Elasticsearch - A hands-on introduction

Alexander Reelsen

Community Advocate

alex@elastic.co | @spinscale

Agenda

- What is the Elastic Stack
- Elasticsearch introduction
- Elasticsearch practical demo
- Integrating Elasticsearch into your application

Product Overview

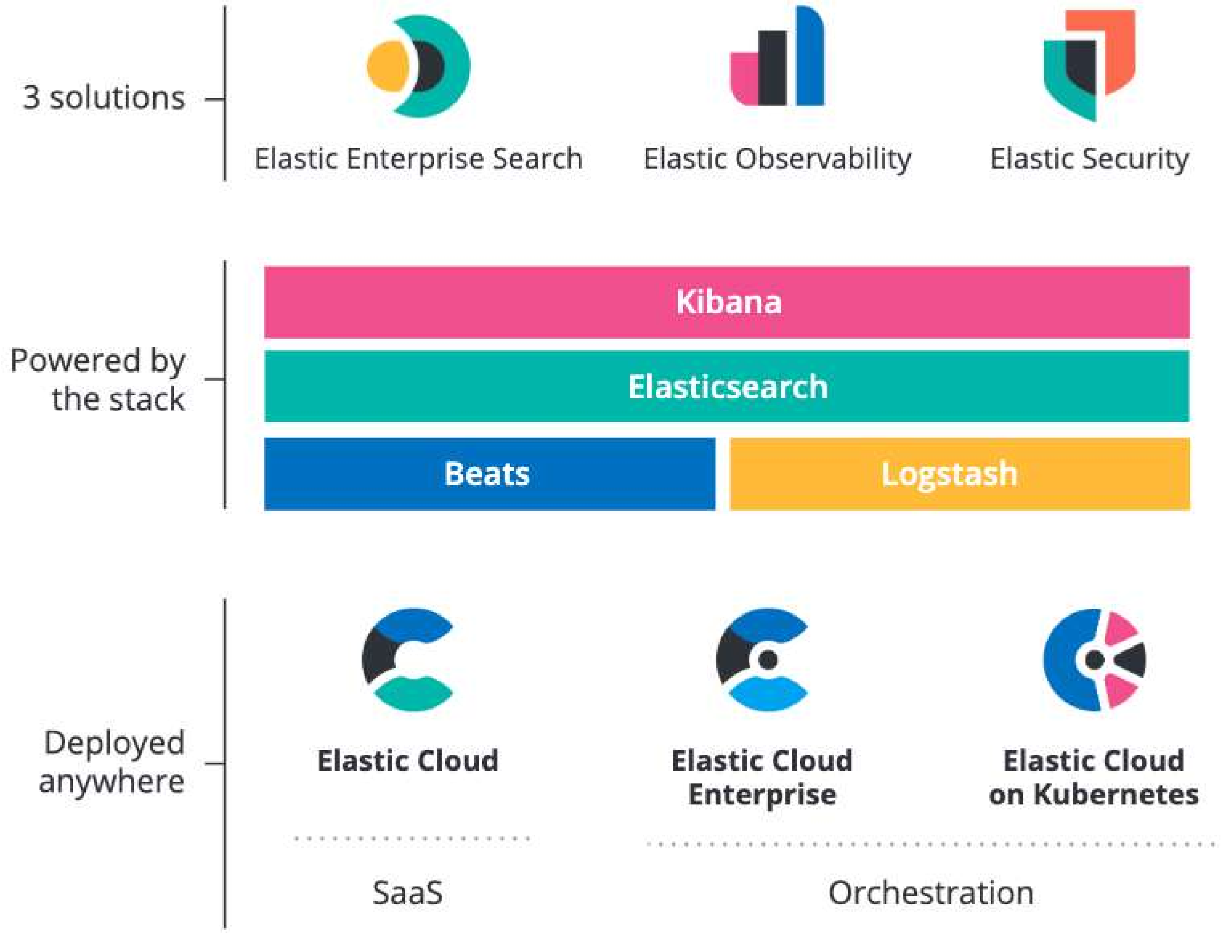

3 solutions

Elastic Enterprise Search

Elastic Observability

Elastic Security

Powered by the stack

Kibana

Elasticsearch

Beats

Logstash

Deployed anywhere

Elastic Cloud

Elastic Cloud Enterprise

Elastic Cloud on Kubernetes

SaaS

Orchestration

elastic

Solutions

on top of the Elastic Stack

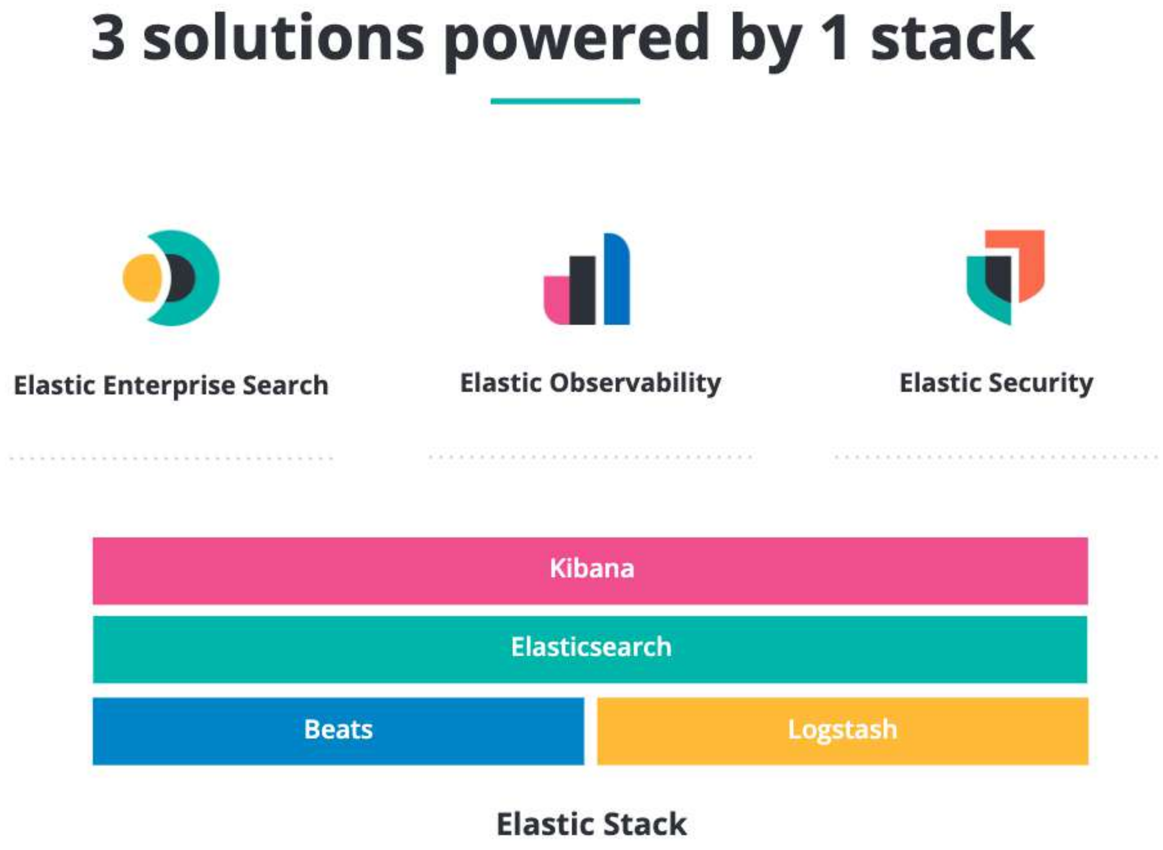

3 solutions powered by 1 stack

Elastic Enterprise Search

Elastic Observability

Elastic Security

Kibana

Elasticsearch

Beats

Logstash

Elastic Stack

Elastic Stack

building & lego blocks

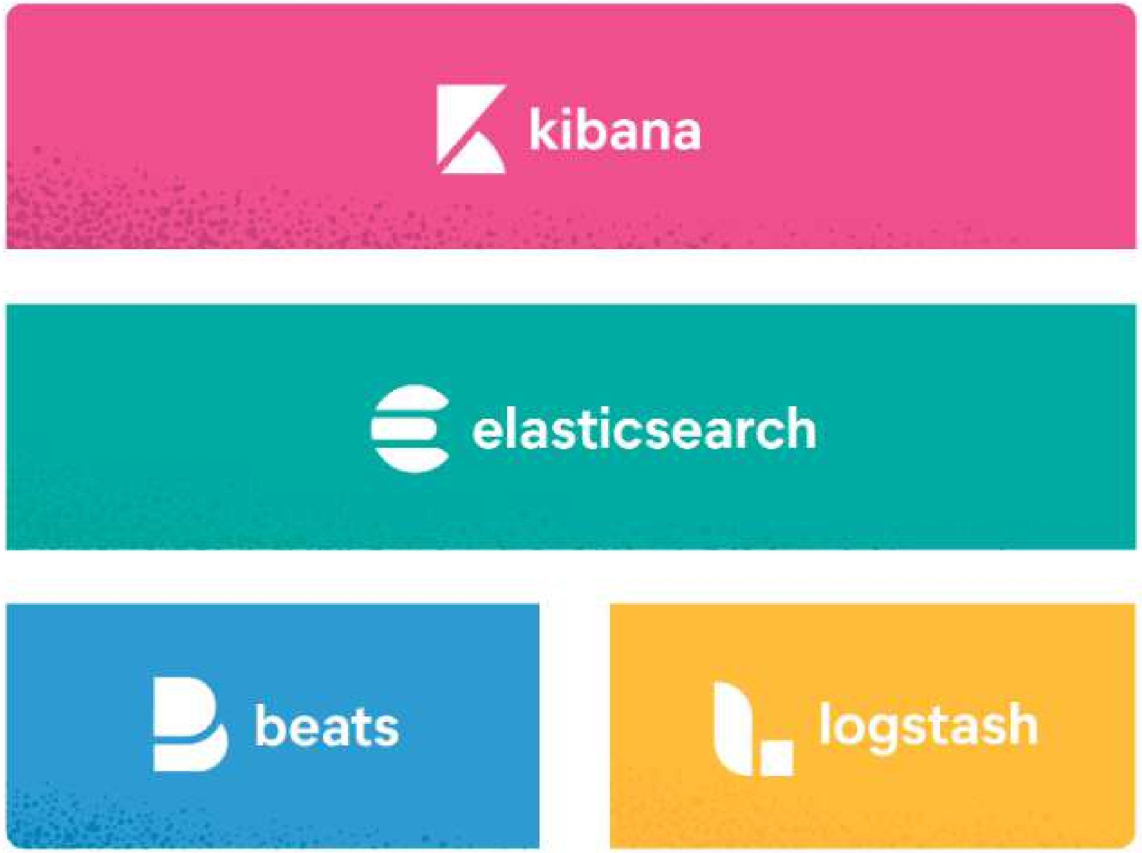

kibana

elasticsearch

beats

logstash

Deployment options

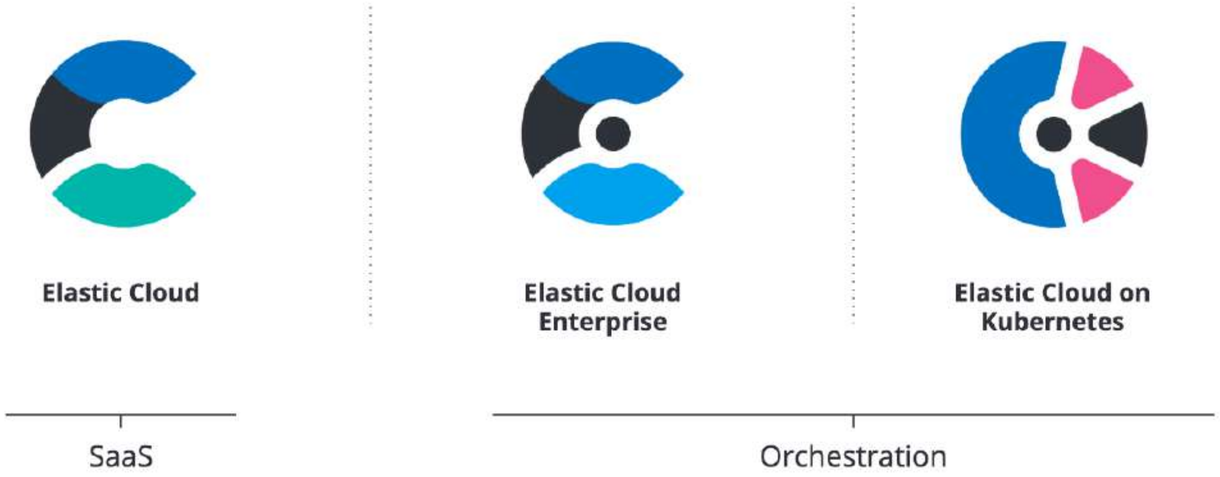

Elastic Cloud

Elastic Cloud Enterprise

Elastic Cloud on Kubernetes

SaaS

Orchestration

Licensing

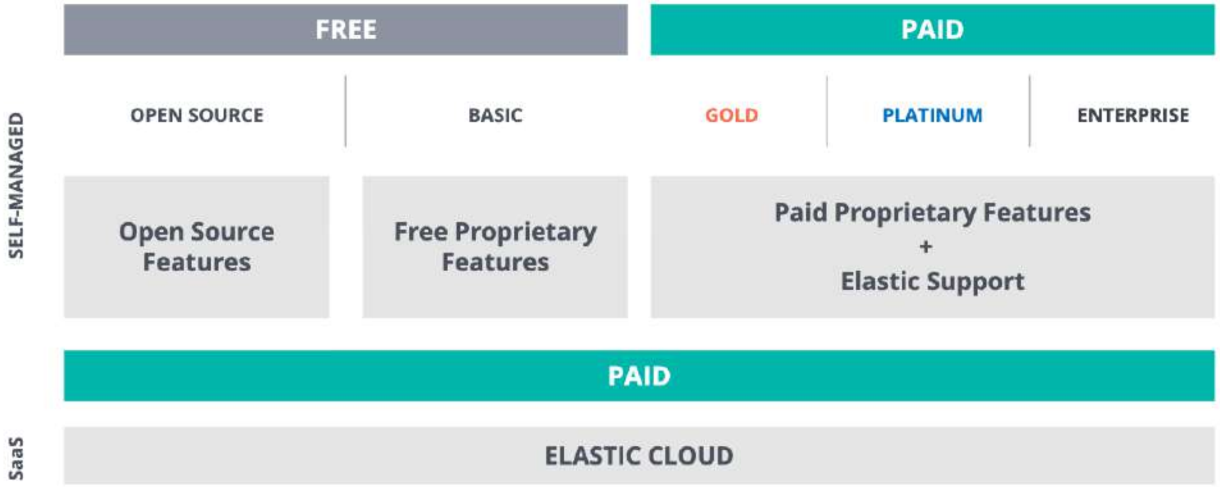

SELF-MANAGED

FREE		PAID		
OPEN SOURCE	BASIC	GOLD	PLATINUM	ENTERPRISE
Open Source Features	Free Proprietary Features	Paid Proprietary Features + Elastic Support		

SaaS

PAID

ELASTIC CLOUD

Elastic Stack

building & lego blocks

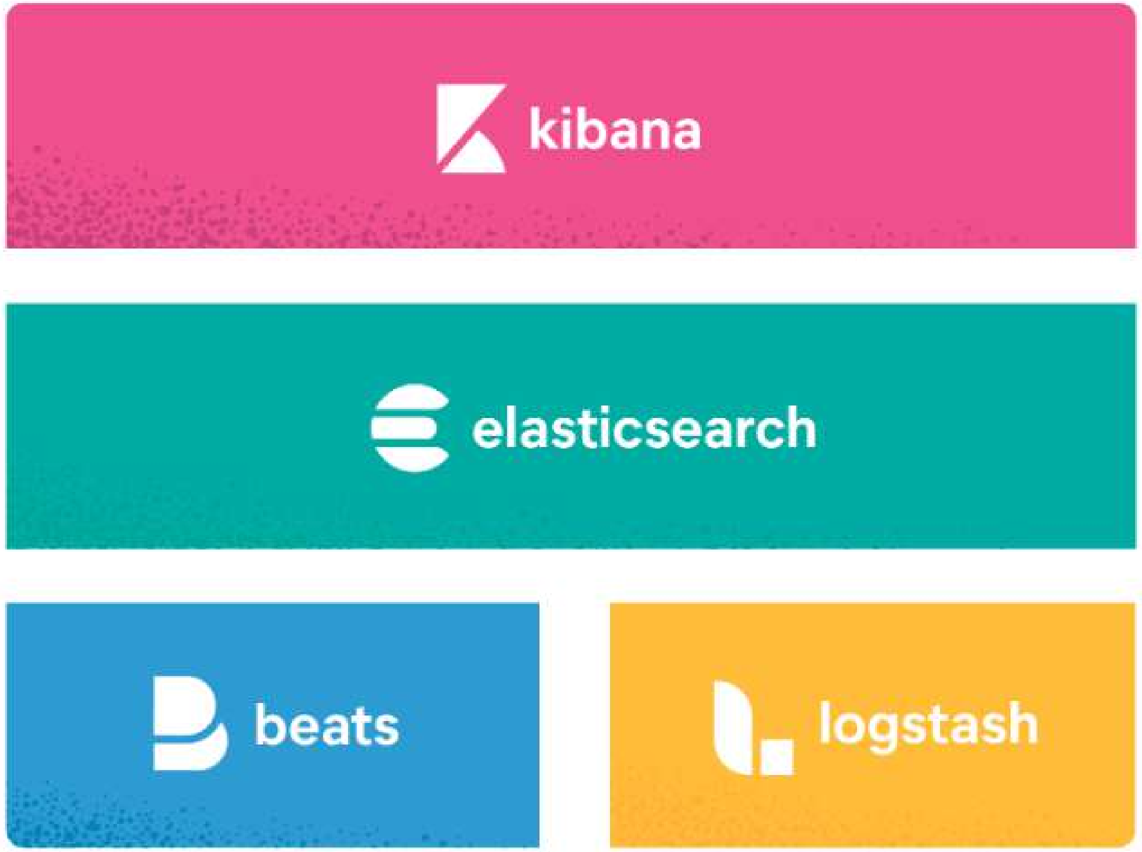

kibana

elasticsearch

beats

logstash

Elasticsearch in 10 seconds

- Search Engine (FTS, Analytics, Geo), near real-time
- Distributed, scalable, highly available, resilient
- Interface: HTTP & JSON
- Heart of the Elastic Stack (Kibana, Logstash, Beats)

Installation & Start

```
# https://www.elastic.co/downloads/elasticsearch
wget https://artifacts.elastic.co/downloads/elasticsearch/elasticsearch-7.7.0-darwin-x86_64.tar.gz
# wget https://artifacts.elastic.co/downloads/elasticsearch/elasticsearch-7.7.0-linux-x86_64.tar.gz
# wget https://artifacts.elastic.co/downloads/elasticsearch/elasticsearch-7.7.0-windows-x86_64.zip

tar zxf elasticsearch-7.7.0-darwin-x86_64.tar.gz
cd elasticsearch-7.7.0

./bin/elasticsearch
```

```
wget https://artifacts.elastic.co/downloads/kibana/kibana-7.7.0-darwin-x86_64.tar.gz
# wget https://artifacts.elastic.co/downloads/kibana/kibana-7.7.0-linux-x86_64.tar.gz
# wget https://artifacts.elastic.co/downloads/kibana/kibana-7.7.0-windows-x86_64.zip

tar zxf kibana-7.7.0-darwin-x86_64.tar.gz
cd kibana-7.7.0
./bin/kibana
```

Point your browser to http://localhost:5601/

Click Dev-Tools

Samples in Kibana

Samples in Github

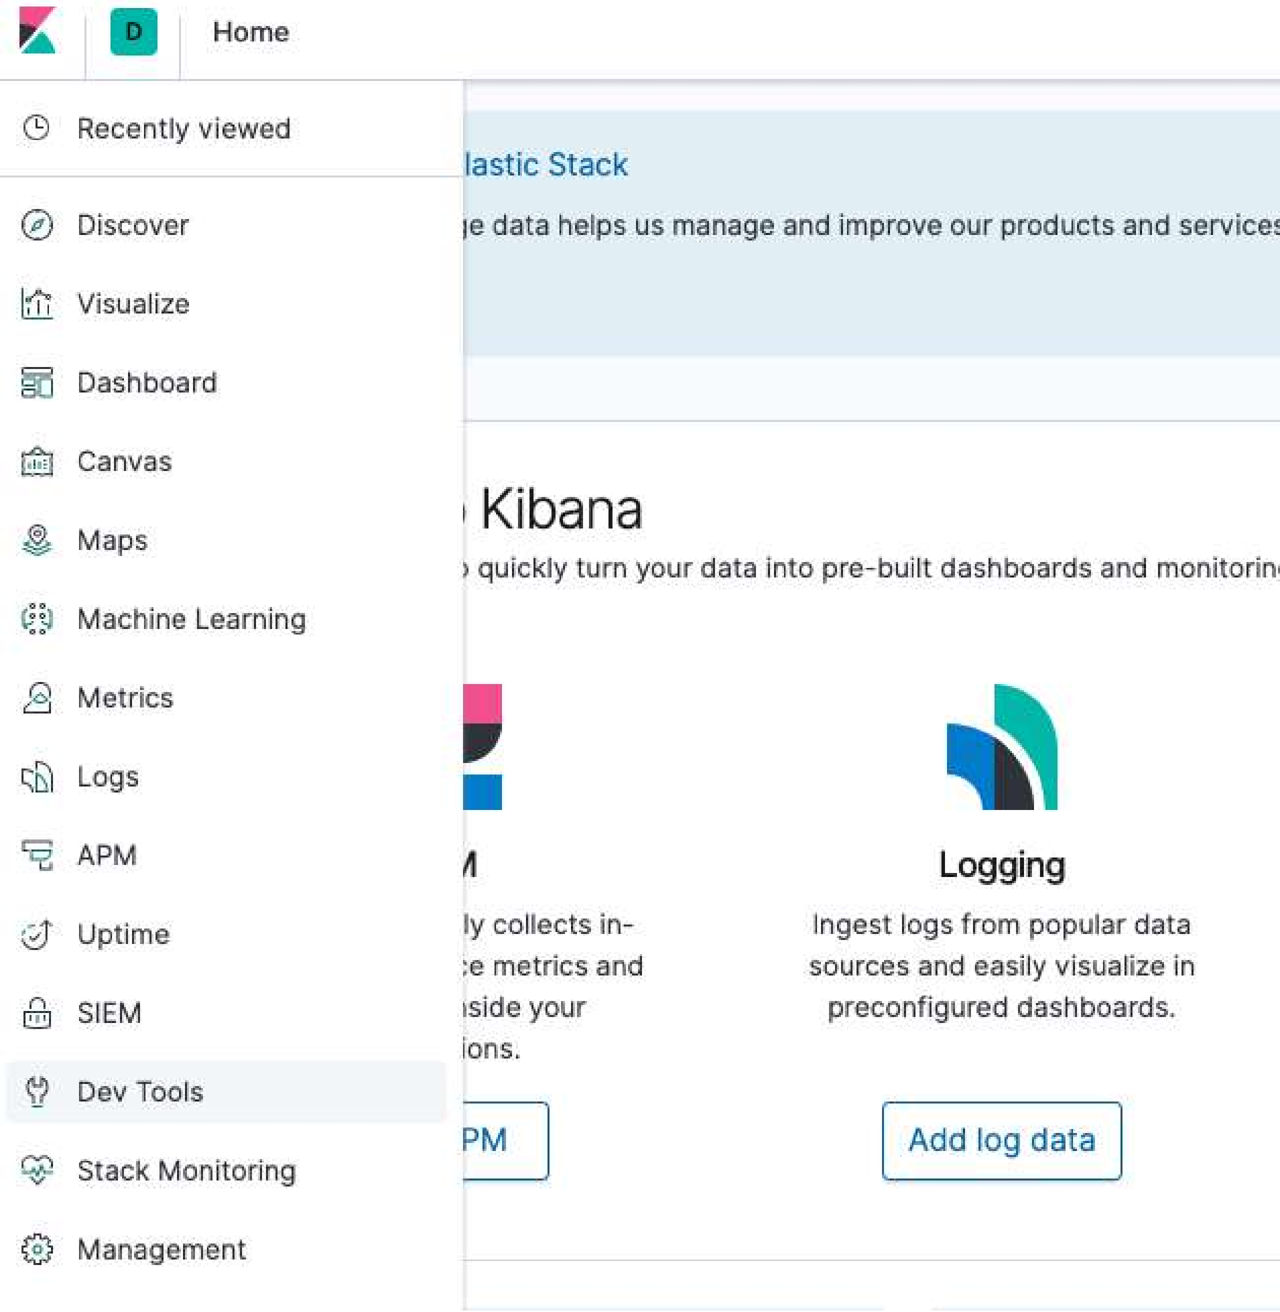

lastic Stack

ge data helps us manage and improve our products and services

Kibana

quickly turn your data into pre-built dashboards and monitorin

M

ly collects in-
ce metrics and
side your
ions.

PM

Logging

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

Add log data

elastic

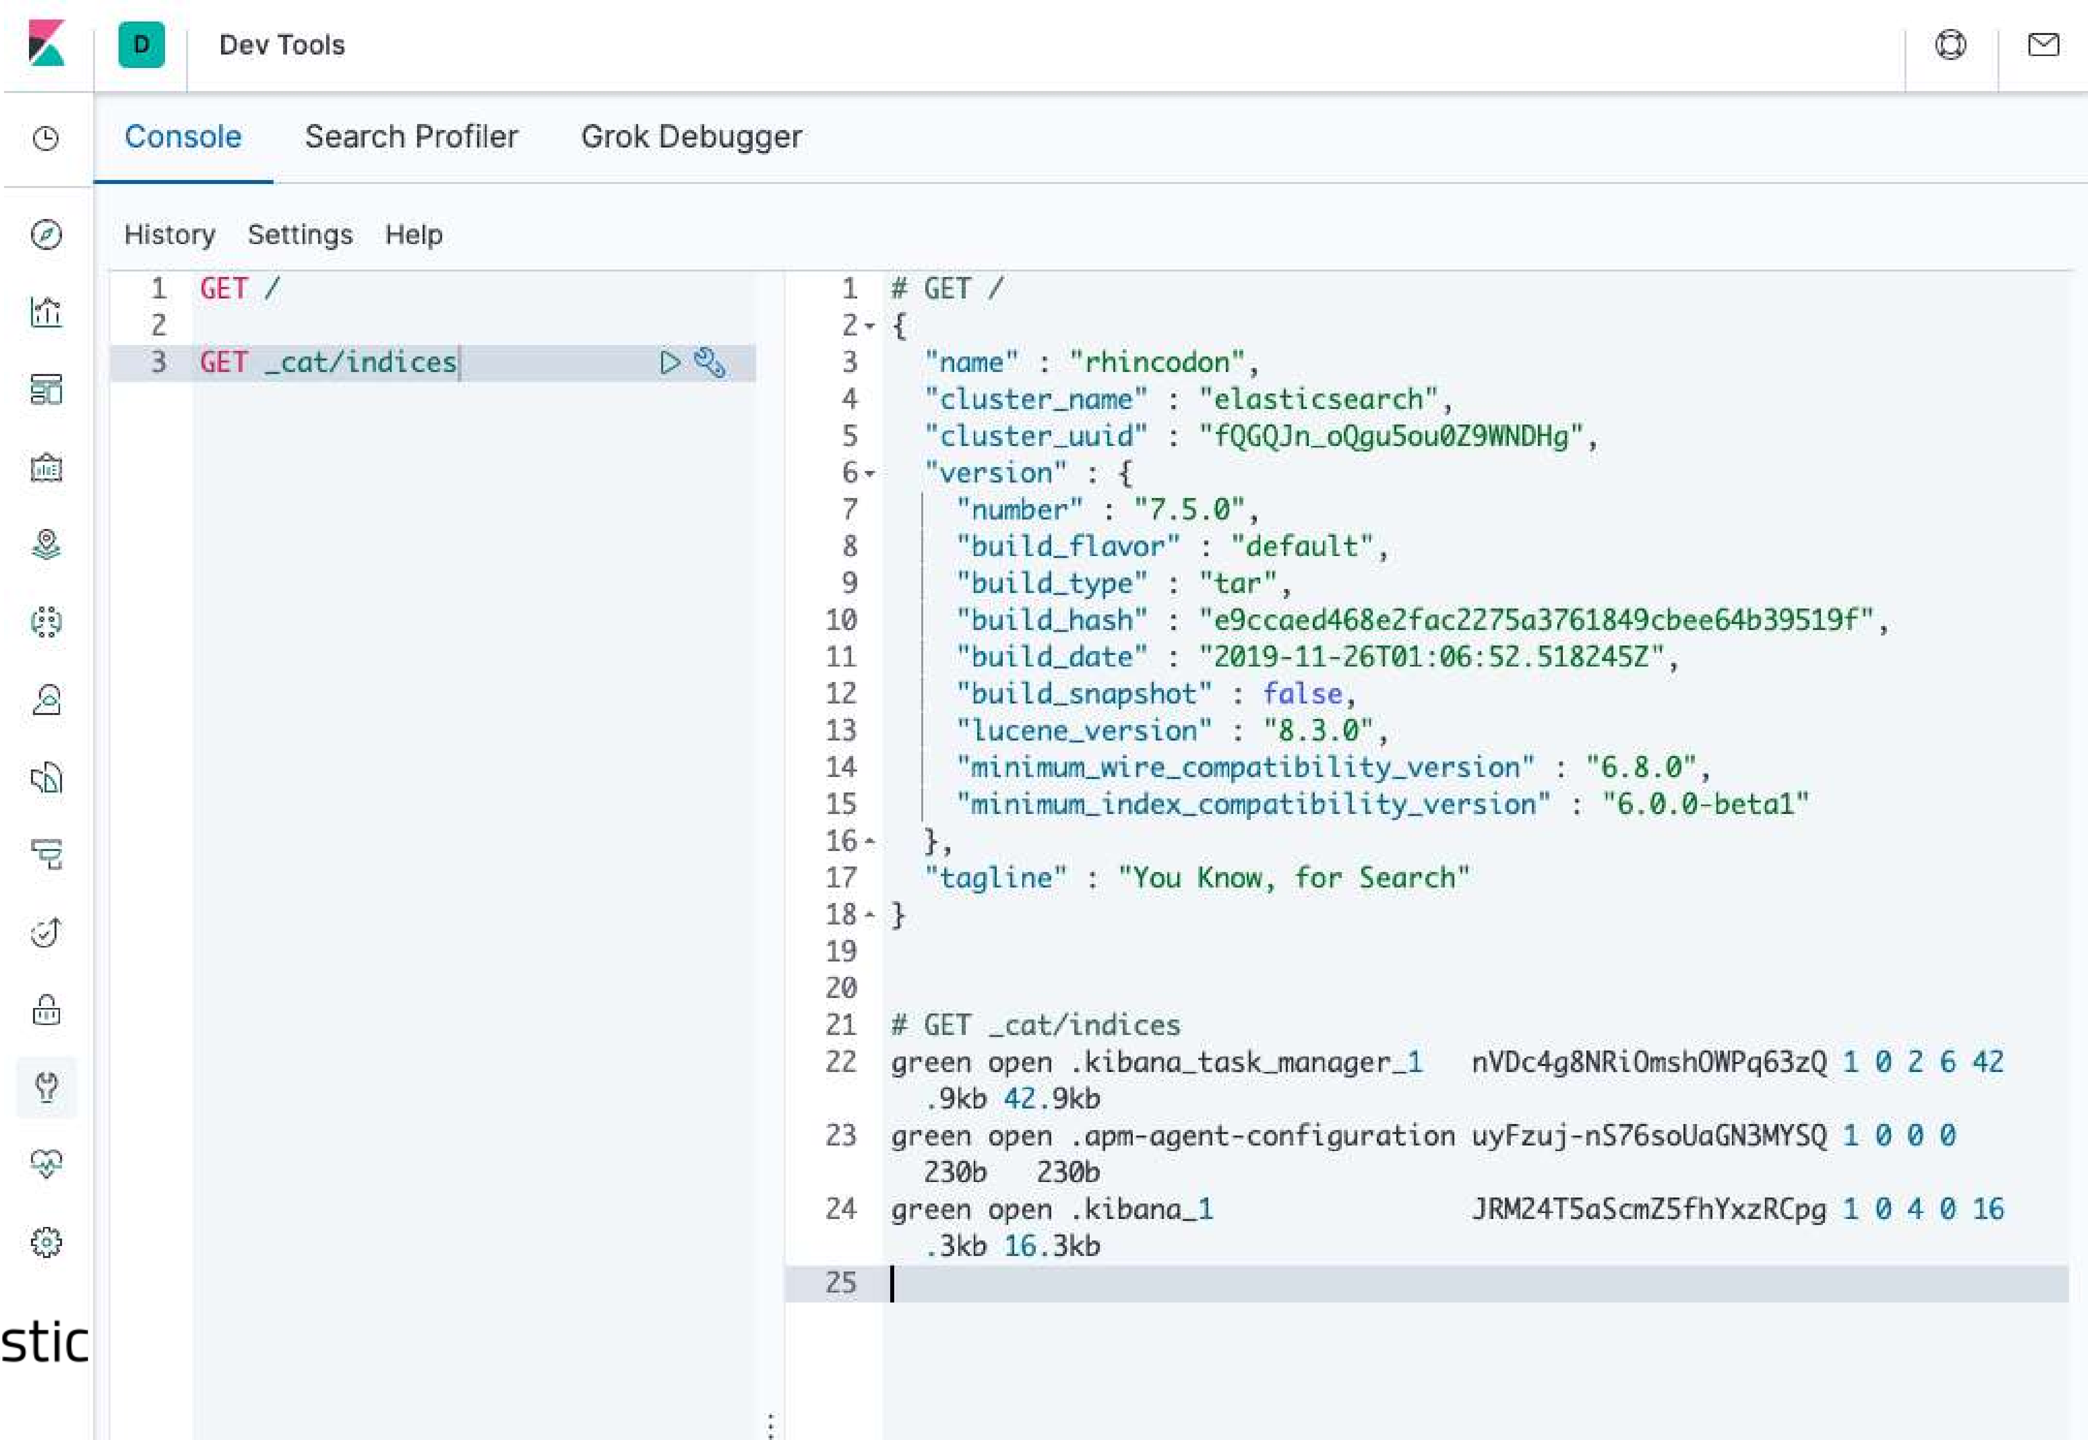

```
GET /

GET _cat/indices
```

```
# GET /
{
  "name" : "rhincodon",
  "cluster_name" : "elasticsearch",
  "cluster_uuid" : "fQGQJn_oQgu5ou0Z9WNDHg",
  "version" : {
    "number" : "7.5.0",
    "build_flavor" : "default",
    "build_type" : "tar",
    "build_hash" : "e9ccaed468e2fac2275a3761849cbee64b39519f",
    "build_date" : "2019-11-26T01:06:52.518245Z",
    "build_snapshot" : false,
    "lucene_version" : "8.3.0",
    "minimum_wire_compatibility_version" : "6.8.0",
    "minimum_index_compatibility_version" : "6.0.0-beta1"
  },
  "tagline" : "You Know, for Search"
}


# GET _cat/indices
green open .kibana_task_manager_1   nVDc4g8NRiOmshOWPq63zQ 1 0 2 6 42
  .9kb 42.9kb
green open .apm-agent-configuration uyFzuj-nS76soUaGN3MYSQ 1 0 0 0
  230b   230b
green open .kibana_1                JRM24T5aScmZ5fhYxzRCpg 1 0 4 0 16
  .3kb 16.3kb
```

Demo

Indexing, Mapping & Enrichment

- Index API
- Bulk API
- Put Mapping API
- Datatypes
- Enrichment

Document search & Aggregations

- Query DSL
- Search API
- Aggregations

Administration tasks

- Snapshot and restore
- Reindexing
- ILM
- Monitoring
- Frozen Indices
- Securing a cluster

Elasticsearch Clients

- Not just glorified HTTP clients
- Retry after failure
- Sniffing
- Bulk helpers
- Java, JavaScript, Ruby, Go, .NET, PHP, Perl, Python, Rust

Elasticsearch is distributed!

- Scaling reads, scaling writes, ensuring high availibility
- Run as single node or hundreds of nodes together
- Users should never care if they query/index against a small or big cluster
- Add a new node, Elasticsearch will balance data & queries automatically
- Specialized roles (master, data, ingest, ml, voting only)
- Orchestration becomes more important as use-case clusters might be easier to maintain & upgrade than the one big cluster

More, more, more...

- More Queries, aggregations & data types
- Text analysis (phonetic search, search as you type)
- ILM, rollup, transform, frozen indices
- Security
- Alerting
- SQL
- Machine Learning
- Stack Monitoring
- Major version upgrades & deprecations
- Solutions (Observability, Enterprise Search, Security)

Summary

- Understanding search is hard
- Use the reference documentation
- Ask your users about expectations, do not guess!

Next steps

Check out https://demo.elastic.co

Check out Observability

- Uptime
- Metrics
- Logs
- APM

Uptime

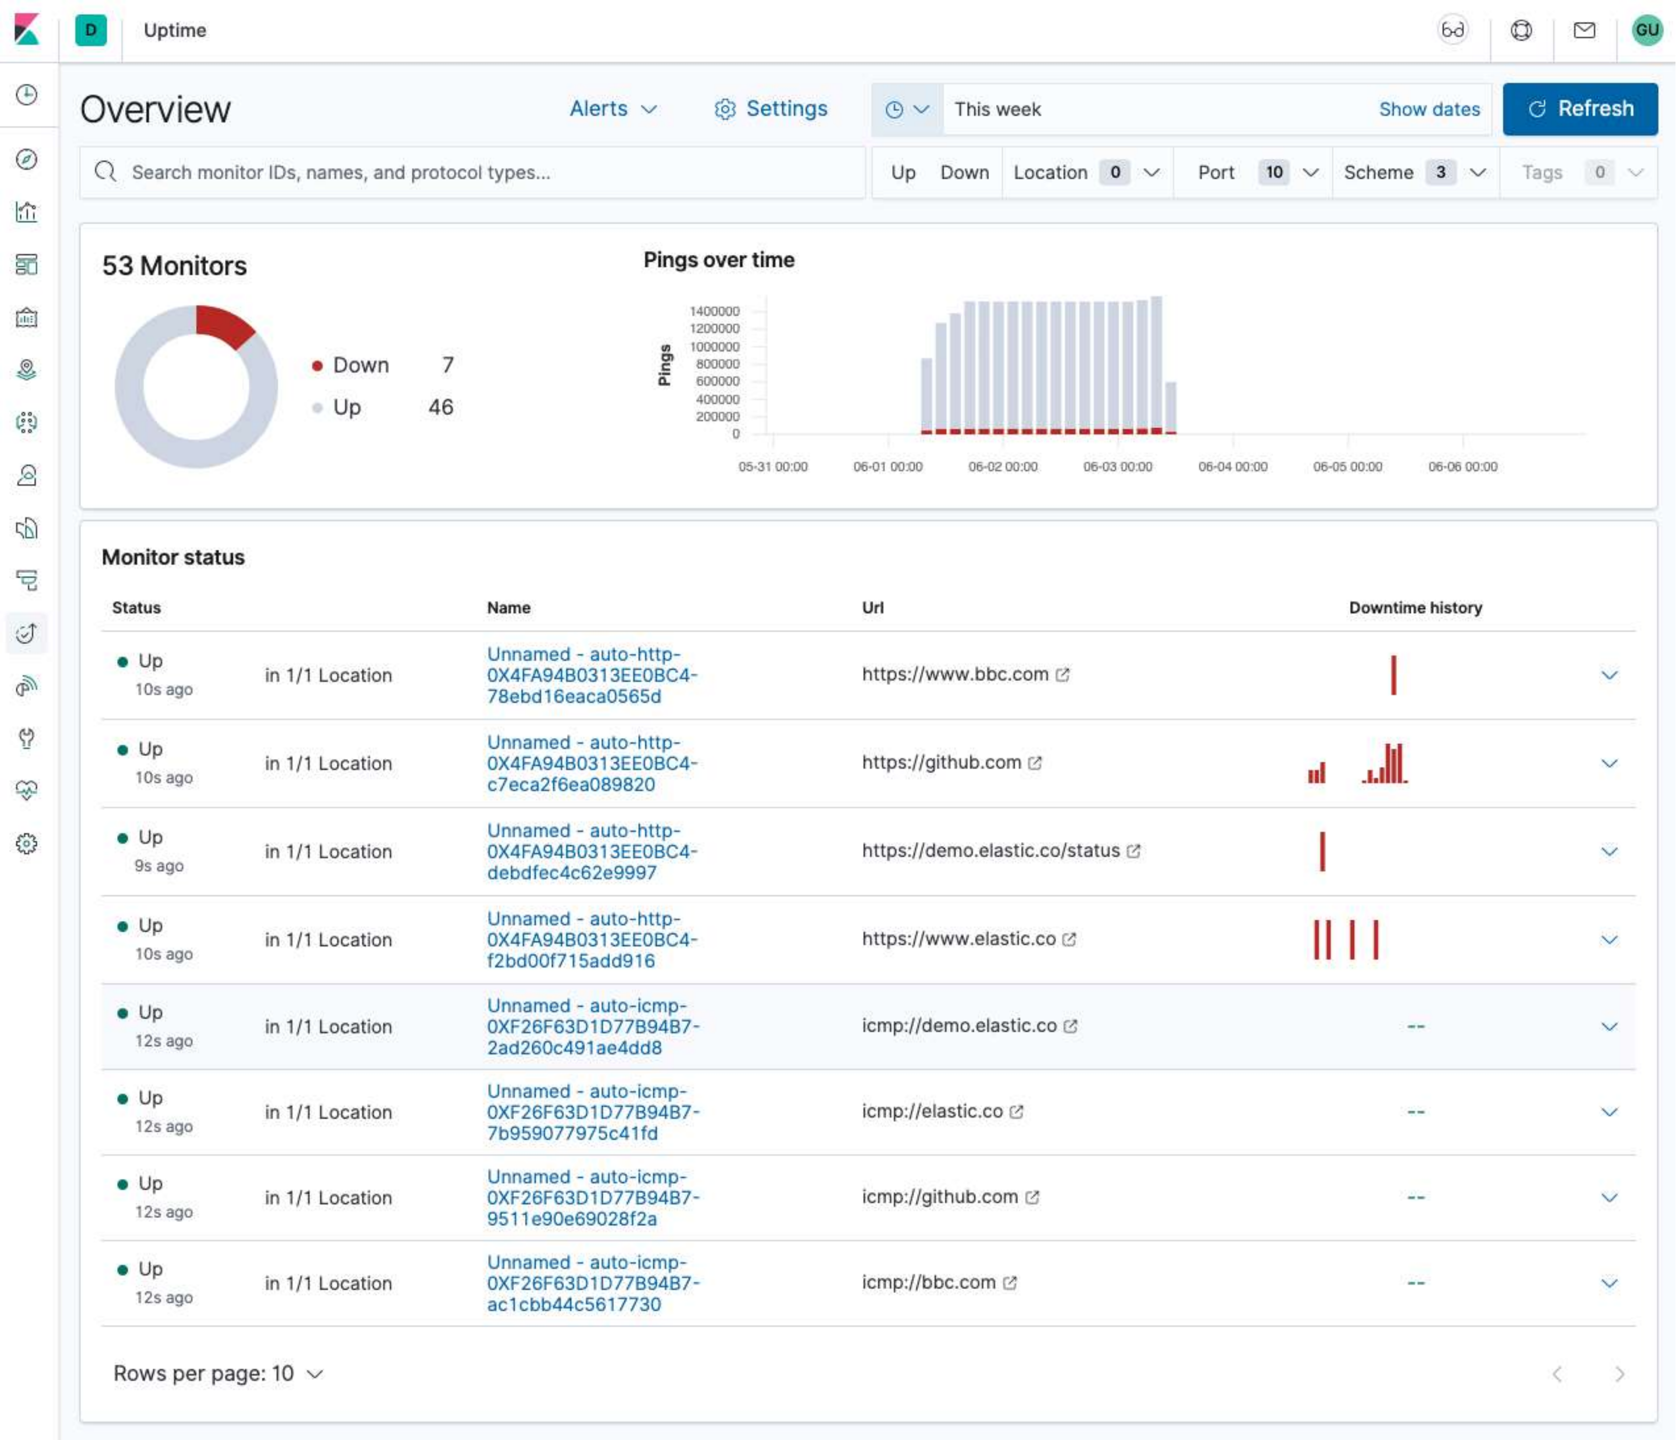

53 Monitors

● Down	7
● Up	46

Monitor status

Status		Name	Url	Downtime history
● Up 10s ago	in 1/1 Location	Unnamed - auto-http-0X4FA94B0313EE0BC4-78ebd16eaca0565d	https://www.bbc.com	
● Up 10s ago	in 1/1 Location	Unnamed - auto-http-0X4FA94B0313EE0BC4-c7eca2f6ea089820	https://github.com	
● Up 9s ago	in 1/1 Location	Unnamed - auto-http-0X4FA94B0313EE0BC4-debdfec4c62e9997	https://demo.elastic.co/status	
● Up 10s ago	in 1/1 Location	Unnamed - auto-http-0X4FA94B0313EE0BC4-f2bd00f715add916	https://www.elastic.co	
● Up 12s ago	in 1/1 Location	Unnamed - auto-icmp-0XF26F63D1D77B94B7-2ad260c491ae4dd8	icmp://demo.elastic.co	--
● Up 12s ago	in 1/1 Location	Unnamed - auto-icmp-0XF26F63D1D77B94B7-7b959077975c41fd	icmp://elastic.co	--
● Up 12s ago	in 1/1 Location	Unnamed - auto-icmp-0XF26F63D1D77B94B7-9511e90e69028f2a	icmp://github.com	--
● Up 12s ago	in 1/1 Location	Unnamed - auto-icmp-0XF26F63D1D77B94B7-ac1cbb44c5617730	icmp://bbc.com	--

Rows per page: 10

Metrics

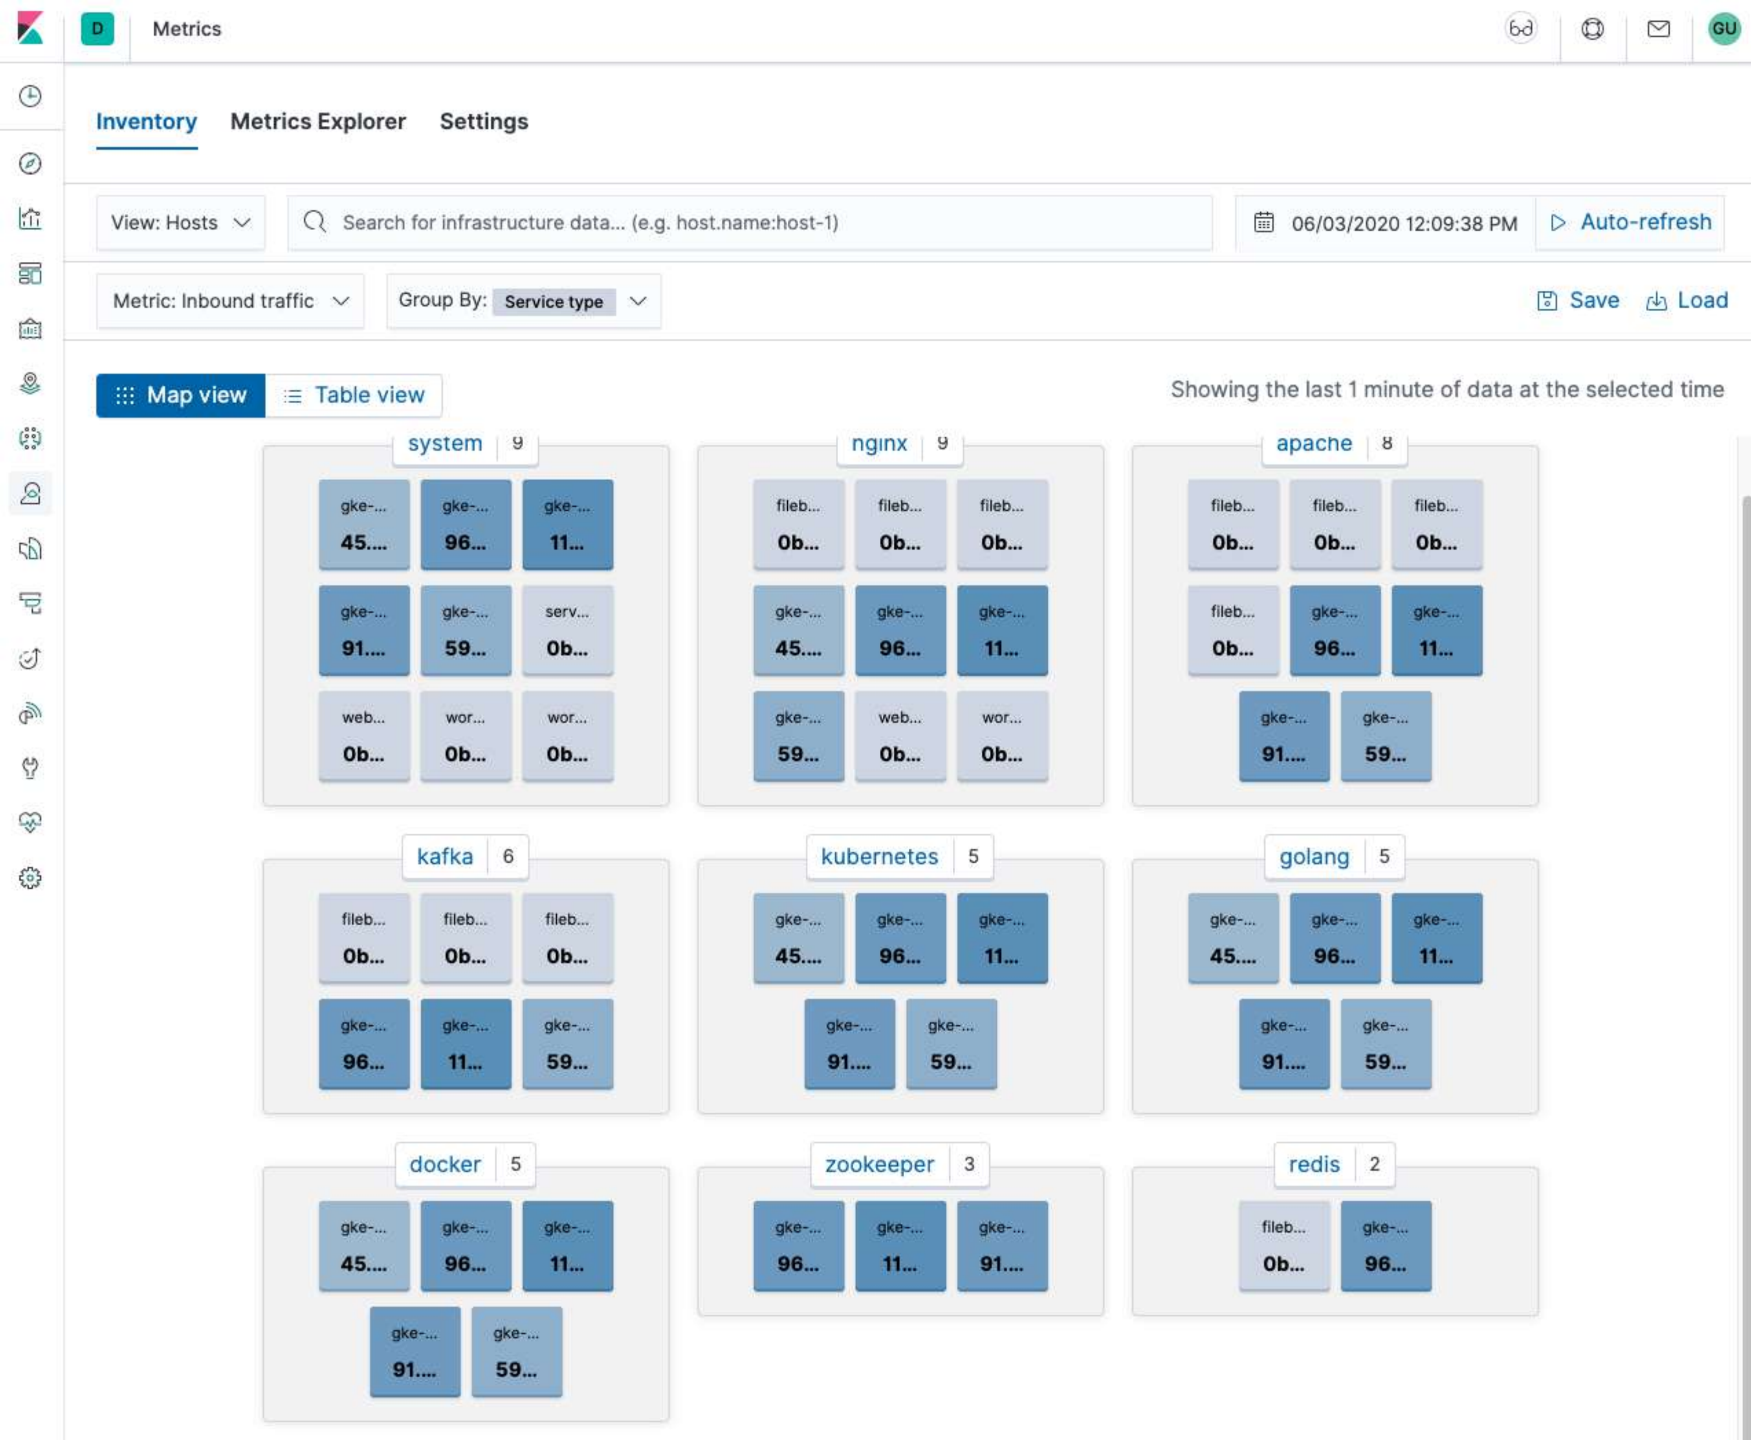

View: Hosts

06/03/2020 12:09:38 PM

Metric: Inbound traffic

Group By: Service type

Showing the last 1 minute of data at the selected time

Logs

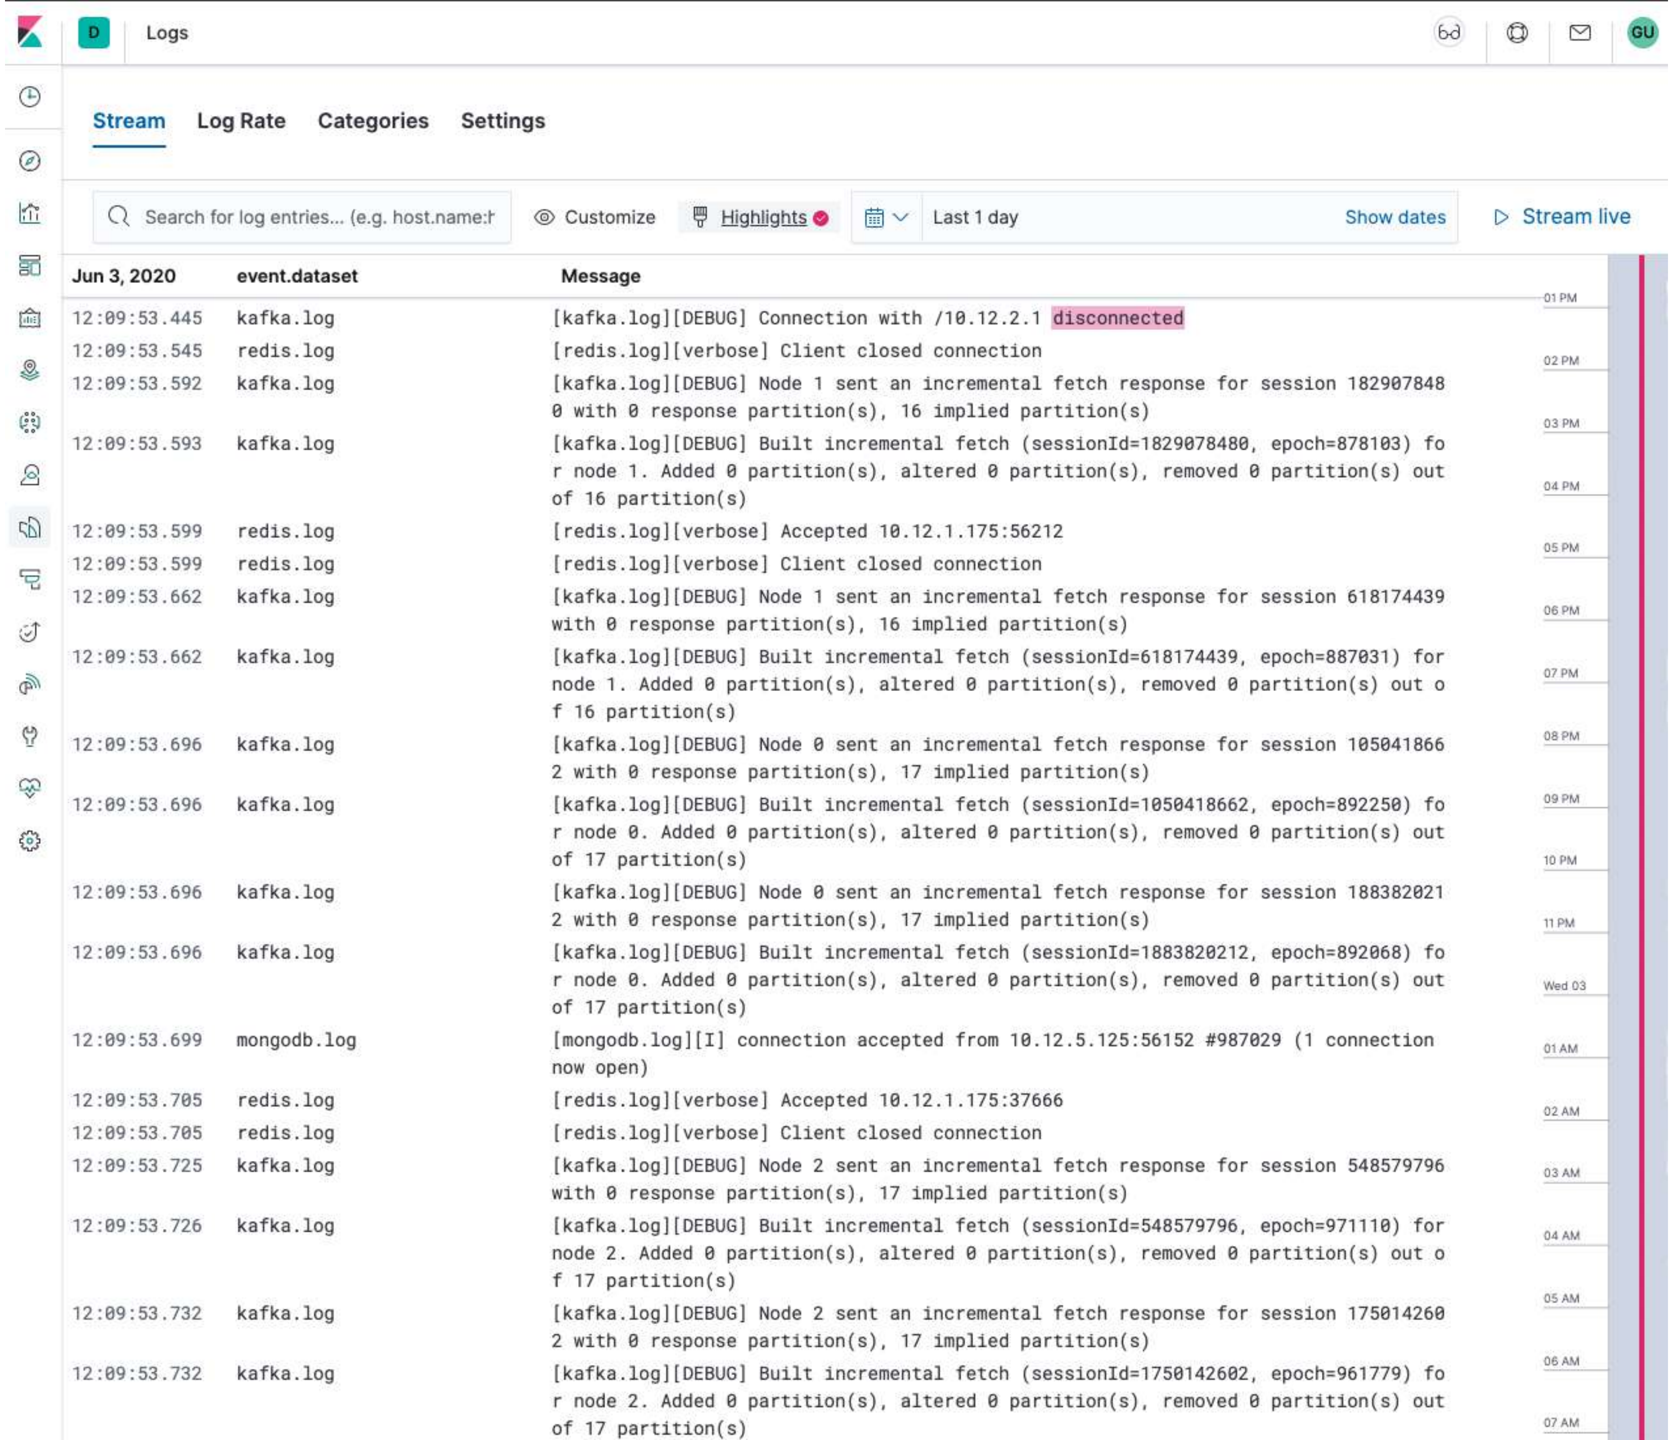

```
Jun 3, 2020   event.dataset   Message
12:09:53.445  kafka.log       [kafka.log][DEBUG] Connection with /10.12.2.1 disconnected
12:09:53.545  redis.log       [redis.log][verbose] Client closed connection
12:09:53.592  kafka.log       [kafka.log][DEBUG] Node 1 sent an incremental fetch response for session 182907848
                              0 with 0 response partition(s), 16 implied partition(s)
12:09:53.593  kafka.log       [kafka.log][DEBUG] Built incremental fetch (sessionId=1829078480, epoch=878103) fo
                              r node 1. Added 0 partition(s), altered 0 partition(s), removed 0 partition(s) out
                              of 16 partition(s)
12:09:53.599  redis.log       [redis.log][verbose] Accepted 10.12.1.175:56212
12:09:53.599  redis.log       [redis.log][verbose] Client closed connection
12:09:53.662  kafka.log       [kafka.log][DEBUG] Node 1 sent an incremental fetch response for session 618174439
                              with 0 response partition(s), 16 implied partition(s)
12:09:53.662  kafka.log       [kafka.log][DEBUG] Built incremental fetch (sessionId=618174439, epoch=887031) for
                              node 1. Added 0 partition(s), altered 0 partition(s), removed 0 partition(s) out o
                              f 16 partition(s)
12:09:53.696  kafka.log       [kafka.log][DEBUG] Node 0 sent an incremental fetch response for session 105041866
                              2 with 0 response partition(s), 17 implied partition(s)
12:09:53.696  kafka.log       [kafka.log][DEBUG] Built incremental fetch (sessionId=1050418662, epoch=892250) fo
                              r node 0. Added 0 partition(s), altered 0 partition(s), removed 0 partition(s) out
                              of 17 partition(s)
12:09:53.696  kafka.log       [kafka.log][DEBUG] Node 0 sent an incremental fetch response for session 188382021
                              2 with 0 response partition(s), 17 implied partition(s)
12:09:53.696  kafka.log       [kafka.log][DEBUG] Built incremental fetch (sessionId=1883820212, epoch=892068) fo
                              r node 0. Added 0 partition(s), altered 0 partition(s), removed 0 partition(s) out
                              of 17 partition(s)
12:09:53.699  mongodb.log     [mongodb.log][I] connection accepted from 10.12.5.125:56152 #987029 (1 connection
                              now open)
12:09:53.705  redis.log       [redis.log][verbose] Accepted 10.12.1.175:37666
12:09:53.705  redis.log       [redis.log][verbose] Client closed connection
12:09:53.725  kafka.log       [kafka.log][DEBUG] Node 2 sent an incremental fetch response for session 548579796
                              with 0 response partition(s), 17 implied partition(s)
12:09:53.726  kafka.log       [kafka.log][DEBUG] Built incremental fetch (sessionId=548579796, epoch=971110) for
                              node 2. Added 0 partition(s), altered 0 partition(s), removed 0 partition(s) out o
                              f 17 partition(s)
12:09:53.732  kafka.log       [kafka.log][DEBUG] Node 2 sent an incremental fetch response for session 175014260
                              2 with 0 response partition(s), 17 implied partition(s)
12:09:53.732  kafka.log       [kafka.log][DEBUG] Built incremental fetch (sessionId=1750142602, epoch=961779) fo
                              r node 2. Added 0 partition(s), altered 0 partition(s), removed 0 partition(s) out
                              of 17 partition(s)
```

elastic

APM

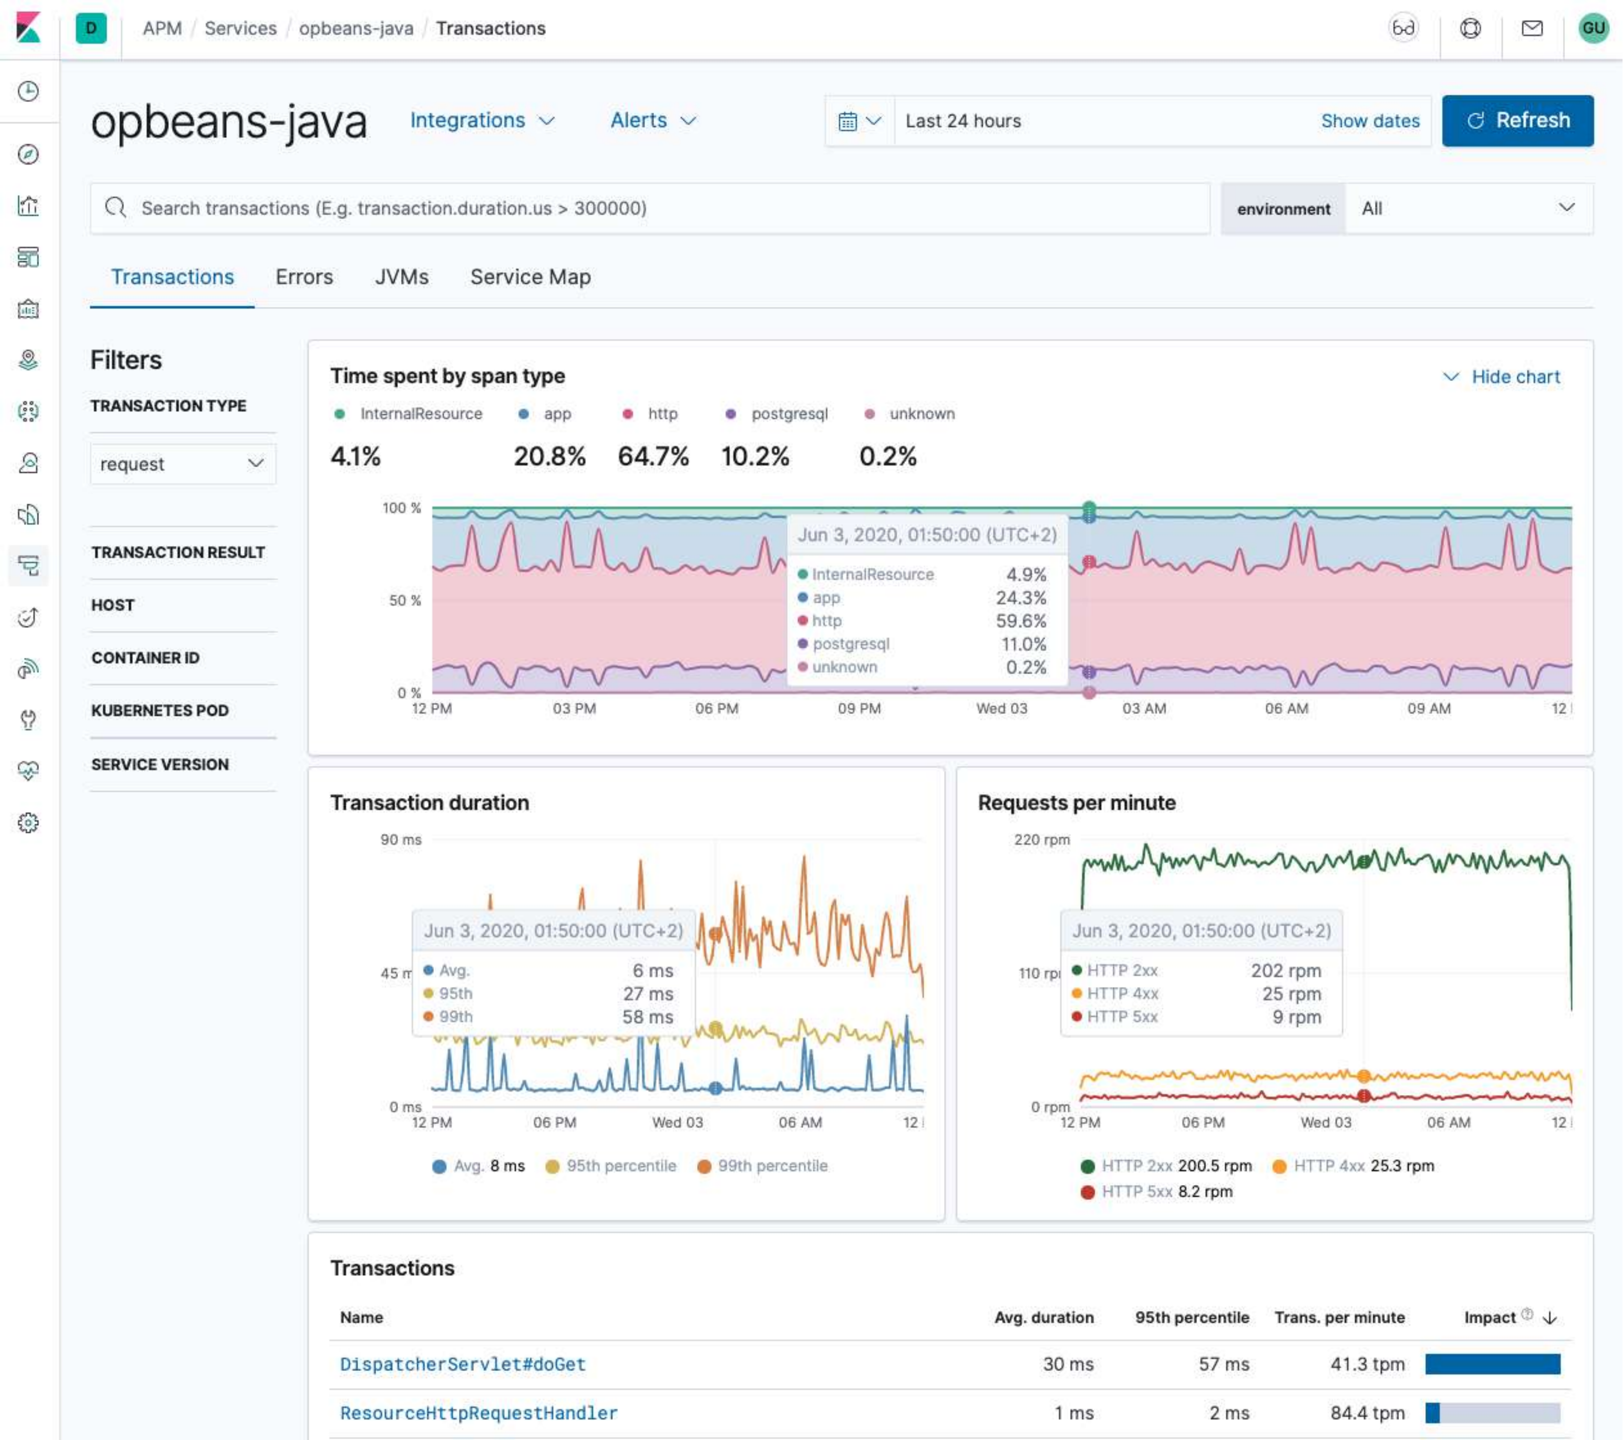

Transactions

Name	Avg. duration	95th percentile	Trans. per minute	Impact
DispatcherServlet#doGet	30 ms	57 ms	41.3 tpm	
ResourceHttpRequestHandler	1 ms	2 ms	84.4 tpm	

elastic

APM

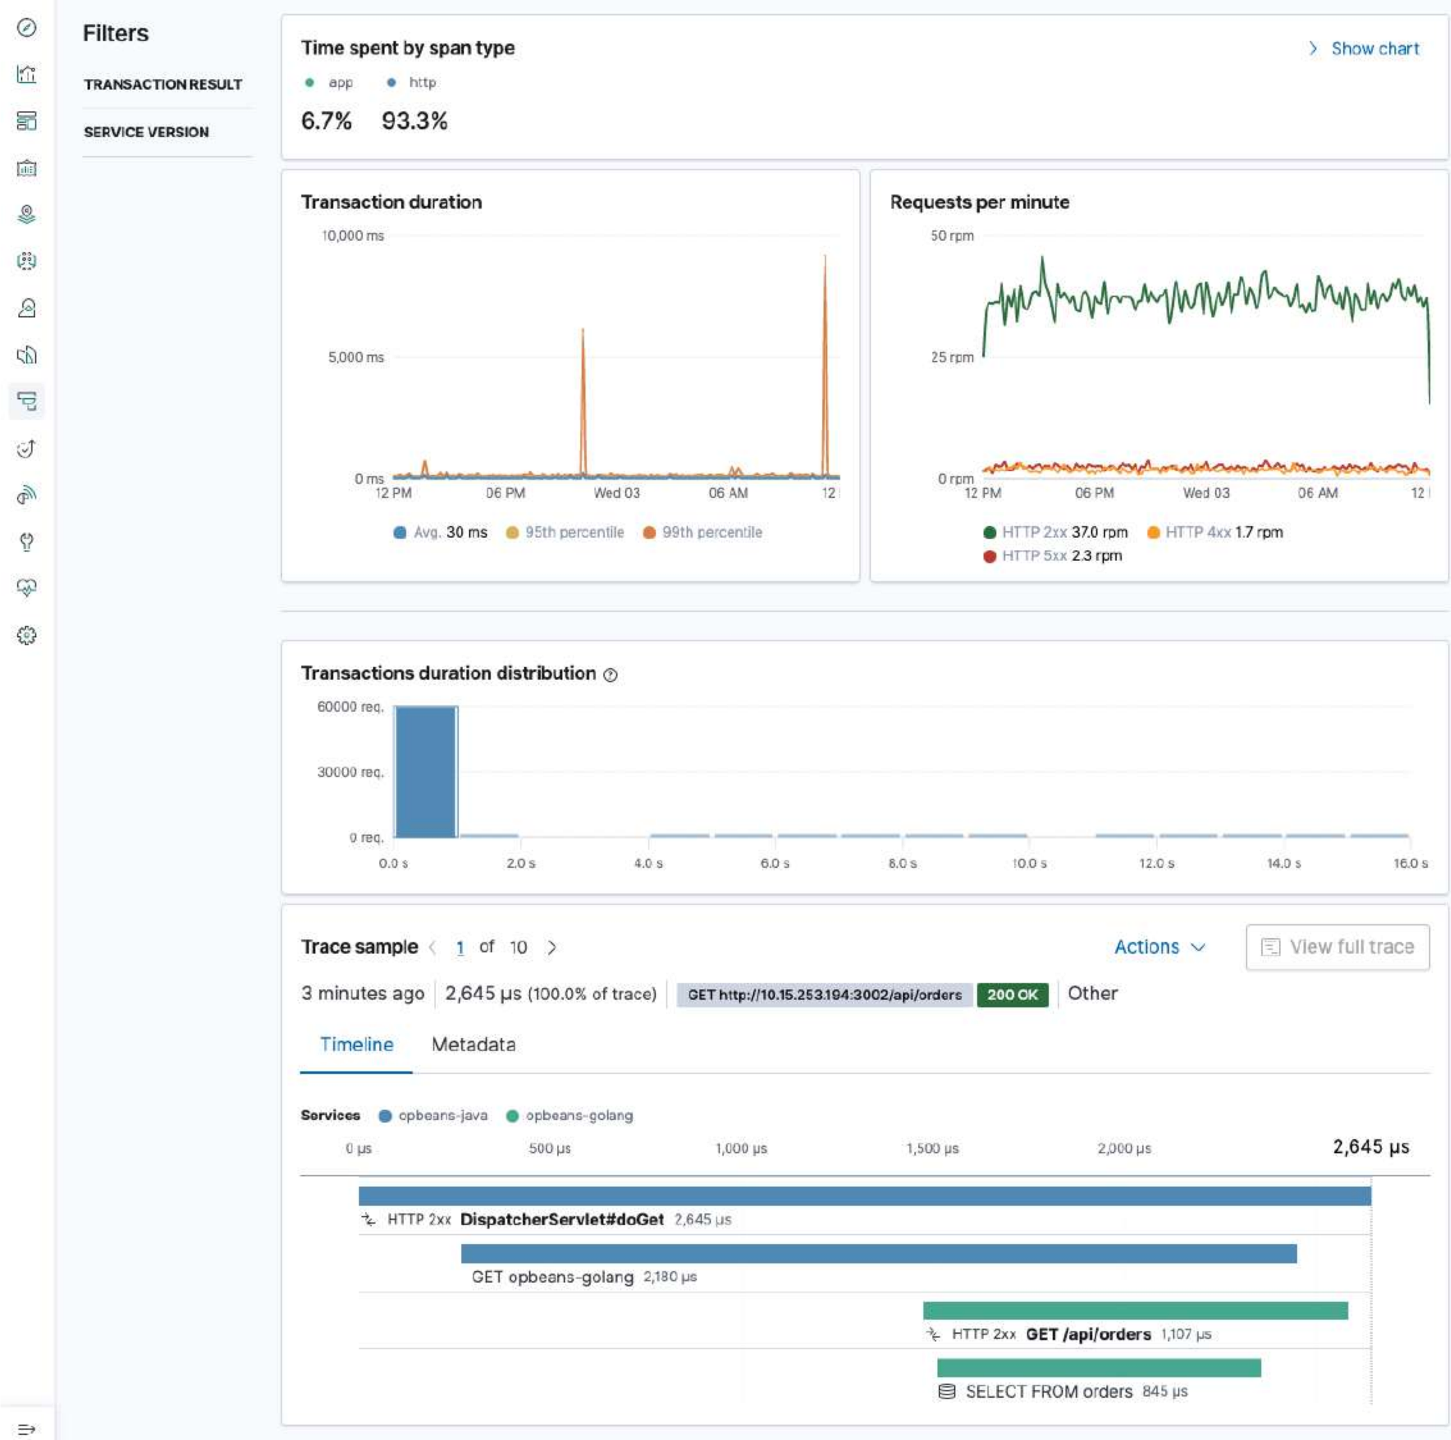
Filters
TRANSACTION RESULT
SERVICE VERSION
Time spent by span type
Show chart
● app ● http
6.7% 93.3%
Transaction duration
10,000 ms
5,000 ms
0 ms
12 PM 06 PM Wed 03 06 AM 12
● Avg. 30 ms ● 95th percentile ● 99th percentile
Requests per minute
50 rpm
25 rpm
0 rpm
12 PM 06 PM Wed 03 06 AM 12
● HTTP 2xx 37.0 rpm ● HTTP 4xx 1.7 rpm
● HTTP 5xx 2.3 rpm
Transactions duration distribution
60000 req.
30000 req.
0 req.
0.0 s 2.0 s 4.0 s 6.0 s 8.0 s 10.0 s 12.0 s 14.0 s 16.0 s
Trace sample 1 of 10
Actions
View full trace
3 minutes ago
2,645 μs (100.0% of trace)
GET http://10.15.253.194:3002/api/orders
200 OK
Other
Timeline
Metadata
Services ● opbeans-java ● opbeans-golang
0 μs 500 μs 1,000 μs 1,500 μs 2,000 μs 2,645 μs
HTTP 2xx DispatcherServlet#doGet 2,645 μs
GET opbeans-golang 2,180 μs
HTTP 2xx GET /api/orders 1,107 μs
SELECT FROM orders 845 μs

Check out Security

- SIEM
- Endpoint Security

SIEM

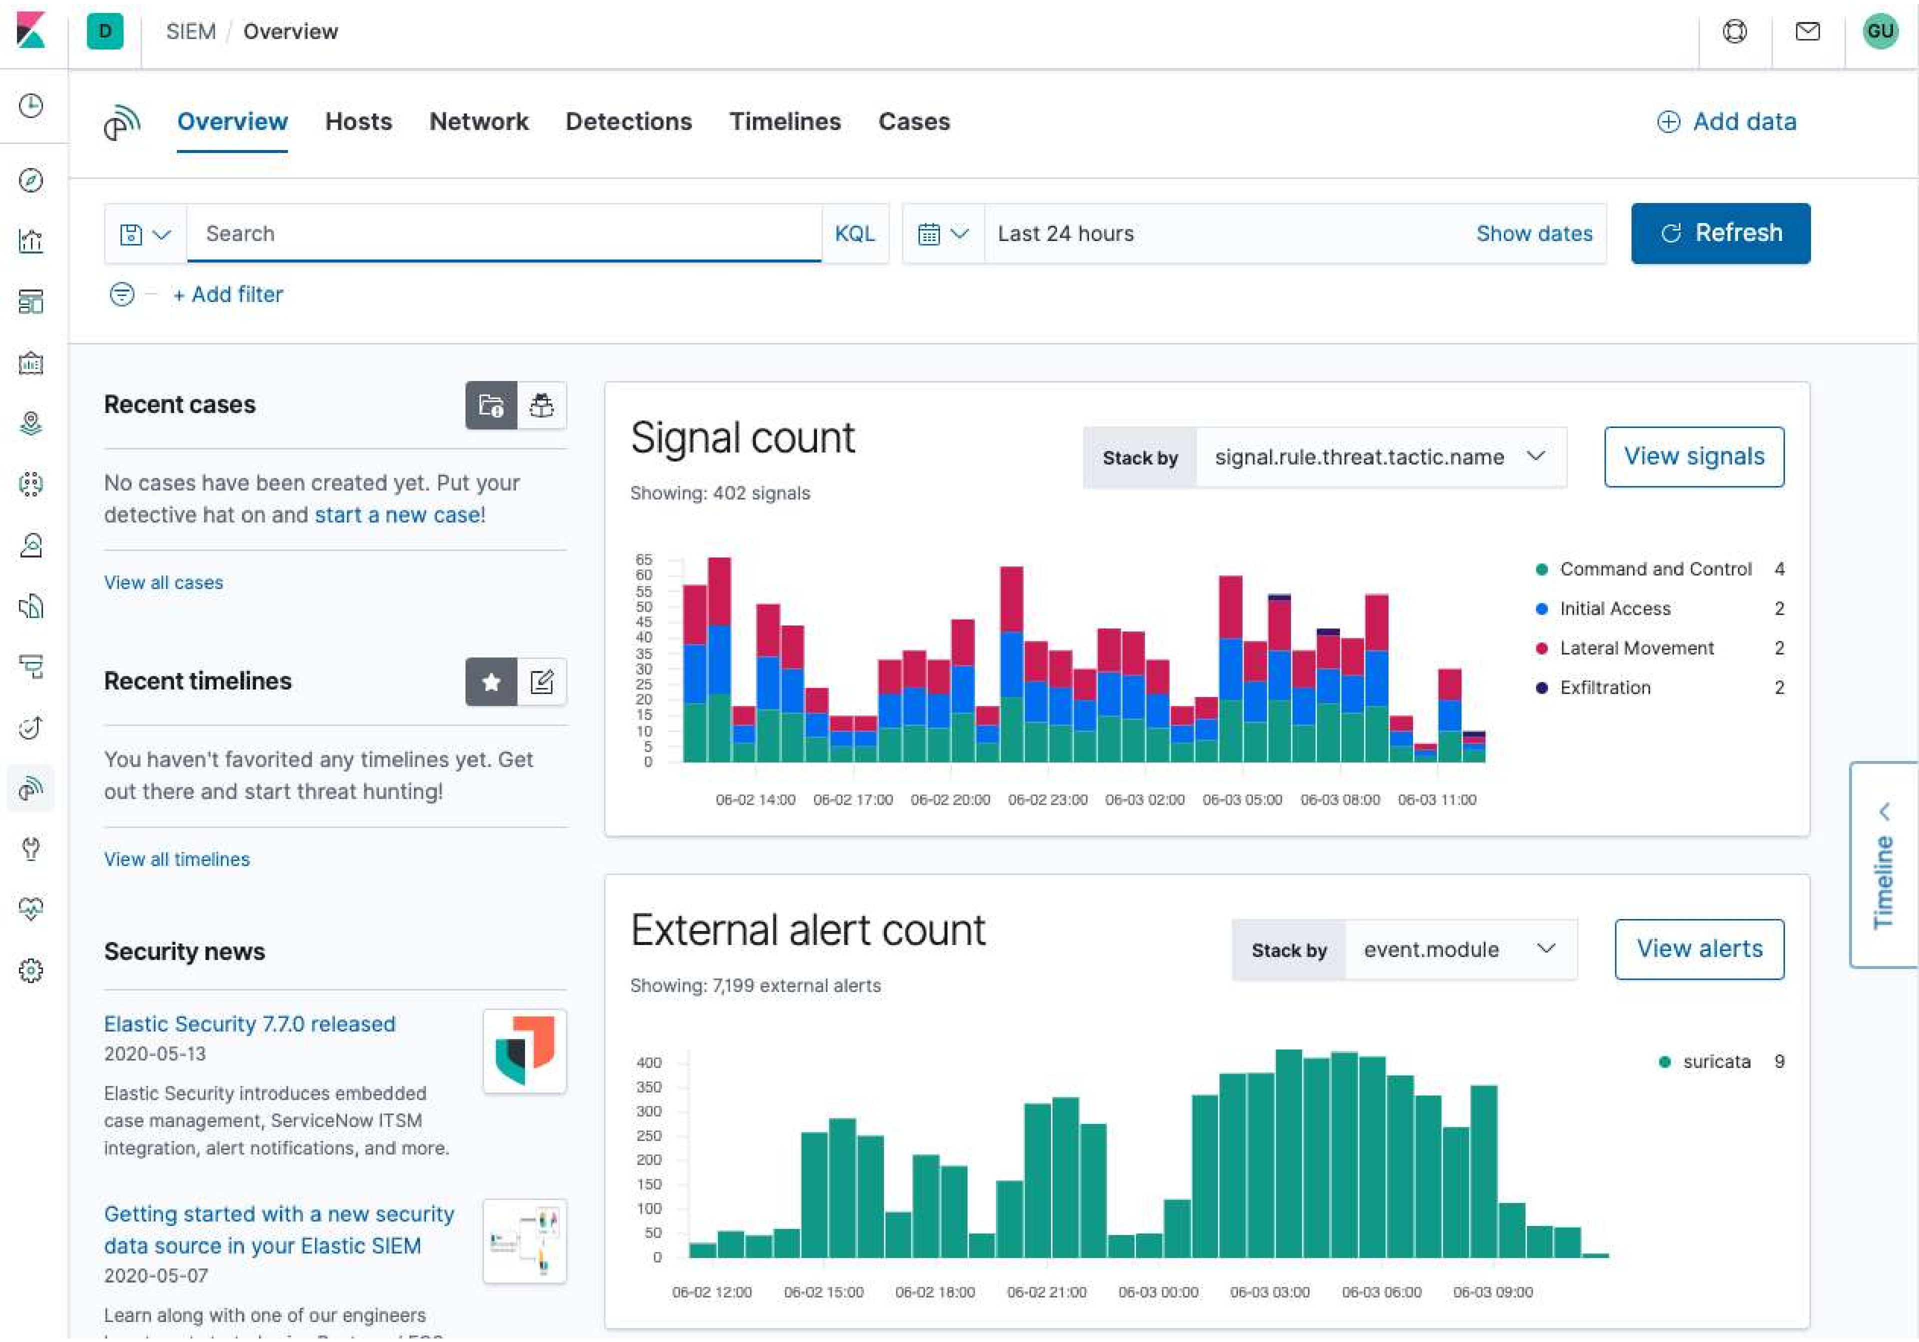

SIEM / Overview

Overview | Hosts | Network | Detections | Timelines | Cases

Add data

Search | KQL | Last 24 hours | Show dates | Refresh

+ Add filter

Recent cases

No cases have been created yet. Put your detective hat on and start a new case!

View all cases

Recent timelines

You haven't favorited any timelines yet. Get out there and start threat hunting!

View all timelines

Security news

Elastic Security 7.7.0 released
2020-05-13
Elastic Security introduces embedded case management, ServiceNow ITSM integration, alert notifications, and more.

Getting started with a new security data source in your Elastic SIEM
2020-05-07
Learn along with one of our engineers

Signal count

Stack by: signal.rule.threat.tactic.name

View signals

Showing: 402 signals

External alert count

Stack by: event.module

View alerts

Showing: 7,199 external alerts

Check out Enterprise Search

- Workplace Search
- App Search

App Search

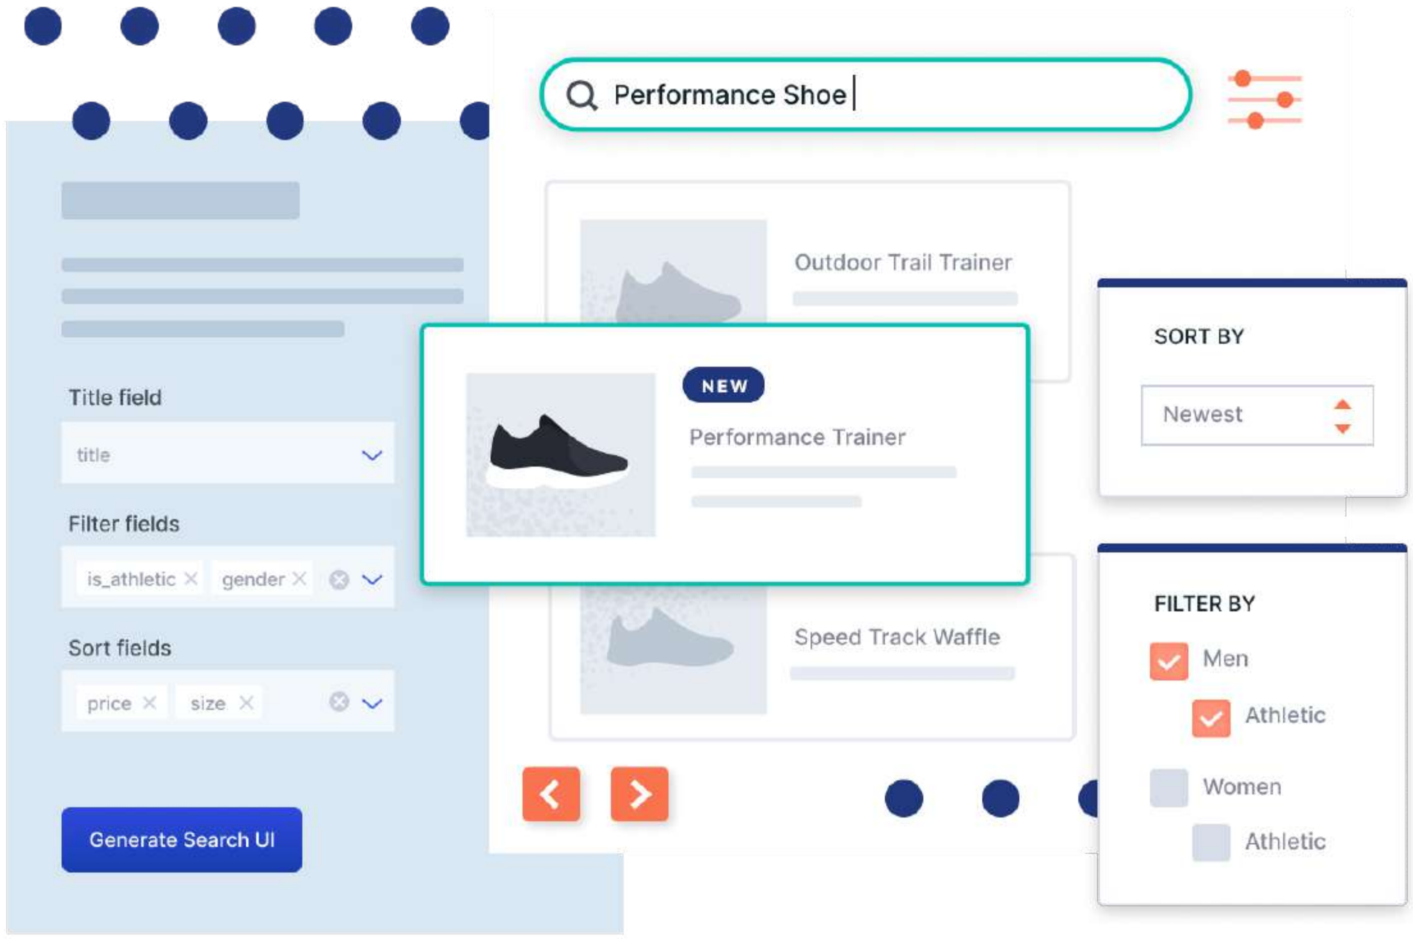

Title field: title

Filter fields: is_athletic, gender

Sort fields: price, size

Performance Shoe

SORT BY: Newest

FILTER BY

☑ Men

☑ Athletic

□ Women

□ Athletic

Workplace Search

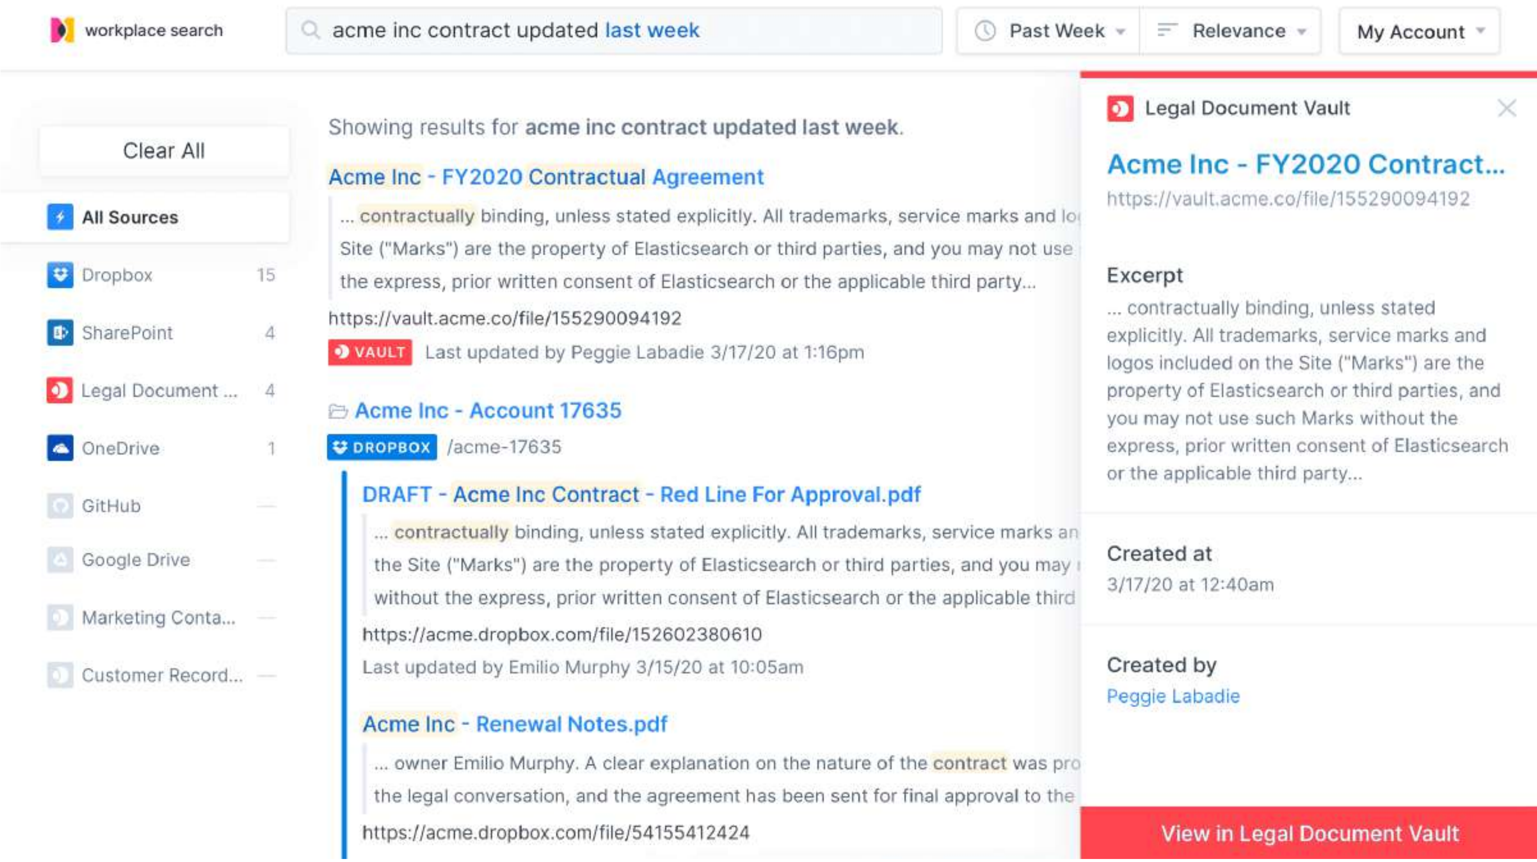

acme inc contract updated last week

Showing results for **acme inc contract updated last week**.

Acme Inc - FY2020 Contractual Agreement

... contractually binding, unless stated explicitly. All trademarks, service marks and lo
Site ("Marks") are the property of Elasticsearch or third parties, and you may not use
the express, prior written consent of Elasticsearch or the applicable third party...

https://vault.acme.co/file/155290094192

Last updated by Peggie Labadie 3/17/20 at 1:16pm

Acme Inc - Account 17635

/acme-17635

DRAFT - Acme Inc Contract - Red Line For Approval.pdf

... contractually binding, unless stated explicitly. All trademarks, service marks an
the Site ("Marks") are the property of Elasticsearch or third parties, and you may
without the express, prior written consent of Elasticsearch or the applicable third

https://acme.dropbox.com/file/152602380610

Last updated by Emilio Murphy 3/15/20 at 10:05am

Acme Inc - Renewal Notes.pdf

... owner Emilio Murphy. A clear explanation on the nature of the contract was pro
the legal conversation, and the agreement has been sent for final approval to the

https://acme.dropbox.com/file/54155412424

Legal Document Vault

Acme Inc - FY2020 Contract...

https://vault.acme.co/file/155290094192

Excerpt

... contractually binding, unless stated explicitly. All trademarks, service marks and logos included on the Site ("Marks") are the property of Elasticsearch or third parties, and you may not use such Marks without the express, prior written consent of Elasticsearch or the applicable third party...

Created at

3/17/20 at 12:40am

Created by

Peggie Labadie

elastic

Connectors

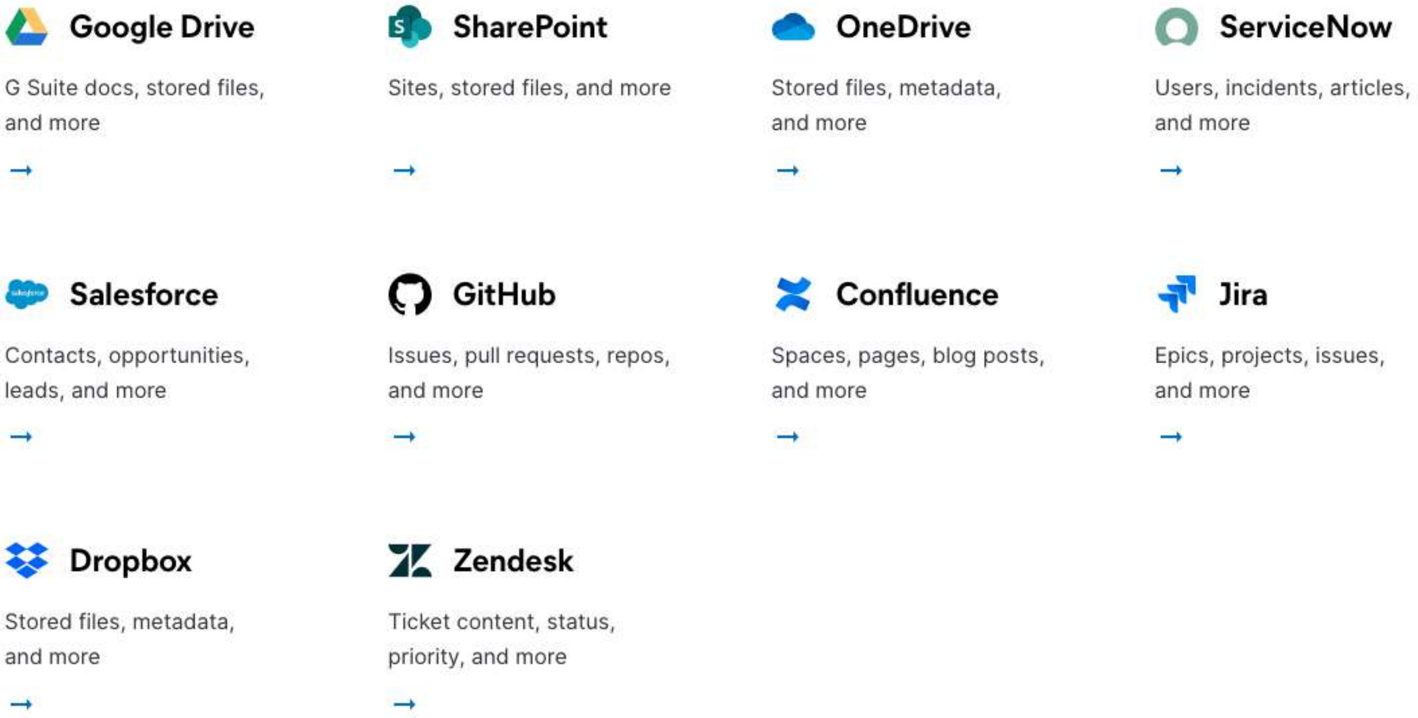

Google Drive

G Suite docs, stored files, and more

→

SharePoint

Sites, stored files, and more

→

OneDrive

Stored files, metadata, and more

→

ServiceNow

Users, incidents, articles, and more

→

Salesforce

Contacts, opportunities, leads, and more

→

GitHub

Issues, pull requests, repos, and more

→

Confluence

Spaces, pages, blog posts, and more

→

Jira

Epics, projects, issues, and more

→

Dropbox

Stored files, metadata, and more

→

Zendesk

Ticket content, status, priority, and more

→

elastic

Getting more help

Discuss Forum

https://discuss.elastic.co

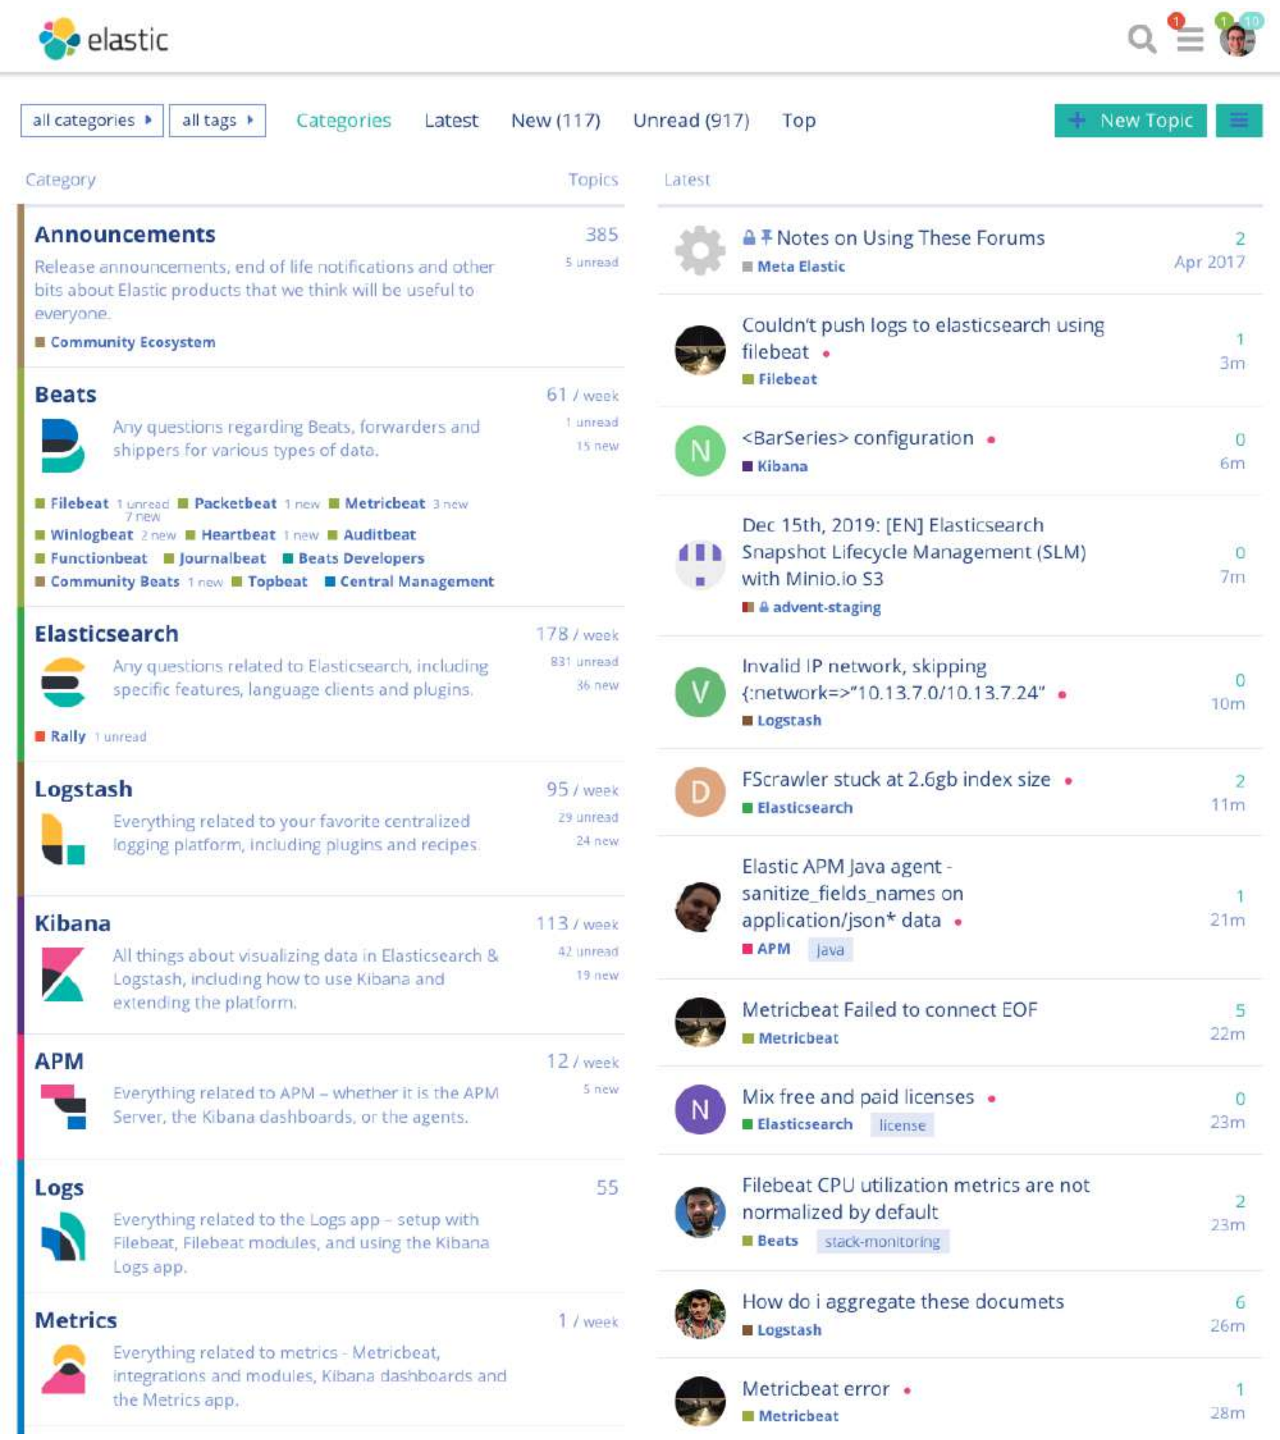

Category	Topics
Announcements Release announcements, end of life notifications and other bits about Elastic products that we think will be useful to everyone. ■ Community Ecosystem	385 5 unread
Beats Any questions regarding Beats, forwarders and shippers for various types of data. ■ Filebeat 1 unread 7 new ■ Packetbeat 1 new ■ Metricbeat 3 new ■ Winlogbeat 2 new ■ Heartbeat 1 new ■ Auditbeat ■ Functionbeat ■ Journalbeat ■ Beats Developers ■ Community Beats 1 new ■ Topbeat ■ Central Management	61 / week 1 unread 15 new
Elasticsearch Any questions related to Elasticsearch, including specific features, language clients and plugins. ■ Rally 1 unread	178 / week 831 unread 36 new
Logstash Everything related to your favorite centralized logging platform, including plugins and recipes.	95 / week 29 unread 24 new
Kibana All things about visualizing data in Elasticsearch & Logstash, including how to use Kibana and extending the platform.	113 / week 42 unread 19 new
APM Everything related to APM – whether it is the APM Server, the Kibana dashboards, or the agents.	12 / week 5 new
Logs Everything related to the Logs app – setup with Filebeat, Filebeat modules, and using the Kibana Logs app.	55
Metrics Everything related to metrics - Metricbeat, integrations and modules, Kibana dashboards and the Metrics app.	1 / week

Latest		
Notes on Using These Forums ■ Meta Elastic	2	Apr 2017
Couldn't push logs to elasticsearch using filebeat • ■ Filebeat	1	3m
<BarSeries> configuration • ■ Kibana	0	6m
Dec 15th, 2019: [EN] Elasticsearch Snapshot Lifecycle Management (SLM) with Minio.io S3 ■ advent-staging	0	7m
Invalid IP network, skipping {:network=>"10.13.7.0/10.13.7.24" • ■ Logstash	0	10m
FScrawler stuck at 2.6gb index size • ■ Elasticsearch	2	11m
Elastic APM Java agent - sanitize_fields_names on application/json* data • ■ APM java	1	21m
Metricbeat Failed to connect EOF ■ Metricbeat	5	22m
Mix free and paid licenses • ■ Elasticsearch license	0	23m
Filebeat CPU utilization metrics are not normalized by default ■ Beats stack-monitoring	2	23m
How do i aggregate these documets ■ Logstash	6	26m
Metricbeat error • ■ Metricbeat	1	28m

Community & Meetups

https://community.elastic.co

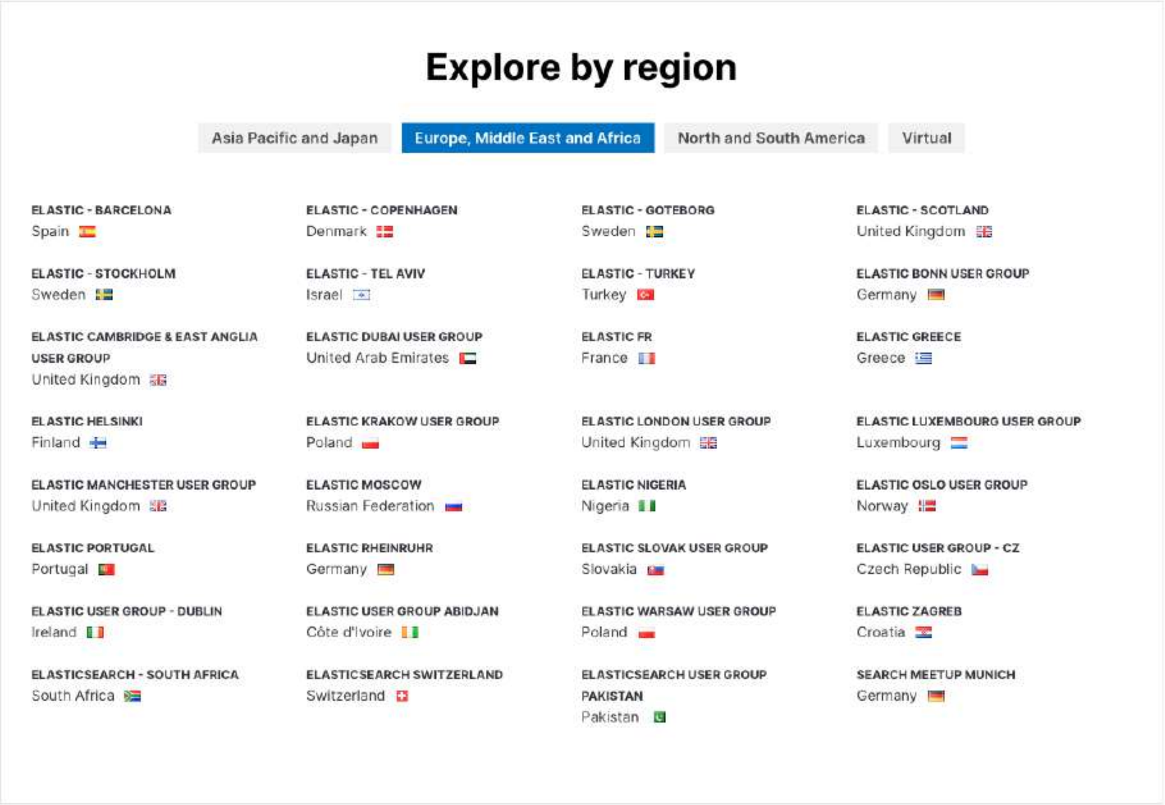

Explore by region

Asia Pacific and Japan | Europe, Middle East and Africa | North and South America | Virtual

ELASTIC - BARCELONA
Spain

ELASTIC - COPENHAGEN
Denmark

ELASTIC - GOTEBORG
Sweden

ELASTIC - SCOTLAND
United Kingdom

ELASTIC - STOCKHOLM
Sweden

ELASTIC - TEL AVIV
Israel

ELASTIC - TURKEY
Turkey

ELASTIC BONN USER GROUP
Germany

ELASTIC CAMBRIDGE & EAST ANGLIA USER GROUP
United Kingdom

ELASTIC DUBAI USER GROUP
United Arab Emirates

ELASTIC FR
France

ELASTIC GREECE
Greece

ELASTIC HELSINKI
Finland

ELASTIC KRAKOW USER GROUP
Poland

ELASTIC LONDON USER GROUP
United Kingdom

ELASTIC LUXEMBOURG USER GROUP
Luxembourg

ELASTIC MANCHESTER USER GROUP
United Kingdom

ELASTIC MOSCOW
Russian Federation

ELASTIC NIGERIA
Nigeria

ELASTIC OSLO USER GROUP
Norway

ELASTIC PORTUGAL
Portugal

ELASTIC RHEINRUHR
Germany

ELASTIC SLOVAK USER GROUP
Slovakia

ELASTIC USER GROUP - CZ
Czech Republic

ELASTIC USER GROUP - DUBLIN
Ireland

ELASTIC USER GROUP ABIDJAN
Côte d'Ivoire

ELASTIC WARSAW USER GROUP
Poland

ELASTIC ZAGREB
Croatia

ELASTICSEARCH - SOUTH AFRICA
South Africa

ELASTICSEARCH SWITZERLAND
Switzerland

ELASTICSEARCH USER GROUP PAKISTAN
Pakistan

SEARCH MEETUP MUNICH
Germany

Official Elastic Training

https://training.elastic.co

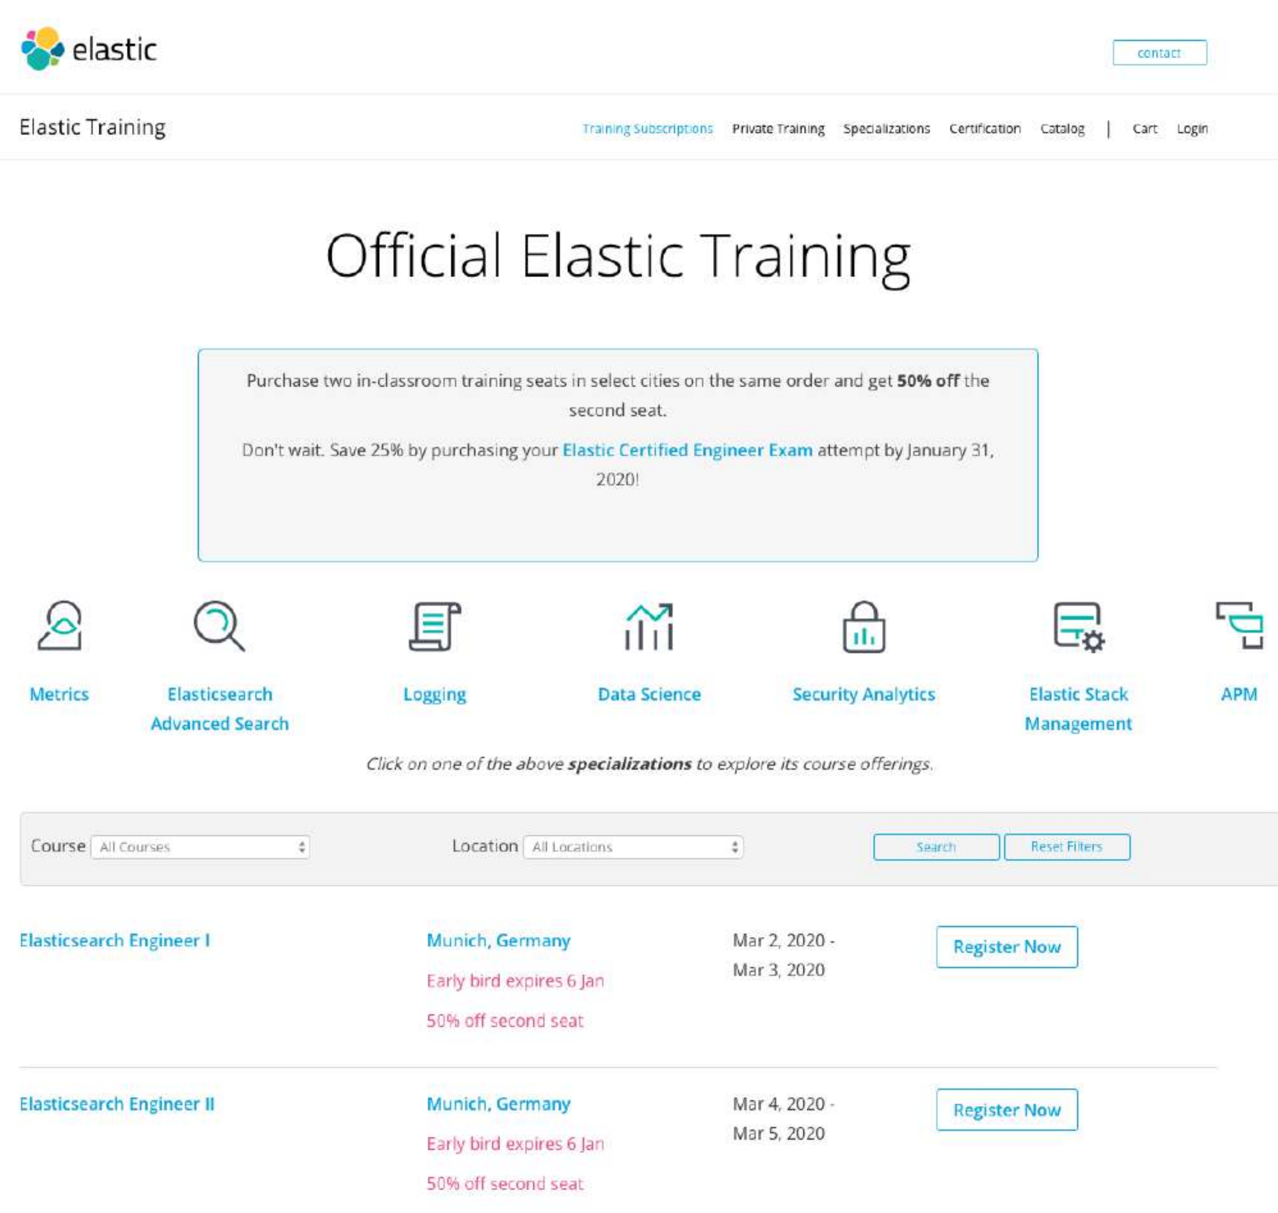

elastic

contact

Elastic Training

Training Subscriptions · Private Training · Specializations · Certification · Catalog | Cart · Login

Official Elastic Training

Purchase two in-classroom training seats in select cities on the same order and get **50% off** the second seat.

Don't wait. Save 25% by purchasing your Elastic Certified Engineer Exam attempt by January 31, 2020!

*Click on one of the above **specializations** to explore its course offerings.*

Course: All Courses · Location: All Locations · Search · Reset Filters

Elasticsearch Engineer I	Munich, Germany Early bird expires 6 Jan 50% off second seat	Mar 2, 2020 - Mar 3, 2020	Register Now
Elasticsearch Engineer II	Munich, Germany Early bird expires 6 Jan 50% off second seat	Mar 4, 2020 - Mar 5, 2020	Register Now

elastic

Thanks for listening

Q & A

Alexander Reelsen

Community Advocate

alex@elastic.co | @spinscale