El Elastic Stack en el mundo DevOps

NowBit

John Guzmán

+57 301 6998783

Elastic Engineer Advanced

john@nowbit.co

Jhon Guzmán

Ingeniero Avanzado de Elastic

john@nowbit.co
+57 301-699-8783
linkedin.com/in/jaguzmanb1/

Elastic Stack

De forma fácil y segura, recolecta datos desde cualquier origen, en cualquier formato, para que sea información buscable, analizable y visualizable en tiempo real.

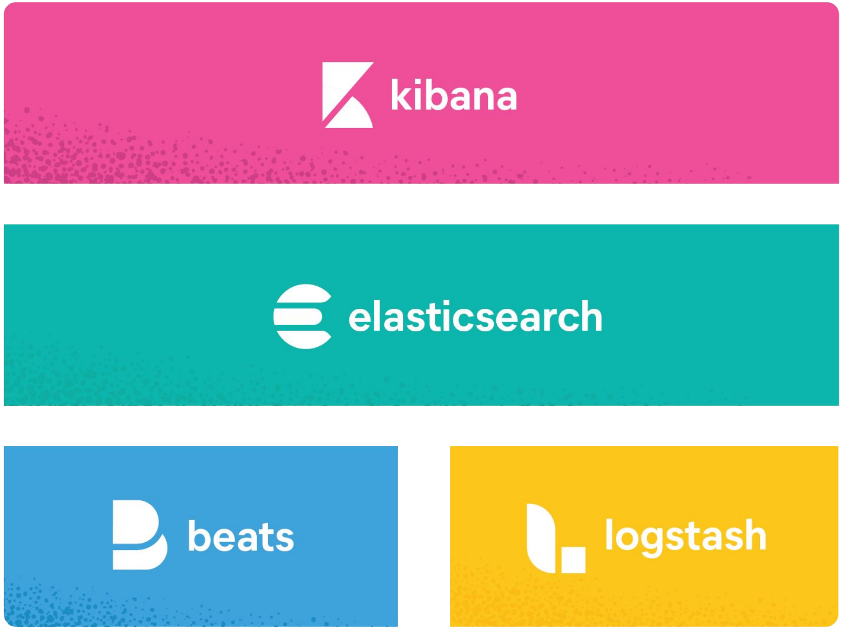

kibana

elasticsearch

beats

logstash

Tres soluciones, una misma tecnología

Elastic Enterprise Search

Elastic Observability

Elastic Security

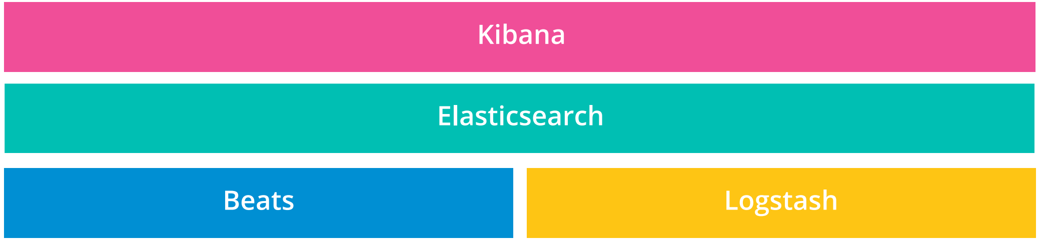

Elastic Stack

Telefonica

entel

Agilitix
beyond productivity

Azure

Red Aval

Banco de Bogotá

GoDaddy

P&G

BLIZZARD
ENTERTAINMENT

CenturyLink

Observabilidad

Rides

Logs

Metrics

APM

Uptime

elastic

https://www.elastic.co/customers/success-stories?usecase=observability

Infraestructuras...

- En la nube.
- Entornos on premise.
- Contenedores.

- Microservicios.
- Monolitos.
- SOA.

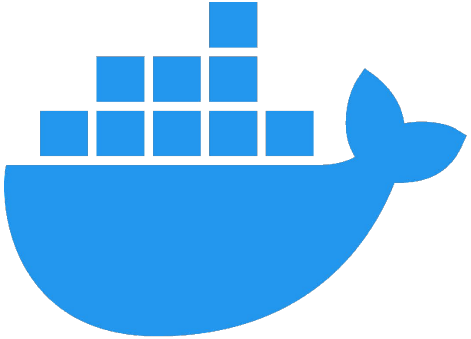

Tipos de información.

Se puede generar todo tipo de información

- Logs de aplicaciones.
- Logs de servidores.
- Información de tráfico de paquetes de red.
- Métricas de bases de datos o aplicaciones.
- Información de negocio.

Beats

Agentes de datos ligeros.

- Código abierto.
- Envían datos de cientos o miles de máquinas y sistemas a Logstash o Elasticsearch.
- Cientos de aportes de la comunidad.
- Integración con Kibana.

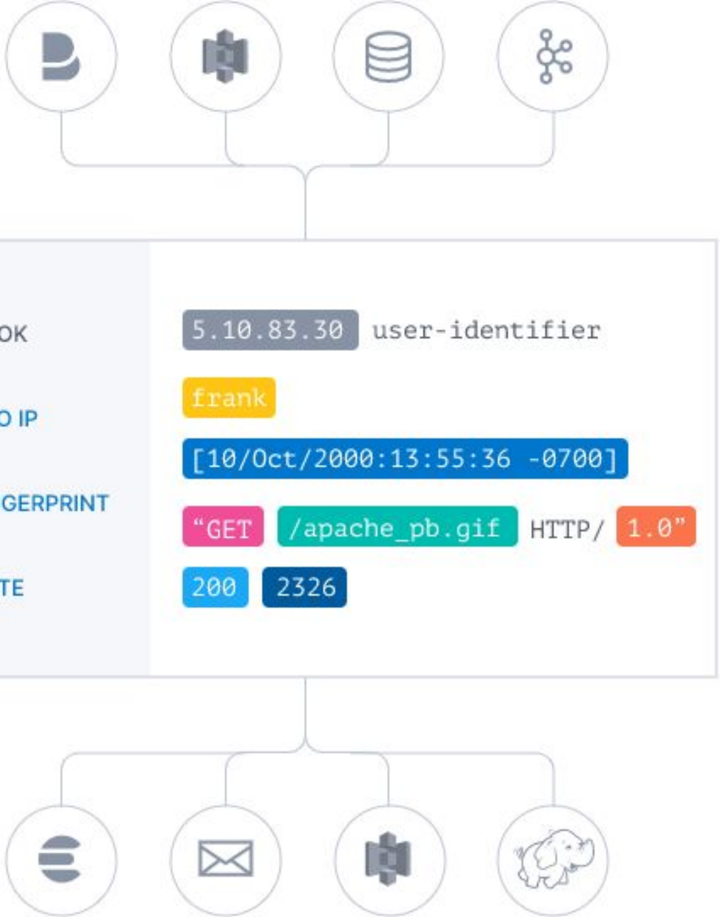

GROK

GEO IP

FINGERPRINT

DATE

```
5.10.83.30 user-identifier
frank
[10/Oct/2000:13:55:36 -0700]
“GET /apache_pb.gif HTTP/ 1.0”
200 2326
```

Logstash (ETL)

Pipeline de procesamiento de datos.

- Código abierto.
- Compatible con un montón de orígenes de información.
- Transforma e ingesta información al gusto.
- Ingesta información en múltiples destinos a la vez.
- Enriquece información.

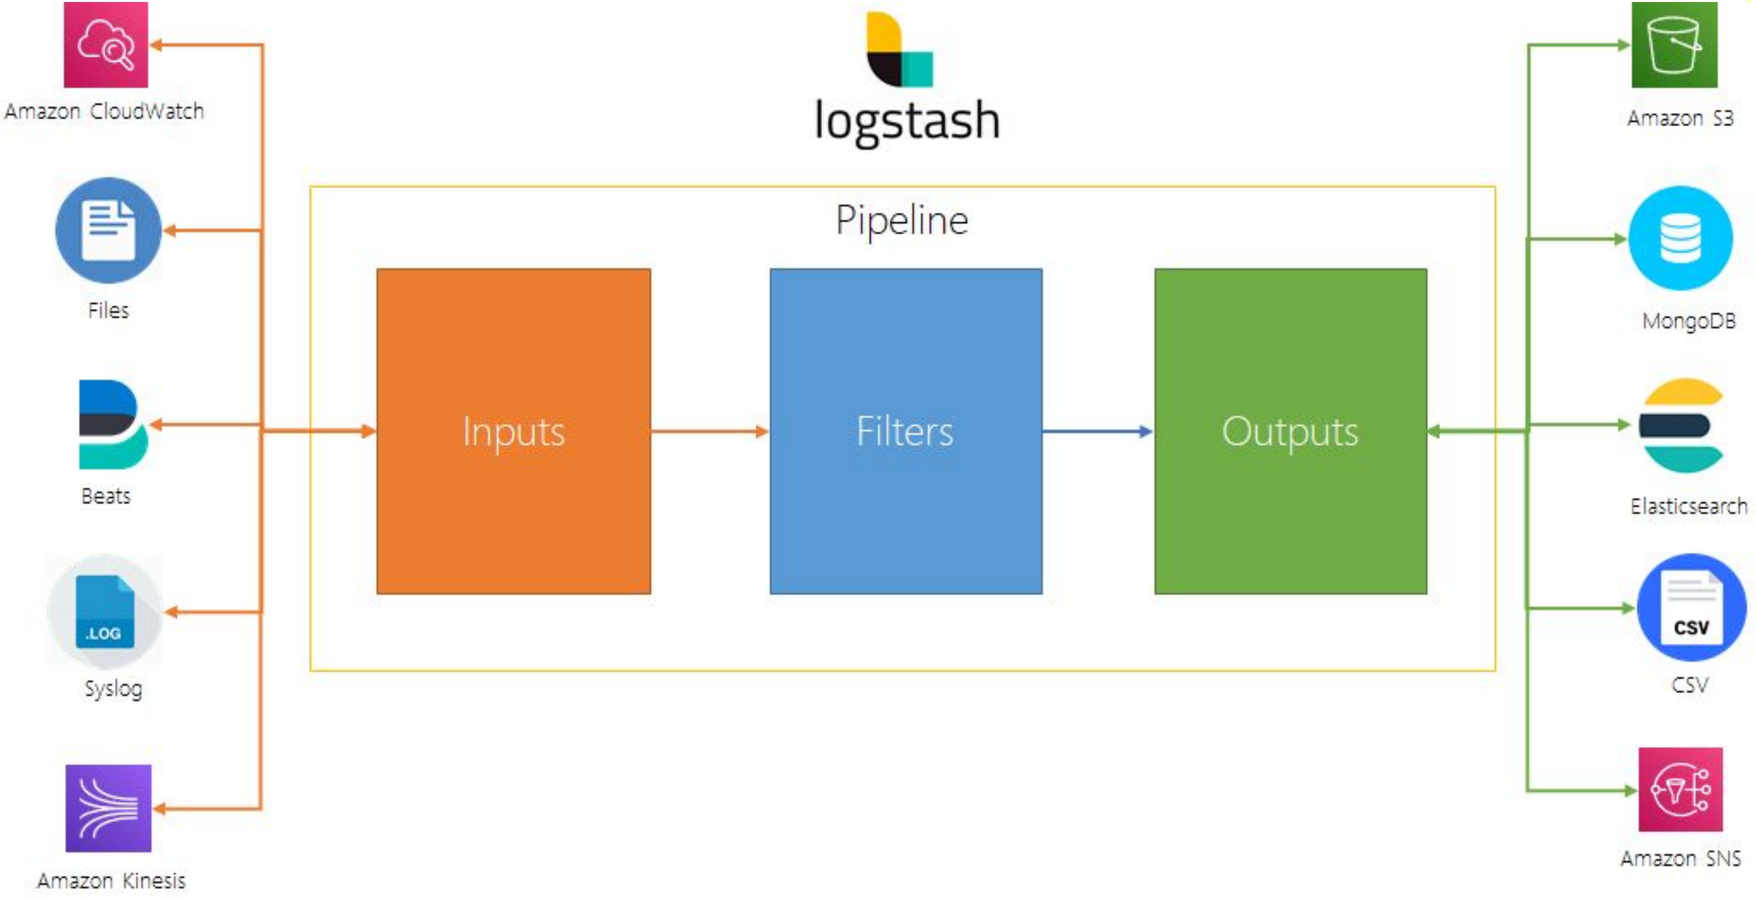
logstash
Pipeline
Inputs
Filters
Outputs
Amazon CloudWatch
Files
Beats
Syslog
.LOG
Amazon Kinesis
Amazon S3
MongoDB
Elasticsearch
CSV
CSV
Amazon SNS

elastic

Logstash (ETL)

Pipeline de procesamiento de datos.

- Azure event hubs
- Beats
- CloudWatch
- Couch DB
- Elasticsearch
- Exec
- File
- HTTP
- jdbc
- Kafka
- Kinesis
- Log4j
- Meetup
- Rabbitmq
- Redis
- S3
- Salesforce
- Twitter

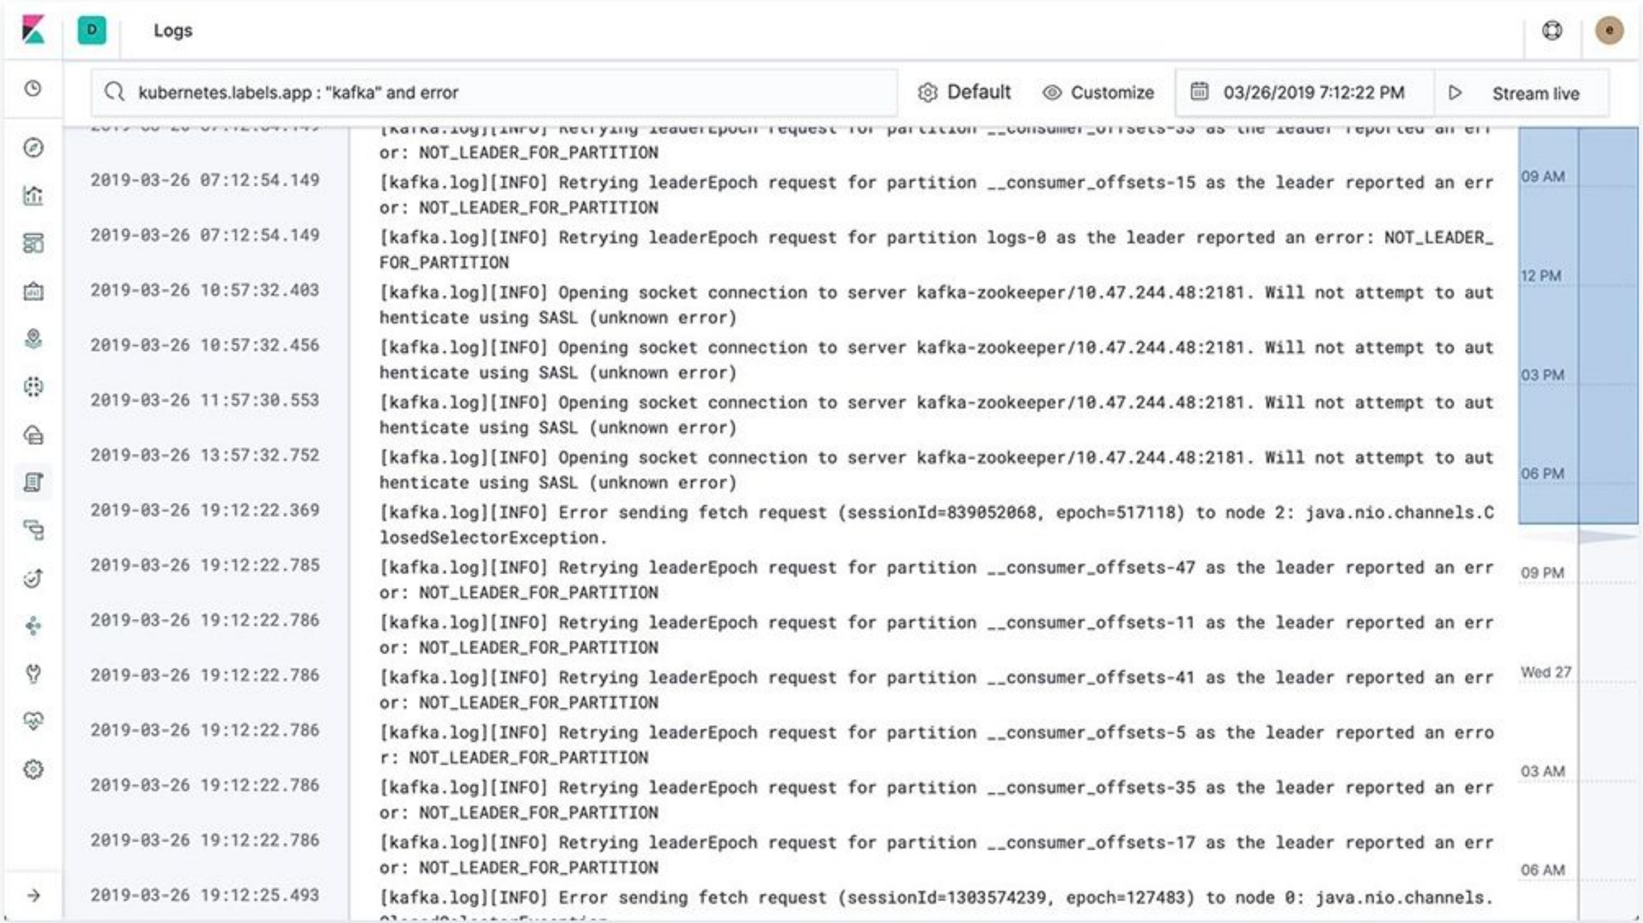

kubernetes.labels.app : "kafka" and error

03/26/2019 7:12:22 PM

```
2019-03-26 07:12:54.149   [kafka.log][INFO] Retrying leaderEpoch request for partition __consumer_offsets-33 as the leader reported an err
                          or: NOT_LEADER_FOR_PARTITION
2019-03-26 07:12:54.149   [kafka.log][INFO] Retrying leaderEpoch request for partition __consumer_offsets-15 as the leader reported an err
                          or: NOT_LEADER_FOR_PARTITION
2019-03-26 07:12:54.149   [kafka.log][INFO] Retrying leaderEpoch request for partition logs-0 as the leader reported an error: NOT_LEADER_
                          FOR_PARTITION
2019-03-26 10:57:32.403   [kafka.log][INFO] Opening socket connection to server kafka-zookeeper/10.47.244.48:2181. Will not attempt to aut
                          henticate using SASL (unknown error)
2019-03-26 10:57:32.456   [kafka.log][INFO] Opening socket connection to server kafka-zookeeper/10.47.244.48:2181. Will not attempt to aut
                          henticate using SASL (unknown error)
2019-03-26 11:57:30.553   [kafka.log][INFO] Opening socket connection to server kafka-zookeeper/10.47.244.48:2181. Will not attempt to aut
                          henticate using SASL (unknown error)
2019-03-26 13:57:32.752   [kafka.log][INFO] Opening socket connection to server kafka-zookeeper/10.47.244.48:2181. Will not attempt to aut
                          henticate using SASL (unknown error)
2019-03-26 19:12:22.369   [kafka.log][INFO] Error sending fetch request (sessionId=839052068, epoch=517118) to node 2: java.nio.channels.C
                          losedSelectorException.
2019-03-26 19:12:22.785   [kafka.log][INFO] Retrying leaderEpoch request for partition __consumer_offsets-47 as the leader reported an err
                          or: NOT_LEADER_FOR_PARTITION
2019-03-26 19:12:22.786   [kafka.log][INFO] Retrying leaderEpoch request for partition __consumer_offsets-11 as the leader reported an err
                          or: NOT_LEADER_FOR_PARTITION
2019-03-26 19:12:22.786   [kafka.log][INFO] Retrying leaderEpoch request for partition __consumer_offsets-41 as the leader reported an err
                          or: NOT_LEADER_FOR_PARTITION
2019-03-26 19:12:22.786   [kafka.log][INFO] Retrying leaderEpoch request for partition __consumer_offsets-5 as the leader reported an erro
                          r: NOT_LEADER_FOR_PARTITION
2019-03-26 19:12:22.786   [kafka.log][INFO] Retrying leaderEpoch request for partition __consumer_offsets-35 as the leader reported an err
                          or: NOT_LEADER_FOR_PARTITION
2019-03-26 19:12:22.786   [kafka.log][INFO] Retrying leaderEpoch request for partition __consumer_offsets-17 as the leader reported an err
                          or: NOT_LEADER_FOR_PARTITION
2019-03-26 19:12:25.493   [kafka.log][INFO] Error sending fetch request (sessionId=1303574239, epoch=127483) to node 0: java.nio.channels.
                          ClosedSelectorException
```

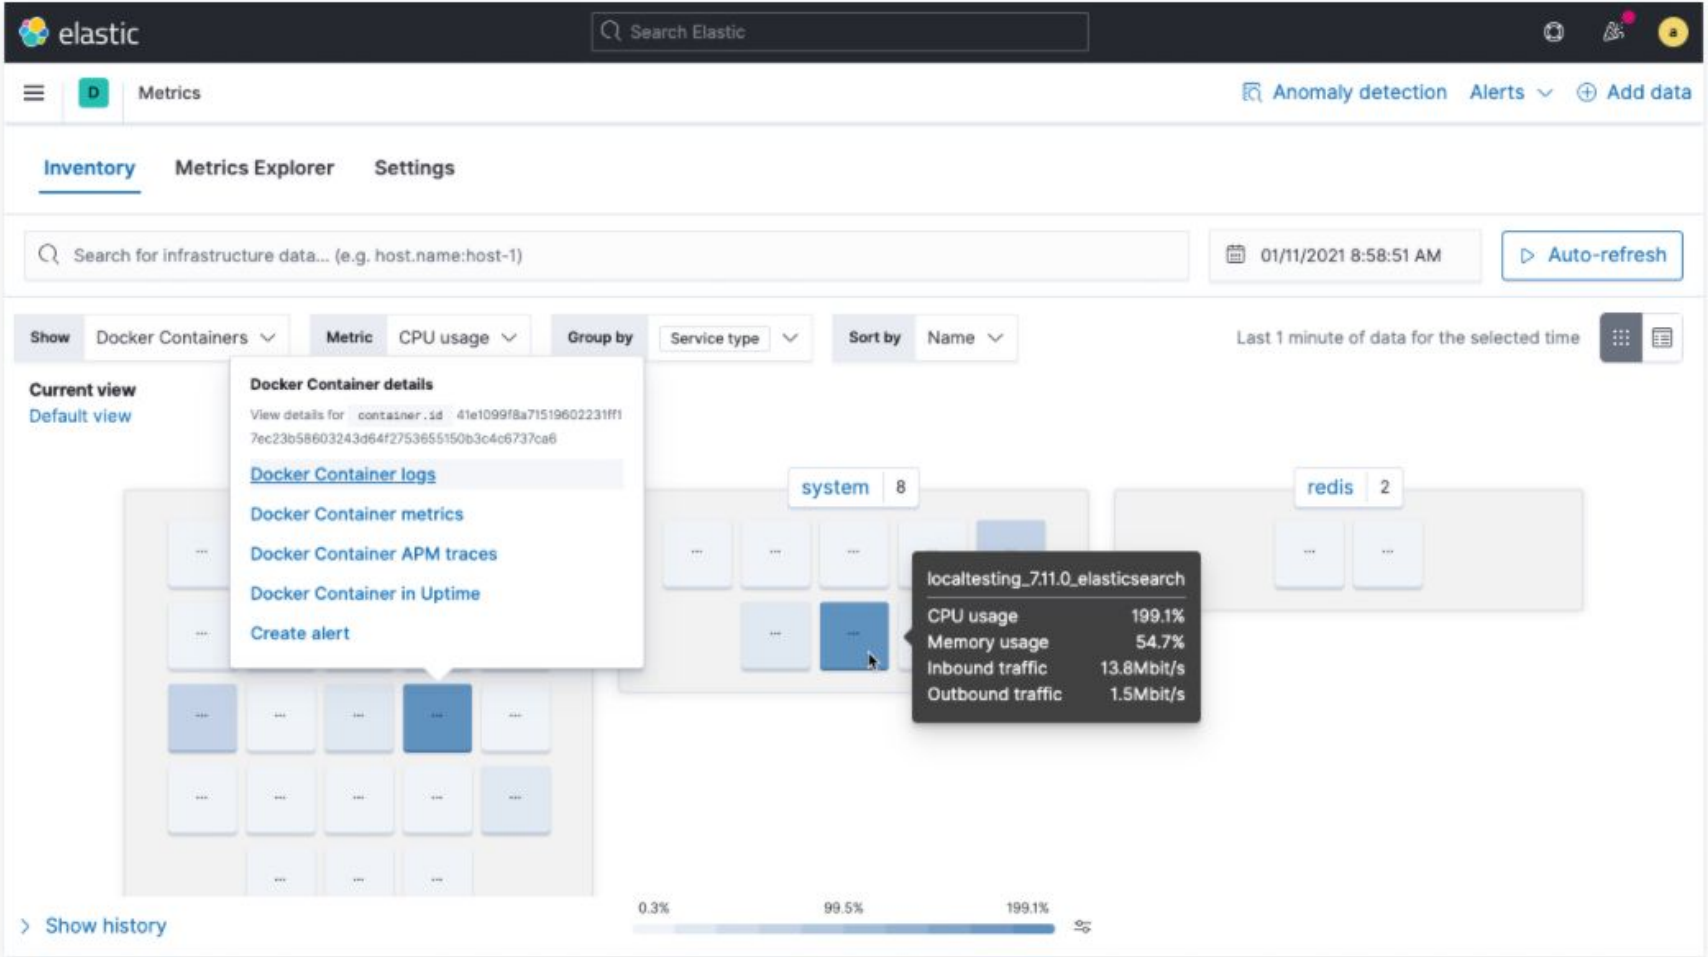

Search for infrastructure data... (e.g. host.name:host-1)

01/11/2021 8:58:51 AM

Show: Docker Containers

Metric: CPU usage

Group by: Service type

Sort by: Name

Last 1 minute of data for the selected time

Docker Container details

View details for `container.id` 41e1099f8a71519602231ff17ec23b58603243d64f2753655150b3c4c6737ca6

- Docker Container logs
- Docker Container metrics
- Docker Container APM traces
- Docker Container in Uptime
- Create alert

localtesting_7.11.0_elasticsearch

CPU usage	199.1%
Memory usage	54.7%
Inbound traffic	13.8Mbit/s
Outbound traffic	1.5Mbit/s

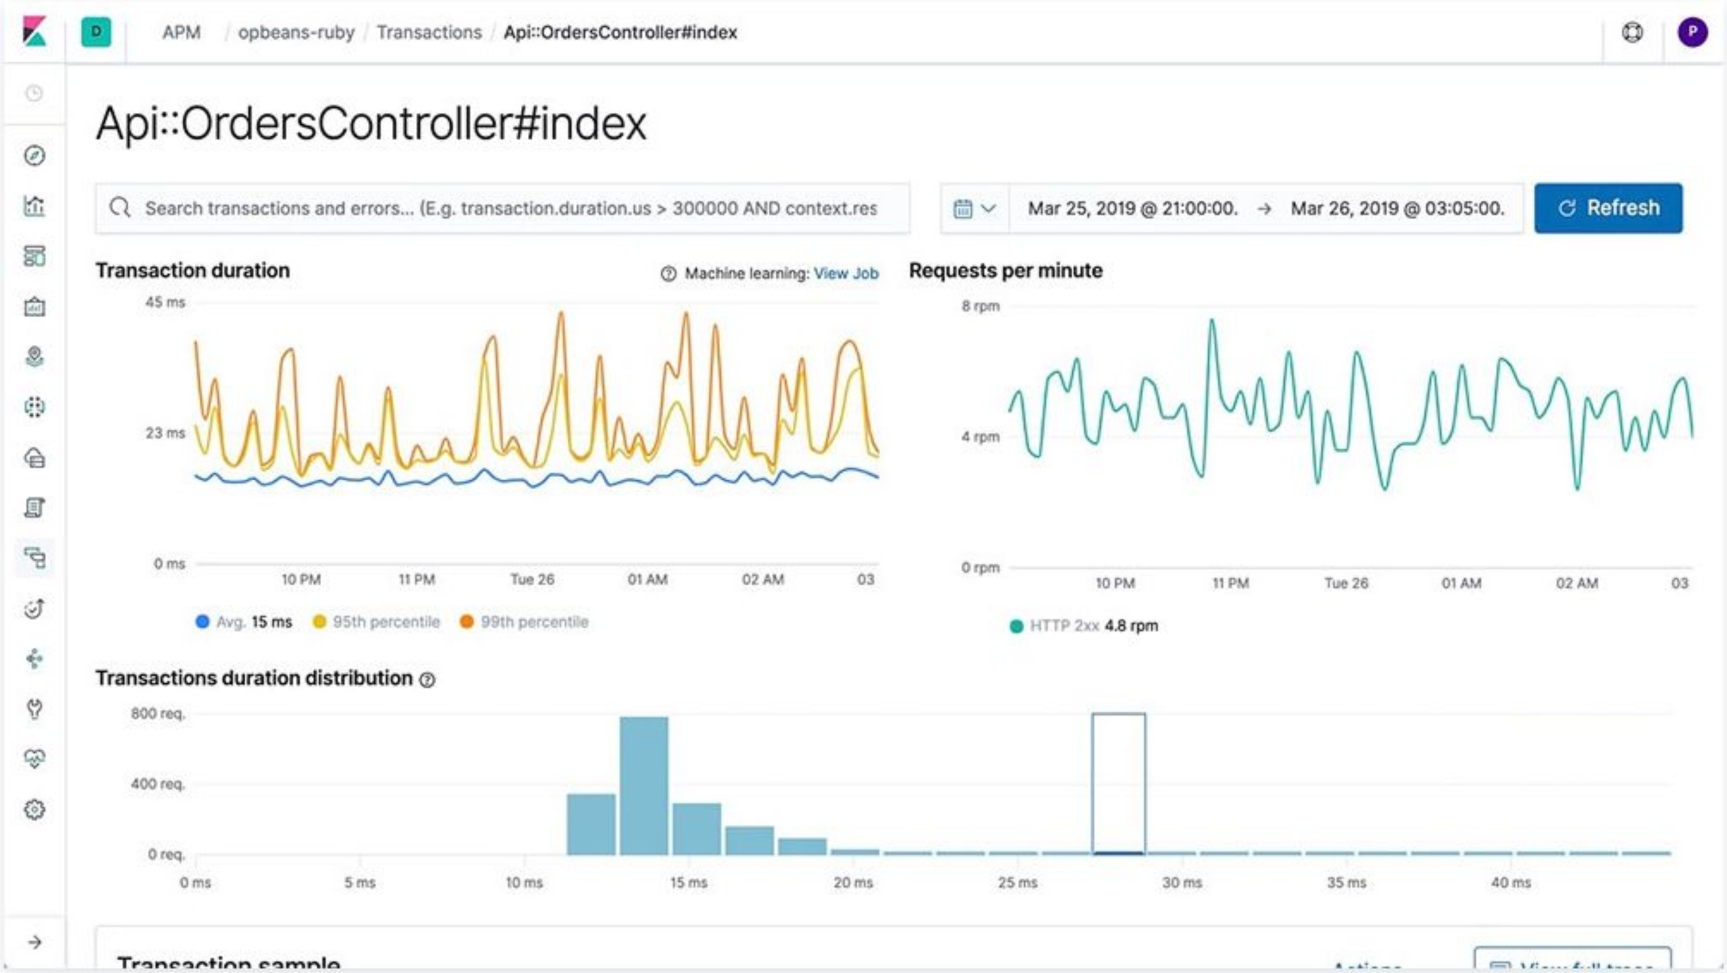

APM / opbeans-ruby / Transactions / Api::OrdersController#index

Api::OrdersController#index

Search transactions and errors... (E.g. transaction.duration.us > 300000 AND context.res

Mar 25, 2019 @ 21:00:00. → Mar 26, 2019 @ 03:05:00.

Refresh

Transaction duration

Machine learning: View Job

Requests per minute

Transactions duration distribution

Transaction sample

Actions

View full trace

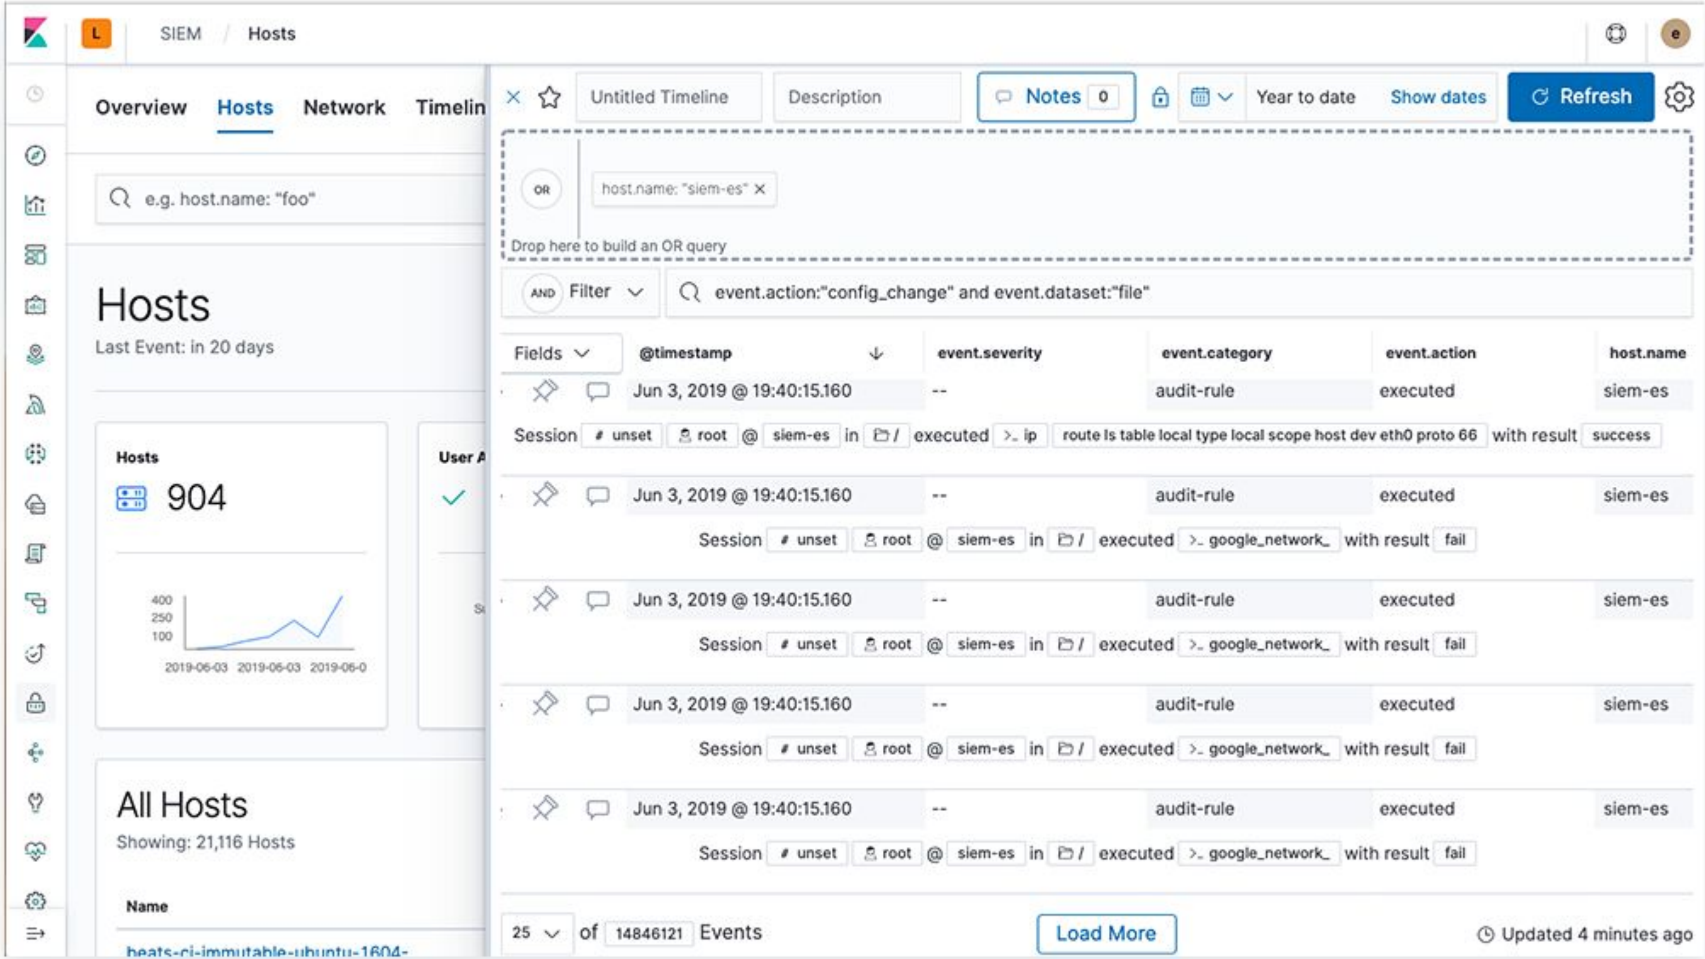

e.g. host.name: "foo"

host.name: "siem-es"

event.action:"config_change" and event.dataset:"file"

@timestamp	event.severity	event.category	event.action	host.name
Jun 3, 2019 @ 19:40:15.160	--	audit-rule	executed	siem-es

Session unset root @ siem-es in / executed ip route ls table local type local scope host dev eth0 proto 66 with result success

@timestamp	event.severity	event.category	event.action	host.name
Jun 3, 2019 @ 19:40:15.160	--	audit-rule	executed	siem-es

Session unset root @ siem-es in / executed google_network_ with result fail

@timestamp	event.severity	event.category	event.action	host.name
Jun 3, 2019 @ 19:40:15.160	--	audit-rule	executed	siem-es

Session unset root @ siem-es in / executed google_network_ with result fail

@timestamp	event.severity	event.category	event.action	host.name
Jun 3, 2019 @ 19:40:15.160	--	audit-rule	executed	siem-es

Session unset root @ siem-es in / executed google_network_ with result fail

@timestamp	event.severity	event.category	event.action	host.name
Jun 3, 2019 @ 19:40:15.160	--	audit-rule	executed	siem-es

Session unset root @ siem-es in / executed google_network_ with result fail

25 of 14846121 Events

Elasticsearch

El corazón del Elastic Stack, gratuito y abierto.

- Código abierto.
- Motor de búsqueda y analítica RESTful.
- Operaciones en tiempo real.
- Escalable y distribuido.
- Útil para un montón de casos de uso.

ESCALA

Distribuido por diseño

VELOCIDAD

Encuentra información en milisegundos

RELEVANCIA

Obtén resultados con alta relevancia

Let's analyze the web logs by country…

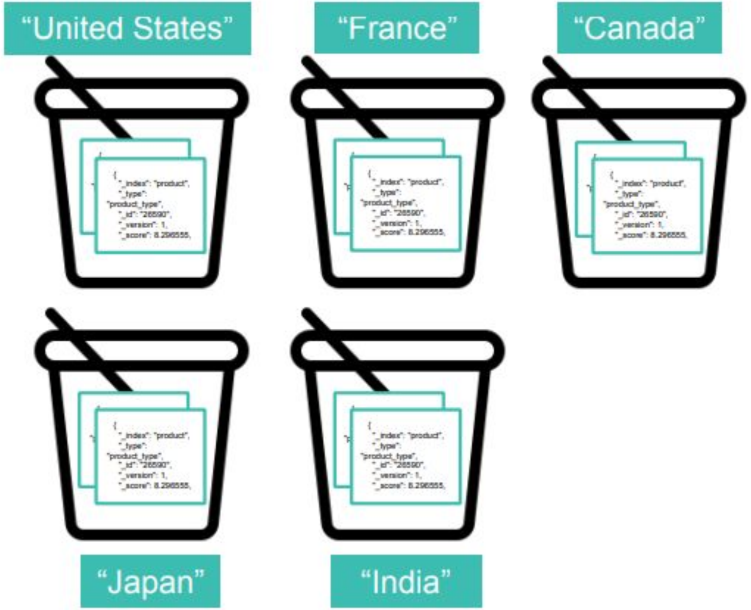

elastic

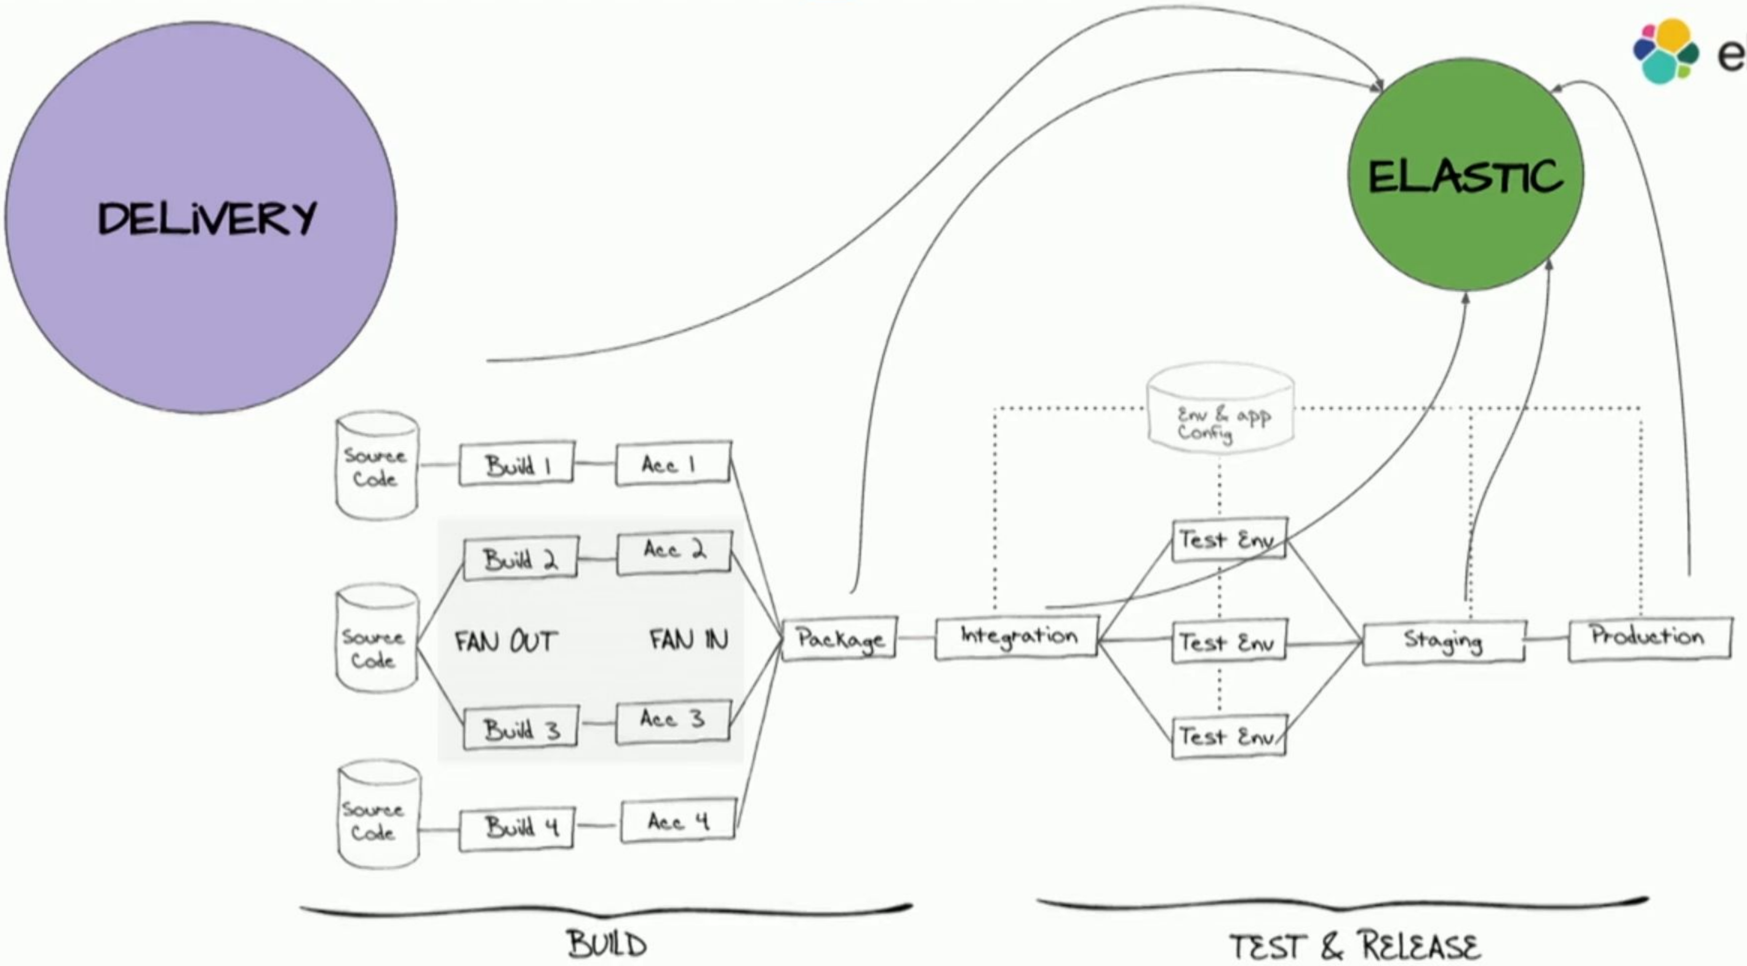

Image courtesy of GoCD

@WalmartLabs

AVALON

Kibana

La ventana al Elastic Stack.

- Herramienta de administración del Elastic Stack.
- Visualizaciones y dashboards en tiempo real.
- Escalable y distribuido.
- Código abierto.
- Útil para un montón de casos de uso.

NowBit

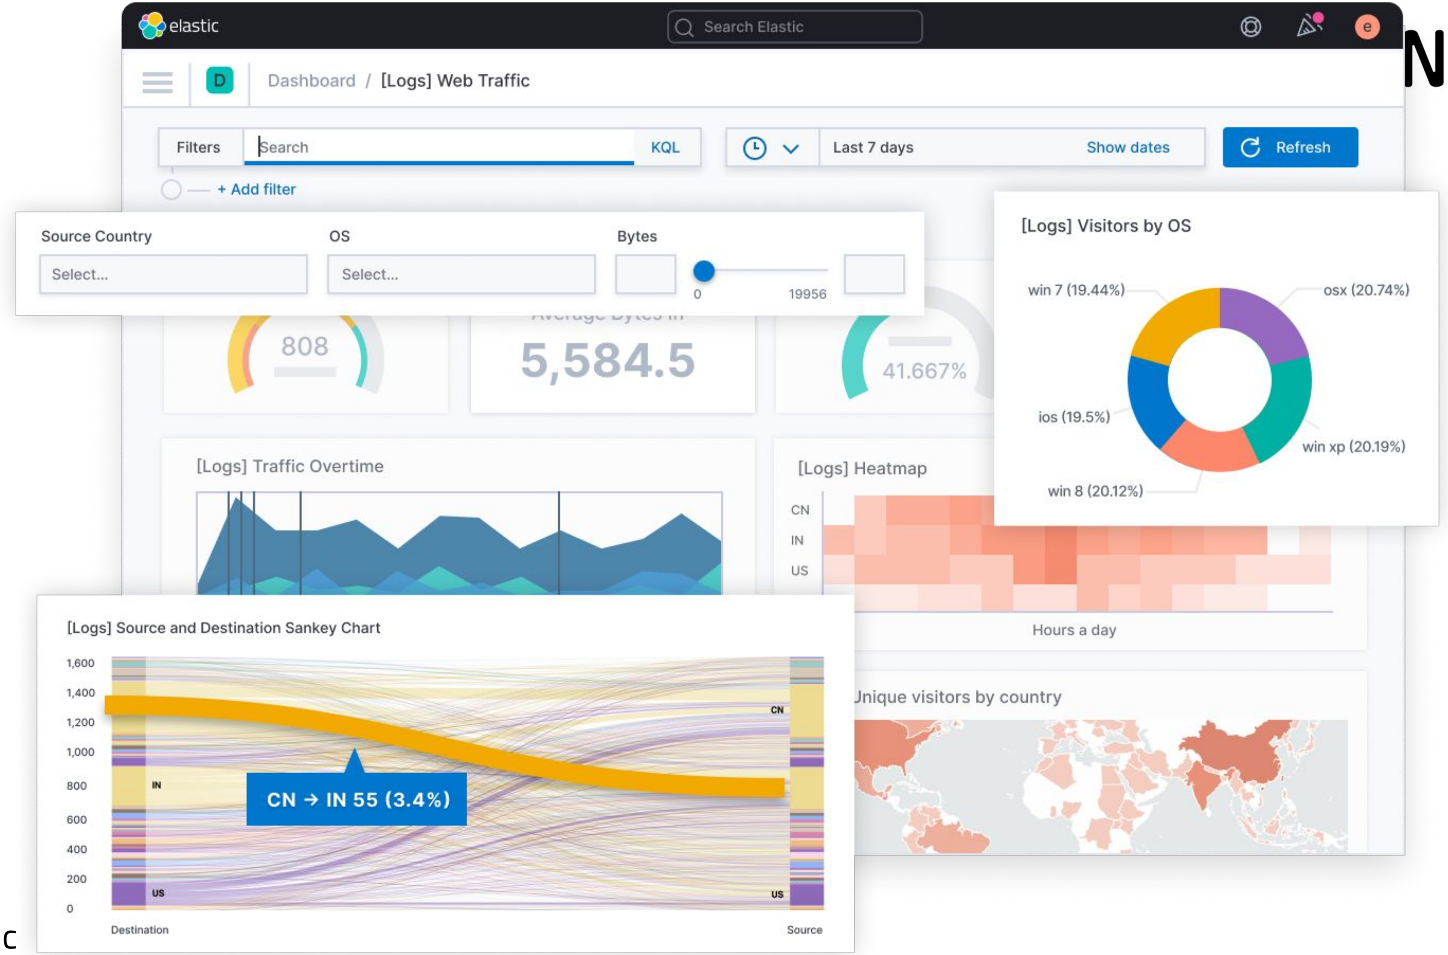

elastic

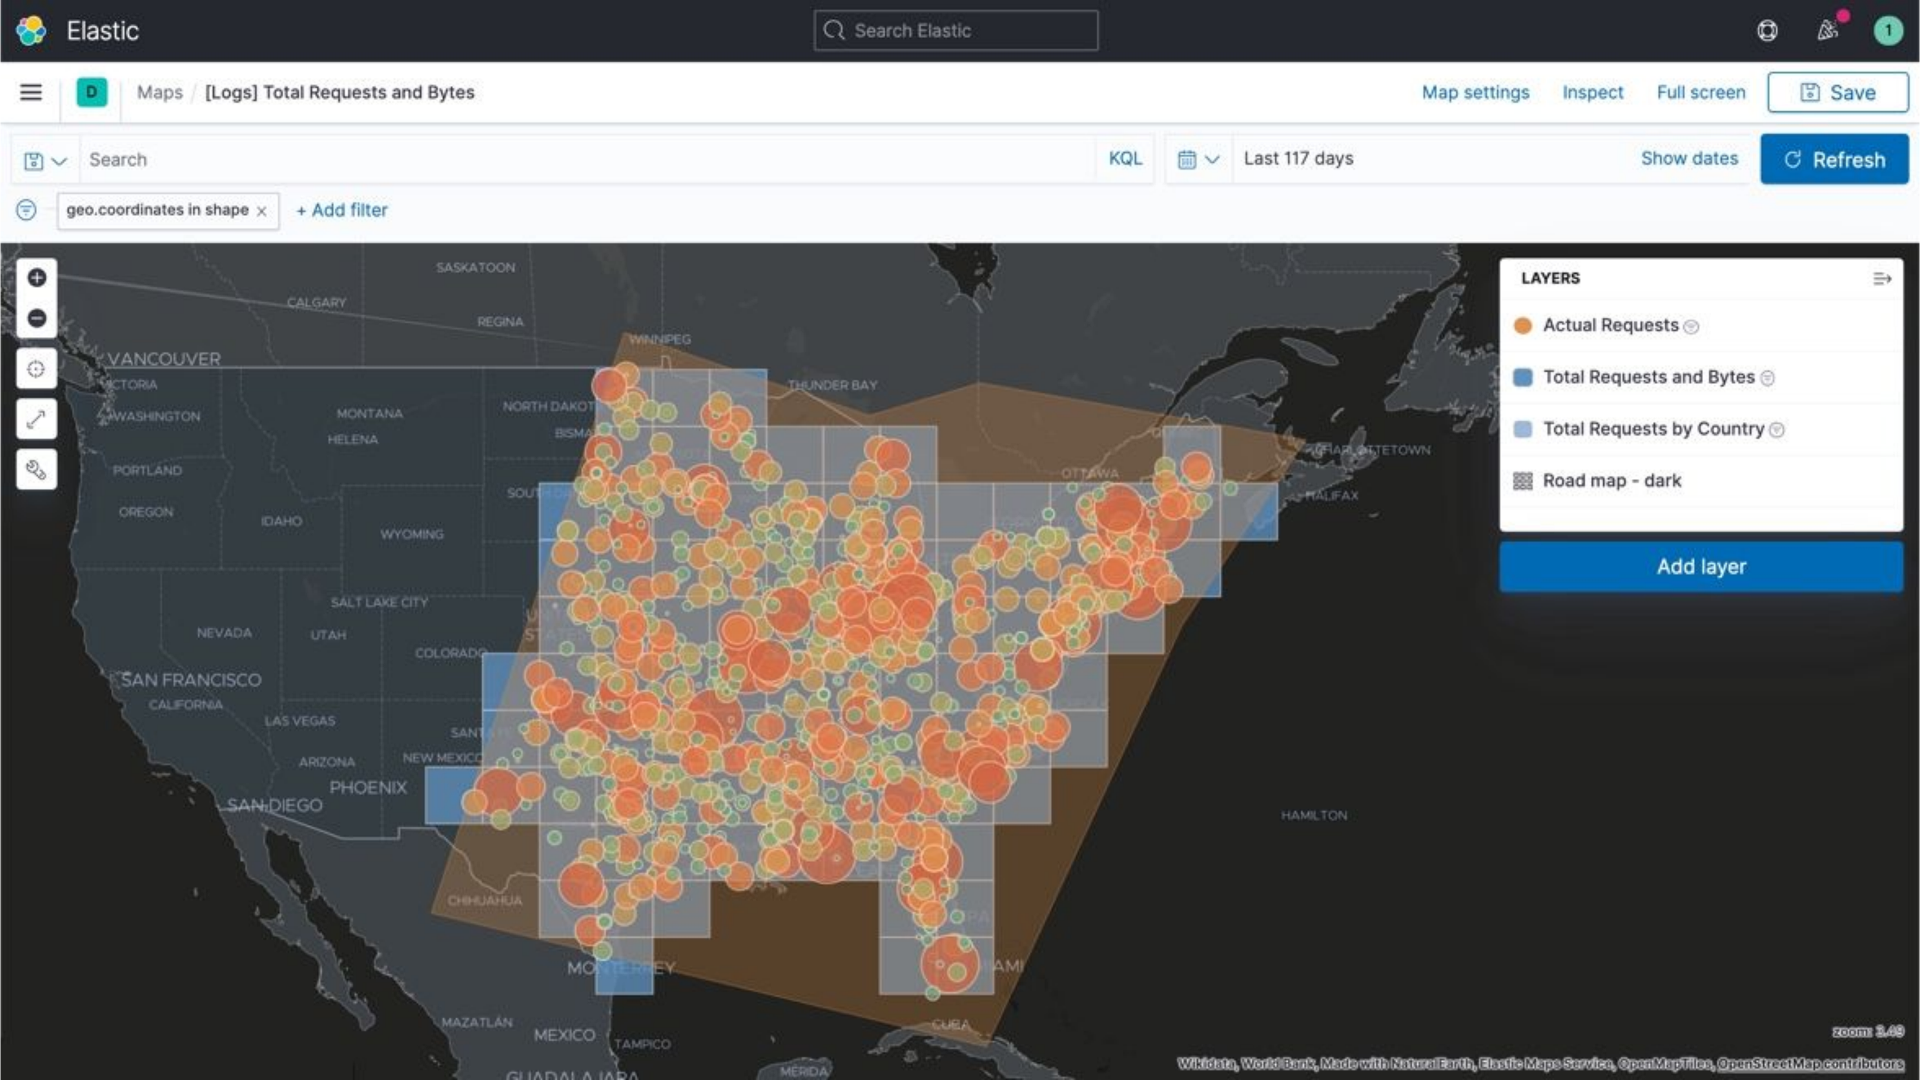
Elastic
Search Elastic
Maps / [Logs] Total Requests and Bytes
Map settings
Inspect
Full screen
Save
Search
KQL
Last 117 days
Show dates
Refresh
geo.coordinates in shape ×
+ Add filter
LAYERS
● Actual Requests
■ Total Requests and Bytes
■ Total Requests by Country
▦ Road map - dark
Add layer
SASKATOON
CALGARY
REGINA
WINNIPEG
VANCOUVER
VICTORIA
WASHINGTON
MONTANA
HELENA
NORTH DAKOTA
BISMARCK
THUNDER BAY
PORTLAND
OREGON
IDAHO
WYOMING
SOUTH DAKOTA
OTTAWA
CHARLOTTETOWN
HALIFAX
SALT LAKE CITY
NEVADA
UTAH
COLORADO
SAN FRANCISCO
CALIFORNIA
LAS VEGAS
SANTA FE
ARIZONA
NEW MEXICO
PHOENIX
SAN DIEGO
CHIHUAHUA
MONTERREY
MIAMI
HAMILTON
MAZATLÁN
MEXICO
TAMPICO
CUBA
MÉRIDA
GUADALAJARA
zoom: 3.49
Wikidata, World Bank, Made with NaturalEarth, Elastic Maps Service, OpenMapTiles, OpenStreetMap contributors

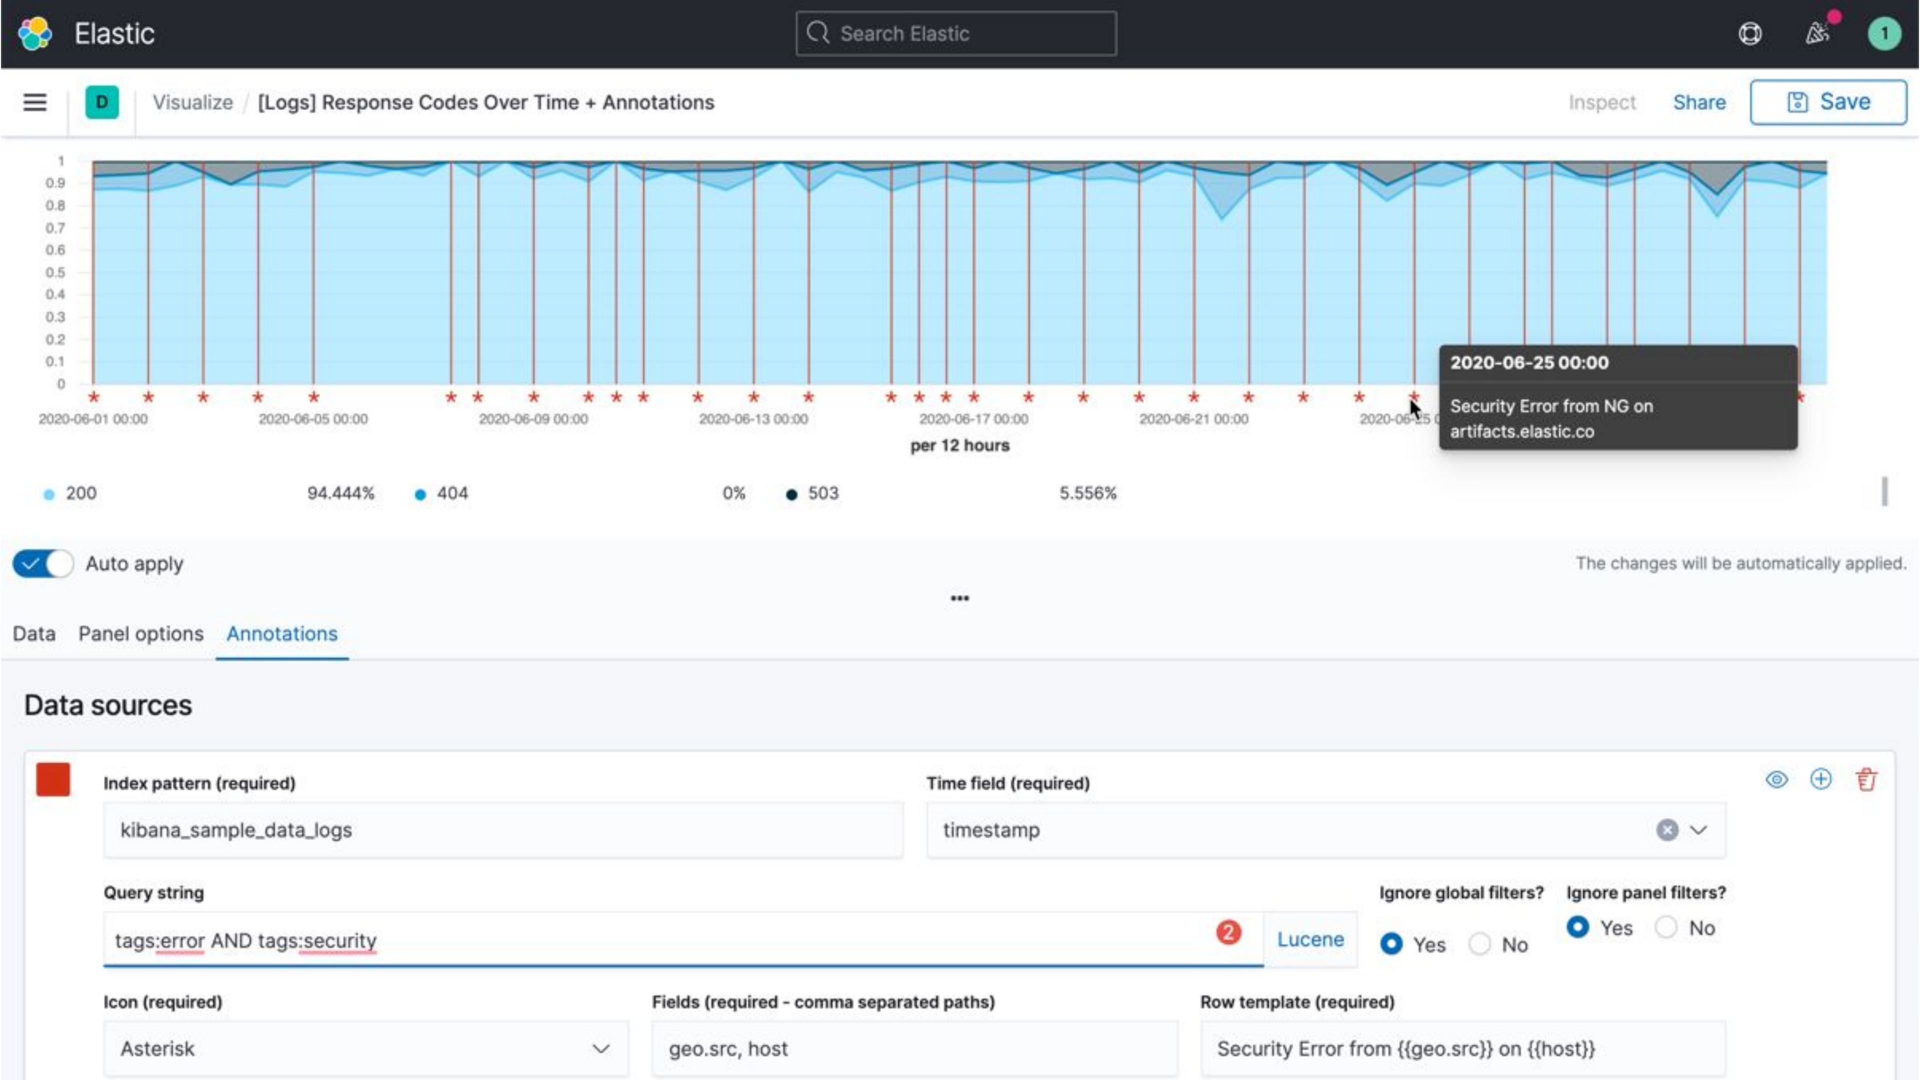

Visualize / [Logs] Response Codes Over Time + Annotations

2020-06-25 00:00

Security Error from NG on artifacts.elastic.co

● 200	94.444%	● 404	0%	● 503	5.556%

Auto apply

The changes will be automatically applied.

Data sources

Index pattern (required)
kibana_sample_data_logs

Time field (required)
timestamp

Query string
tags:error AND tags:security

Ignore global filters?
● Yes ○ No

Ignore panel filters?
● Yes ○ No

Icon (required)
Asterisk

Fields (required - comma separated paths)
geo.src, host

Row template (required)
Security Error from {{geo.src}} on {{host}}

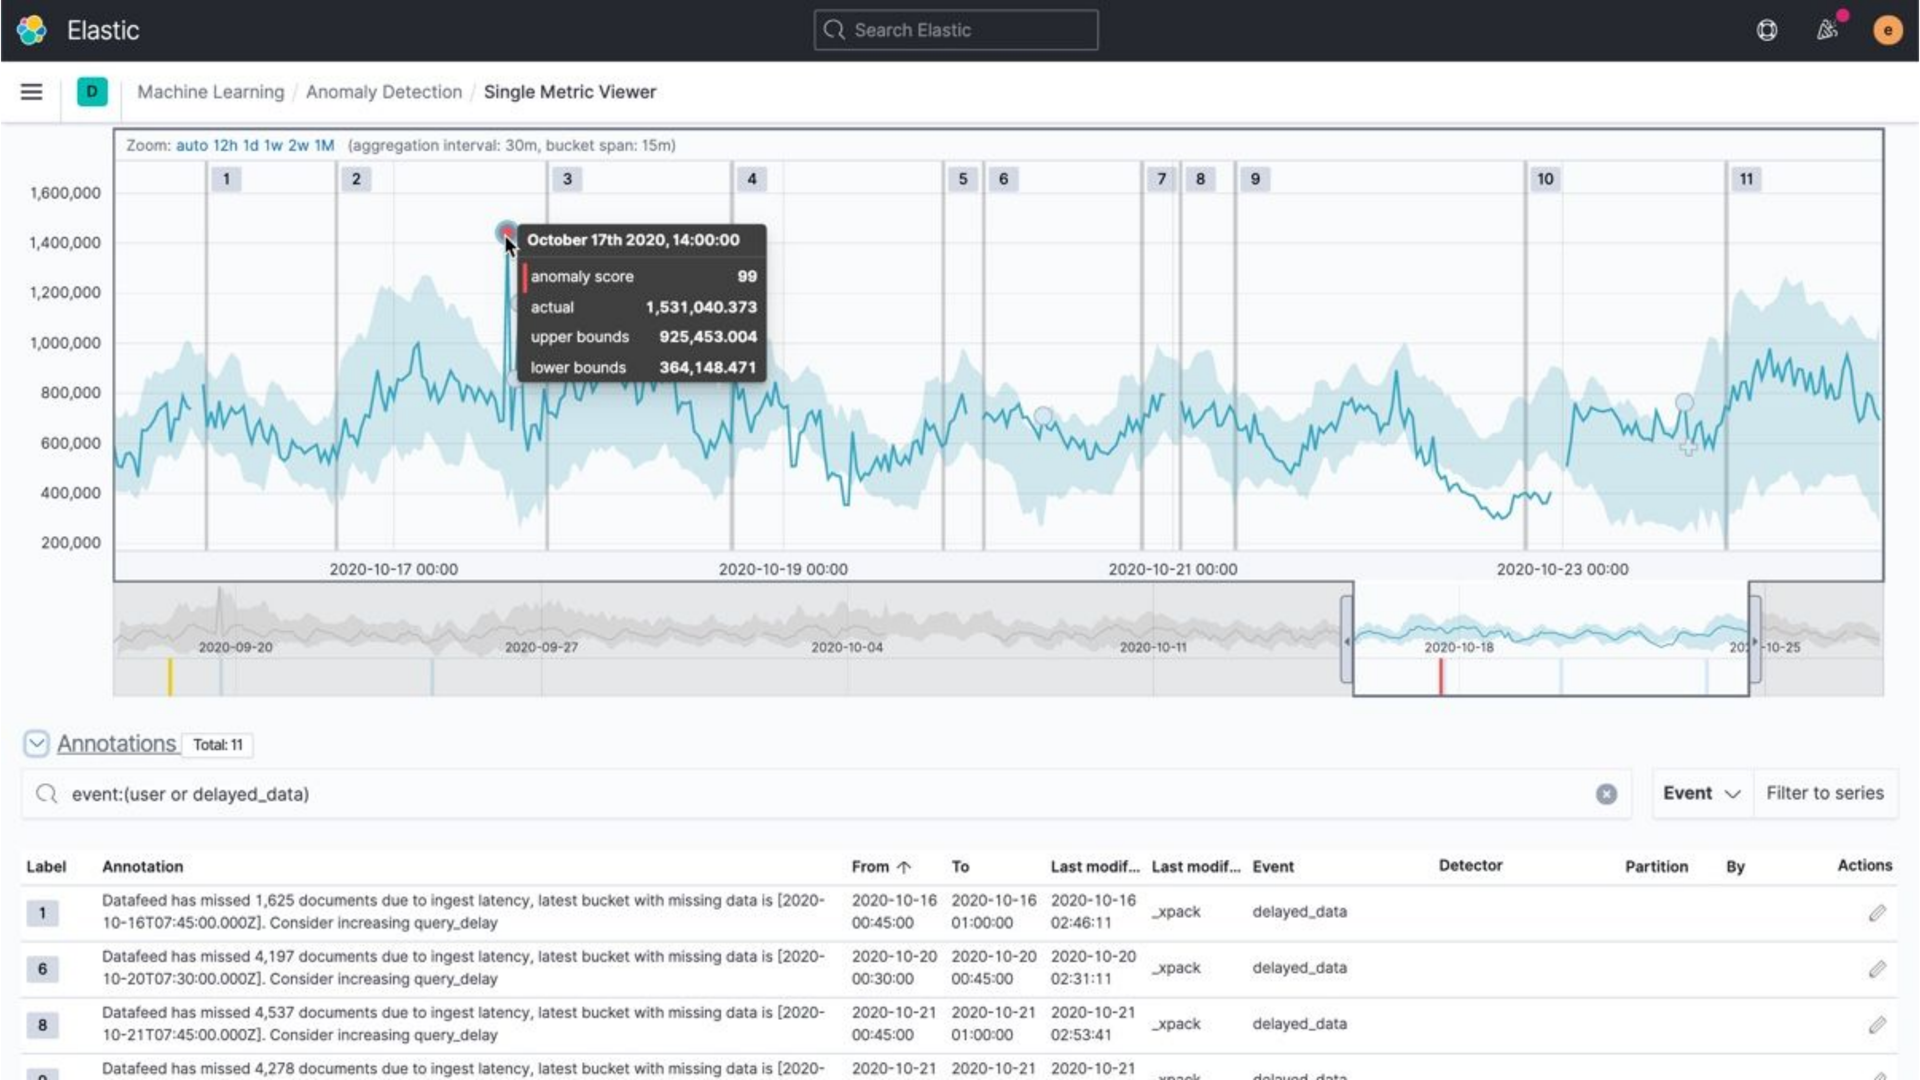

October 17th 2020, 14:00:00

anomaly score	99
actual	1,531,040.373
upper bounds	925,453.004
lower bounds	364,148.471

Annotations Total: 11

event:(user or delayed_data)

Label	Annotation	From ↑	To	Last modif...	Last modif...	Event	Detector	Partition	By	Actions
1	Datafeed has missed 1,625 documents due to ingest latency, latest bucket with missing data is [2020-10-16T07:45:00.000Z]. Consider increasing query_delay	2020-10-16 00:45:00	2020-10-16 01:00:00	2020-10-16 02:46:11	_xpack	delayed_data				
6	Datafeed has missed 4,197 documents due to ingest latency, latest bucket with missing data is [2020-10-20T07:30:00.000Z]. Consider increasing query_delay	2020-10-20 00:30:00	2020-10-20 00:45:00	2020-10-20 02:31:11	_xpack	delayed_data				
8	Datafeed has missed 4,537 documents due to ingest latency, latest bucket with missing data is [2020-10-21T07:45:00.000Z]. Consider increasing query_delay	2020-10-21 00:45:00	2020-10-21 01:00:00	2020-10-21 02:53:41	_xpack	delayed_data				
9	Datafeed has missed 4,278 documents due to ingest latency, latest bucket with missing data is [2020-	2020-10-21	2020-10-21	2020-10-21	_xpack	delayed_data				

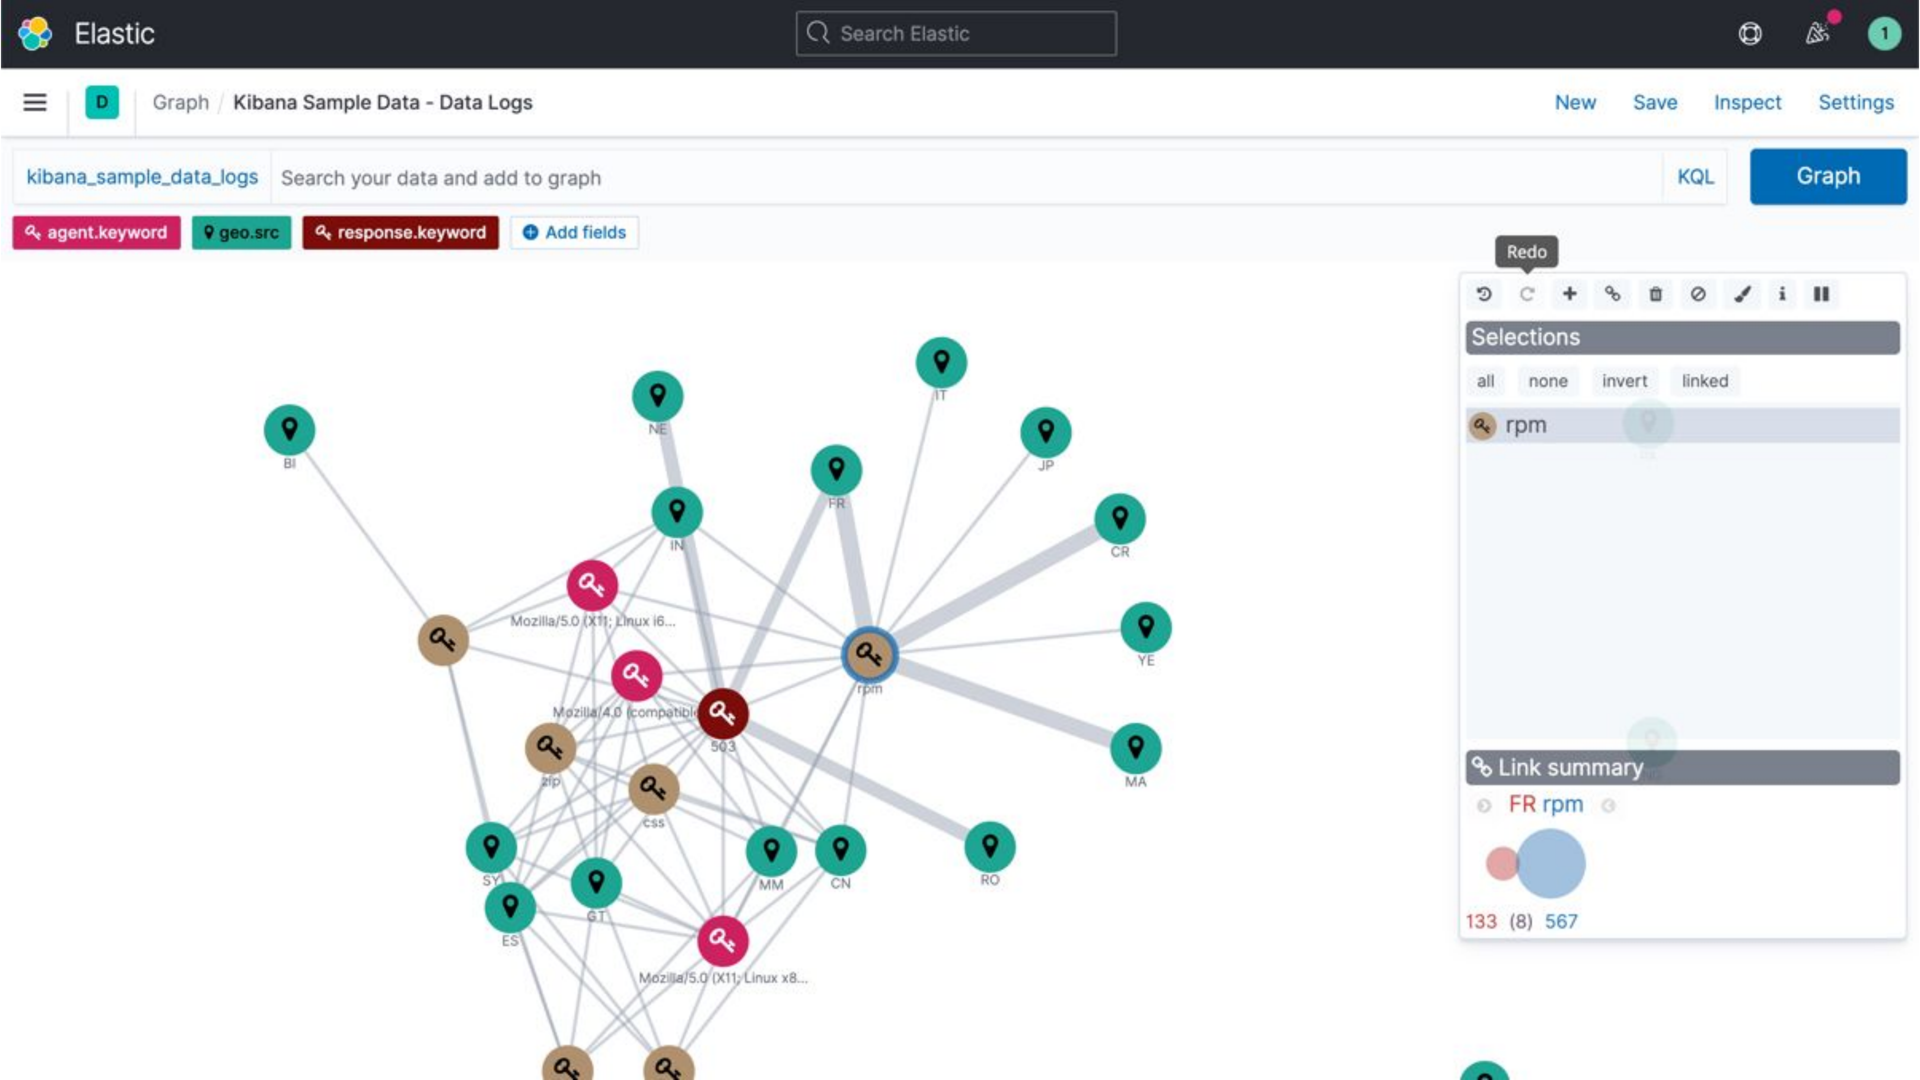
Elastic
Search Elastic
1
D
Graph / Kibana Sample Data - Data Logs
New
Save
Inspect
Settings
kibana_sample_data_logs
Search your data and add to graph
KQL
Graph
agent.keyword
geo.src
response.keyword
Add fields
Redo
Selections
all
none
invert
linked
rpm
Link summary
FR rpm
133 (8) 567
BI
NE
IT
JP
FR
IN
CR
YE
MA
RO
rpm
503
zip
css
SY
GT
ES
MM
CN
Mozilla/5.0 (X11; Linux i6...
Mozilla/4.0 (compatible
Mozilla/5.0 (X11; Linux x8...

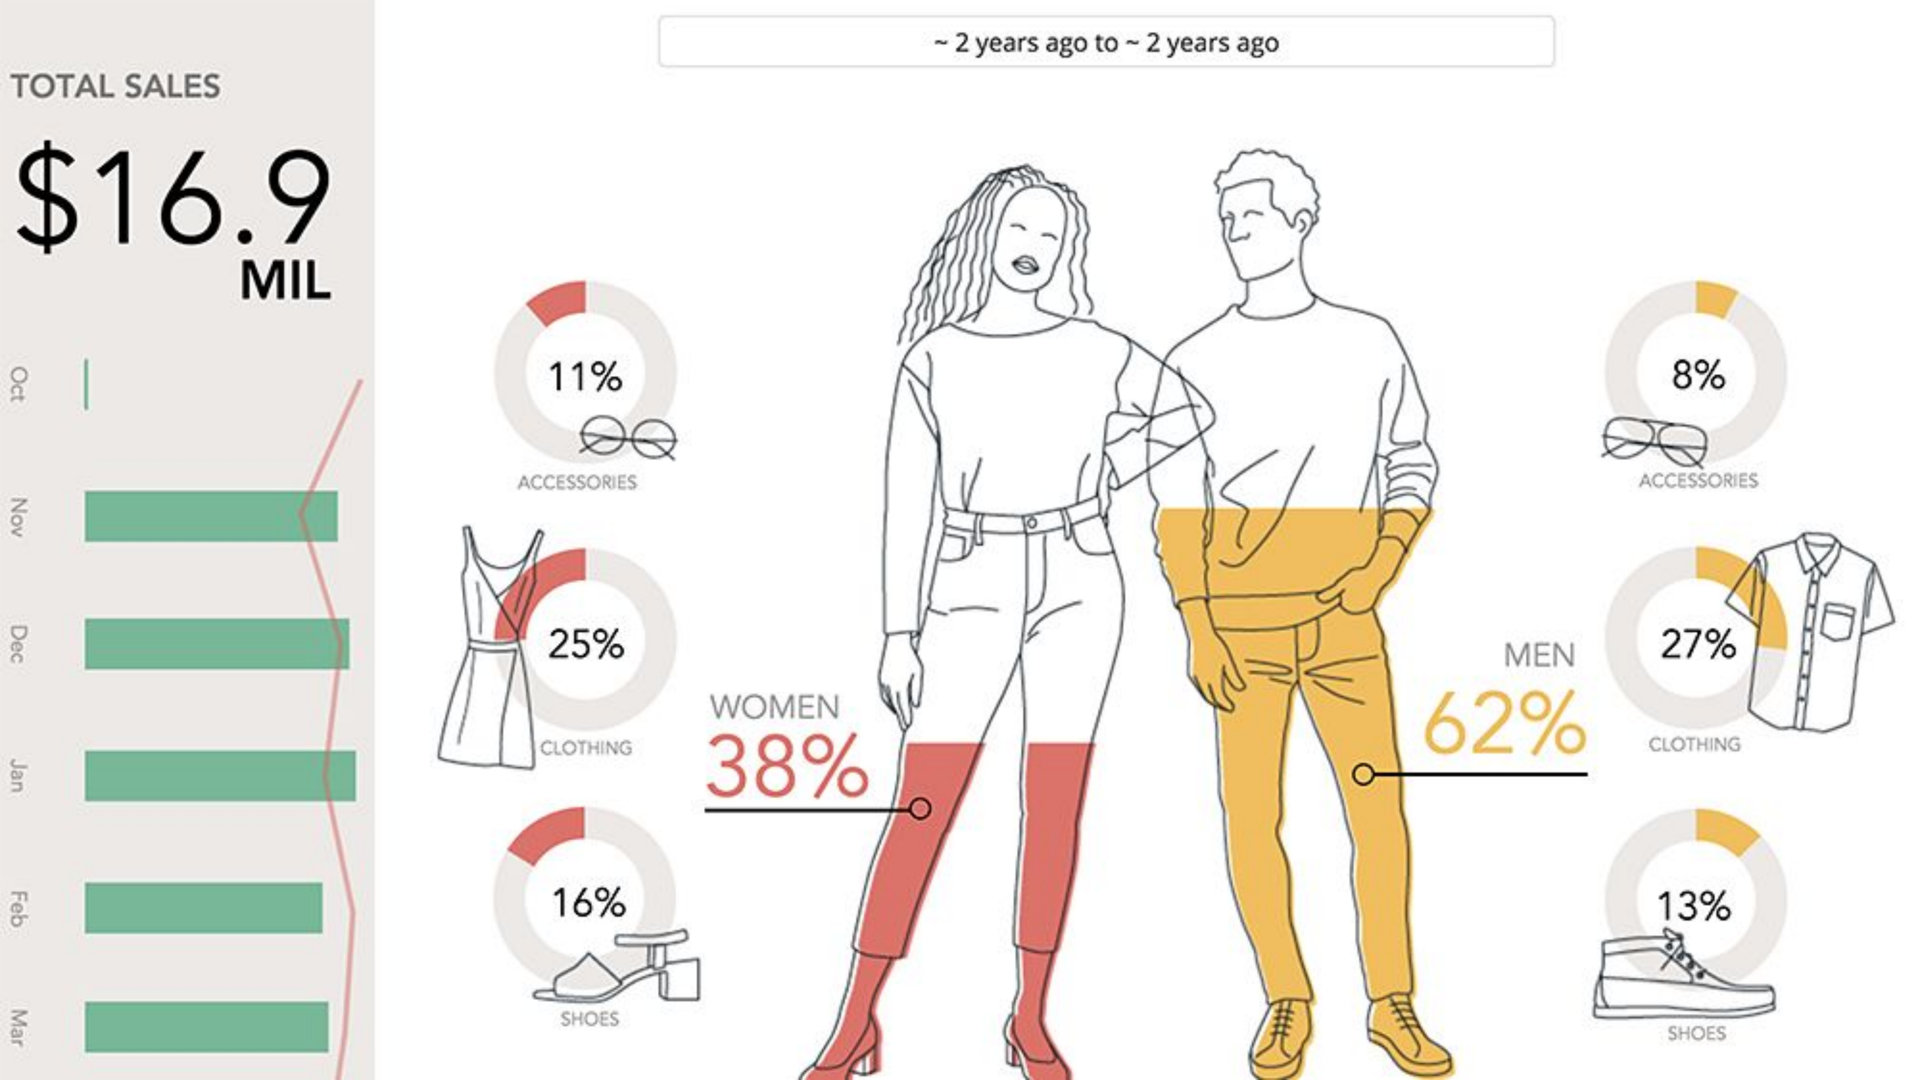

~ 2 years ago to ~ 2 years ago

TOTAL SALES

$16.9 MIL

Subscription Options

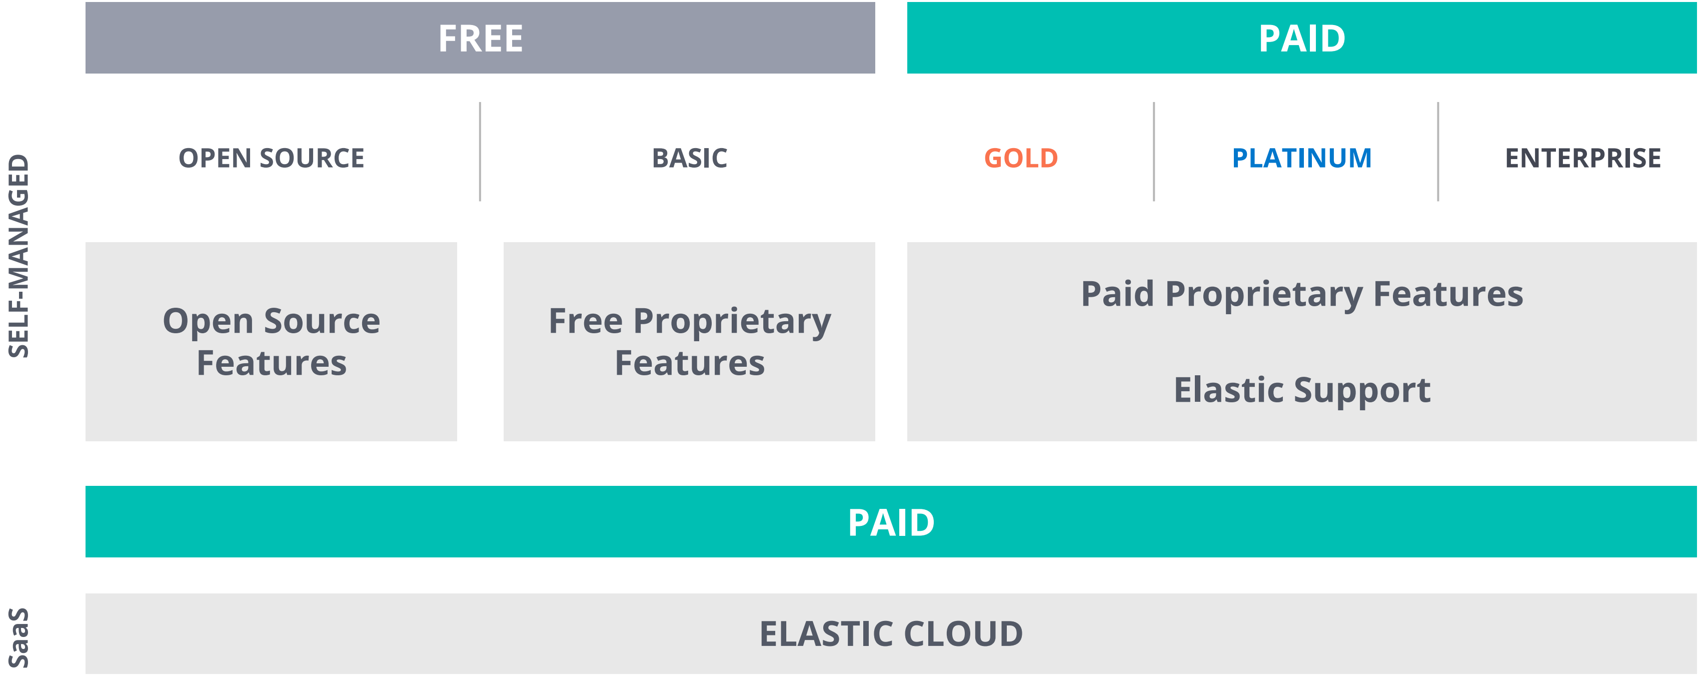

	FREE		PAID		
	OPEN SOURCE	BASIC	GOLD	PLATINUM	ENTERPRISE
SELF-MANAGED	Open Source Features	Free Proprietary Features	Paid Proprietary Features Elastic Support		
	PAID				
SaaS	ELASTIC CLOUD				

NowBit

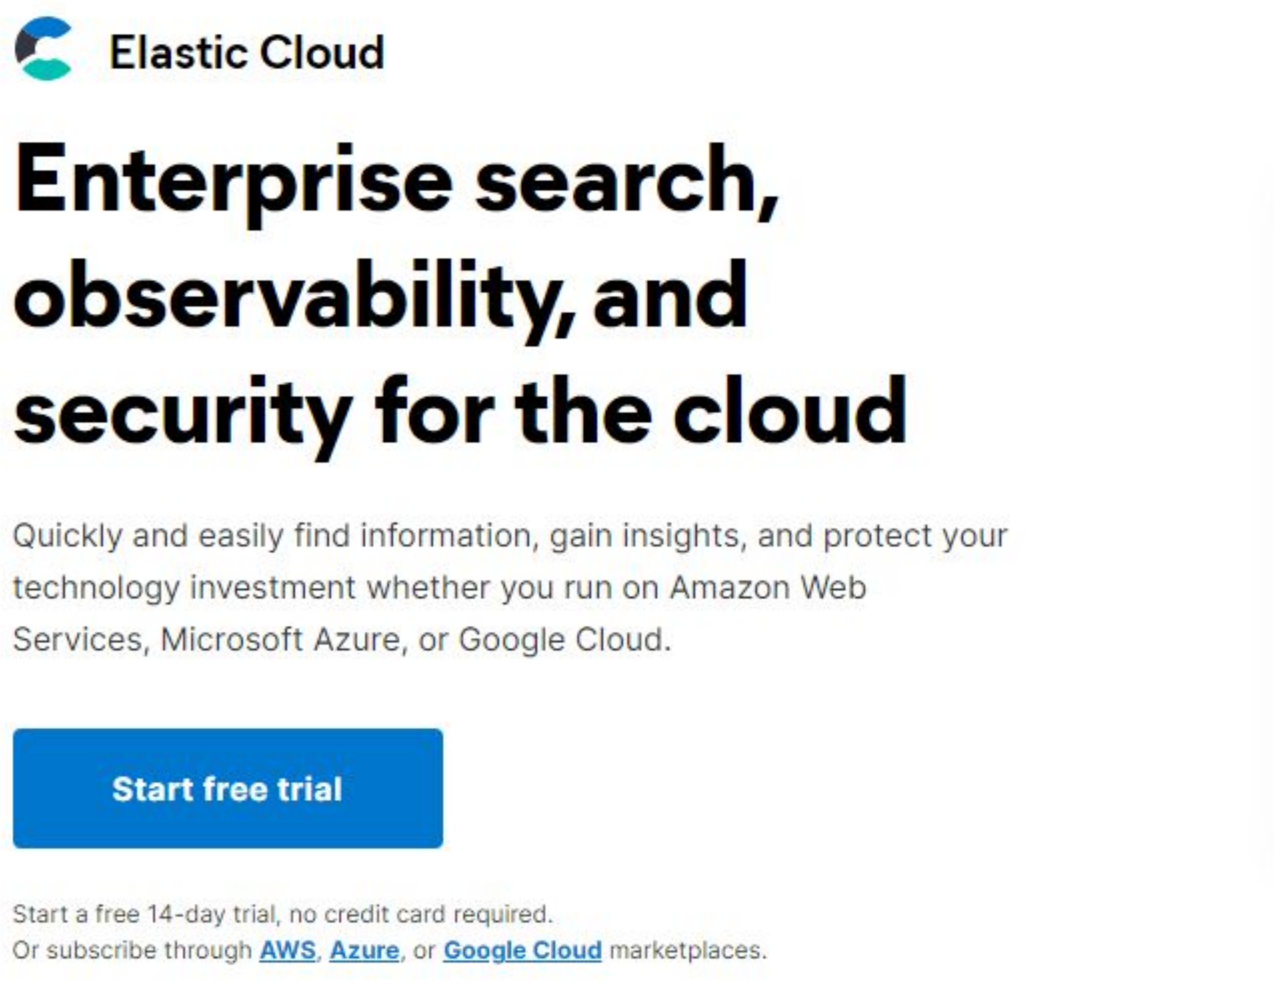

Elastic Cloud

Enterprise search, observability, and security for the cloud

Quickly and easily find information, gain insights, and protect your technology investment whether you run on Amazon Web Services, Microsoft Azure, or Google Cloud.

Start a free 14-day trial, no credit card required.
Or subscribe through AWS, Azure, or Google Cloud marketplaces.

elastic

Nombres de campos

Cuando tenemos información de múltiples orígenes puede suceder lo siguiente:

- Campos con la misma información pero con nombres o tipos de datos distintos.
- Por ejemplo: user, username, nginx.access.user_name

Esto tiene un gran impacto a través de múltiples índices.

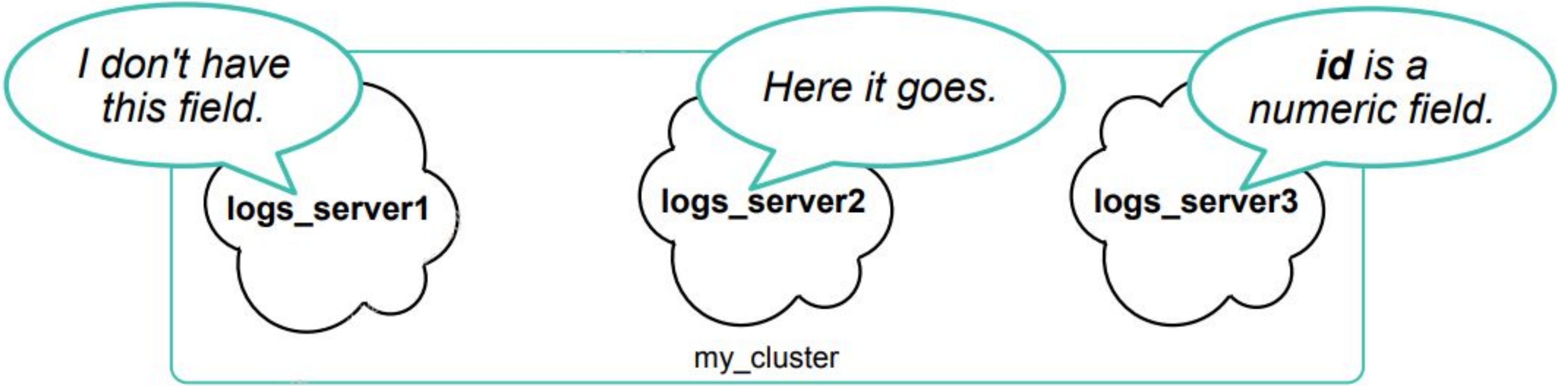

Elastic Common Schema

- Es una especificación open source .
  - Que define un conjunto común de campos de documentos.
- Está diseñado para admitir modelos de datos uniformes.
  - Especifica nombres de campos y tipos de datos de Elasticsearch para cada campo.
  - Provee descripciones y ejemplos de uso.

Elastic Common Schema

- El ECS define dos niveles de campos: Core y Extended.
- **Core**
    - Campos que son comunes a través de la mayoría de casos de uso.
- **Extended**
    - Campos que complementan la información de los core fields.
    - Contienen información más específica.

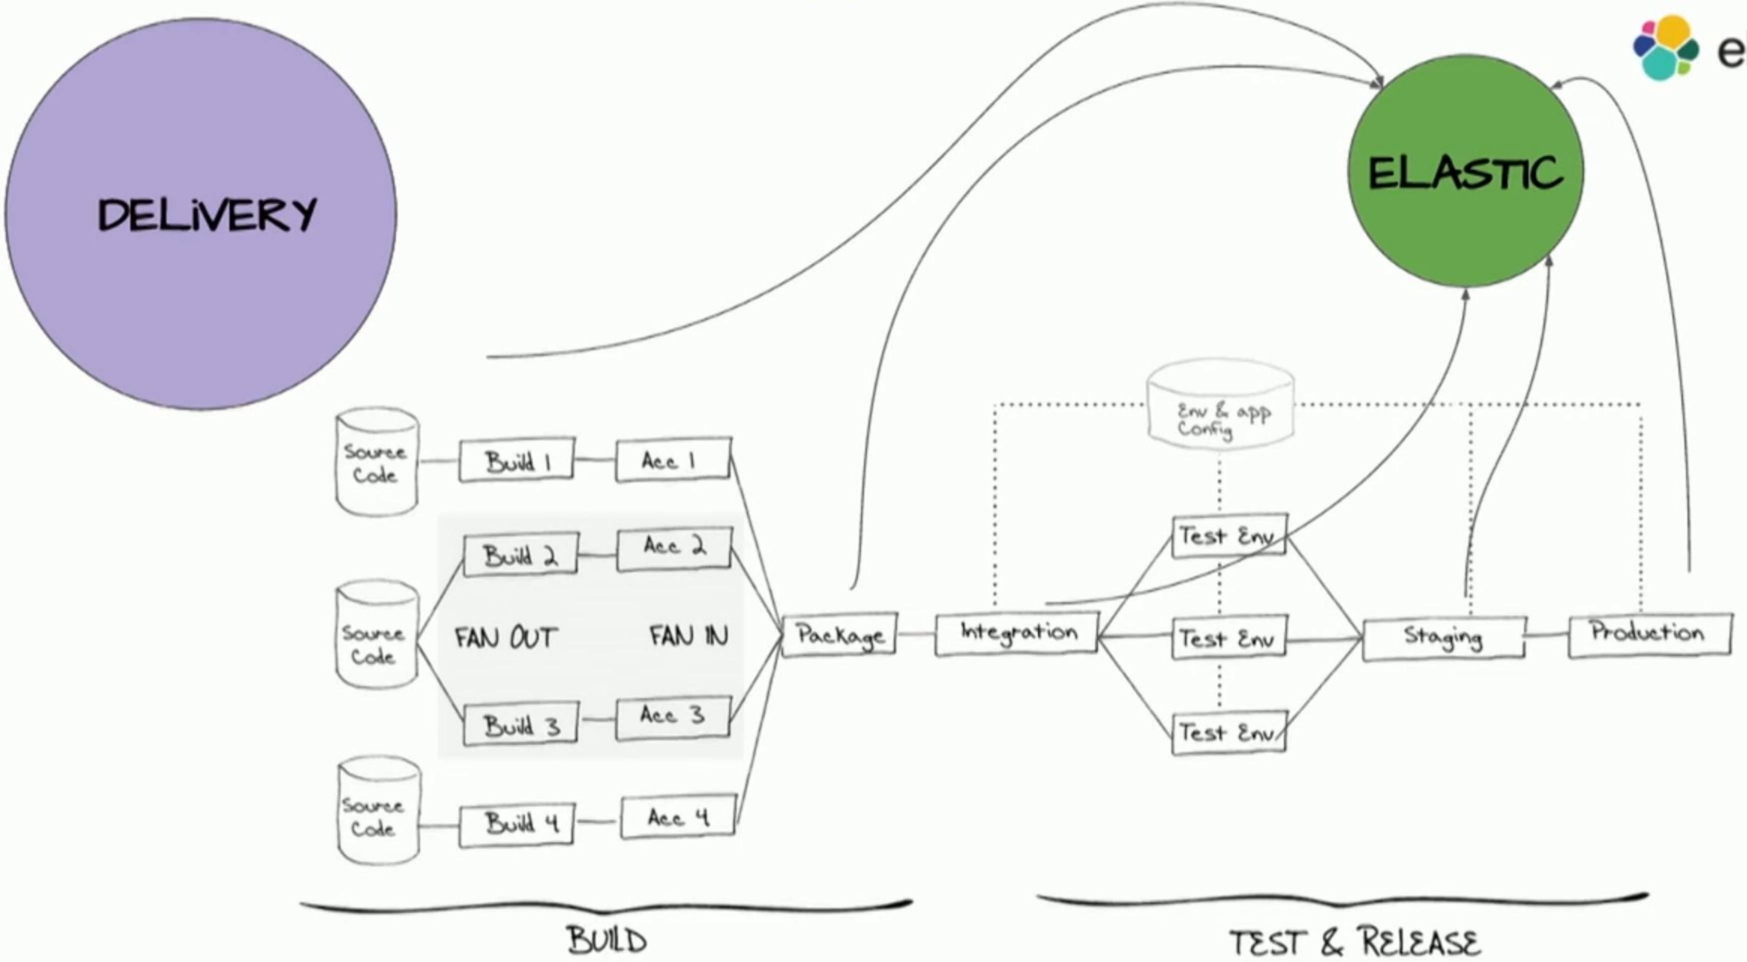

Image courtesy of GoCD

NowBit

¡Gracias!

NowBit, expertos en buscar, observar y proteger.
www.nowbit.co