Web mapping with Elasticsearch

Jorge Sanz | jorge.sanz@elastic.co

June 2023

Web mapping with Elasticsearch

Jorge Sanz | jorge.sanz@elastic.co

June 2023

🎓 Cartography and Geodesy Engineer

👩‍💻 GIS consultant at Prodevelop

🗺️ Solutions Engineer & Support Manager at CARTO

📊 Principal Software Engineer at Elastic

Jorge Sanz

Forward Looking Statements and Non-GAAP Disclaimer

This presentation and the accompanying oral presentation contain forward-looking statements that involve substantial risk and uncertainties, which include, but are not limited to, our expected financial results for the fiscal quarter ending July 31, 2020 and the fiscal year ending April 30, 2021, our expectations regarding the impact of the COVID-19 pandemic, our customer base, potential market and growth opportunities, and our go-to-market strategy. These forward-looking statements are subject to the safe harbor provisions under the Private Securities Litigation Reform Act of 1995. In some cases, you can identify forward-looking statements because they contain words such as “may,” “will,” “should,” “would,” “expects,” “plans,” “anticipates,” “could,” “intends,” “target,” “projects,” “contemplates,” “believes,” “estimates,” “predicts,” “potential” or “continue” or the negative of these words or other similar terms or expressions that concern our expectations, strategy, plans or intentions. Our expectations and beliefs in light of currently available information regarding these matters may not materialize. Actual outcomes and results may differ materially from those contemplated by these forward-looking statements due to uncertainties, risks, and changes in circumstances, including but not limited to those related to: the impact of COVID-19 on our business, operations, hiring and financial results, and on businesses of our customers and partners, including the effect of governmental lockdowns, restrictions and new regulations; our future financial performance, including our expectations regarding our revenue, cost of revenue, gross profit or gross margin, operating expenses (which include changes in sales and marketing, research and development and general and administrative expenses), and our ability to achieve and maintain future profitability; our ability to continue to deliver and improve our offerings and successfully develop new offerings, including security-related product offerings and SaaS offerings; customer acceptance and purchase of our existing offerings and new offerings, including the expansion and adoption of our SaaS offerings; our ability to maintain and expand our user and customer base; the impact of foreign currency exchange rate and interest rate fluctuations on our results; our international expansion strategy; our operating results and cash flows; our beliefs and objectives for future operations; the sufficiency of our capital resources; our ability to successfully execute our go-to-market strategy and expand in our existing markets and into new markets; and general market, political, economic and business conditions (including developments and volatility arising from the COVID-19 pandemic).

Any additional or unforeseen effect from the COVID-19 pandemic may exacerbate these risks. Additional risks and uncertainties that could cause actual outcomes and results to differ materially are included in our filings with the Securities and Exchange Commission (the “SEC”), including the quarterly report on Form 10-Q for the quarter ended January 31, 2020 and any subsequent reports filed with the SEC. SEC filings are available on the Investor Relations section of Elastic’s website at ir.elastic.co and the SEC’s website at www.sec.gov. Elastic assumes no obligation to, and does not currently intend to, update any such forward-looking statements, except as required by law.

In addition to GAAP financial information, this presentation and the accompanying oral presentation include certain non-GAAP financial measures. See the Appendix for a reconciliation of all historical non-GAAP financial measures to their nearest GAAP equivalent.

elastic

Agenda

What are we covering in the next two hours:

- Introductions: Elastic, Elasticsearch, and Kibana (for geo)
- Web mapping with Elasticsearch
  - Rendering documents and aggregated data
  - Searching and drilling down with dates and any text
  - Geospatial filtering

What are we not covering today:

- Advanced web processing or UI (TypeScript, React, and such)
- Generative AI

Elastic, a search company

Search. Observe. Protect

elastic

Today we live in an *always on world*

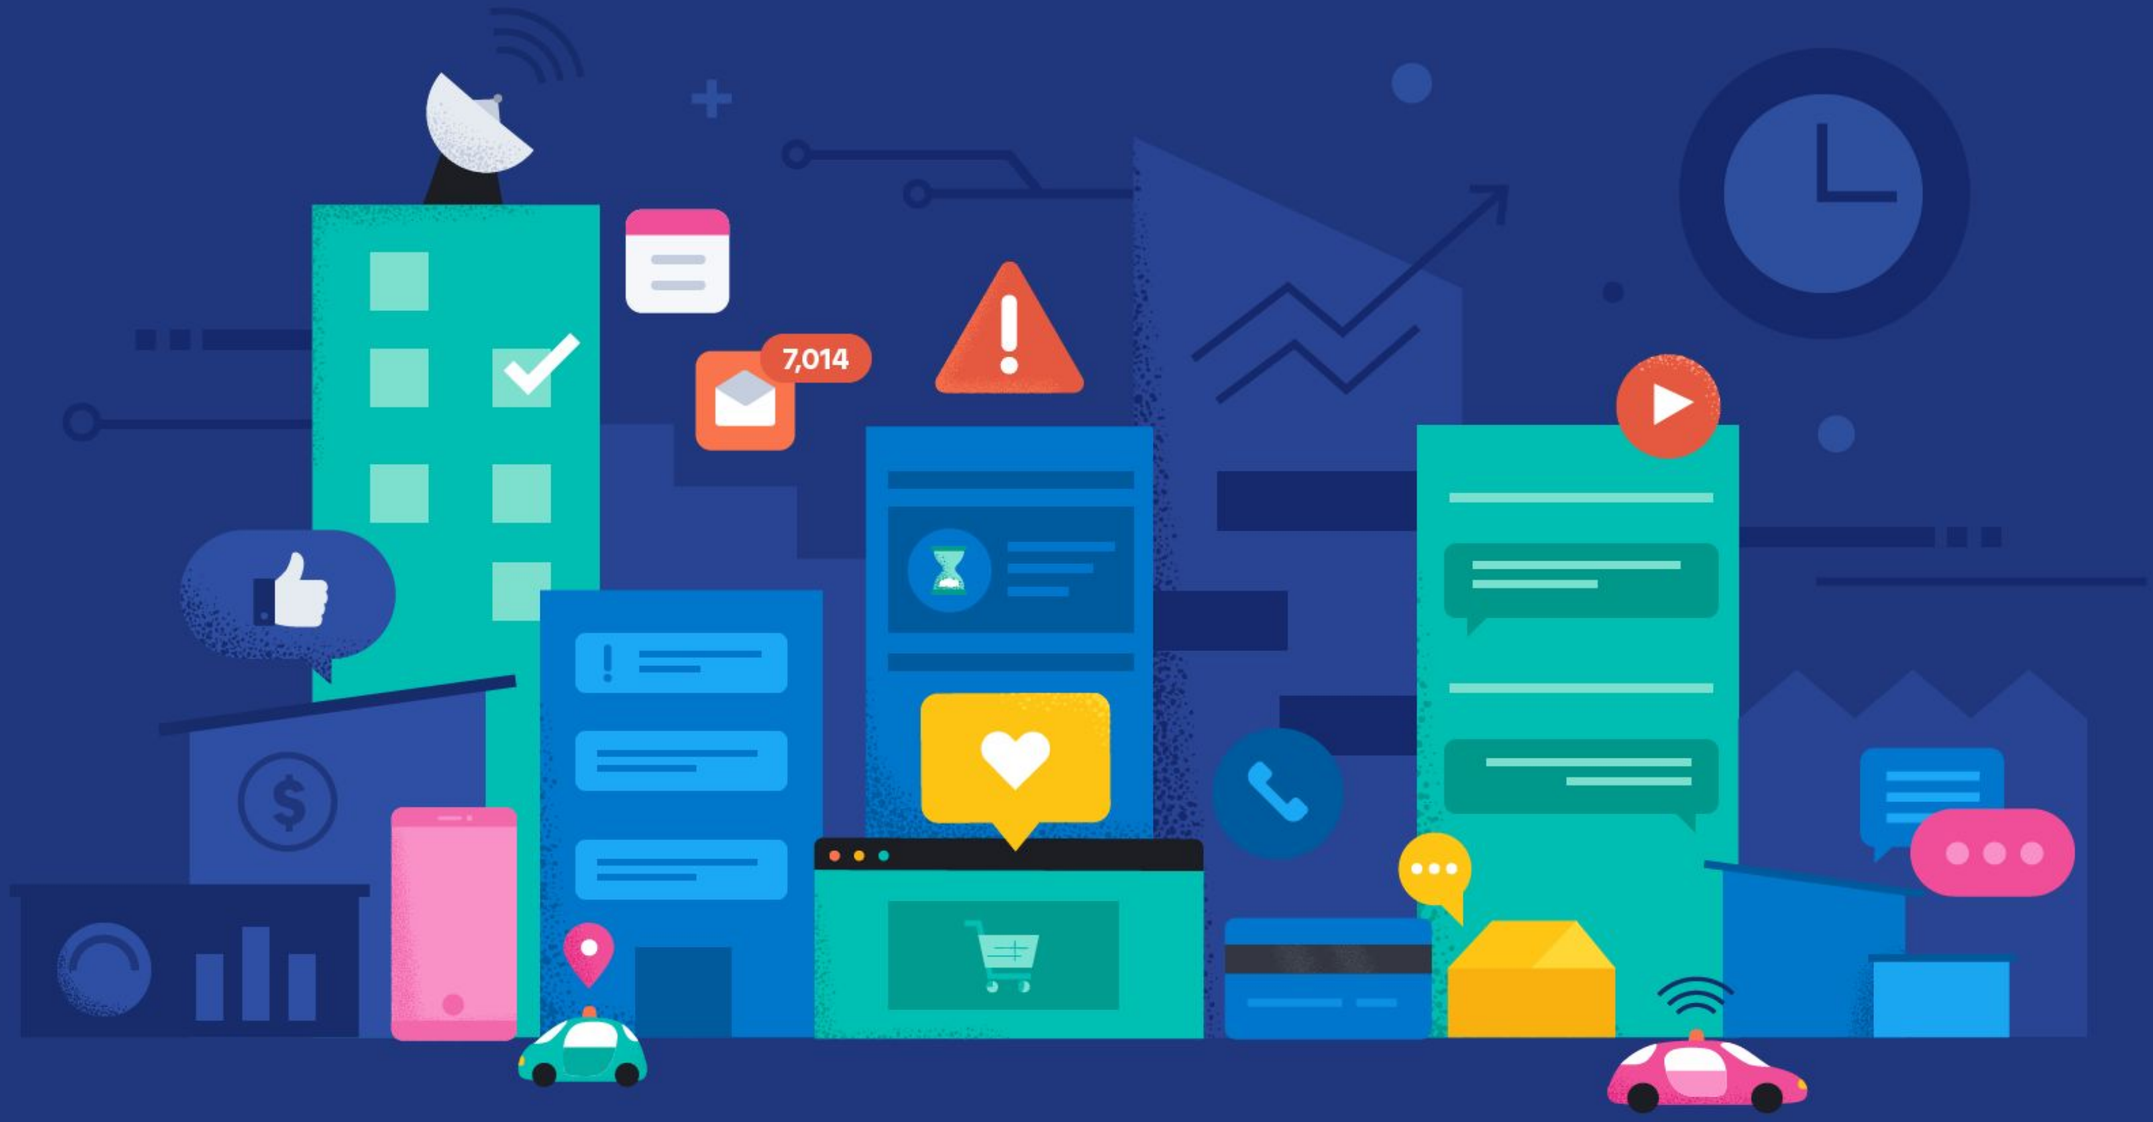

A world characterized by real challenges

Content is becoming **harder** to find

Enterprise IT is becoming more **complex**

Cyber threats are becoming more **sophisticated**

elastic

A world characterized by endless data

480EB

1 EB = 1000 PB = 1,000,000 TB

Data produced daily by 2025

IDC: Global DataSphere and StorageSphere Forecasts

elastic

Meet Elastic

Elastic helps the world's leading organizations **accelerate results that matter** by putting data to work with the **power of search**.

elastic

Elastic at a glance

NYSE: ESTC

Founded in 2012

3000+
employees

50+
countries with employees

17,900+
subscriptions

54%
of Fortune 500 companies trust Elastic

elastic

Elastic at a glance

NYSE: ESTC

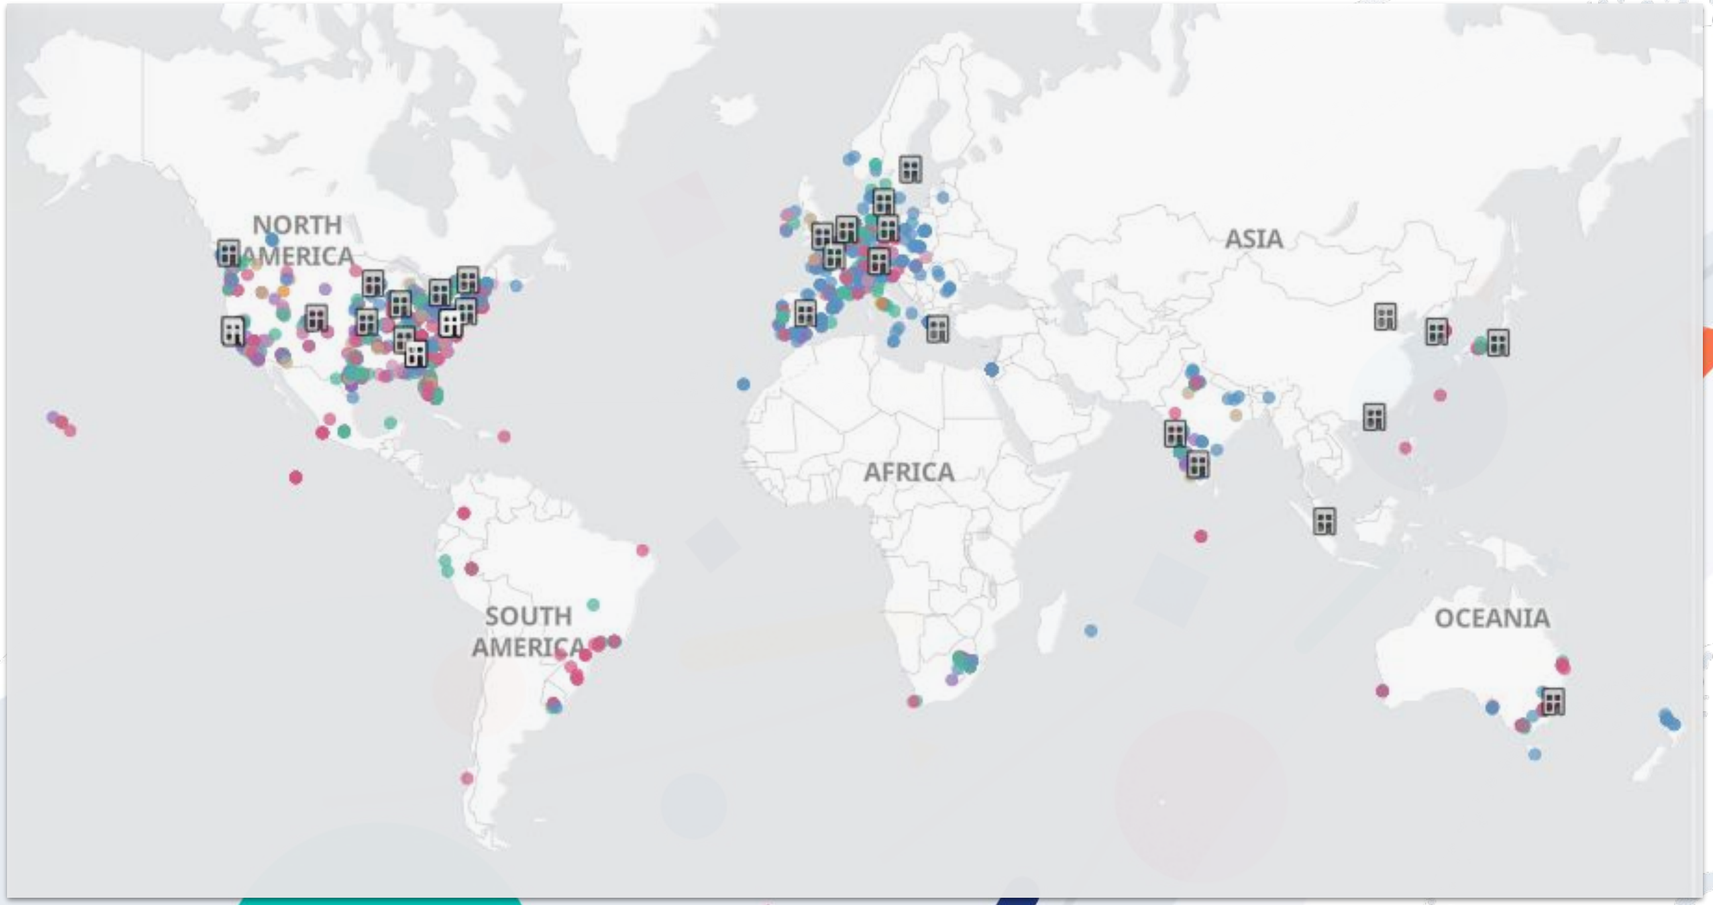

elastic

The Elastic Search Platform *is for everyone*

TECHNOLOGY	FINANCE	TELCO	CONSUMER	HEALTHCARE	PUBLIC SECTOR	AUTOMOTIVE / TRANSPORTATION	RETAIL
Adobe	BARCLAYS	orange	Uber	VITAS Healthcare	IU	VOLVO Volvo Group	H-E-B
CISCO	ZURICH	Bell	Grab	UCLA Health	OAK RIDGE National Laboratory	Airbus	THE HOME DEPOT
workday	USAA	SoftBank	tinder	Yale NewHaven Health	United States Census Bureau	Travelport	ebay
Microsoft	collector bank	verizon	ACTIVISION BLIZZARD	MAYO CLINIC	JPL Jet Propulsion Laboratory	CSX TRANSPORTATION	Kroger
INGRAM MICRO	Postbank	T-Mobile	lyft	Pfizer	WILSON NORTH CAROLINA		Walgreens

elastic

Communit
y

https://github.com/elastic

https://ela.st/slack

https://discuss.elastic.co

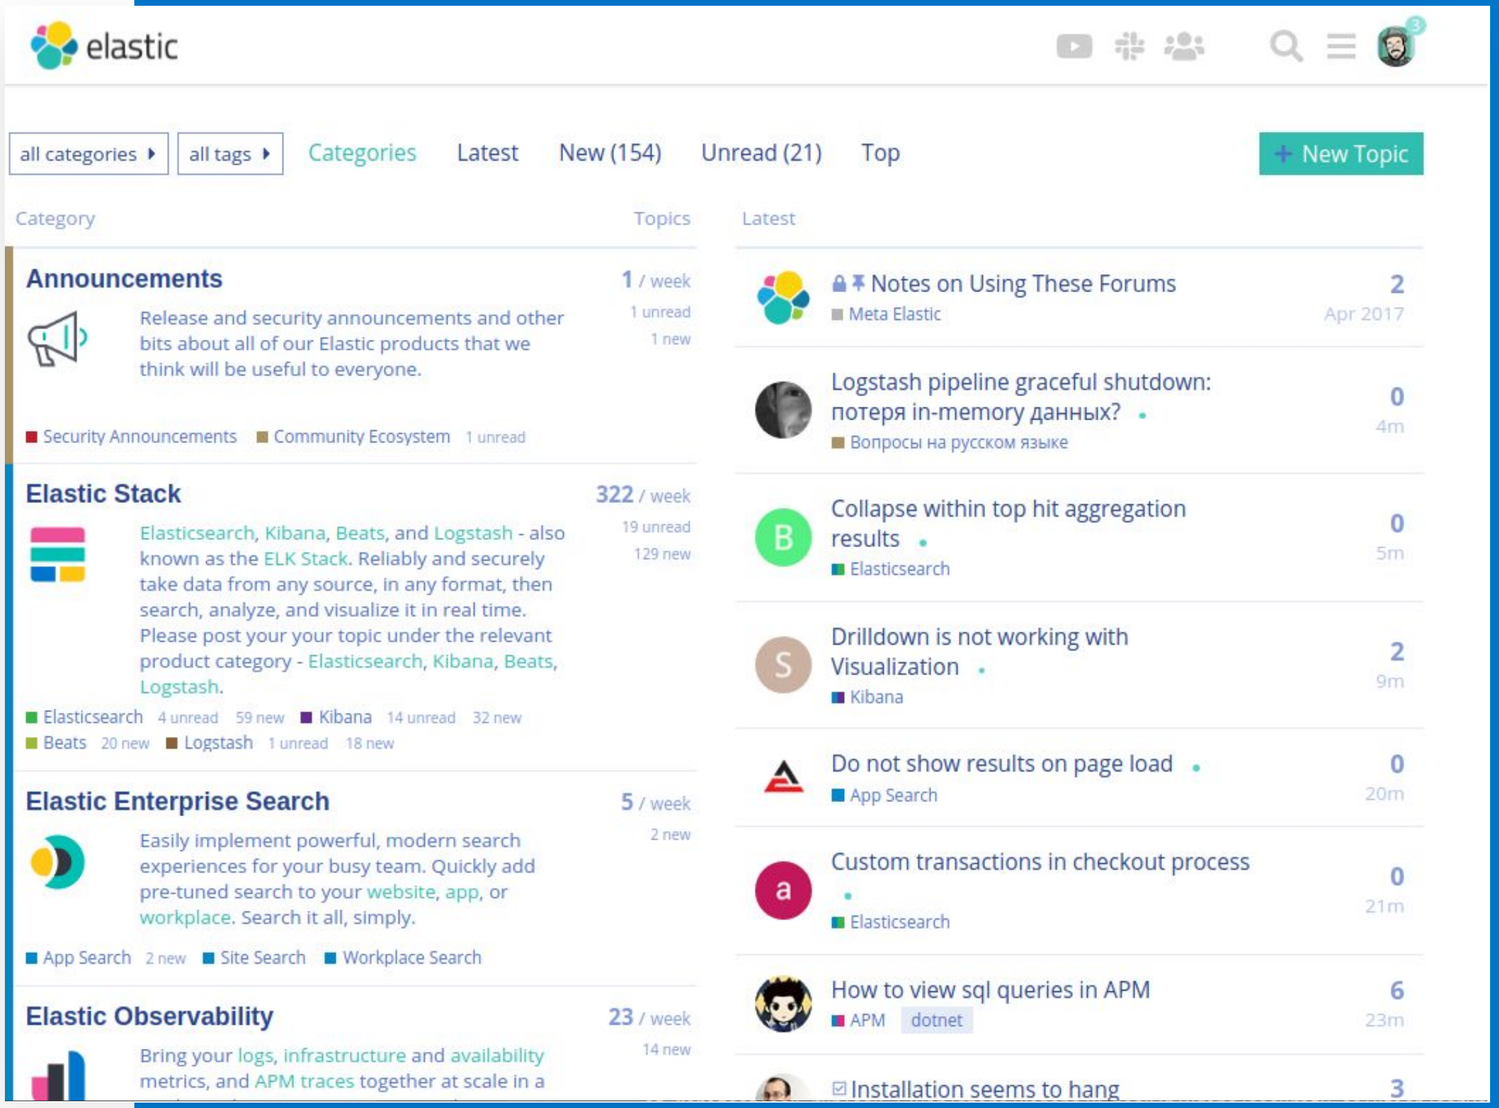

The Elastic Search Platform

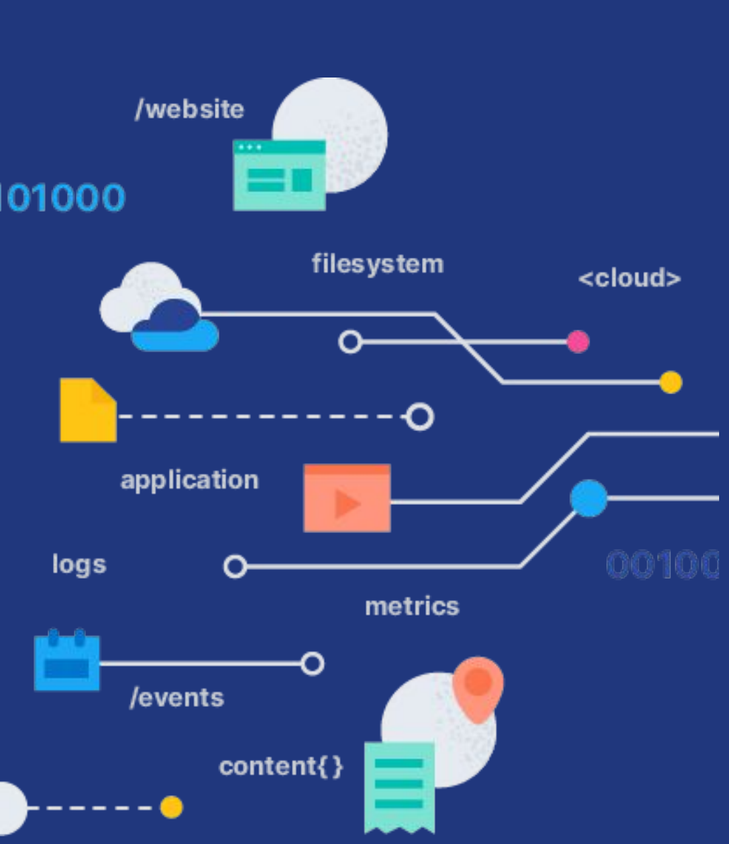

Enterprise Search

Observability

Security

Kibana
Explore, Visualize, Engage

Elasticsearch
Store, Search, Analyze

Integrations
Connect, Collect, Alert

Public cloud

Hybrid

On-premises

elastic

ELASTIC ENTERPRISE SEARCH

Search everything, anywhere

Easily implement powerful, modern search experiences across your website, app, or digital workplace. Search it all, simply.

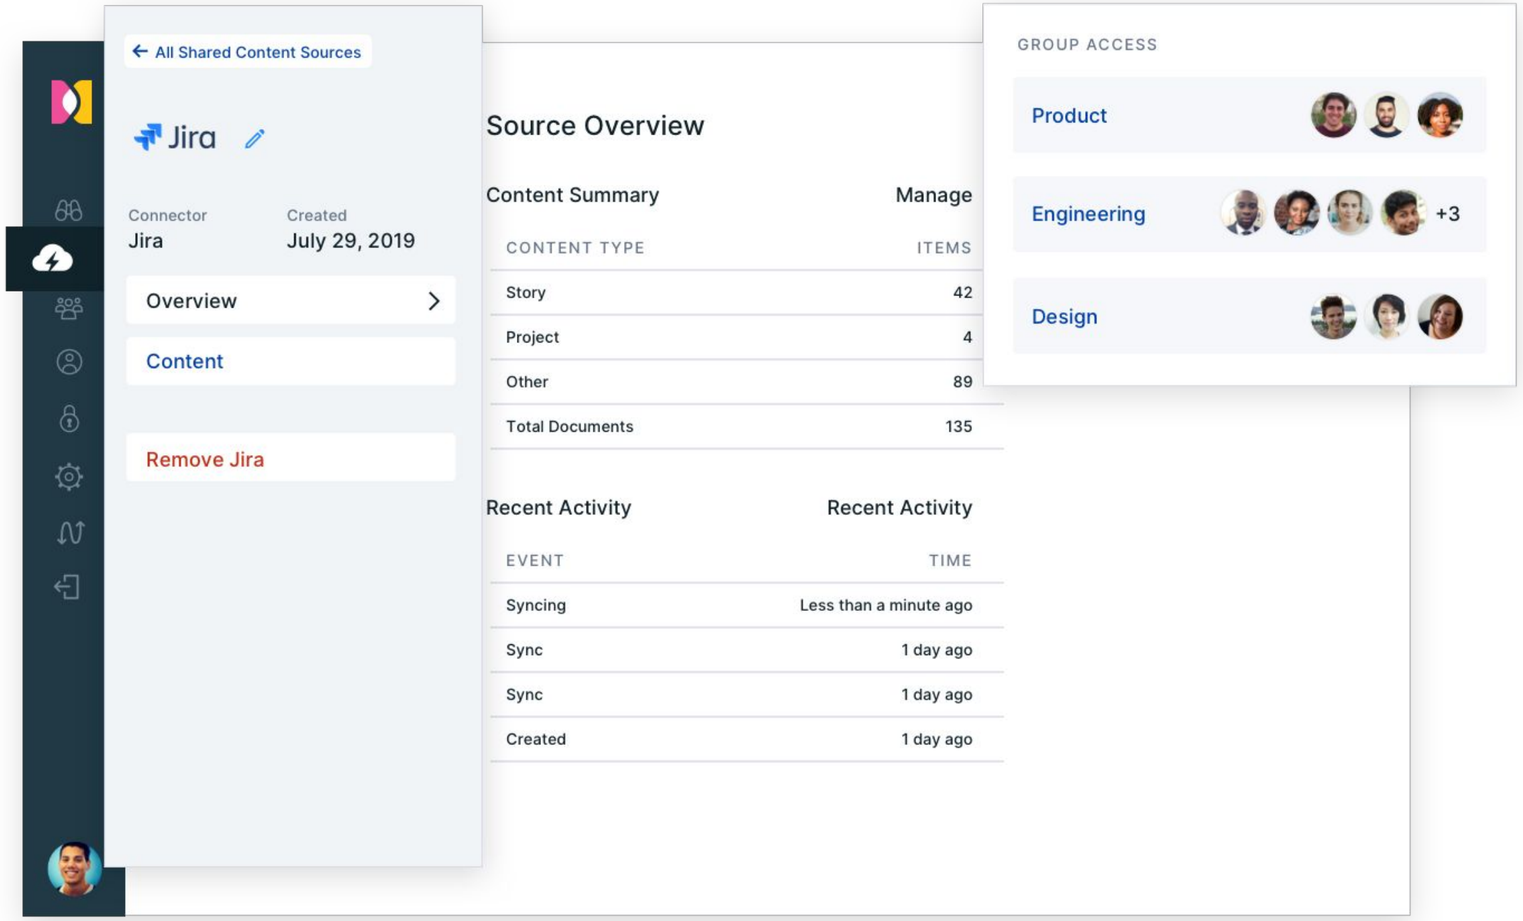

ELASTIC OBSERVABILITY

Unified visibility across your entire ecosystem

Bring your logs, metrics, and traces together into a single stack so you can monitor, detect, and react to events with speed.

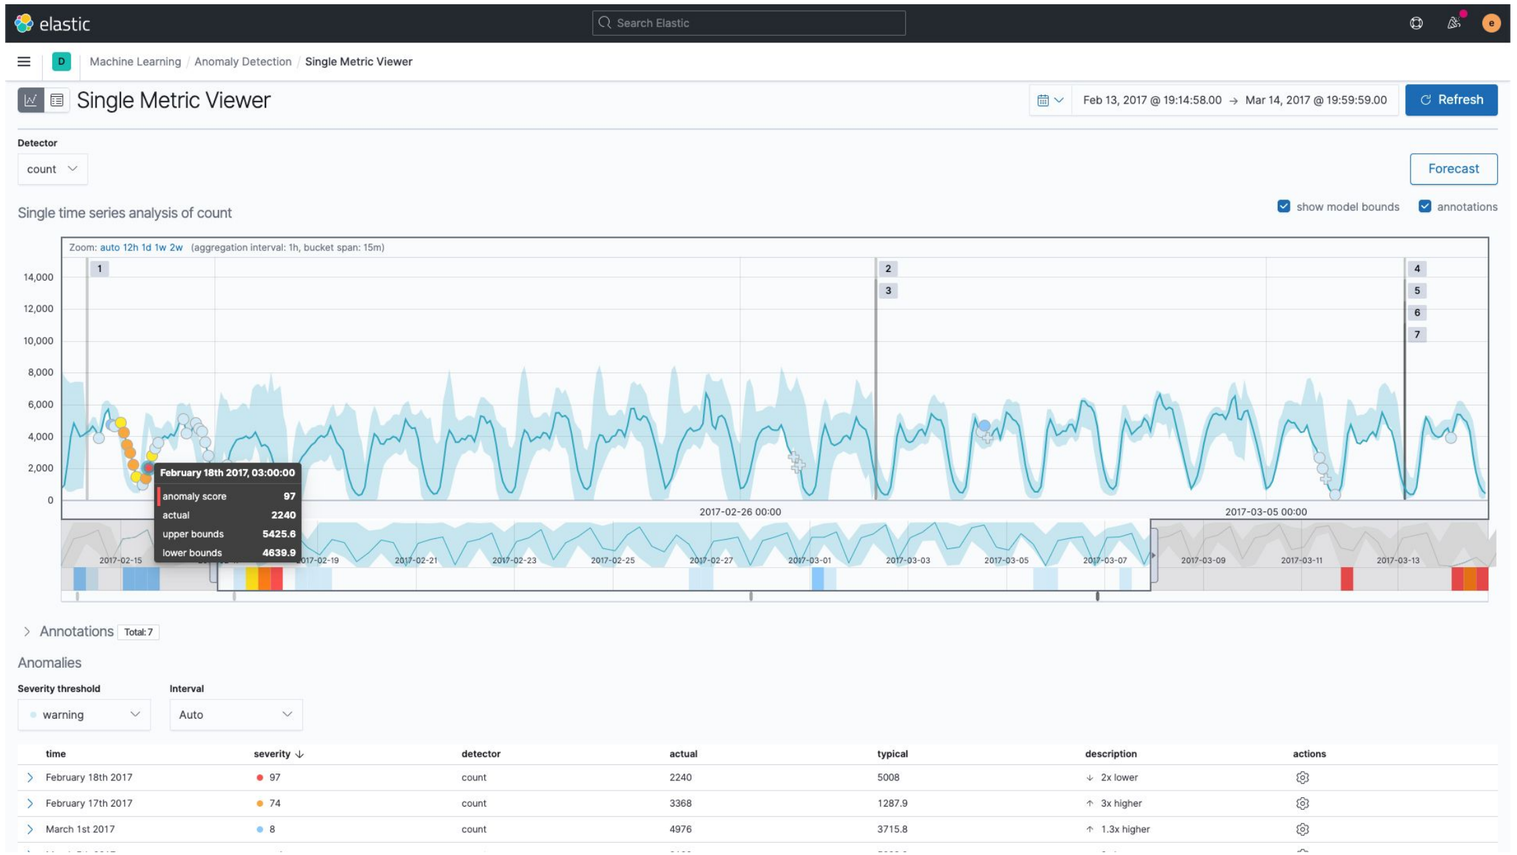

ELASTIC SECURITY

Security how it should be: open

Elastic Security integrates endpoint security, SIEM, and Cloud Security to give you prevention, collection, detection, and response capabilities for unified protection across your infrastructure.

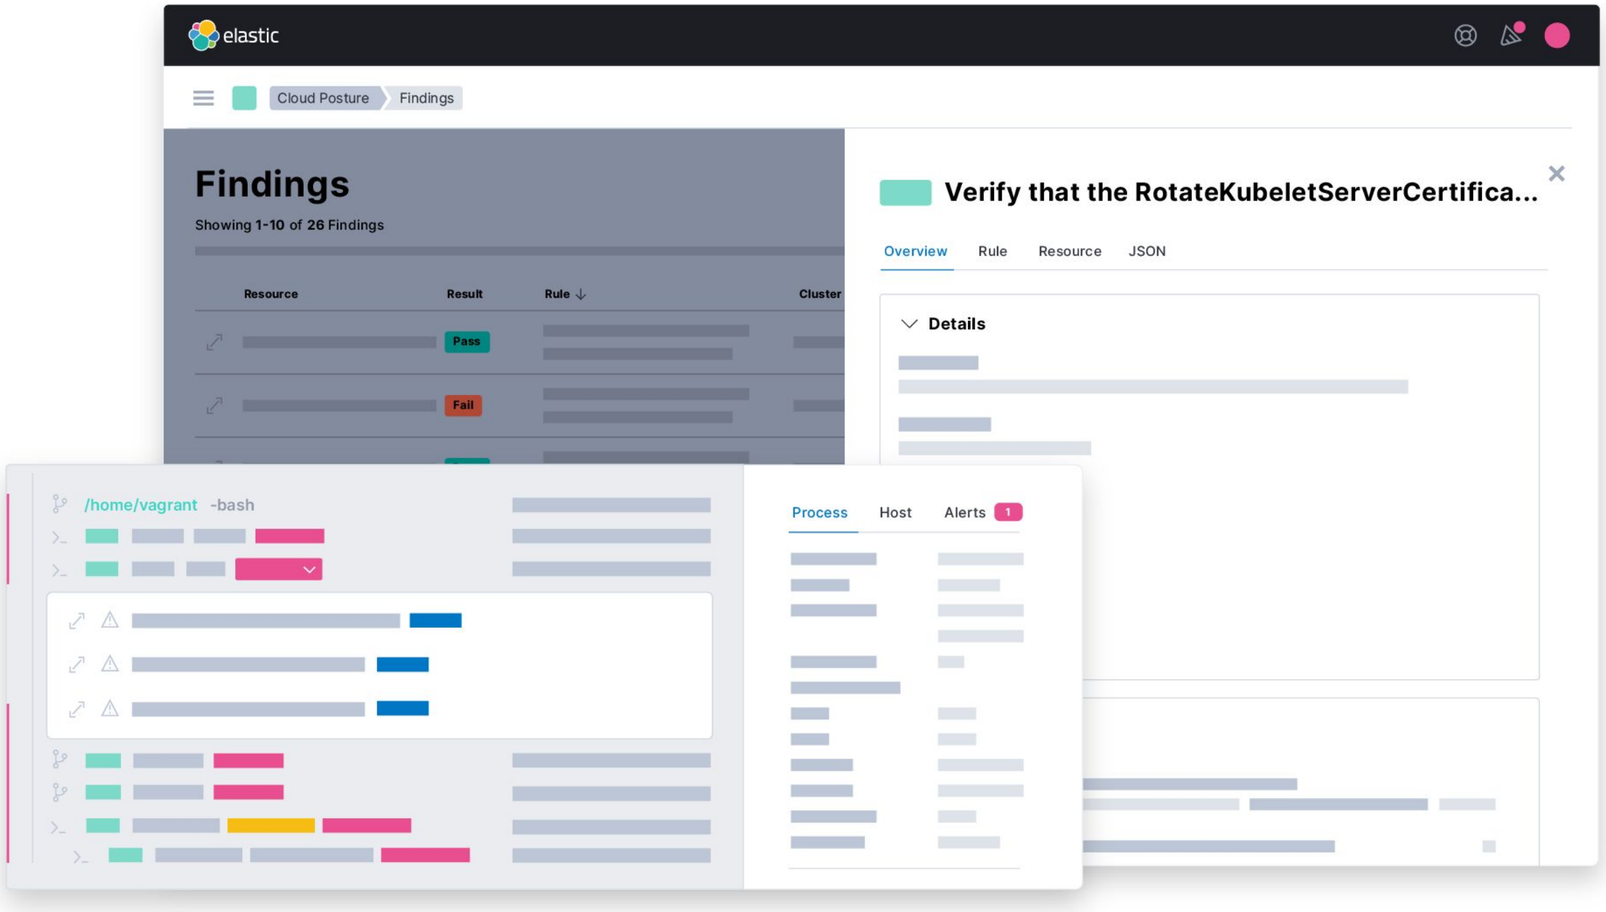

Elasticsearch and Kibana intro

elastic

Elastic **Stack**

Ingest, Store, Search, Visualise

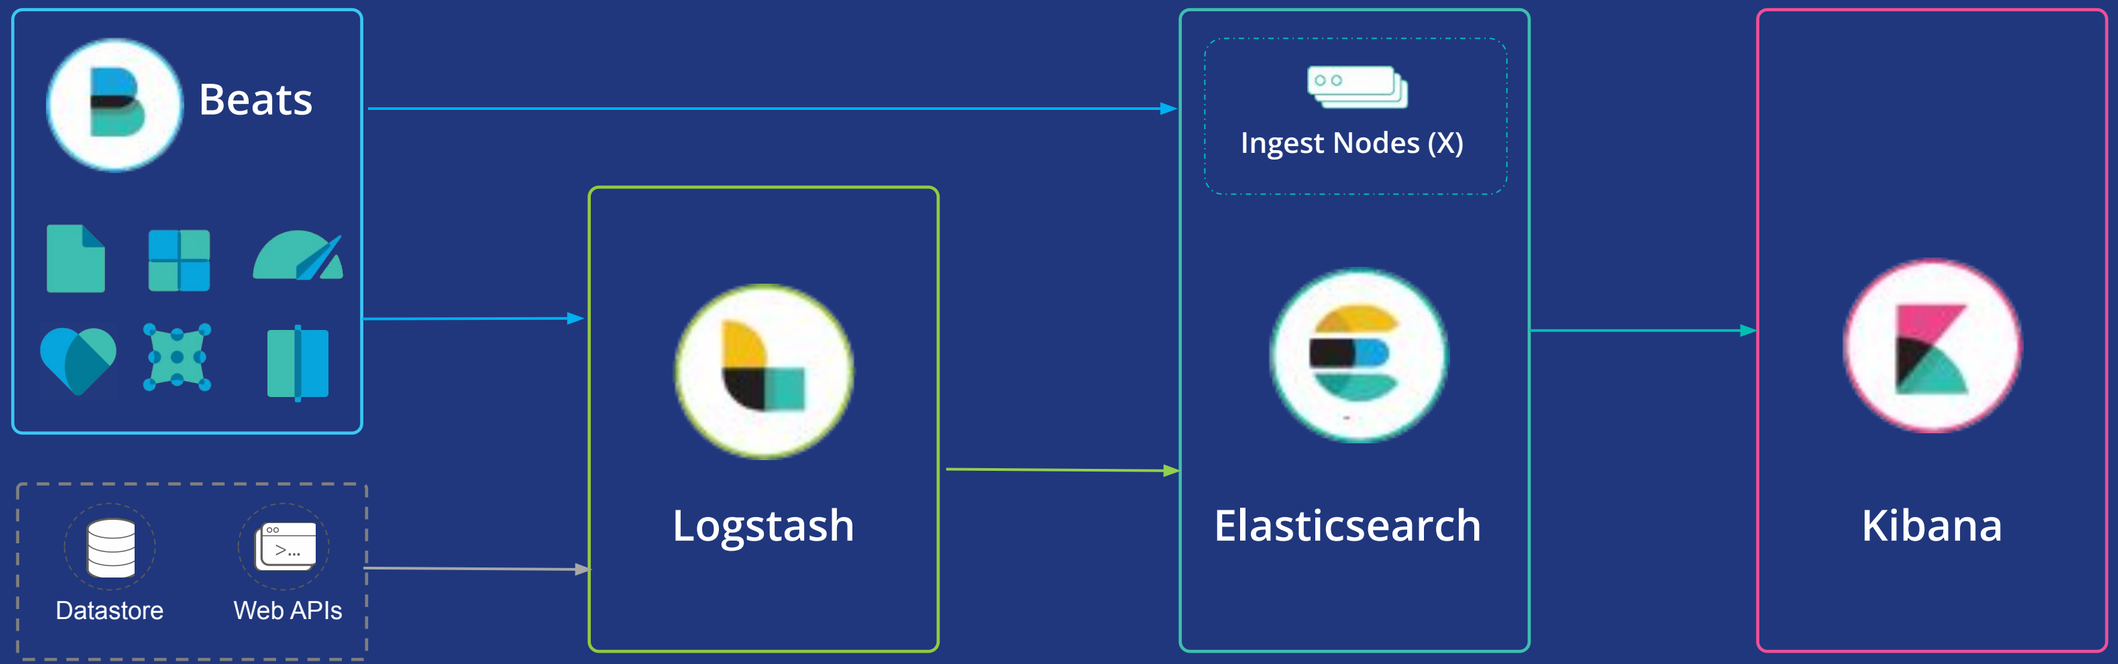

elastic

Elasticsearch

All data is welcome

elastic

Elasticsearch components

Cluster

Node

Shard

Index

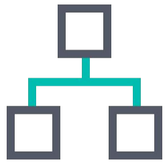

Mapping

Document

Field

elastic

Communicating with Elasticsearch

- All communication thr
  HTTP endpoints 📖
- JSON
- REST methods: GET, I
  POST, DELETE
- _cat API for human re
  display 📖

```
~
→ curl -s --user "${ELASTIC_USER}:${ELASTIC_PASSWORD}" "${ELASTIC_HOST}/" | jq
{
```

```
~
→ time curl -s --user "${ELASTIC_USER}:${ELASTIC_PASSWORD}" \
  "${ELASTIC_HOST}/geonames/_search?q=valencia&size=20" |\
  jq -c ".hits.hits[]._source | { g: .location, c: .UNLOCODEName, n: .Name}"
{"g":{"lon":"122.71","lat":"10.668"},"c":"Philippines","n":"East Valencia"}
{"g":{"lon":"30.27718","lat":"-23.87011"},"c":"South Africa","n":"Valencia Estate"}
{"g":{"lon":"-87.45963","lat":"19.69255"},"c":"Mexico","n":"Valencia"}
{"g":{"lon":"-81.41667","lat":"8.06667"},"c":"Panama","n":"La Valencia"}
{"g":{"lon":"-61.19993","lat":"10.64988"},"c":"Trinidad And Tobago","n":"Valencia"}
{"g":{"lon":"-103.4128","lat":"26.29734"},"c":"Mexico","n":"Valencia"}
{"g":{"lon":"-97.90795","lat":"21.584"},"c":"Mexico","n":"La Valencia"}
{"g":{"lon":"-97.55388","lat":"18.65448"},"c":"Mexico","n":"Valencia"}
{"g":{"lon":"-75.11332","lat":"9.13451"},"c":"Colombia","n":"Valencia"}
{"g":{"lon":"-76.6136","lat":"2.44189"},"c":"Colombia","n":"Valencia"}
{"g":{"lon":"-78.4","lat":"-0.36667"},"c":"Ecuador","n":"Hacienda Valencia"}
{"g":{"lon":"-61.1668","lat":"10.68233"},"c":"Trinidad And Tobago","n":"Ward of Valencia"}
{"g":{"lon":"-102.35591","lat":"29.33355"},"c":"Mexico","n":"Valencia"}
{"g":{"lon":"125.0","lat":"7.95"},"c":"Philippines","n":"City of Valencia"}
{"g":{"lon":"-109.80707","lat":"29.09612"},"c":"Mexico","n":"Valencia"}
{"g":{"lon":"124.19428","lat":"13.58267"},"c":"Philippines","n":"Valencia"}
{"g":{"lon":"123.62489","lat":"10.14994"},"c":"Philippines","n":"Valencia"}
{"g":{"lon":"123.39093","lat":"9.7588"},"c":"Philippines","n":"Valencia"}
{"g":{"lon":"121.0378","lat":"14.6104"},"c":"Philippines","n":"Valencia"}
{"g":{"lon":"121.6537","lat":"14.065"},"c":"Philippines","n":"Valencia"}
curl -s --user "${ELASTIC_USER}:${ELASTIC_PASSWORD}"  0,07s user 0,00s system 11% cpu 0,635 to
jq -c ".hits.hits[]._source | { g: .location, c: .UNLOCODEName, n: .Name}"  0,02s user 0,00s sy
```

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

API for 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

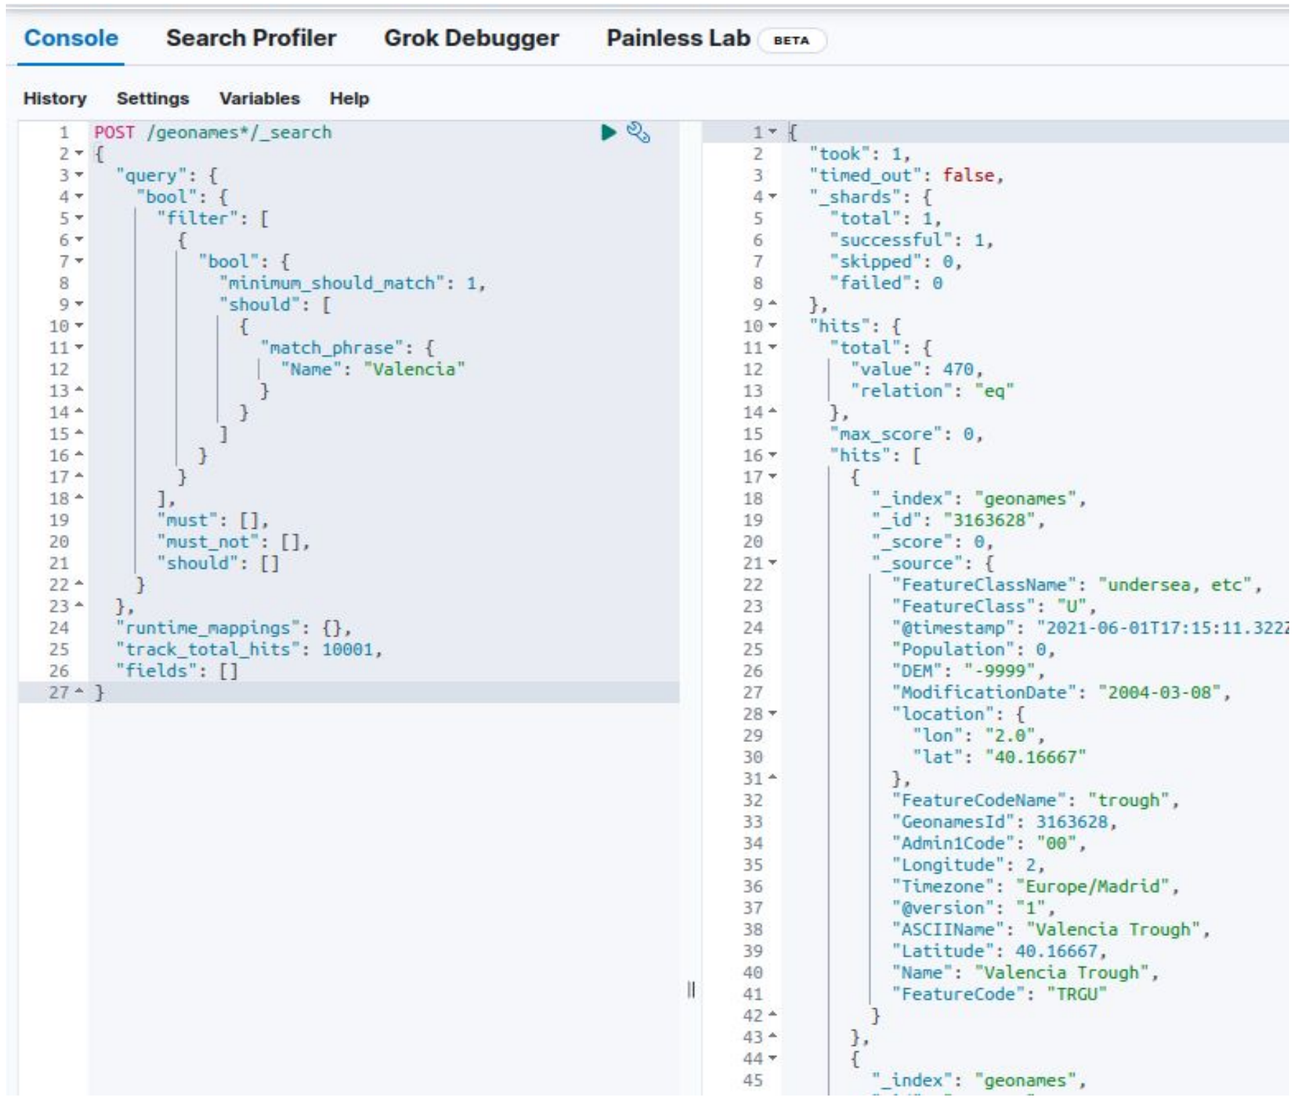

Search

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 ...)

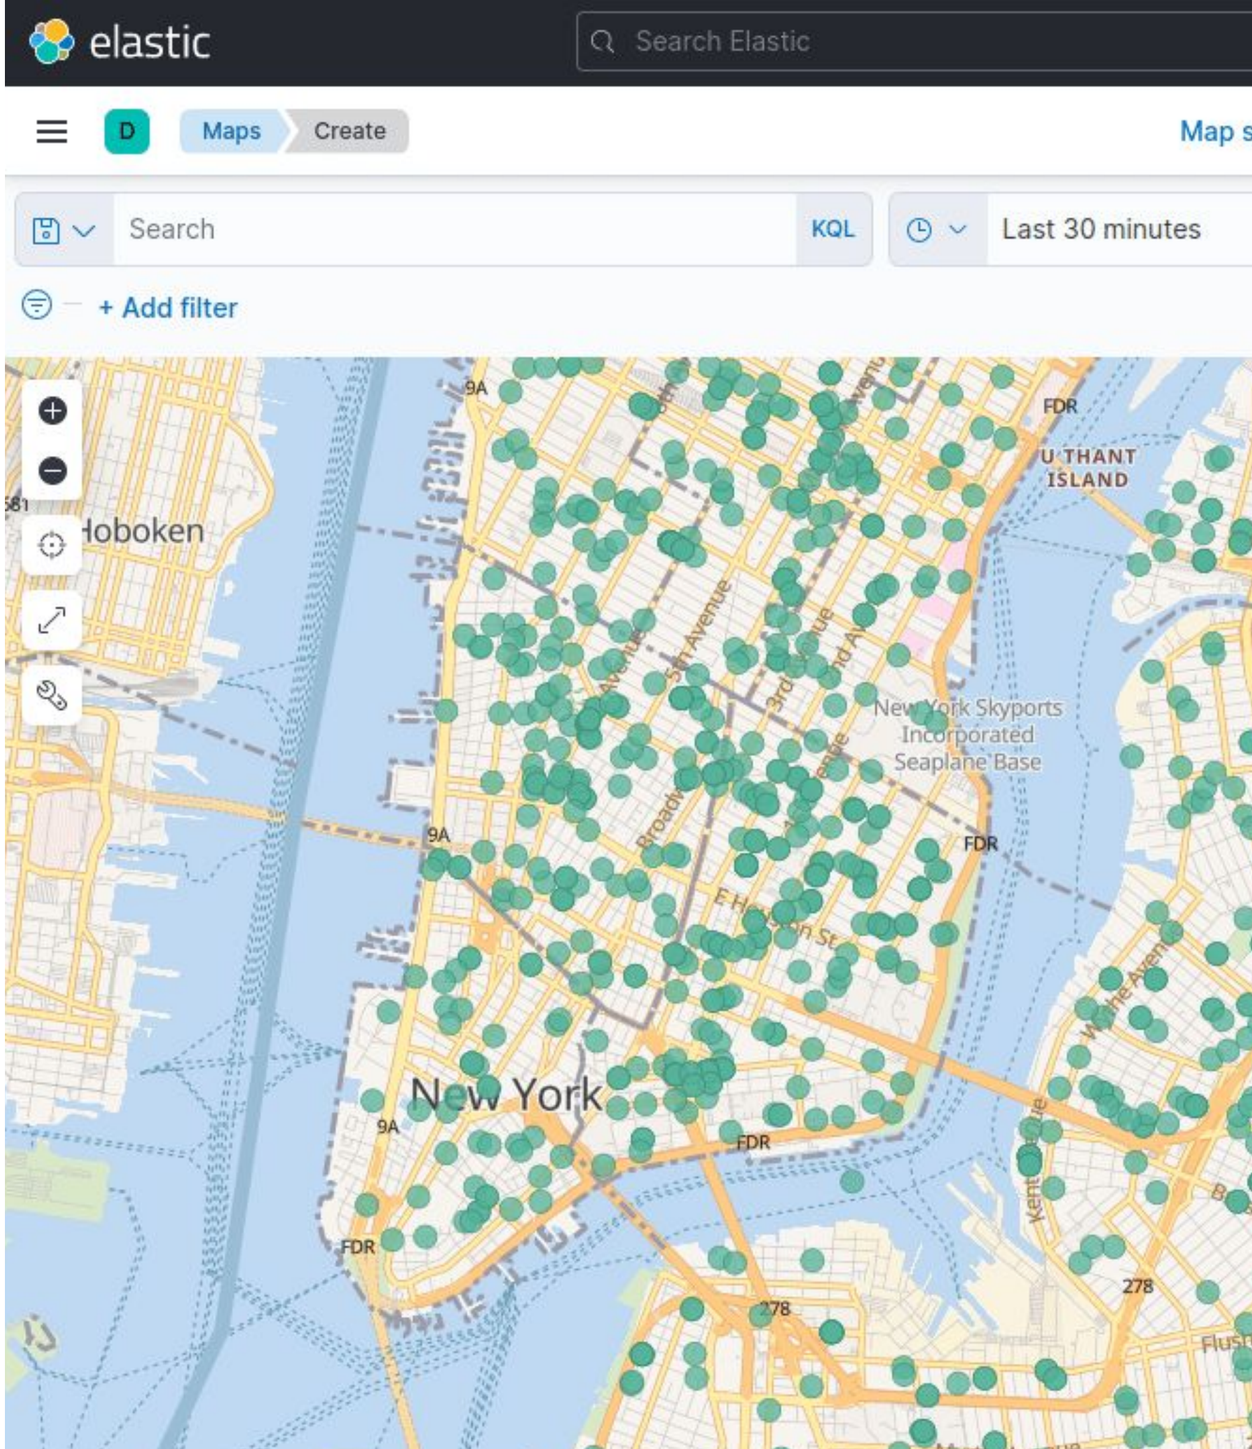

Aggregate

Geo Bucket

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

Geo Metric

- Centroid 📖
- Bounds 📖
- Geoline 📖

Aggregate non-geo using geo filters

- Huge range of aggregations 📖

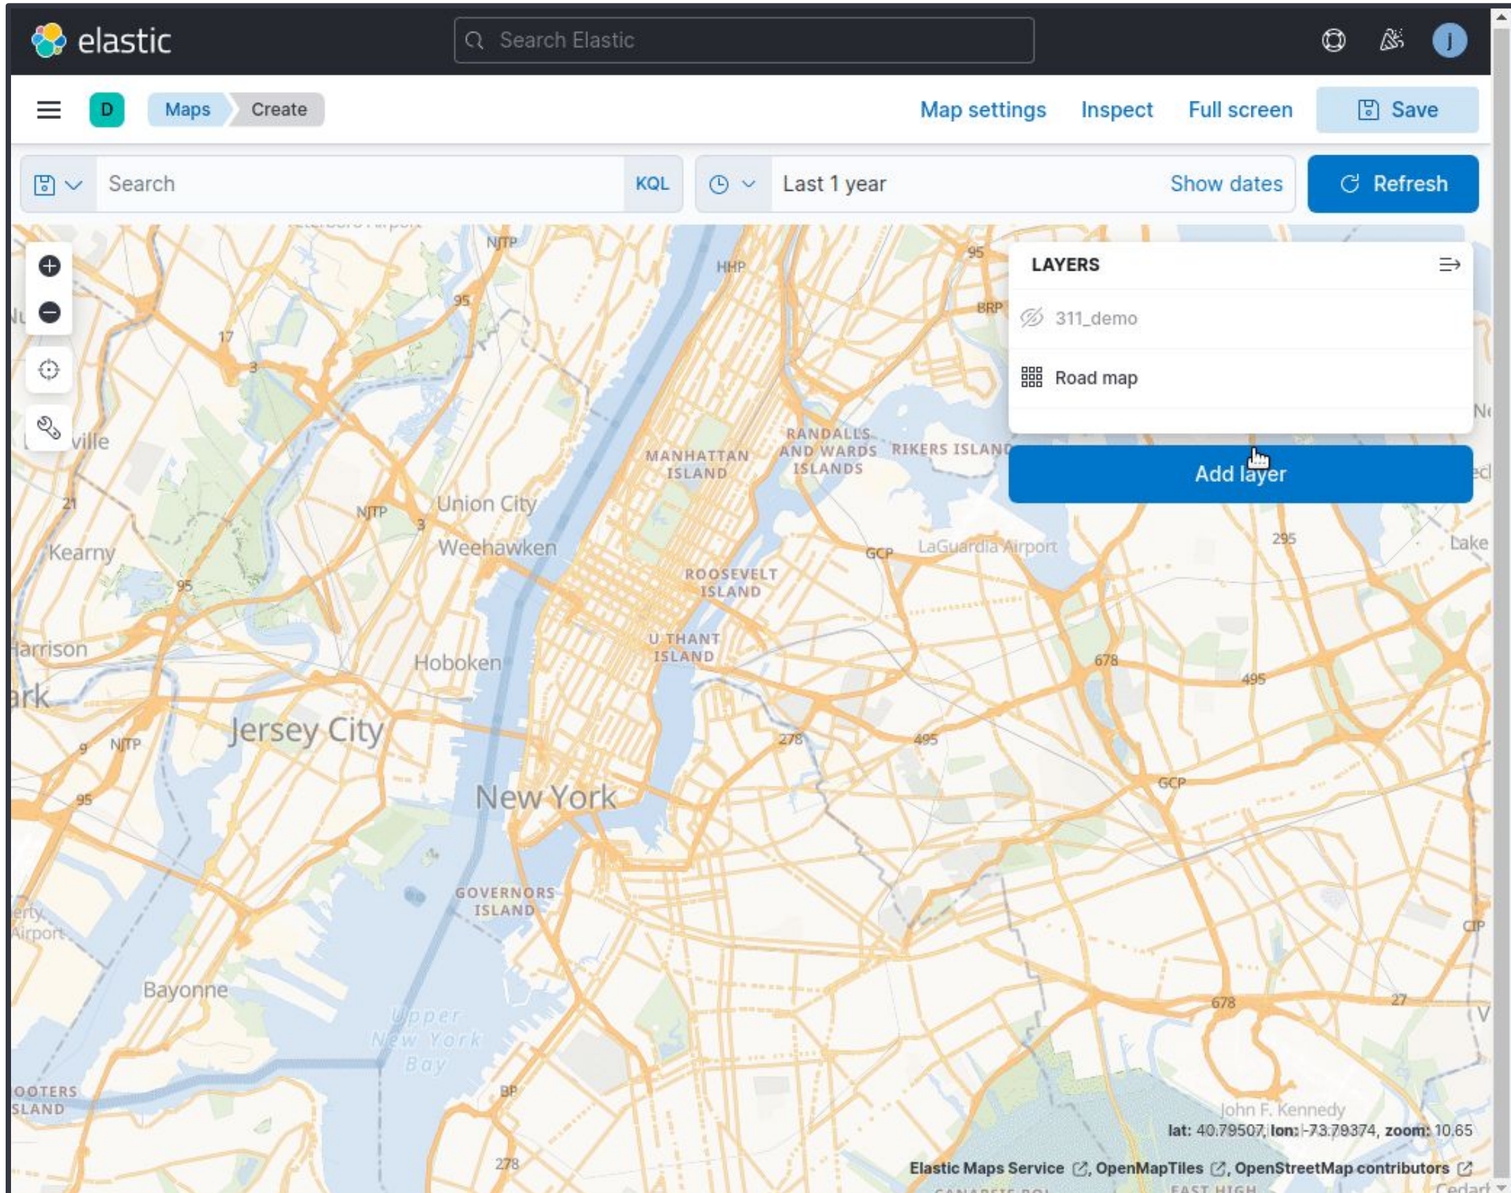

Kibana

Some basic concepts about Kibana

elastic

Data Analysis with Kibana

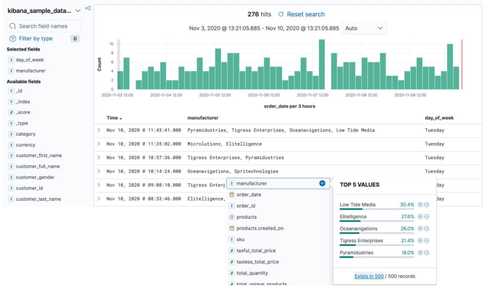

Discover

Lens

Dashboards & DrillDown

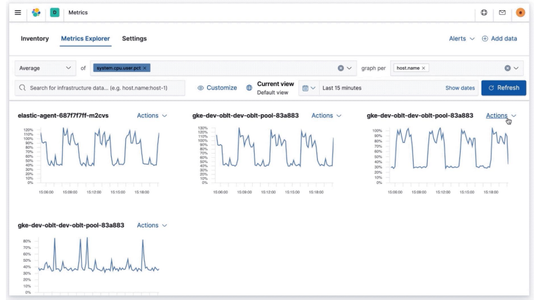

Alerting & Actions

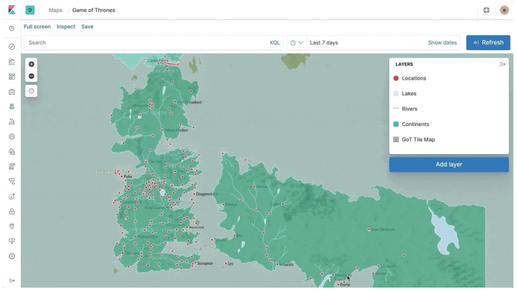

Maps

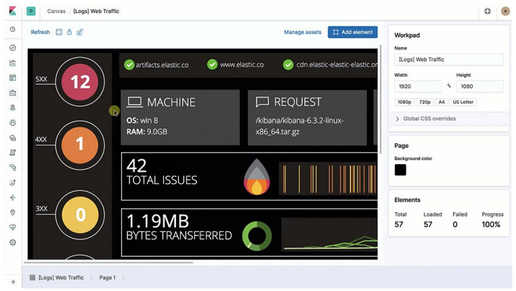

Canvas

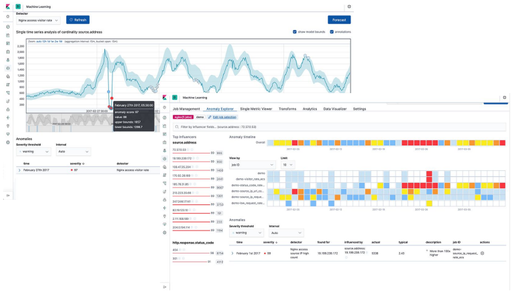

Machine Learning

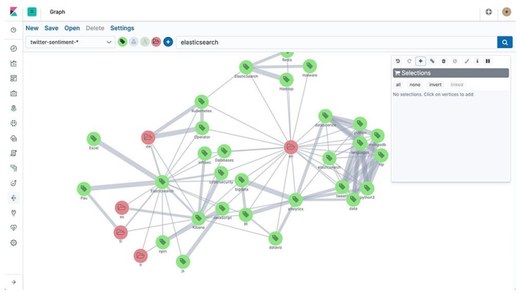

Graph

and much more ...

elastic

Who uses Kibana?

- **Anyone** trying to make sense of data
- **Business** analysts
- **Data** scientists
- Log/metrics **analysts**
- **Security** analysts
- Data service **providers**

elastic

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.

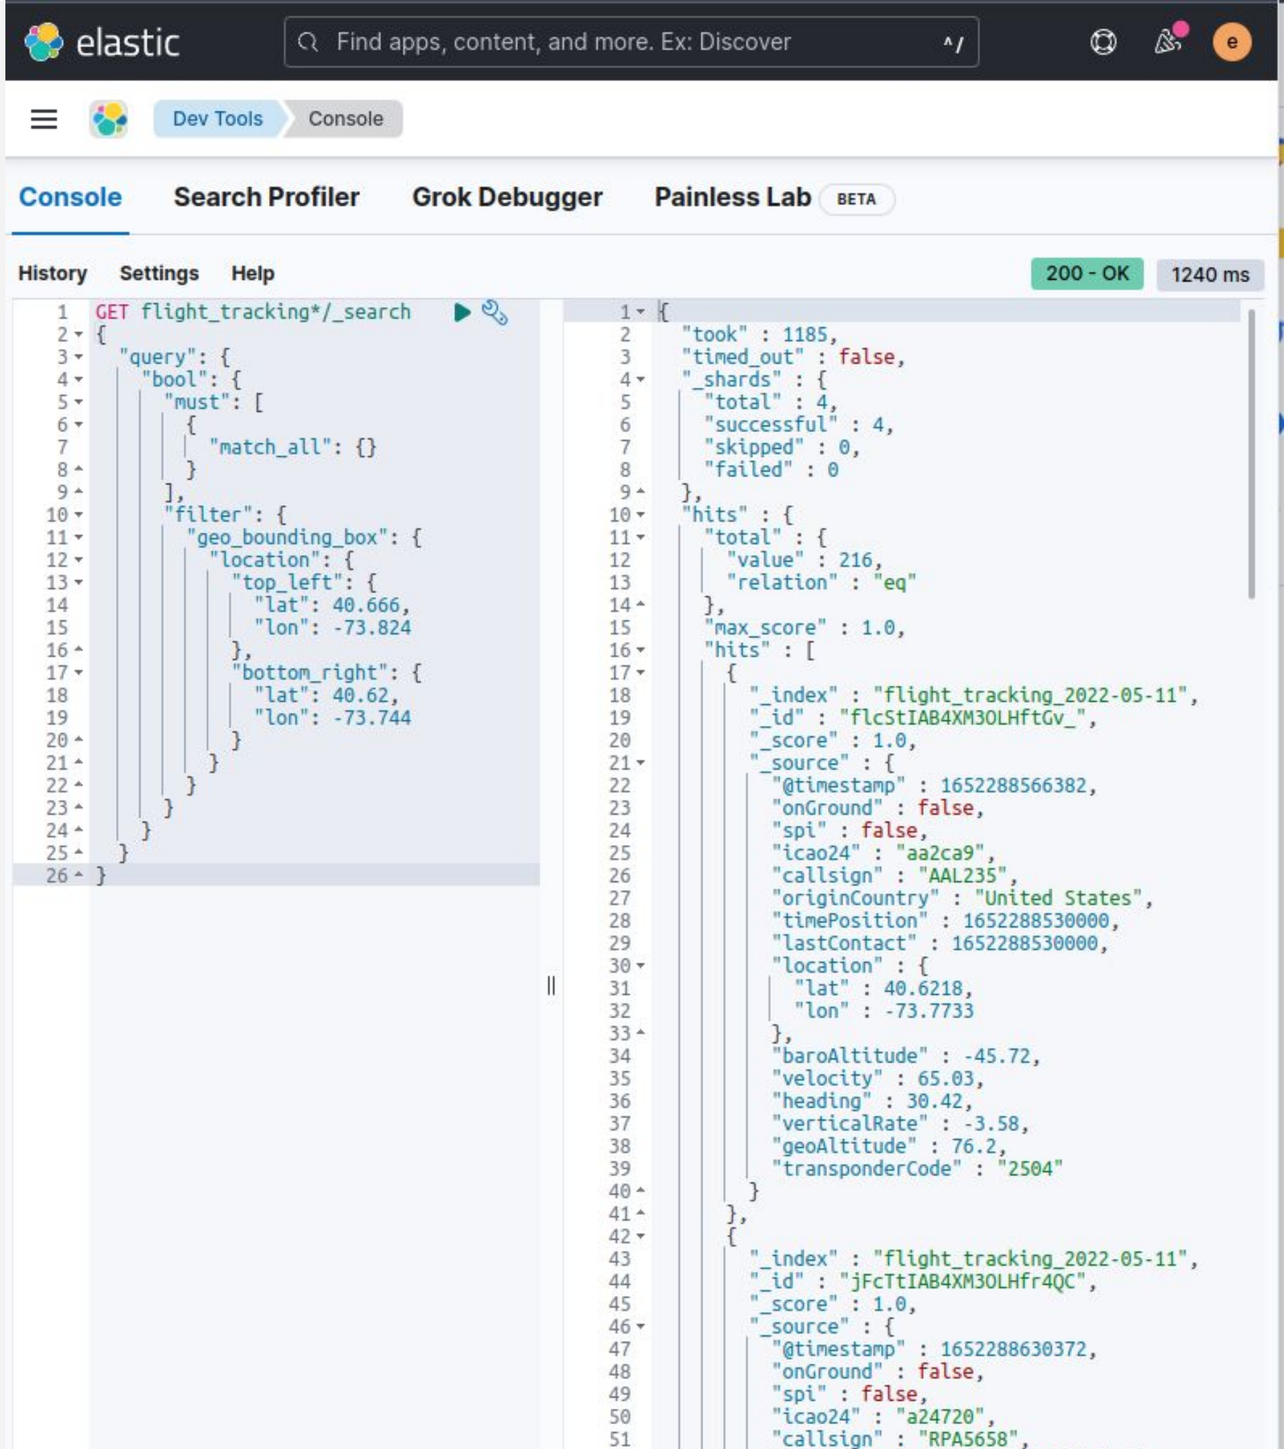

Data Views

- 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

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

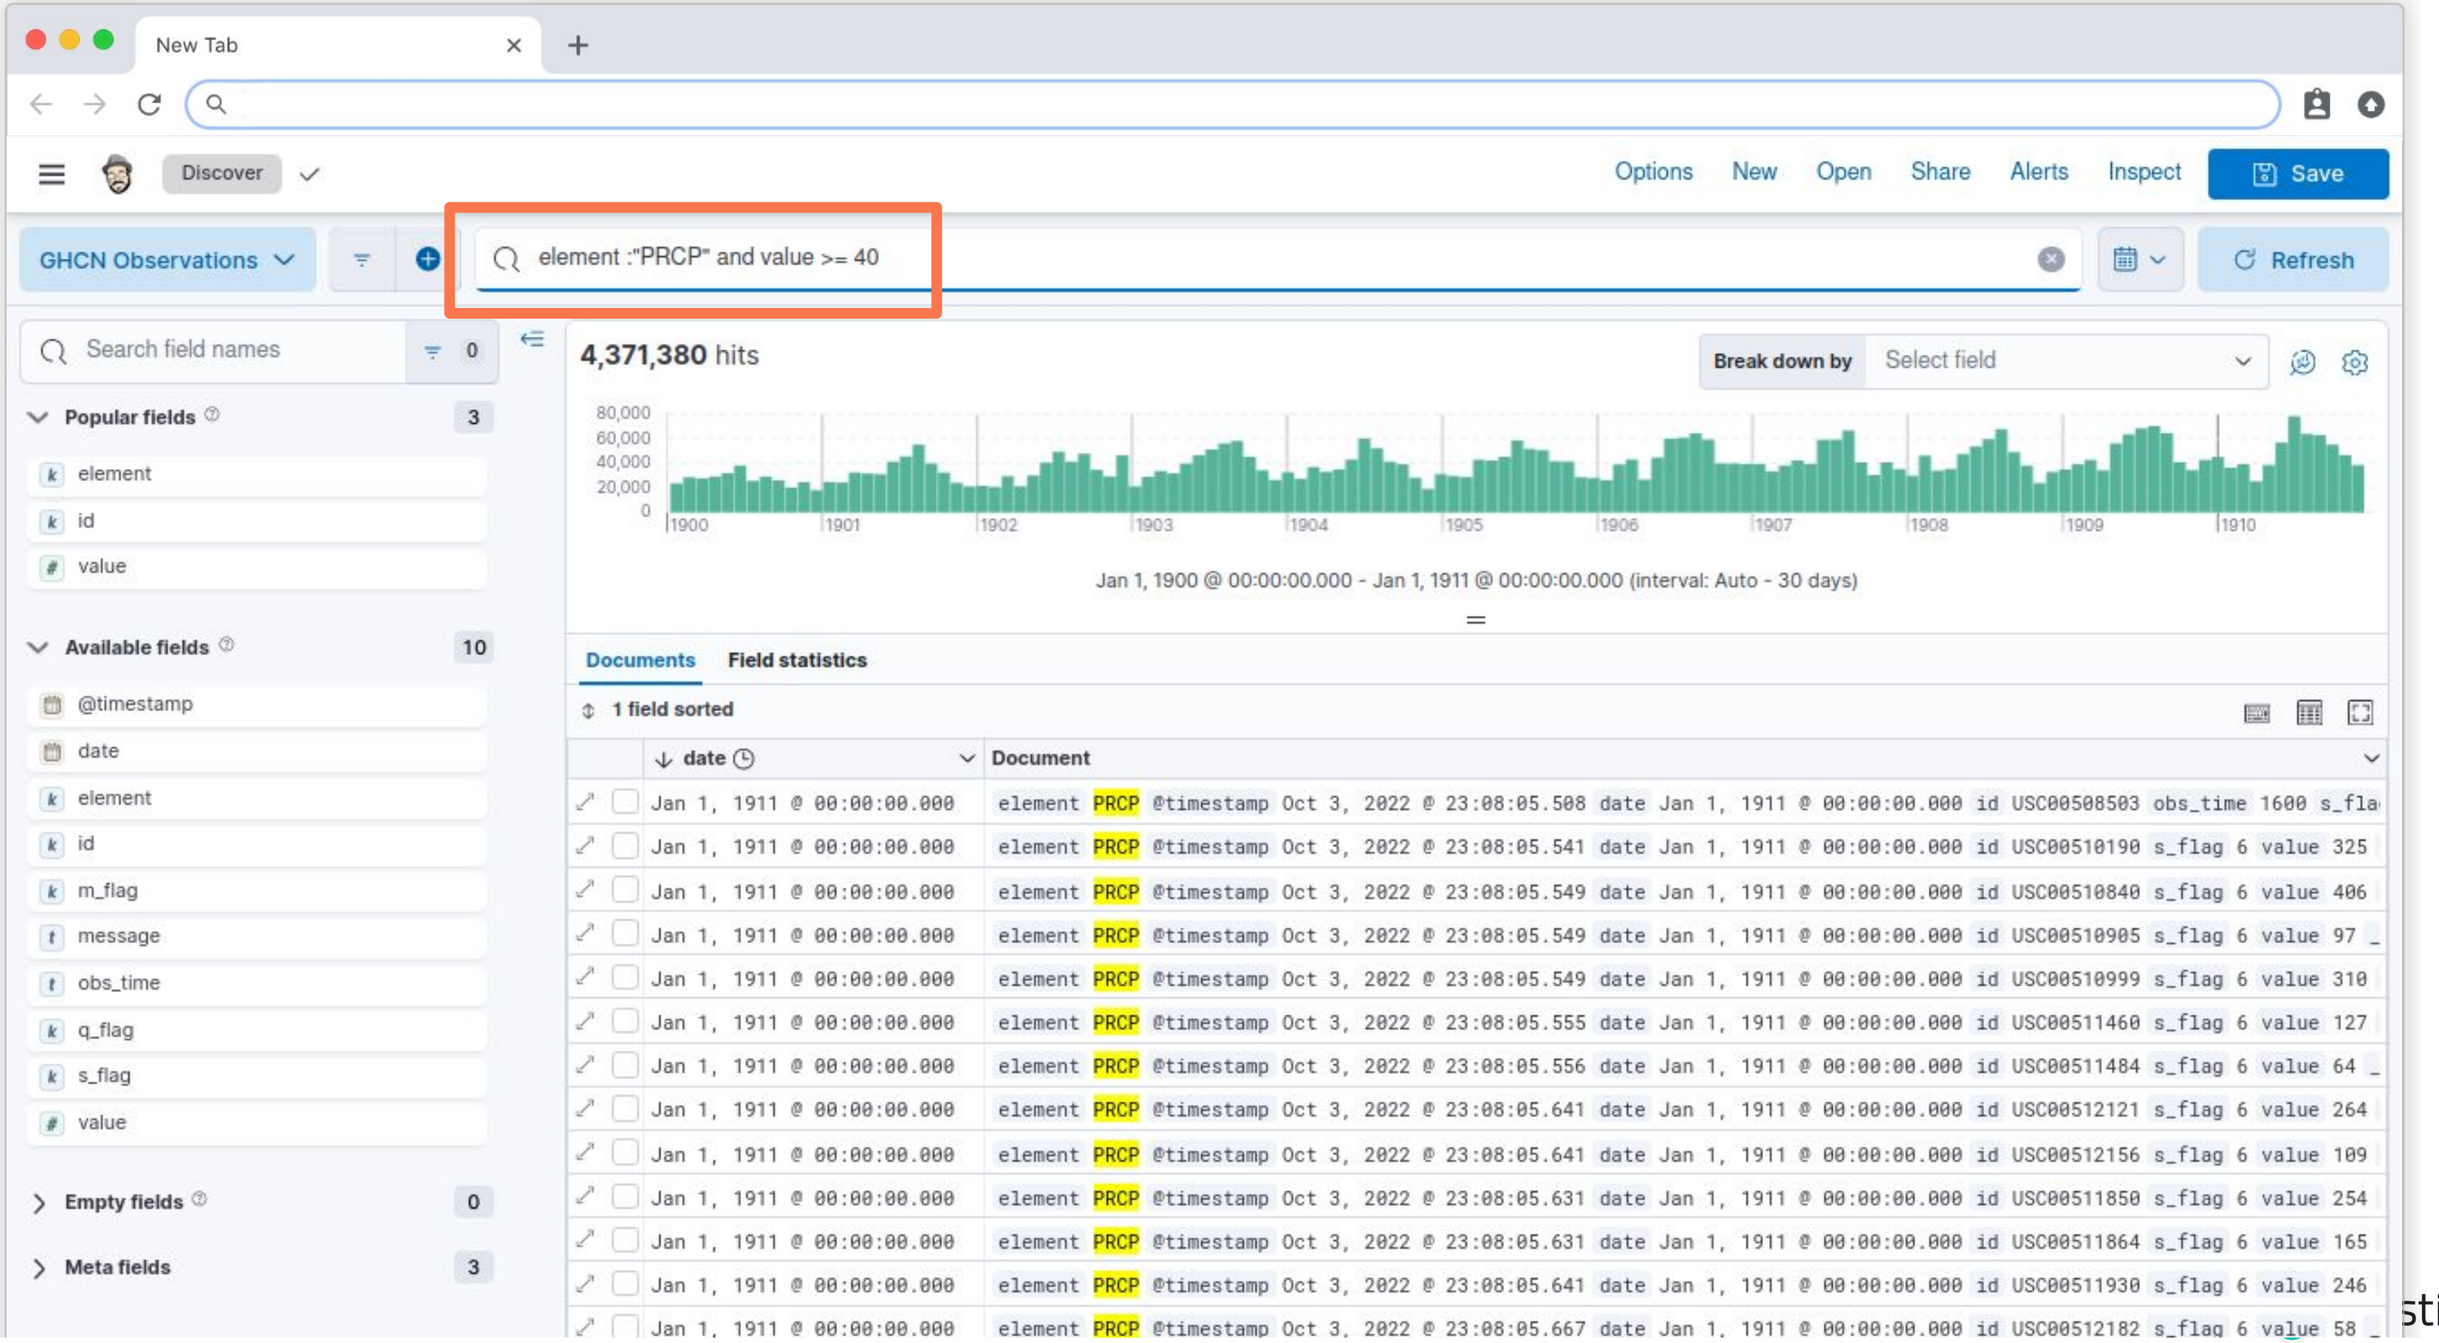
New Tab
Discover
Options
New
Open
Share
Alerts
Inspect
Save
GHCN Observations
element :"PRCP" and value >= 40
Refresh
Search field names
Popular fields
element
id
value
Available fields
@timestamp
date
element
id
m_flag
message
obs_time
q_flag
s_flag
value
Empty fields
Meta fields
4,371,380 hits
Break down by
Select field
80,000
60,000
40,000
20,000
0
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
Jan 1, 1900 @ 00:00:00.000 - Jan 1, 1911 @ 00:00:00.000 (interval: Auto - 30 days)
Documents
Field statistics
1 field sorted
date
Document
Jan 1, 1911 @ 00:00:00.000 element PRCP @timestamp Oct 3, 2022 @ 23:08:05.508 date Jan 1, 1911 @ 00:00:00.000 id USC00508503 obs_time 1600 s_fla
Jan 1, 1911 @ 00:00:00.000 element PRCP @timestamp Oct 3, 2022 @ 23:08:05.541 date Jan 1, 1911 @ 00:00:00.000 id USC00510190 s_flag 6 value 325
Jan 1, 1911 @ 00:00:00.000 element PRCP @timestamp Oct 3, 2022 @ 23:08:05.549 date Jan 1, 1911 @ 00:00:00.000 id USC00510840 s_flag 6 value 406
Jan 1, 1911 @ 00:00:00.000 element PRCP @timestamp Oct 3, 2022 @ 23:08:05.549 date Jan 1, 1911 @ 00:00:00.000 id USC00510905 s_flag 6 value 97 _
Jan 1, 1911 @ 00:00:00.000 element PRCP @timestamp Oct 3, 2022 @ 23:08:05.549 date Jan 1, 1911 @ 00:00:00.000 id USC00510999 s_flag 6 value 310
Jan 1, 1911 @ 00:00:00.000 element PRCP @timestamp Oct 3, 2022 @ 23:08:05.555 date Jan 1, 1911 @ 00:00:00.000 id USC00511460 s_flag 6 value 127
Jan 1, 1911 @ 00:00:00.000 element PRCP @timestamp Oct 3, 2022 @ 23:08:05.556 date Jan 1, 1911 @ 00:00:00.000 id USC00511484 s_flag 6 value 64 _
Jan 1, 1911 @ 00:00:00.000 element PRCP @timestamp Oct 3, 2022 @ 23:08:05.641 date Jan 1, 1911 @ 00:00:00.000 id USC00512121 s_flag 6 value 264
Jan 1, 1911 @ 00:00:00.000 element PRCP @timestamp Oct 3, 2022 @ 23:08:05.641 date Jan 1, 1911 @ 00:00:00.000 id USC00512156 s_flag 6 value 109
Jan 1, 1911 @ 00:00:00.000 element PRCP @timestamp Oct 3, 2022 @ 23:08:05.631 date Jan 1, 1911 @ 00:00:00.000 id USC00511850 s_flag 6 value 254
Jan 1, 1911 @ 00:00:00.000 element PRCP @timestamp Oct 3, 2022 @ 23:08:05.631 date Jan 1, 1911 @ 00:00:00.000 id USC00511864 s_flag 6 value 165
Jan 1, 1911 @ 00:00:00.000 element PRCP @timestamp Oct 3, 2022 @ 23:08:05.641 date Jan 1, 1911 @ 00:00:00.000 id USC00511930 s_flag 6 value 246
Jan 1, 1911 @ 00:00:00.000 element PRCP @timestamp Oct 3, 2022 @ 23:08:05.667 date Jan 1, 1911 @ 00:00:00.000 id USC00512182 s_flag 6 value 58 _

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

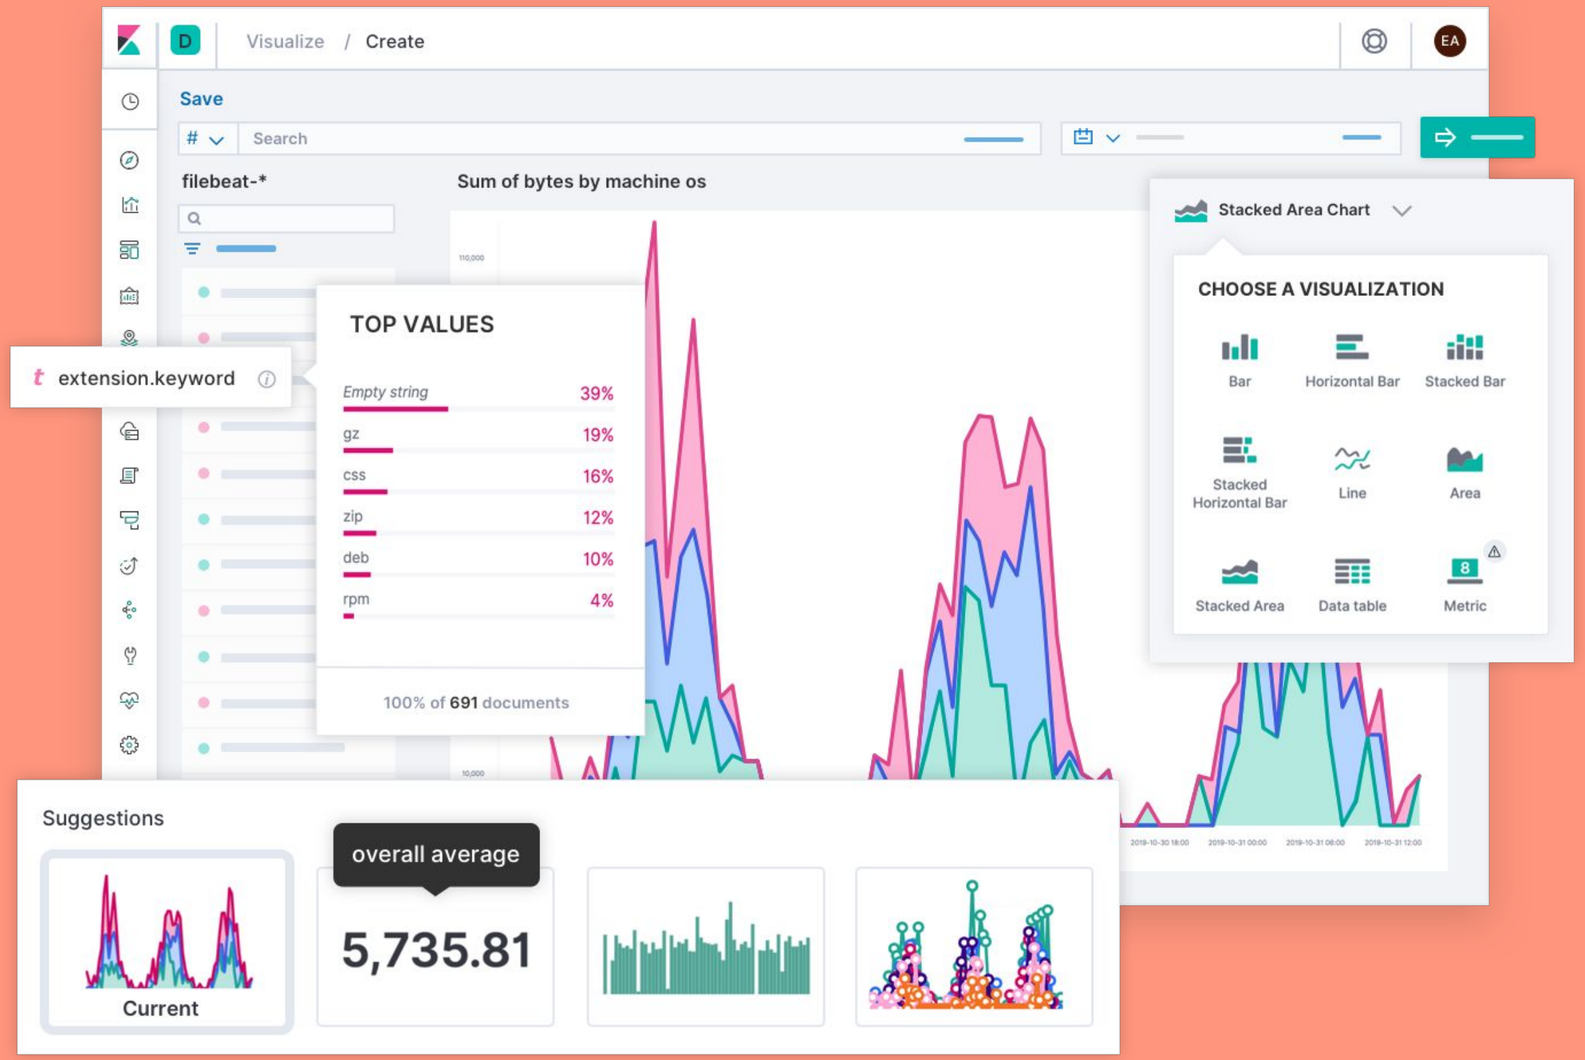

Dashboards

- 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

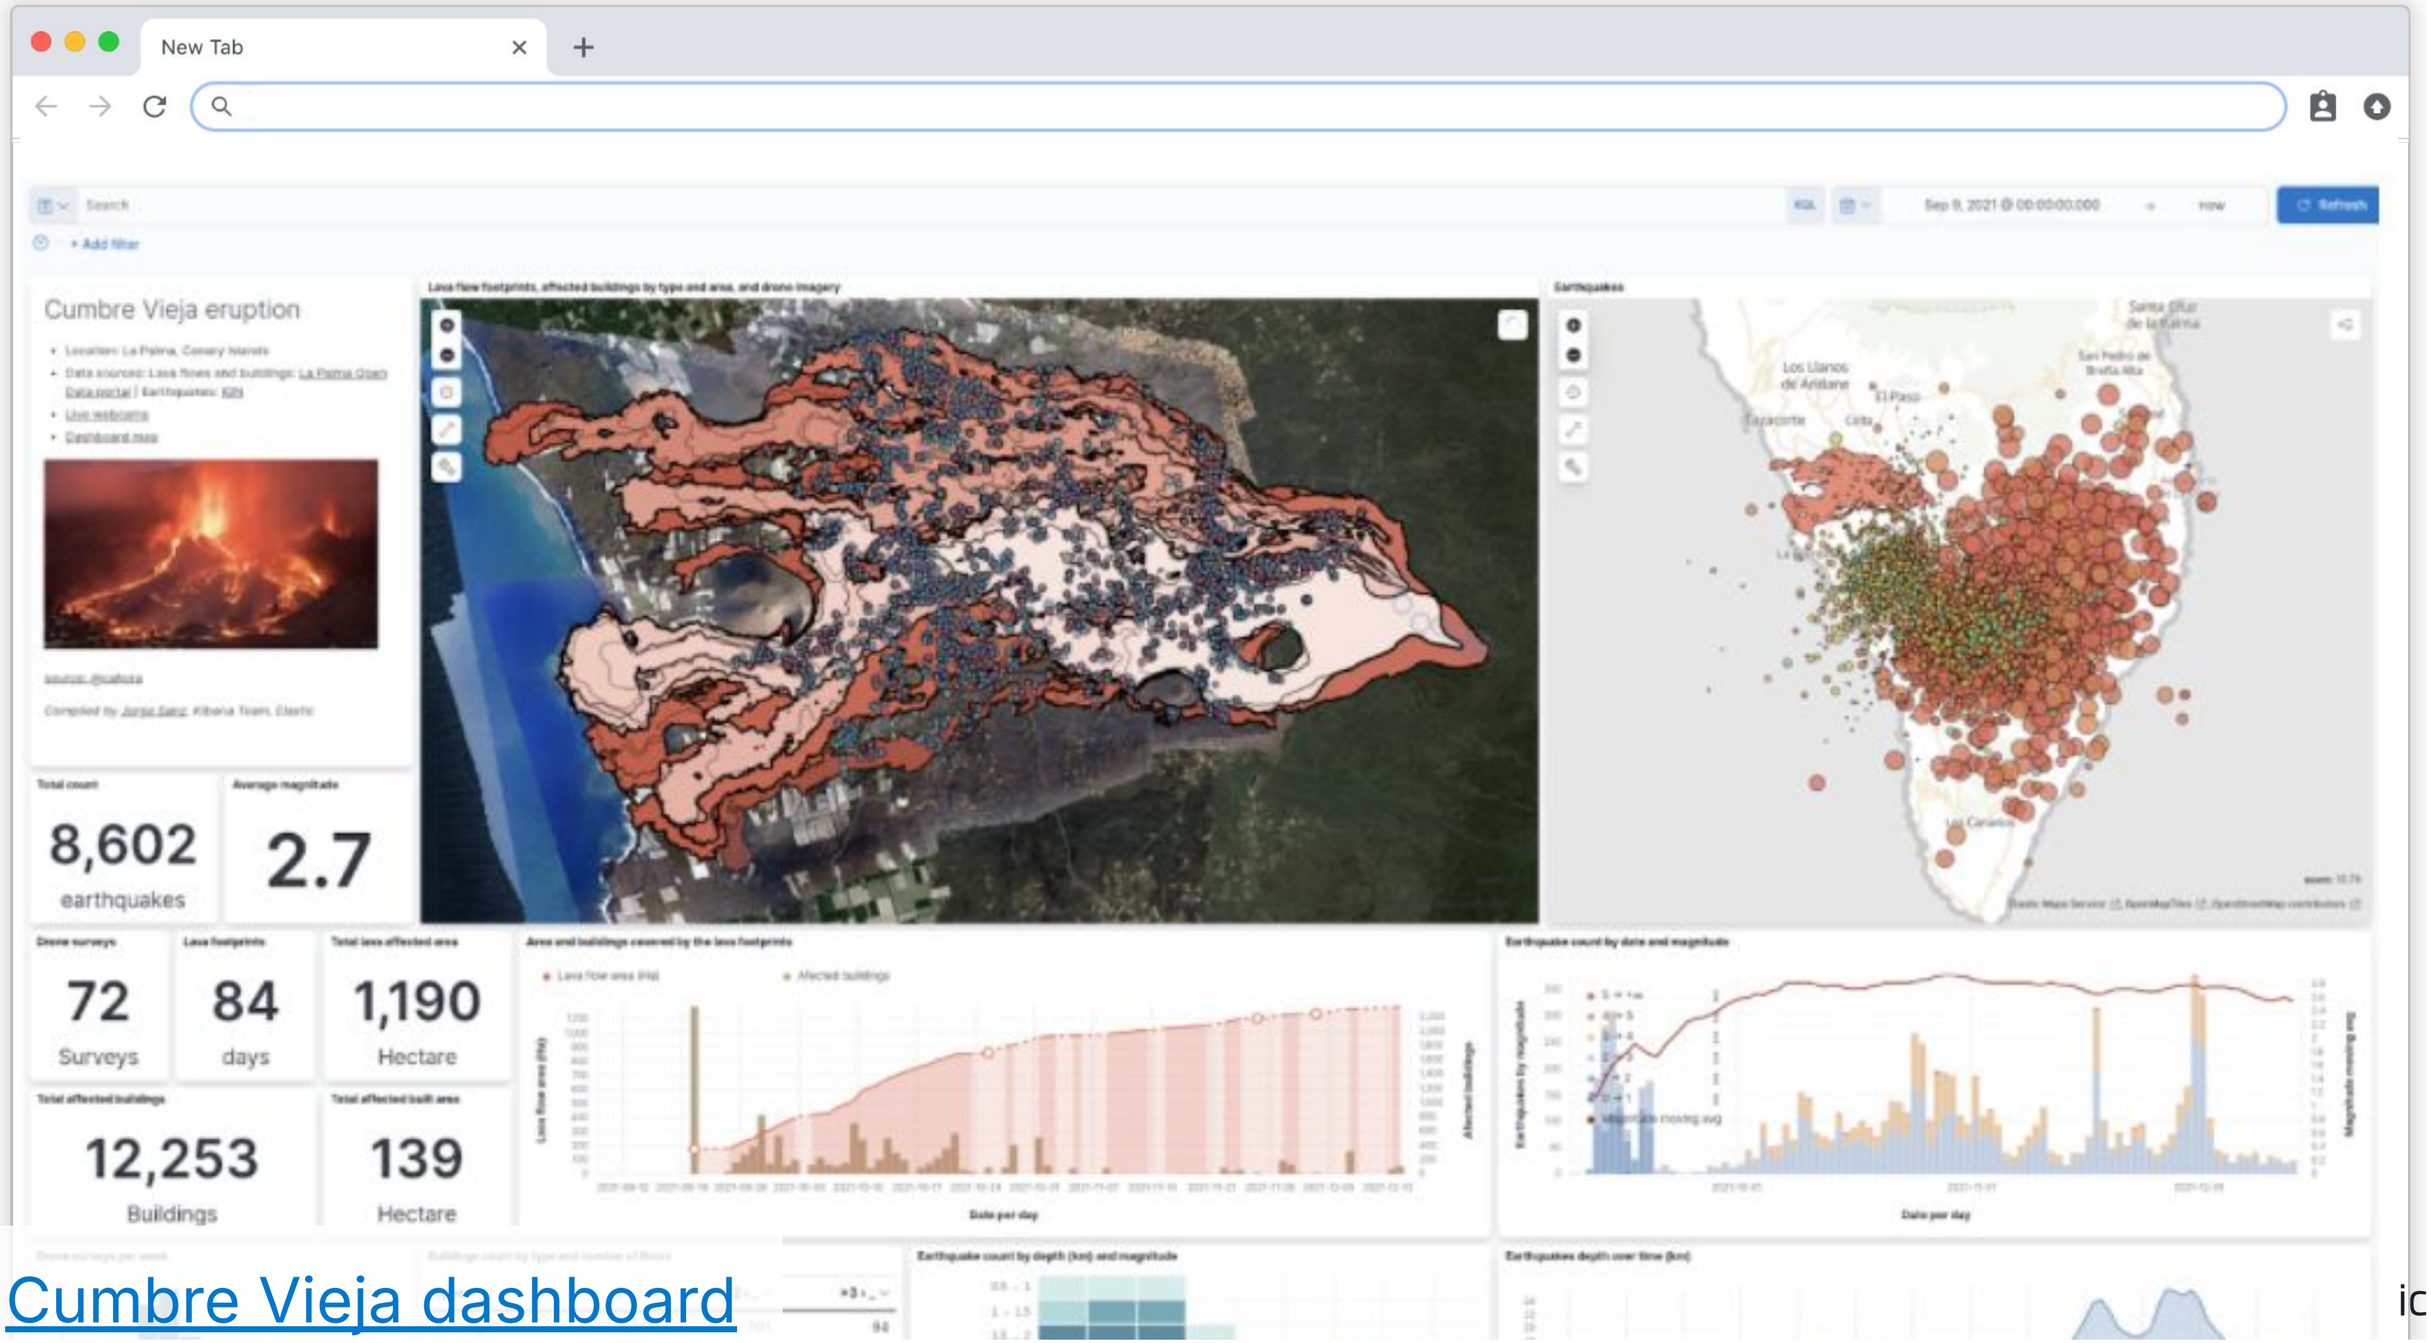

[Cumbre Vieja dashboard]

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

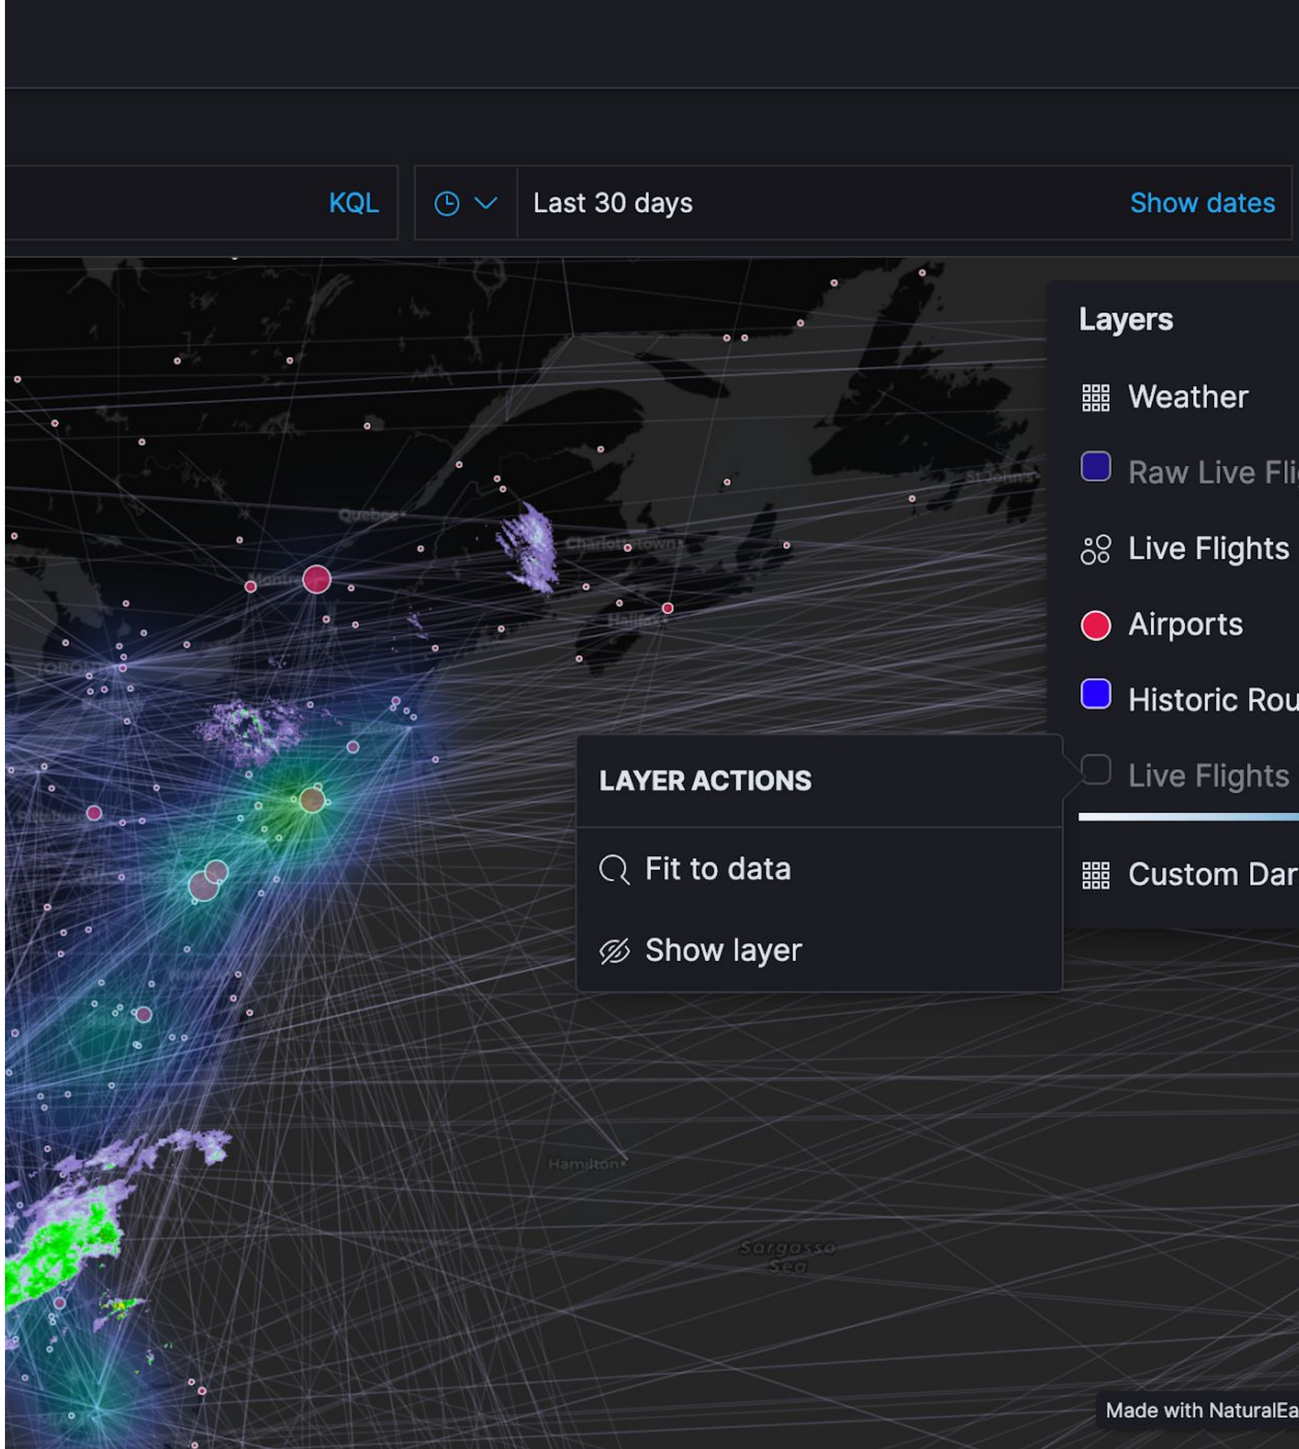

Quick web mapping intro

elastic

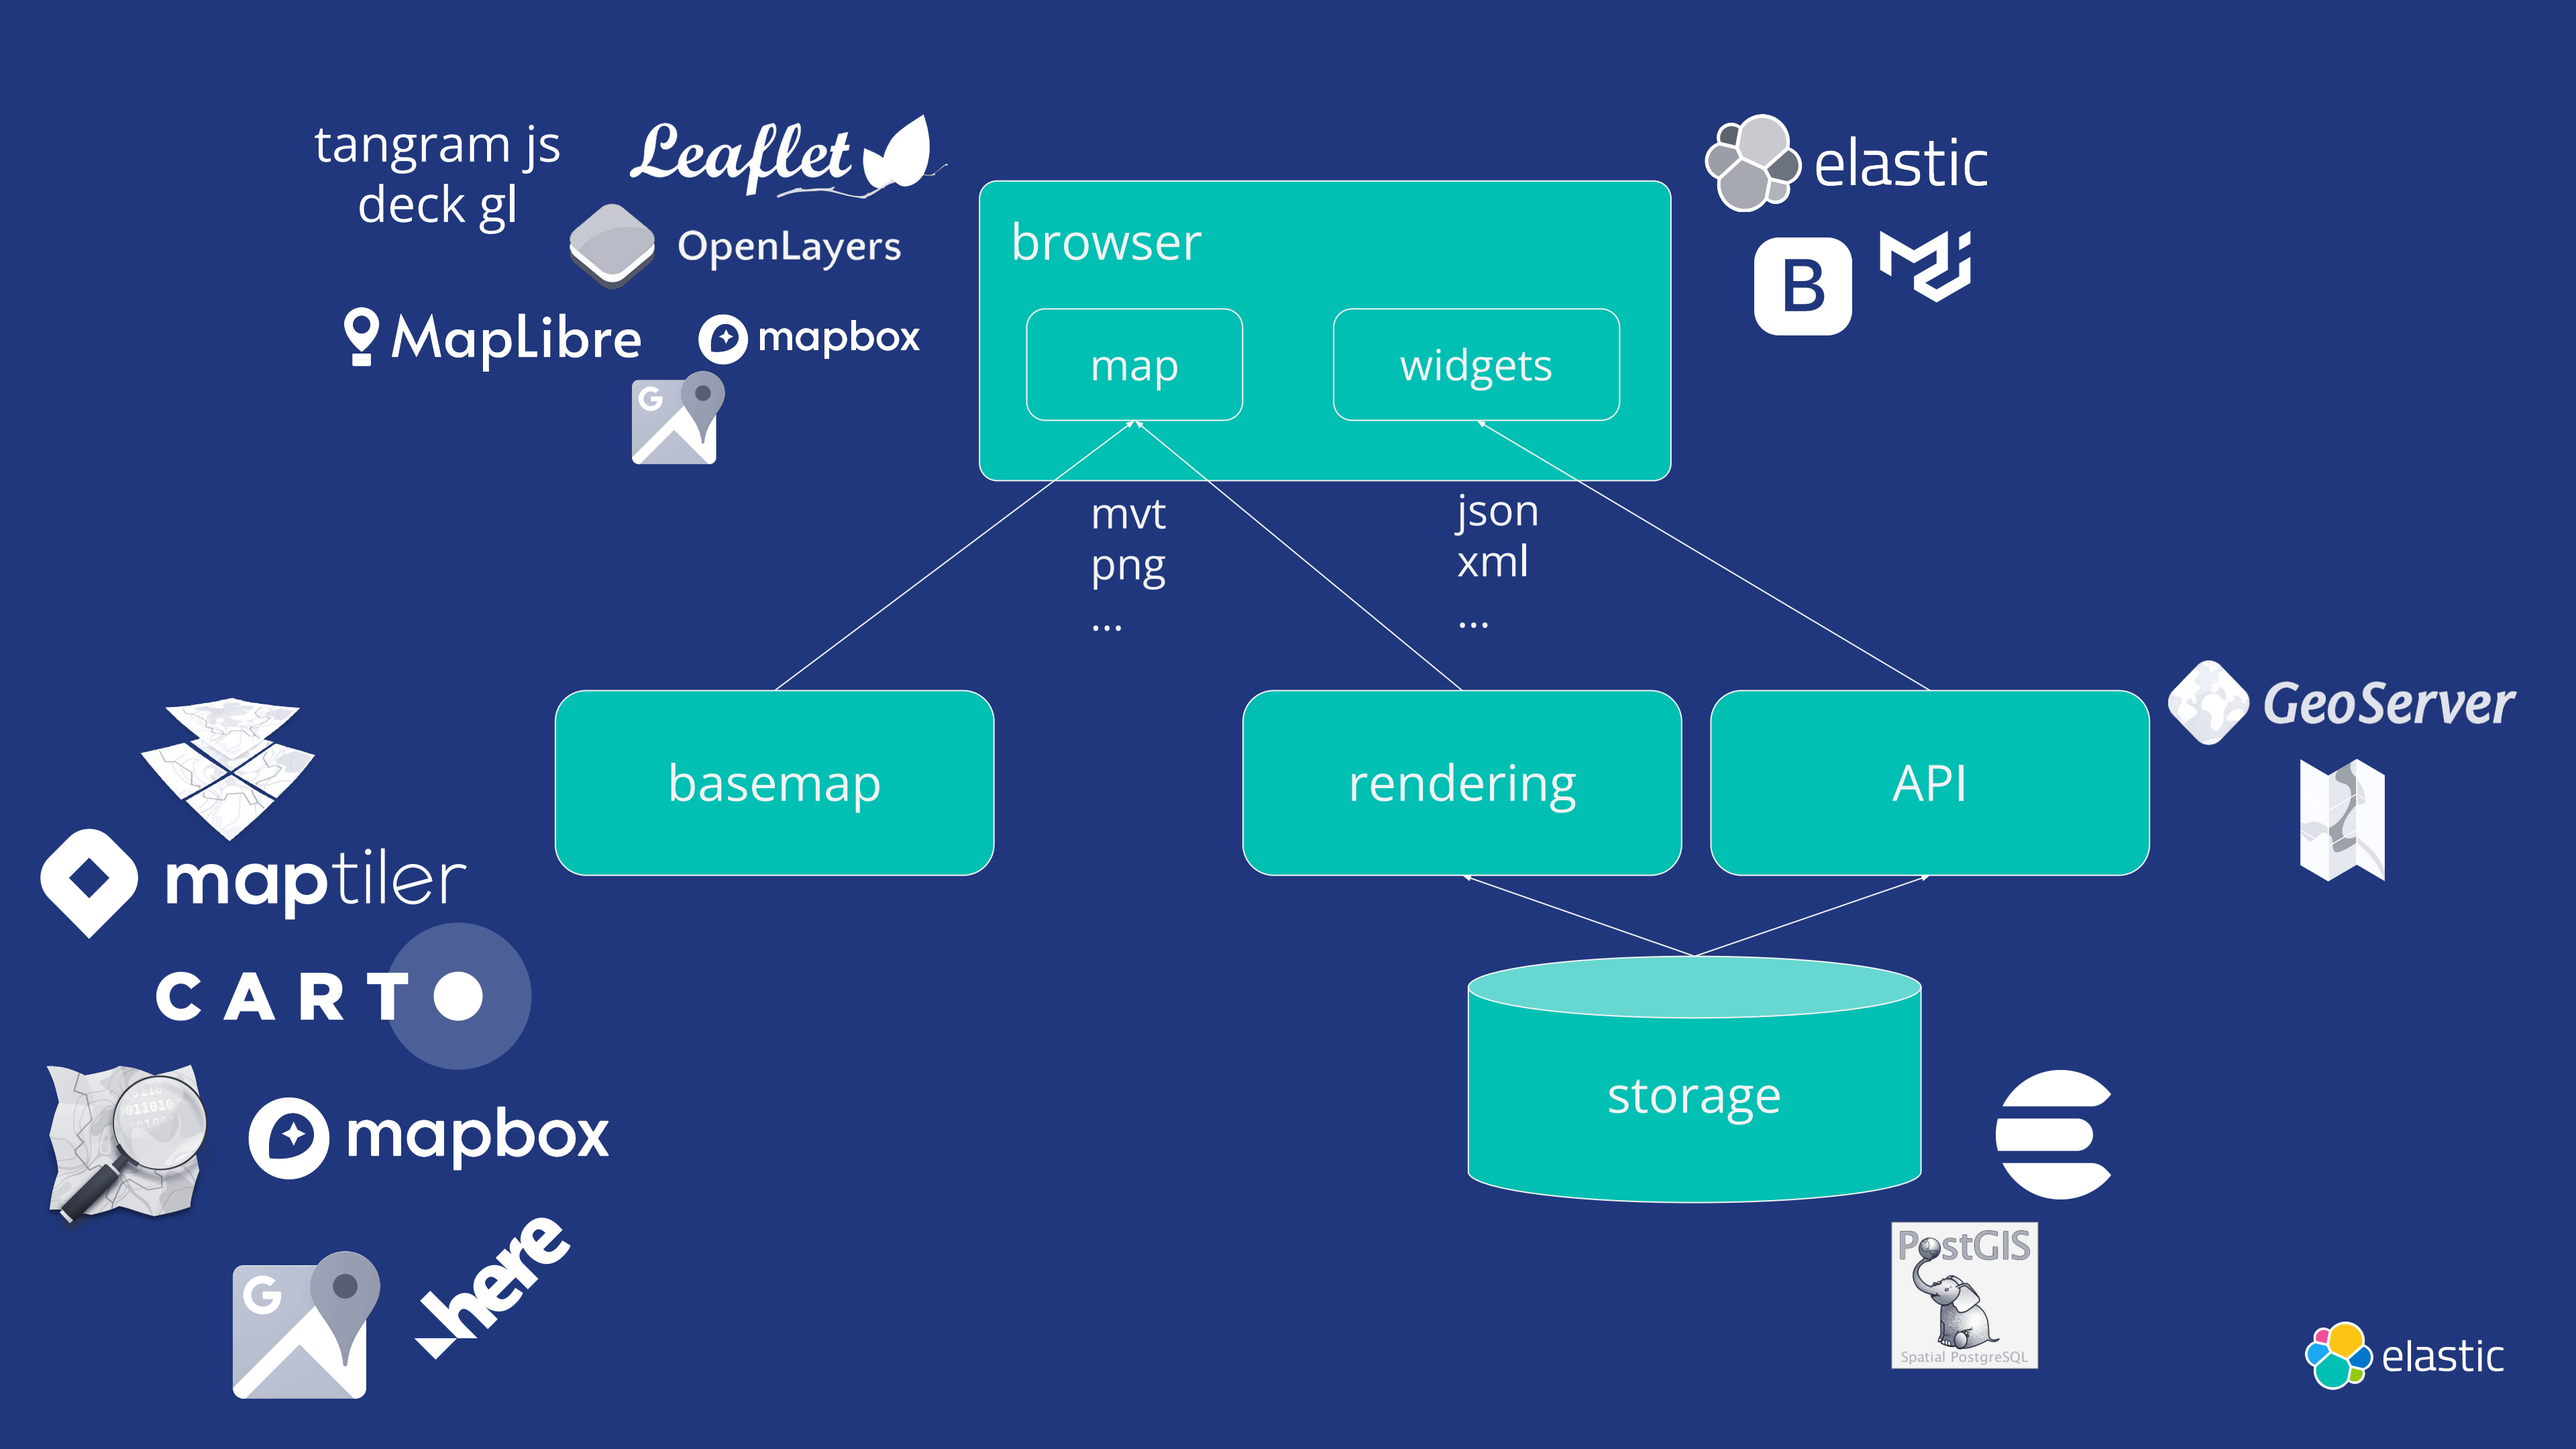
tangram js
deck gl
Leaflet
OpenLayers
MapLibre
mapbox
elastic
browser
map
widgets
mvt
png
...
json
xml
...
basemap
rendering
API
GeoServer
maptiler
CARTO
mapbox
here
storage
PostGIS
Spatial PostgreSQL
elastic

From last year workshop....

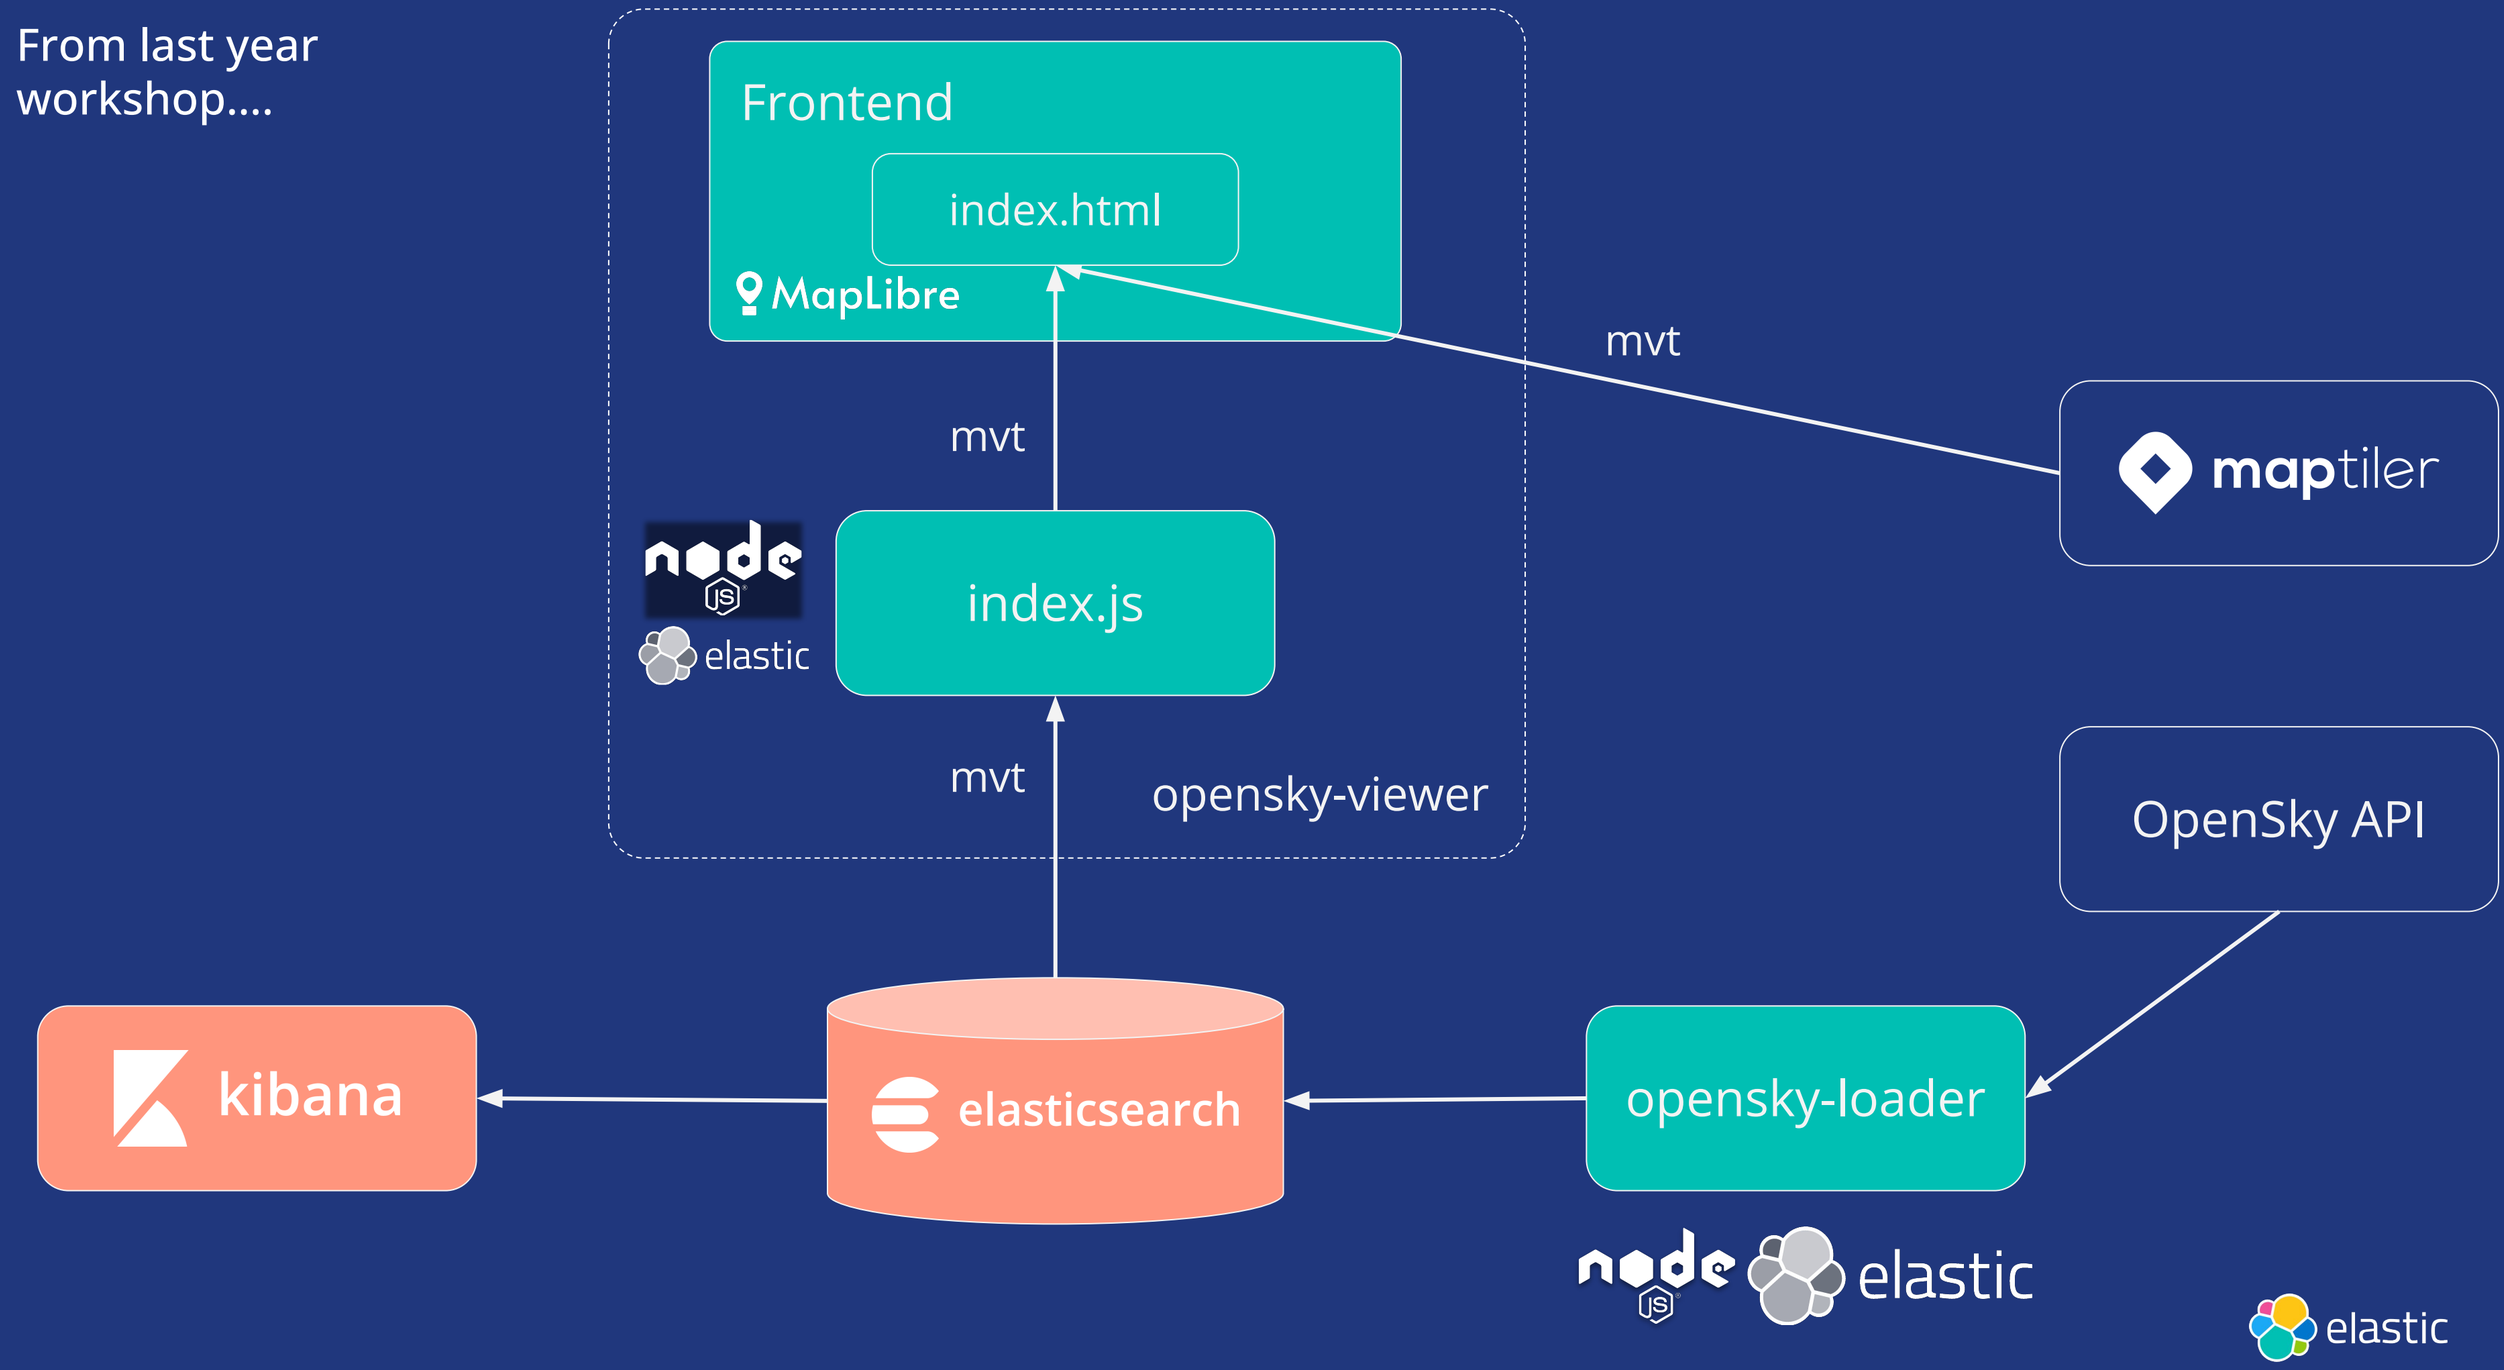

Today

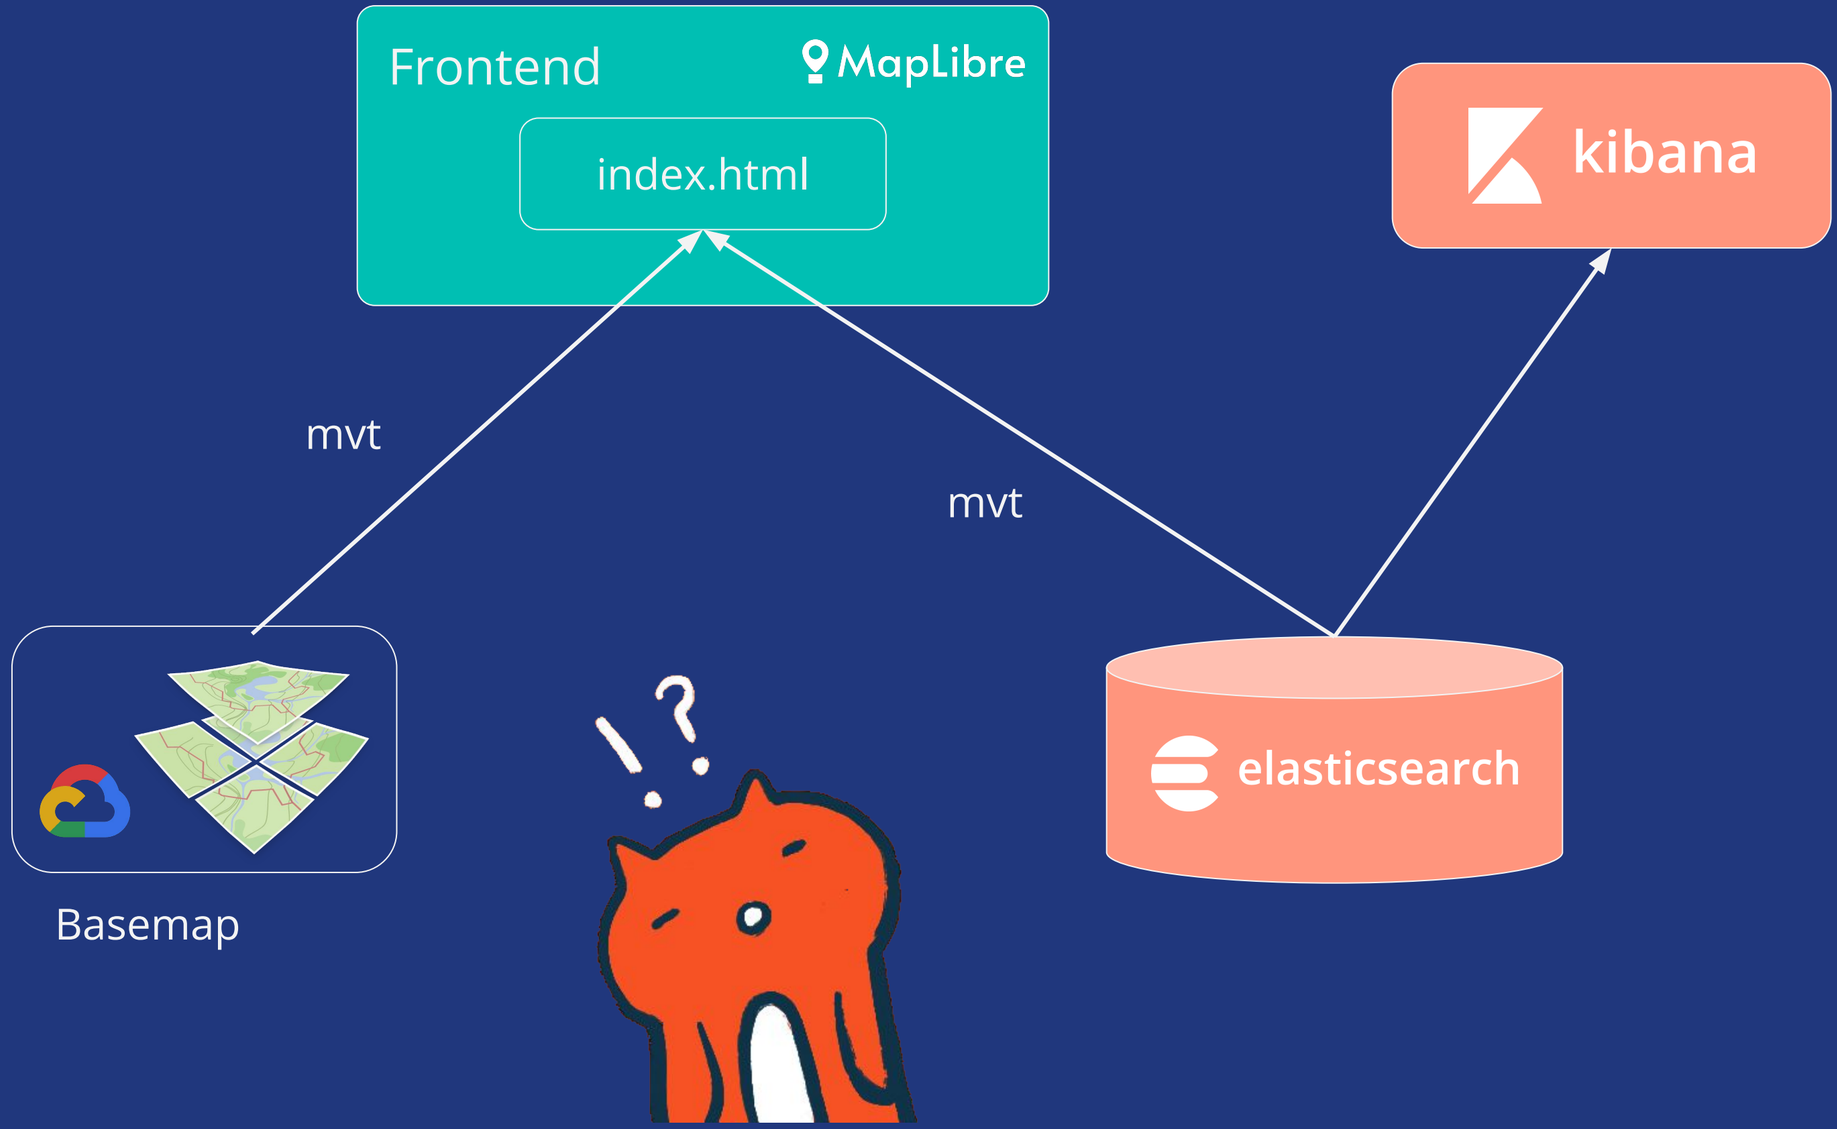

elastic

Simplifying our setup thanks to:

Changing a basemaps provider by a PMTiles file

Use planetiler and pmtiles to conveniently generate a basemap file suited for this project and removing the need for a basemaps provider.

Connecting to Elasticsearch with API keys

Setting up our cluster to allow well scoped API key requests.

Leveraging MapLibre transformRequest feature

With this feature, we can include a payload in our vector tile requests to perform arbitrarily complex search and aggregation queries.

elastic

Replacing a maps provider by our own vector tiles

- Use planetiler to quickly and easily generate vector tiles in the PMTiles format
  - All **Catalonia** region is generated in around **90s**
  - **Denmark** in less than 4 minutes

```
$ java -jar planetiler.jar openmaptiles \
  --download --keep_unzipped=true --area=cataluña \
  --output=data/catalonia.pmtiles
```

- For this workshop we will use a file that combines:
  - A **full planet** for zooms **1 to 6**
  - **New York City** for **zooms > 6**

elastic

Replacing a maps provider by our own vector tiles

- Adapt the OSM Bright style to consume this file

```
    ],
    "zoom": 1,
    "bearing": 0,
    "pitch": 0,
    "sources": {
      "openmaptiles": {
        "type": "vector",
        "url": "pmtiles://https://storage.googleapis.com/jsanz-bucket/planet/planet-nyc.pmtiles"
      }
    },
```

- Add the JavaScript `pmtiles` library and enable the protocol

```
/* initialize pmtiles support */
let protocol = new pmtiles.Protocol();
maplibregl.addProtocol("pmtiles",protocol.tile);
```

elastic

Exposing Elasticsearch to the Internet securely

Elasticsearch can be accessed anonymously 📖

```
xpack.security.authc:
  anonymous:
    username: anonymous_user 1
    roles: role1, role2 2
    authz_exception: true 3
```

Instead, we will use an API key 📖 to read dedicated indices

```
POST /_security/api_key
{
  "name": "workshop-api-key",
  "expiration": "5d",
  "role_descriptors": {
    "workshop": {
      "index": [
      {
        "names": [
          "geonames",
          "311",
          "nyc_boroughs"
        ],
        "privileges": [
          "read",
          "view_index_metadata"
        ],
        "field_security": {
          "grant": [
            "*"
          ]
        }
```

elastic

Exposing Elasticsearch to the Internet securely

CORS is disabled by default 📖

```
http.cors:
    enabled : true
    allow-origin: "*"
    allow-methods: OPTIONS, HEAD, GET, POST
    allow-headers: "X-Requested-With, Content-Type, Content-Length, Authorization, Accept, Us
```

Our cluster is ready to accept API key requests 🥳

```
$ ELASTIC_HOST="https://your-cluster-url"
$ ELASTIC_APIKEY="your-encoded-name-and-api-key-here"
$ curl -H "Authorization: ApiKey ${ELASTIC_APIKEY}" \
  "${ELASTIC_HOST}/geonames/_count?pretty=true"
```

elastic

How to make Elasticsearch queries from Maplibre?

Vector Tile servers understand GET

In general templates for querying vector tile servers contain all parameters in the URL like

`http://myserver/{z}/{x}/{y}.[pbf|png]?search={query}`

Elasticsearch `_mvt` and `_search` endpoints really needs a payload

Some parameters are allowed in the URL but most of the extensive capabilities for searching are only available as payloads on POST requests

Maplibre allows to "hack" the requests for tiles

The `transformRequest` map creation option allows arbitrary changes to each HTTP request, even changing the method.

elastic

How to make Elasticsearch queries from Maplibre?

```
const map = new maplibregl.Map({
    container: 'map',
    style: MAP_STYLE,
    center: [-73.95, 40.7],
    zoom: 10,
    hash: true,
    transformRequest: function (url, resourceType) {
        /* This function enriches the HTTP request to include
        the ES search body, change to a POST request, and include
        the Content-Type and authorization headers */
        if (resourceType == 'Tile' && url.startsWith(ES_HOST)) {
            return {
                url: url,
                method: 'POST',
                headers: {
                    'Content-Type': 'application/json',
                    'Authorization': `ApiKey ${ES_APIKEY}`
                },
                type: 'arrayBuffer',
                body: JSON.stringify(ES_SEARCH_BODY)
            }
        }
    }
});
```

elastic

Laboratory

Let's draw some maps, finally!

elastic

Set up

Get the **code** we will review together

- Easiest way, hit the “Remix” button at this glitch project:
  - https://glitch.com/edit/#!/living-crocus-vermicelli
- If you are quick with git(hub) and node:
  - Clone the repo jsanz/elastic-workshop
  - Run the project inside lab/vector-tile-viewer
  - Play with HTML documents inside pages folder

Open **Kibana** with the read-only credentials shared on the session and check:

- Dev Console
- Discover
- Maps
- Dashboards

elastic

Just a basemap

- Code
- Create a map with the OSM bright style

```
<div id='map'></div>
<script>
        const map = new maplibregl.Map({
            container: 'map',
            style: MAP_STYLE,
            center: [-73.95, 40.7],
            zoom: 10,
            hash: true
        });

        map.addControl(new maplibregl.NavigationControl());

</script>
```

- Extra: consider updating the style to point to other tilesets available in the same folder: catalonia.pmtiles, denmark.pmtiles, kosovo.pmtiles, andorra.pmtiles

elastic

First documents from Elasticsearch

- Code
- Define a query and a new vector layer
- Count geometries from the rendered features

- Extra:
  - consider checking with Discover the data to select another time frame
  - Render geonames dataset

elastic

Documents from Elasticsearch themed by complaint type

- Code
- Extend the query to filter by terms
- Thematic mapping with Maplibre styling

- Extra:
  - Render by another field
  - Use other conditions to filter your data (Maps, Discover and their *Inspect tool* will be your friends for this)

elastic

Documents from Elasticsearch themed by complaint type

- Code
- Add more fields to the vector tiles responses
- Include a basic popup implementation

- Extra:
  - Play with the pop up template
  - Tooltip on hover?

elastic

Search and filter documents from Elasticsearch

- Code
- Add a simple form with a text input
- Update the ES body to include the search query
- Reload the layer

- Extra:
  - Refine the search to query only a single field
  - Add more fields to extend the query (date filter, agency selector, etc.)

elastic

Render aggregated data into H3 hexagons

- Code
- Update the query with `grid_agg`, `grid_type` properties
- Pop up now works as a tooltip showing the count
- New problem: the legend is fixed but our counts heavily depend on the zoom level!

- Extra:
  - Do you notice an outlier?
  - Play with the `grid_precision` query property

elastic

Adapt the legend to the zoom level using basic stats

- Code
- Add a listener to the `zoomend` event to update the style based in the new maximum for the aggregation count and reload the layer.
- Include a "must_not → geo_distance" filter to remove those picky new yorkers

- Extra:
  - Use Kibana Maps to identify that outlier
  - Find other filter types to remove it

elastic

Add a new metric: cardinality

- Code
- Include a new aggregation to compute how many different complaint types exist per hexagon.
- Increase the number of steps in the legend for better display

- Extra:
  - Find other interesting metrics available with Kibana Maps.

elastic

Aggregate using mercator tiles

- Code
- Change the grid_agg type to geotile (square mercator tiles).
- This aggregation type is much faster

- Extra:
  - Play with the max and min zoom levels of the application
  - Play with different color schemes and get help from Kibana Maps (https://colorbrewer2.org is a classic)

elastic

Aggregate using a heat map

- Code
- A heatmap is in reality a grid with a custom styling
- Increase the grid_precision to 8

- Extra:
  - Figure out why it crashes on Firefox 😢
  - Play with different color schemes and heatmap parameters

elastic

Closing

Some final remarks

- Everything we covered in this session is available in the free Basic license (cloud or self-hosted).
- This setup is OK for controlled environments
- A proper backend to interact with Elasticsearch is the recommended approach
- Elasticsearch offers a wide variety of capabilities
  - `geo_line` aggregation and `geo_shape` aggregation (hex and tile)
  - Time Series, data streams, ingest pipelines, transforms, Cross Cluster Search, Cross Cluster Replication, ...
  - Including Artificial Intelligence and large language models!

elastic

Thanks!

Webmapping with Elasticsearch

Jorge Sanz | jorge.sanz@elastic.co