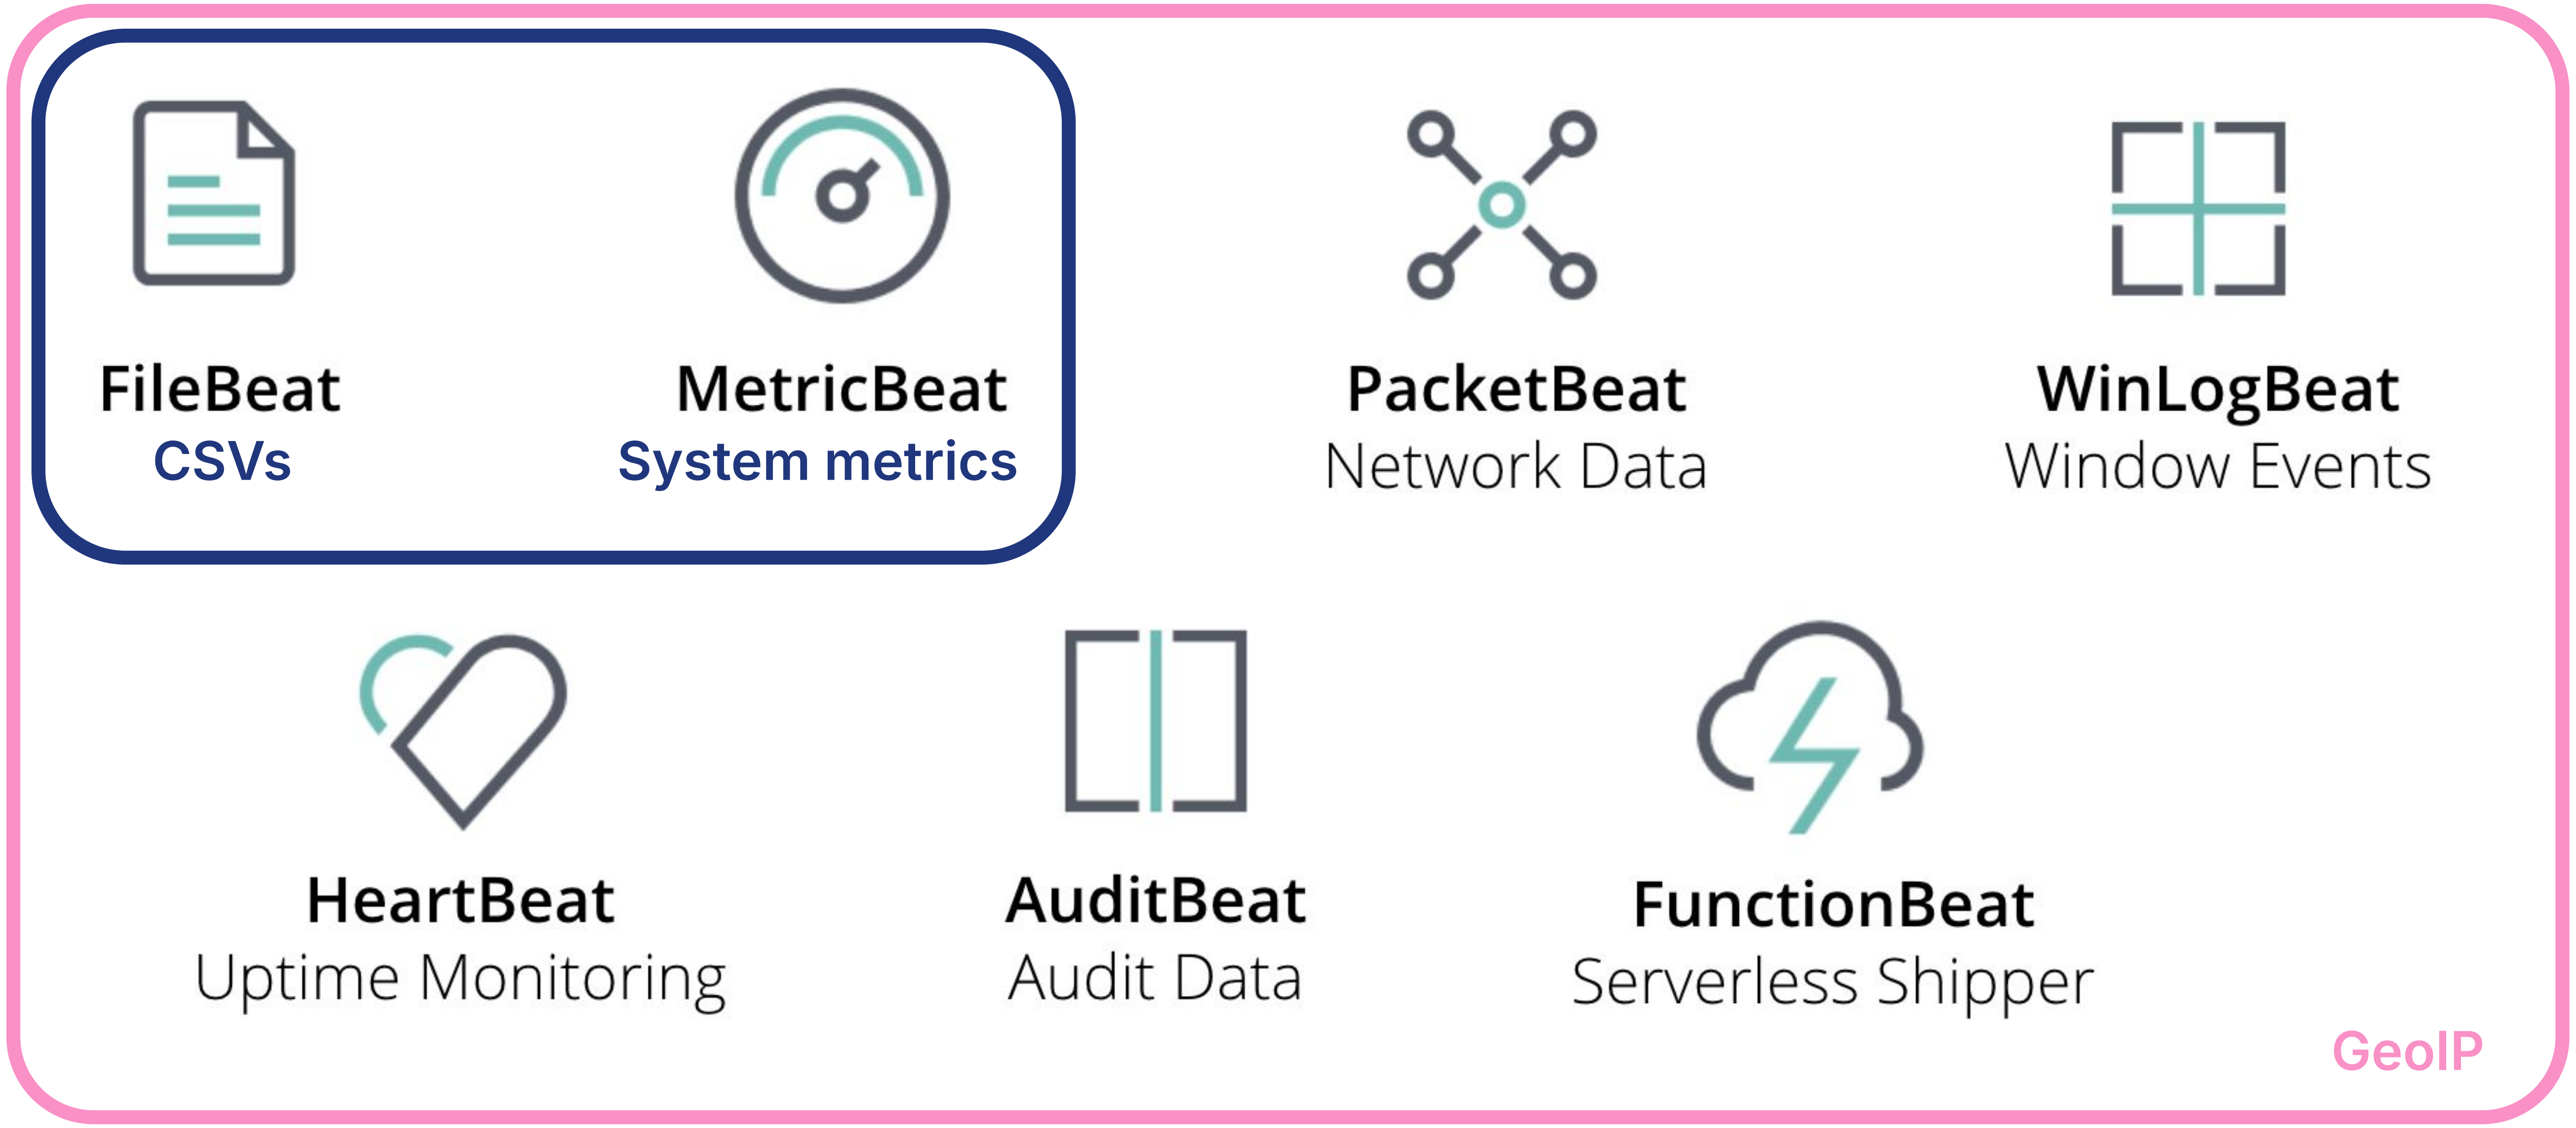

FileBeat
CSVs

MetricBeat
System metrics

PacketBeat
Network Data

WinLogBeat
Window Events

HeartBeat
Uptime Monitoring

AuditBeat
Audit Data

FunctionBeat
Serverless Shipper

GeoIP

Logstash

ETL for Elasticsearch

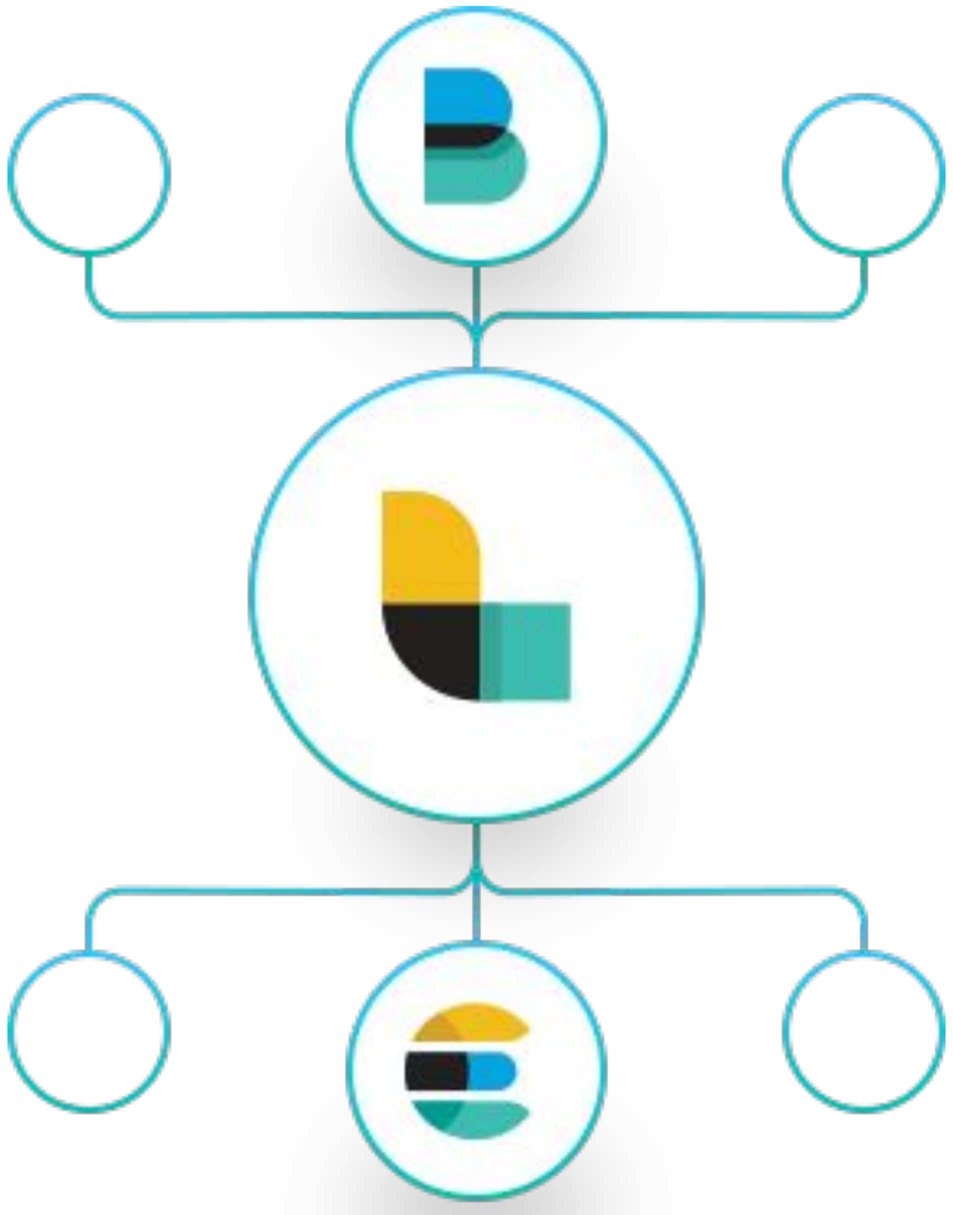

Ingest data of all shapes, sizes, and sources	Parse and dynamically transform data	Transport data to any output
Secure and encrypt data inputs	Build your own pipelines	Lots of plugins

elastic

Logstash

Normalize and Enrich Data before Indexing

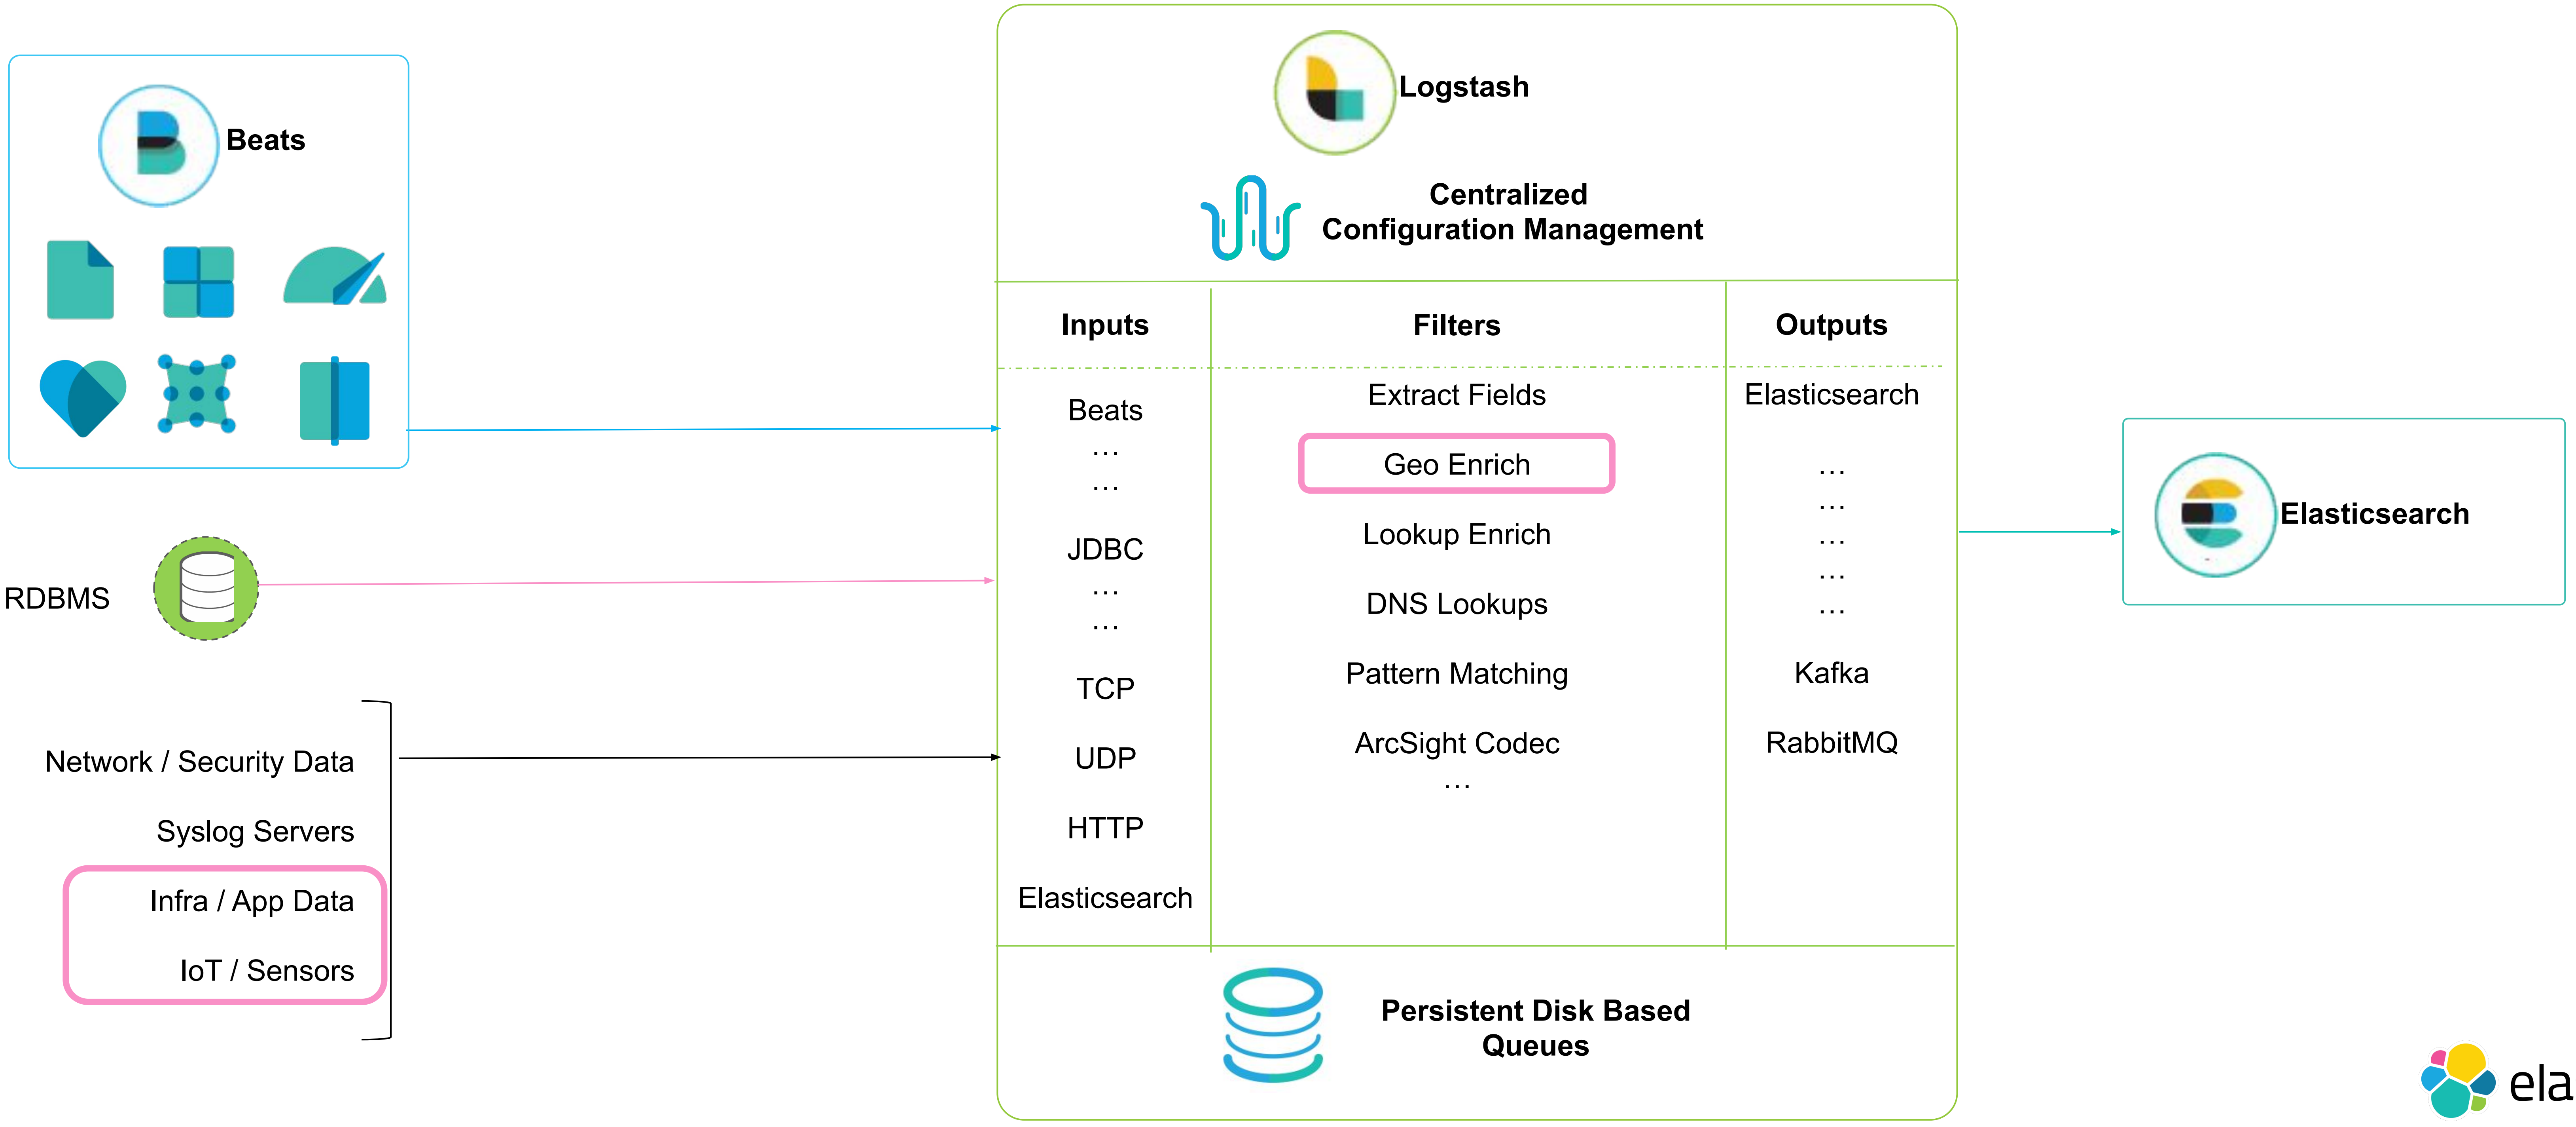

Elastic **Stack**

Ingest, Store, Search, Visualise

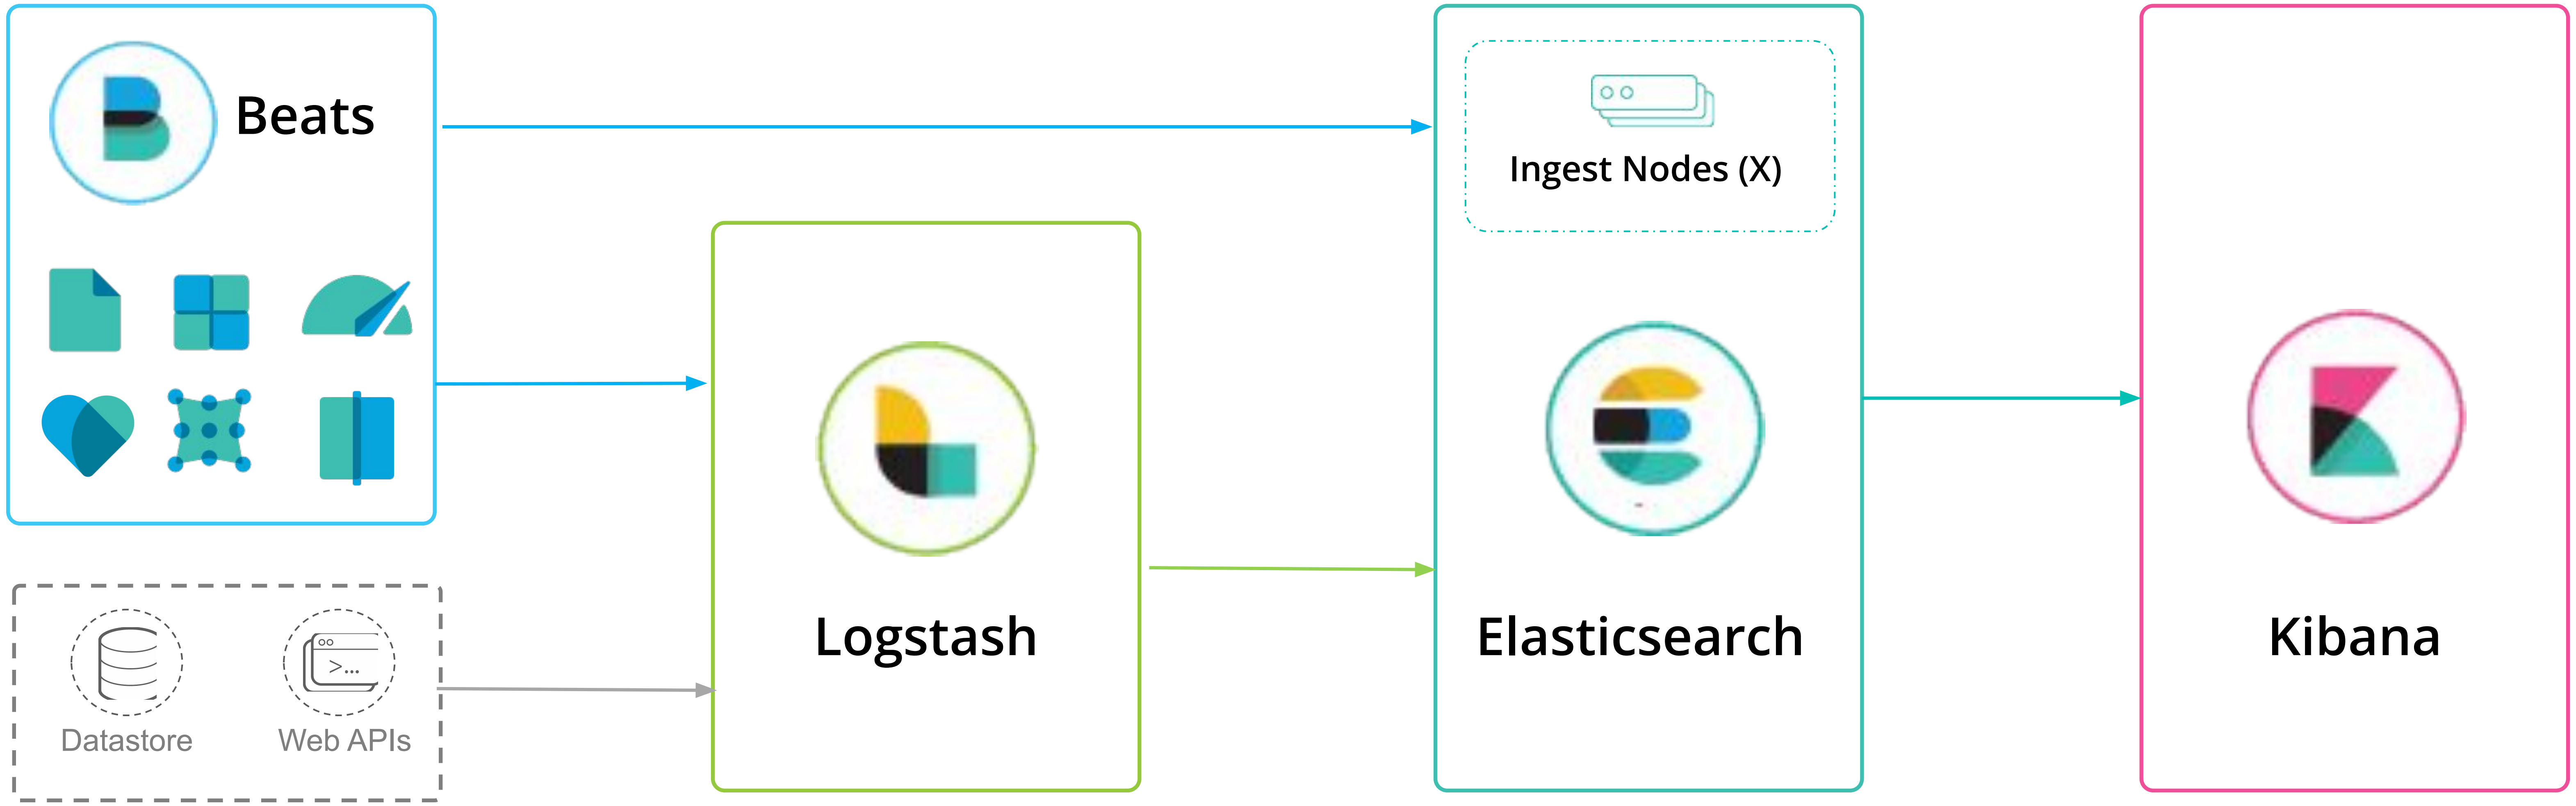

Agenda

1. Introduction to the Elastic stack
2. Lab setup
3. Kibana introduction
4. Elastics Maps
5. Elasticsearch Geo
6. Web mapping and Elasticsearch

Run the stack from Elastic Cloud or your laptop

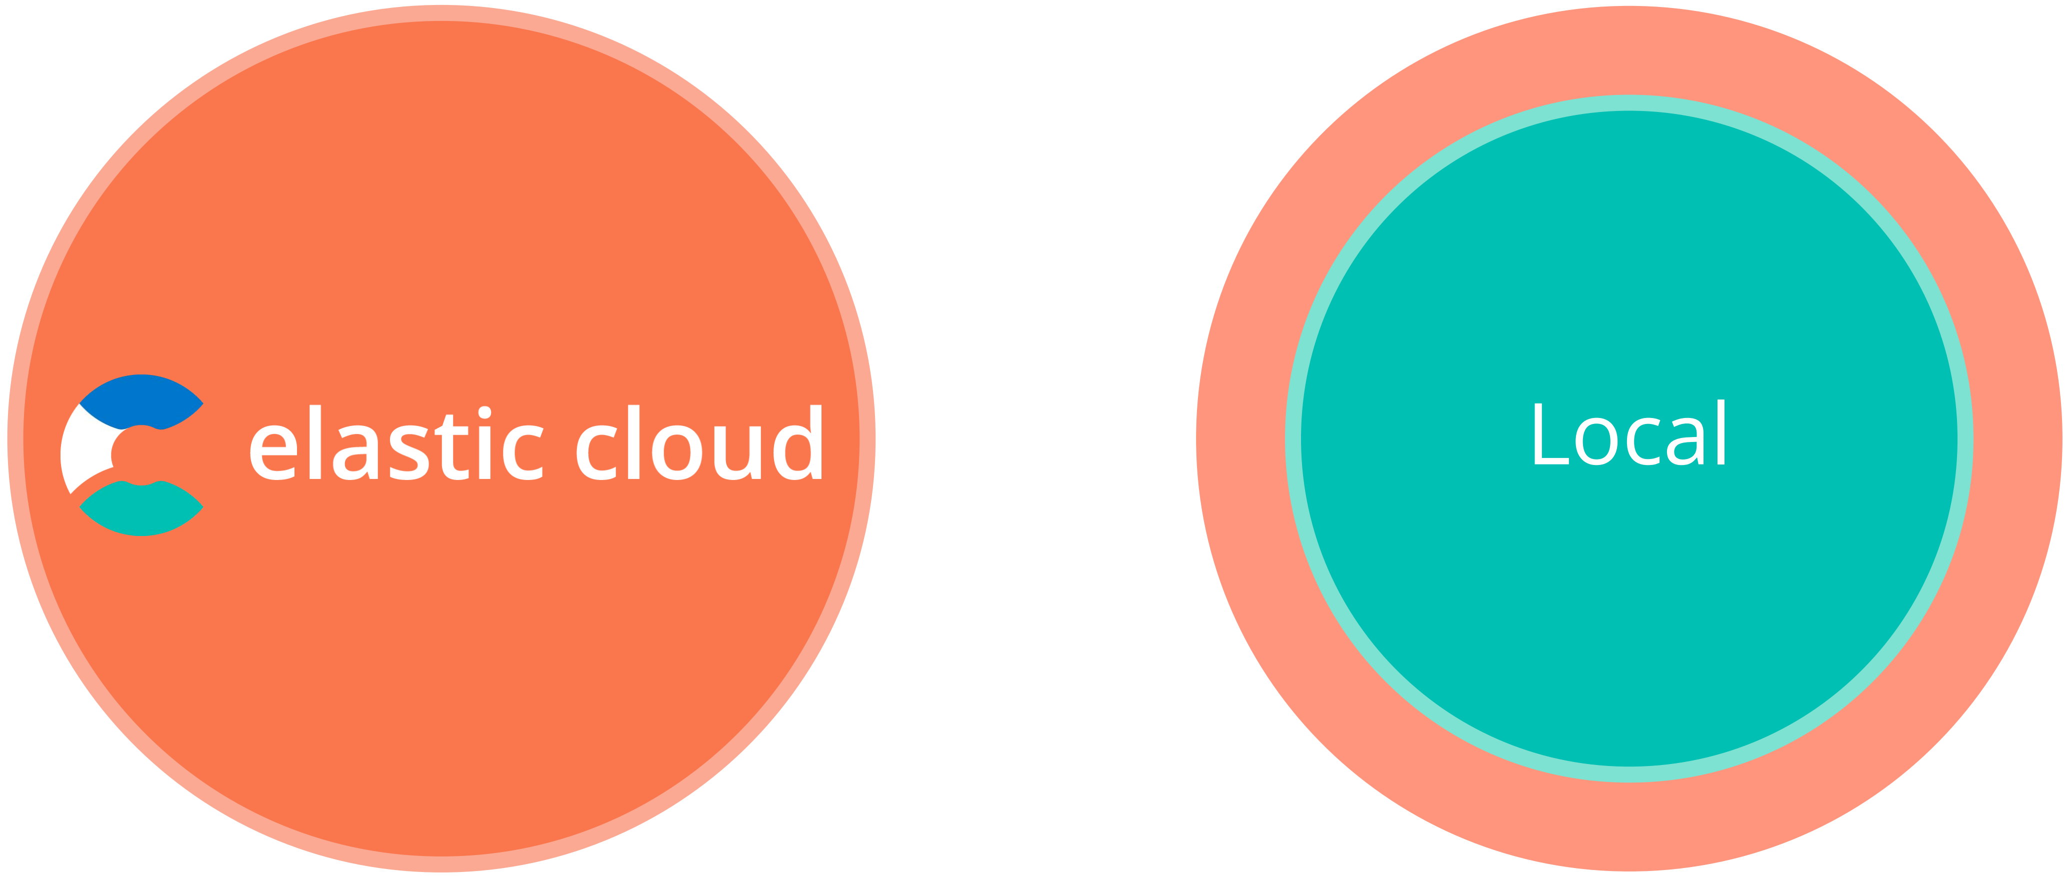

elastic cloud

Local

Elastic Cloud trial account

- You can create a two week free trial account at **cloud.elastic.co**
- Create a new *deployment* (**details**)
  - Default settings are fine
- Loading some data (open sky):
  - Use Node.js or Docker to upload real-time data with the opensky-loader script (**details**)
  - Or you can upload a **static dataset** or even generate a new one (**details**) as CSV or GeoJSON

Using Docker Compose

- Prerequisite: install Docker and the Compose plugin
- Download the lab: https://github.com/jsanz/elastic-workshop

```
git clone https://github.com/jsanz/elastic-workshop.git
```

- Adjust `.env` parameters
  - OpenSky viewer port (for example, port 80 is not available in OSGeo Live)
    - `OPENSKY_VIEWER_PORT=8080`
  - (Optional) change elastic passwords
    - `ELASTIC_PASSWORD=changeme`
    - `KIBANA_PASSWORD=changeme`
  - (Optional) set up OpenSky credentials (1000 requests for users, 100 for anonymous)
    - `OPENSKY_USER=`
    - `OPENSKY_PASSWORD=`

Start up

```
$ cd elastic-workshop/lab

$ sudo sysctl -w vm.max_map_count=262144 # adjust virtual memory (only Linux)

$ docker compose pull # download Elastic stack images

$ docker compose build # build the local opensky images

$ docker compose up -d # start all the containers
```

Other useful commands

```
$ docker compose ps # shows status info

$ docker compose logs -f kibana # show logs

$ docker compose down # shut down all containers and services

$ docker compose restart opensky-viewer # resets a container

$ docker compose stop opensky-viewer # stops a container

$ docker volume ls # list all the volumes (Docker hard disks)

$ docker volume rm [volume-name] # deletes a volume
```

Checking the status of the containers

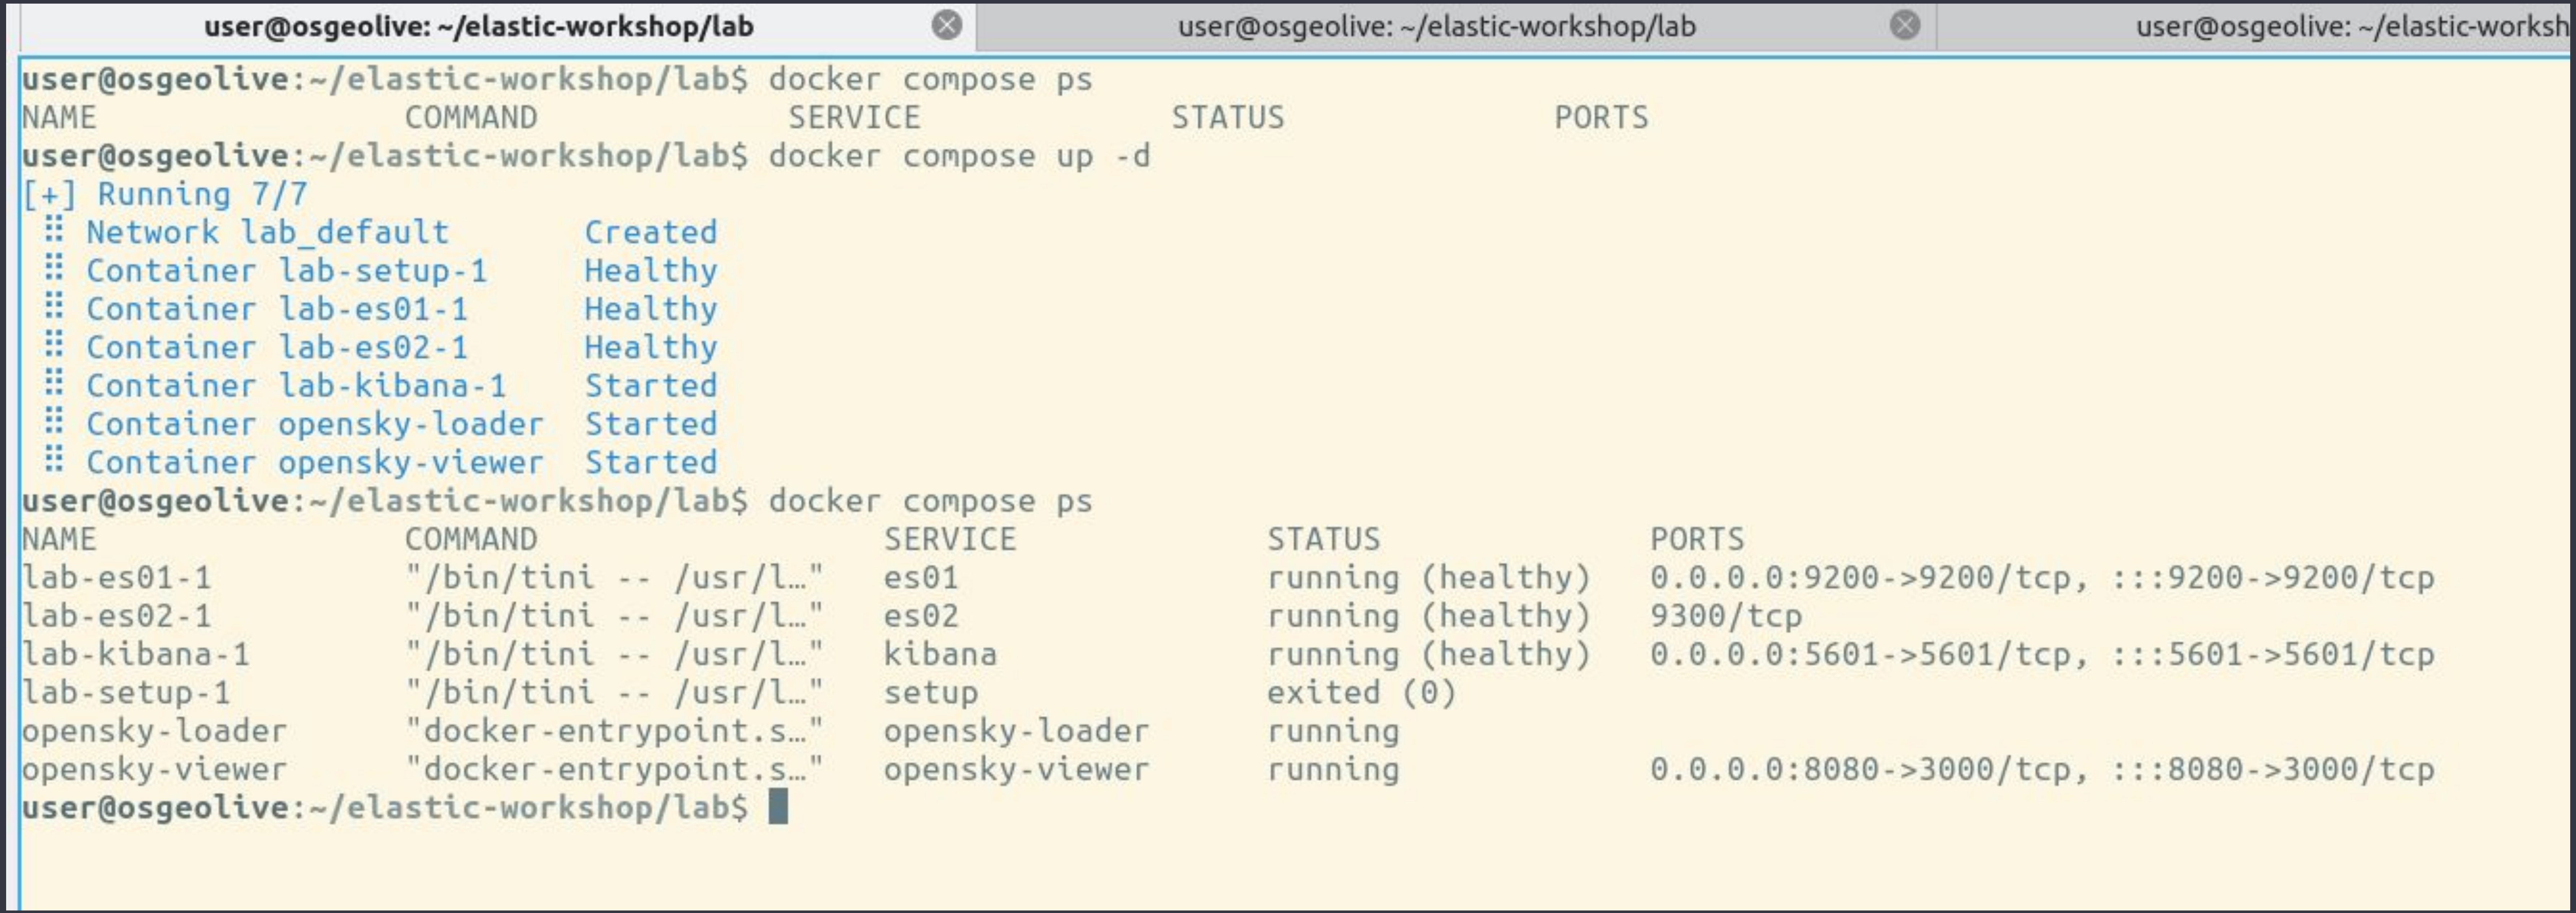
user@osgeolive: ~/elastic-workshop/lab
user@osgeolive: ~/elastic-workshop/lab
user@osgeolive: ~/elastic-worksh

```
user@osgeolive:~/elastic-workshop/lab$ docker compose ps
NAME                COMMAND             SERVICE             STATUS              PORTS
user@osgeolive:~/elastic-workshop/lab$ docker compose up -d
[+] Running 7/7
 ⠿ Network lab_default        Created
 ⠿ Container lab-setup-1      Healthy
 ⠿ Container lab-es01-1       Healthy
 ⠿ Container lab-es02-1       Healthy
 ⠿ Container lab-kibana-1     Started
 ⠿ Container opensky-loader   Started
 ⠿ Container opensky-viewer   Started
user@osgeolive:~/elastic-workshop/lab$ docker compose ps
NAME                COMMAND                  SERVICE             STATUS              PORTS
lab-es01-1          "/bin/tini -- /usr/l…"   es01                running (healthy)   0.0.0.0:9200->9200/tcp, :::9200->9200/tcp
lab-es02-1          "/bin/tini -- /usr/l…"   es02                running (healthy)   9300/tcp
lab-kibana-1        "/bin/tini -- /usr/l…"   kibana              running (healthy)   0.0.0.0:5601->5601/tcp, :::5601->5601/tcp
lab-setup-1         "/bin/tini -- /usr/l…"   setup               exited (0)
opensky-loader      "docker-entrypoint.s…"   opensky-loader      running
opensky-viewer      "docker-entrypoint.s…"   opensky-viewer      running             0.0.0.0:8080->3000/tcp, :::8080->3000/tcp
user@osgeolive:~/elastic-workshop/lab$ █
```

What do we have?

- `es01` & `es02`: Elasticsearch cluster, `https://localhost:9200`
- `opensky-loader`: data loader
- `kibana`: `http://localhost:5601`
- `opensky-viewer`: sample consuming app, `http://localhost:80`

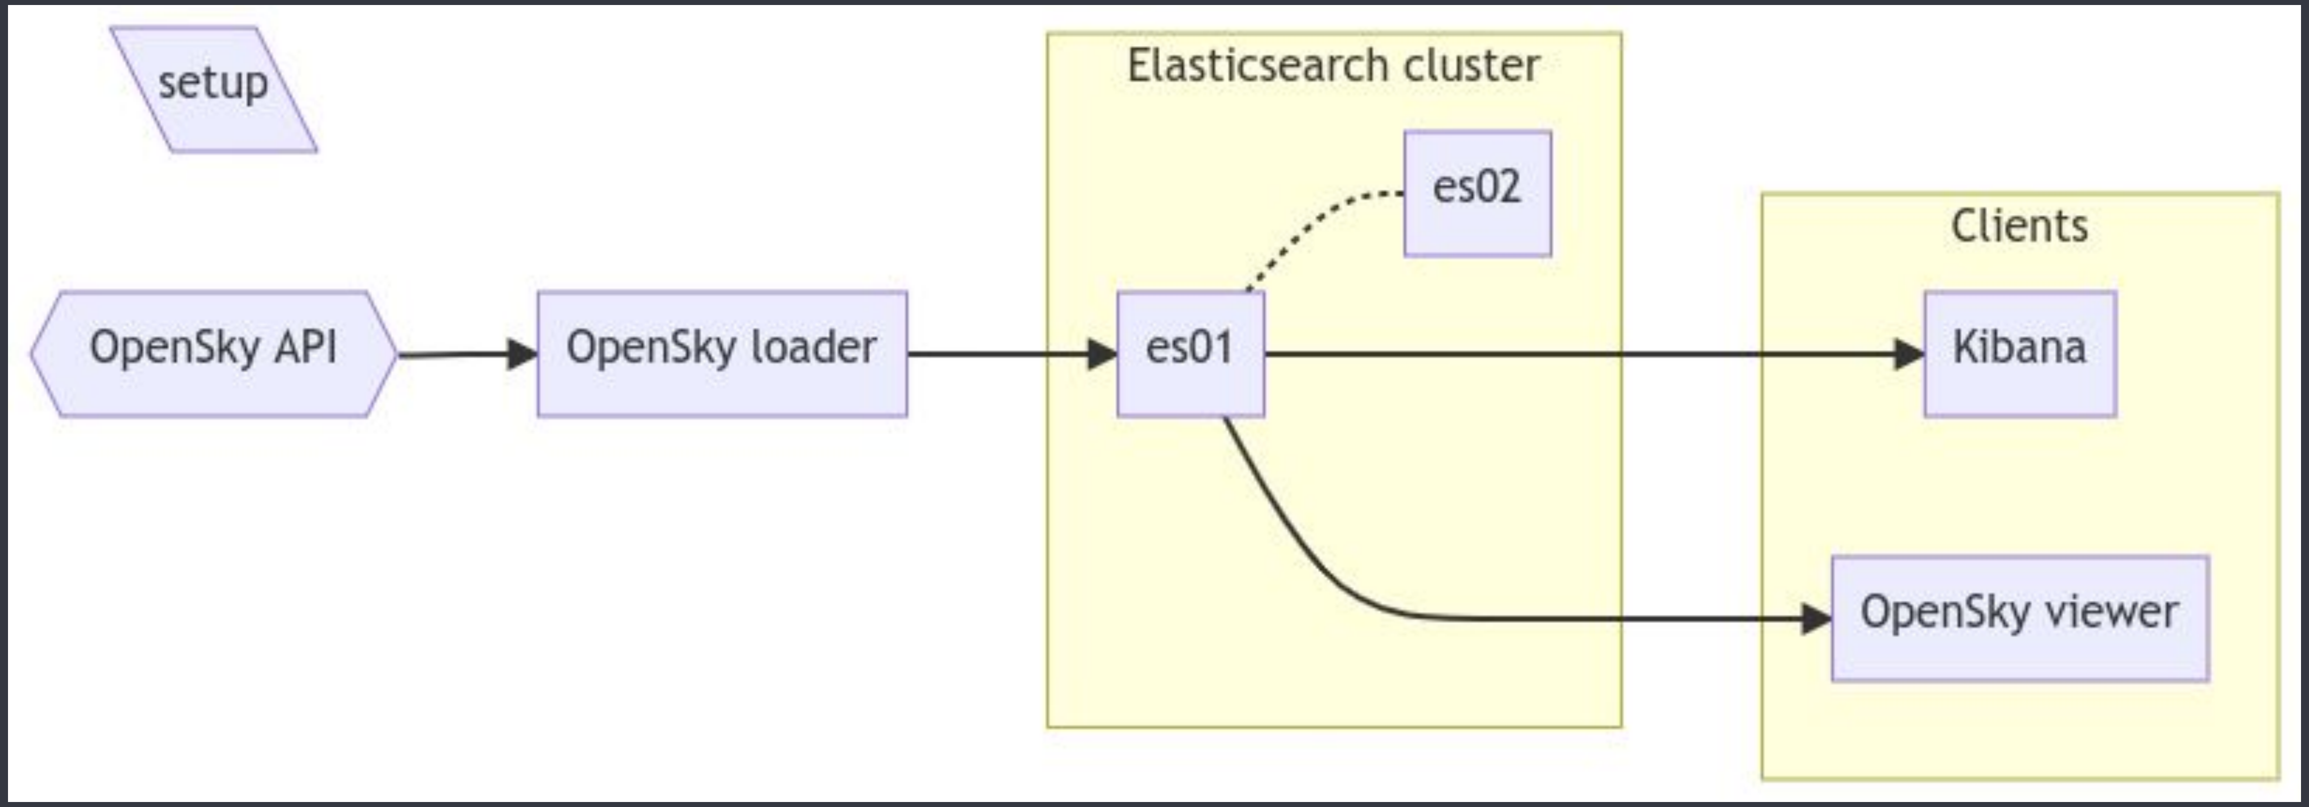

Agenda

1. Introduction to the Elastic stack
2. Lab setup
3. Kibana introduction
4. Elastics Maps
5. Elasticsearch Geo
6. Web mapping and Elasticsearch

Kibana

Some basic concepts about Kibana

Developer Tools

Console

Allows to run Elasticsearch queries with autocomplete, code formatting, history, etc.

Search profiler

Shows statistics about query performance.

Grok debugger

Helps creating grok expressions for Logstash.

Painless lab

An environment to test painless scripts.

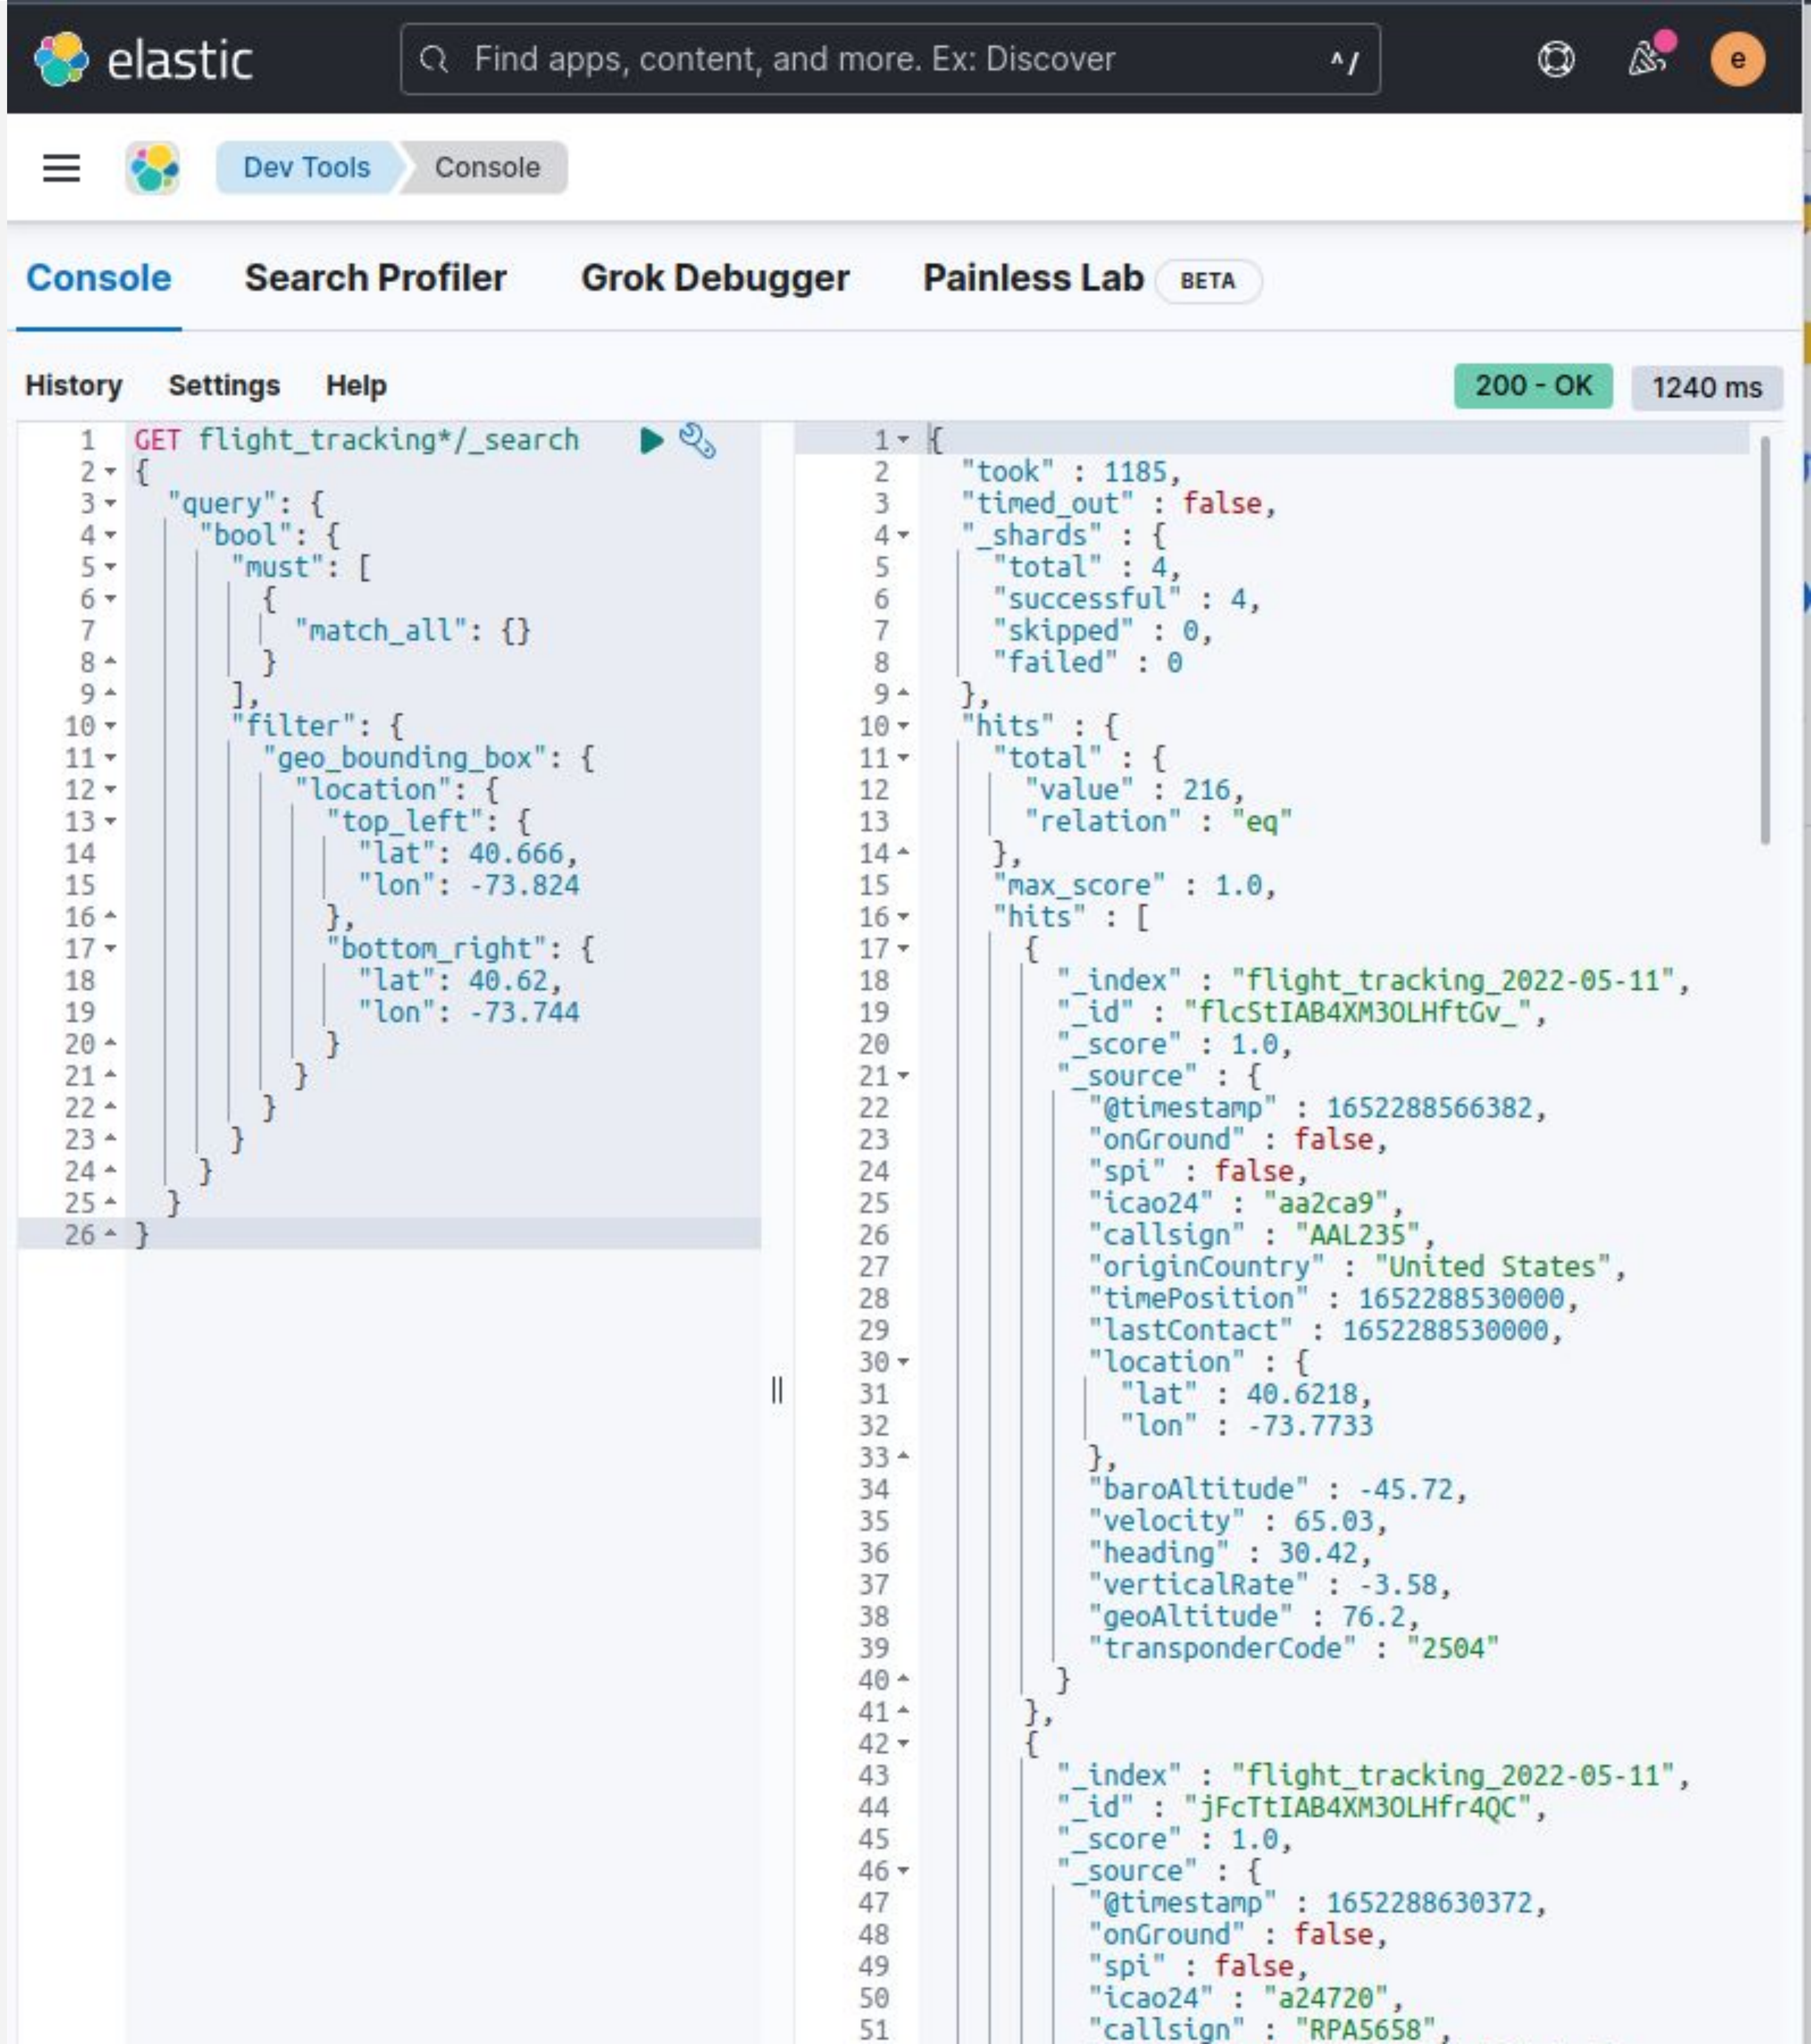

```
GET flight_tracking*/_search
{
  "query": {
    "bool": {
      "must": [
        {
          "match_all": {}
        }
      ],
      "filter": {
        "geo_bounding_box": {
          "location": {
            "top_left": {
              "lat": 40.666,
              "lon": -73.824
            },
            "bottom_right": {
              "lat": 40.62,
              "lon": -73.744
            }
          }
        }
      }
    }
  }
}
```

```
{
  "took" : 1185,
  "timed_out" : false,
  "_shards" : {
    "total" : 4,
    "successful" : 4,
    "skipped" : 0,
    "failed" : 0
  },
  "hits" : {
    "total" : {
      "value" : 216,
      "relation" : "eq"
    },
    "max_score" : 1.0,
    "hits" : [
      {
        "_index" : "flight_tracking_2022-05-11",
        "_id" : "flcStIAB4XM3OLHftGv_",
        "_score" : 1.0,
        "_source" : {
          "@timestamp" : 1652288566382,
          "onGround" : false,
          "spi" : false,
          "icao24" : "aa2ca9",
          "callsign" : "AAL235",
          "originCountry" : "United States",
          "timePosition" : 1652288530000,
          "lastContact" : 1652288530000,
          "location" : {
            "lat" : 40.6218,
            "lon" : -73.7733
          },
          "baroAltitude" : -45.72,
          "velocity" : 65.03,
          "heading" : 30.42,
          "verticalRate" : -3.58,
          "geoAltitude" : 76.2,
          "transponderCode" : "2504"
        }
      },
      {
        "_index" : "flight_tracking_2022-05-11",
        "_id" : "jFcTtIAB4XM3OLHfr4QC",
        "_score" : 1.0,
        "_source" : {
          "@timestamp" : 1652288630372,
          "onGround" : false,
          "spi" : false,
          "icao24" : "a24720",
          "callsign" : "RPA5658",
```

Data Views

- Logic component that **gathers** indices using a name **pattern**
  - `my_application_logs_*`
- Defines field **formatters**: number, currency, image, URL, …
- Defines **temporal field** for filtering (optional)
- **Runtime fields** for query time computations

Creating a Data View

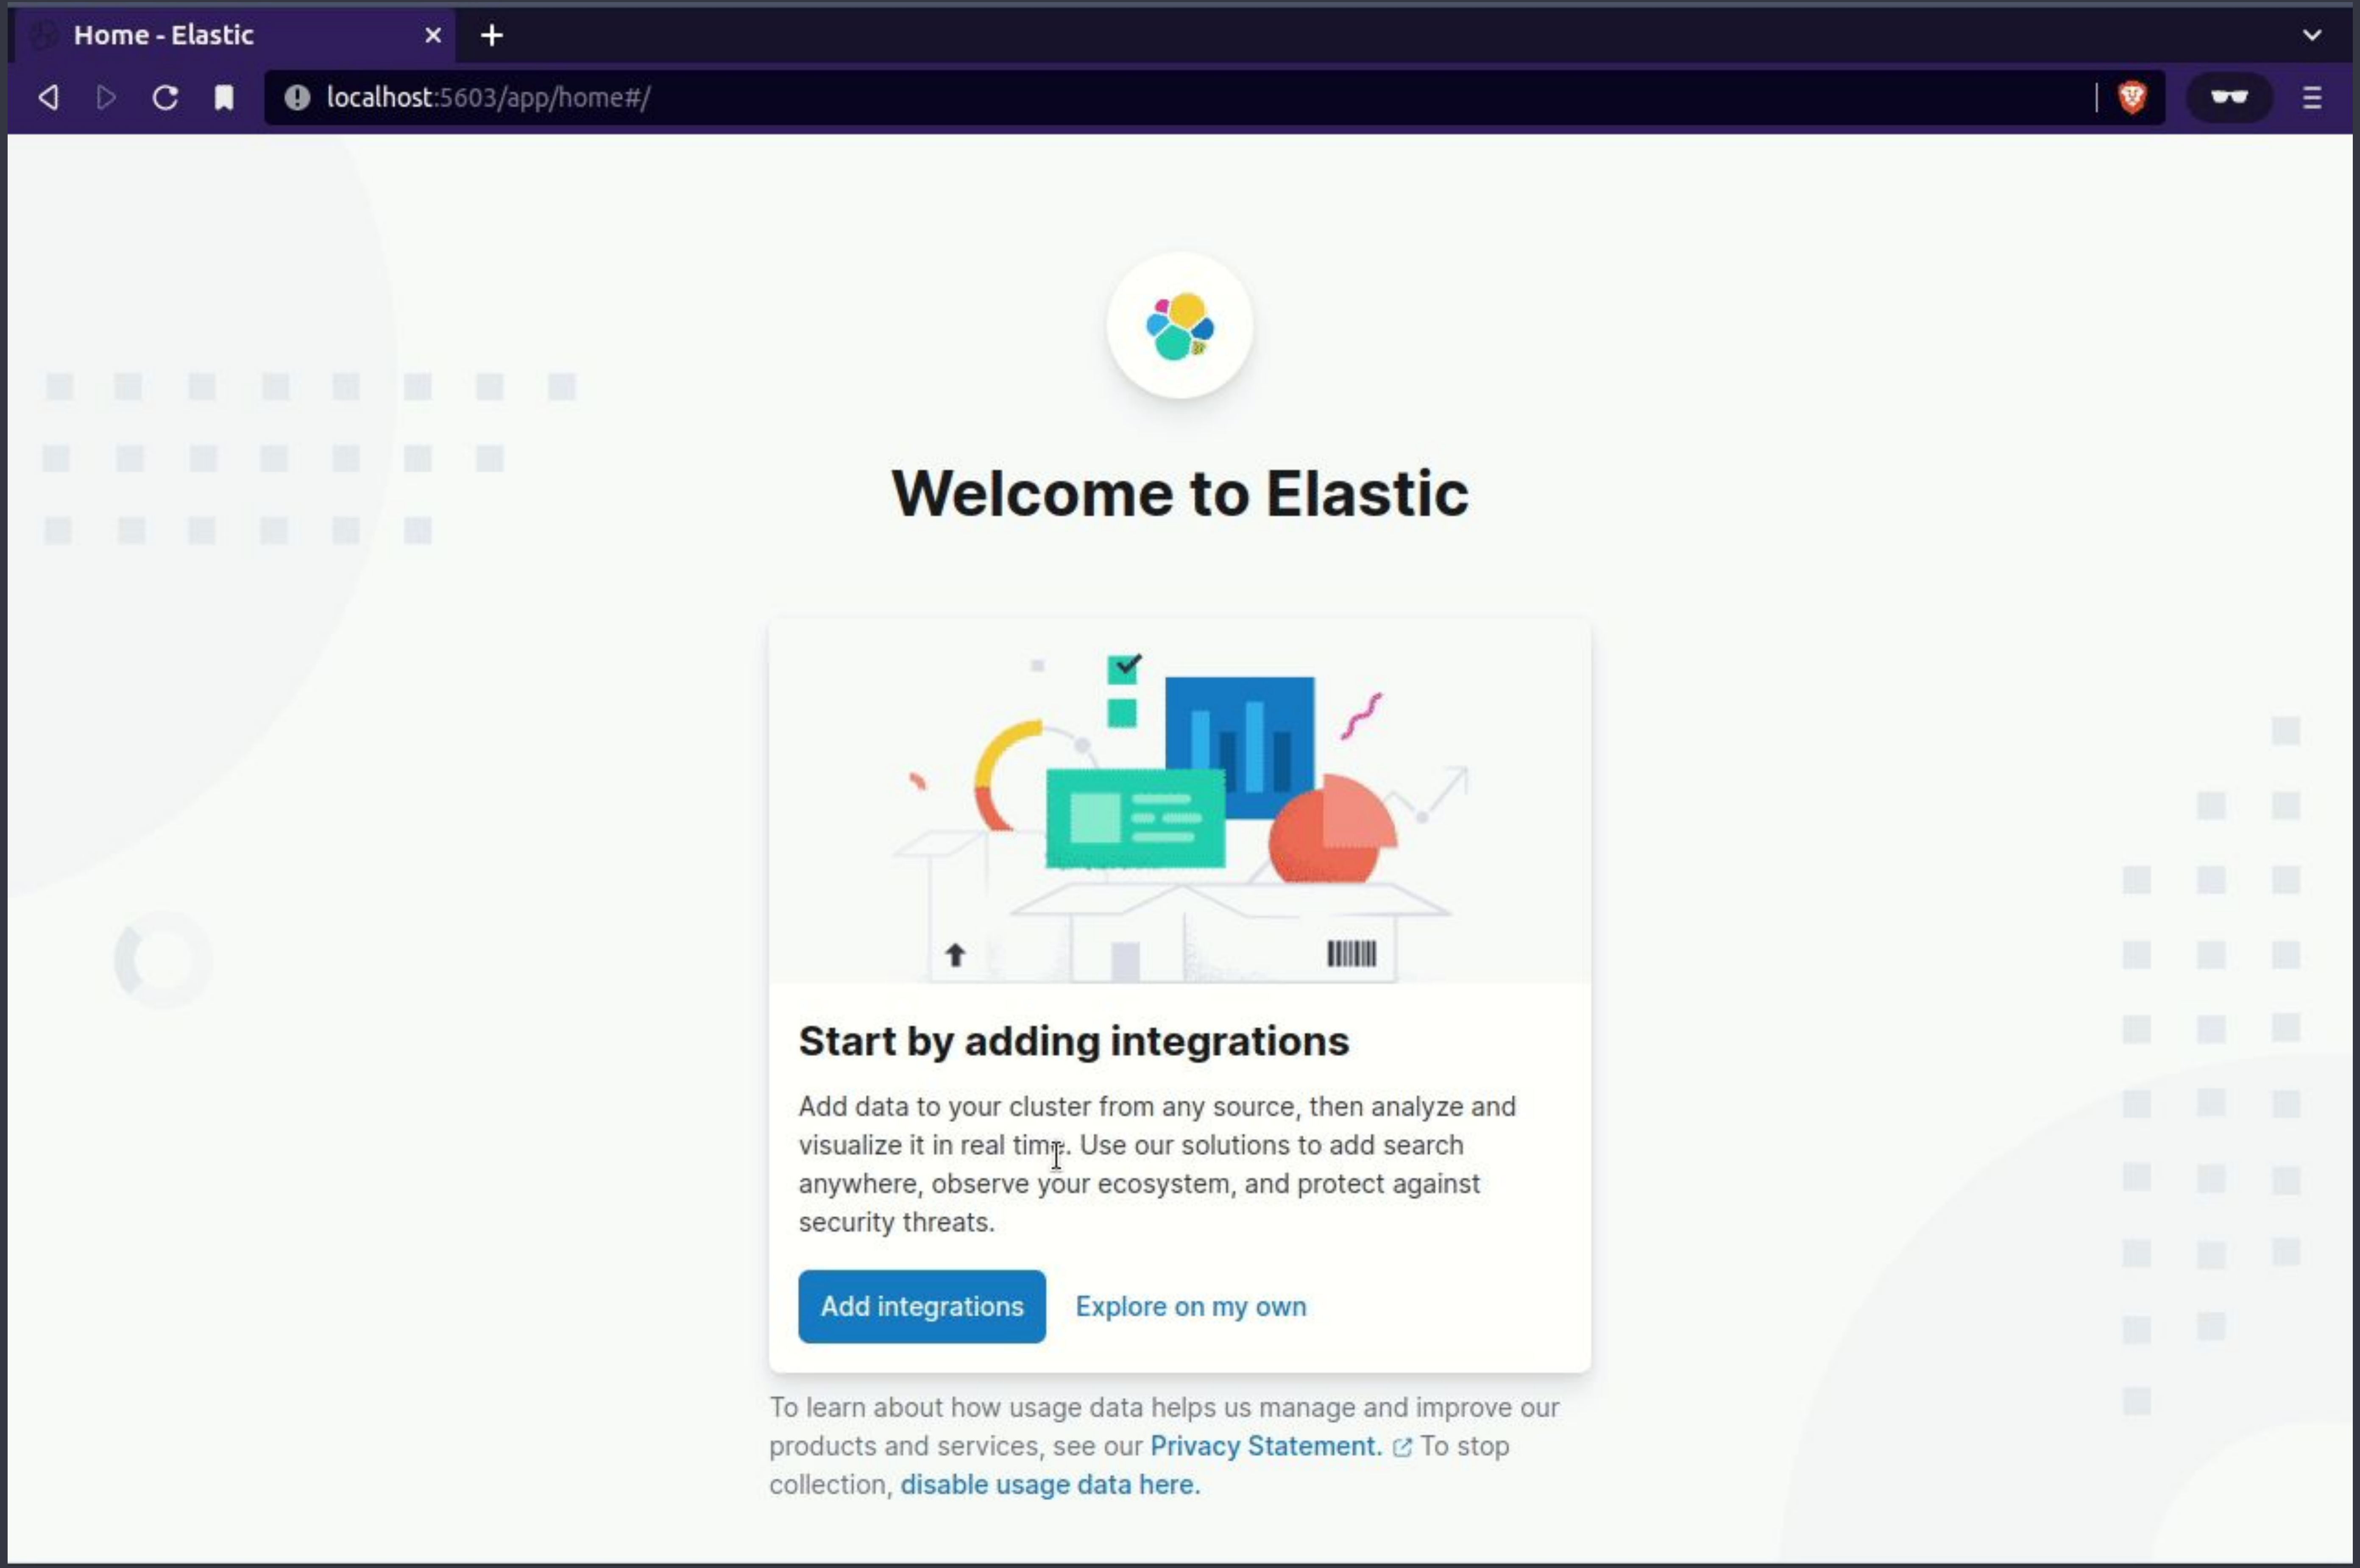

Welcome to Elastic

Start by adding integrations

Add data to your cluster from any source, then analyze and visualize it in real time. Use our solutions to add search anywhere, observe your ecosystem, and protect against security threats.

To learn about how usage data helps us manage and improve our products and services, see our Privacy Statement. To stop collection, disable usage data here.

Discover

- Quick **exploration** tool
- **Time range** and automatic **refresh***
- **Search bar** using Kibana Query Language or Lucene*
- **Filters***
- Table view with custom **columns**
- Field **statistics**
- **Inspect** tool: statistics, complete query and response
- **Save** your search to be used later on dashboards

* shared UI with other Kibana applications

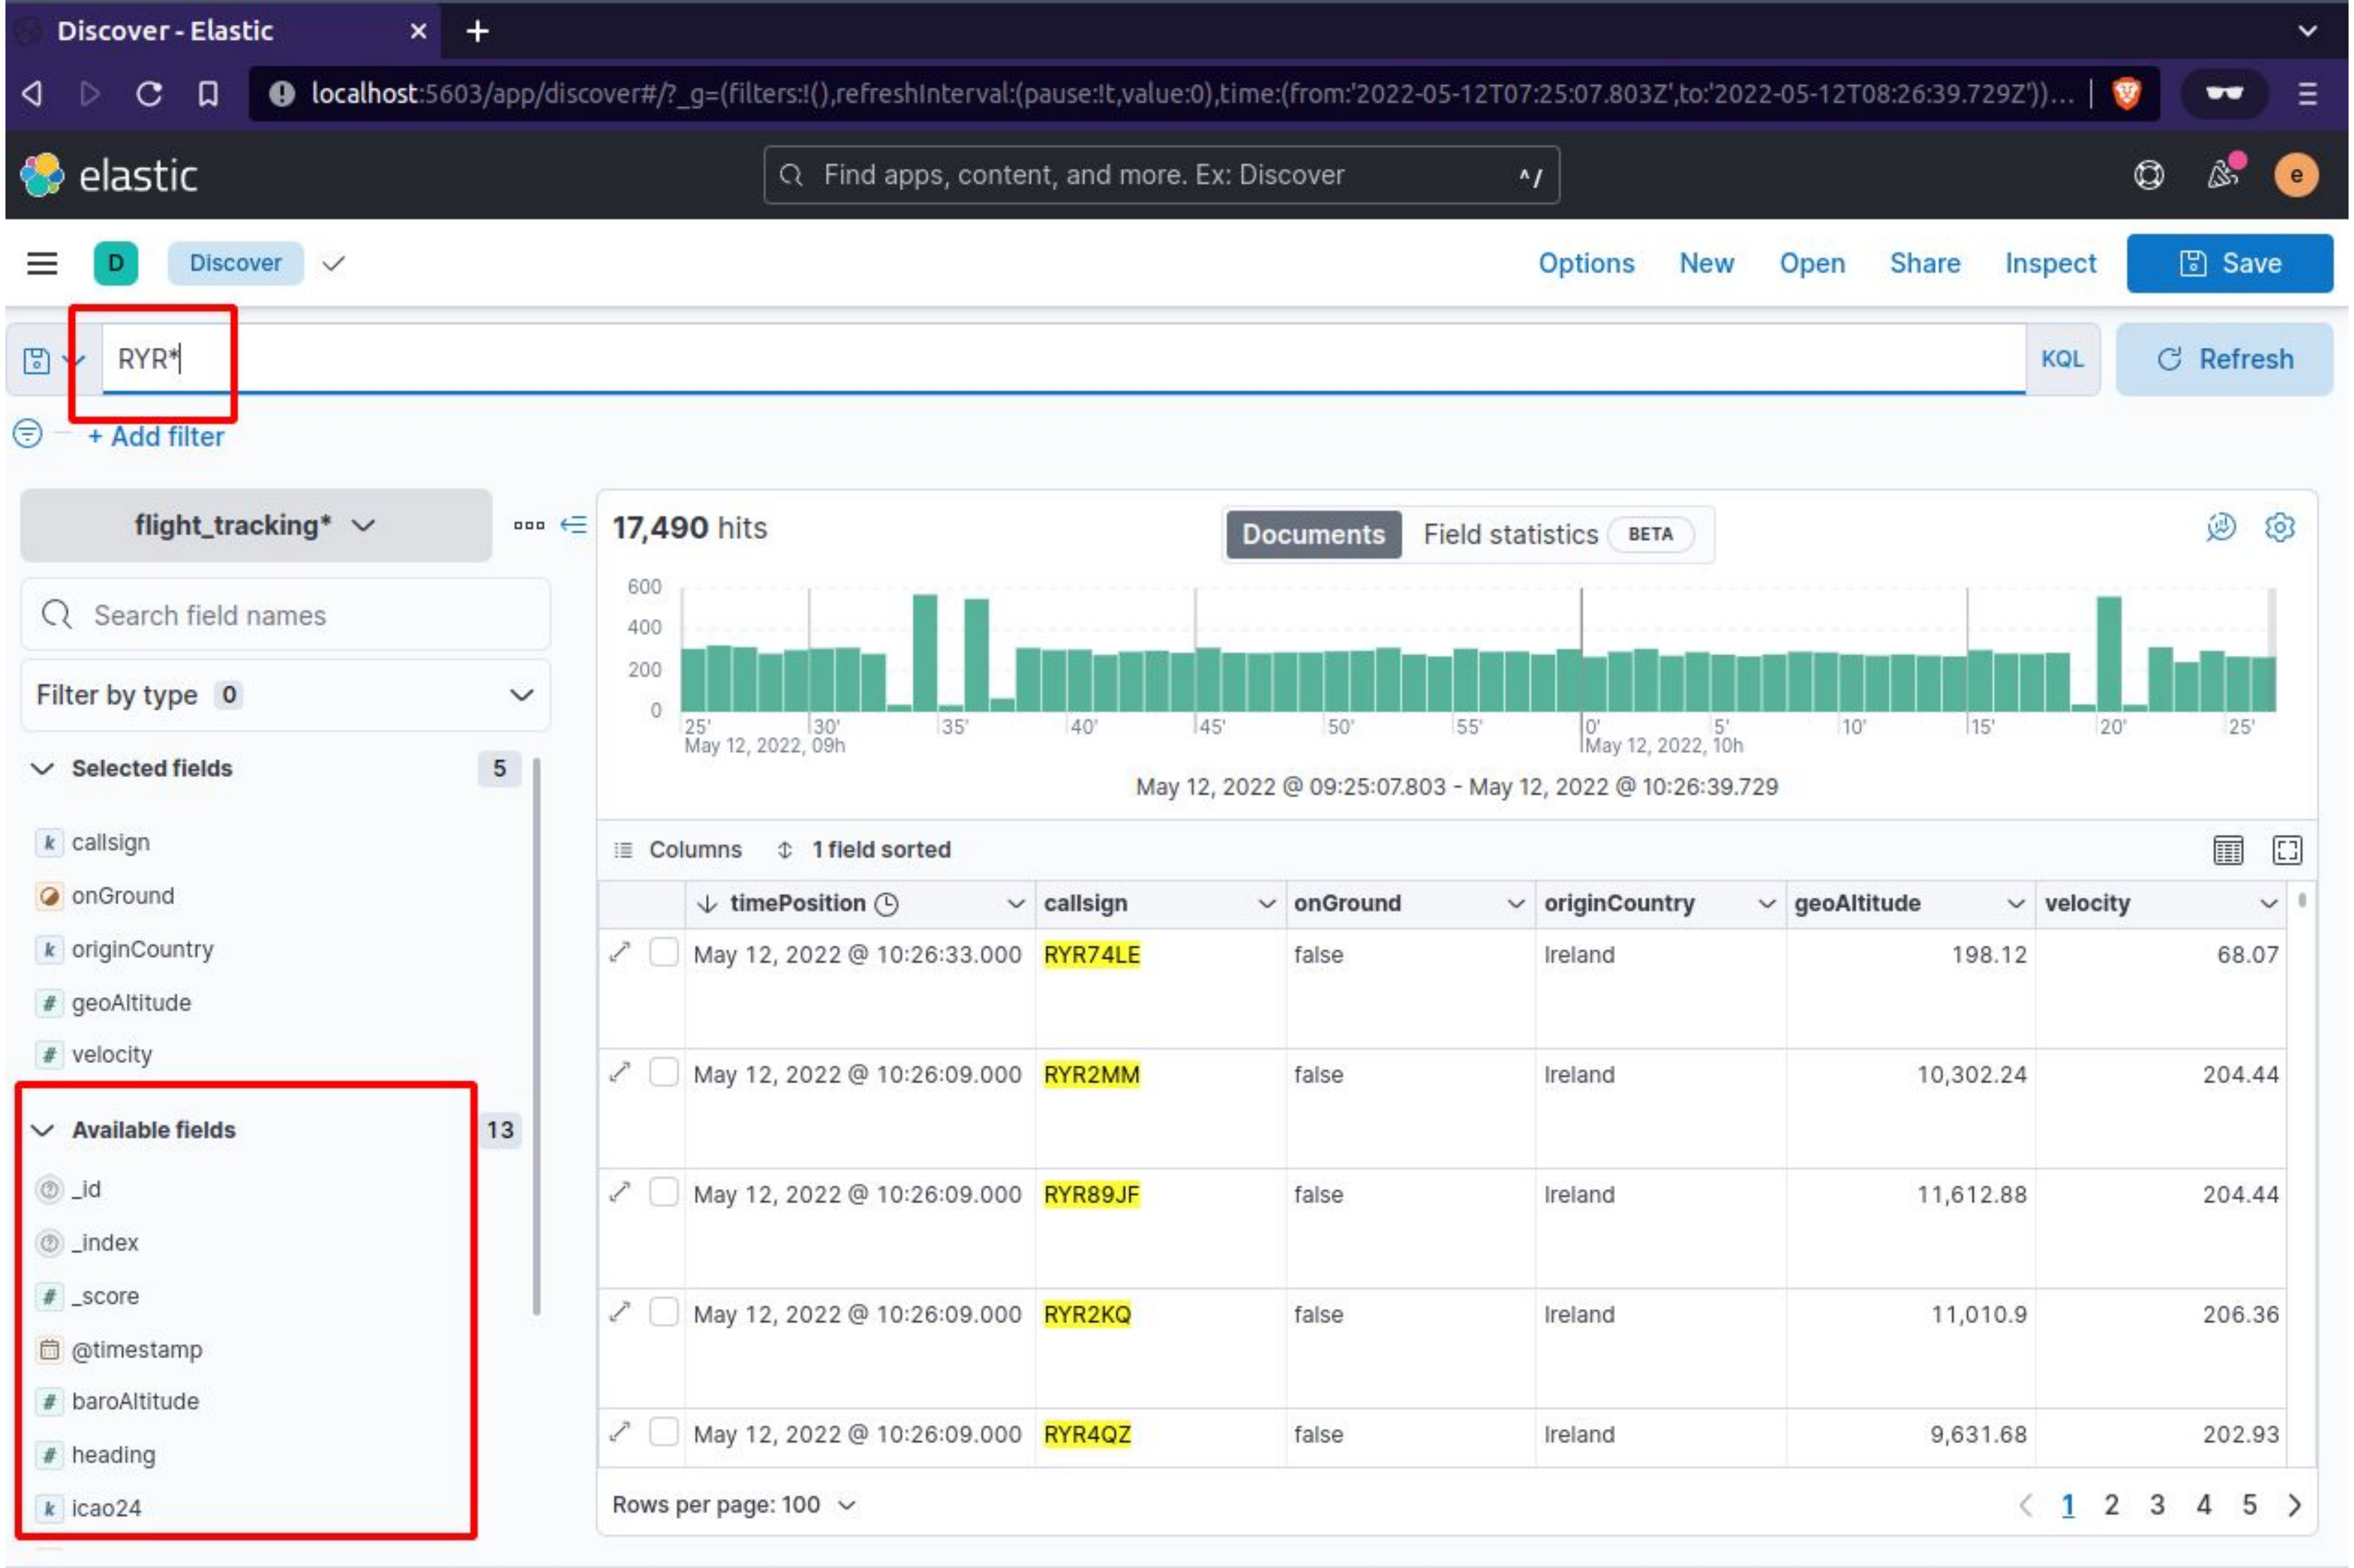

RYR*

17,490 hits

May 12, 2022 @ 09:25:07.803 - May 12, 2022 @ 10:26:39.729

Selected fields 5

- callsign
- onGround
- originCountry
- geoAltitude
- velocity

Available fields 13

- _id
- _index
- _score
- @timestamp
- baroAltitude
- heading
- icao24

timePosition	callsign	onGround	originCountry	geoAltitude	velocity
May 12, 2022 @ 10:26:33.000	RYR74LE	false	Ireland	198.12	68.07
May 12, 2022 @ 10:26:09.000	RYR2MM	false	Ireland	10,302.24	204.44
May 12, 2022 @ 10:26:09.000	RYR89JF	false	Ireland	11,612.88	204.44
May 12, 2022 @ 10:26:09.000	RYR2KQ	false	Ireland	11,010.9	206.36
May 12, 2022 @ 10:26:09.000	RYR4QZ	false	Ireland	9,631.68	202.93

Lens

Your data in front of you

Explore your fields with a single click

Drag and drop

Go from nothing to visual insights with a single mouse gesture.

Smart suggestions

Let Lens help guide your analysis with useful chart suggestions

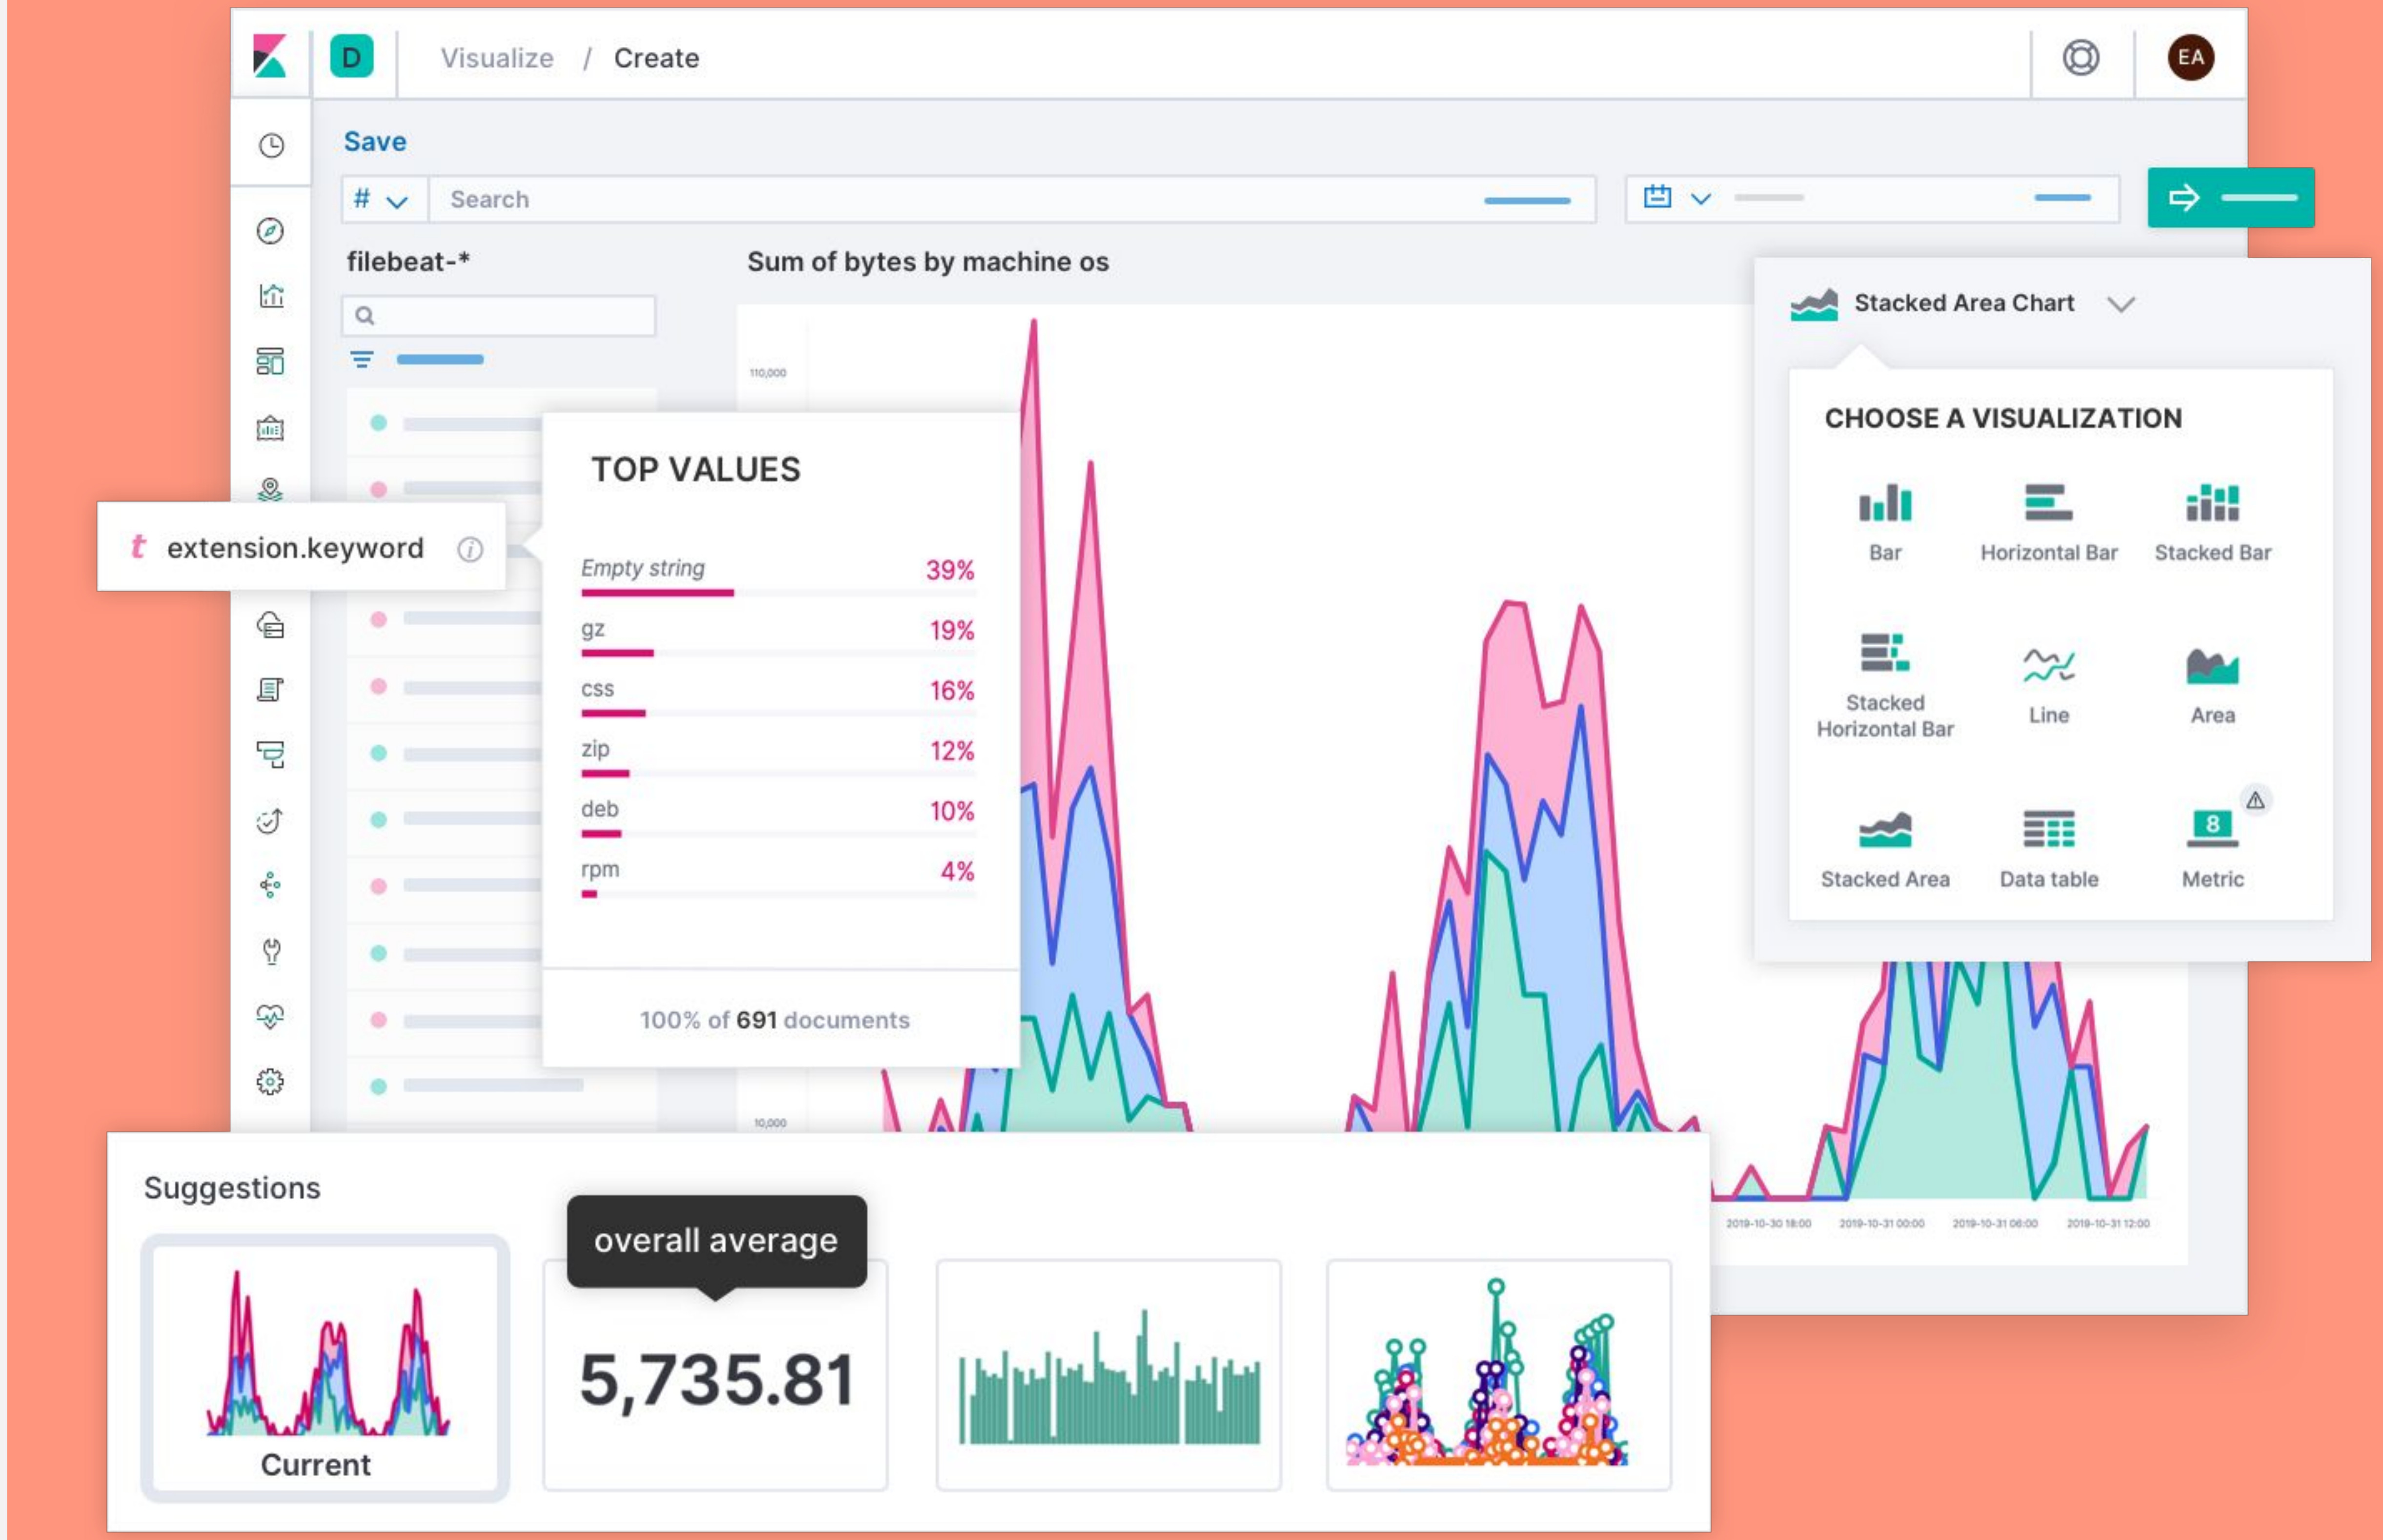

TOP VALUES

Empty string	39%
gz	19%
css	16%
zip	12%
deb	10%
rpm	4%

100% of **691** documents

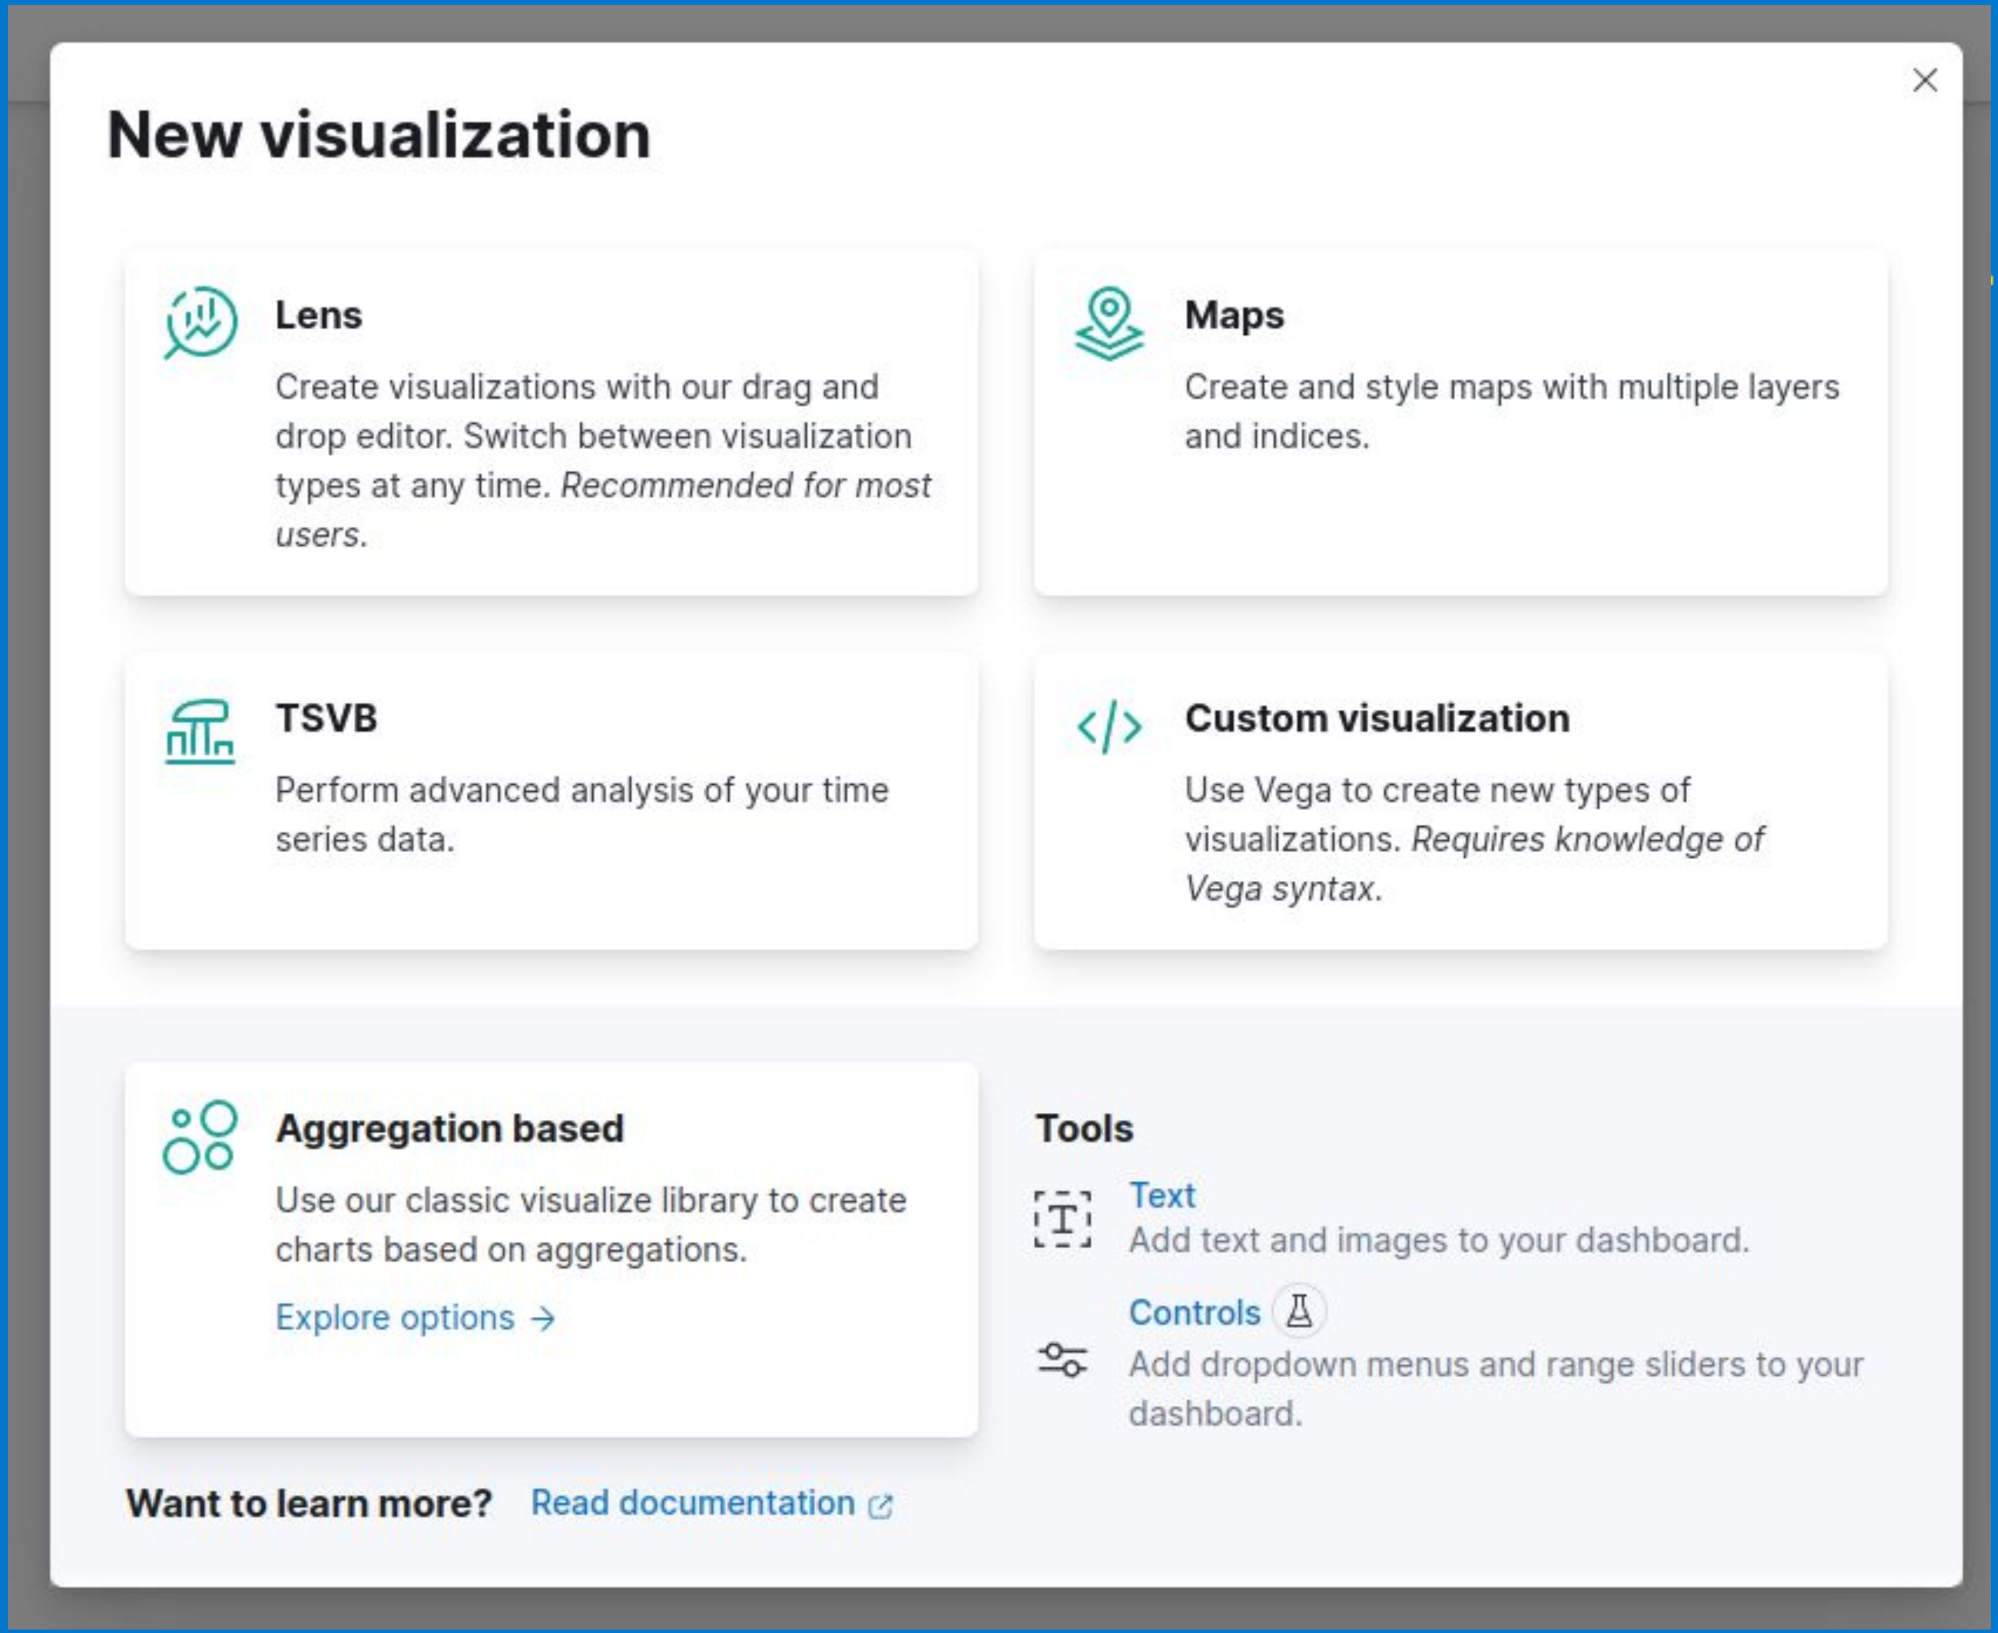

New visualization

Lens

Create visualizations with our drag and drop editor. Switch between visualization types at any time. *Recommended for most users.*

Maps

Create and style maps with multiple layers and indices.

TSVB

Perform advanced analysis of your time series data.

Custom visualization

Use Vega to create new types of visualizations. *Requires knowledge of Vega syntax.*

Aggregation based

Use our classic visualize library to create charts based on aggregations.

Explore options →

Tools

Text
Add text and images to your dashboard.

Controls
Add dropdown menus and range sliders to your dashboard.

Want to learn more? Read documentation

… and more

Dashboards

- Combine multiple visualizations: **panels**
- Time Range + Search Bar + Filters
- Panels can use filters to perform **drill downs**
- Panels can have **custom** time ranges
- **Share**
- **Export** to PDF or PNG

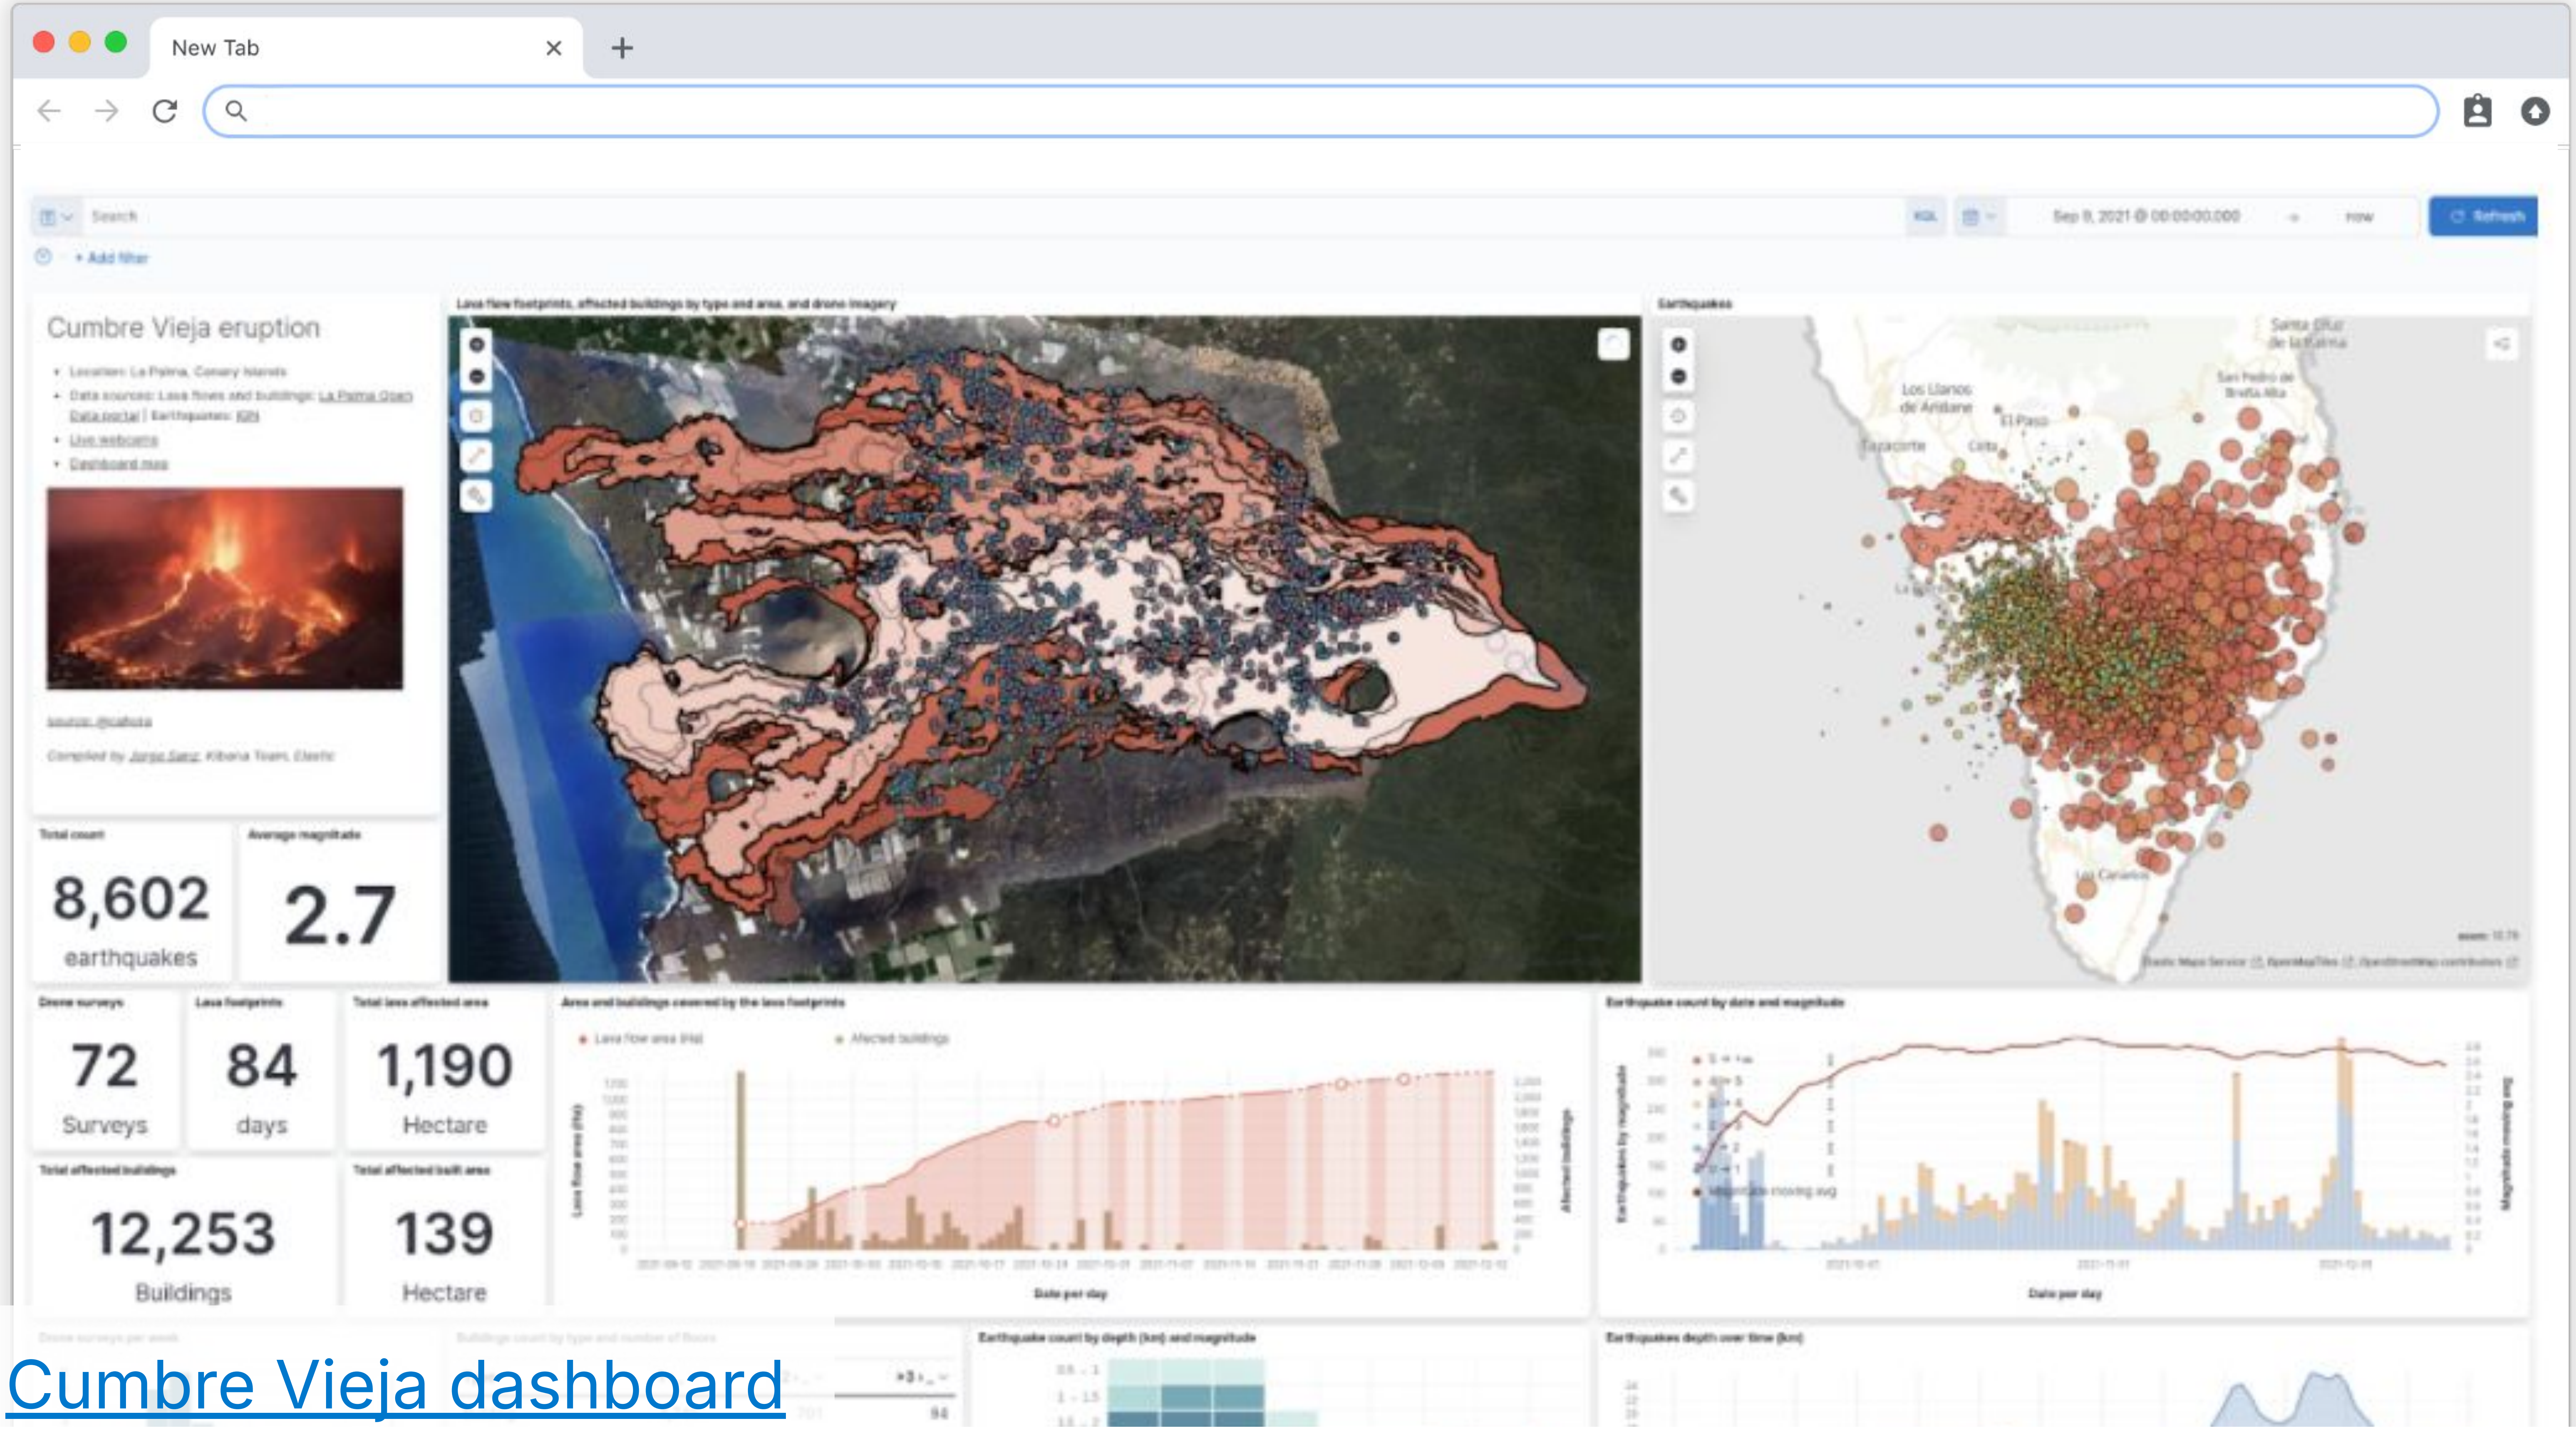

Cumbre Vieja eruption

- Location: La Palma, Canary Islands
- Data sources: Lava flows and buildings: La Palma Open Data portal | Earthquakes: IGN
- Live webcams
- Dashboard map

source: @cahora

Compiled by Jorge Sanz, Kibana Team, Elastic

Total count: 8,602 earthquakes

Average magnitude: 2.7

Drone surveys: 72 Surveys

Lava footprints: 84 days

Total lava affected area: 1,190 Hectare

Total affected buildings: 12,253 Buildings

Total affected built area: 139 Hectare

[Cumbre Vieja dashboard](#)

Ingest with Kibana: CSV file upload

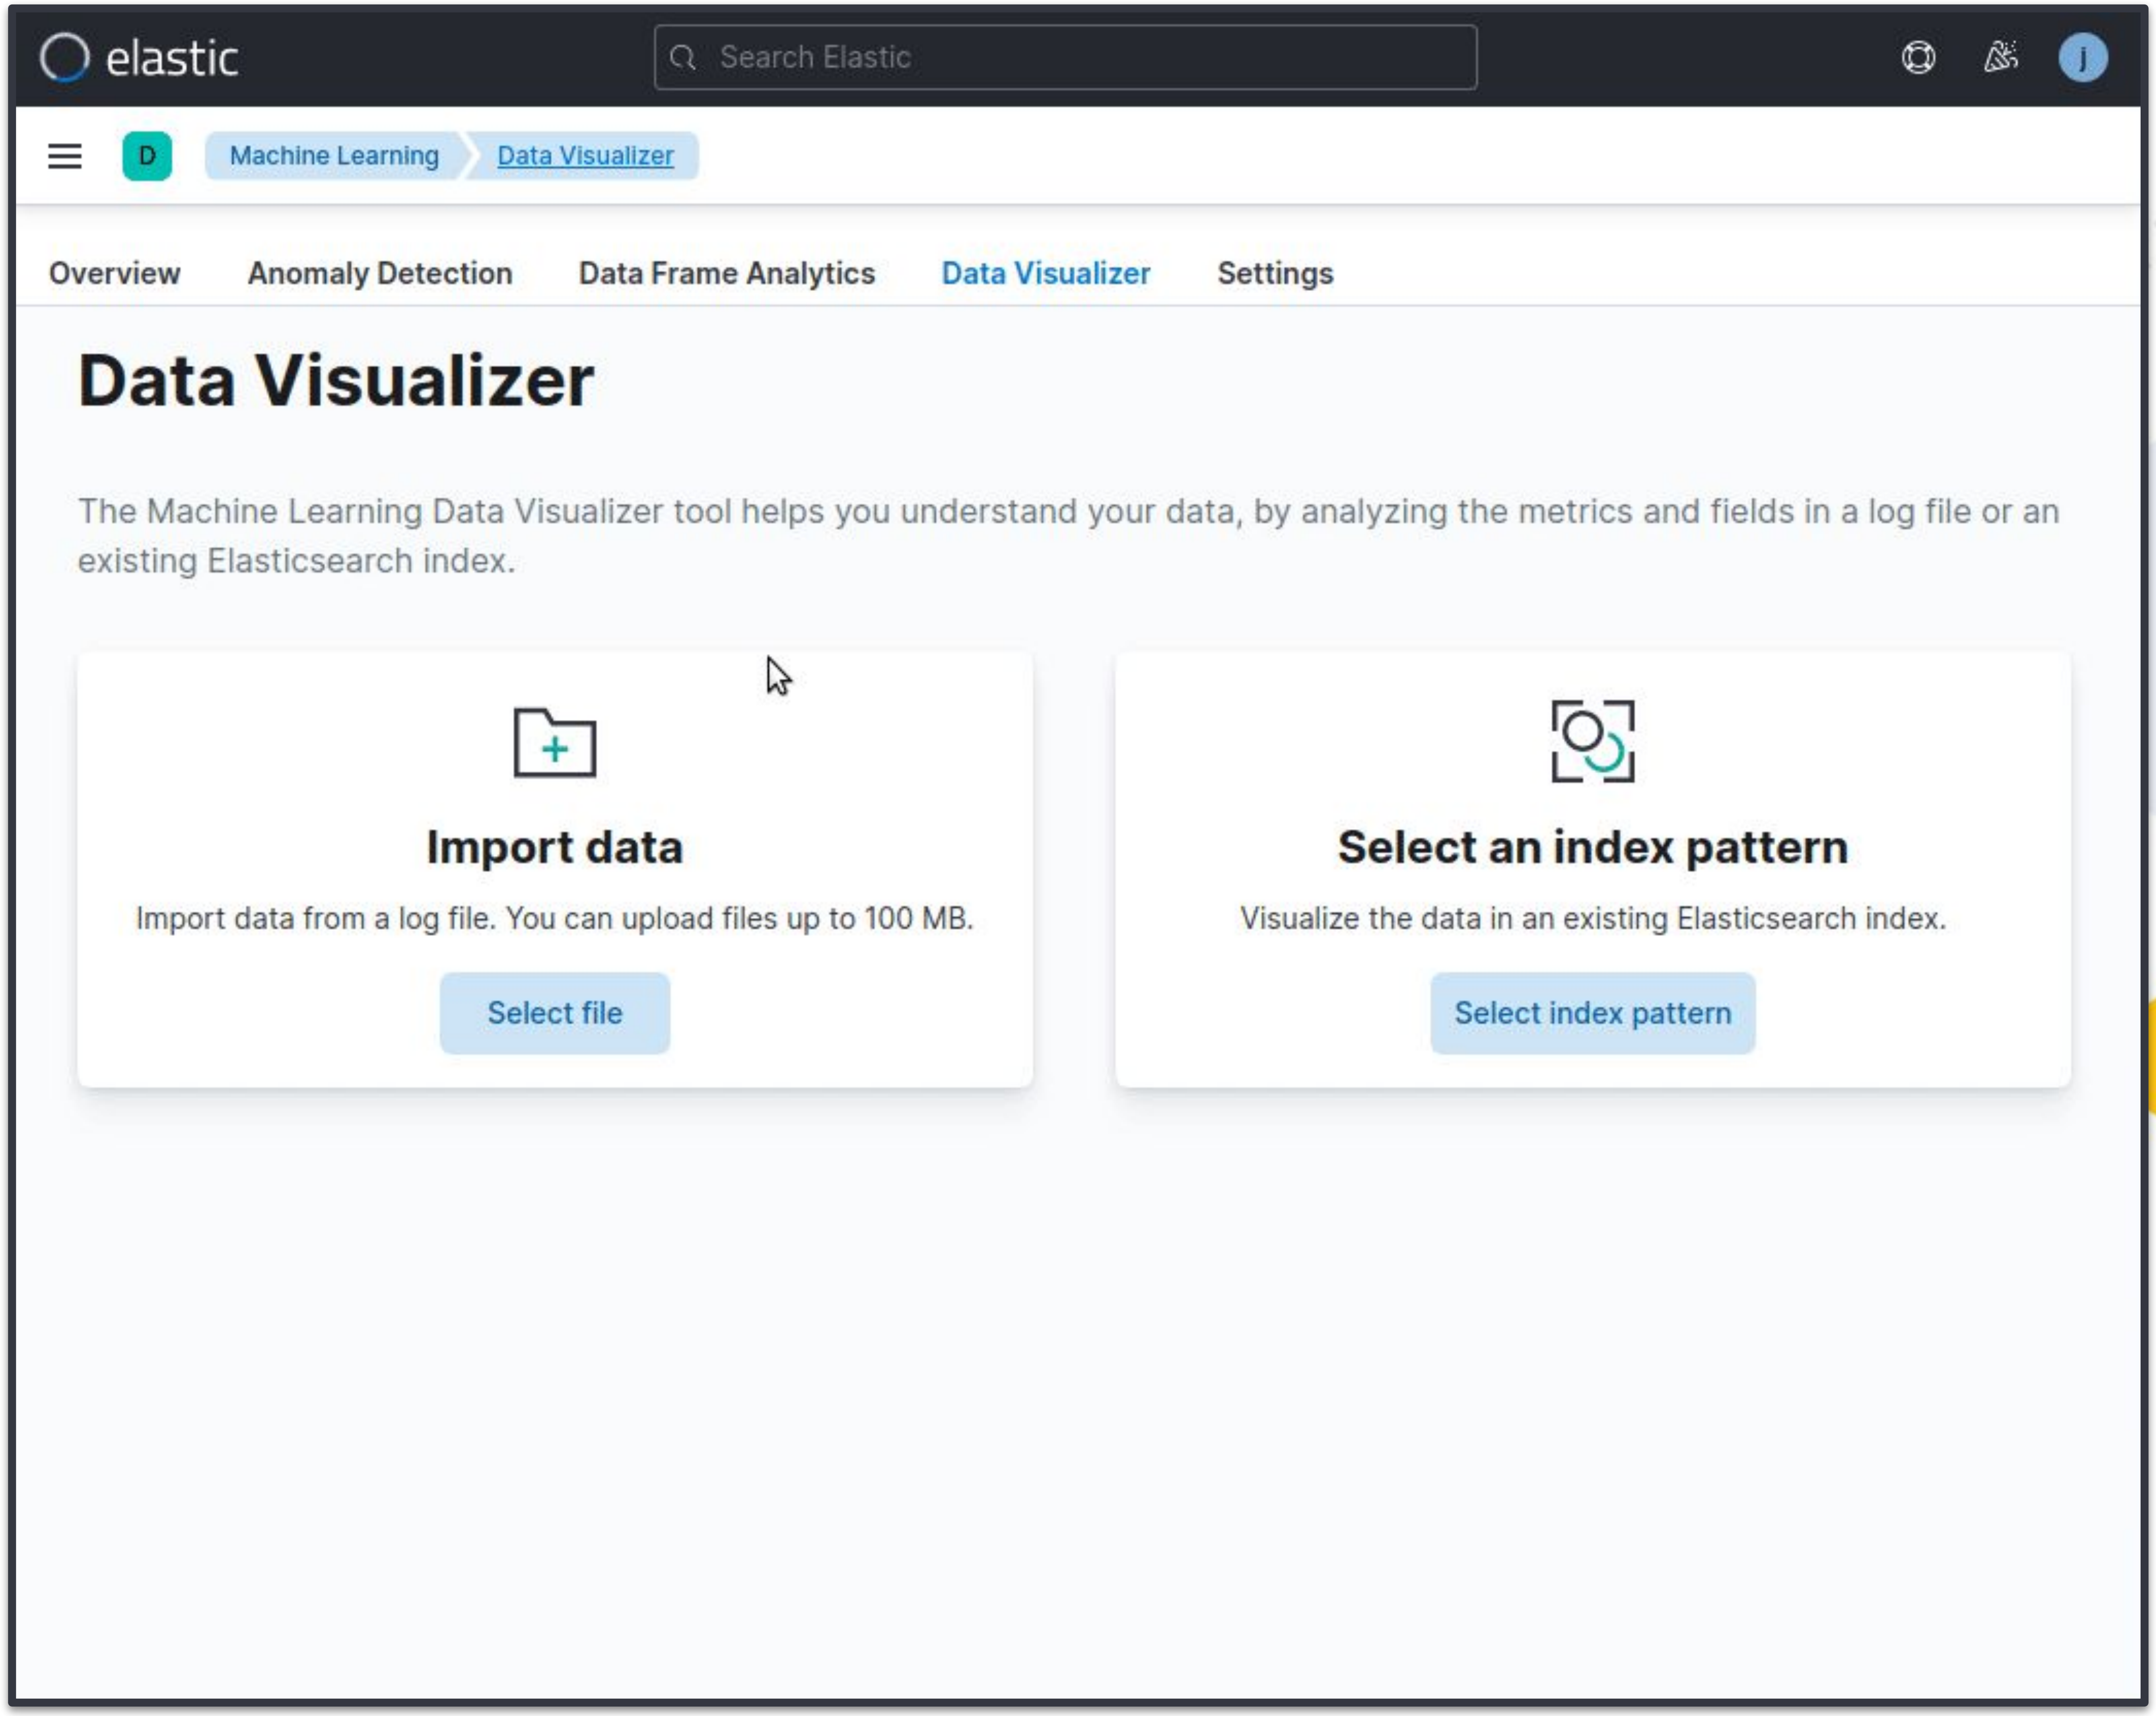

Data Visualizer

The Machine Learning Data Visualizer tool helps you understand your data, by analyzing the metrics and fields in a log file or an existing Elasticsearch index.

Import data

Import data from a log file. You can upload files up to 100 MB.

Select an index pattern

Visualize the data in an existing Elasticsearch index.

Agenda

1. Introduction to the Elastic stack
2. Lab setup
3. Kibana introduction
4. Elastics Maps
5. Elasticsearch Geo
6. Web mapping and Elasticsearch

Elastic Maps

Kibana approach to Geographical Information Systems

Elastic Maps

OOTB Geo Analytics interface within Kibana

- Friendly user experience
- Aggregations: heat map, clustering, grids, geoline
- Data driven styling
- Tools for drawing, filtering, measuring
- Add layers from external tile servers
- Used alone or in dashboards or Canvas workpads
- Embedded in other Kibana solution applications

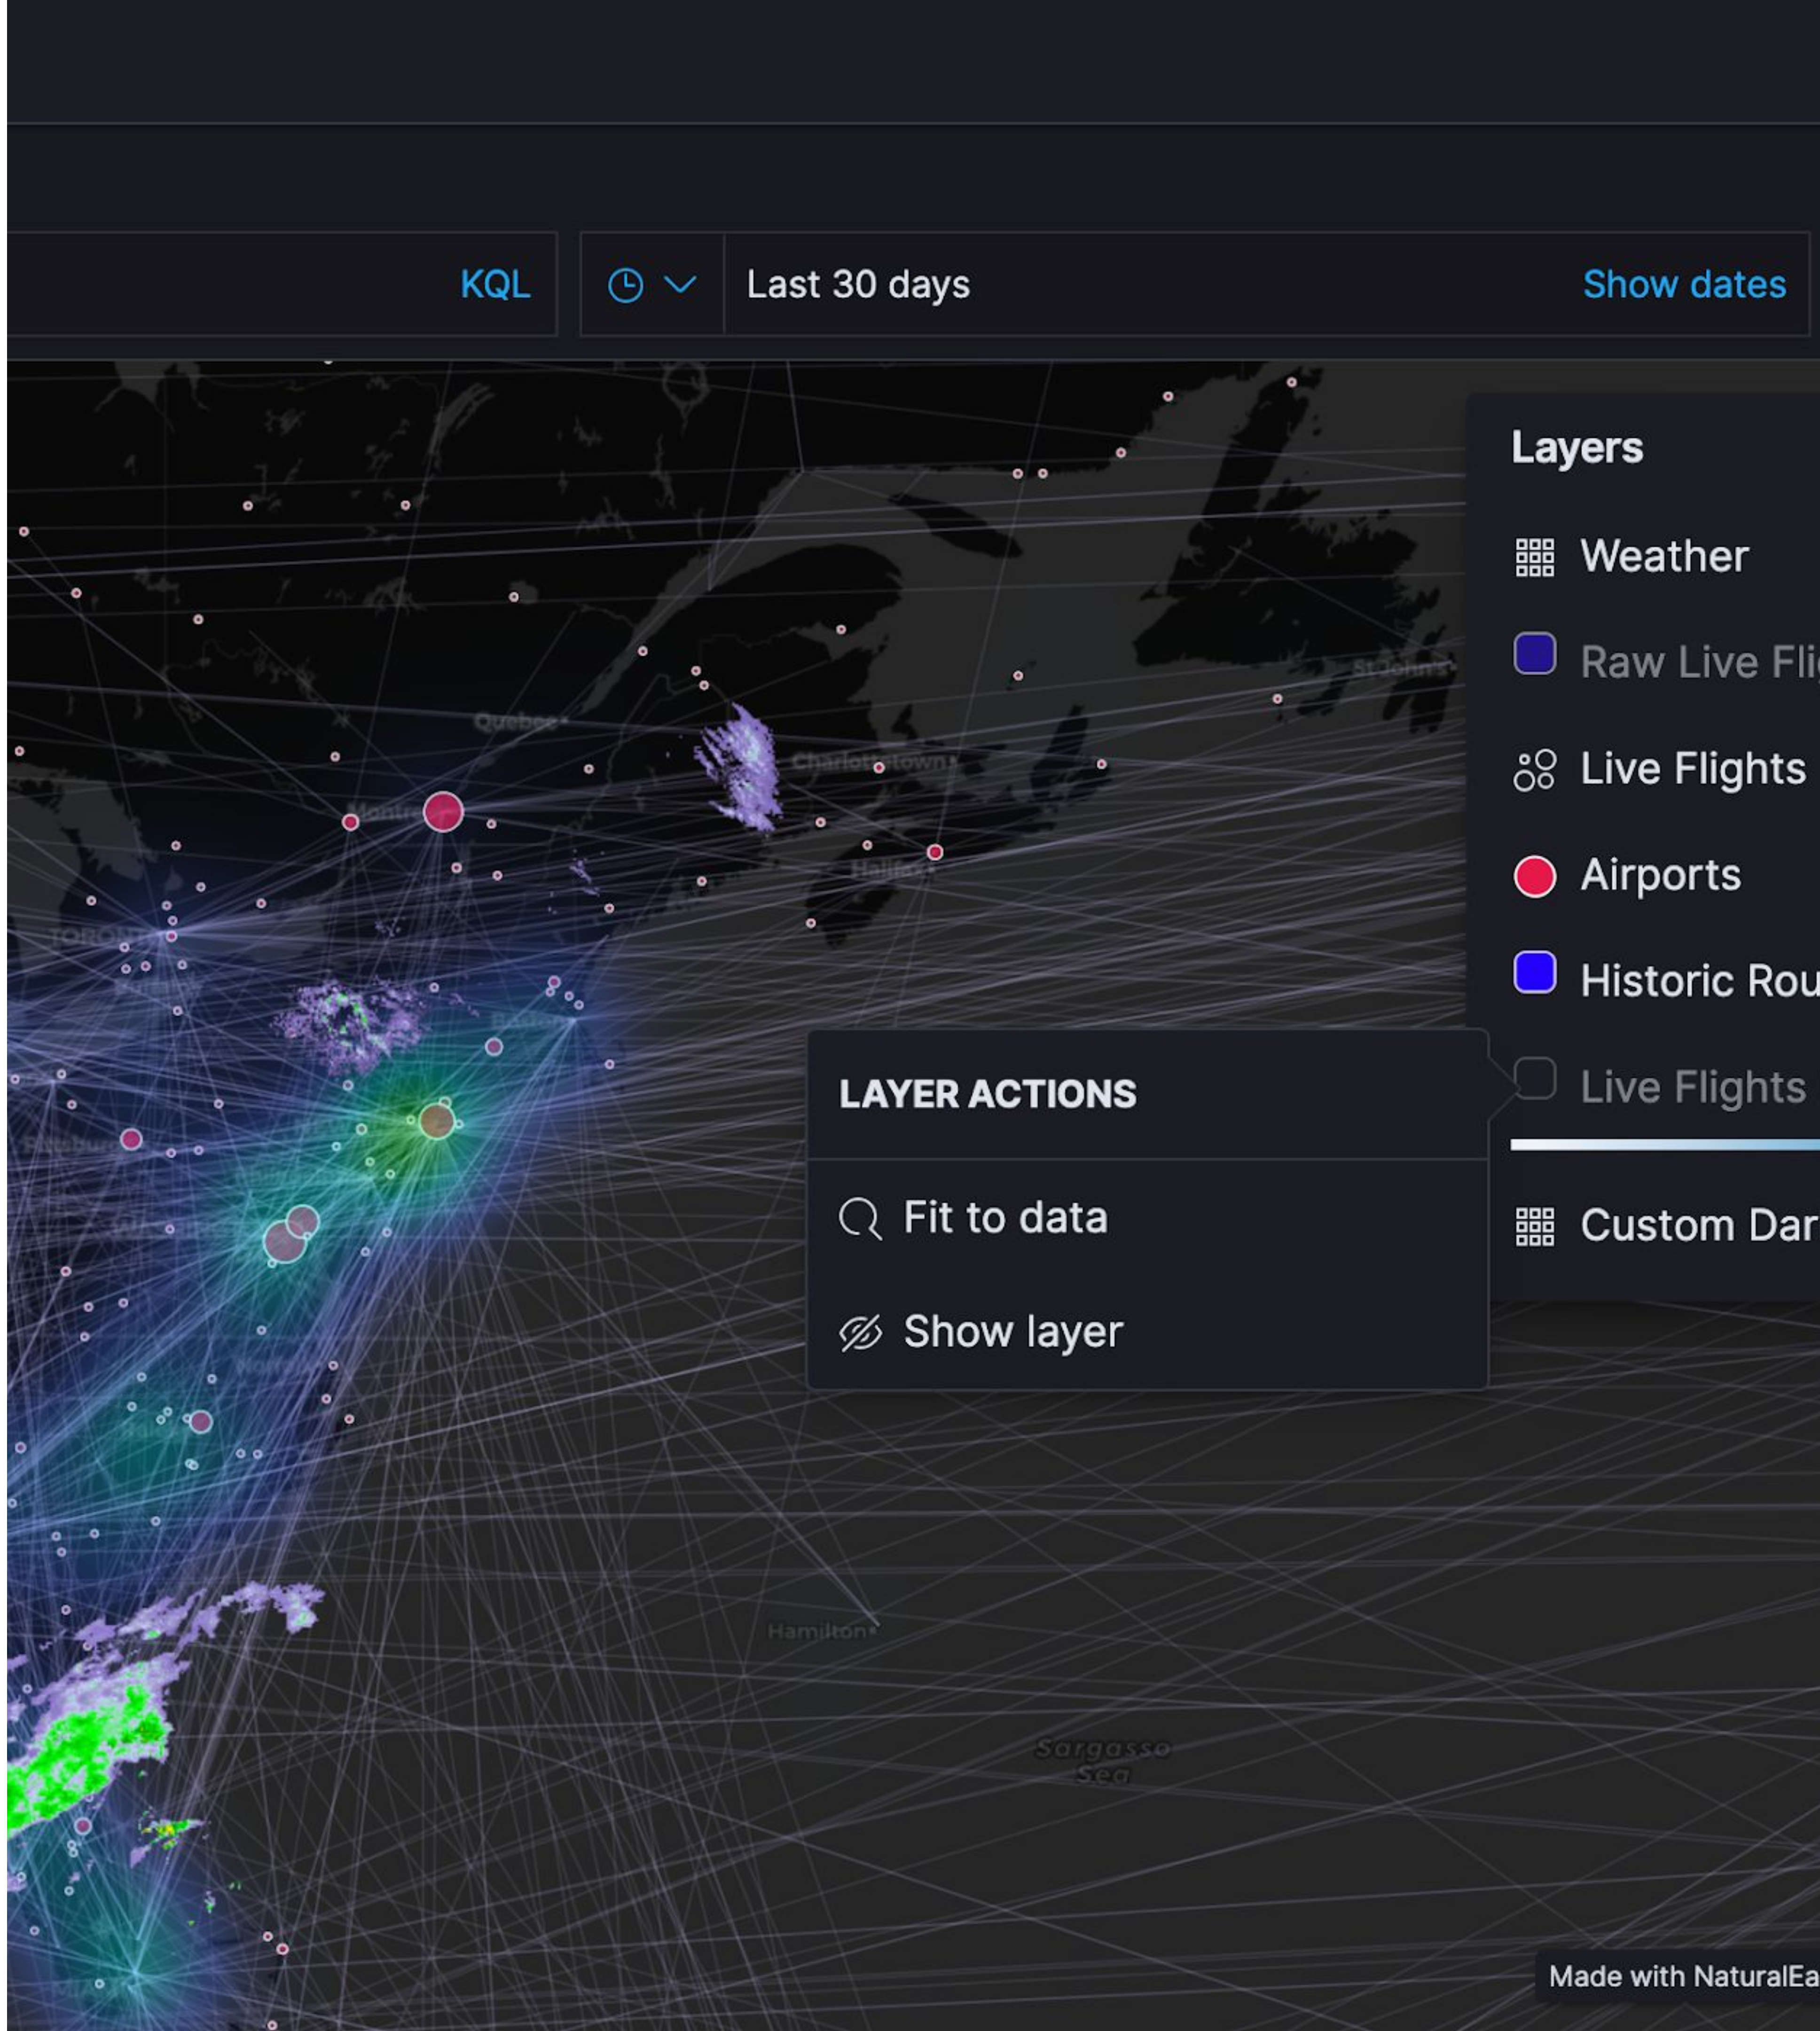

Elastic Maps Service

maps.elastic.co

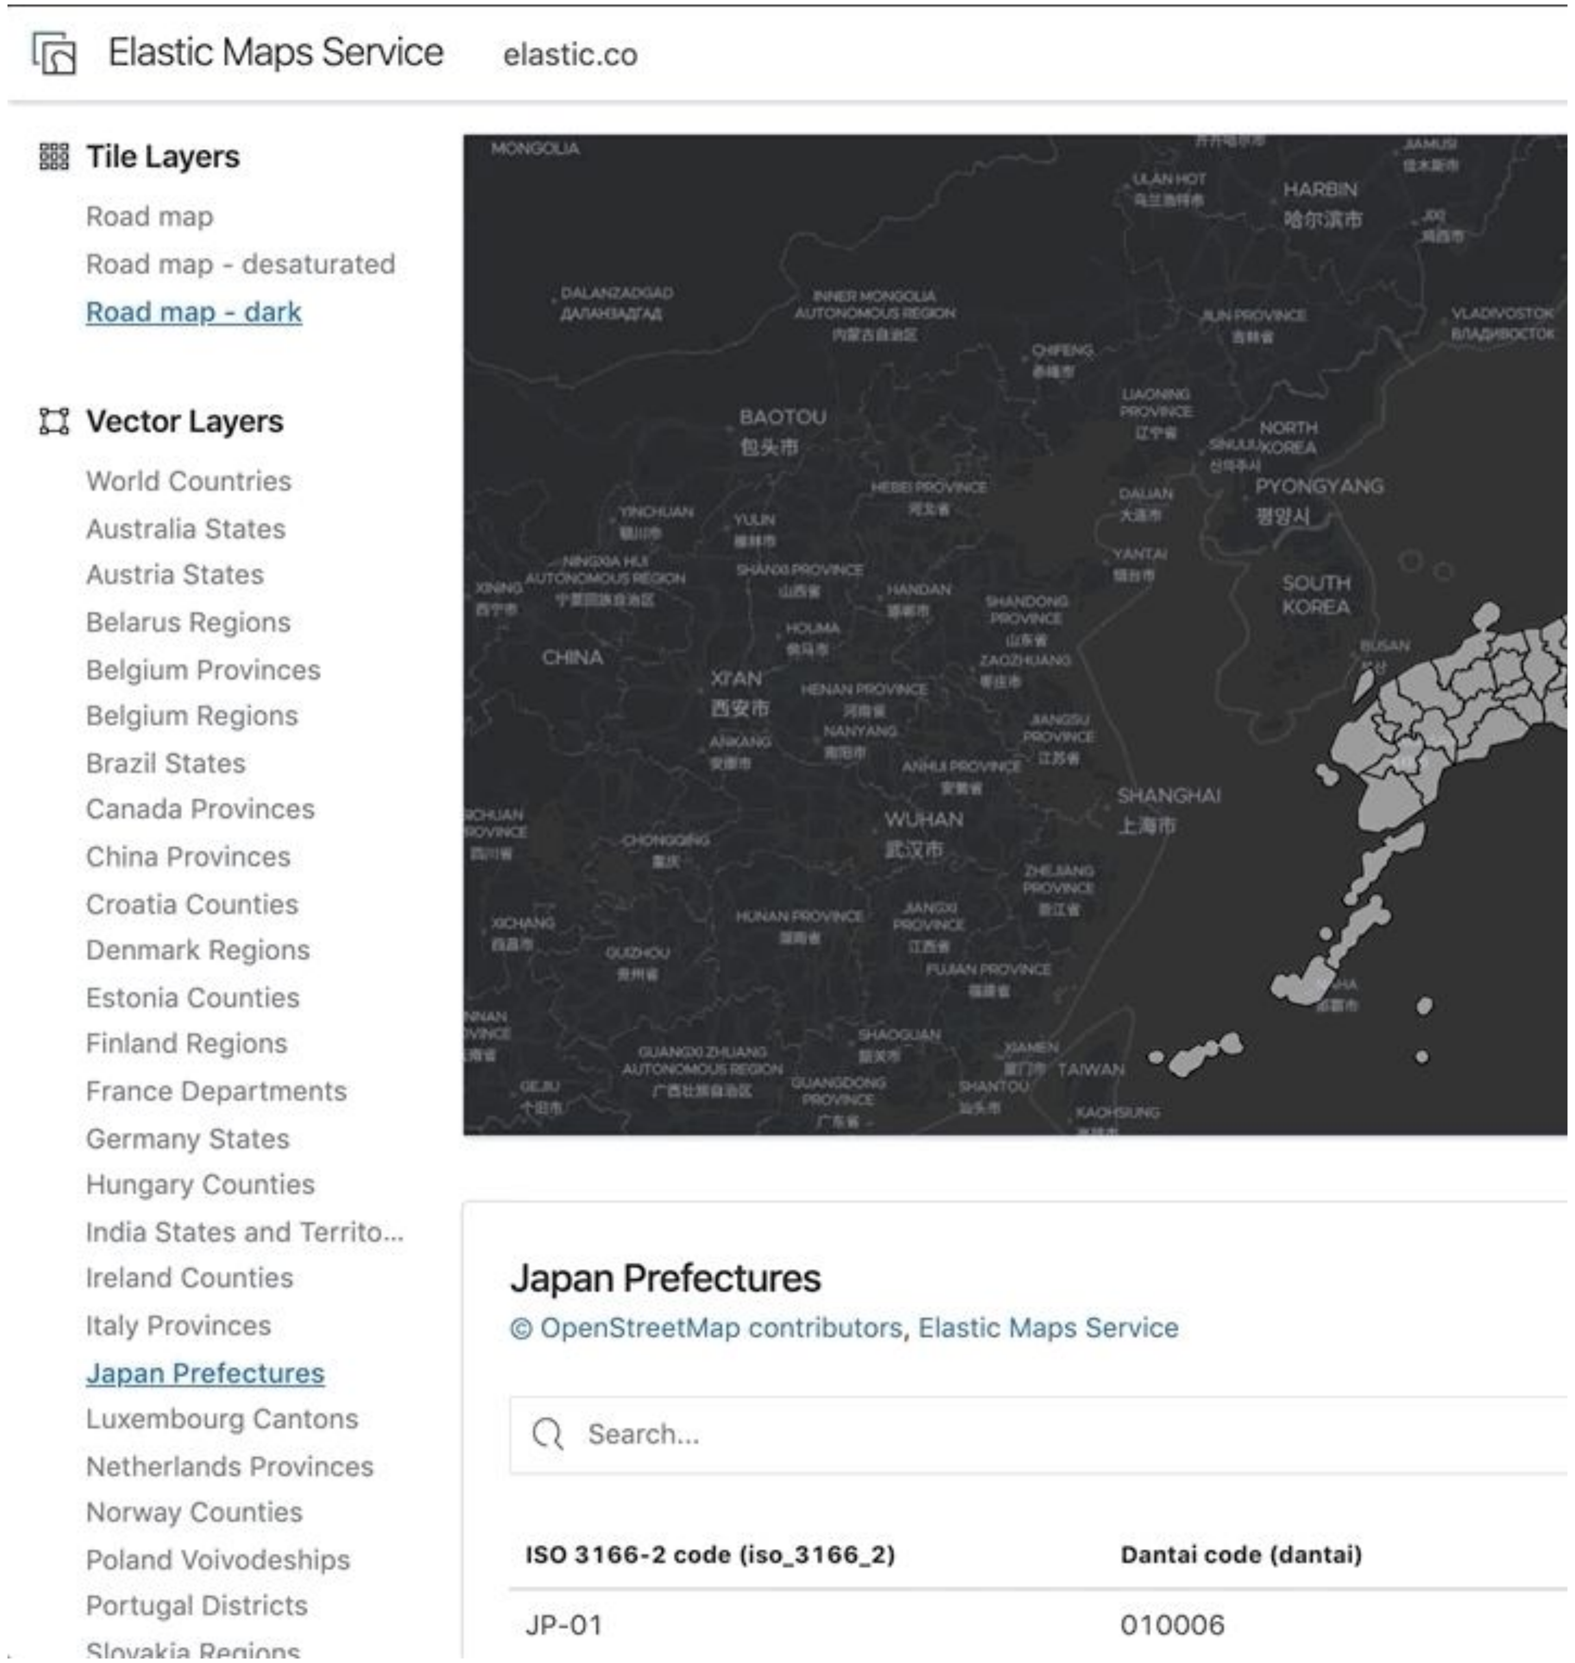

Japan Prefectures

© OpenStreetMap contributors, Elastic Maps Service

ISO 3166-2 code (iso_3166_2)	Dantai code (dantai)
JP-01	010006

- Based in OSM and OpenMapTiles
- 18 zoom levels worldwide
- Three stiles: dark, light, classic
- Administrative boundaries

Ingest with Kibana: GeoJSON and Shapefile upload

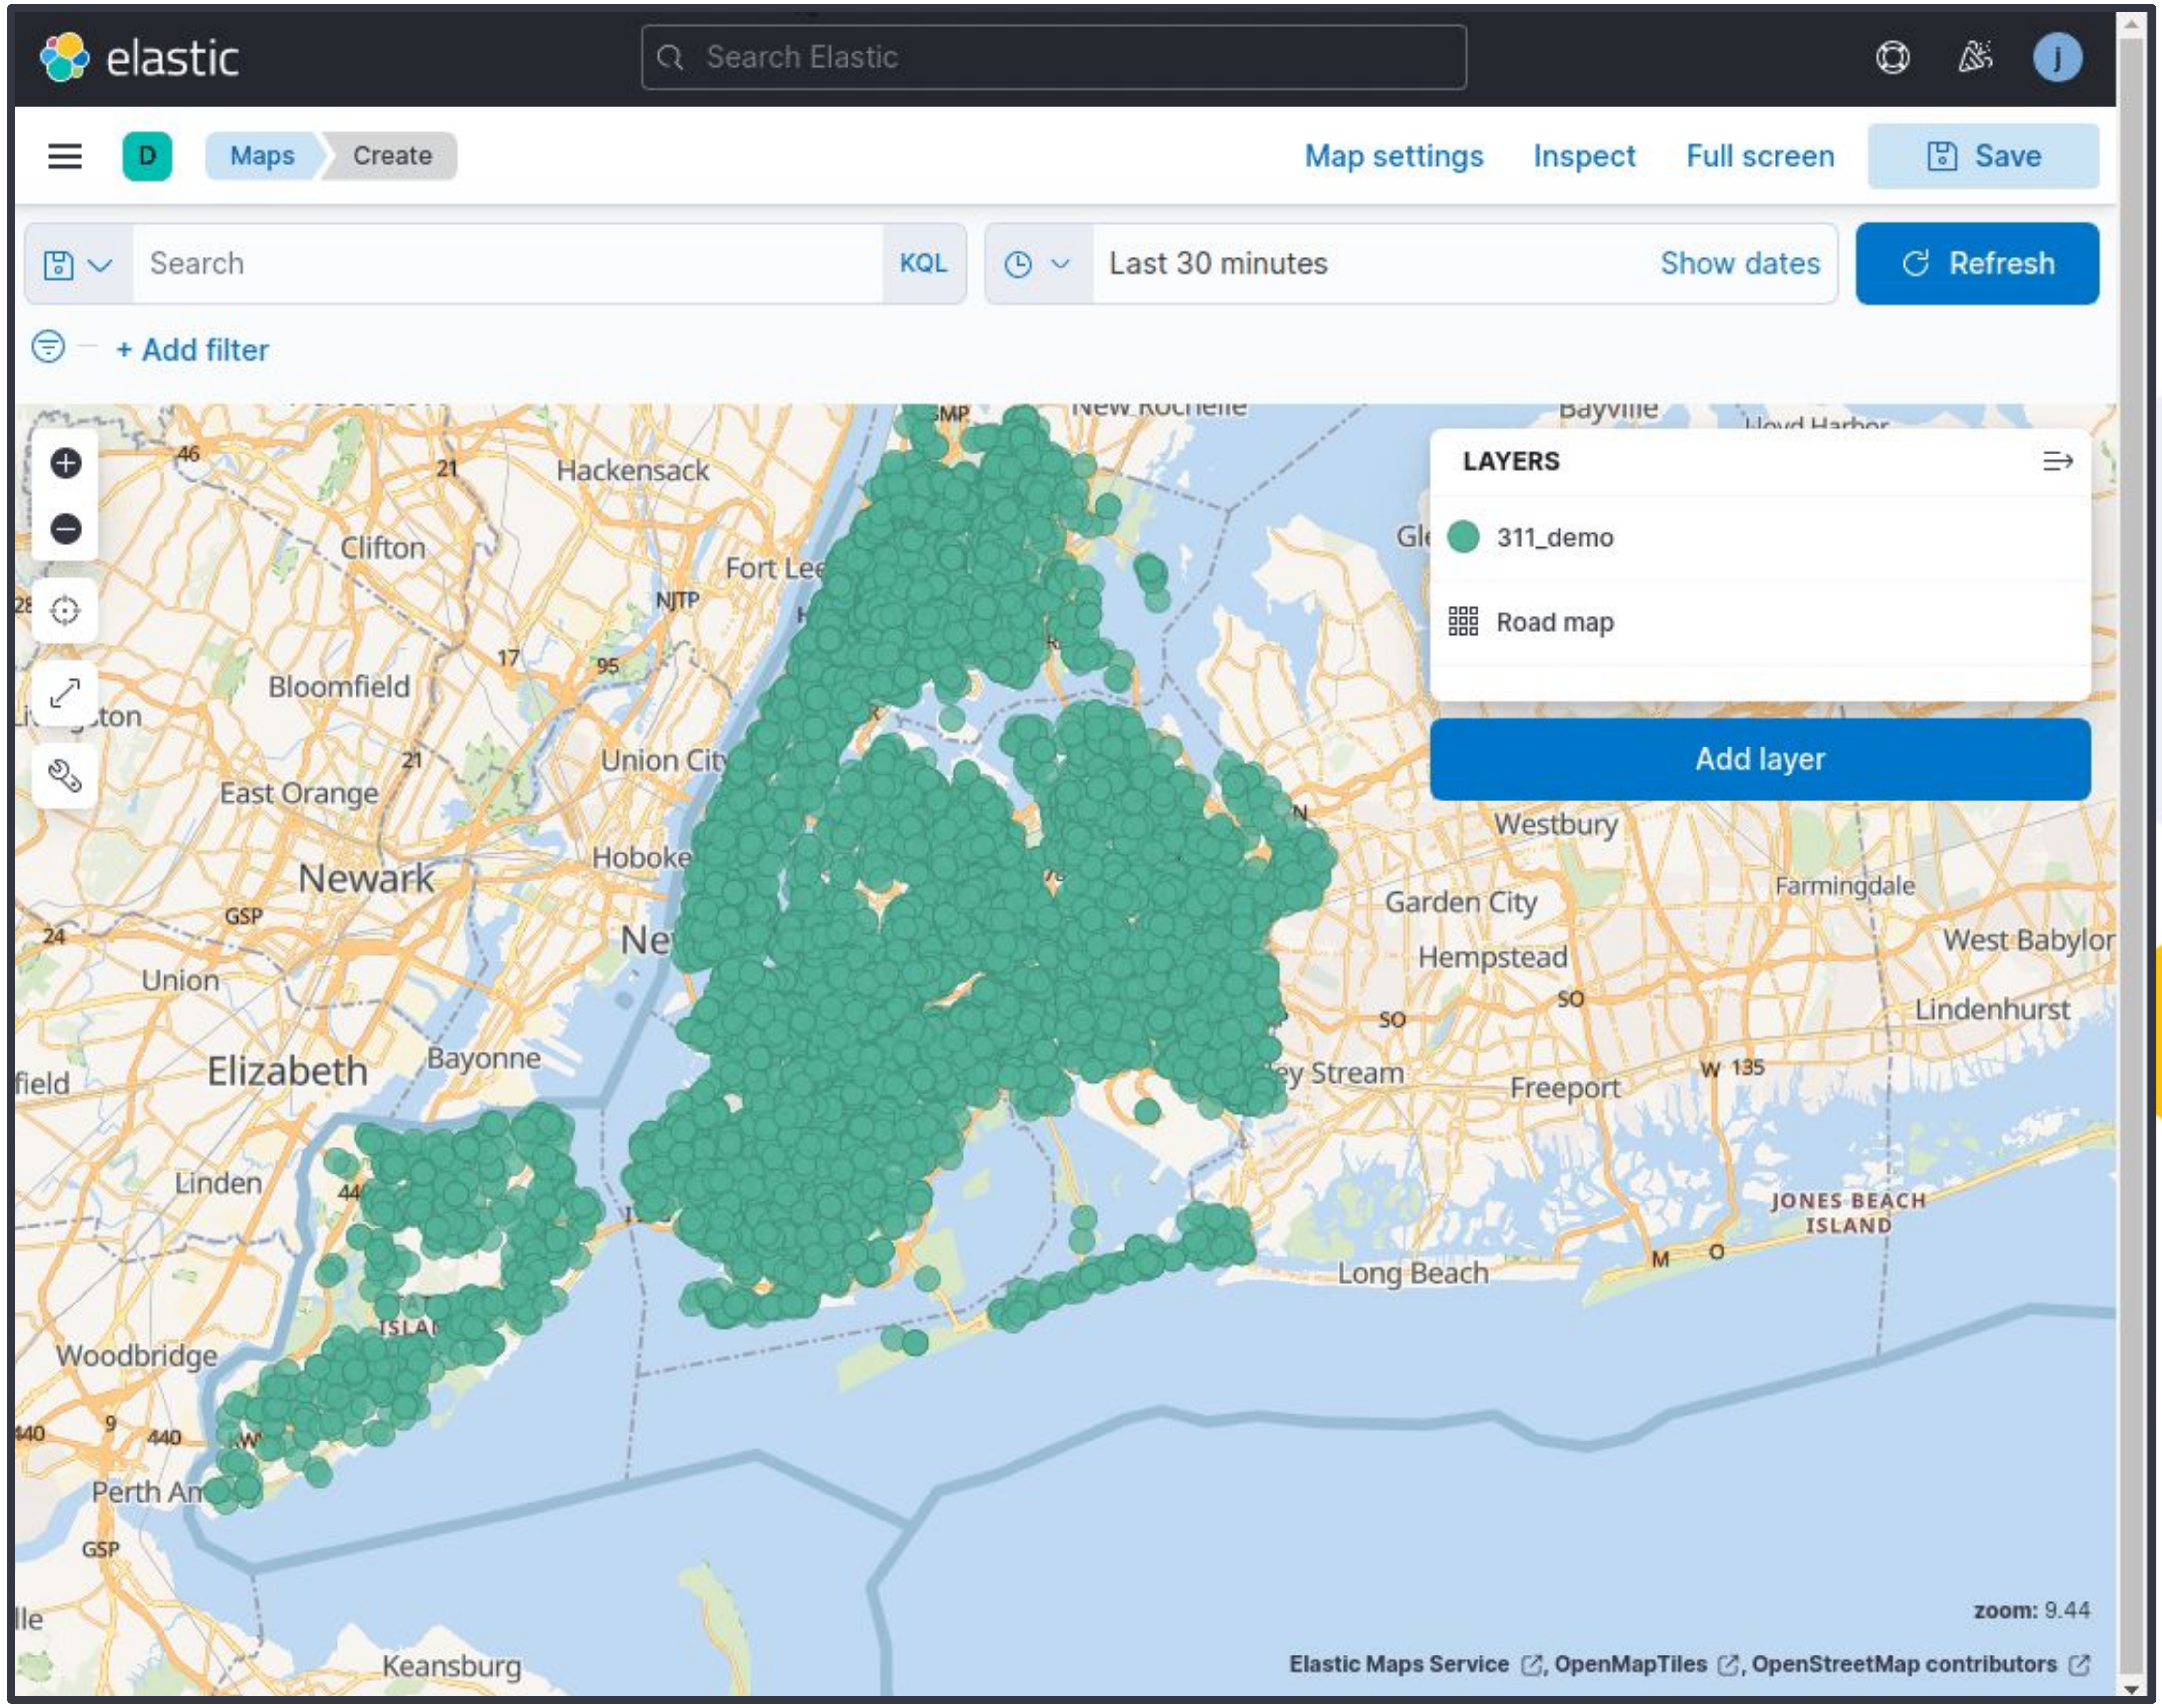

- Docs 📖
- Tutorial 📖

Alert

- Index areas of interest or draw in Kibana Maps
- Setup Tracking Containment alert
- Alert on:
  - Entered
  - Exited
  - Crossed
- Link to actions
  - Email
  - Slack/ MS Teams
  - Jiira
  - ...

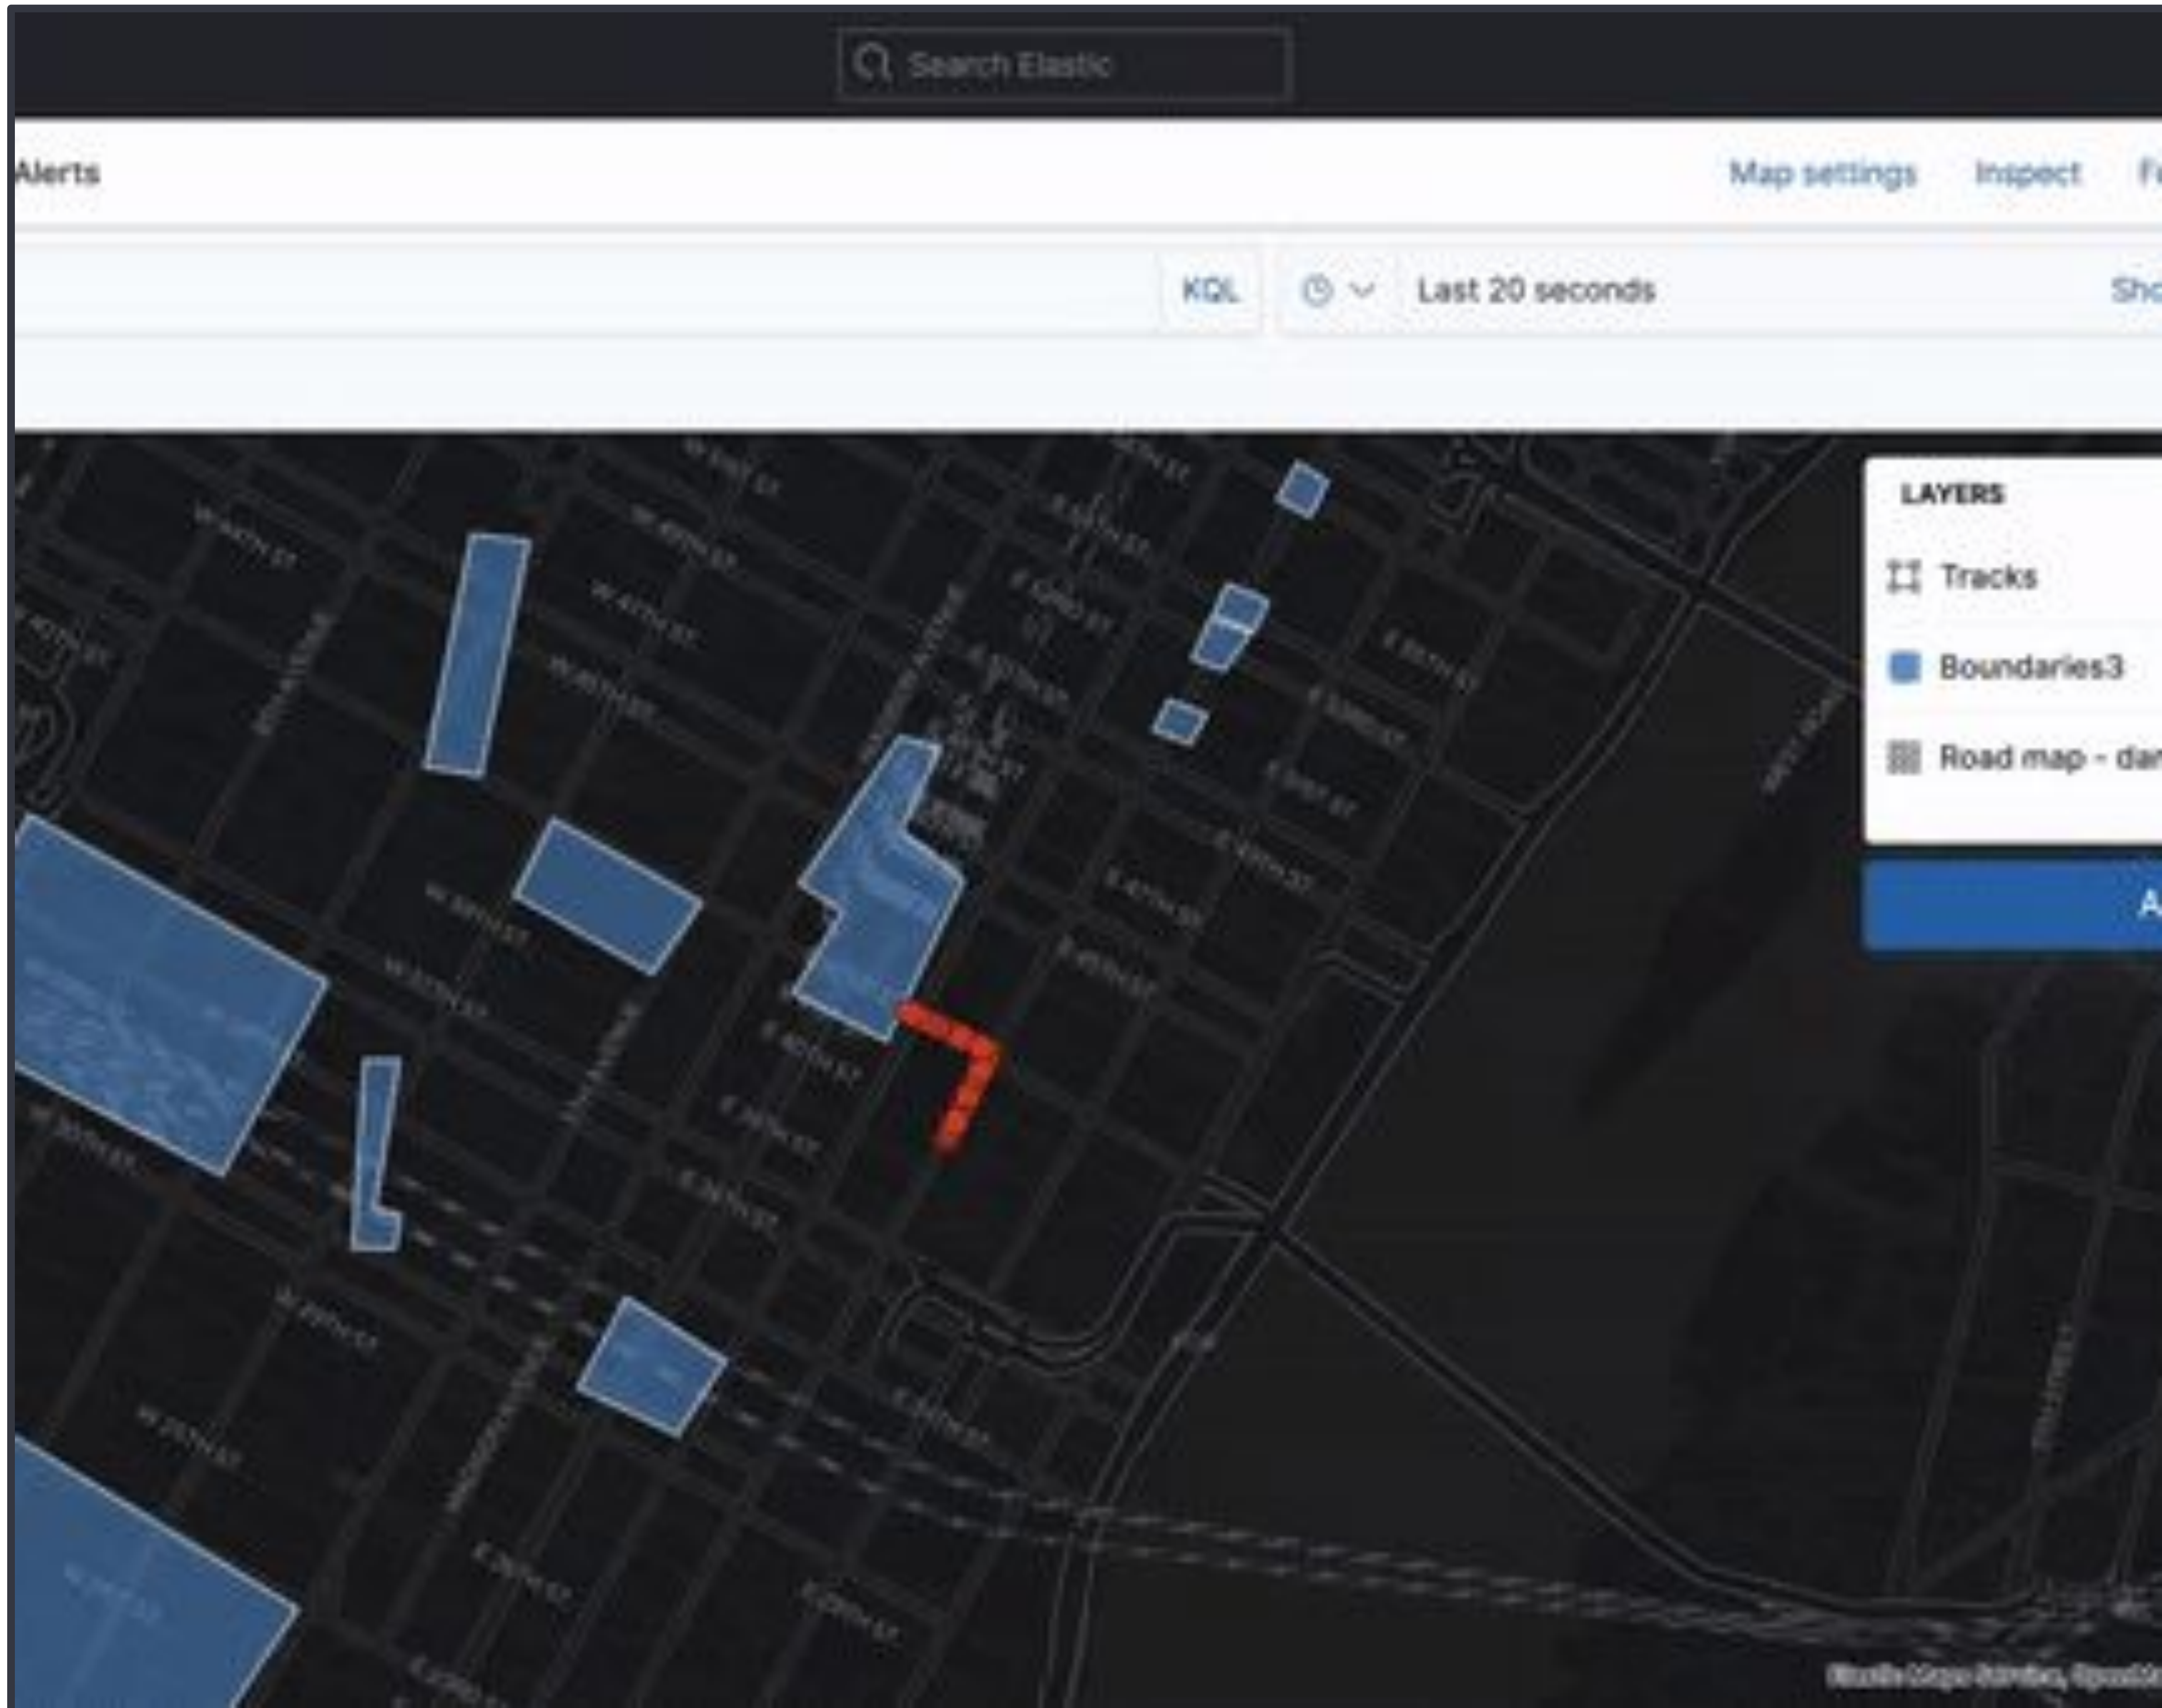

Example

"Slack me when one of our busses leaves the city boundary "

Time to practice

- Create a **new map**
- Upload data to your *cluster*
    - **Ainports** (from the Console or as GeoJSON details)
    - **Positions** (with the loader or from static file)
- Create the **airports** layer
- **Grid (hex, tile or cluster) layer** for positions (zoom 0 to 6)
- **Positions layer** (zoom 7 and below)
- Add **tooltips**, play with data-driven **symbology**, etc
- Group positions by callsign (**tracks**) to render flight paths
- Add the map to a dashboard with some other visualizations like position metrics and histogram, countries treemap, onGround bar, etc...

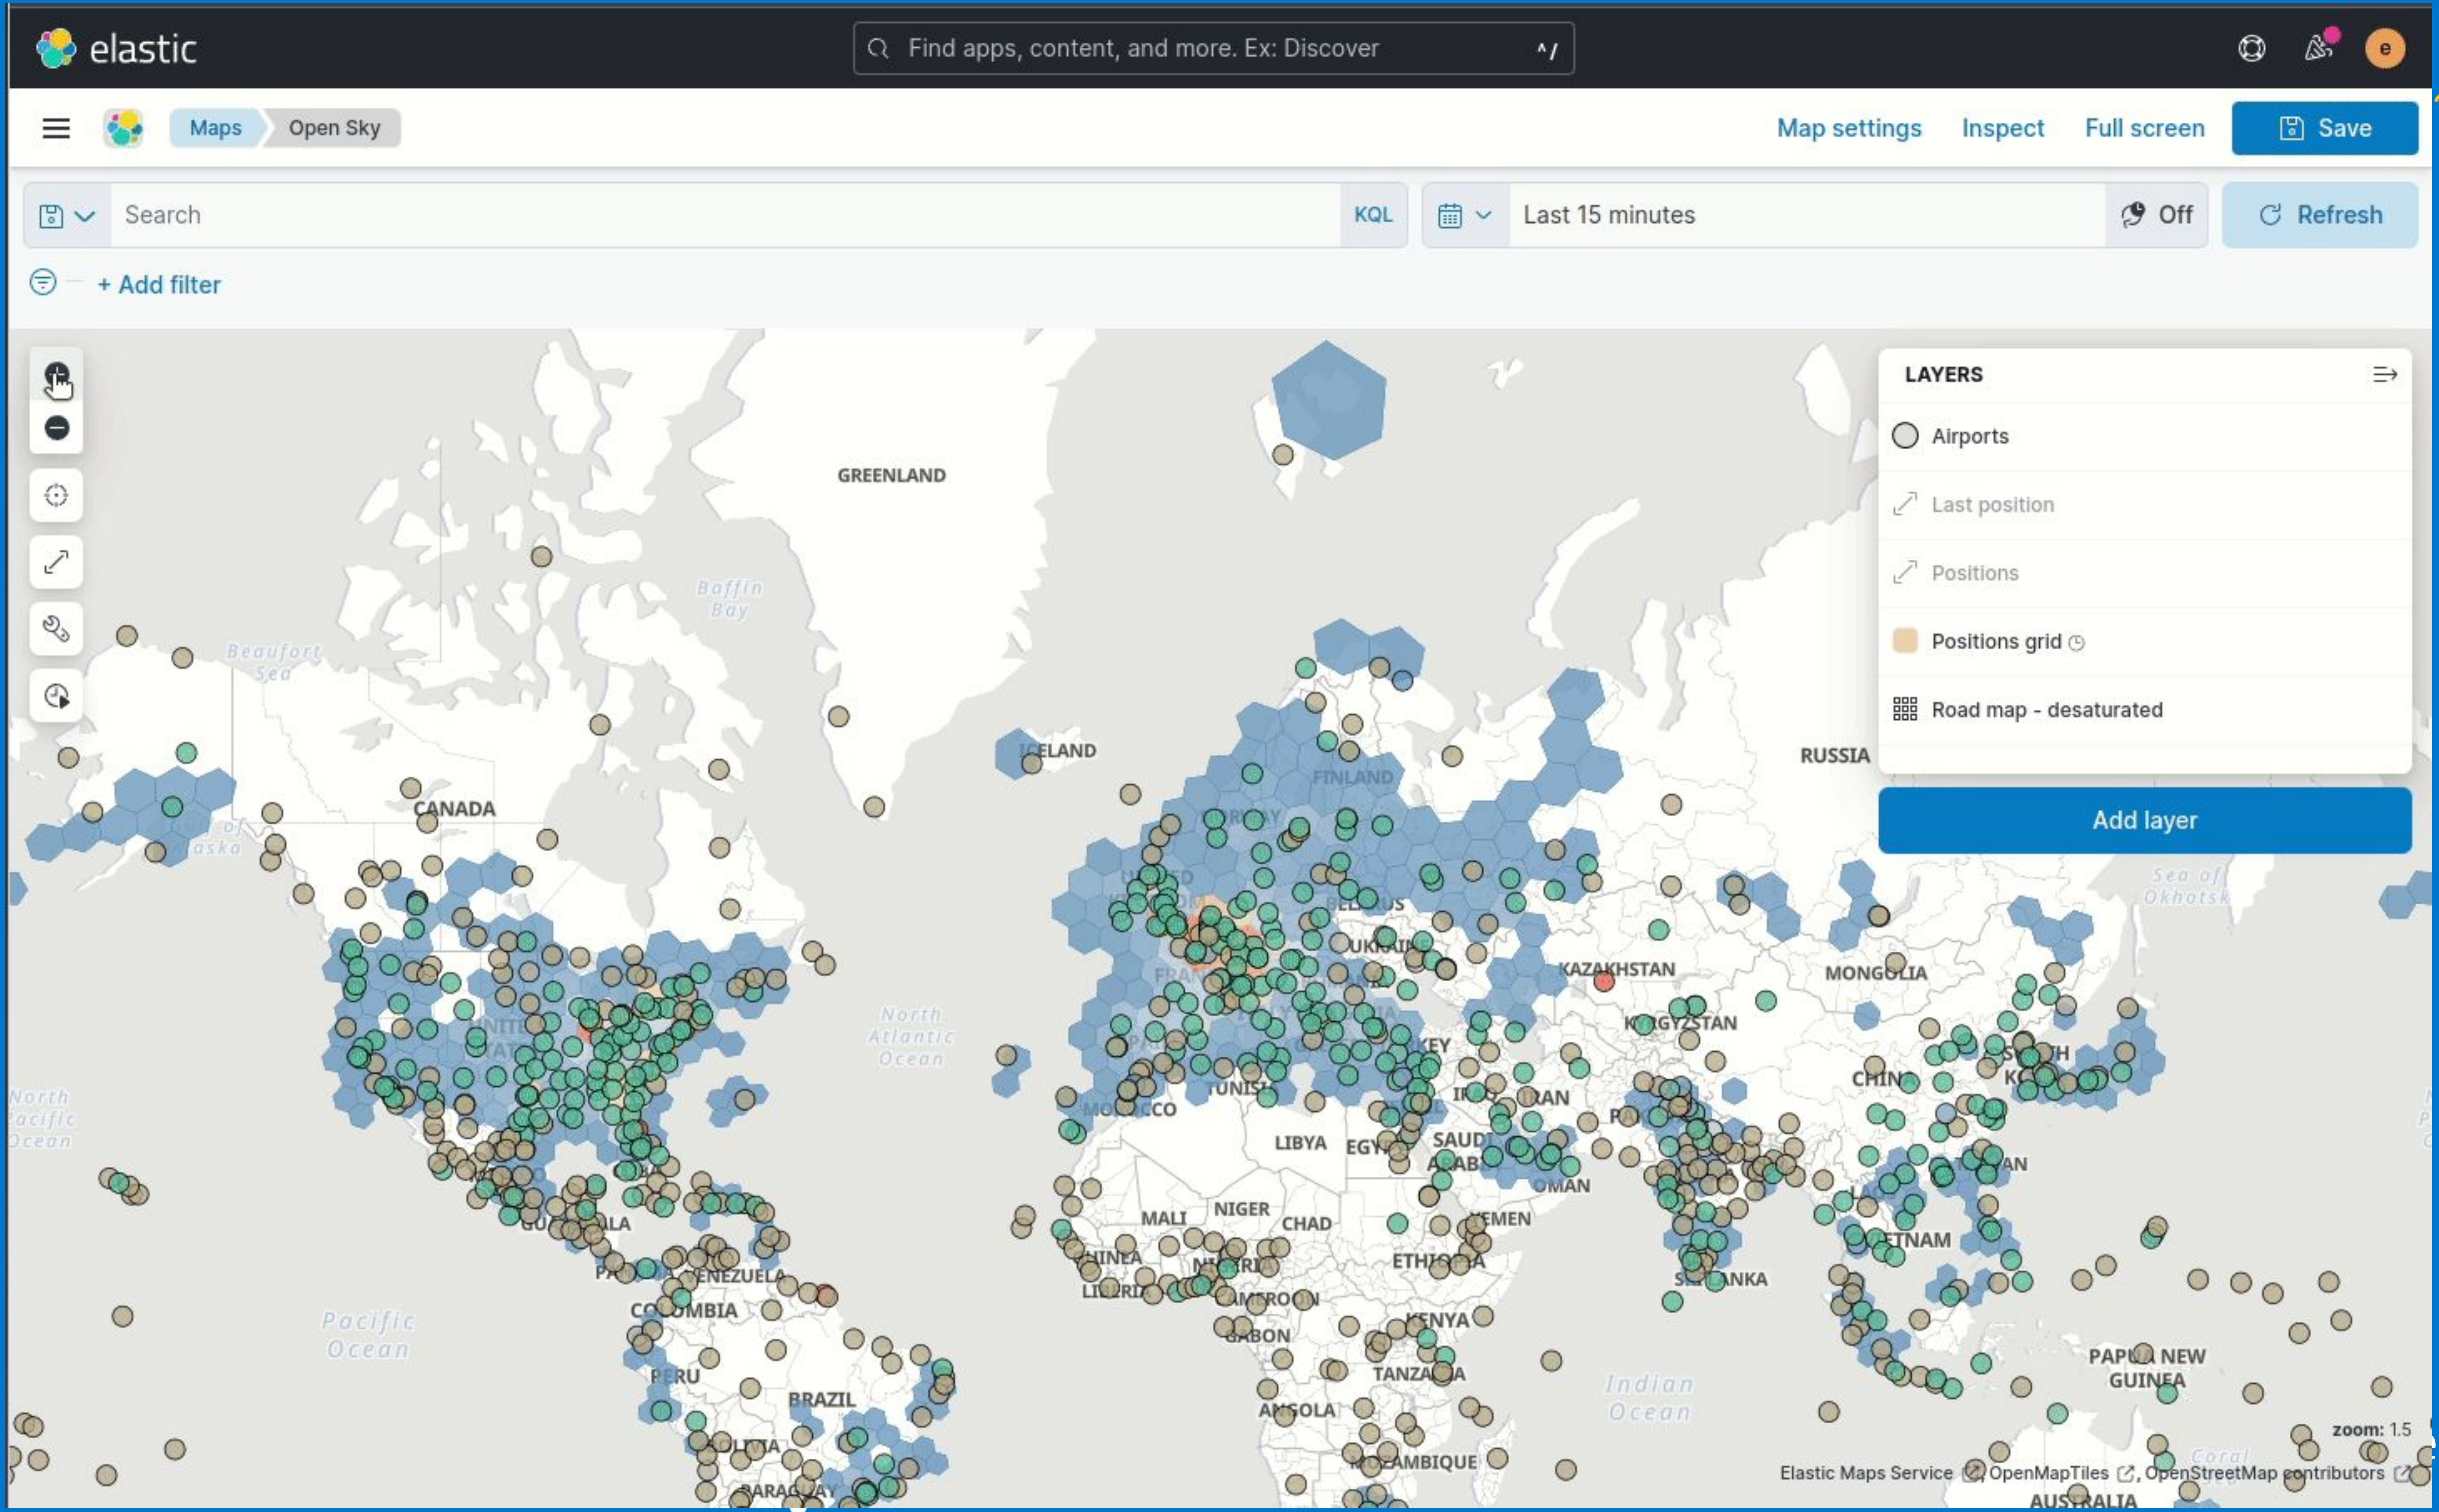
elastic
Find apps, content, and more. Ex: Discover
Maps
Open Sky
Map settings
Inspect
Full screen
Save
Search
KQL
Last 15 minutes
Off
Refresh
+ Add filter
LAYERS
Airports
Last position
Positions
Positions grid
Road map - desaturated
Add layer
GREENLAND
ICELAND
FINLAND
RUSSIA
CANADA
KAZAKHSTAN
KYRGYZSTAN
MONGOLIA
CHINA
IRAN
SAUDI ARABIA
OMAN
YEMEN
LIBYA
EGYPT
MALI
NIGER
CHAD
ETHIOPIA
KENYA
TANZANIA
ANGOLA
GABON
MOZAMBIQUE
COLOMBIA
VENEZUELA
PERU
BRAZIL
BOLIVIA
PARAGUAY
PAPUA NEW GUINEA
AUSTRALIA
VIETNAM
SRI LANKA
North Atlantic Ocean
North Pacific Ocean
Pacific Ocean
Indian Ocean
Baffin Bay
Beaufort Sea
Sea of Okhotsk
zoom: 1.5
Elastic Maps Service, OpenMapTiles, OpenStreetMap contributors

Agenda

1. Introduction to the Elastic stack
2. Lab setup
3. Kibana introduction
4. Elastics Maps
5. Elasticsearch Geo
6. Web mapping and Elasticsearch

Elasticsearch geospatial data types

- `geo_point` 📖
  - A single pair of latitude and longitude **coordinates**
  - Can be inserted as an object, WKT, array, geohash
- `geo_shape` 📖
  - Supports any **lat/lon** geometry type, incl. envelope and circle
  - Inserted with GeoJSON or WKT notation
- `shape` 📖
  - Supports any **cartesian** geometry type
  - Inserted with GeoJSON or WKT notation

Geo Enrichment

Adding a geo dimension to your data

- Transform data at ingest time
- Create `geo_point` from lat/lon fields
- Enrich IP addressed with estimated location
- Lookup location based on another index e.g. postcodes
- Tag documents by matching with polygons in another index e.g. local authority boundaries
- Convenient UI in Kibana

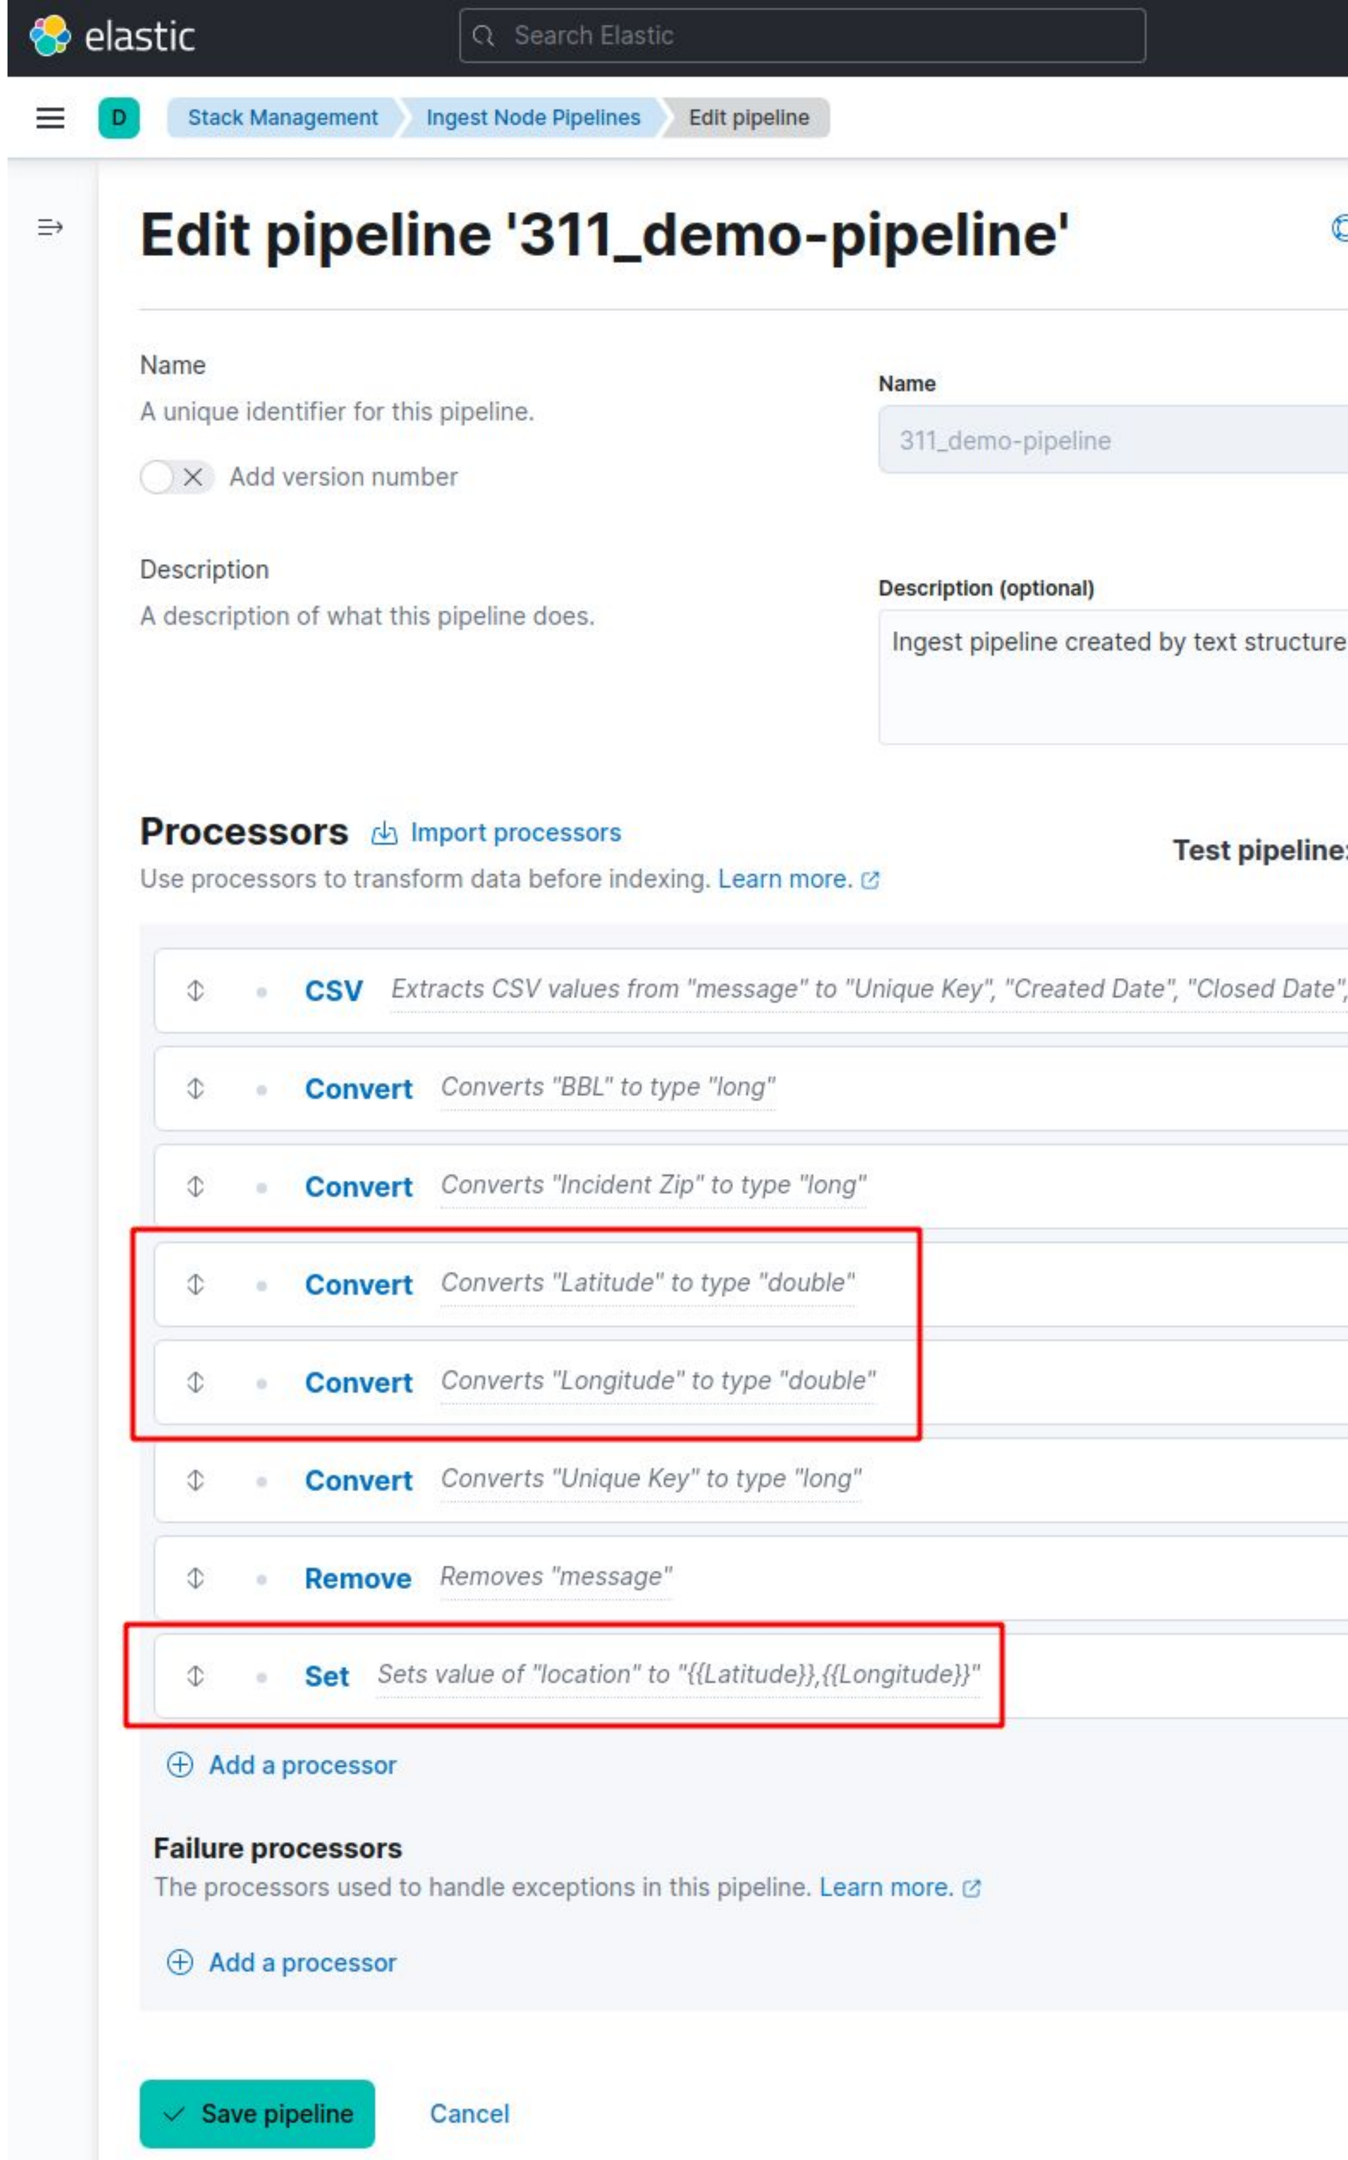

Edit pipeline '311_demo-pipeline'

Name
A unique identifier for this pipeline.
Add version number

Name: 311_demo-pipeline

Description
A description of what this pipeline does.

Description (optional): Ingest pipeline created by text structure

Processors
Use processors to transform data before indexing. Learn more.

- **CSV** *Extracts CSV values from "message" to "Unique Key", "Created Date", "Closed Date",*
- **Convert** *Converts "BBL" to type "long"*
- **Convert** *Converts "Incident Zip" to type "long"*
- **Convert** *Converts "Latitude" to type "double"*
- **Convert** *Converts "Longitude" to type "double"*
- **Convert** *Converts "Unique Key" to type "long"*
- **Remove** *Removes "message"*
- **Set** *Sets value of "location" to "{{Latitude}},{{Longitude}}"*

Failure processors
The processors used to handle exceptions in this pipeline. Learn more.

API or Vector tiles

Integrate in to your own system

Elasticsearch **search API**

- Modern, **REST**-based API
- **JSON** is the default output format

Elasticsearch **Vector Tiles API**

- Output in **protobuffer** format
- Use queries and aggregations to generate standard vector tiles

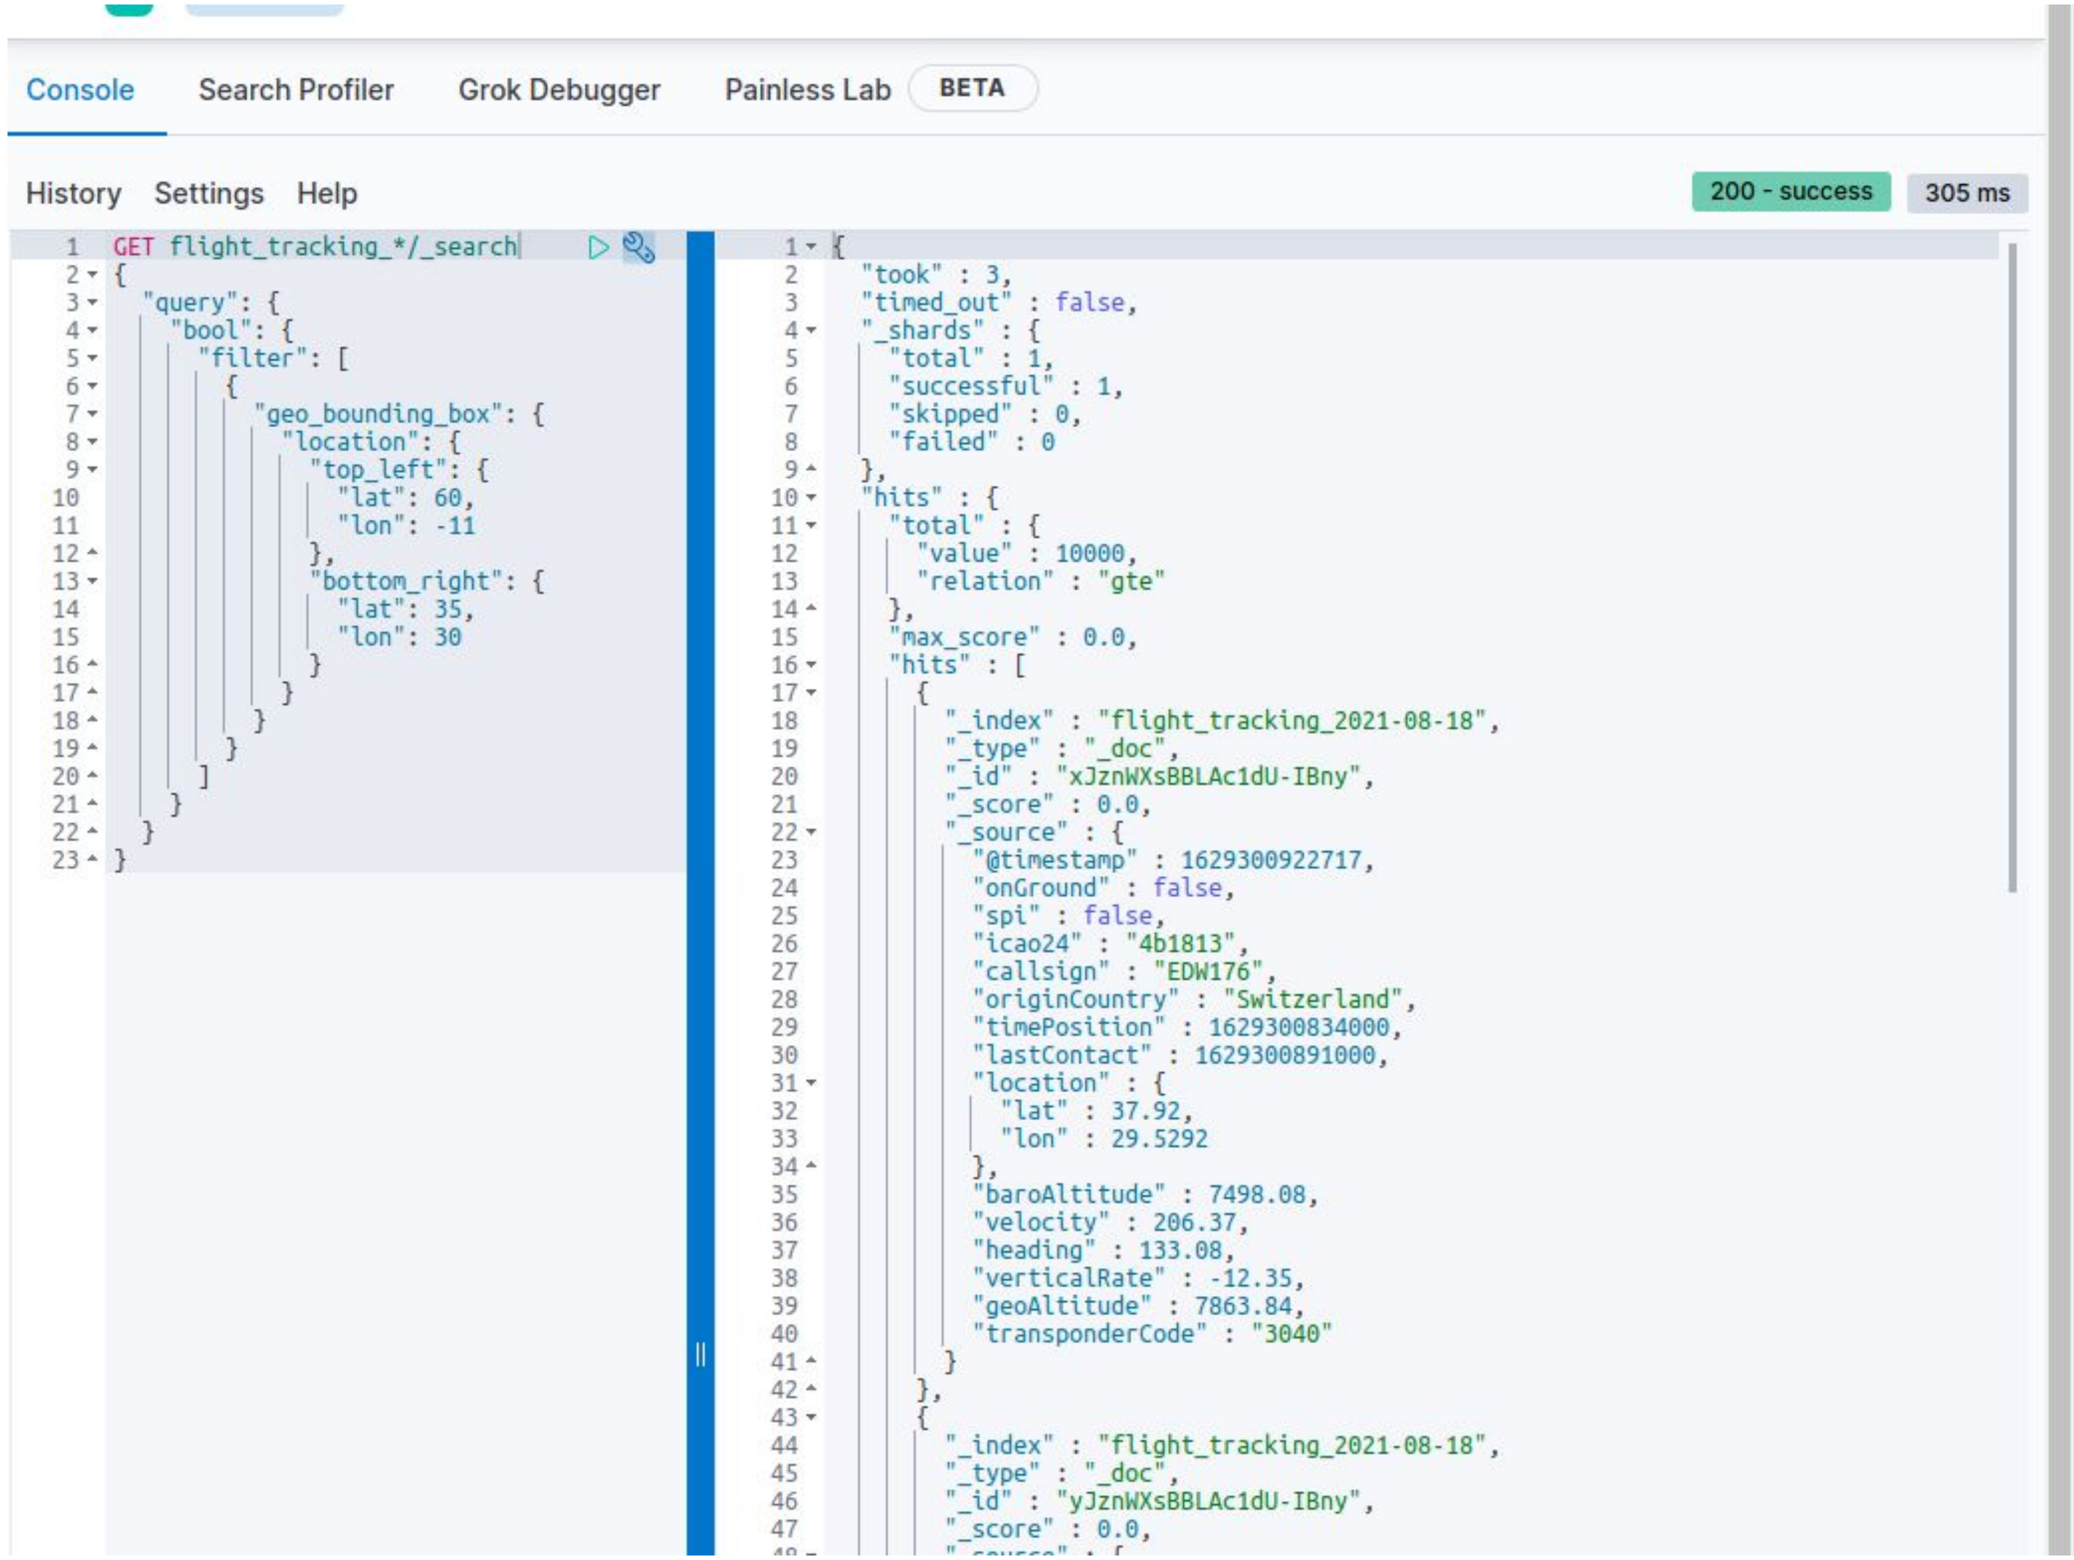

```
GET flight_tracking_*/_search
{
  "query": {
    "bool": {
      "filter": [
        {
          "geo_bounding_box": {
            "location": {
              "top_left": {
                "lat": 60,
                "lon": -11
              },
              "bottom_right": {
                "lat": 35,
                "lon": 30
              }
            }
          }
        }
      ]
    }
  }
}
```

```
{
  "took" : 3,
  "timed_out" : false,
  "_shards" : {
    "total" : 1,
    "successful" : 1,
    "skipped" : 0,
    "failed" : 0
  },
  "hits" : {
    "total" : {
      "value" : 10000,
      "relation" : "gte"
    },
    "max_score" : 0.0,
    "hits" : [
      {
        "_index" : "flight_tracking_2021-08-18",
        "_type" : "_doc",
        "_id" : "xJznWXsBBLAc1dU-IBny",
        "_score" : 0.0,
        "_source" : {
          "@timestamp" : 1629300922717,
          "onGround" : false,
          "spi" : false,
          "icao24" : "4b1813",
          "callsign" : "EDW176",
          "originCountry" : "Switzerland",
          "timePosition" : 1629300834000,
          "lastContact" : 1629300891000,
          "location" : {
            "lat" : 37.92,
            "lon" : 29.5292
          },
          "baroAltitude" : 7498.08,
          "velocity" : 206.37,
          "heading" : 133.08,
          "verticalRate" : -12.35,
          "geoAltitude" : 7863.84,
          "transponderCode" : "3040"
        }
      },
      {
        "_index" : "flight_tracking_2021-08-18",
        "_type" : "_doc",
        "_id" : "yJznWXsBBLAc1dU-IBny",
        "_score" : 0.0,
        "_source" : {
```

Ingest with ogr2ogr

https://gdal.org/drivers/vector/elasticsearch.html

- `ogr2ogr` can read and write into Elasticsearch
- Support for custom mapping definitions
- Blog posts:
    - How to ingest geospatial data into Elasticsearch with GDAL
    - Import OSM data into Elasticsearch with ogr2ogr and Docker

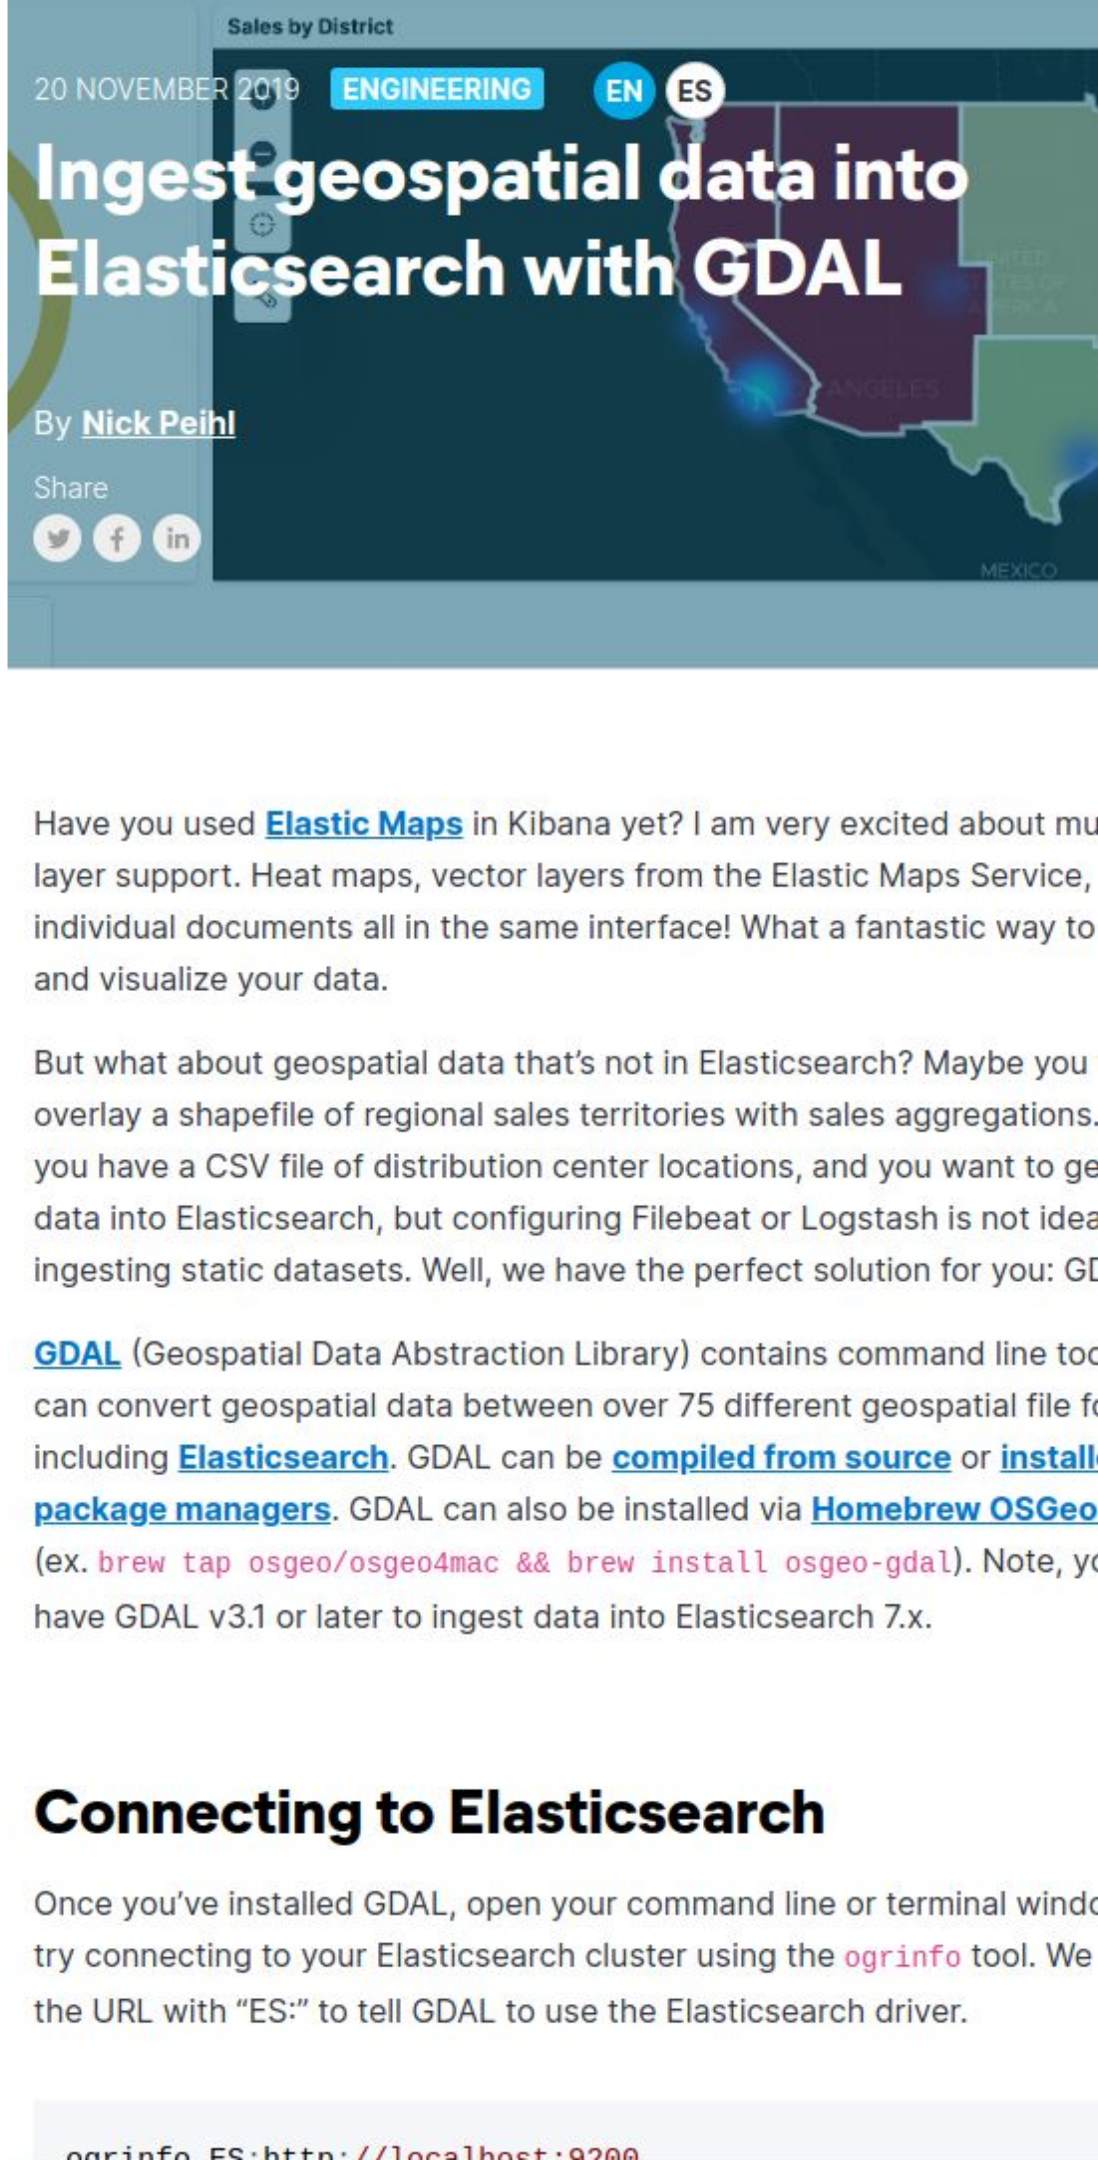

20 NOVEMBER 2019 ENGINEERING EN ES

Ingest geospatial data into Elasticsearch with GDAL

By Nick Peihl

Share

Have you used Elastic Maps in Kibana yet? I am very excited about mu
layer support. Heat maps, vector layers from the Elastic Maps Service,
individual documents all in the same interface! What a fantastic way to
and visualize your data.

But what about geospatial data that's not in Elasticsearch? Maybe you
overlay a shapefile of regional sales territories with sales aggregations.
you have a CSV file of distribution center locations, and you want to ge
data into Elasticsearch, but configuring Filebeat or Logstash is not idea
ingesting static datasets. Well, we have the perfect solution for you: GD

GDAL (Geospatial Data Abstraction Library) contains command line too
can convert geospatial data between over 75 different geospatial file fo
including Elasticsearch. GDAL can be compiled from source or install
package managers. GDAL can also be installed via Homebrew OSGeo
(ex. `brew tap osgeo/osgeo4mac && brew install osgeo-gdal`). Note, yo
have GDAL v3.1 or later to ingest data into Elasticsearch 7.x.

Connecting to Elasticsearch

Once you've installed GDAL, open your command line or terminal windo
try connecting to your Elasticsearch cluster using the `ogrinfo` tool. We
the URL with "ES:" to tell GDAL to use the Elasticsearch driver.

```
ogrinfo ES:http://localhost:9200
```

Search

Filter documents with geospatial relationships

Geo Filters

- Bounding box
- Point and radius
- Polygon
- An indexed geo_shape

Plus every other Elasticsearch filter

- Boolean
- Range (numeric, date, IP)
- Unstructured text (stemming, fuzzy ...)

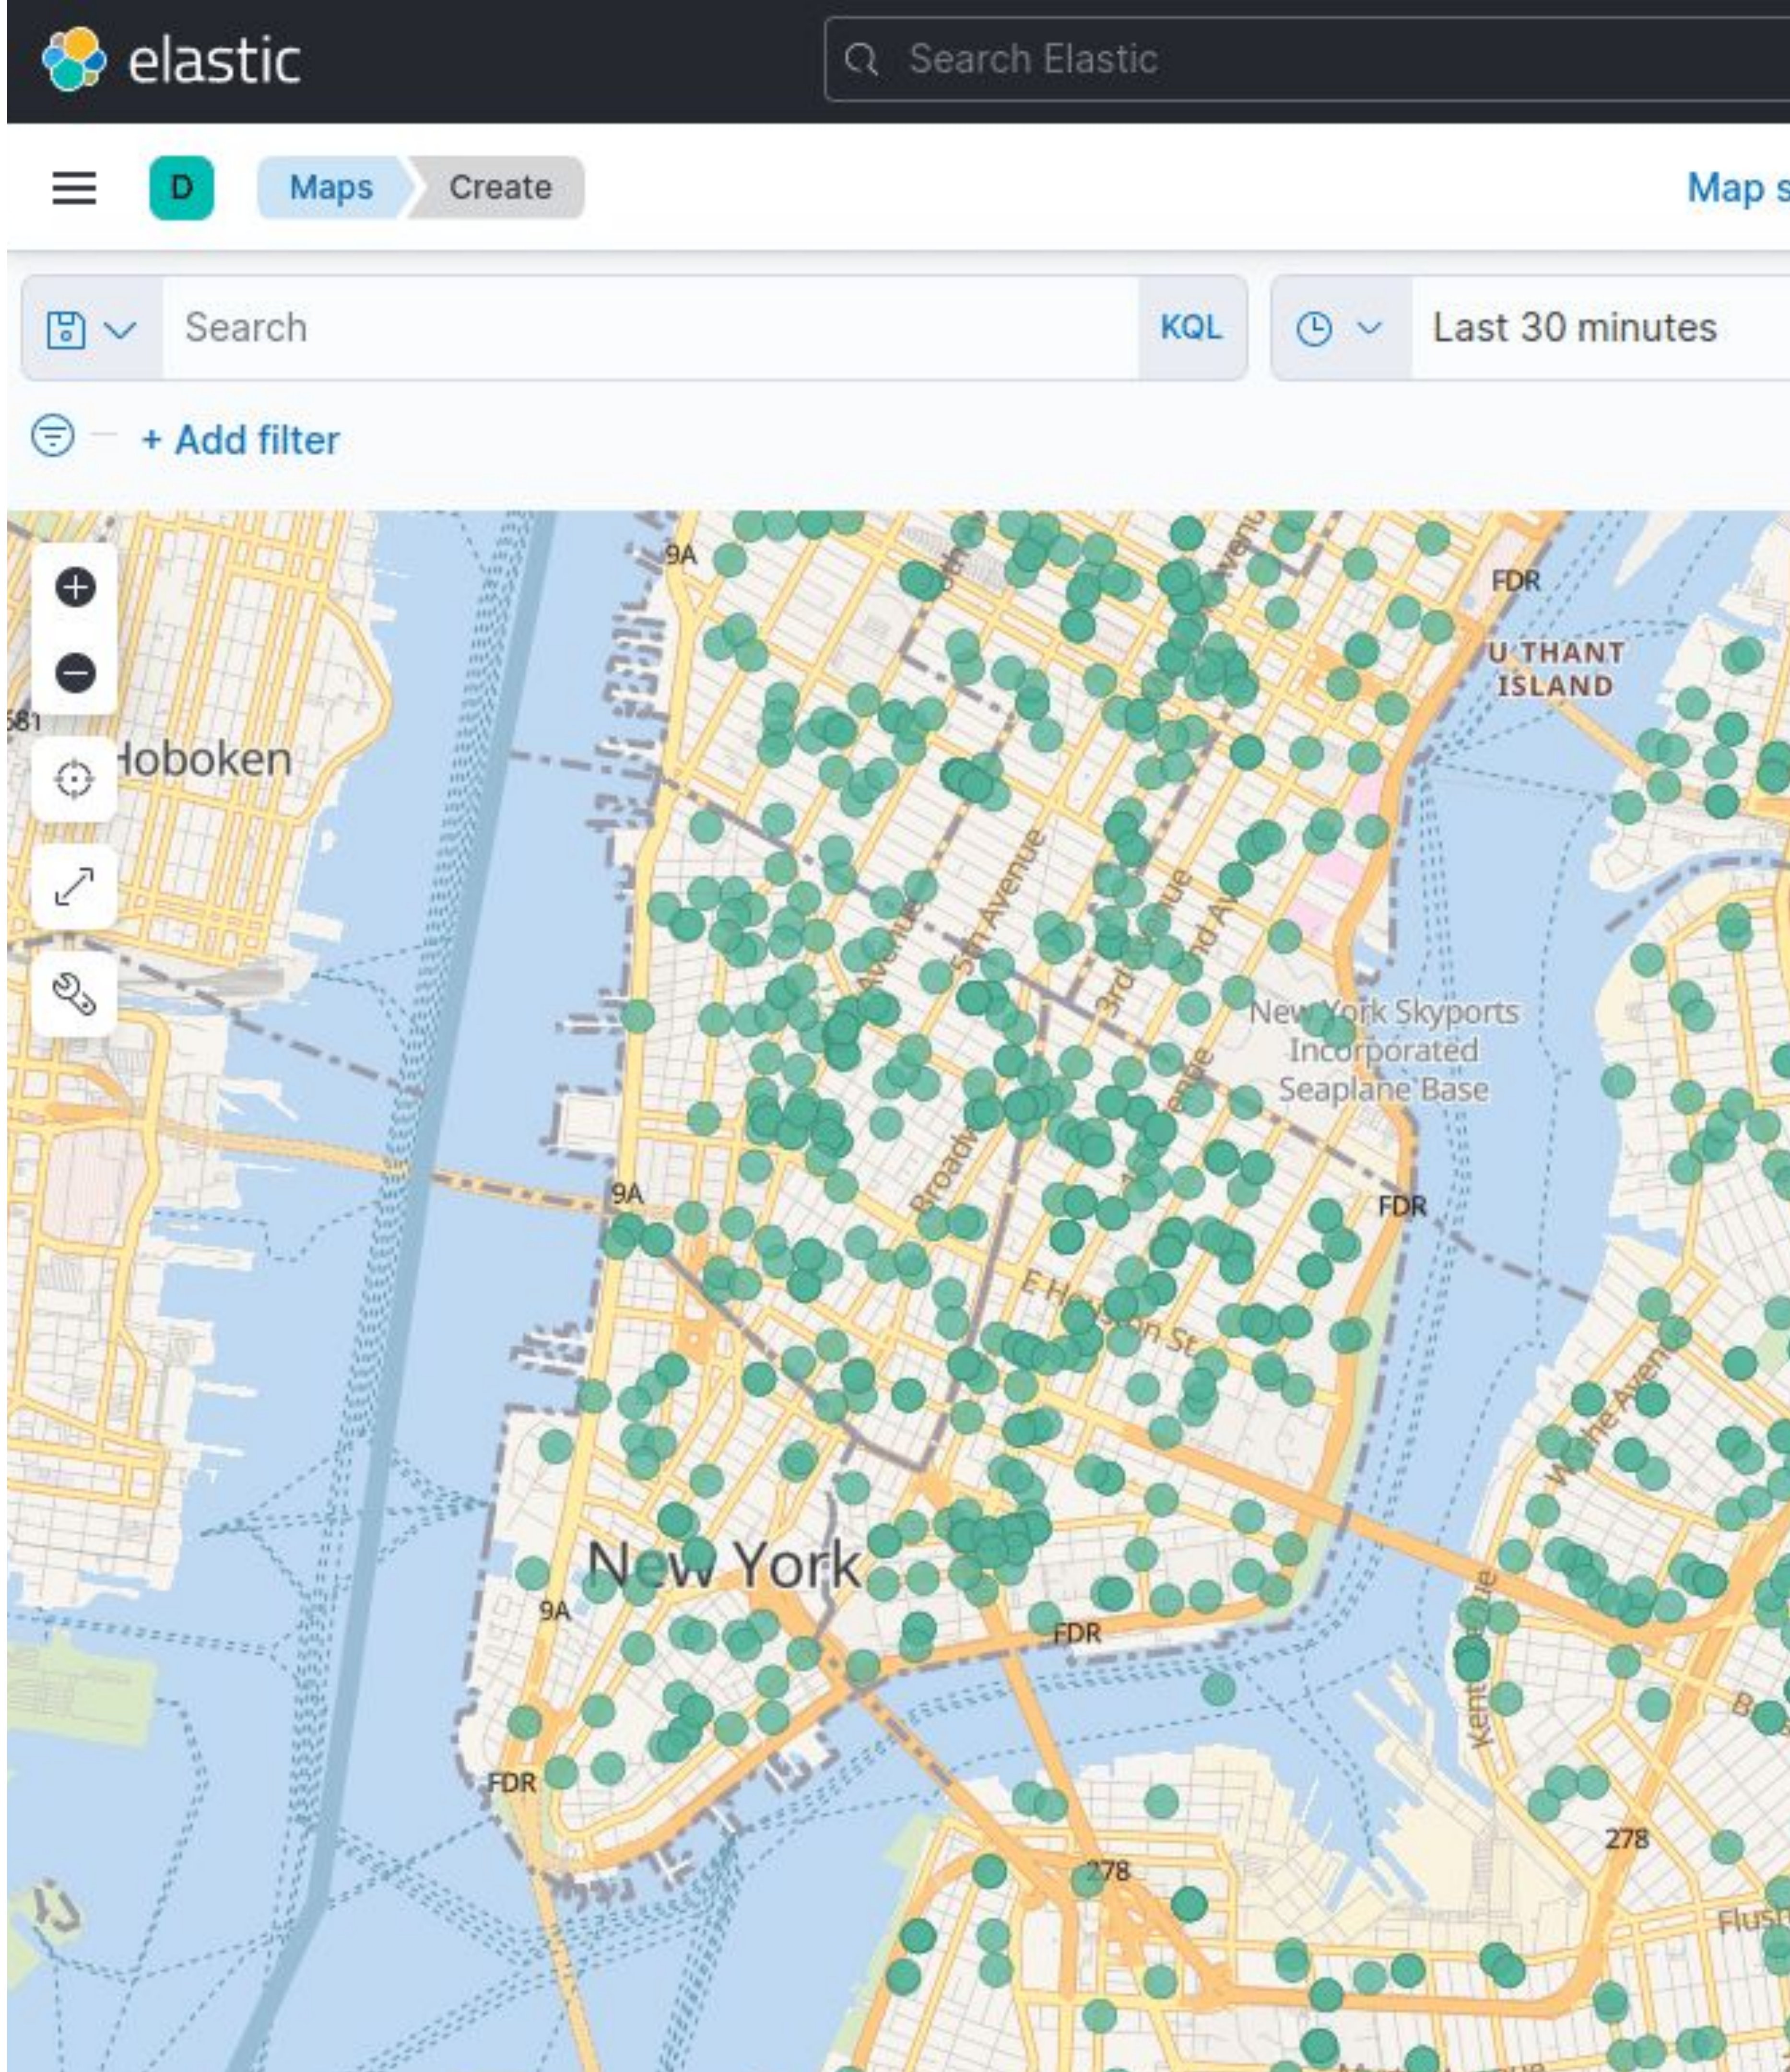

Example

"Show me all subscribers that live within 5 miles of our new gym location, that joined in the last year and have «running» mentioned in their profile"

Aggregate

Geo Bucket

- Distance (rings) 📖
- Hash 📖
- Geotile 📖
- Hex Grid 📖

Geo Metric

- Centroid 📖
- Bounds 📖
- Geoline 📖

Aggregate non-geo using geo filters

- Huge range of aggregations 📖

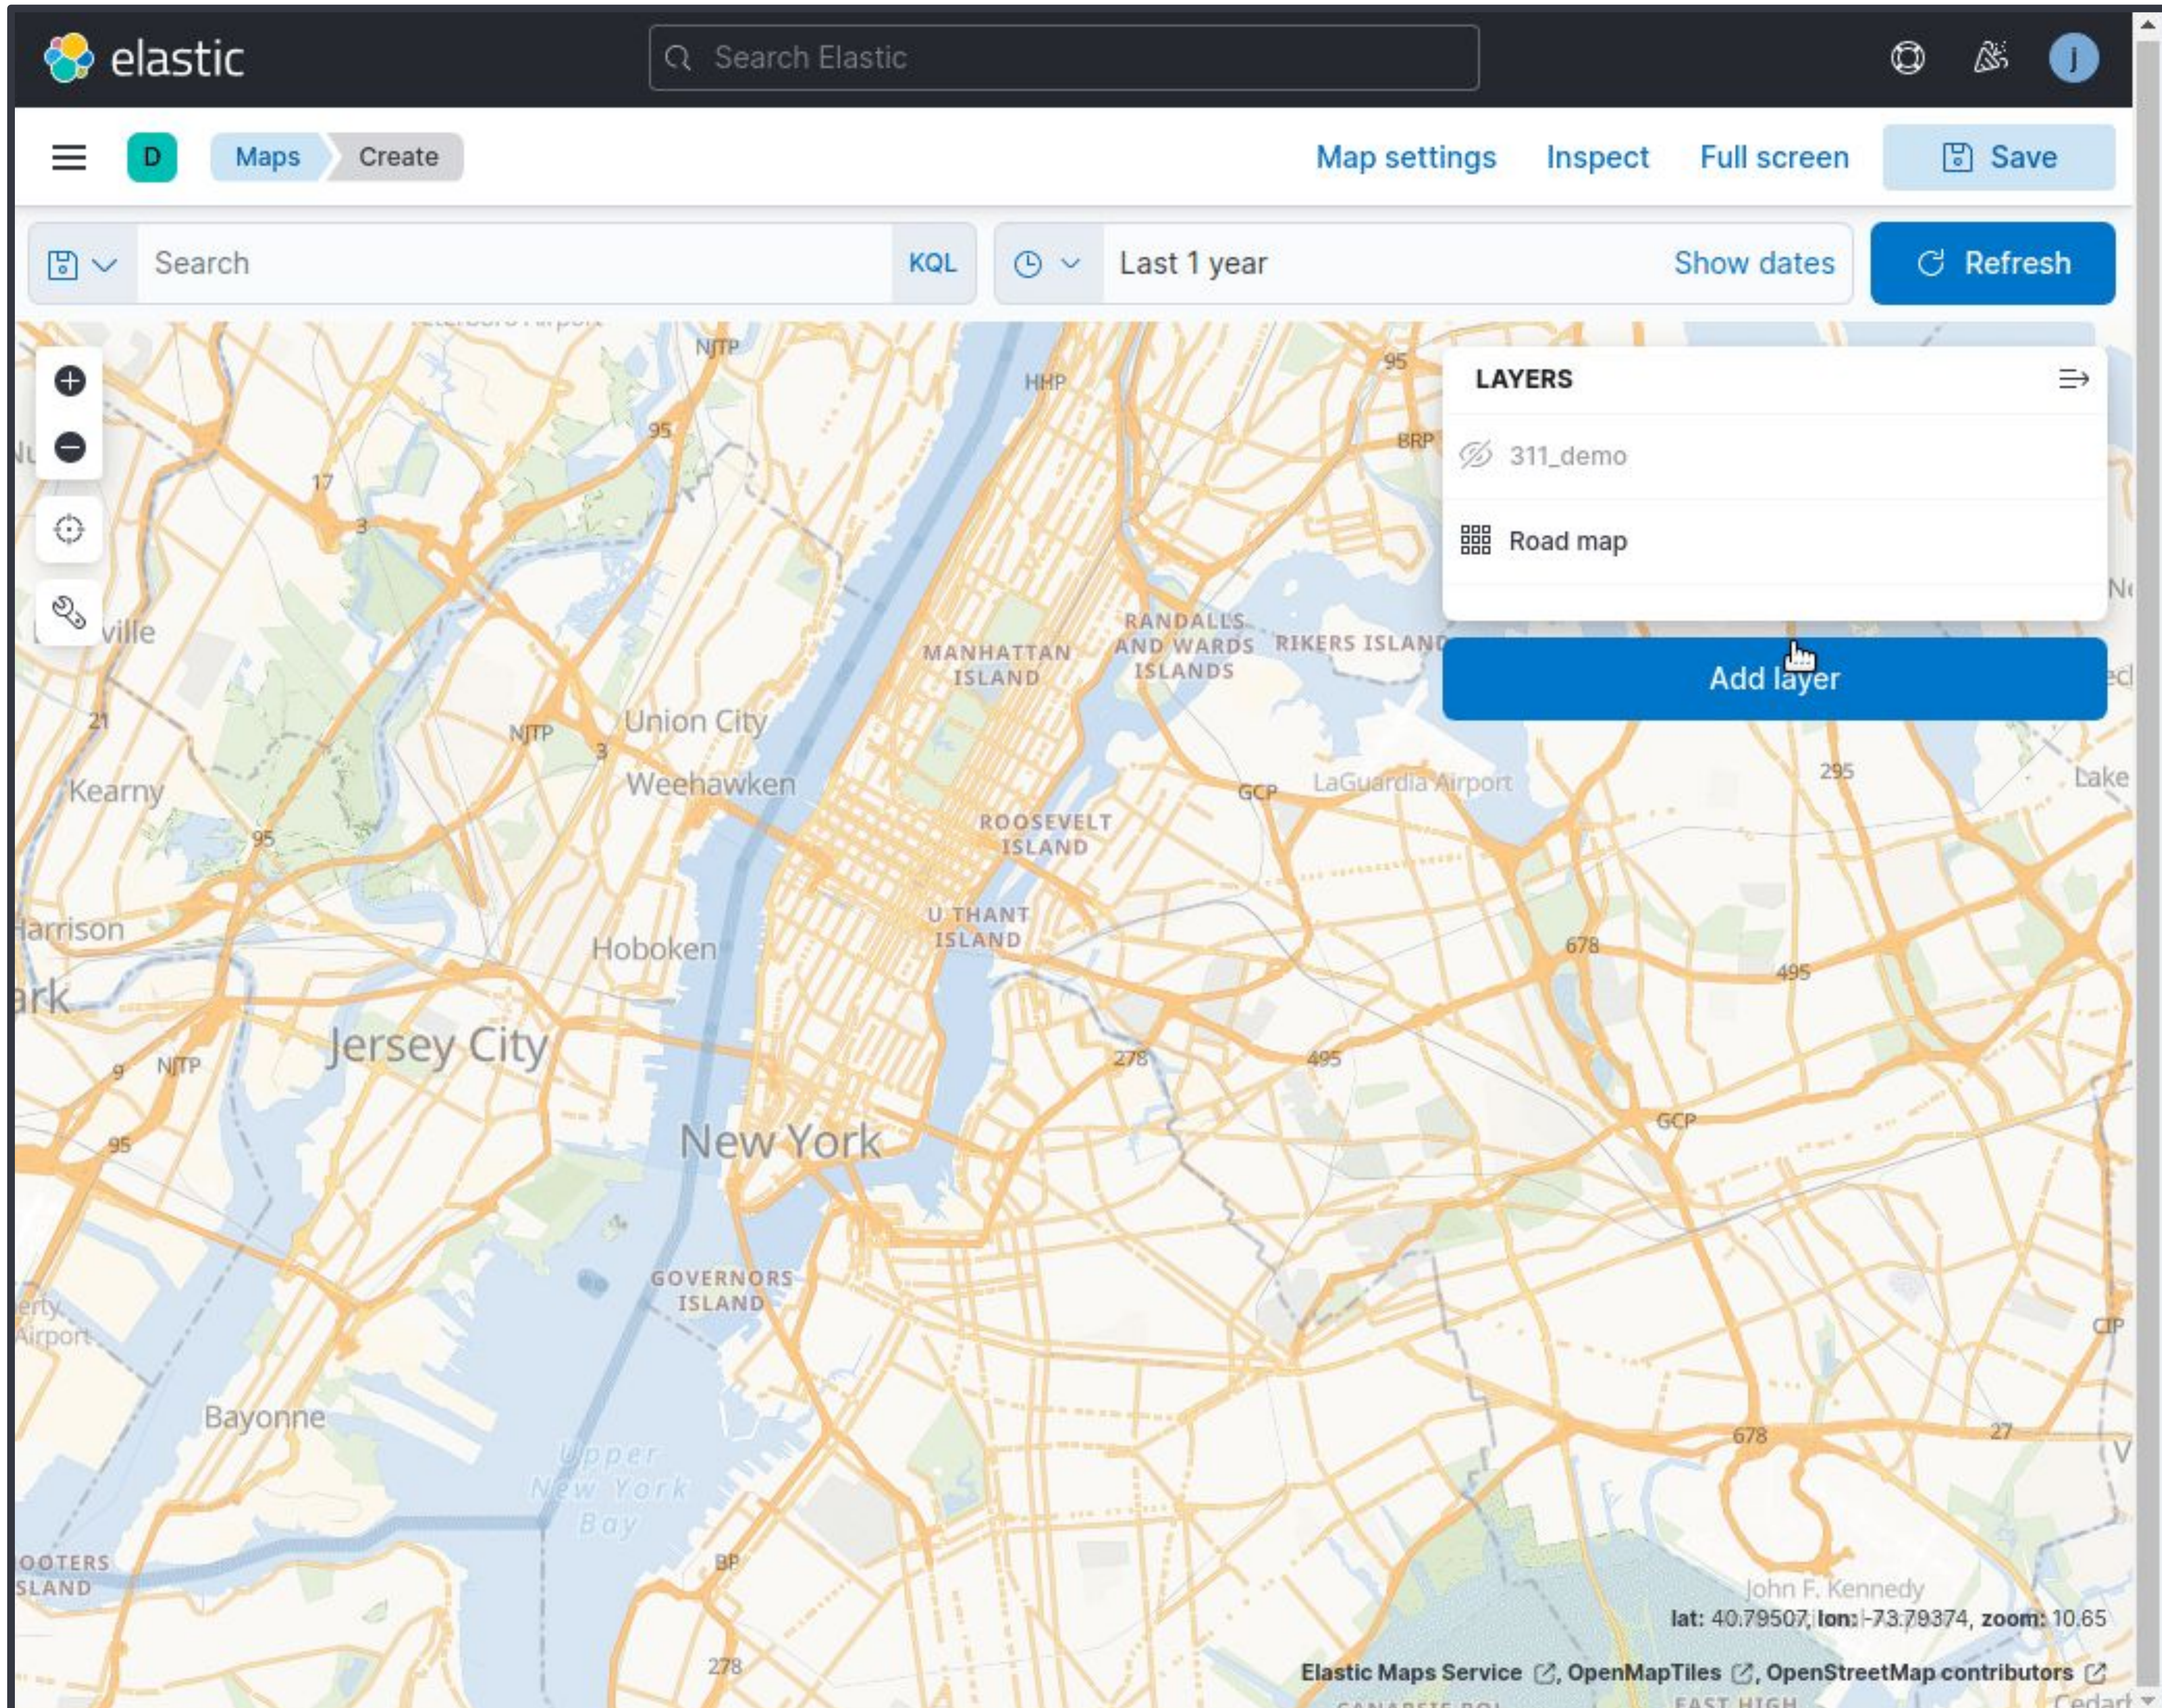

Example

“Provide a geographic heat map of total sales of **blue shirts** for the **last 5 years** ”

OGC servers and Elasticsearch

Expose Elasticsearch indices as OGC services

GeoServer

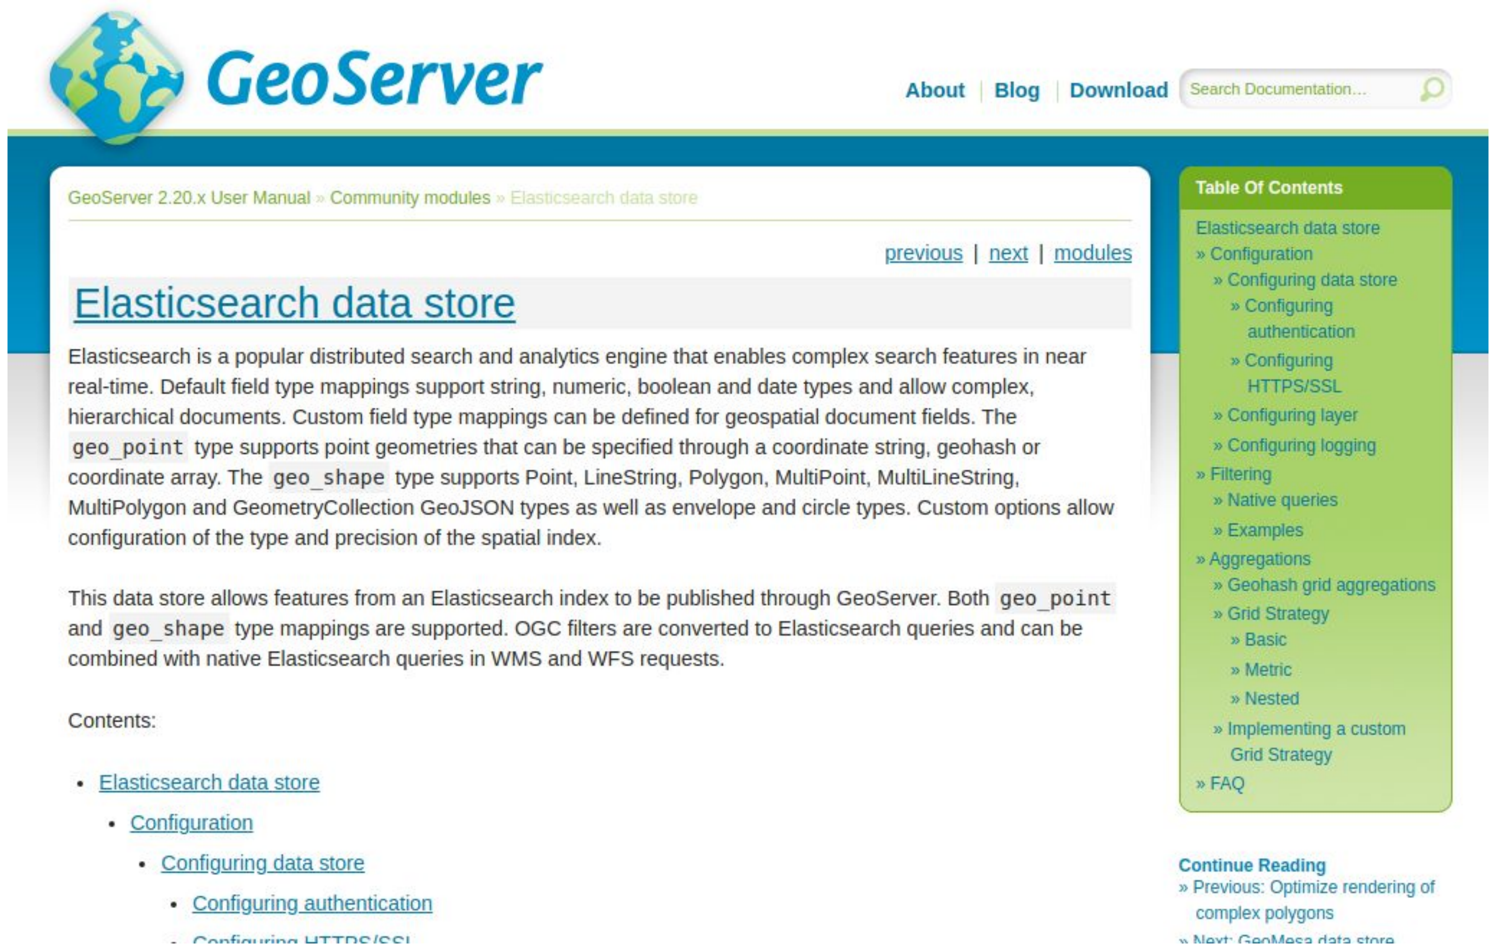

Elasticsearch data store

Elasticsearch is a popular distributed search and analytics engine that enables complex search features in near real-time. Default field type mappings support string, numeric, boolean and date types and allow complex, hierarchical documents. Custom field type mappings can be defined for geospatial document fields. The `geo_point` type supports point geometries that can be specified through a coordinate string, geohash or coordinate array. The `geo_shape` type supports Point, LineString, Polygon, MultiPoint, MultiLineString, MultiPolygon and GeometryCollection GeoJSON types as well as envelope and circle types. Custom options allow configuration of the type and precision of the spatial index.

This data store allows features from an Elasticsearch index to be published through GeoServer. Both `geo_point` and `geo_shape` type mappings are supported. OGC filters are converted to Elasticsearch queries and can be combined with native Elasticsearch queries in WMS and WFS requests.

Contents:

- Elasticsearch data store
  - Configuration
    - Configuring data store
      - Configuring authentication
      - Configuring HTTPS/SSL

pygeoapi

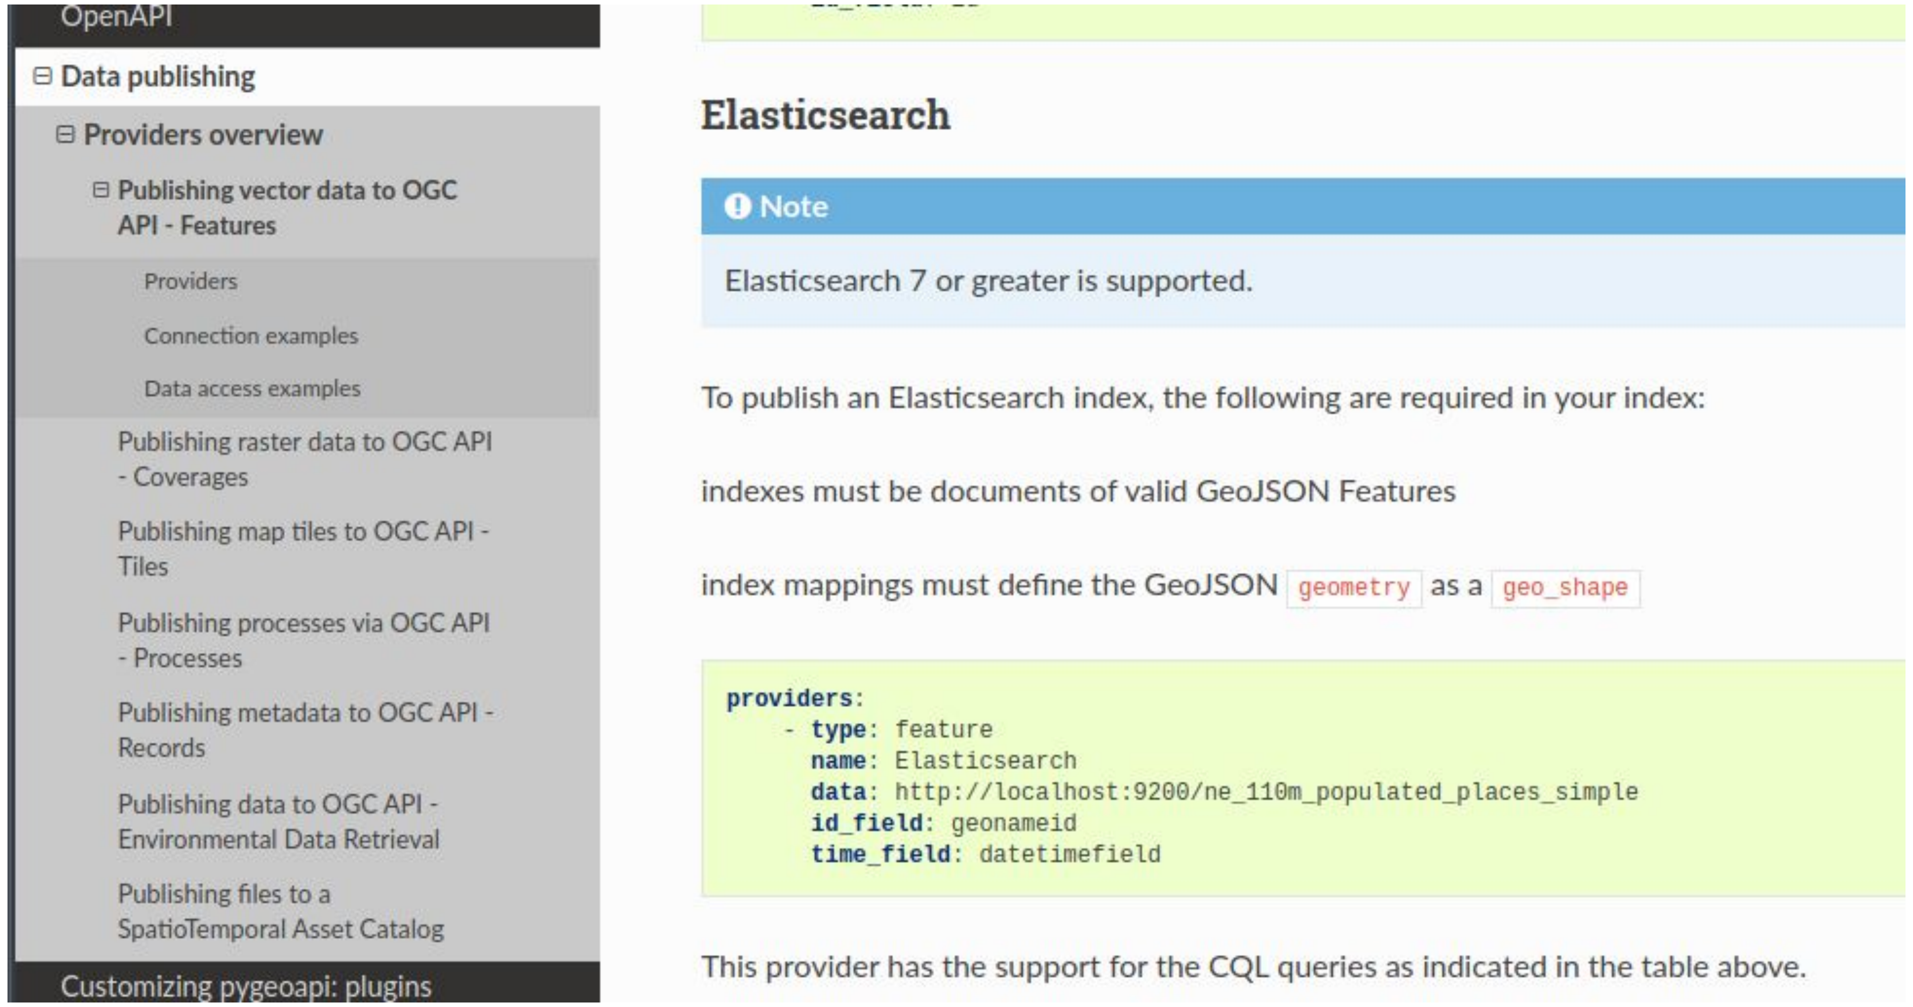
OpenAPI
⊟ Data publishing
⊟ Providers overview
⊟ Publishing vector data to OGC API - Features
Providers
Connection examples
Data access examples
Publishing raster data to OGC API - Coverages
Publishing map tiles to OGC API - Tiles
Publishing processes via OGC API - Processes
Publishing metadata to OGC API - Records
Publishing data to OGC API - Environmental Data Retrieval
Publishing files to a SpatioTemporal Asset Catalog
Customizing pygeoapi: plugins

Elasticsearch

Note

Elasticsearch 7 or greater is supported.

To publish an Elasticsearch index, the following are required in your index:

indexes must be documents of valid GeoJSON Features

index mappings must define the GeoJSON `geometry` as a `geo_shape`

```
providers:
    - type: feature
      name: Elasticsearch
      data: http://localhost:9200/ne_110m_populated_places_simple
      id_field: geonameid
      time_field: datetimefield
```

This provider has the support for the CQL queries as indicated in the table above.

Elasticsearch DSL

- Kibana DevTools Console
  - Or curl, postman, ...
- Create an index
  - Field types (also geospatial)
- Add data to your index
- Search documents
- Aggregate data

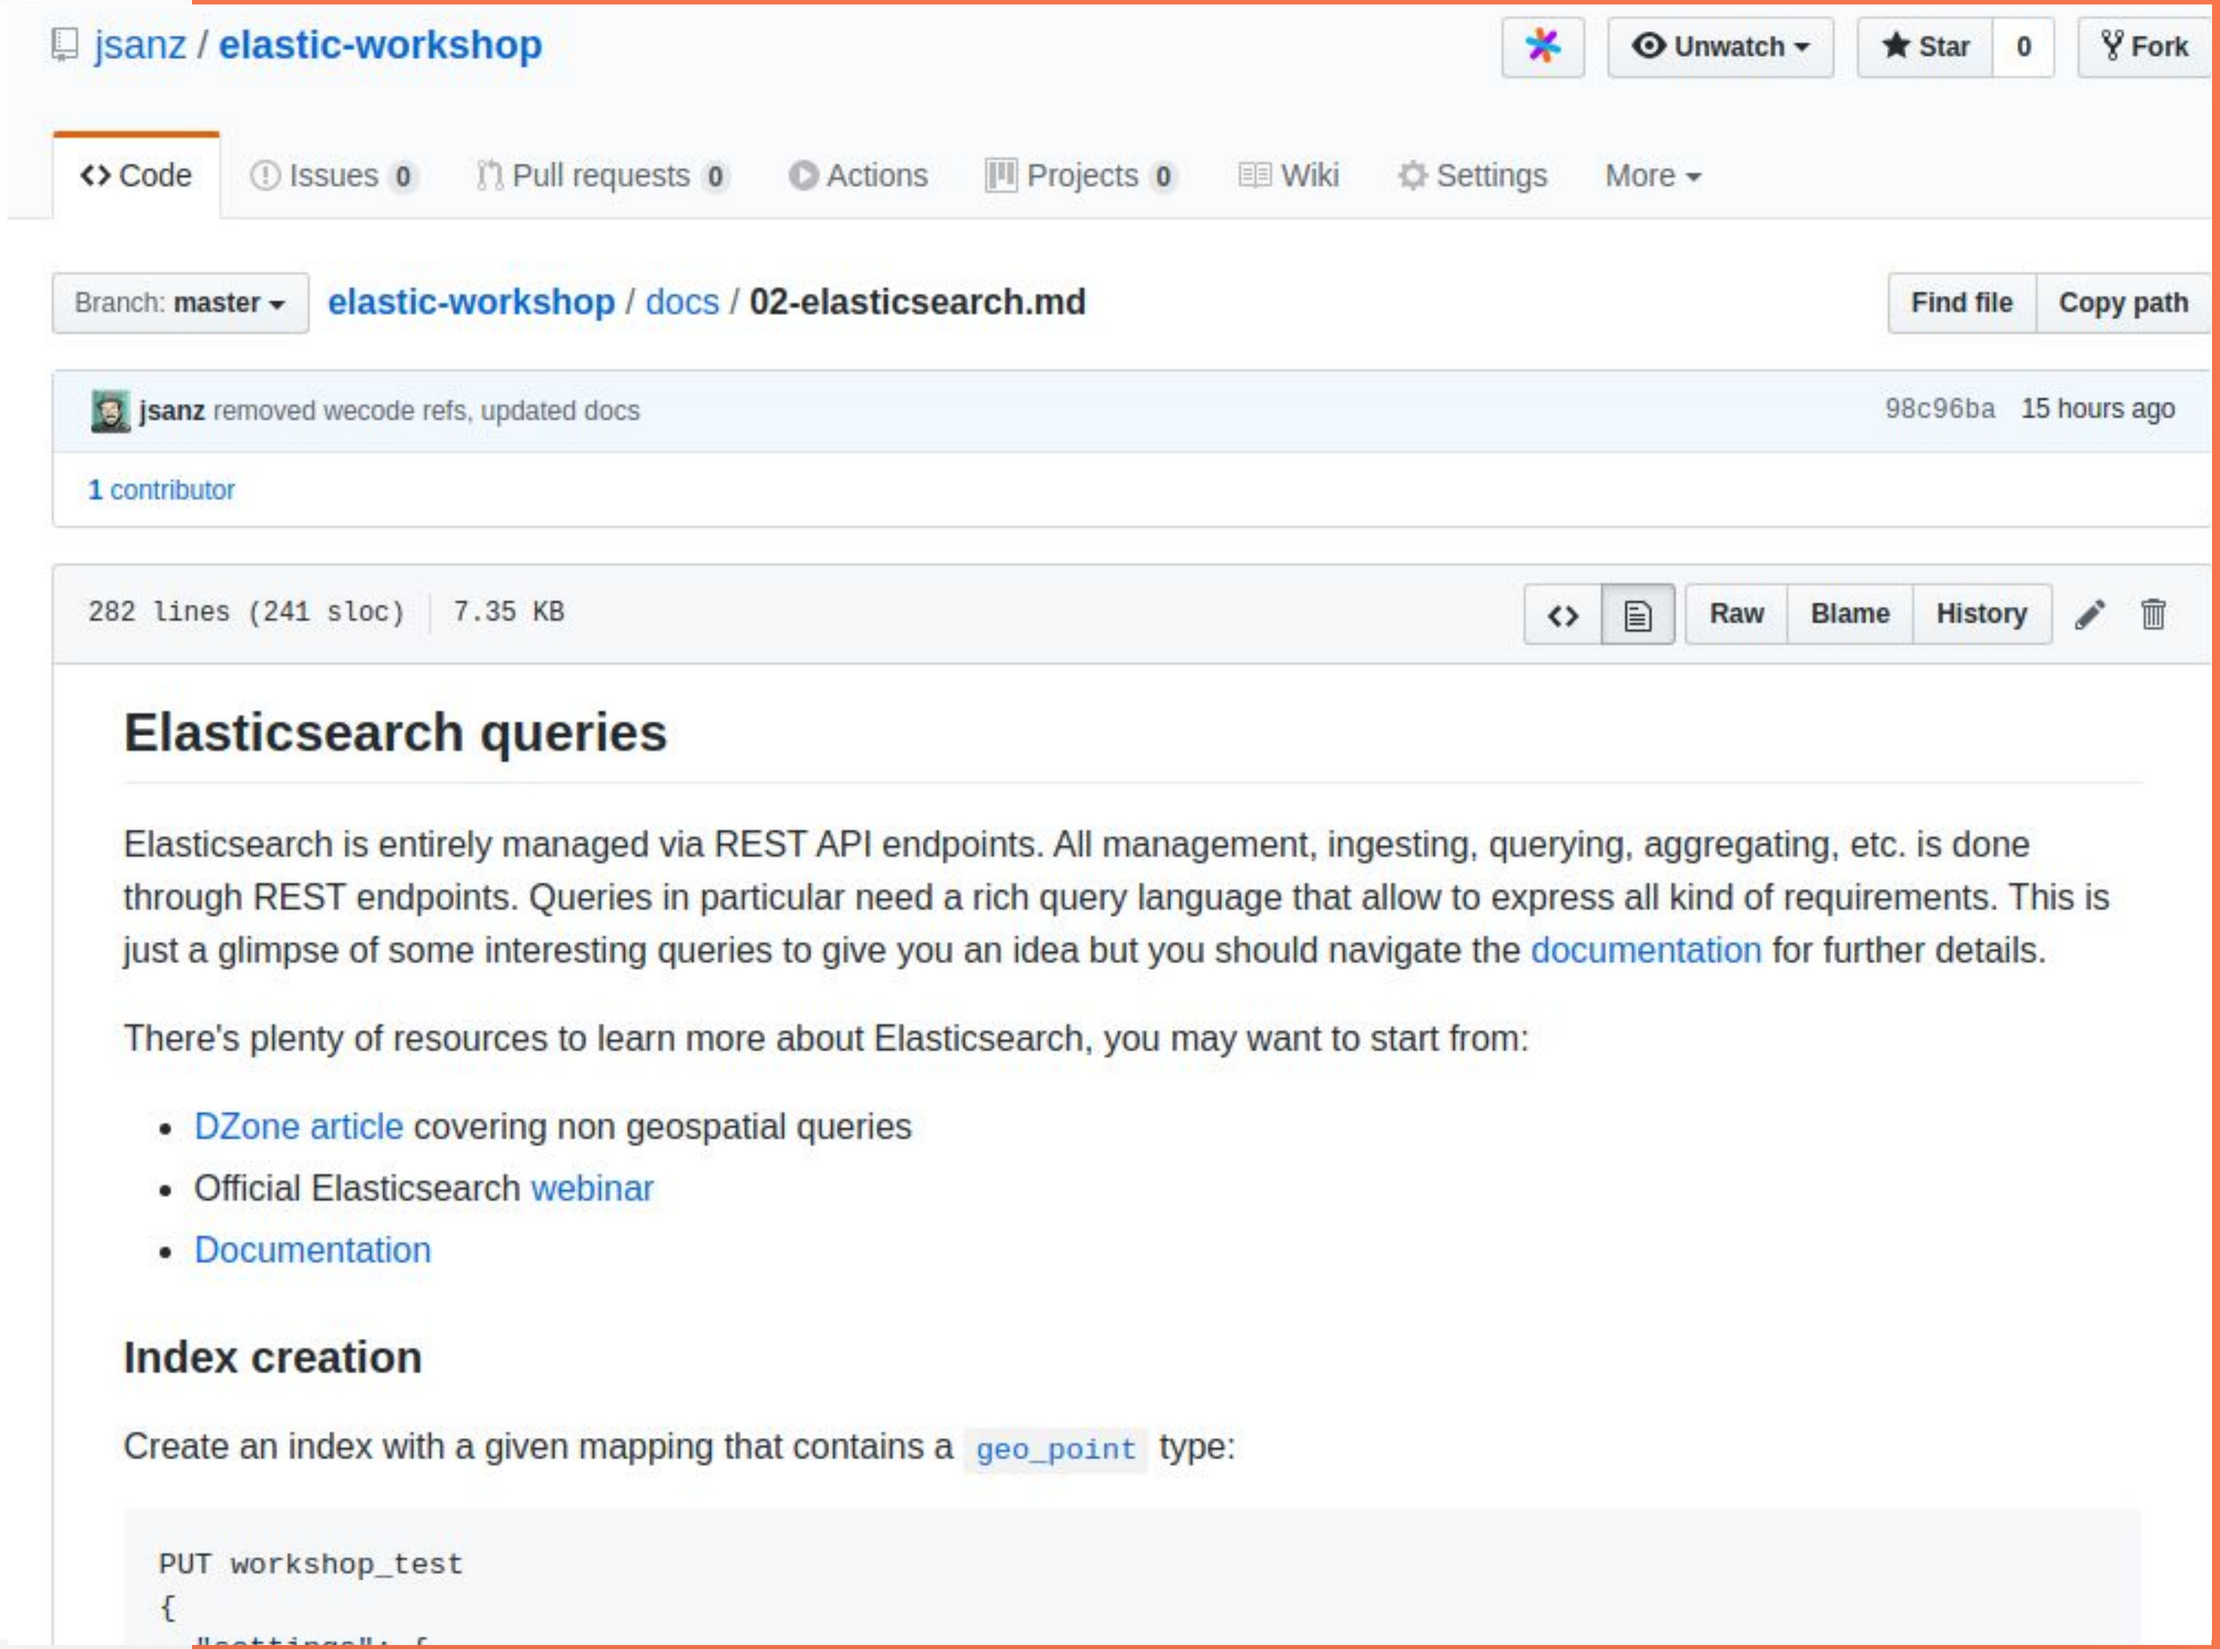

Elasticsearch queries

Elasticsearch is entirely managed via REST API endpoints. All management, ingesting, querying, aggregating, etc. is done through REST endpoints. Queries in particular need a rich query language that allow to express all kind of requirements. This is just a glimpse of some interesting queries to give you an idea but you should navigate the documentation for further details.

There's plenty of resources to learn more about Elasticsearch, you may want to start from:

- DZone article covering non geospatial queries
- Official Elasticsearch webinar
- Documentation

Index creation

Create an index with a given mapping that contains a `geo_point` type:

```
PUT workshop_test
{
  "settings": {
```

Geospatial queries

- Search
  - Find documents by point/radius, bounding box, polygons
- Metric aggregations
  - Find the centroid or the bounding box of your search results
- Aggregate
  - Bucket your results by geospatial definitions like rings or grids

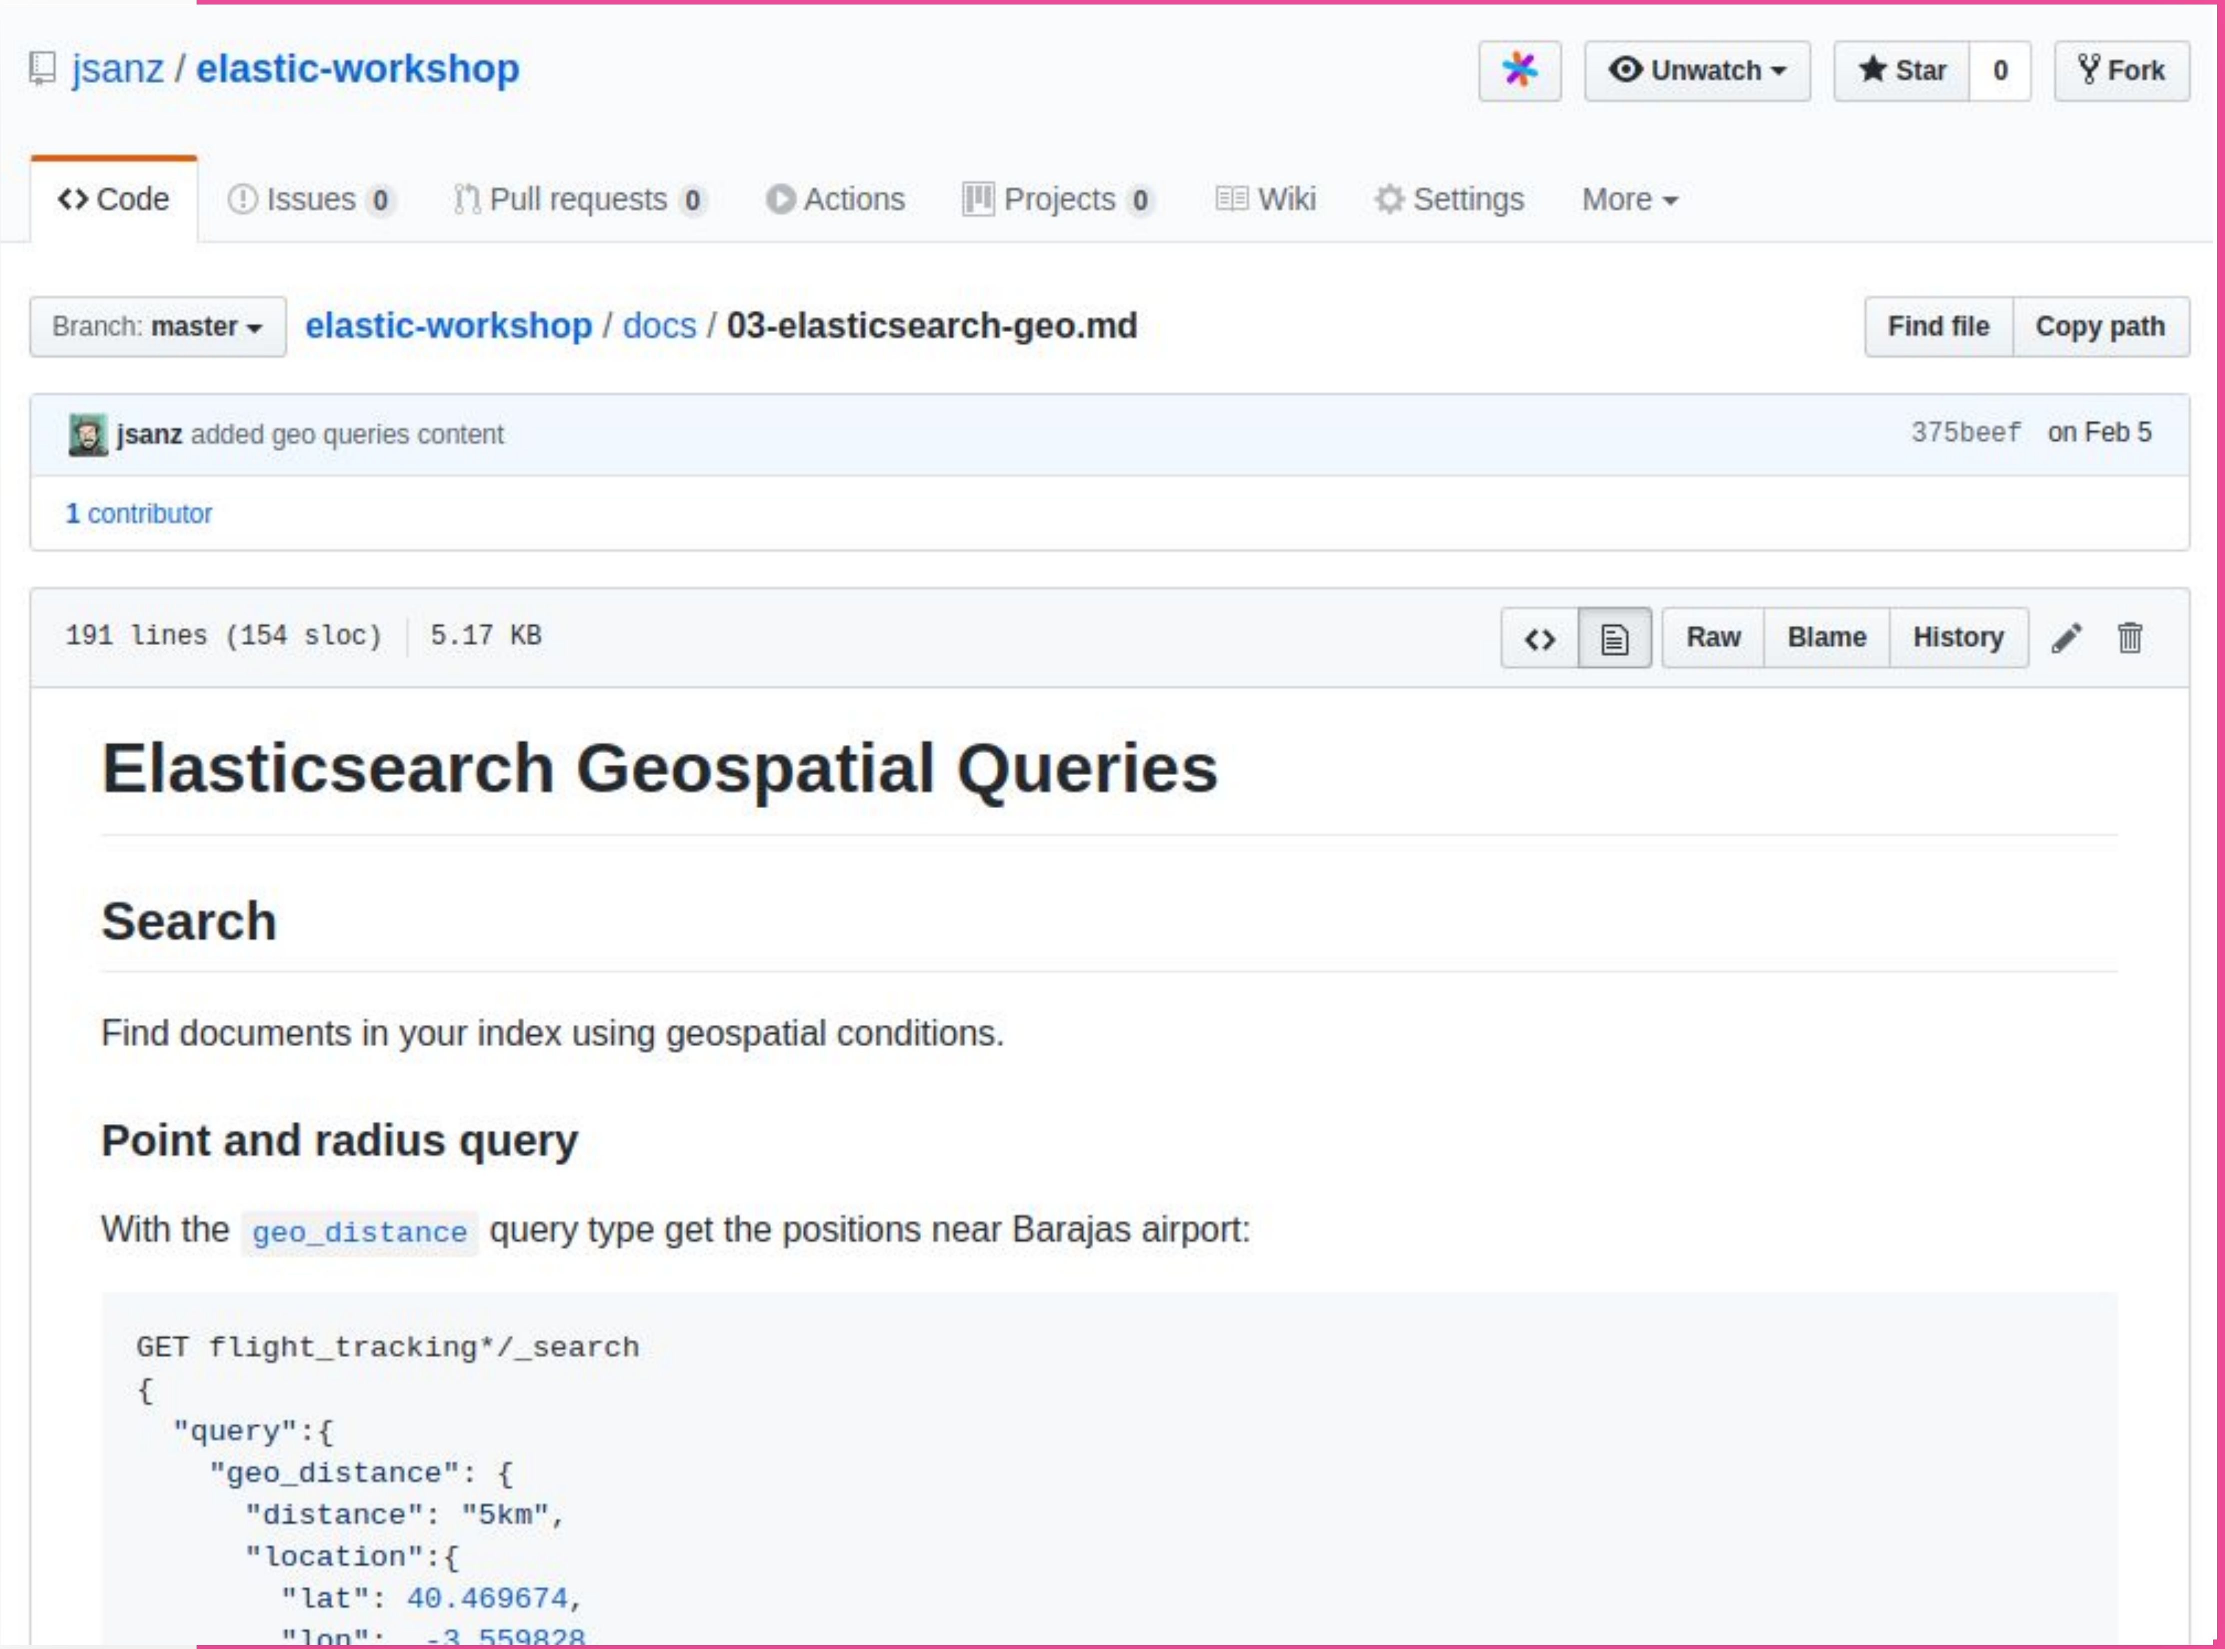

Elasticsearch Geospatial Queries

Search

Find documents in your index using geospatial conditions.

Point and radius query

With the `geo_distance` query type get the positions near Barajas airport:

```
GET flight_tracking*/_search
{
  "query":{
    "geo_distance": {
      "distance": "5km",
      "location":{
        "lat": 40.469674,
        "lon": -3.559828
```

Time to exercise

- Open the Kibana DevTools console
- Test the different queries from the github script
    - Elasticsearch basic queries
    - Geospatial queries

Agenda

1 Introduction to the Elastic stack

2 Lab setup

3 Kibana introduction

4 Elastics Maps

5 Elasticsearch Geo

6 Web mapping and Elasticsearch

Quick webmapping intro

elastic

tangram js
deck gl

Leaflet

OpenLayers

MapLibre

mapbox

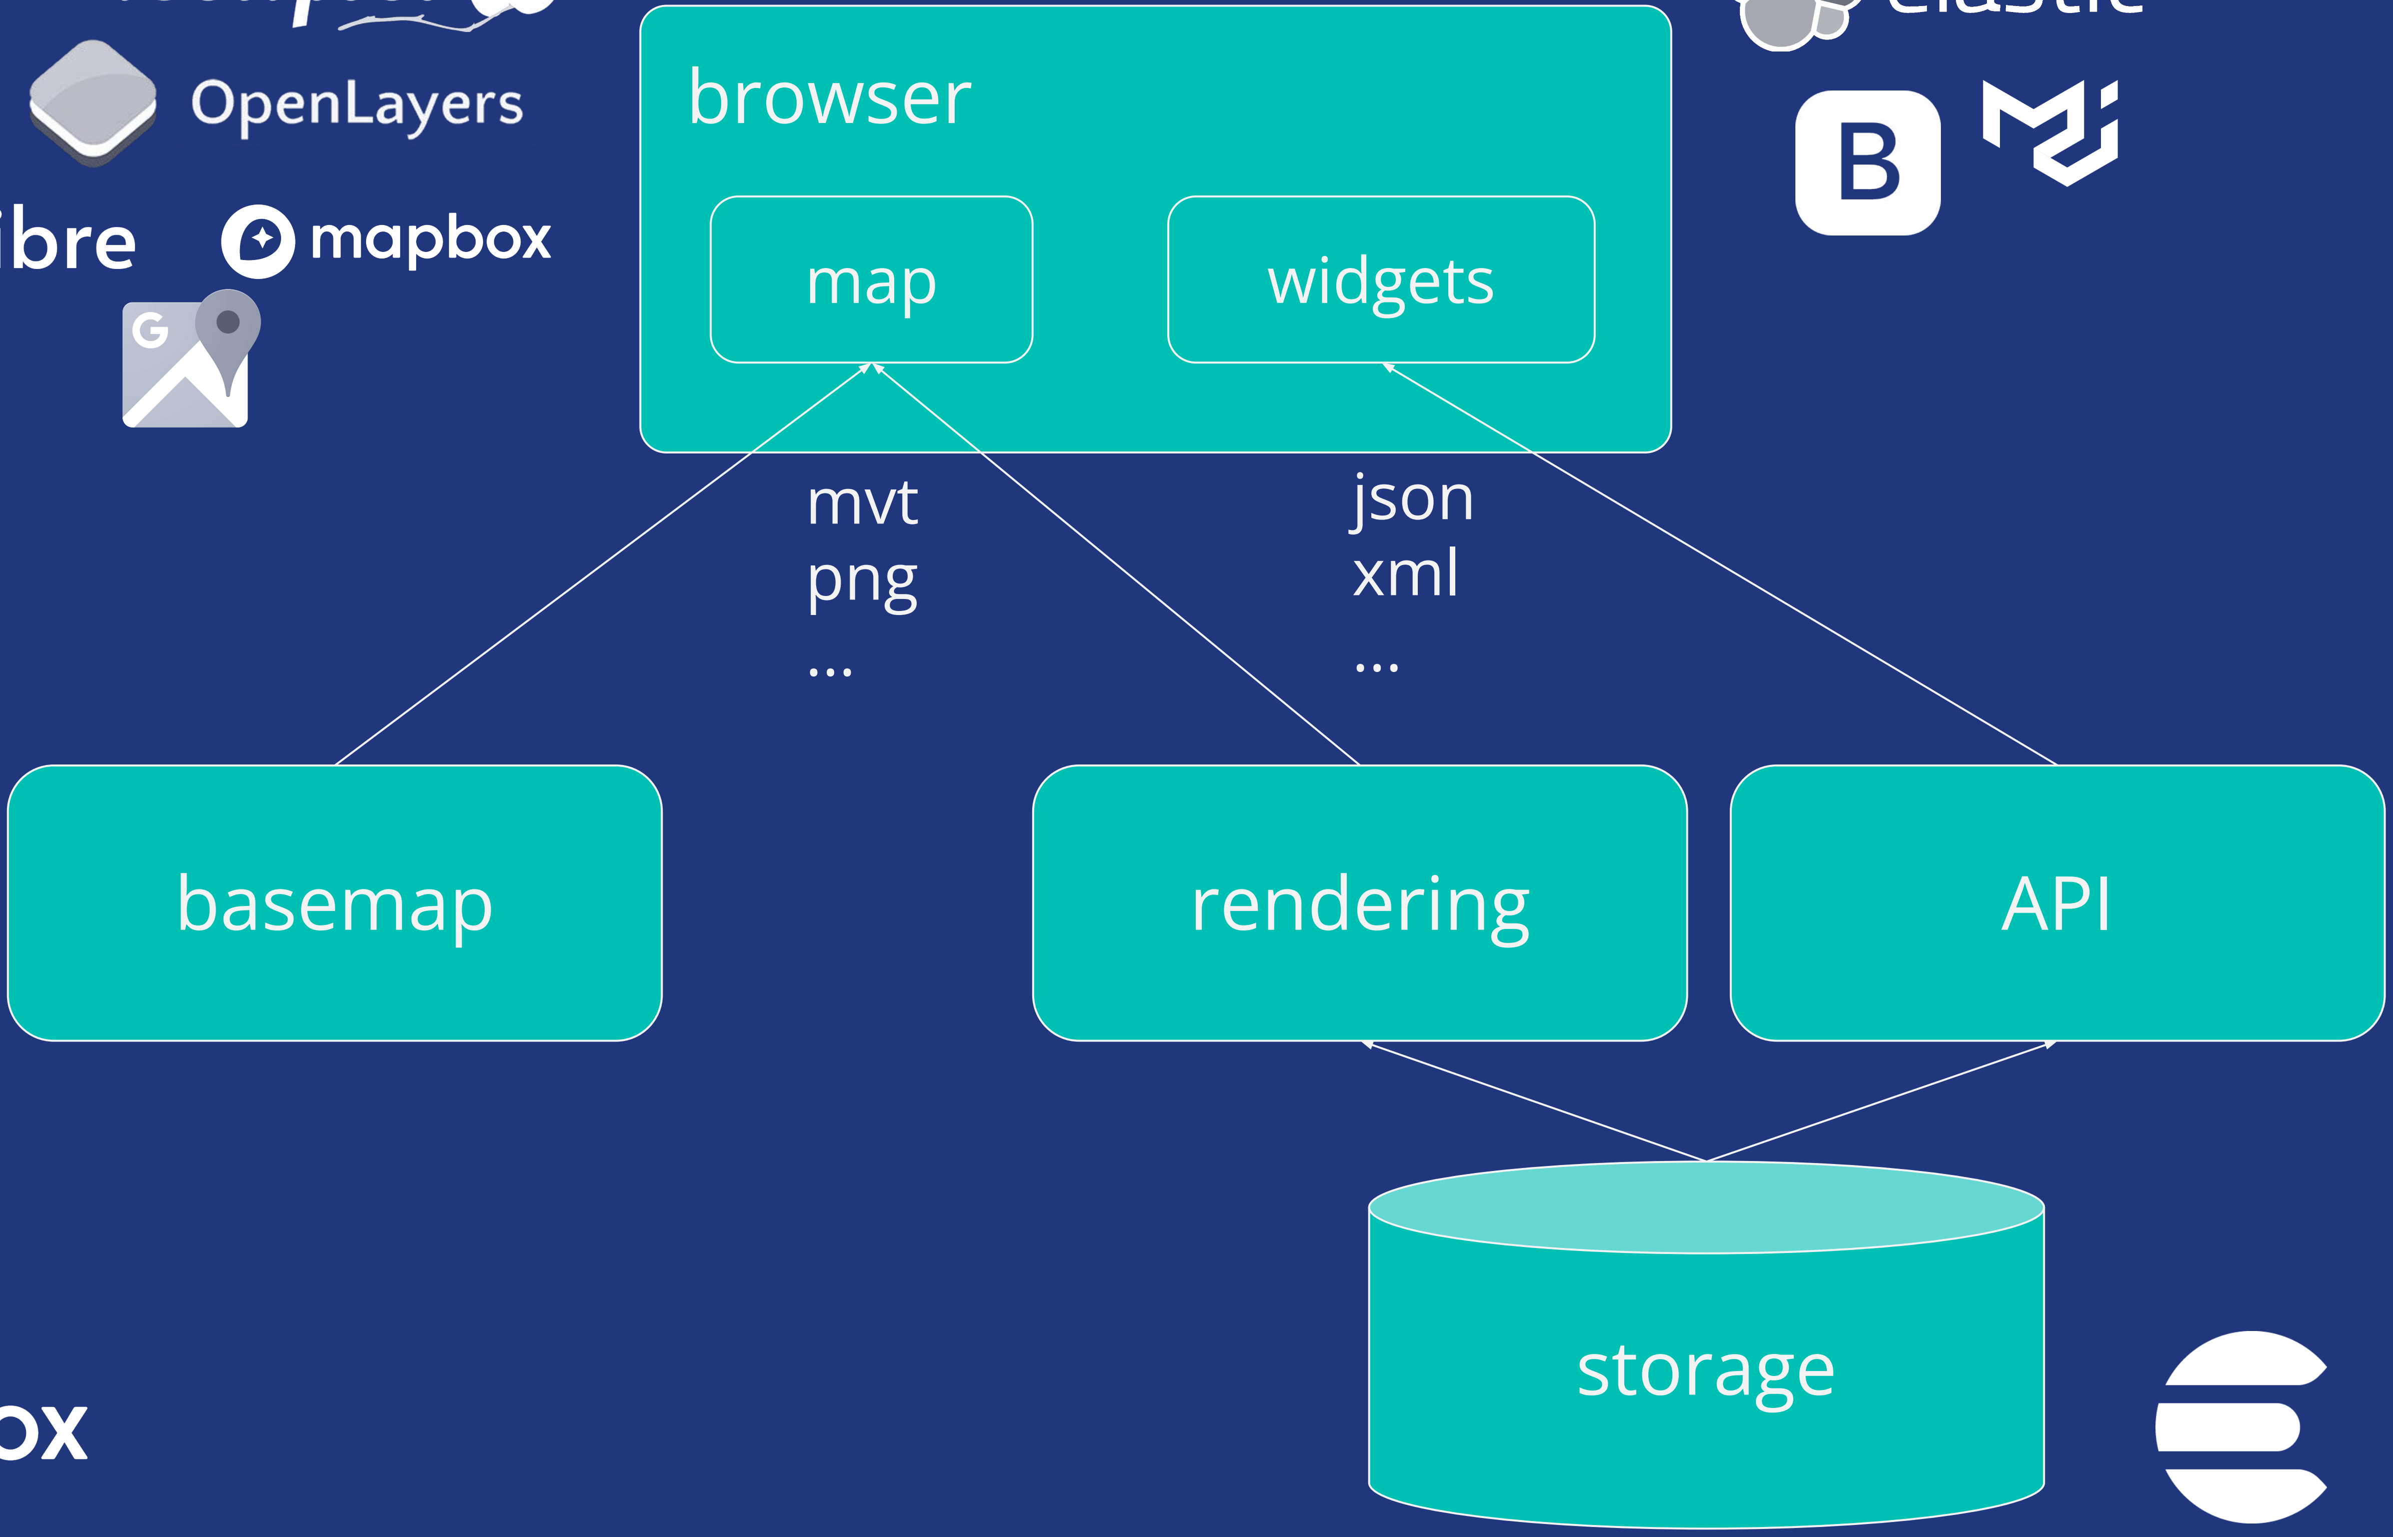

browser

map

widgets

mvt
png
…

json
xml
…

basemap

rendering

API

storage

elastic

GeoServer

maptiler

CARTO

mapbox

here

PostGIS
Spatial PostgreSQL

elastic

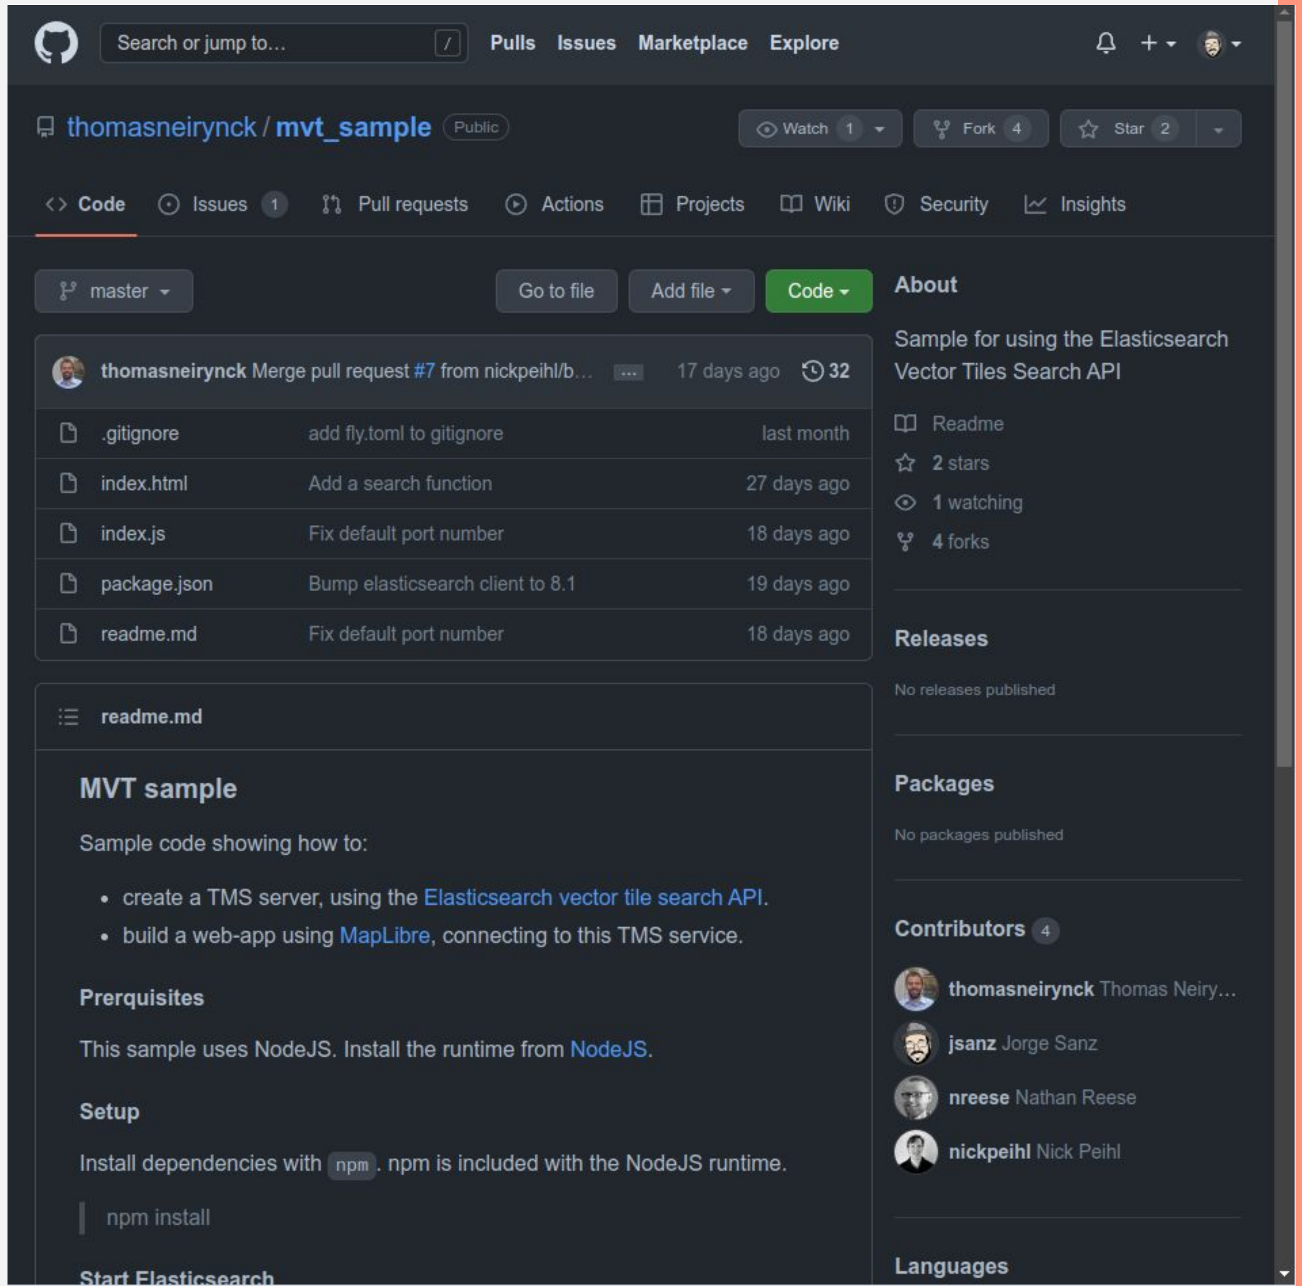
Search or jump to...
Pulls Issues Marketplace Explore
Watch 1 Fork 4 Star 2
Code Issues 1 Pull requests Actions Projects Wiki Security Insights
master
Go to file Add file Code

thomasneirynck / mvt_sample Public

thomasneirynck Merge pull request #7 from nickpeihl/b… 17 days ago 32

.gitignore	add fly.toml to gitignore	last month
index.html	Add a search function	27 days ago
index.js	Fix default port number	18 days ago
package.json	Bump elasticsearch client to 8.1	19 days ago
readme.md	Fix default port number	18 days ago

readme.md

MVT sample

Sample code showing how to:

- create a TMS server, using the Elasticsearch vector tile search API.
- build a web-app using MapLibre, connecting to this TMS service.

Prerquisites

This sample uses NodeJS. Install the runtime from NodeJS.

Setup

Install dependencies with `npm`. npm is included with the NodeJS runtime.

```
npm install
```

Start Elasticsearch

About

Sample for using the Elasticsearch Vector Tiles Search API

Readme
2 stars
1 watching
4 forks

Releases

No releases published

Packages

No packages published

Contributors 4

thomasneirynck Thomas Neiry…
jsanz Jorge Sanz
nreese Nathan Reese
nickpeihl Nick Peihl

Languages

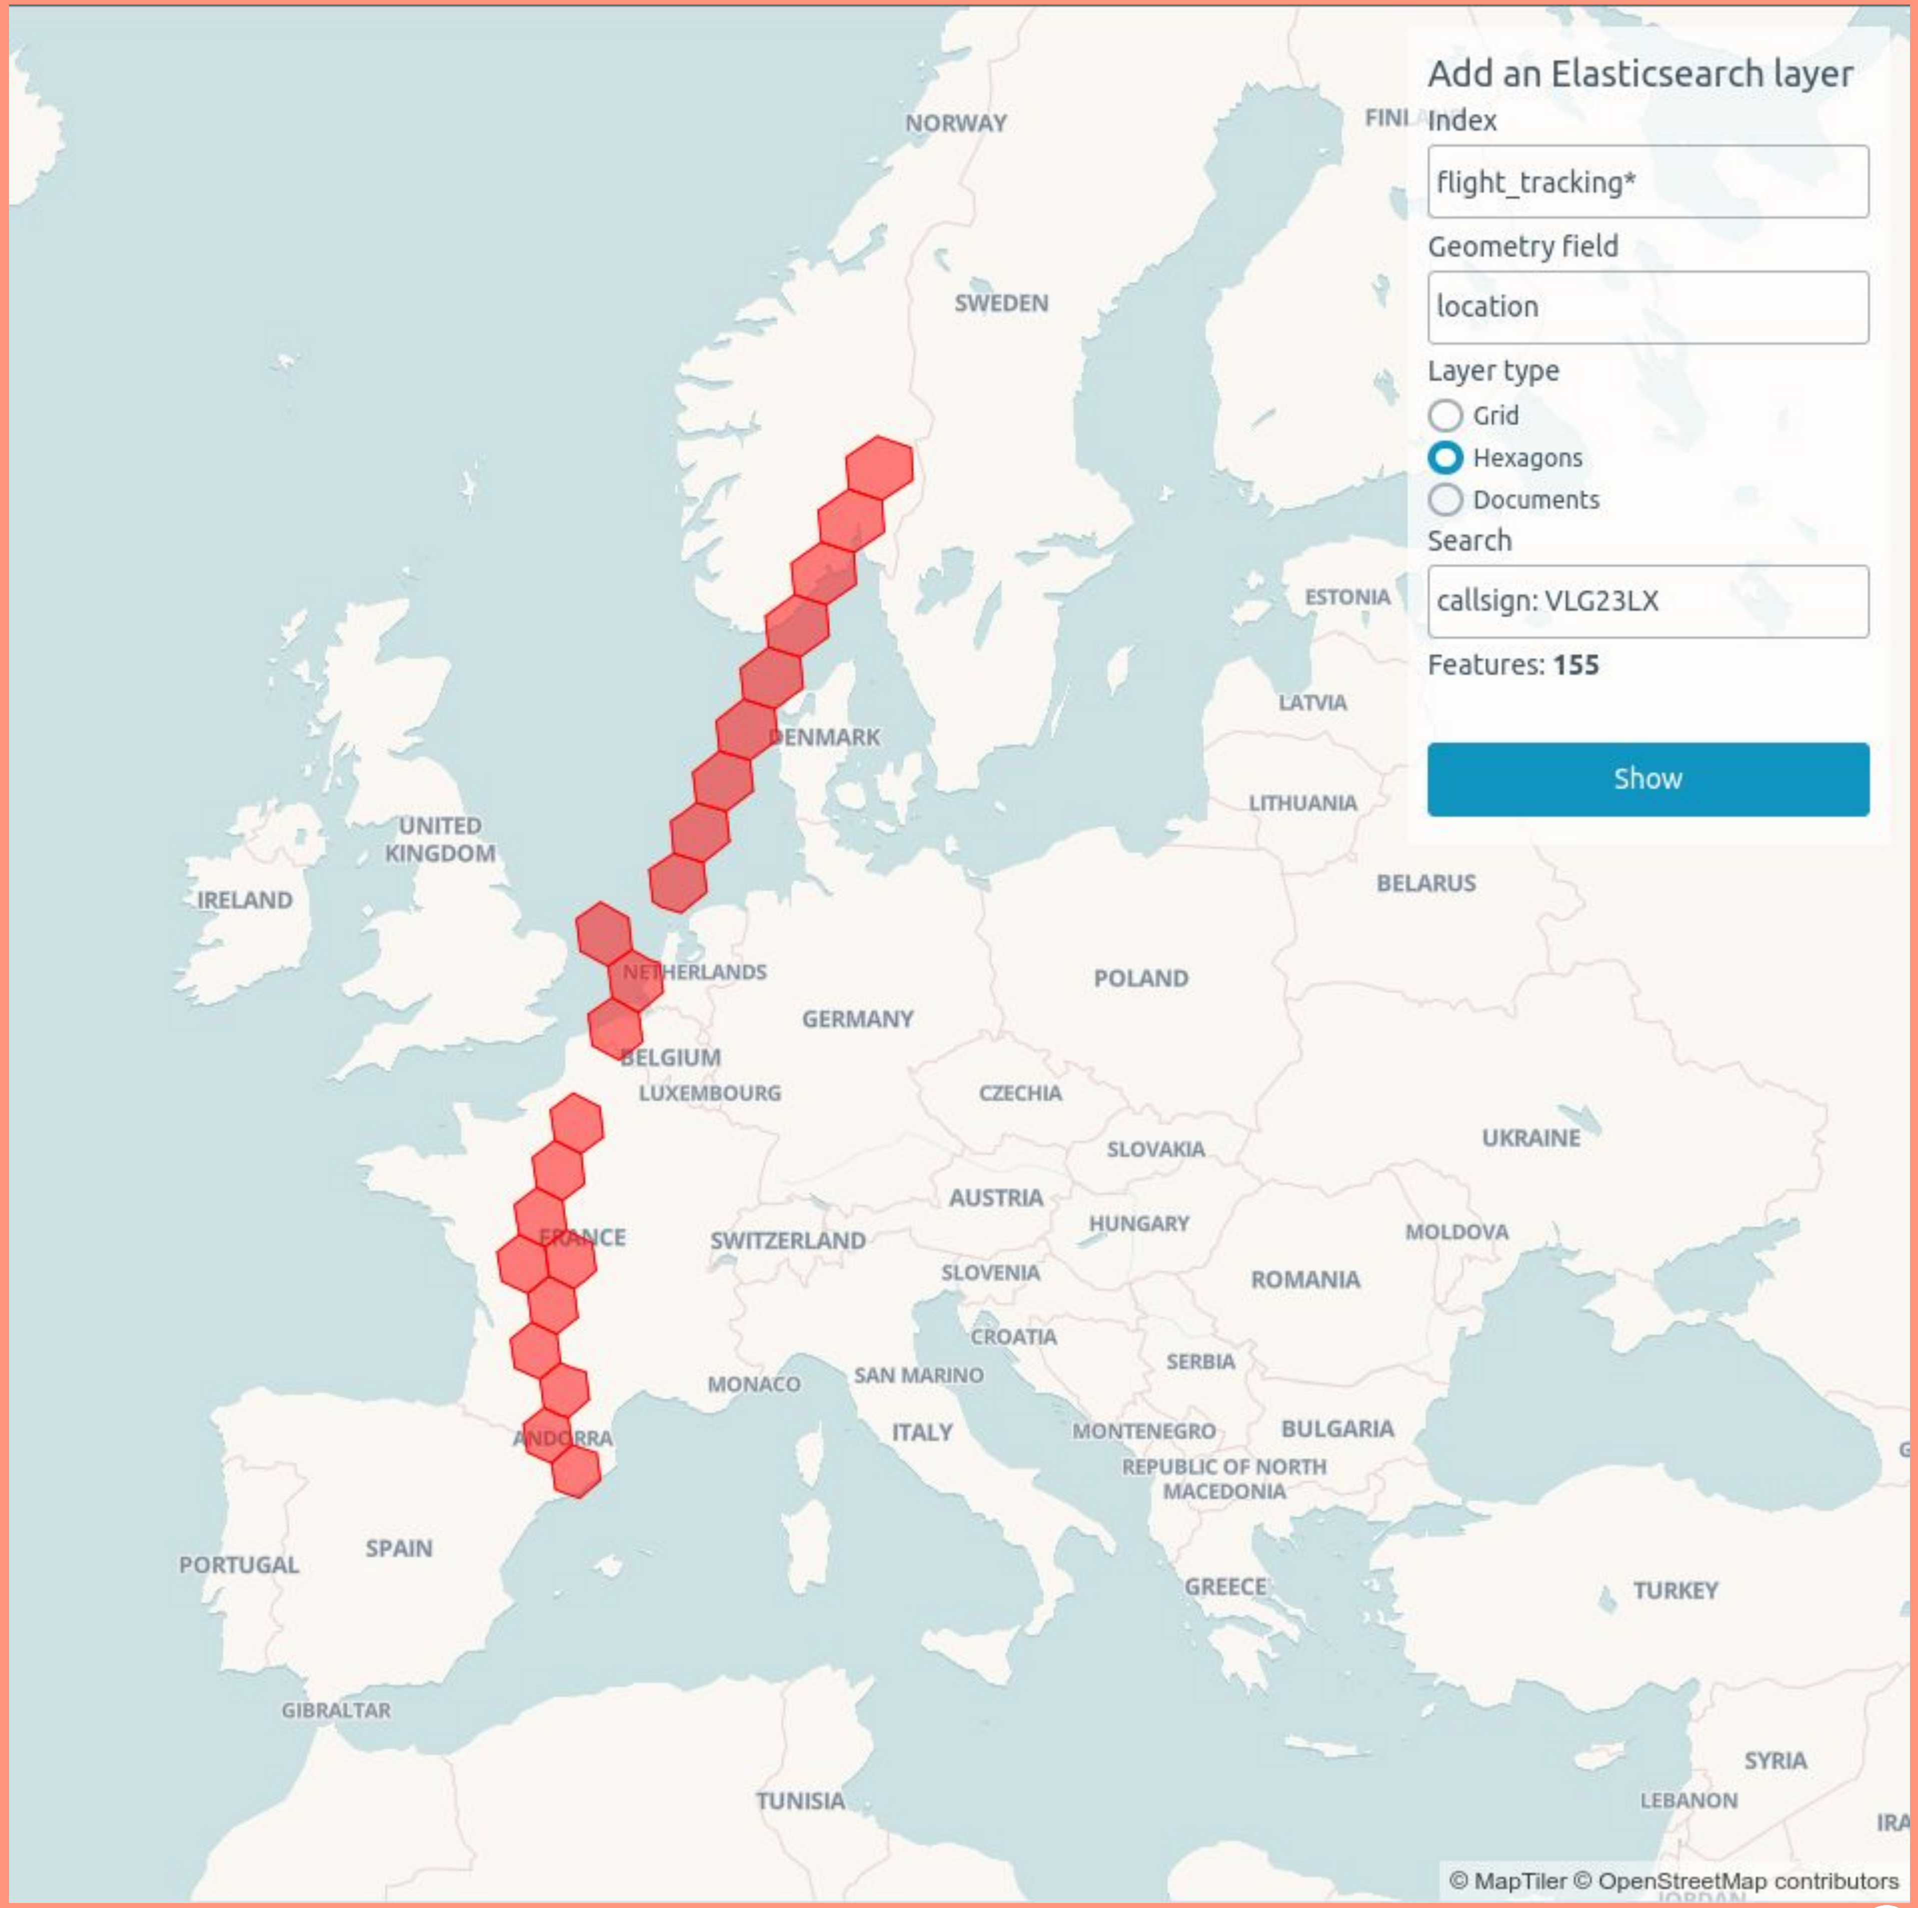

Add an Elasticsearch layer

Index: flight_tracking*

Geometry field: location

Layer type

○ Grid
● Hexagons
○ Documents

Search: callsign: VLG23LX

Features: **155**

© MapTiler © OpenStreetMap contributors

- A tile server: nodejs
- A generic viewer: Maplibre

https://github.com/thomasneirynck/mvt_sample

Vector tiles from Elasticsearch

1. Browser requests a tile
2. Middleware gathers **query parameters**: aggregation, search filters, etc
3. Database performs the query **and returns a vector tile**
4. Middleware **forwards** results to client
5. **Fast** rendering

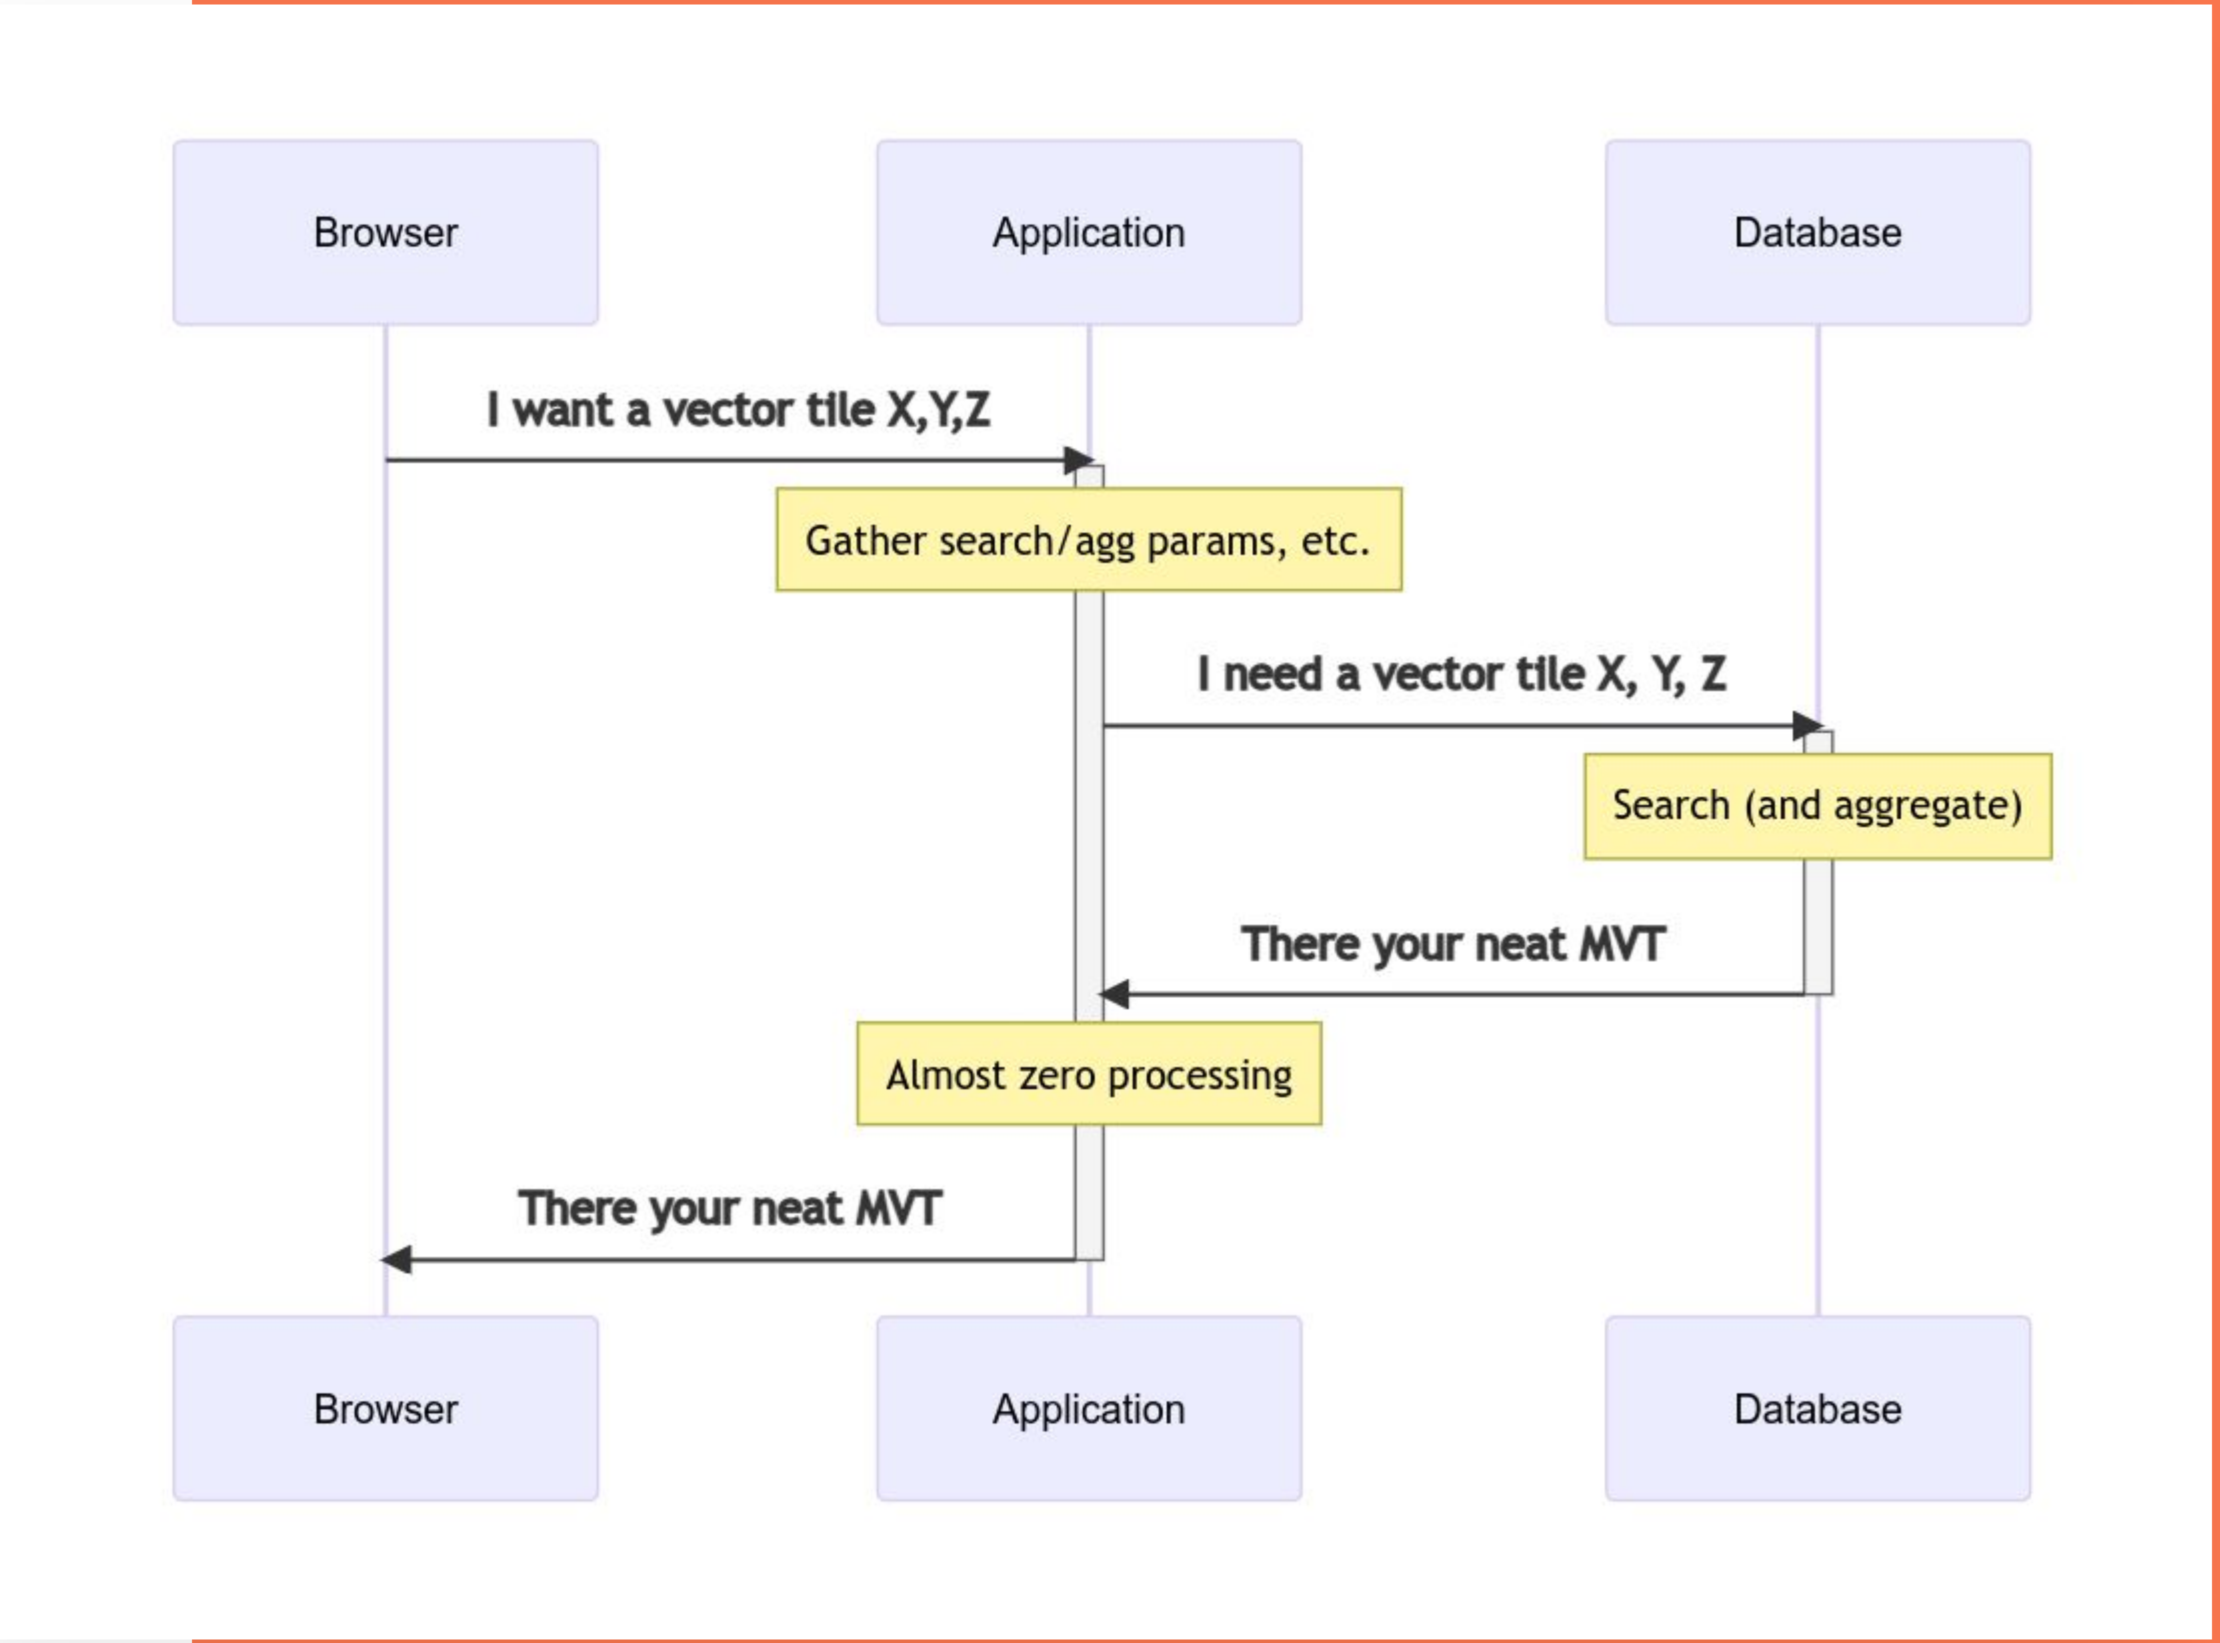

Vector tiles from Elasticsearch

1. `_mvt` endpoint 📖
2. Elasticsearch outputs mapbox **vector tiles** in *protobuff* format
3. Can render up to **10.000** documents per tile
4. Geometries are **simplified**
5. `meta` layer with **details**
6. Optional **label** positions

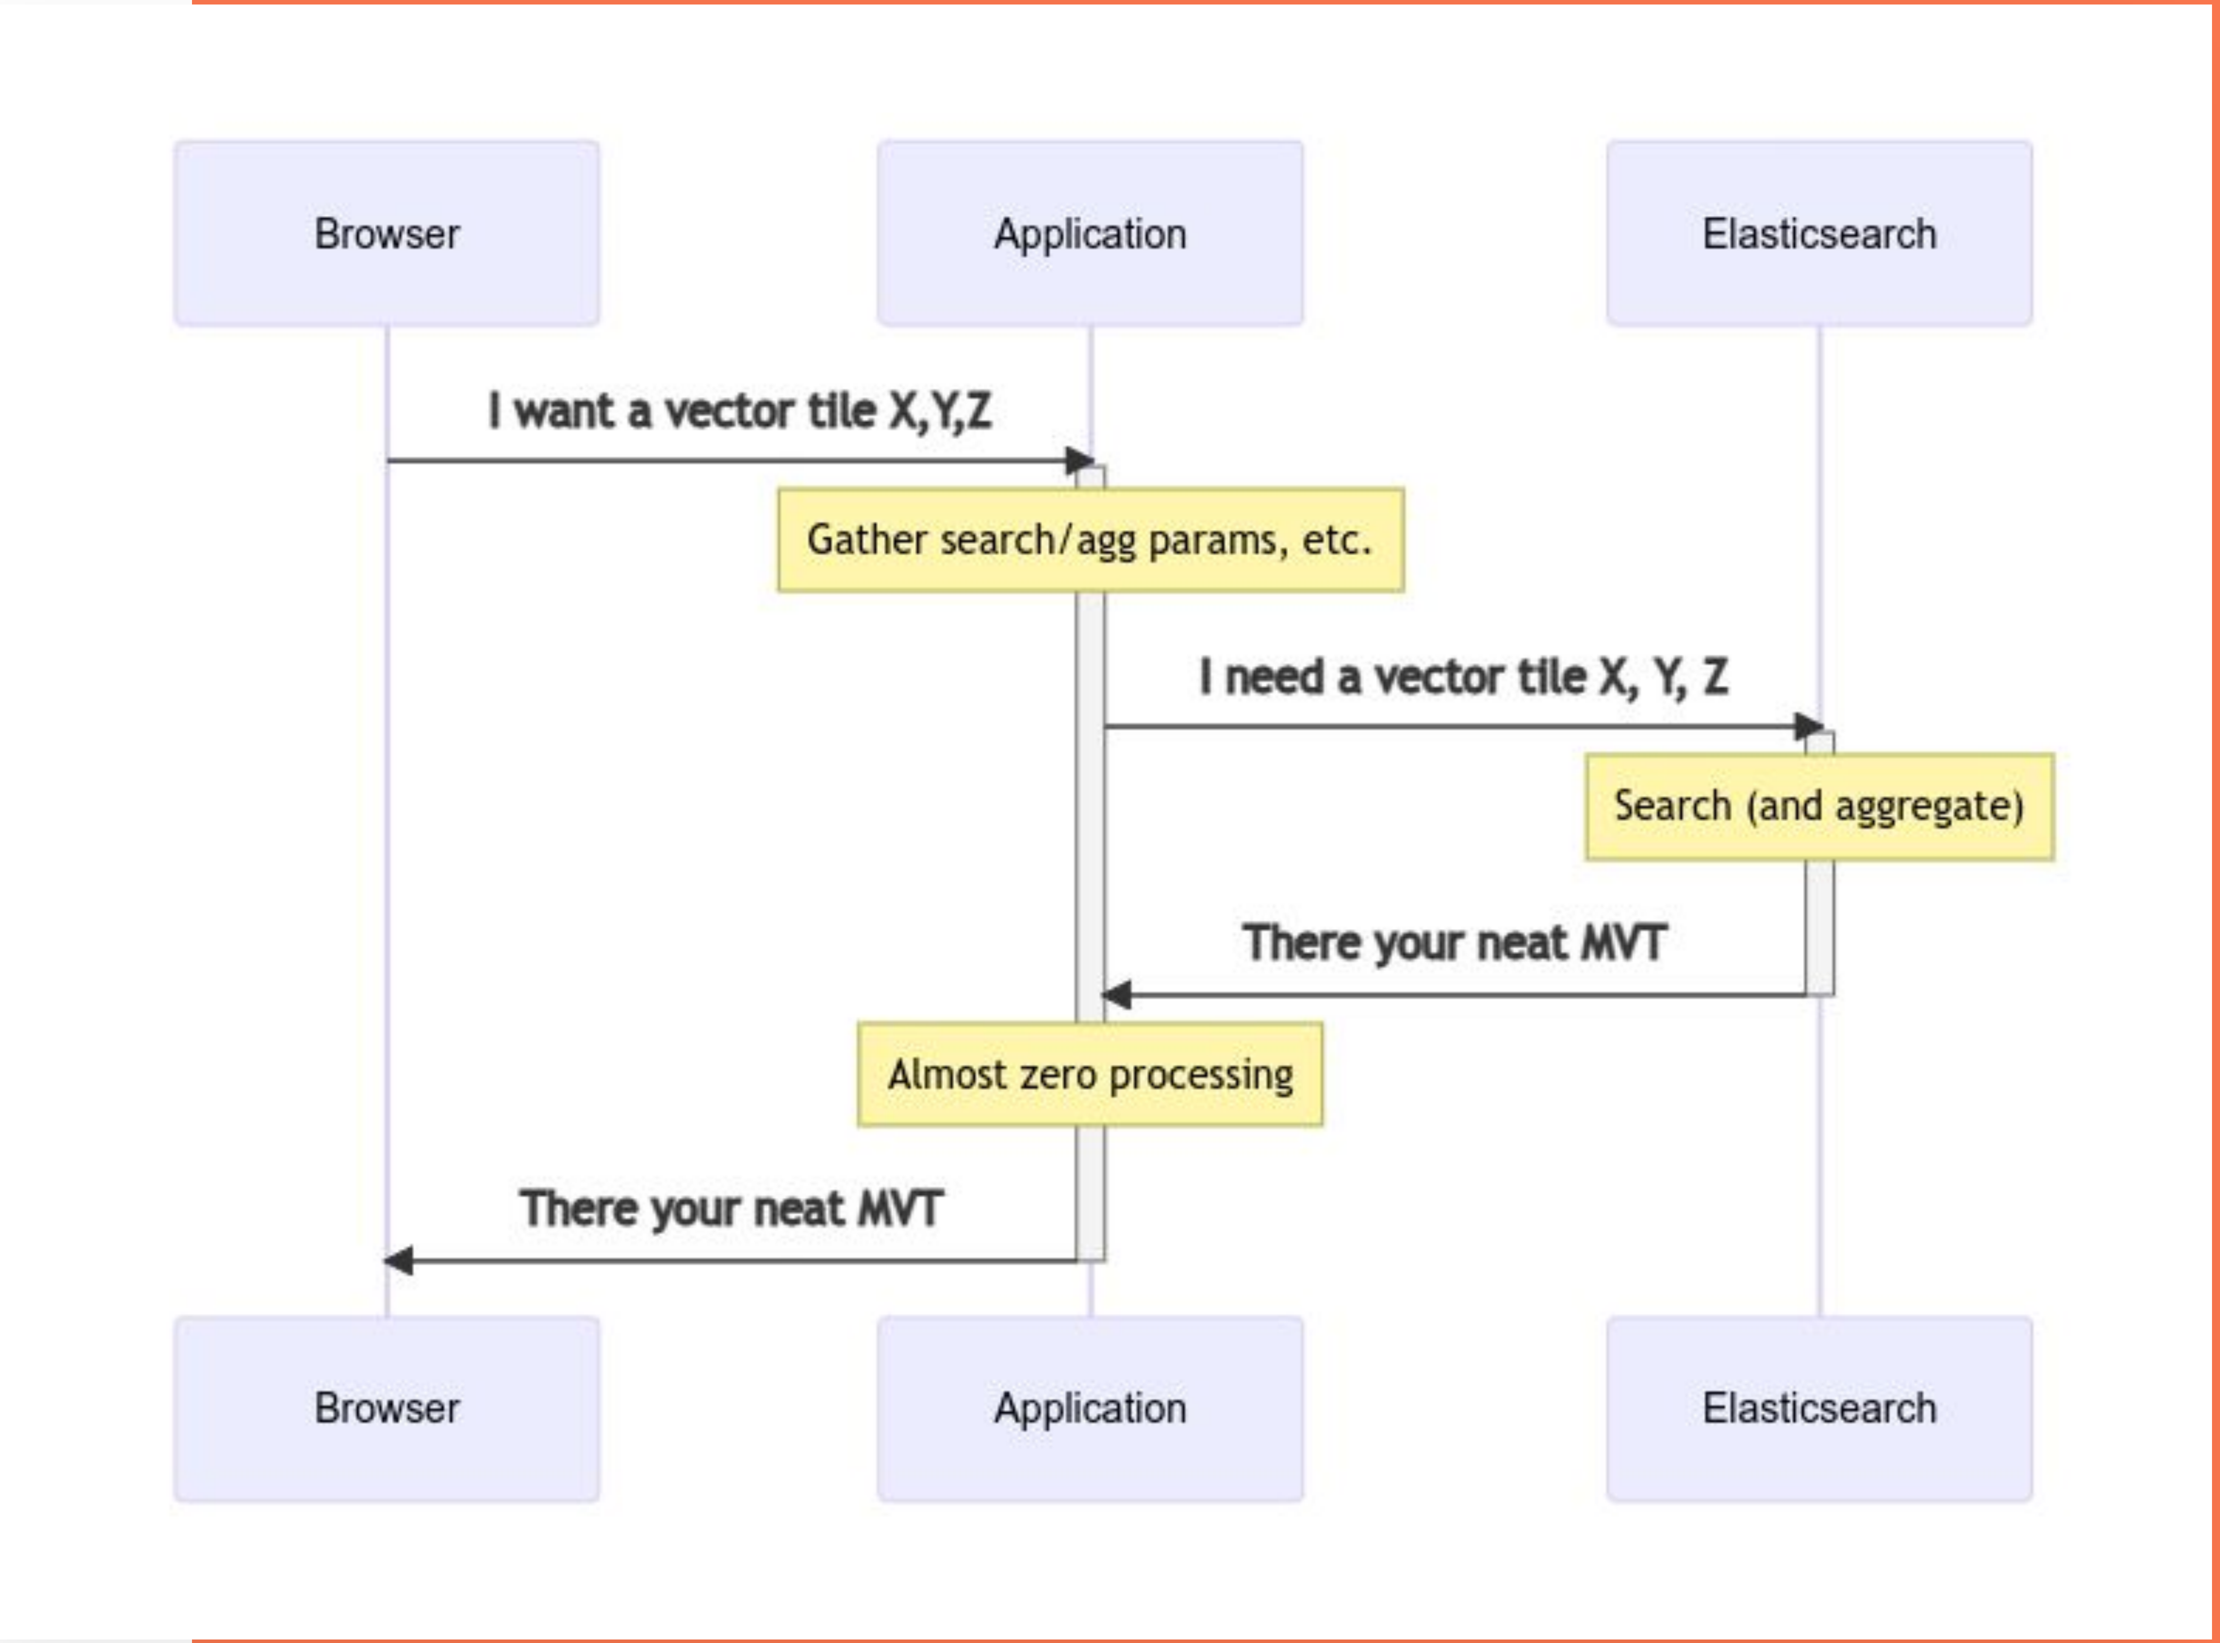

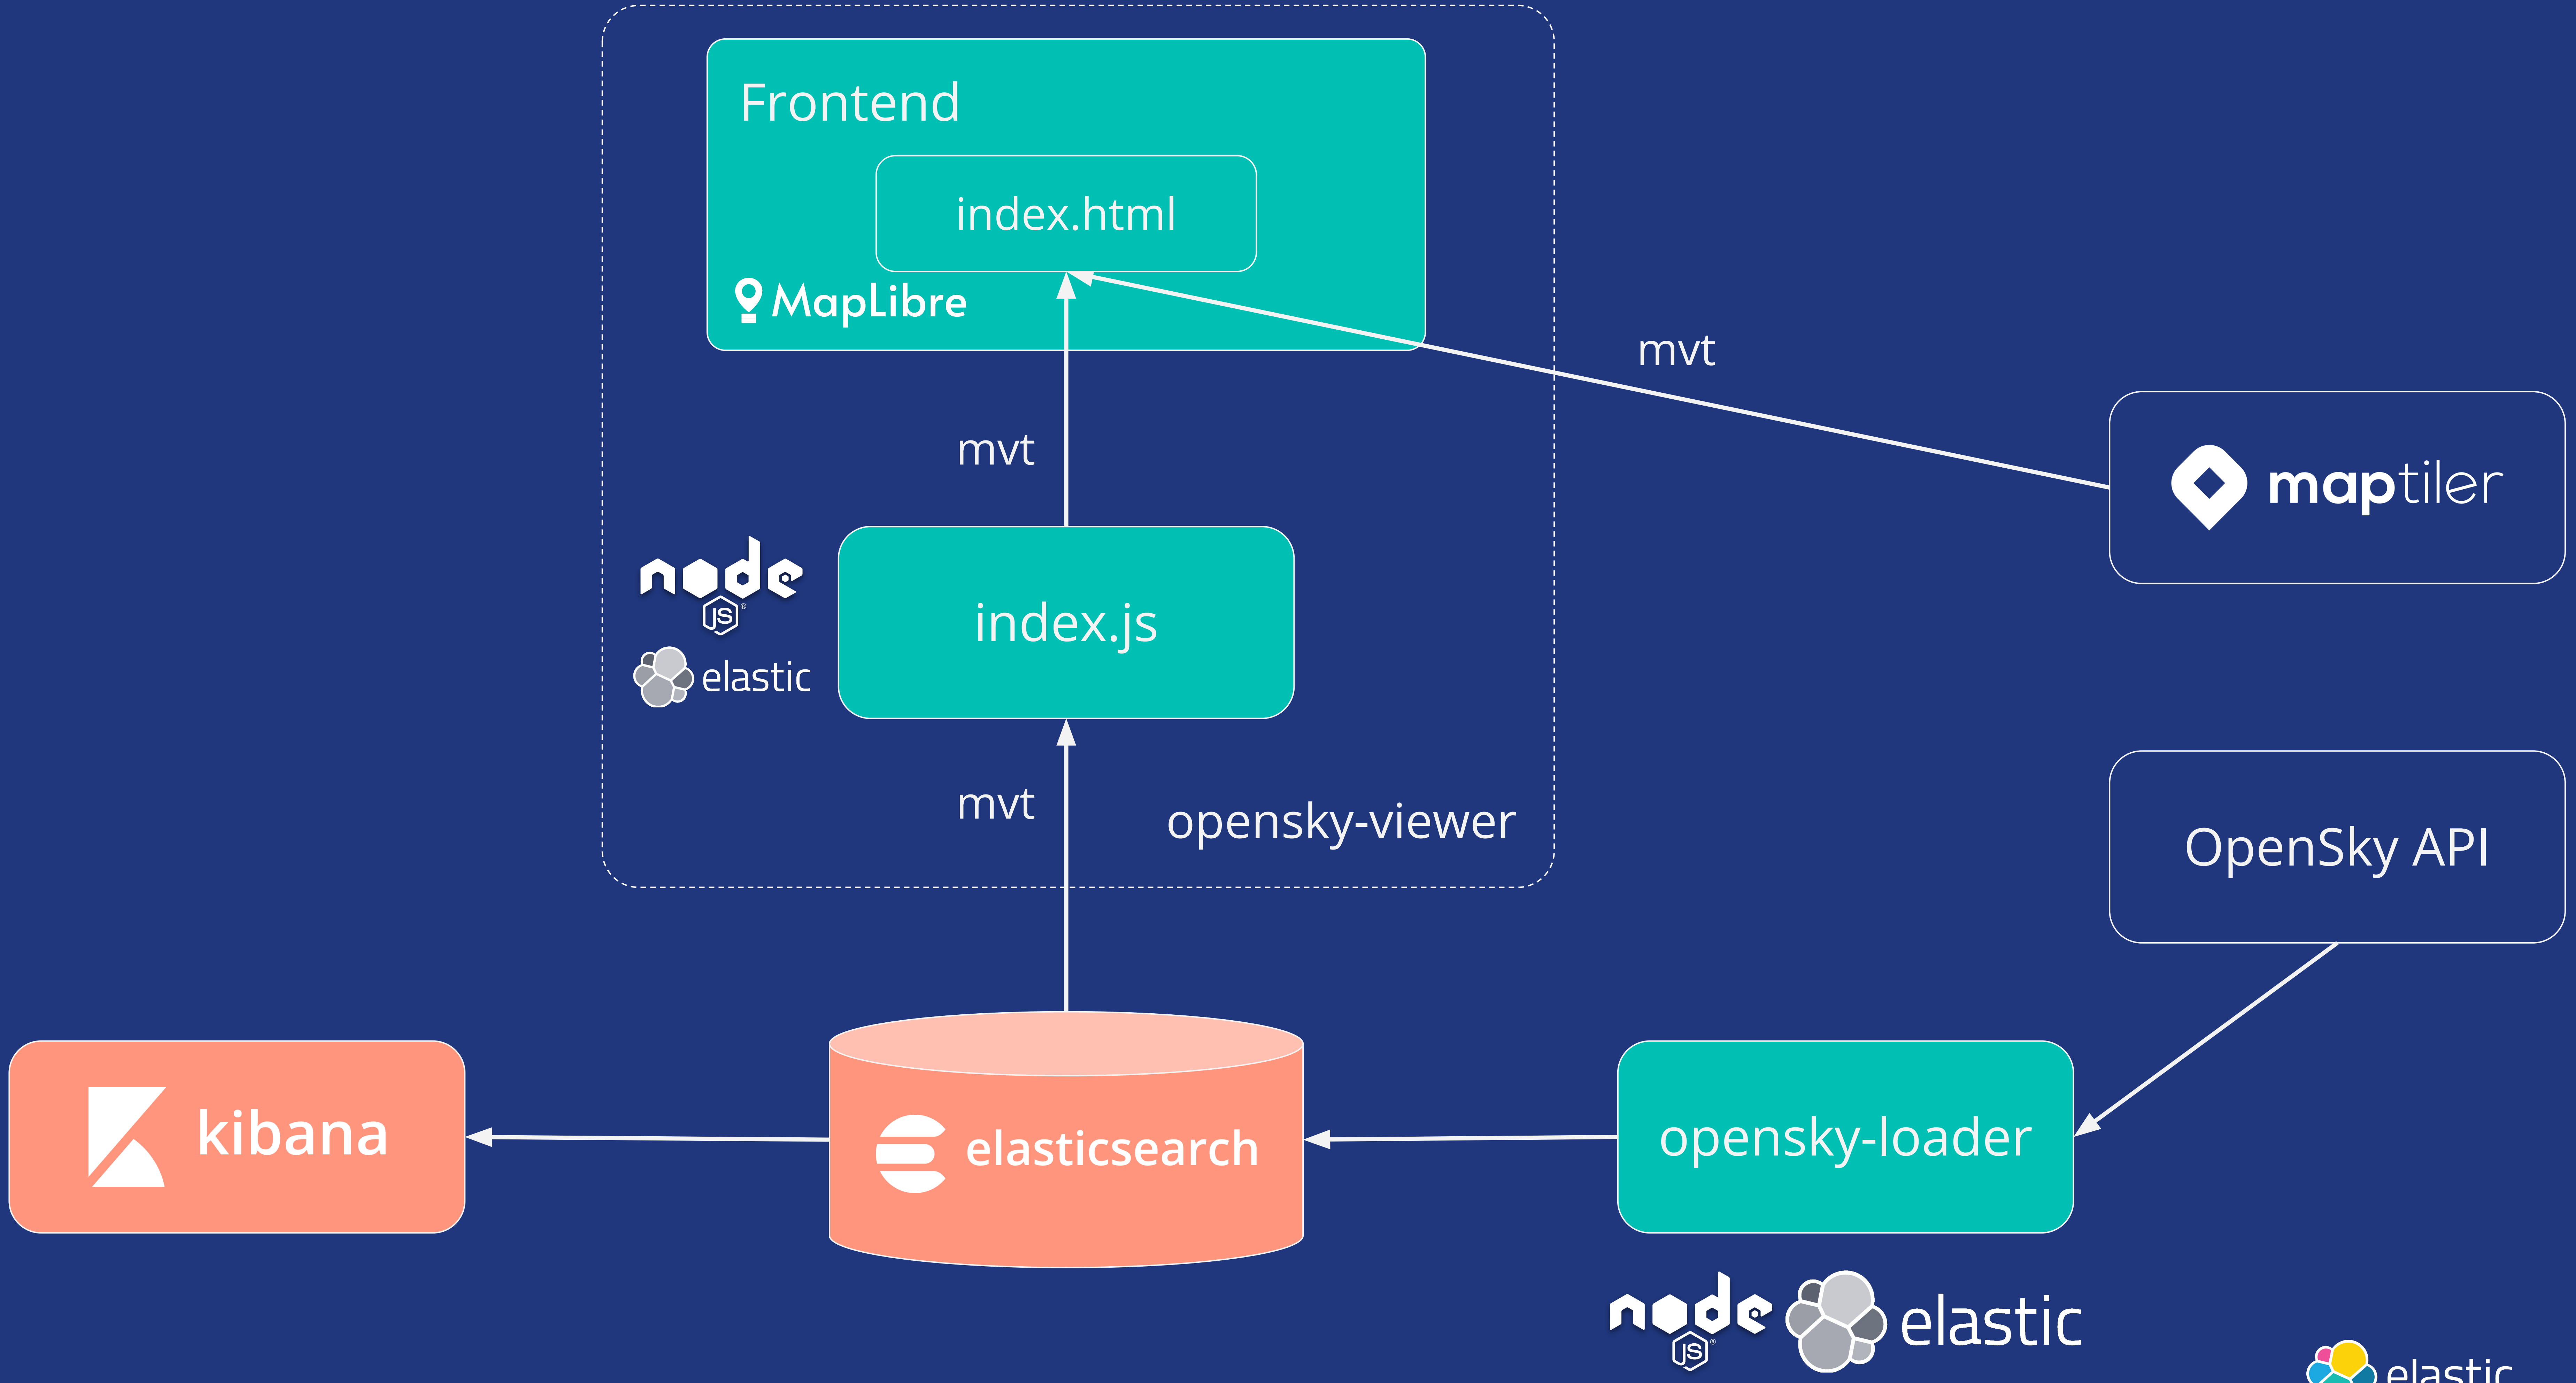
Frontend
index.html
MapLibre
mvt
maptiler
mvt
node JS
elastic
index.js
mvt
opensky-viewer
OpenSky API
kibana
elasticsearch
opensky-loader
node JS
elastic

elastic

Time to exercise: data loading

- /lab/opensky-loader/index.js
- Nodejs application using the elasticsearch JS client
- Review the data load workflow
- How is the index created?
- How is data updated before uploading?
- How is the Bulk API used?

Time to exercise: viewer

https://github.com/thomasneirynck/mvt_sample

Backend

- /lab/opensky-viewer/index.js
- Root route serves the web app
- /tile route controller
- Review the tile parameters
- Check how the Elasticsearch query is built
- Check the additional HTTP headers added

Frontend

- /lab/opensky-viewer/index.html
- Maplibre and a simple form
- Layers for polygons, lines, and points
- Check the Maplibre *source*
- Check the styles
- Review the *callback* for the feature counter

¡Gracias!

Jorge Sanz , Kibana, Elastic
xurxosanz jsanz jorge.sanz@elastic.co
2022-08-22

https://ela.st/foss4g22-workshop

elastic