OSUG

Observatoire des
Sciences de l'Univers
de Grenoble

De geonetwork à elasticsearch

Préambule

3 bonnes raisons pour ne pas faire :

- *Je ne suis pas le plus compétent pour cela*
- *Ce n'est pas mon job*
- *Il aurait fallu qu'il s'intègre dans le processus des référents*

1 mauvaise raison pour faire :

« Parfois, au lieu de faire ce que le monde attend de vous, faites ce que vous avez envie de faire »

Geonetwork à l'OSUG

- geonetwork
  - Catalogue de métadonnées **géoréférencées**

- Existant à l'OSUG
  - Geonetwork envirhonalpes : https://osug-geonetwork.osug.fr
  - Geonetwork « OSUG » data.osug.fr qui moissonne leca-bdgis et https://data.Irstea qui sont tous les deux offline pour des raisons différentes
  - Contenu et version de geonetwork sont obsolete
- Enjeu 1ere étape

  ⇒ mettre en place un nouveau geonetwork avec au départ les utilisateurs identifiés du service auxquels s'agrègeront les autres demandeurs

  ⇒ ouvrir ce service aux interessés et le cas échéant récupérer leurs données. Aujourd'hui :
  - Julien Renaud LECA leca-bdgis.u-ga.fr, Lucie Liger Jardins du Lautaret, Mathilde Ratouis Zone Atelier Alpes, Cecile Pignol Edytem USMB
  - Coté OSUG Charly, Rémi et moi

Geonetwork et OGC-CWS,

OGC : *Open Geospatial Consortium*

- OGC – CWS protocole de moissonnage
  - Les geocatalogues peuvent s'enrichir par moissonnage (harvesting)
  - L'identifiant unique sur les métadonnées et les données permet de garantir l'unicité des métadonnée d'un dataset dans le cas de moissonnage en boucle

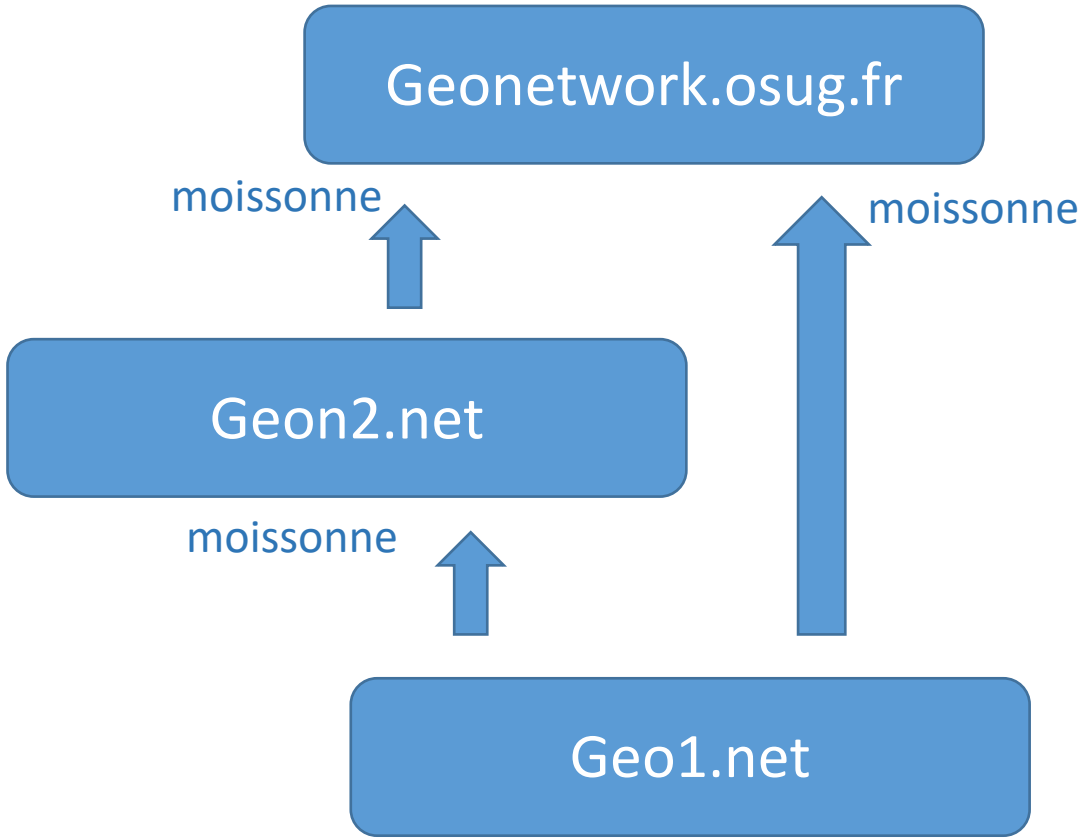

Architecture de geonetwork

https://geonetwork.osug.fr

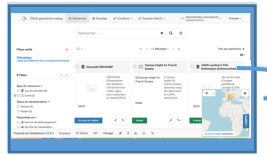

https://geonetwork.osug.fr/geoserver

WMS

Ajout de layer dans les applications telles que googleearth

Ha_proxy

osug-geonetwork-1.u-ga.fr

tomcat

webapps

Geoserver
Serveur cartographique

WMS
et autre

Geonetwork
Catalogue de metadonnées géoréférencées

POSTGRES + POSTGIS

Base de données geonetwork

osug-postgres.u-ga.fr

Elasticsearch

Recherche indexée
Remplace code basé sur lucene

Architecture de geonetwork

https://geonetwork.osug.fr

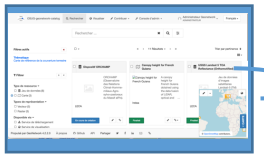

https://geonetwork.osug.fr/geoserver

WMS

Ajout de layer dans les applications telles que googleearth

Ha_proxy

osug-geonetwork-1.u-ga.fr

tomcat

webapps

Geoserver
Serveur cartographique

WMS
et autre

Geonetwork
Catalogue de metadonnées géoréférencées

POSTGRES + POSTGIS

Base de données geonetwork

osug-postgres.u-ga.fr

Elasticsearch

Recherche indexée

Elasticsearch

OSUG

Chaine de traitement des logs

OSUG

Chaine de traitement des logs

- Objectif répondre à un besoin de tableau de bord

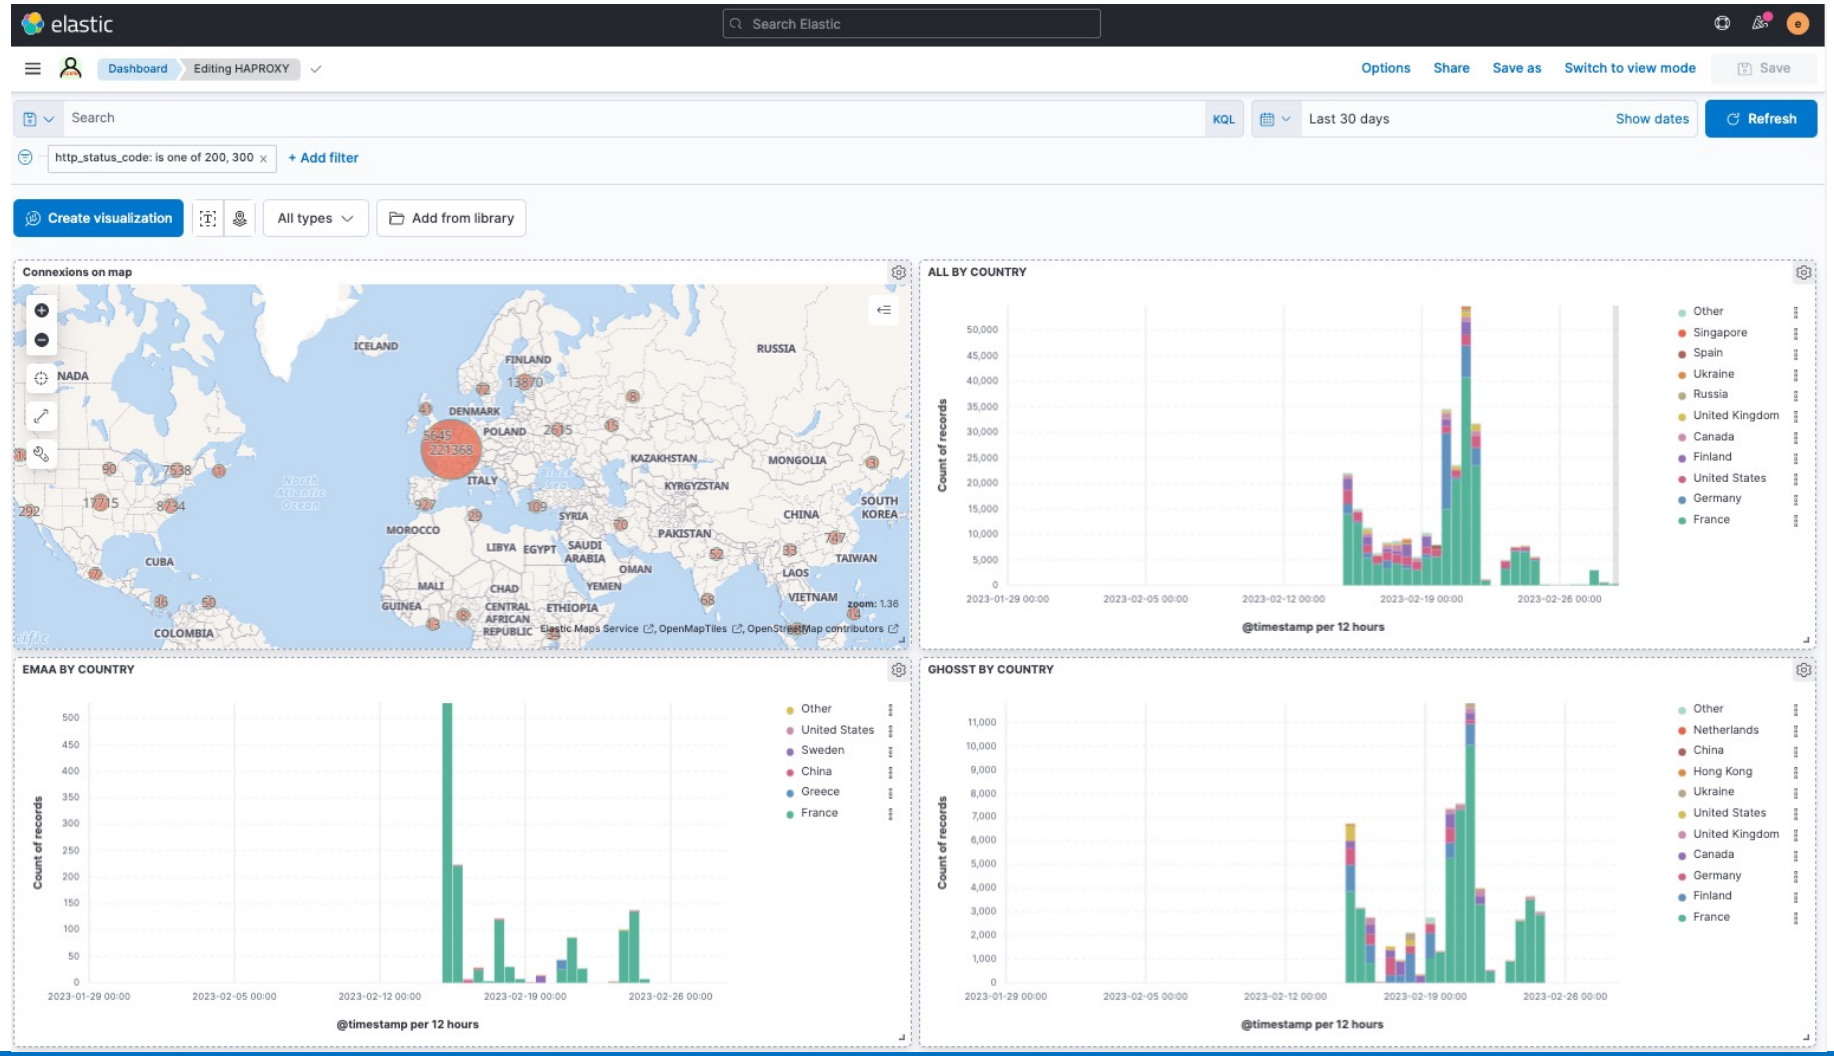

Au départ il y a le log

https://emaa.osug.fr

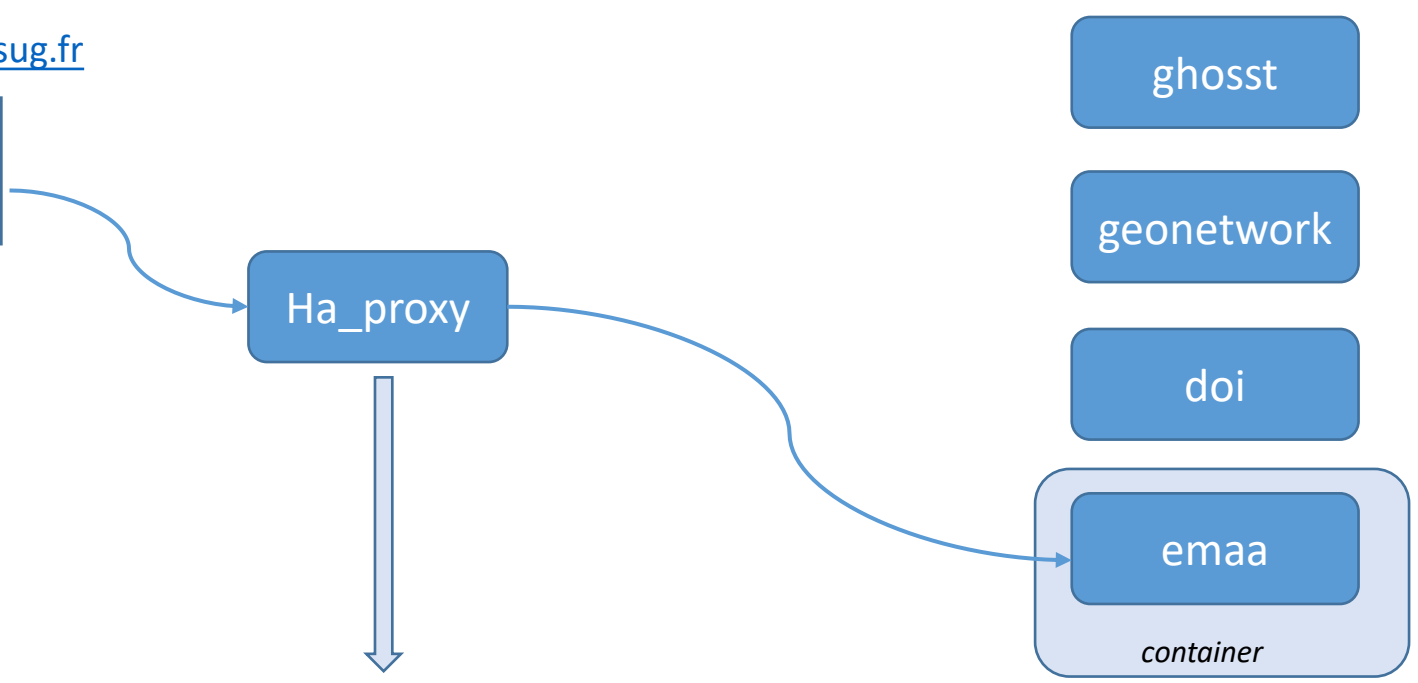

/var/log/ha_proxy.log

…

```
Feb 27 13:13:42 osug-frontend2 haproxy[809770]: 147.171.169.181:52349
[27/Feb/2023:13:13:42.731] http-in~ vsphere_cattle/cattle1 0/0/0/148/148 200 470 - - --VU
6/3/0/0/0 0/0 {emaa.osug.fr|Mozilla/5.0 (Macintosh; Intel Mac OS X 10.15; rv:109.0)
Gecko/20100101 Firefox/112.0} "GET /targets/[13C] HTTP/1.1"
```

Logstash met en forme le log

```
Feb  7 09:22:04 osug-frontend2 haproxy[669021]: 152.77.119.244:53750 [07/Feb/2023:09:22:04.178] http-in~
geonetwork/geonetwork 0/0/0/1/1 304 205 - - ---- 13/7/1/1/0 0/0
{geonetwork.osug.fr|like Gecko) Chrome/108.0.0.0 Safari/537.36}
"GET /geonetwork/static/ng-skos.css?v=e4951c3279b8aff0499cc9fa2beb20e4b1183e12 HTTP/1.1"
```


```
{
  "server_name": "geonetwork",
  "srvconn": "1",
  "time_backend_response": "1",
  ...
  "client_port": "53750",
  "backend_name": "geonetwork",
  ...
  "client_ip": "152.77.119.244",
  "user_agent": "like Gecko) Chrome/108.0.0.0
Safari/537.36",
  "http_status_code": "304",
  "syslog_server": "osug-frontend2",
  "req_host": "geonetwork.osug.fr",
  "month": "Feb",
  "accept_date": "07/Feb/2023:09:22:04.178",
  "frontend_name": "http-in~",
  ...
  "http_request": "/geonetwork/static/ng-
skos.css?v=e4951c3279b8aff0499cc9fa2beb20e4b1183e12",
}
```

OSUG

Logstash, parlez-vous « grok »?

```
filter {
    grok {
      match => { "message" => "%{MONTH:month}\s*%{MONTHDAY:day}\s%{TIME:time}
%{IPORHOST:syslog_server} %{SYSLOGPROG}: %{IP:client_ip}:%{INT:client_port}
\[%{HAPROXYDATE:accept_date}\] %{NOTSPACE:frontend_name}
%{NOTSPACE:backend_name}/%{NOTSPACE:server_name}
%{INT:time_request}/%{INT:time_queue}/%{INT:time_backend_connect}/%{INT:time_bac
kend_response}/%{NOTSPACE:time_duration} %{INT:http_status_code}
%{NOTSPACE:bytes_read} %{DATA:captured_request_cookie}
%{DATA:captured_response_cookie} %{NOTSPACE:termination_state}
%{INT:actconn}/%{INT:feconn}/%{INT:beconn}/%{INT:srvconn}/%{NOTSPACE:retries}
%{INT:srv_queue}/%{INT:backend_queue} \{(?<req_host>[0-9A-Za-
z.]*)\|(?<user_agent>[0-9A-Za-z \)./]*)\}?( )?\"(<BADREQ>|(%{WORD:http_verb}
(%{URIPROTO:http_proto}://)?(?:%{USER:http_user}(?::[^@]*)?@)?(?:%{URIHOST:http_
host})?(?:%{URIPATHPARAM:http_request})?( HTTP/%{NUMBER:http_version})?))?\"(
%{NOTSPACE:tls_version})?( %{NOTSPACE:tls_ciphersuite})?"
}
```

OSUG

Logstash, parlez-vous « grok »?

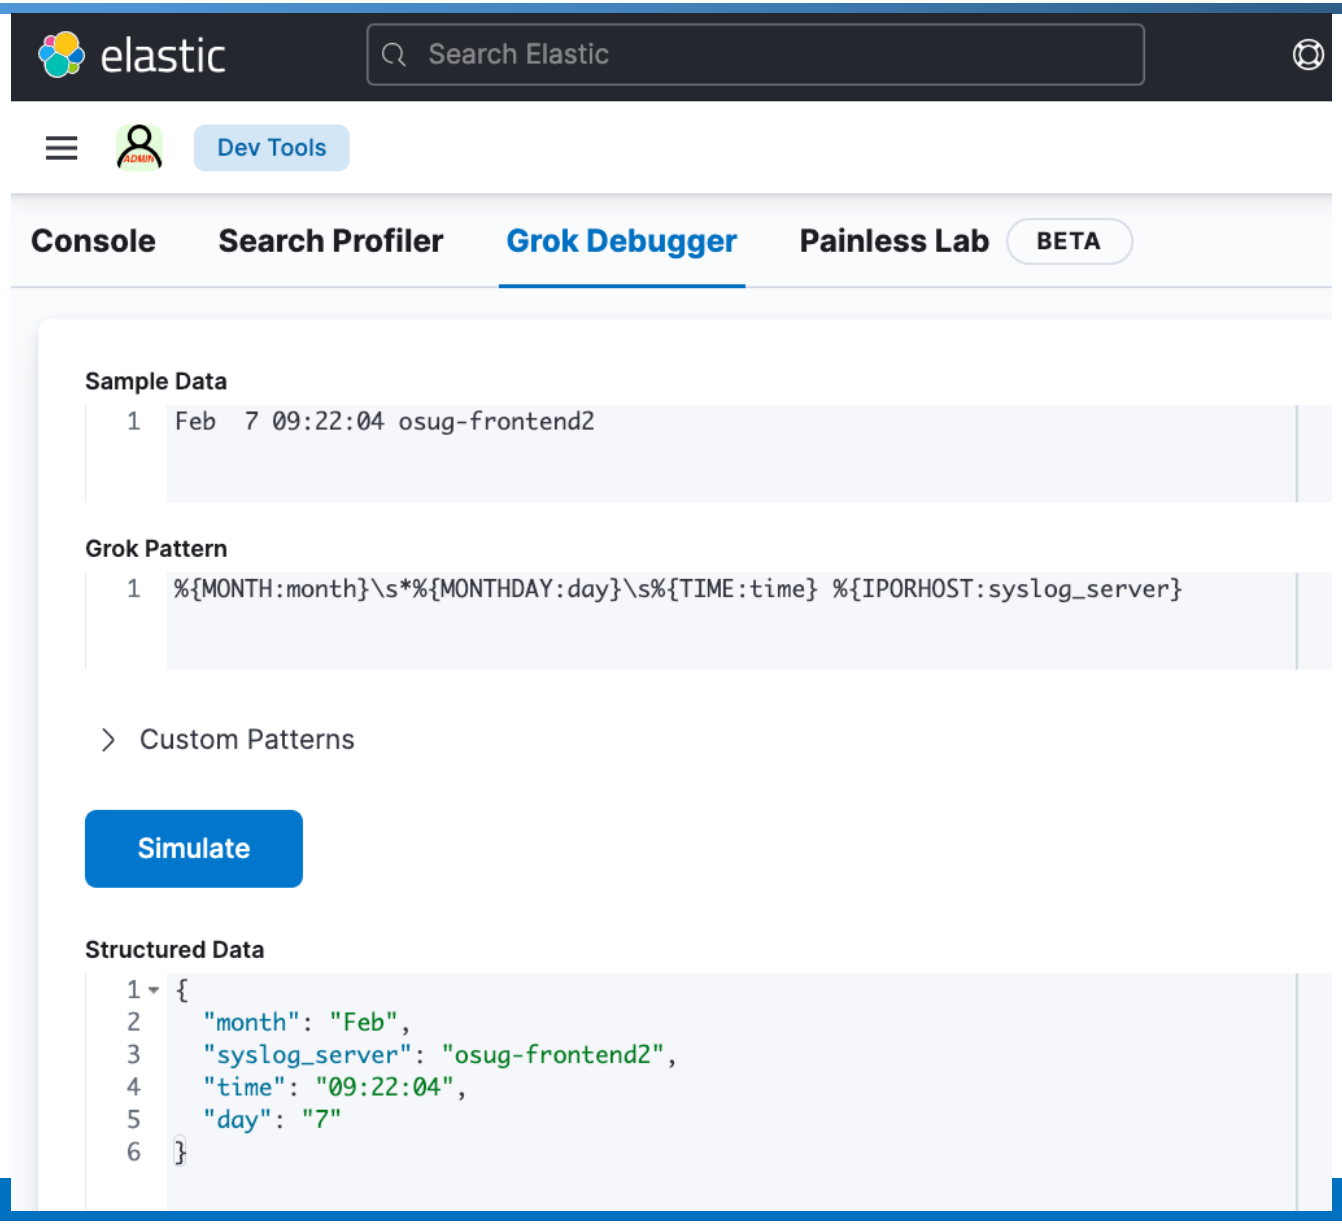

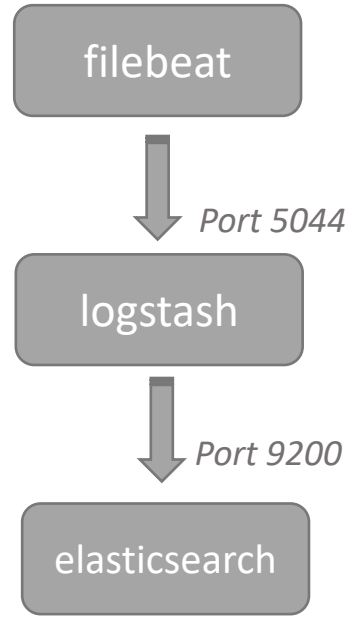

elasticsearch

- Les échanges avec elasticsearch se font via le port 9200 :

```
sysosug@osug-geonetwork-1:~$ curl -u elastic:$PASS "localhost:9200/_cat/indices?v"
health status index                    uuid                   pri rep docs.count docs.deleted store.size pri.store.size
yellow open   gn-features              GXLeTKb7SK-etCRBDPTFEw   1   1       1777            1     47.9mb         47.9mb
yellow open   gn-records               i7JlNDQQR_WHTiCUsydyng   1   1        816          948     10.6mb         10.6mb
yellow open   gn-searchlogs            uJdZ2TwcSeahbYg9SY-bsw   1   1          0            0       226b           226b
yellow open   logstash-2023.02.23-000001 Qi6aIPXfRkaEAv1O4NbfUA   1   1     972150            0    441.5mb        441.5mb
    ...
```

OSUG

elasticsearch

```
curl -u elastic:$PASS  -X GET "localhost:9200/logstash-2023.02.23-000001
                                                /_search?pretty&q=response=emaa"
```

```
"geoip" : {
  "location" : {
    "lon" : 5.7205,
    "lat" : 45.1766
  },
  "postal_code" : "38000",
  "latitude" : 45.1766,
  "city_name" : "Grenoble",
  "country_name" : "France",
  "longitude" : 5.7205,
  "region_name" : "Isère",
  "country_code2" : "FR",
  "ip" : "130.190.106.203",
  "region_code" : "38",
  "country_code3" : "FR",
  "continent_code" : "EU",
  "timezone" : "Europe/Paris"
```

```
...
  "http_status_code" : "301",
  "haproxy_month" : "Feb"
  "http_verb" : "GET",
  "retries" : "0",
  "accept_date" : "21/Feb/2023:11:07:54.643",
  "http_request" : "/public/EMAA/EMAA_para
-H2O_Rotation_18d206fb.html",
  "@timestamp" : "2023-02-21T10:07:54.643Z",
...
```

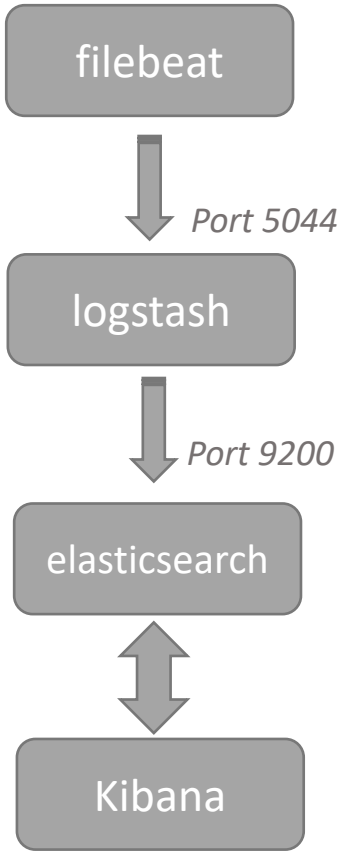

OSUG

Kibana

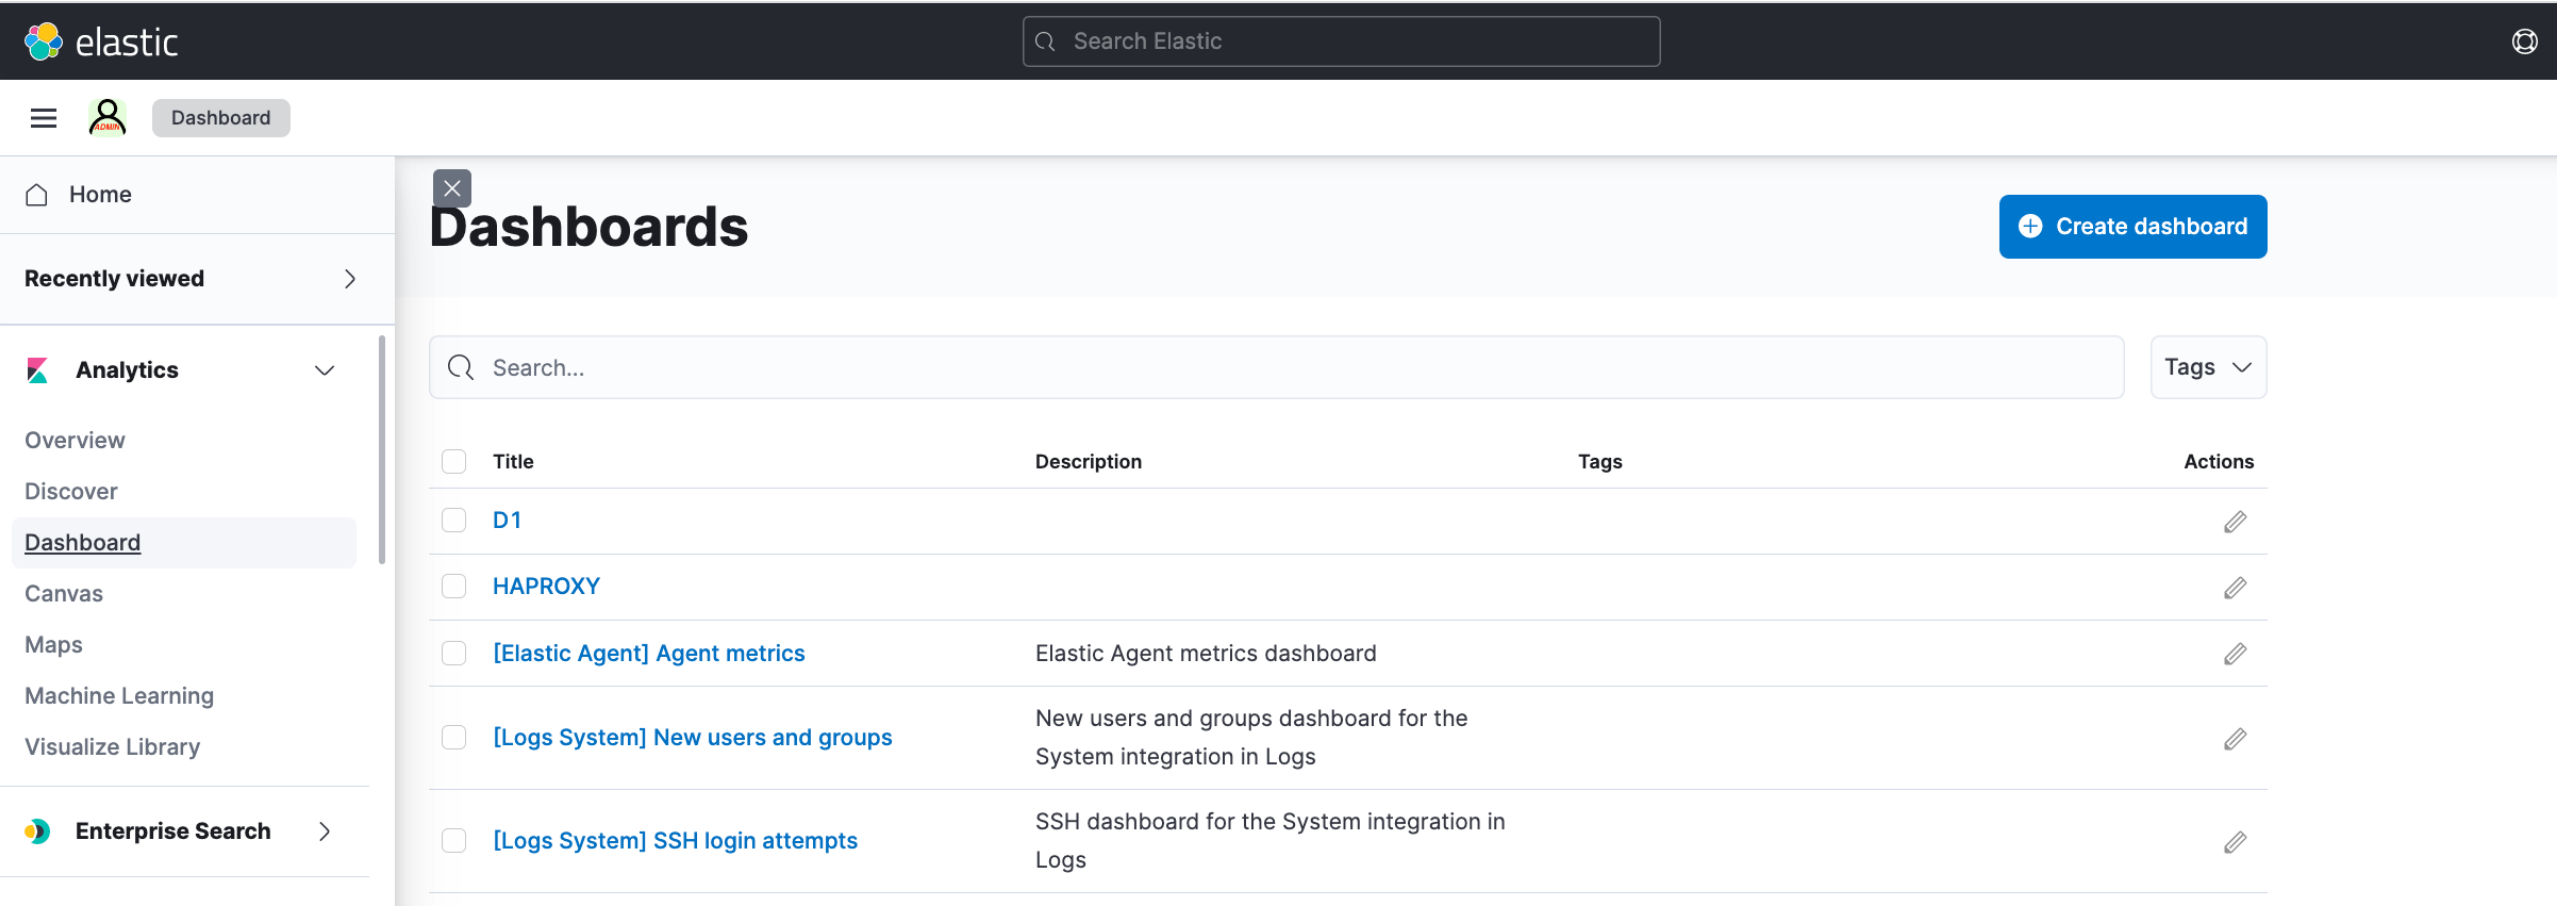
elastic
Search Elastic
Dashboard
Home
Recently viewed
Analytics
Overview
Discover
Dashboard
Canvas
Maps
Machine Learning
Visualize Library
Enterprise Search
Dashboards
Create dashboard
Search...
Tags
Title
Description
Tags
Actions
D1
HAPROXY
[Elastic Agent] Agent metrics
Elastic Agent metrics dashboard
[Logs System] New users and groups
New users and groups dashboard for the System integration in Logs
[Logs System] SSH login attempts
SSH dashboard for the System integration in Logs

OSUG

Kibana

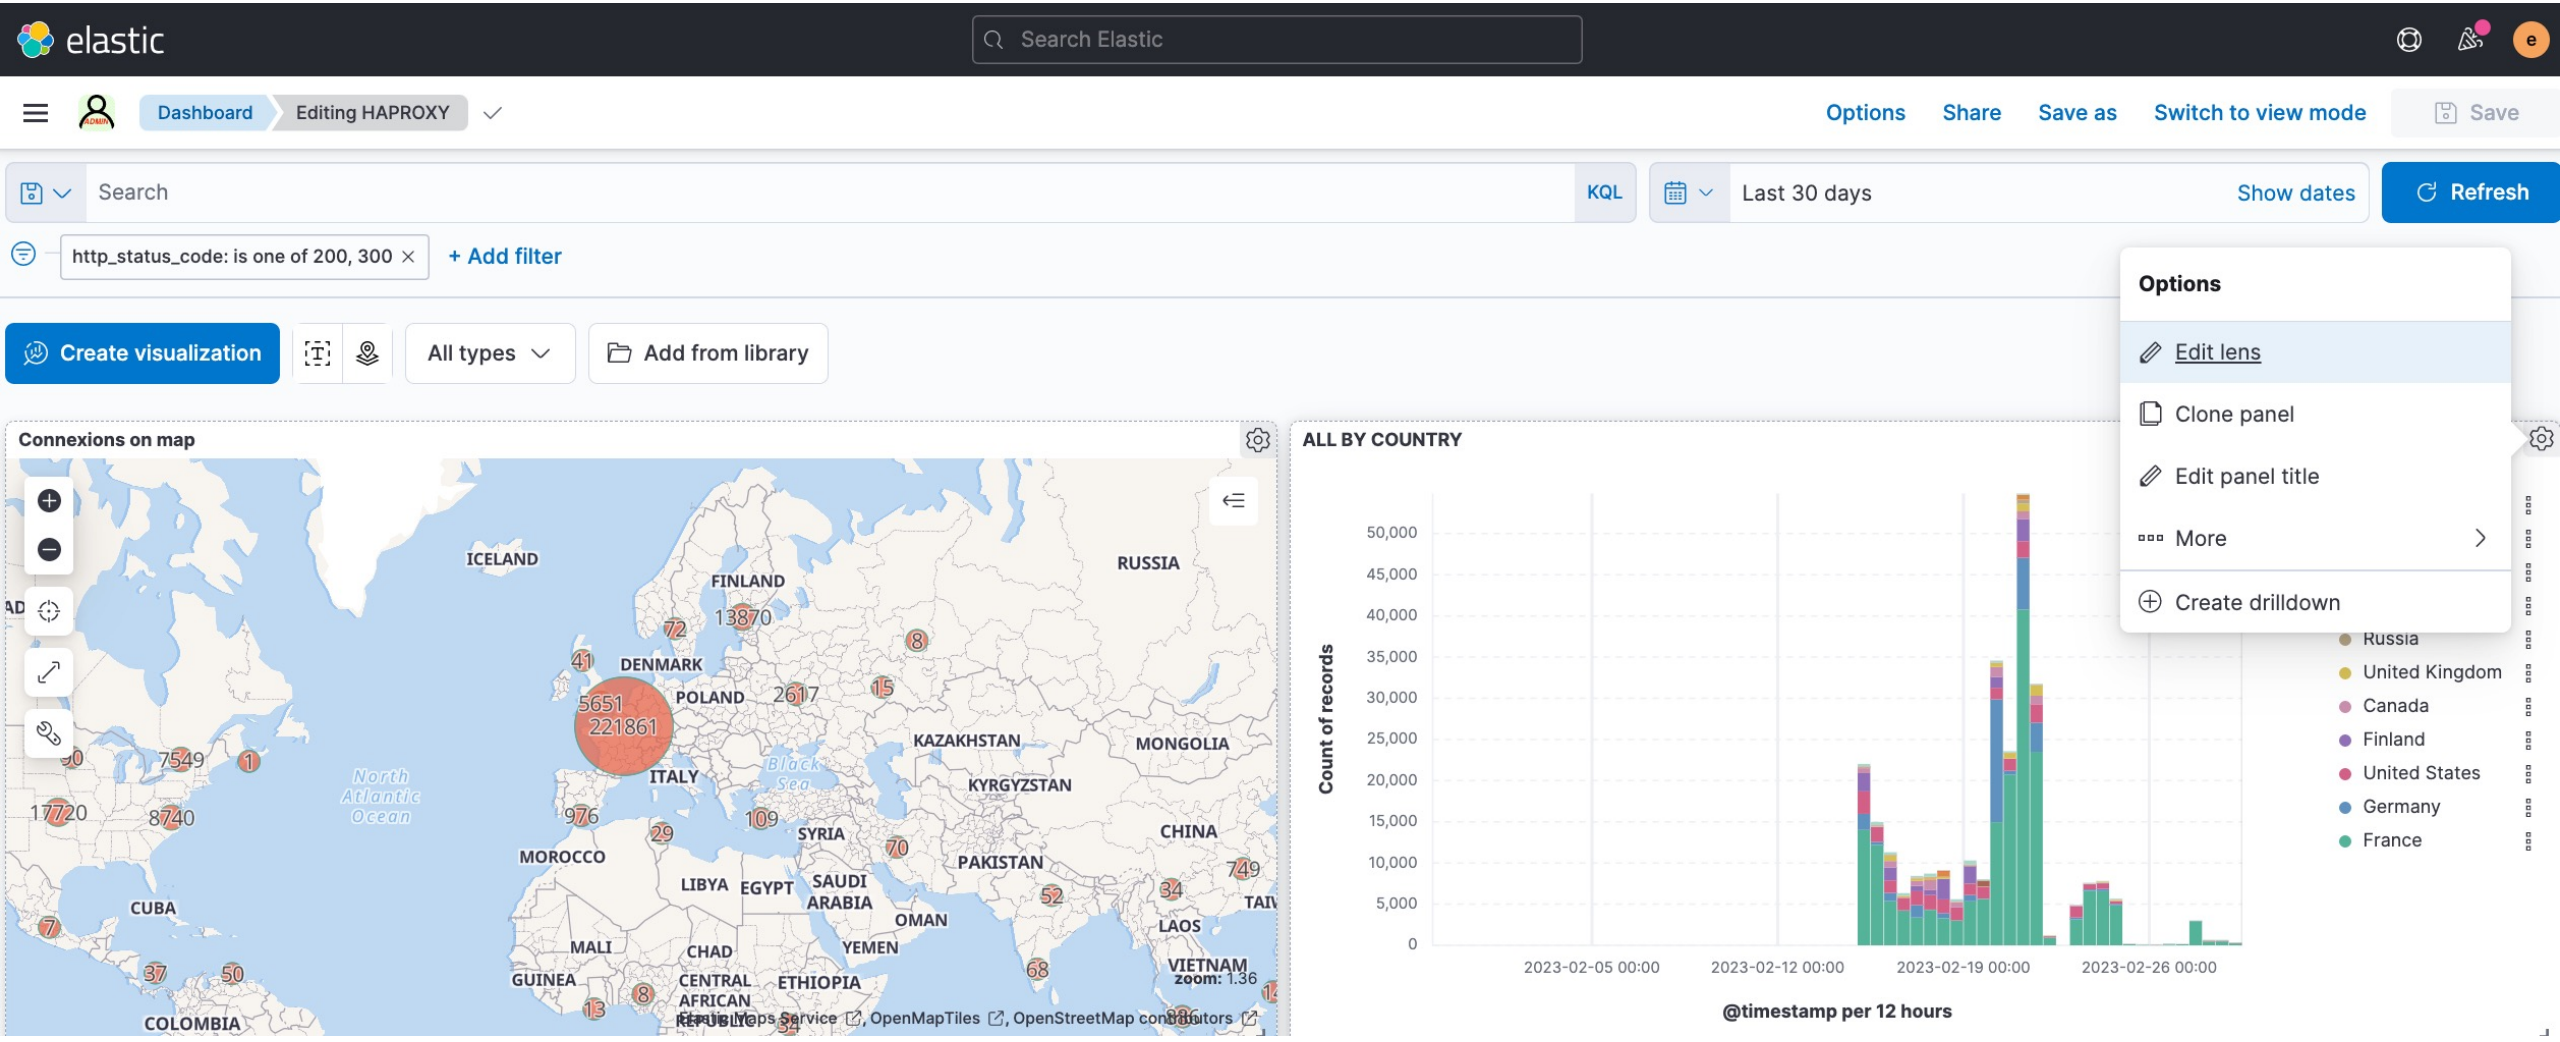
elastic
Search Elastic
Dashboard
Editing HAPROXY
Options
Share
Save as
Switch to view mode
Save
Search
KQL
Last 30 days
Show dates
Refresh
http_status_code: is one of 200, 300
+ Add filter
Create visualization
All types
Add from library
Options
Edit lens
Clone panel
Edit panel title
More
Create drilldown
Connexions on map
ICELAND
FINLAND
RUSSIA
DENMARK
POLAND
KAZAKHSTAN
MONGOLIA
ITALY
KYRGYZSTAN
CHINA
MOROCCO
SYRIA
PAKISTAN
LIBYA
EGYPT
SAUDI ARABIA
OMAN
YEMEN
CUBA
MALI
CHAD
LAOS
VIETNAM
GUINEA
CENTRAL AFRICAN REPUBLIC
ETHIOPIA
COLOMBIA
North Atlantic Ocean
Black Sea
zoom: 1.36
ALL BY COUNTRY
Count of records
50,000
45,000
40,000
35,000
30,000
25,000
20,000
15,000
10,000
5,000
0
2023-02-05 00:00
2023-02-12 00:00
2023-02-19 00:00
2023-02-26 00:00
@timestamp per 12 hours
Russia
United Kingdom
Canada
Finland
United States
Germany
France

OSUG

Kibana

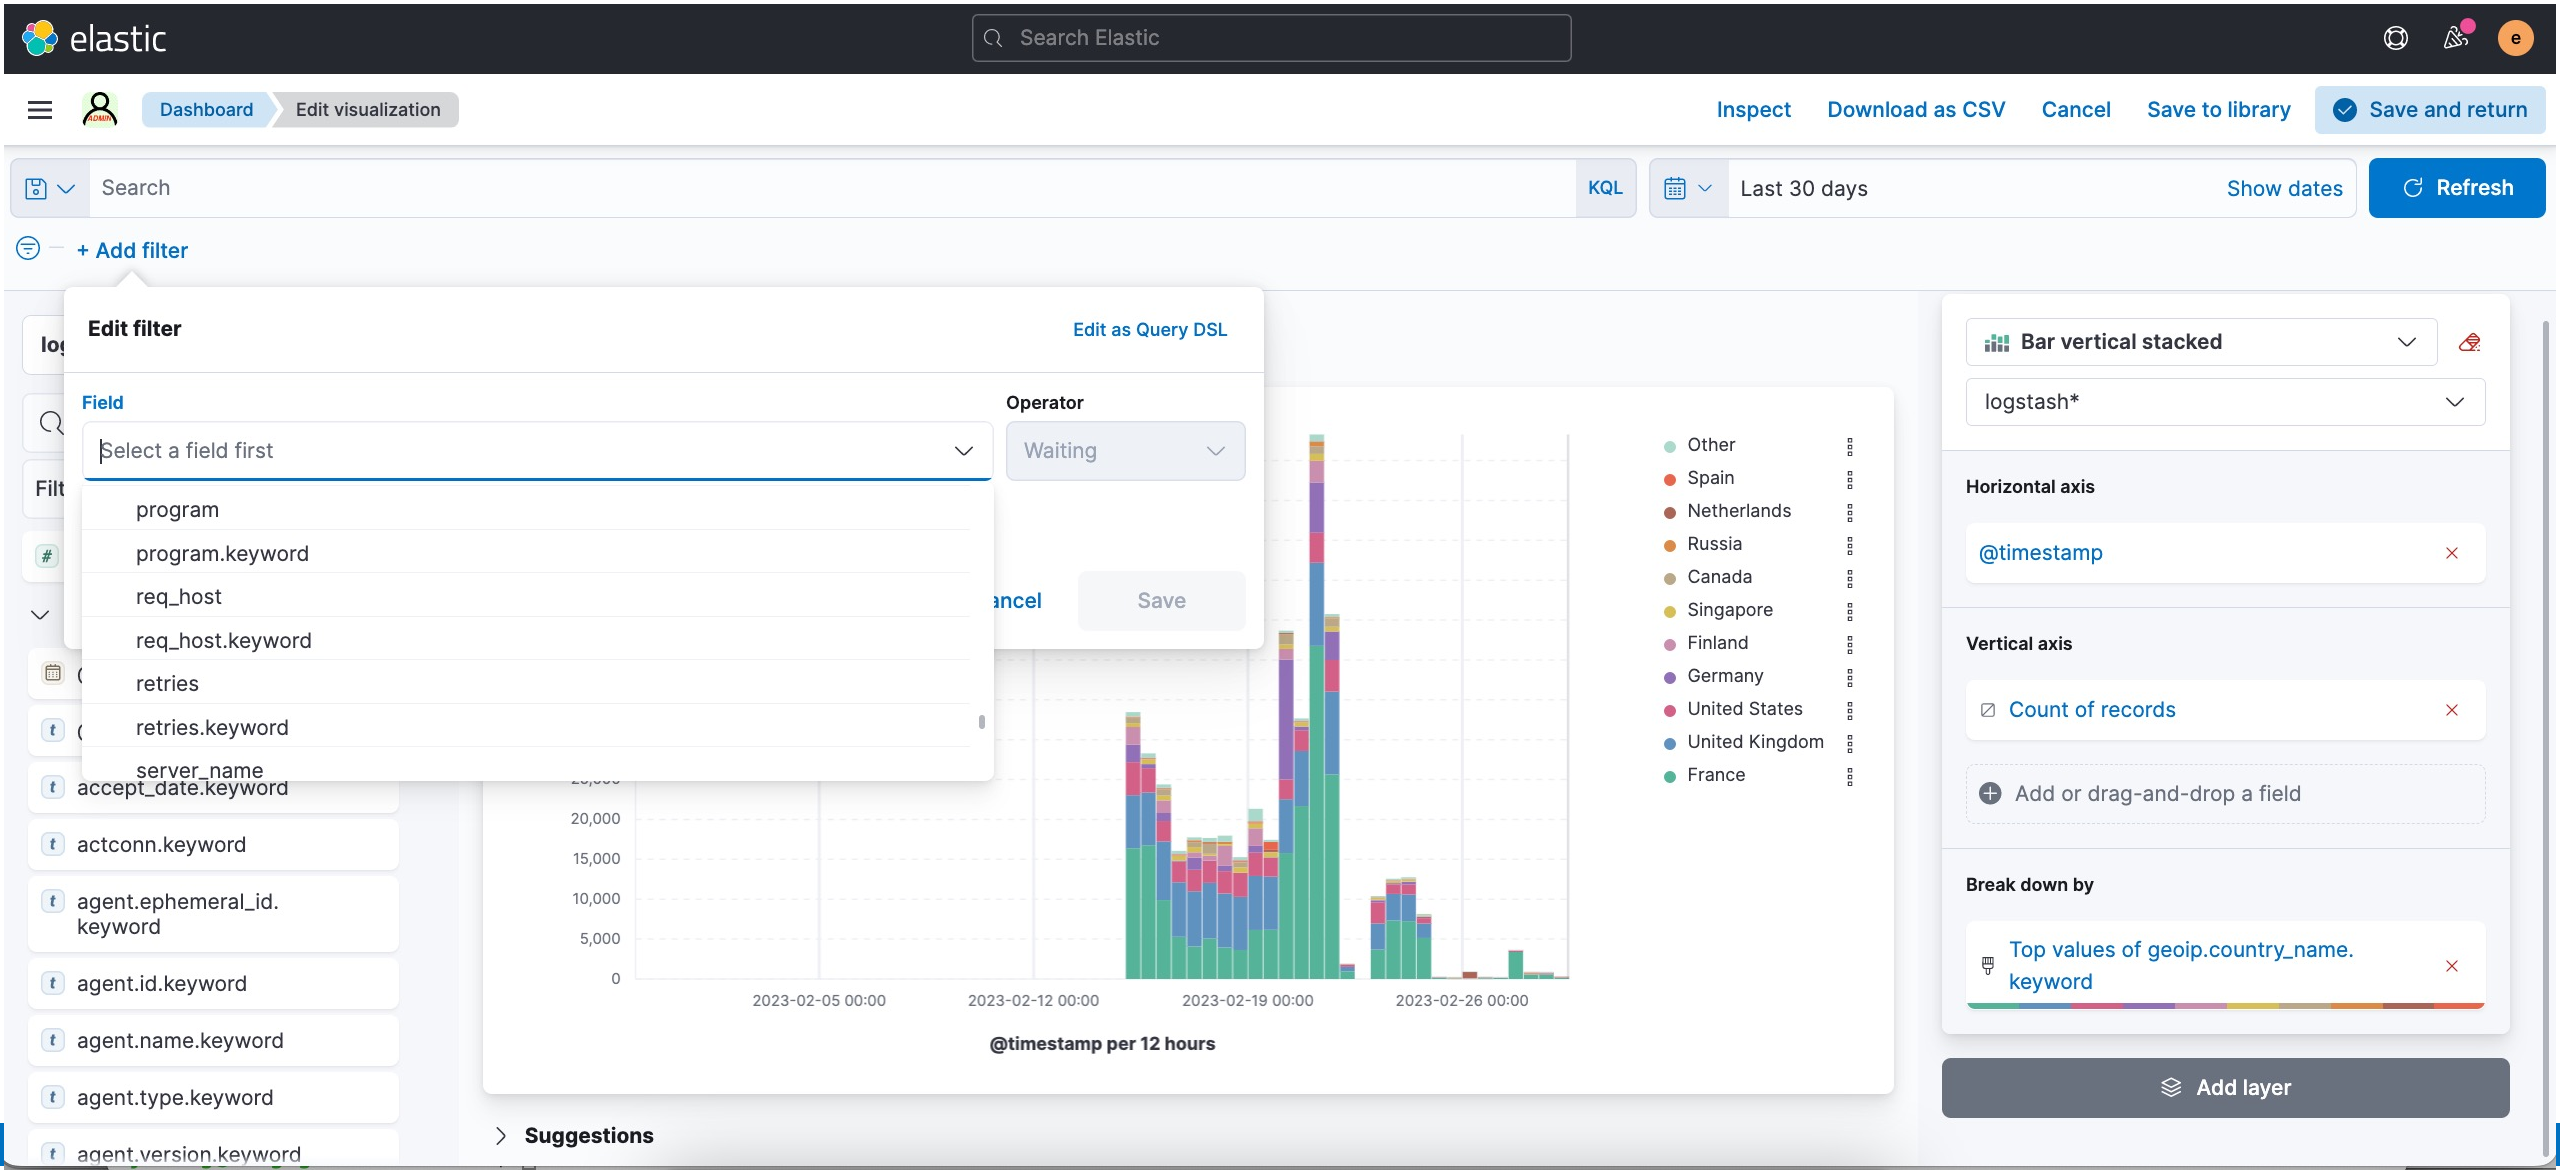
elastic
Search Elastic
Dashboard
Edit visualization
Inspect
Download as CSV
Cancel
Save to library
Save and return
Search
KQL
Last 30 days
Show dates
Refresh
+ Add filter
Edit filter
Edit as Query DSL
Field
Select a field first
Operator
Waiting
program
program.keyword
req_host
req_host.keyword
retries
retries.keyword
server_name
Cancel
Save
accept_date.keyword
actconn.keyword
agent.ephemeral_id.keyword
agent.id.keyword
agent.name.keyword
agent.type.keyword
agent.version.keyword
20,000
15,000
10,000
5,000
0
2023-02-05 00:00
2023-02-12 00:00
2023-02-19 00:00
2023-02-26 00:00
@timestamp per 12 hours
Other
Spain
Netherlands
Russia
Canada
Singapore
Finland
Germany
United States
United Kingdom
France
Bar vertical stacked
logstash*
Horizontal axis
@timestamp
Vertical axis
Count of records
Add or drag-and-drop a field
Break down by
Top values of geoip.country_name.keyword
Add layer
Suggestions

OSUG

Kibana

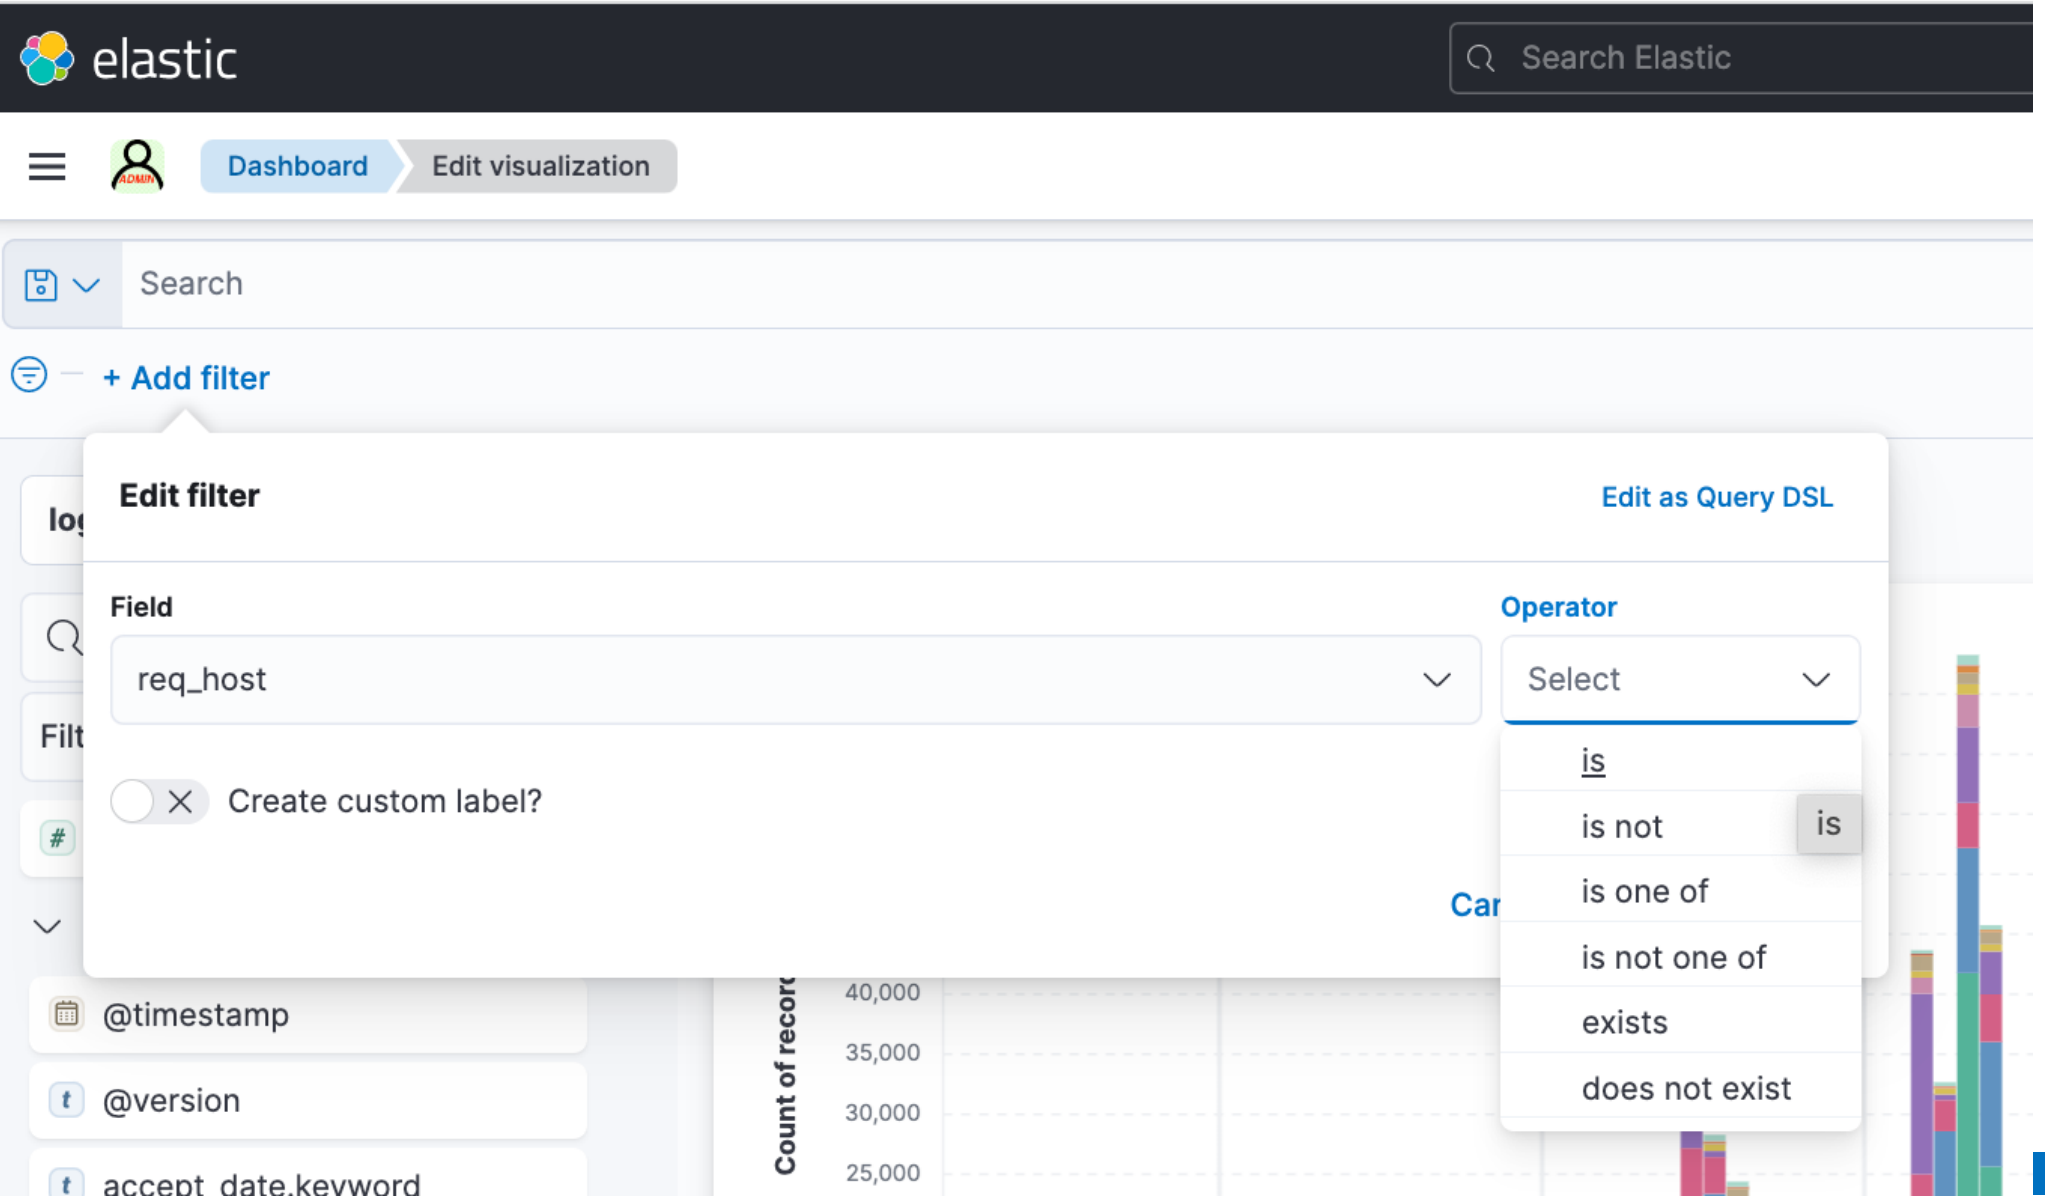
elastic
Search Elastic
Dashboard
Edit visualization
Search
+ Add filter
Edit filter
Edit as Query DSL
Field
req_host
Operator
Select
is
is not
is one of
is not one of
exists
does not exist
Create custom label?
@timestamp
@version
accept_date.keyword
Count of record
40,000
35,000
30,000
25,000

OSUG

Kibana

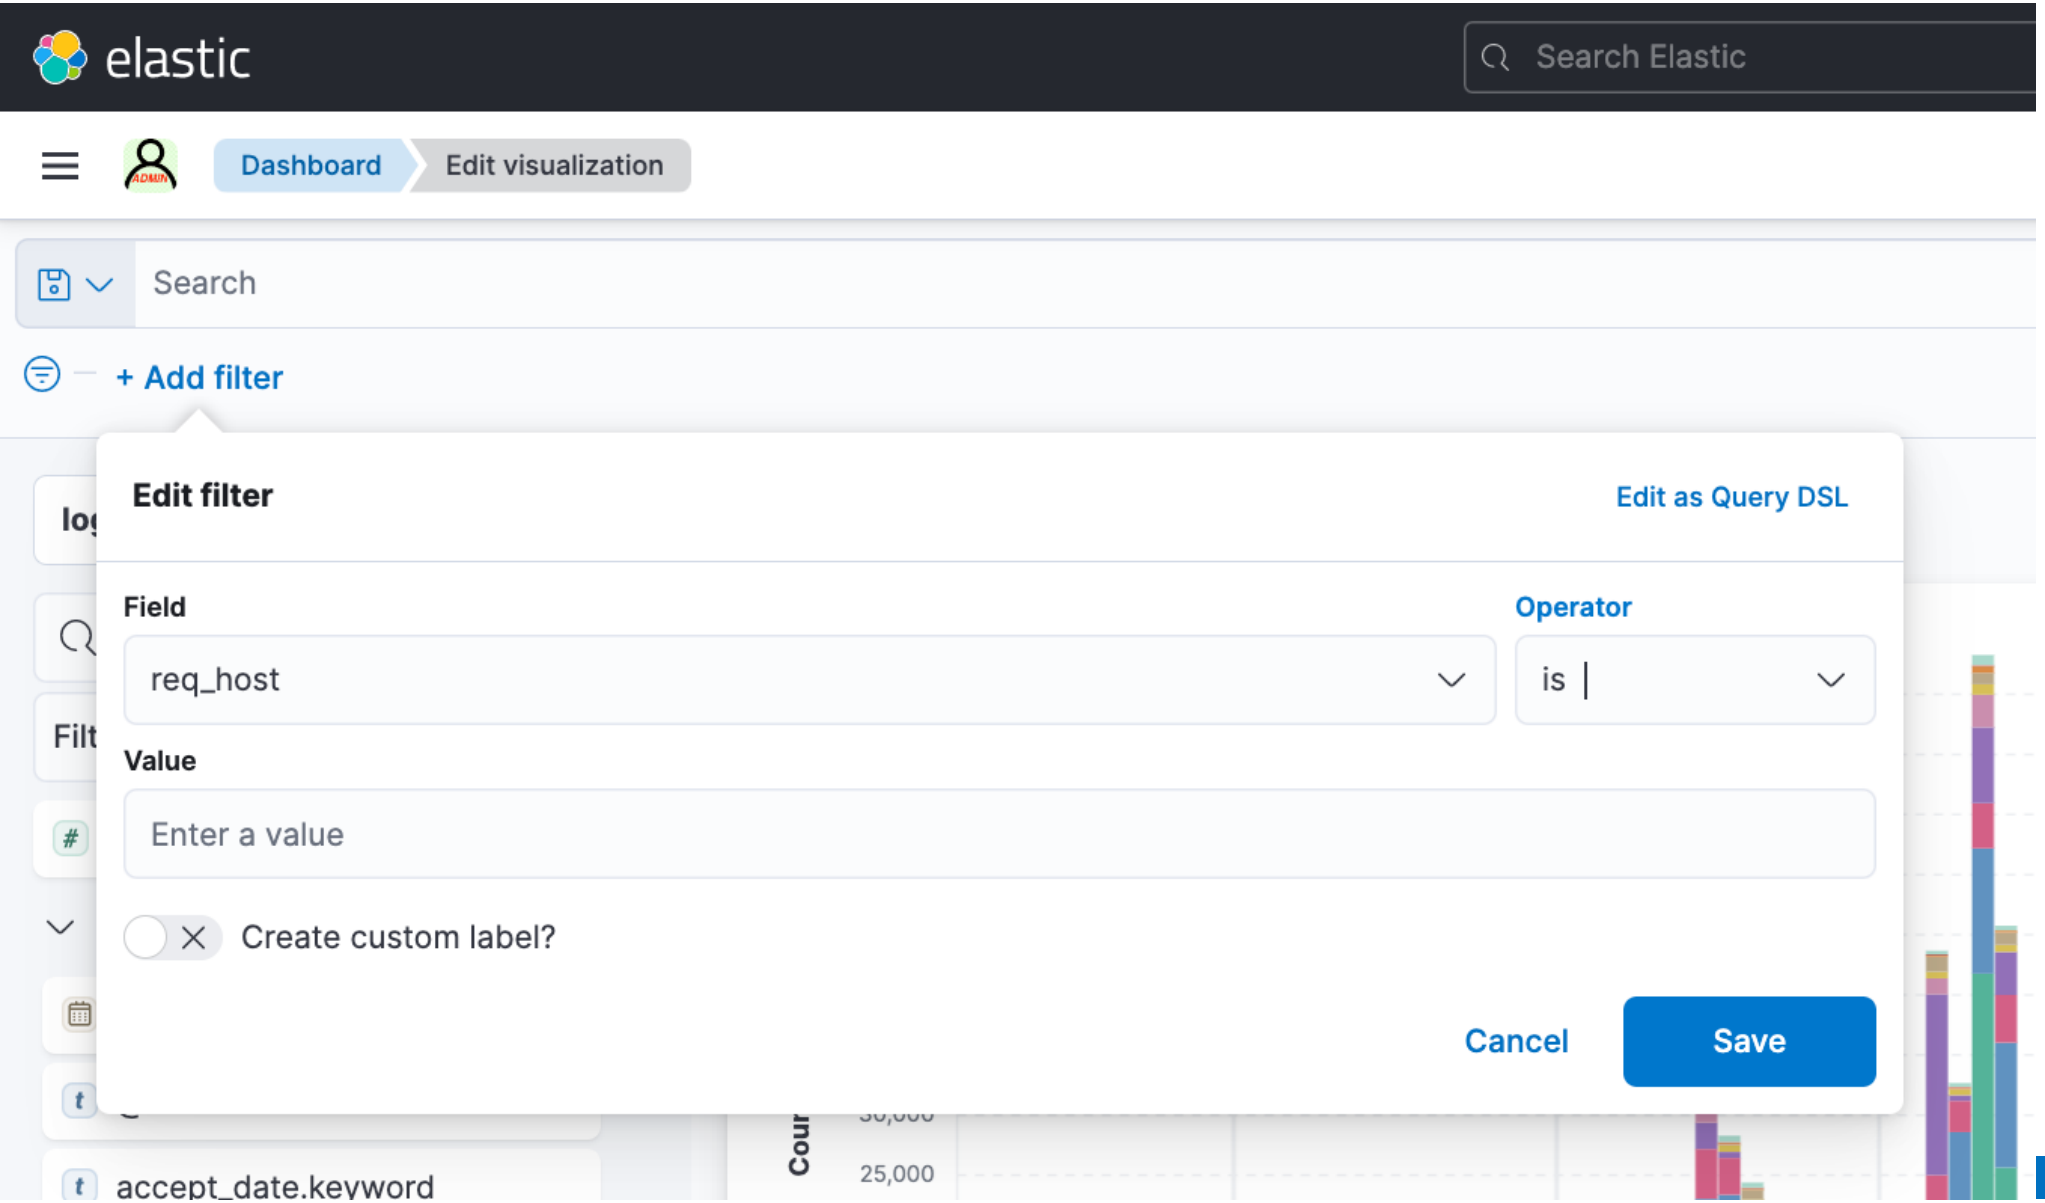
elastic
Search Elastic
Dashboard
Edit visualization
Search
+ Add filter
Edit filter
Edit as Query DSL
Field
req_host
Operator
is
Value
Enter a value
Create custom label?
Cancel
Save
25,000
accept_date.keyword

OSUG

Kibana

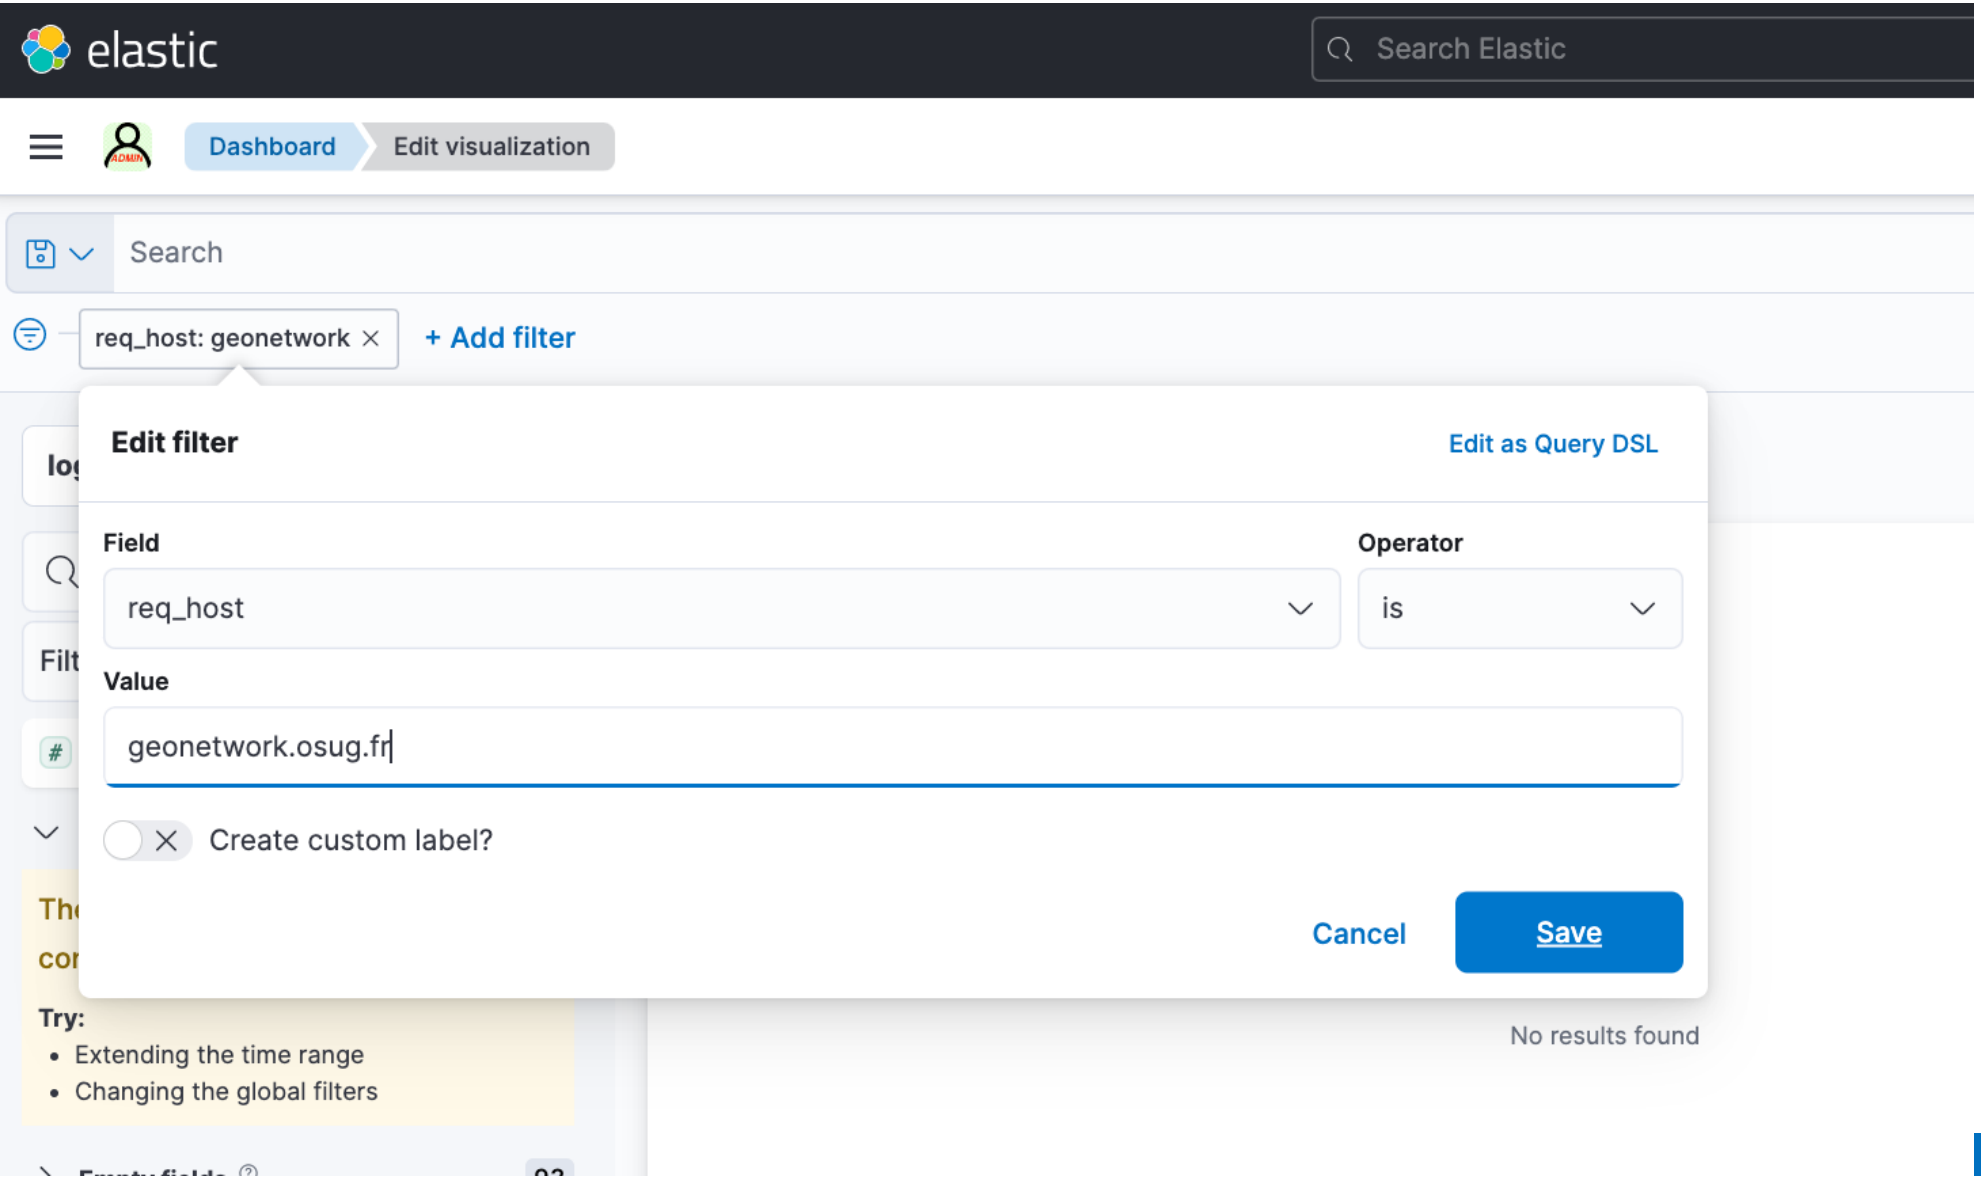
elastic
Search Elastic
Dashboard
Edit visualization
Search
req_host: geonetwork
+ Add filter
Edit filter
Edit as Query DSL
Field
req_host
Operator
is
Value
geonetwork.osug.fr
Create custom label?
Cancel
Save
No results found
Try:
Extending the time range
Changing the global filters

OSUG

Kibana

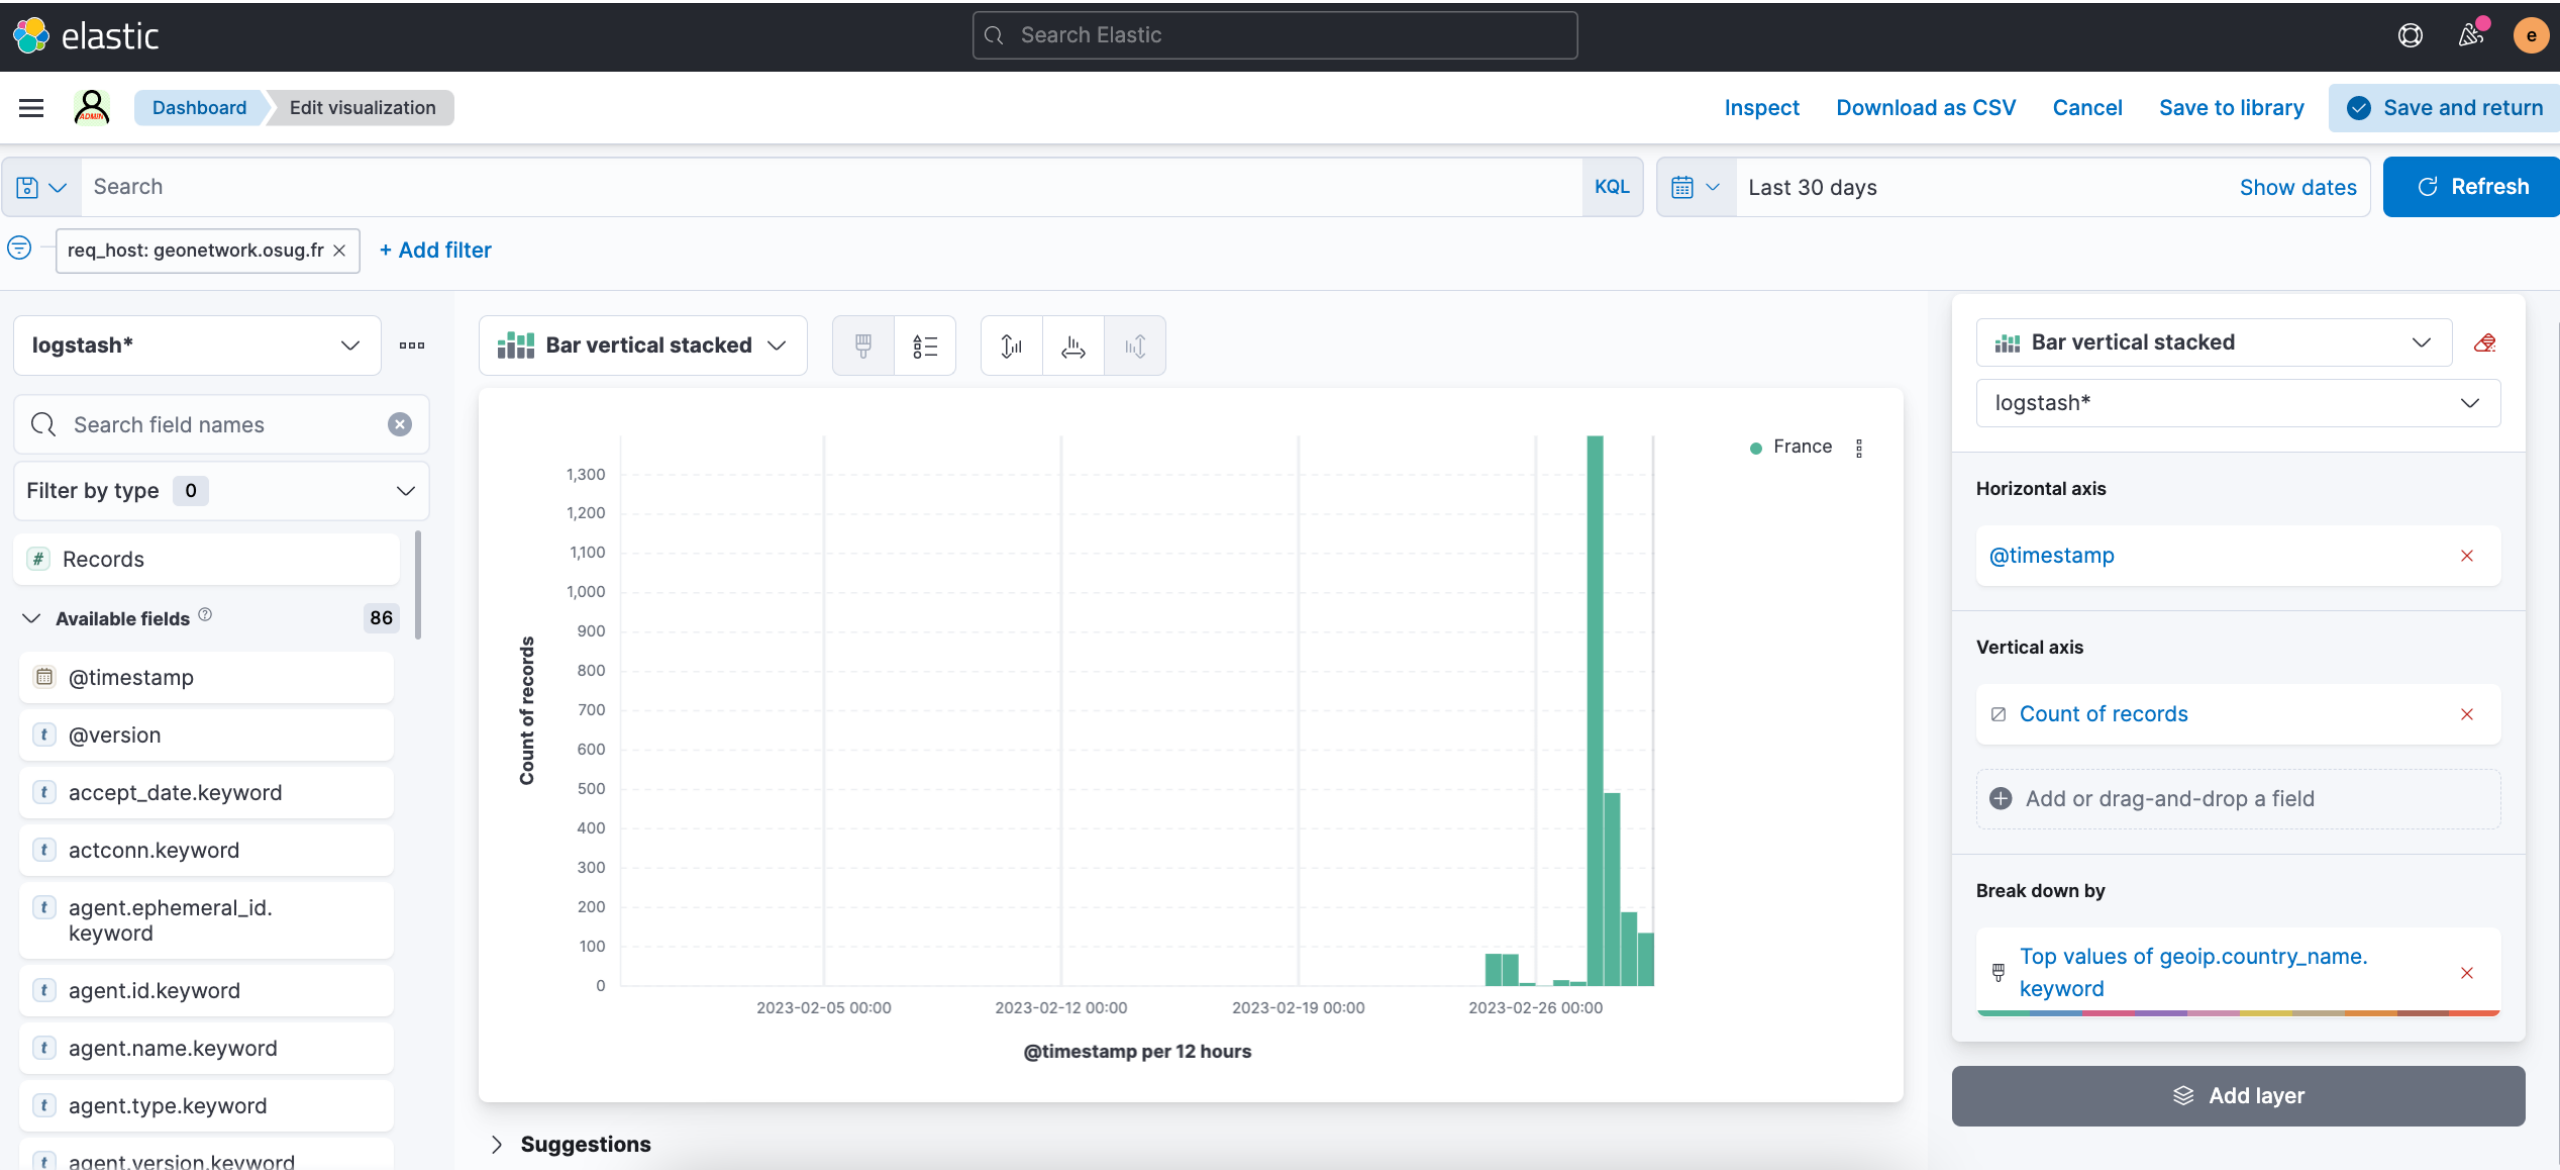
elastic
Search Elastic
Dashboard
Edit visualization
Inspect
Download as CSV
Cancel
Save to library
Save and return
Search
KQL
Last 30 days
Show dates
Refresh
req_host: geonetwork.osug.fr
+ Add filter
logstash*
Search field names
Filter by type 0
Records
Available fields 86
@timestamp
@version
accept_date.keyword
actconn.keyword
agent.ephemeral_id.keyword
agent.id.keyword
agent.name.keyword
agent.type.keyword
agent.version.keyword
Bar vertical stacked
France
Count of records
1,300
1,200
1,100
1,000
900
800
700
600
500
400
300
200
100
0
2023-02-05 00:00
2023-02-12 00:00
2023-02-19 00:00
2023-02-26 00:00
@timestamp per 12 hours
Suggestions
Bar vertical stacked
logstash*
Horizontal axis
@timestamp
Vertical axis
Count of records
Add or drag-and-drop a field
Break down by
Top values of geoip.country_name.keyword
Add layer

TODO

TODO : VM elasticsearch-kibana dédiée

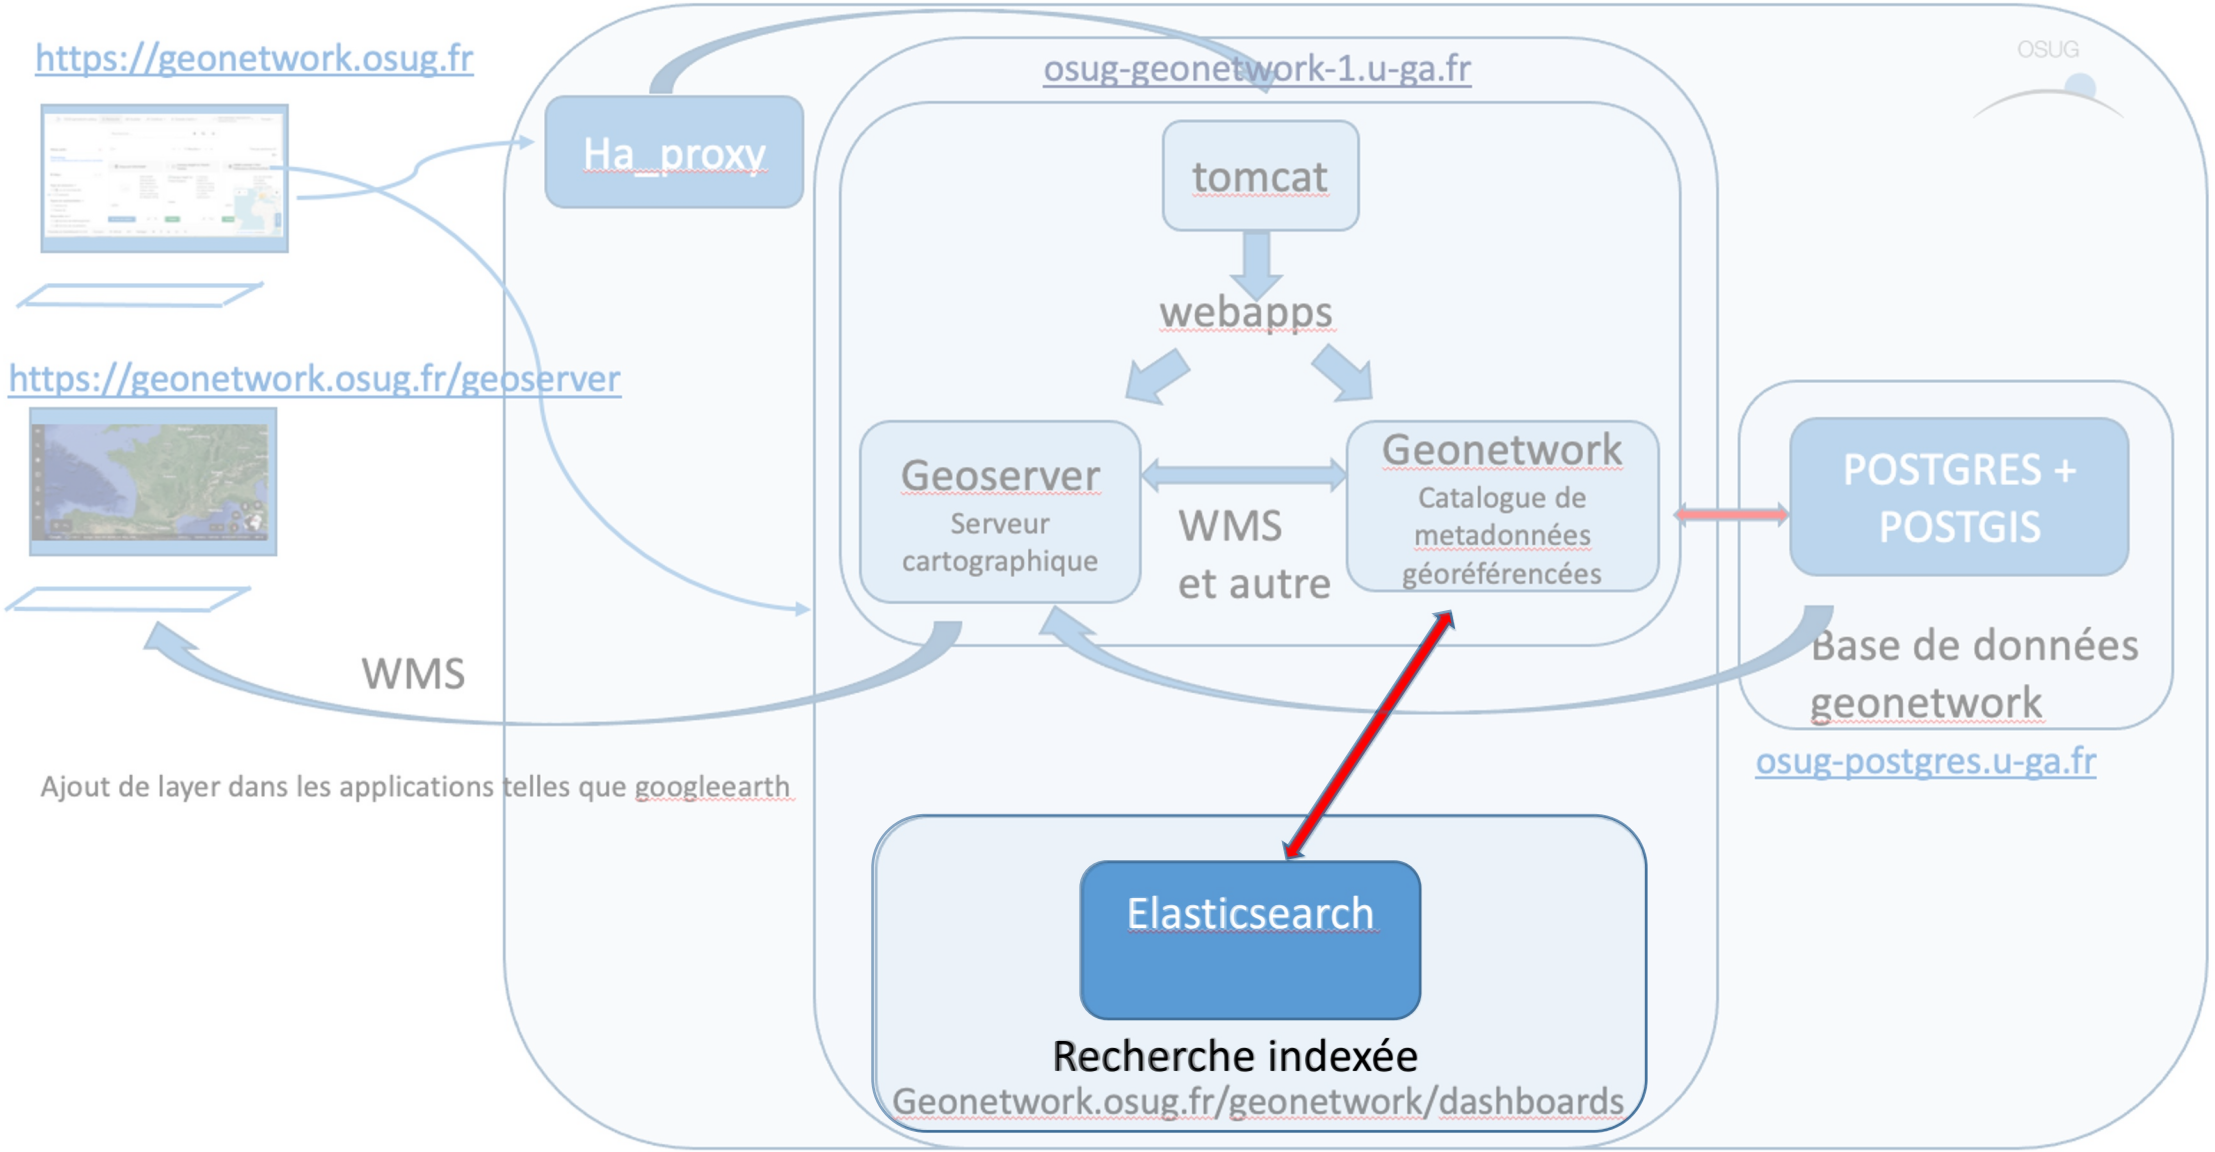

OSUG

TODO : logs des VM hors haproxy

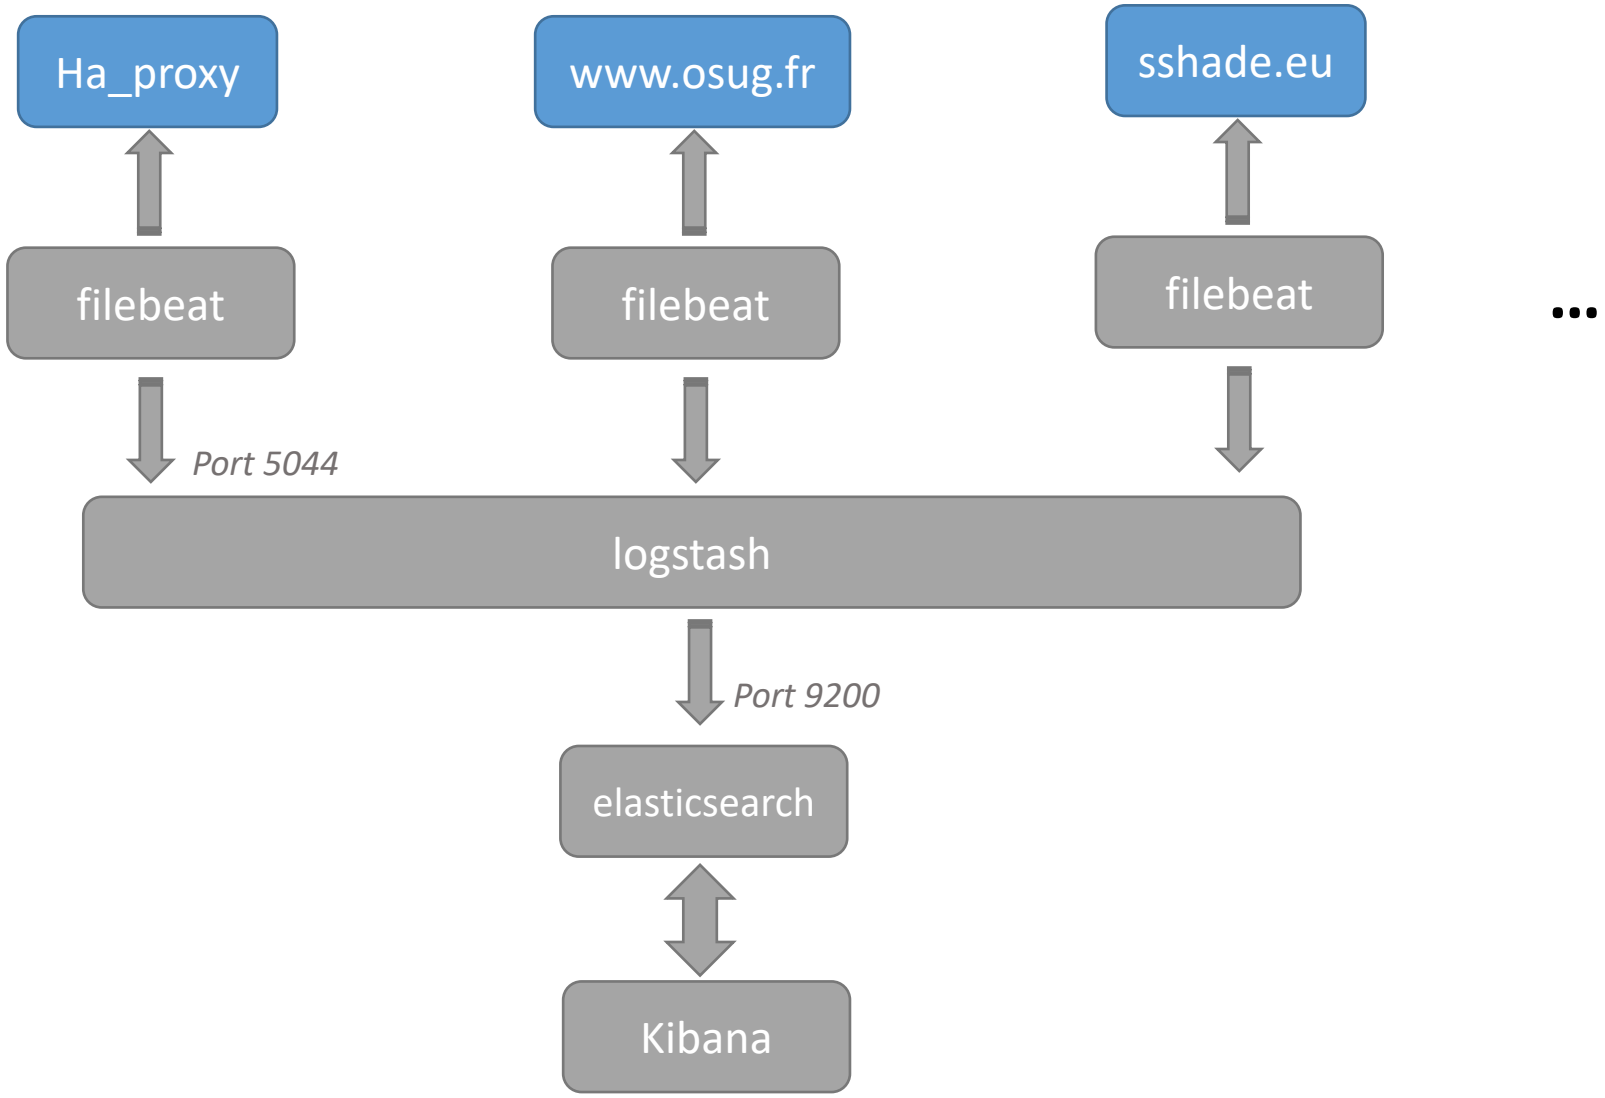

OSUG

TODO : kibana, cloisonnement des rôles

- Définir et créer les rôles
  - Droits sur les index : create, delete, read, write etc.
- Créer des utilisateurs et leur affecter des rôles
- Définir et créer les espaces utilisateurs
  - Pour un espace utilisateur : filtrage des menus disponibles

The end

merci de votre attention