SOUS LE HAUT PATRONAGE DE SA MAJESTÉ LE ROI MOHAMMED VI

DEVOXX™ MOROCCO

November 27, 28 and 29, 2018
Marrakech, Morocco

Monitor Your Java Applications with the Elastic Stack

Philipp Krenn

@xeraa

Elastic

elastic

#DevoxxMa

Powered by : XHUB

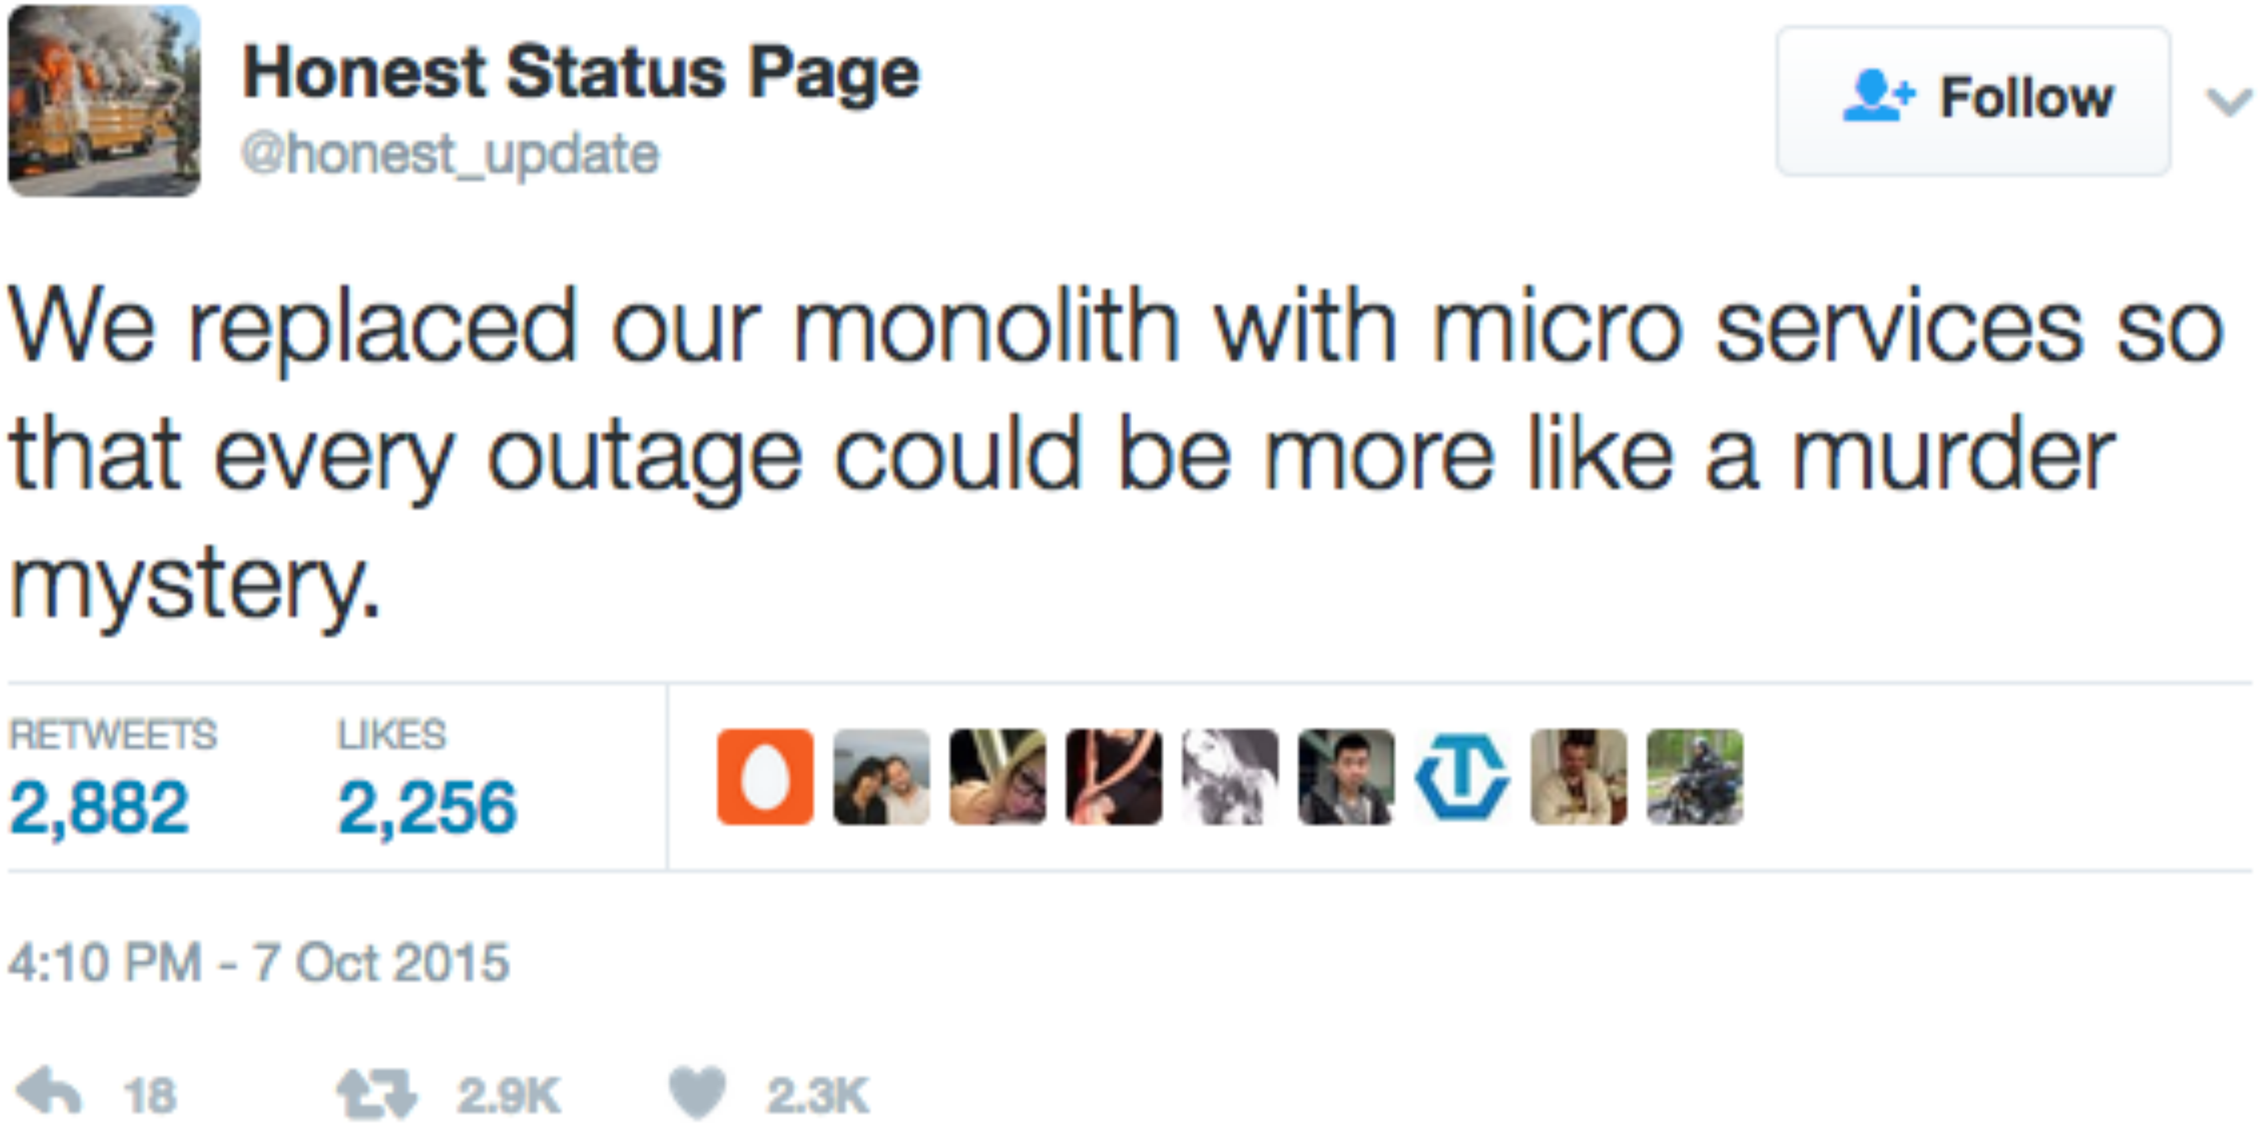

Honest Status Page
@honest_update

We replaced our monolith with micro services so that every outage could be more like a murder mystery.

RETWEETS 2,882
LIKES 2,256

4:10 PM - 7 Oct 2015

elastic

TH-
THIS IS NOT FINE!!
OHMYGOD EVERYTHINGS ON FIRE
WHAT THE HELL IS MY PROBLEM

ALL THE THINGS!

HOW?

🏝️ vs 🗺️

elastic

Developer 🥑

Disclaimer

I build **highly** monitored Hello World apps

Disclaimer

This is **not** a training

https://training.elastic.co

elasticsearch.

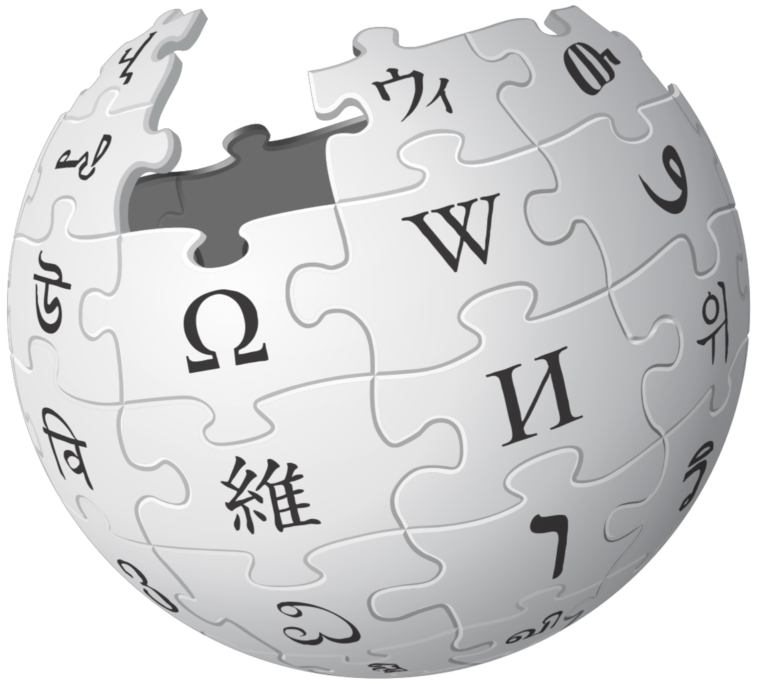

ウィ
W
Ω
И
維

kibana

logstash

lyft

slack

fitbit

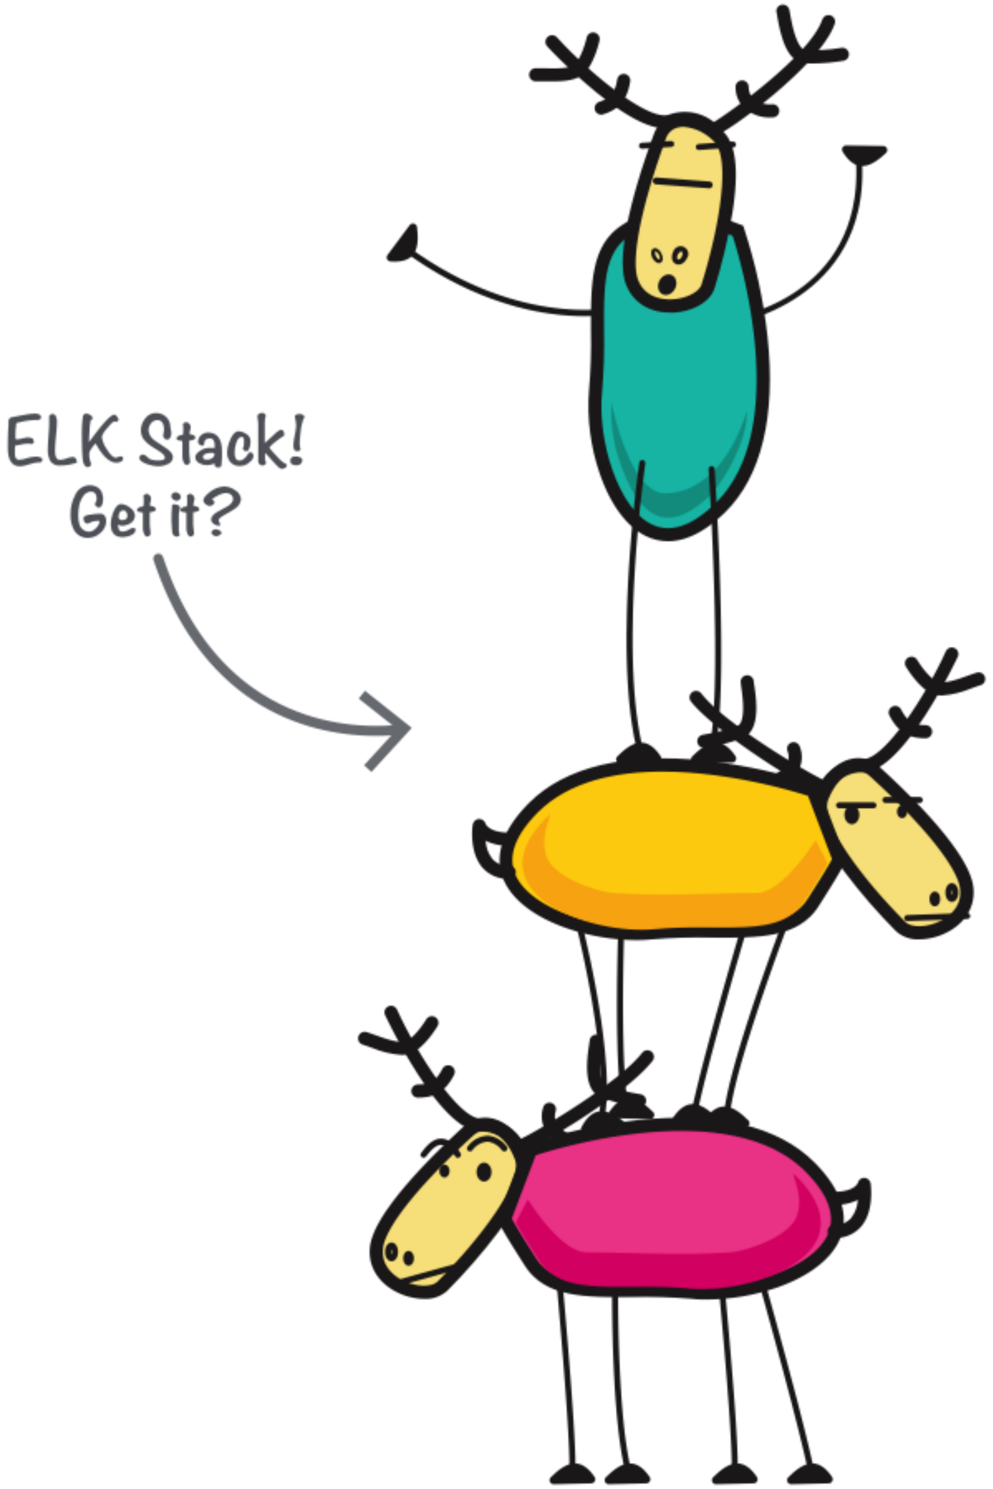

E Elasticsearch

L Logstash

K Kibana

Apparently, I'm an ELKB personality.
Meyers Briggs
Meyers Briggs

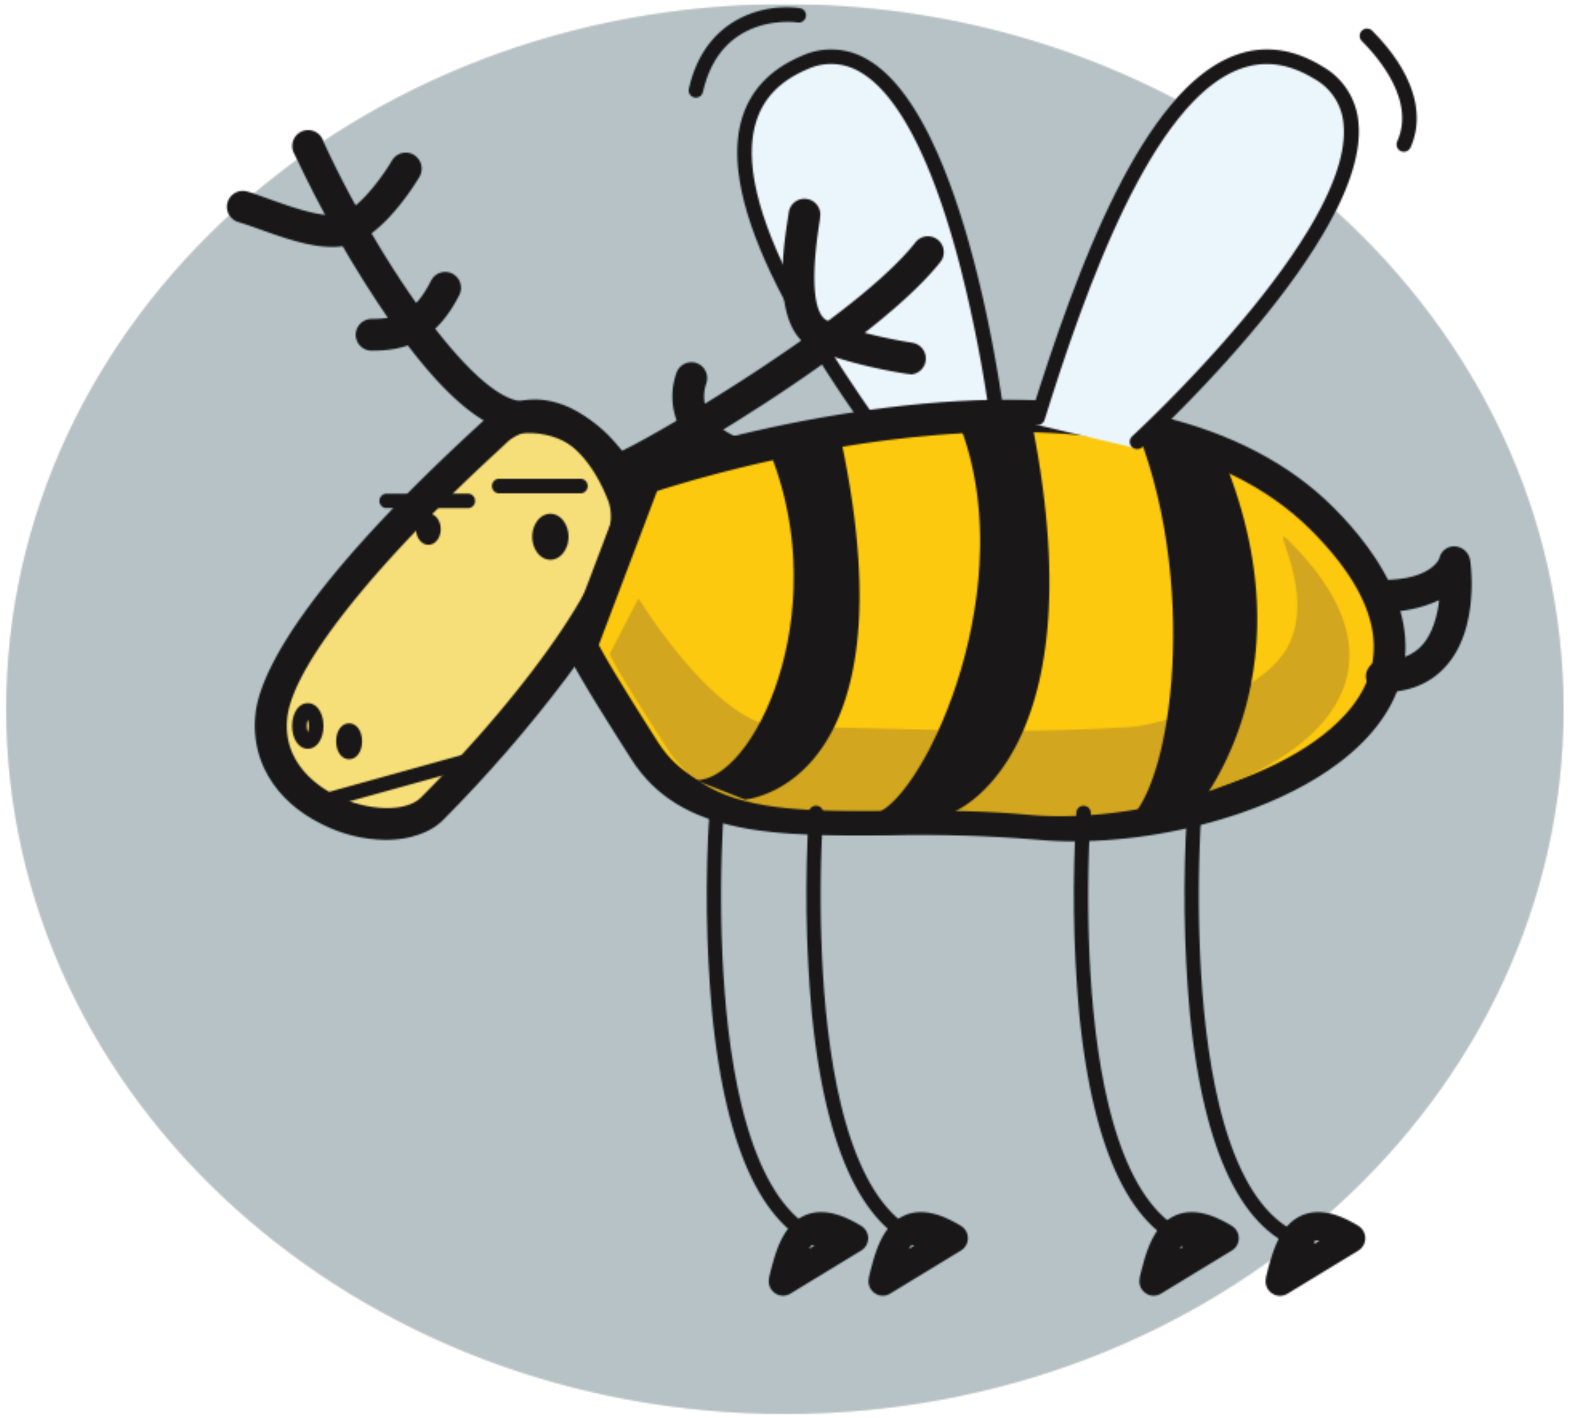

elastic

elastic stack

BLIZZARD®
ENTERTAINMENT

CAR
2GO

ebay

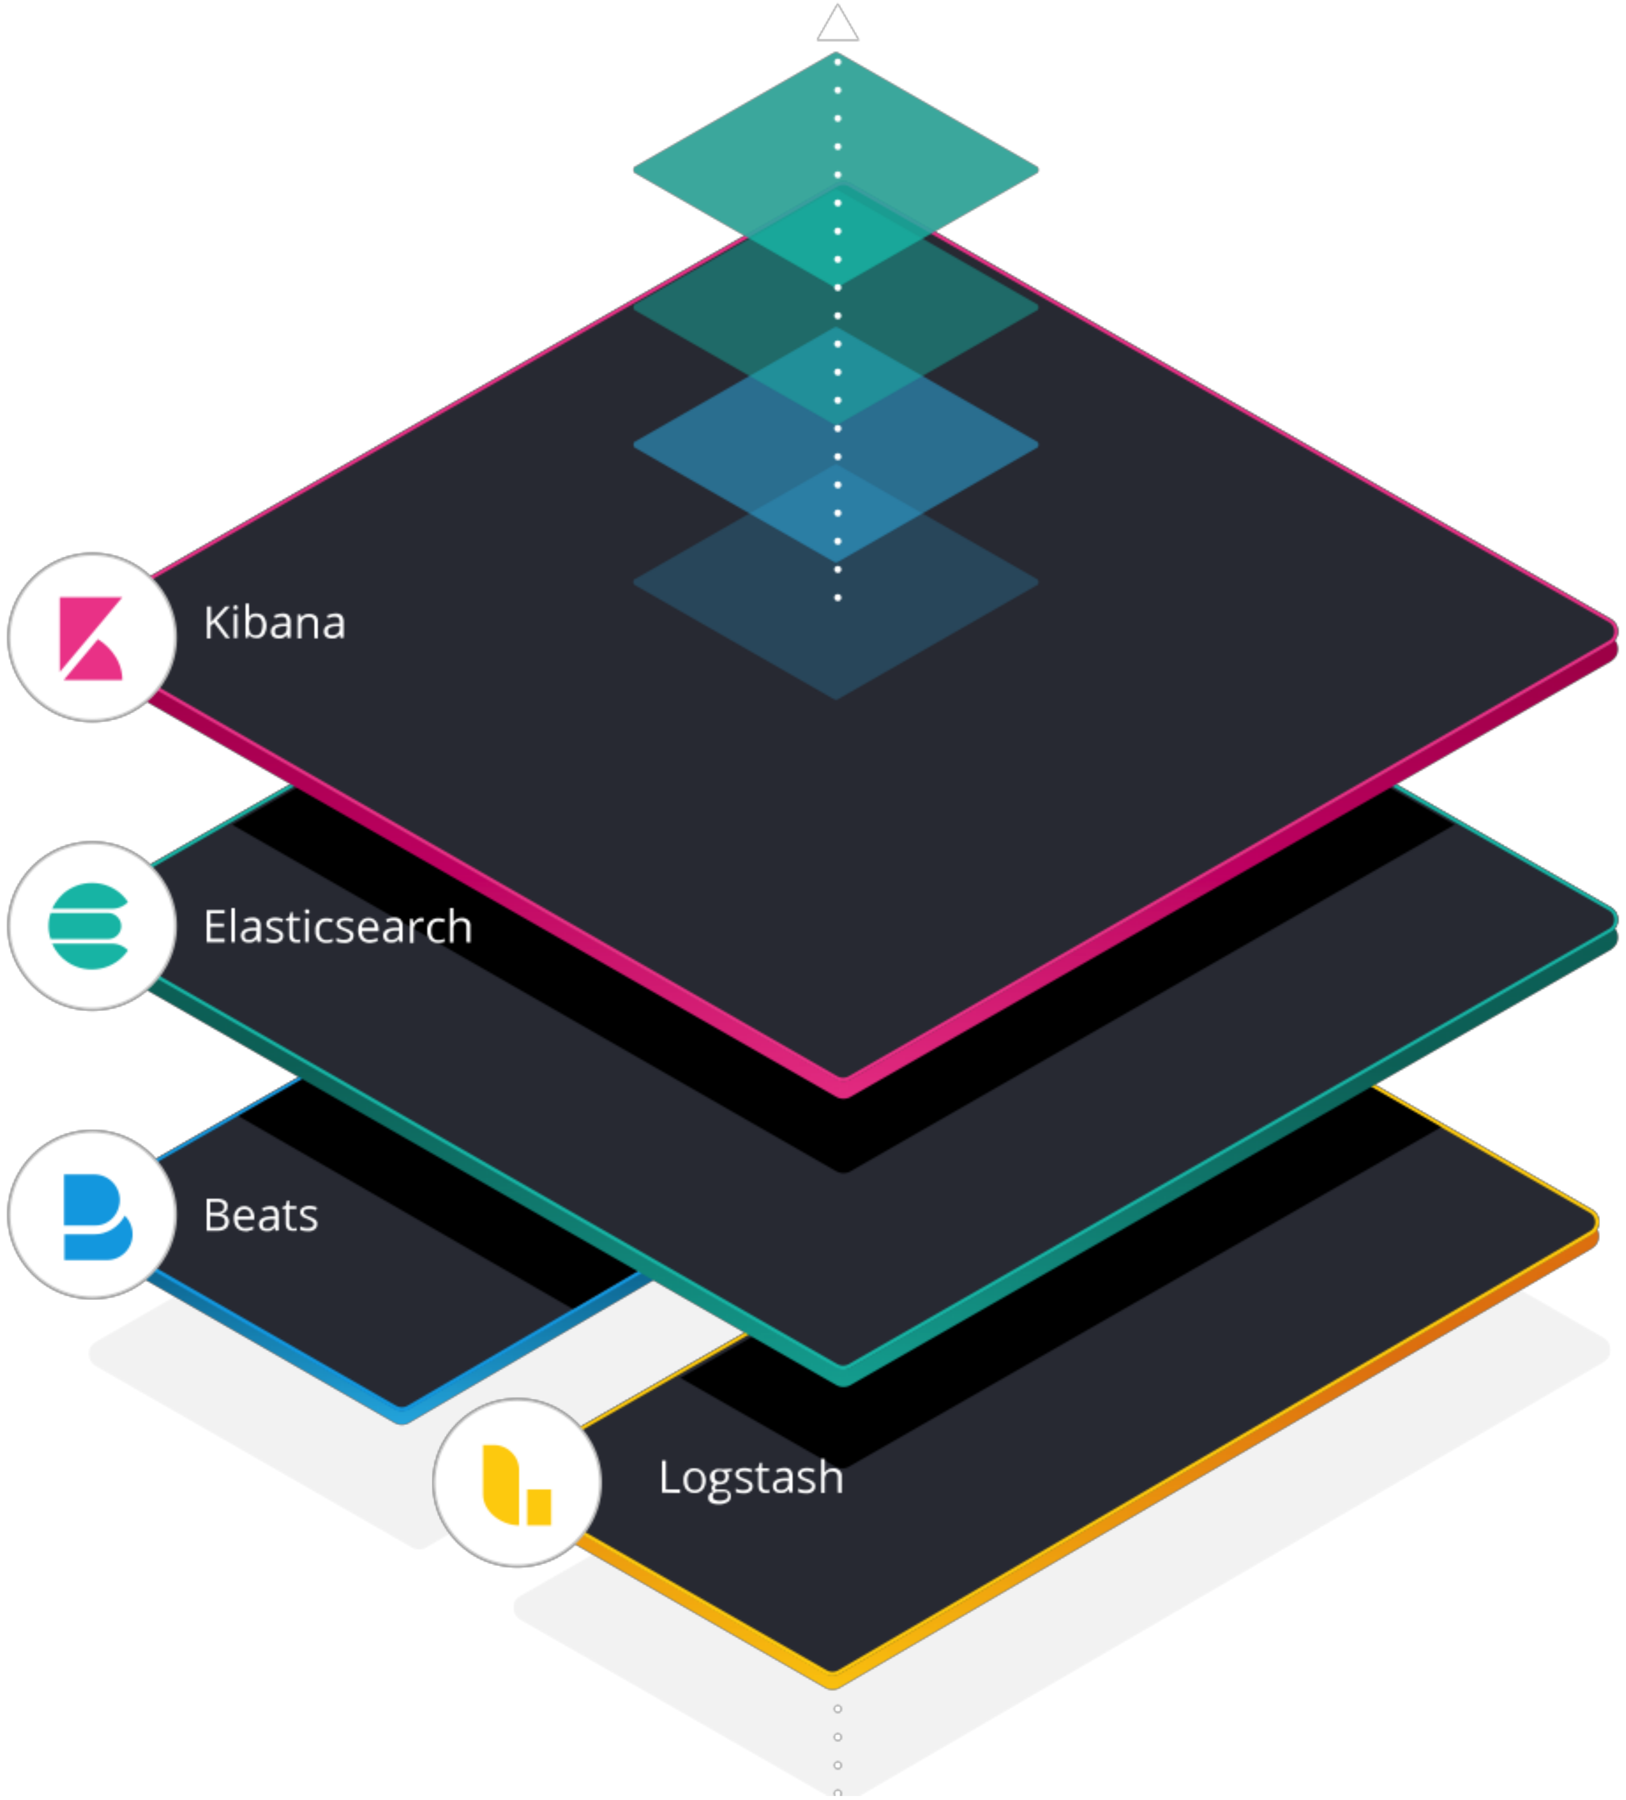
Kibana
Elasticsearch
Beats
Logstash

Licensing

Open Source Apache-2.0

Basic free

Commercial 💸

Code

https://github.com/xeraa/microservice-monitoring

Simple

No discovery, load-balancing,...

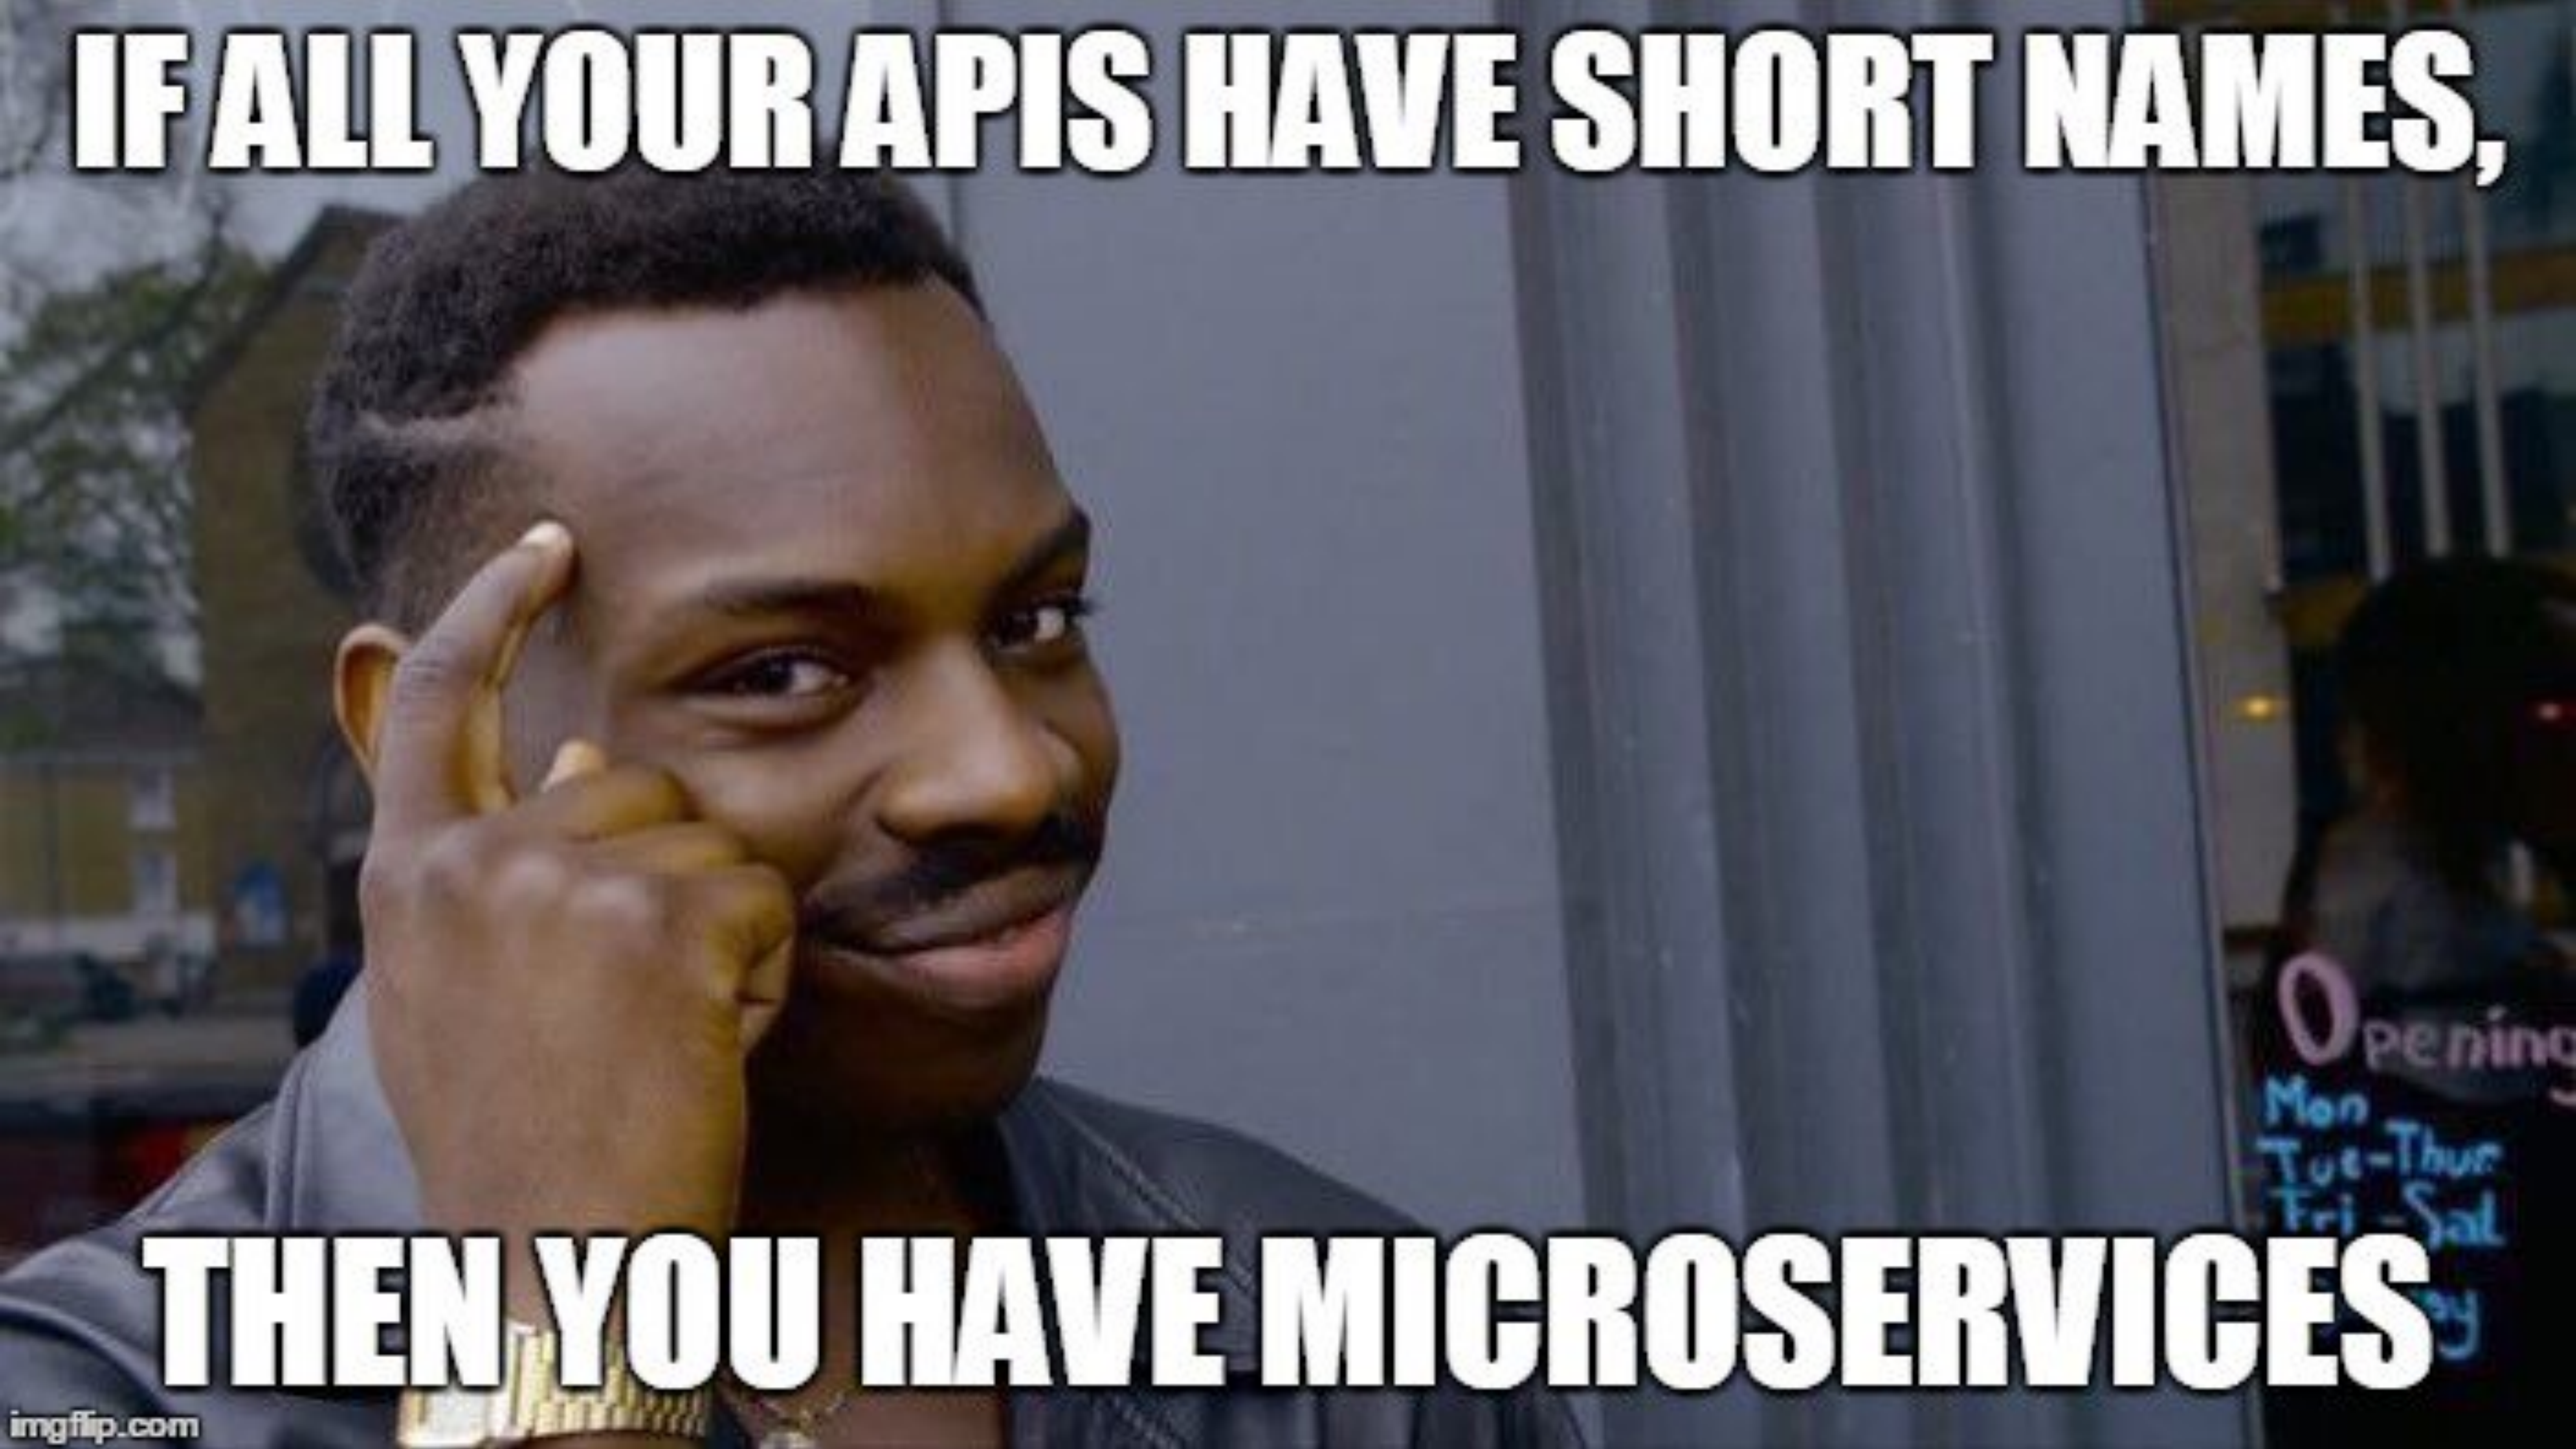
IF ALL YOUR APIS HAVE SHORT NAMES,
THEN YOU HAVE MICROSERVICES
Opening
Mon
Tue-Thur
Fri-Sat
imgflip.com

Amazon Lightsail

TERRAFORM

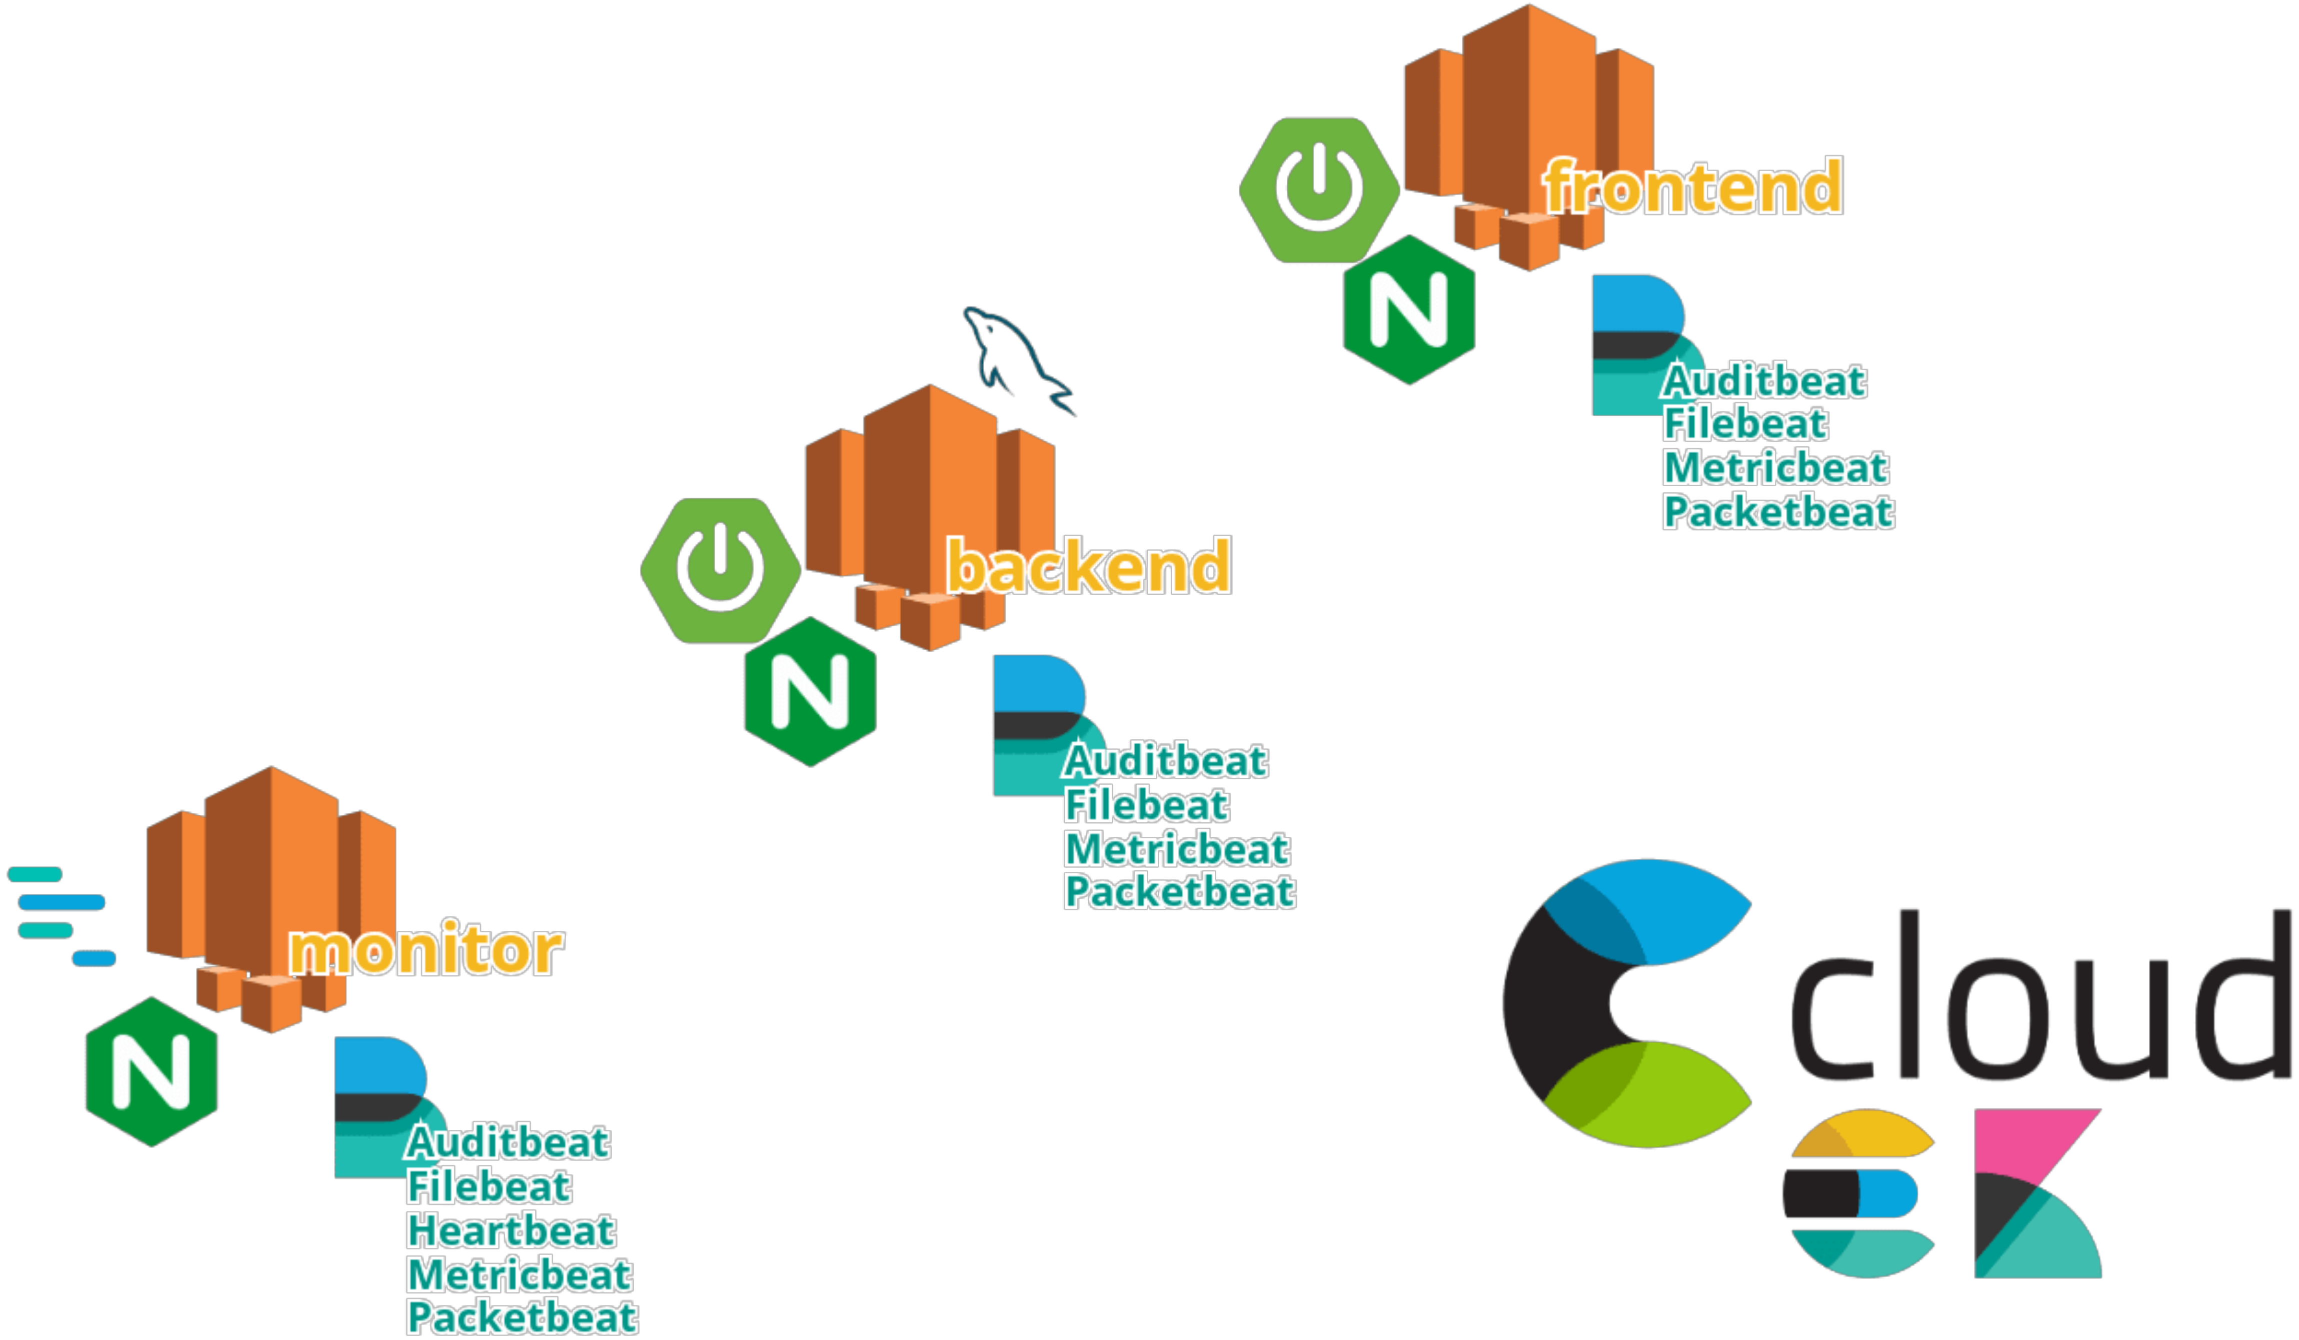

frontend

Auditbeat
Filebeat
Metricbeat
Packetbeat

backend

Auditbeat
Filebeat
Metricbeat
Packetbeat

monitor

Auditbeat
Filebeat
Heartbeat
Metricbeat
Packetbeat

cloud

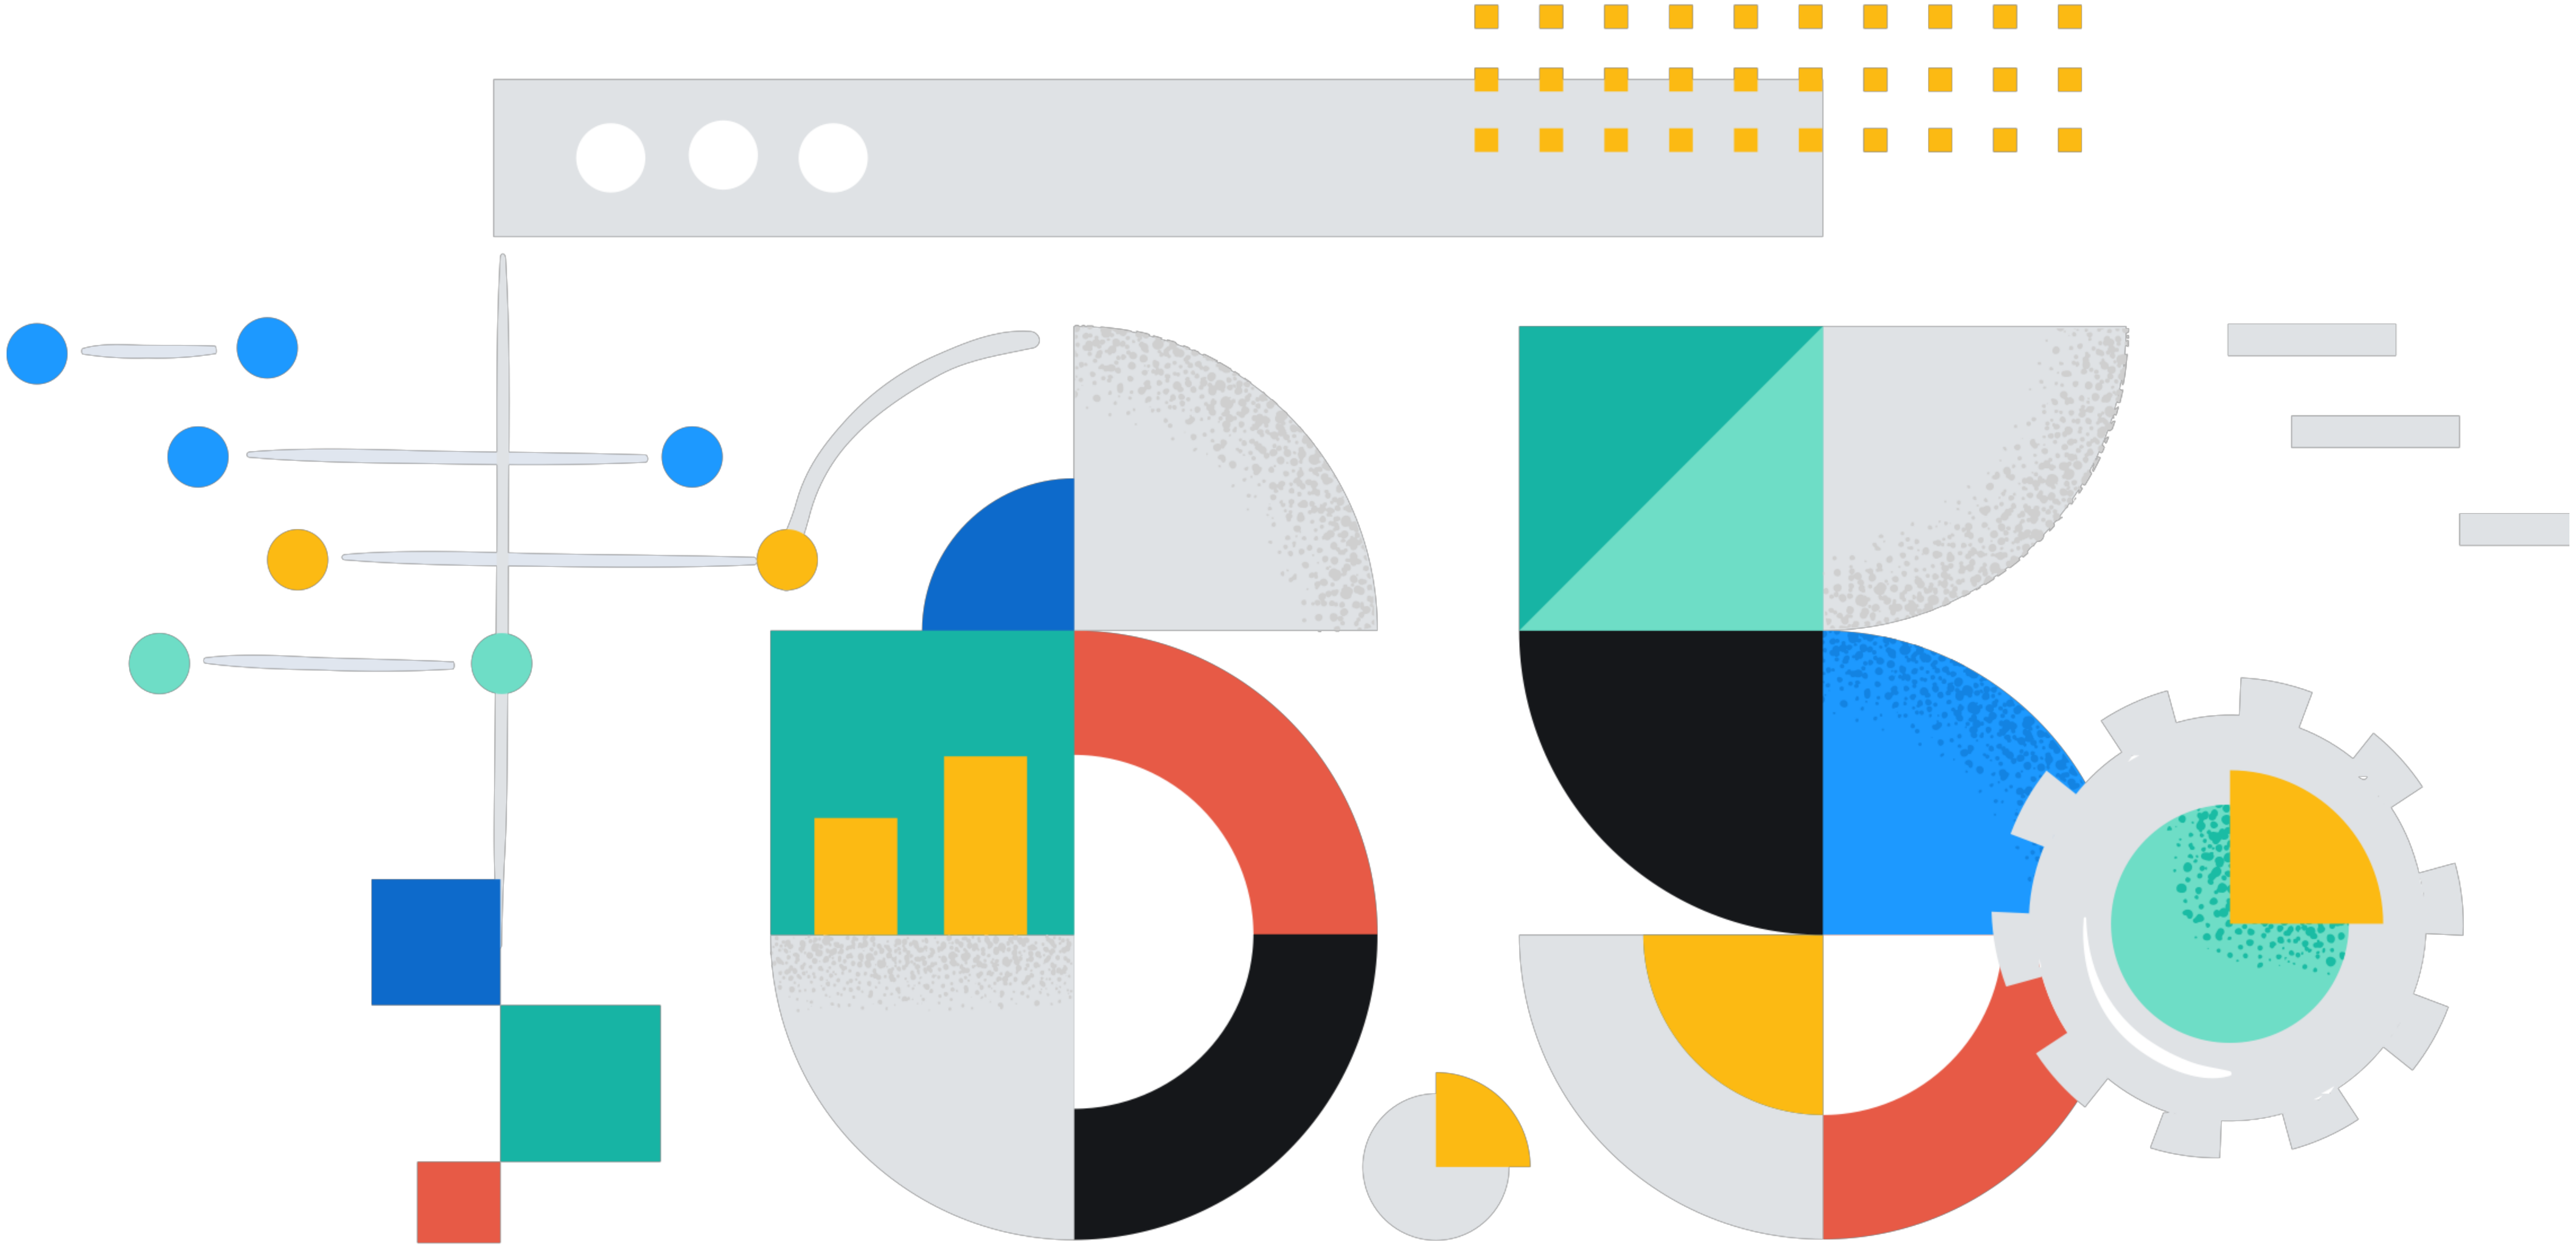

elastic

Kibana Monitoring

Overview of the Elastic Stack components

elastic

Metricbeat System

[Metricbeat System] Overview and [Metricbeat System] Host overview dashboards

See the memory spike every 5min

Time Series Visual Builder

Sum of system.memory.actual.used.bytes

Sum of system.process.memory.rss.bytes grouped by the term system.process.name and moved to the negative y-axis with a Math step

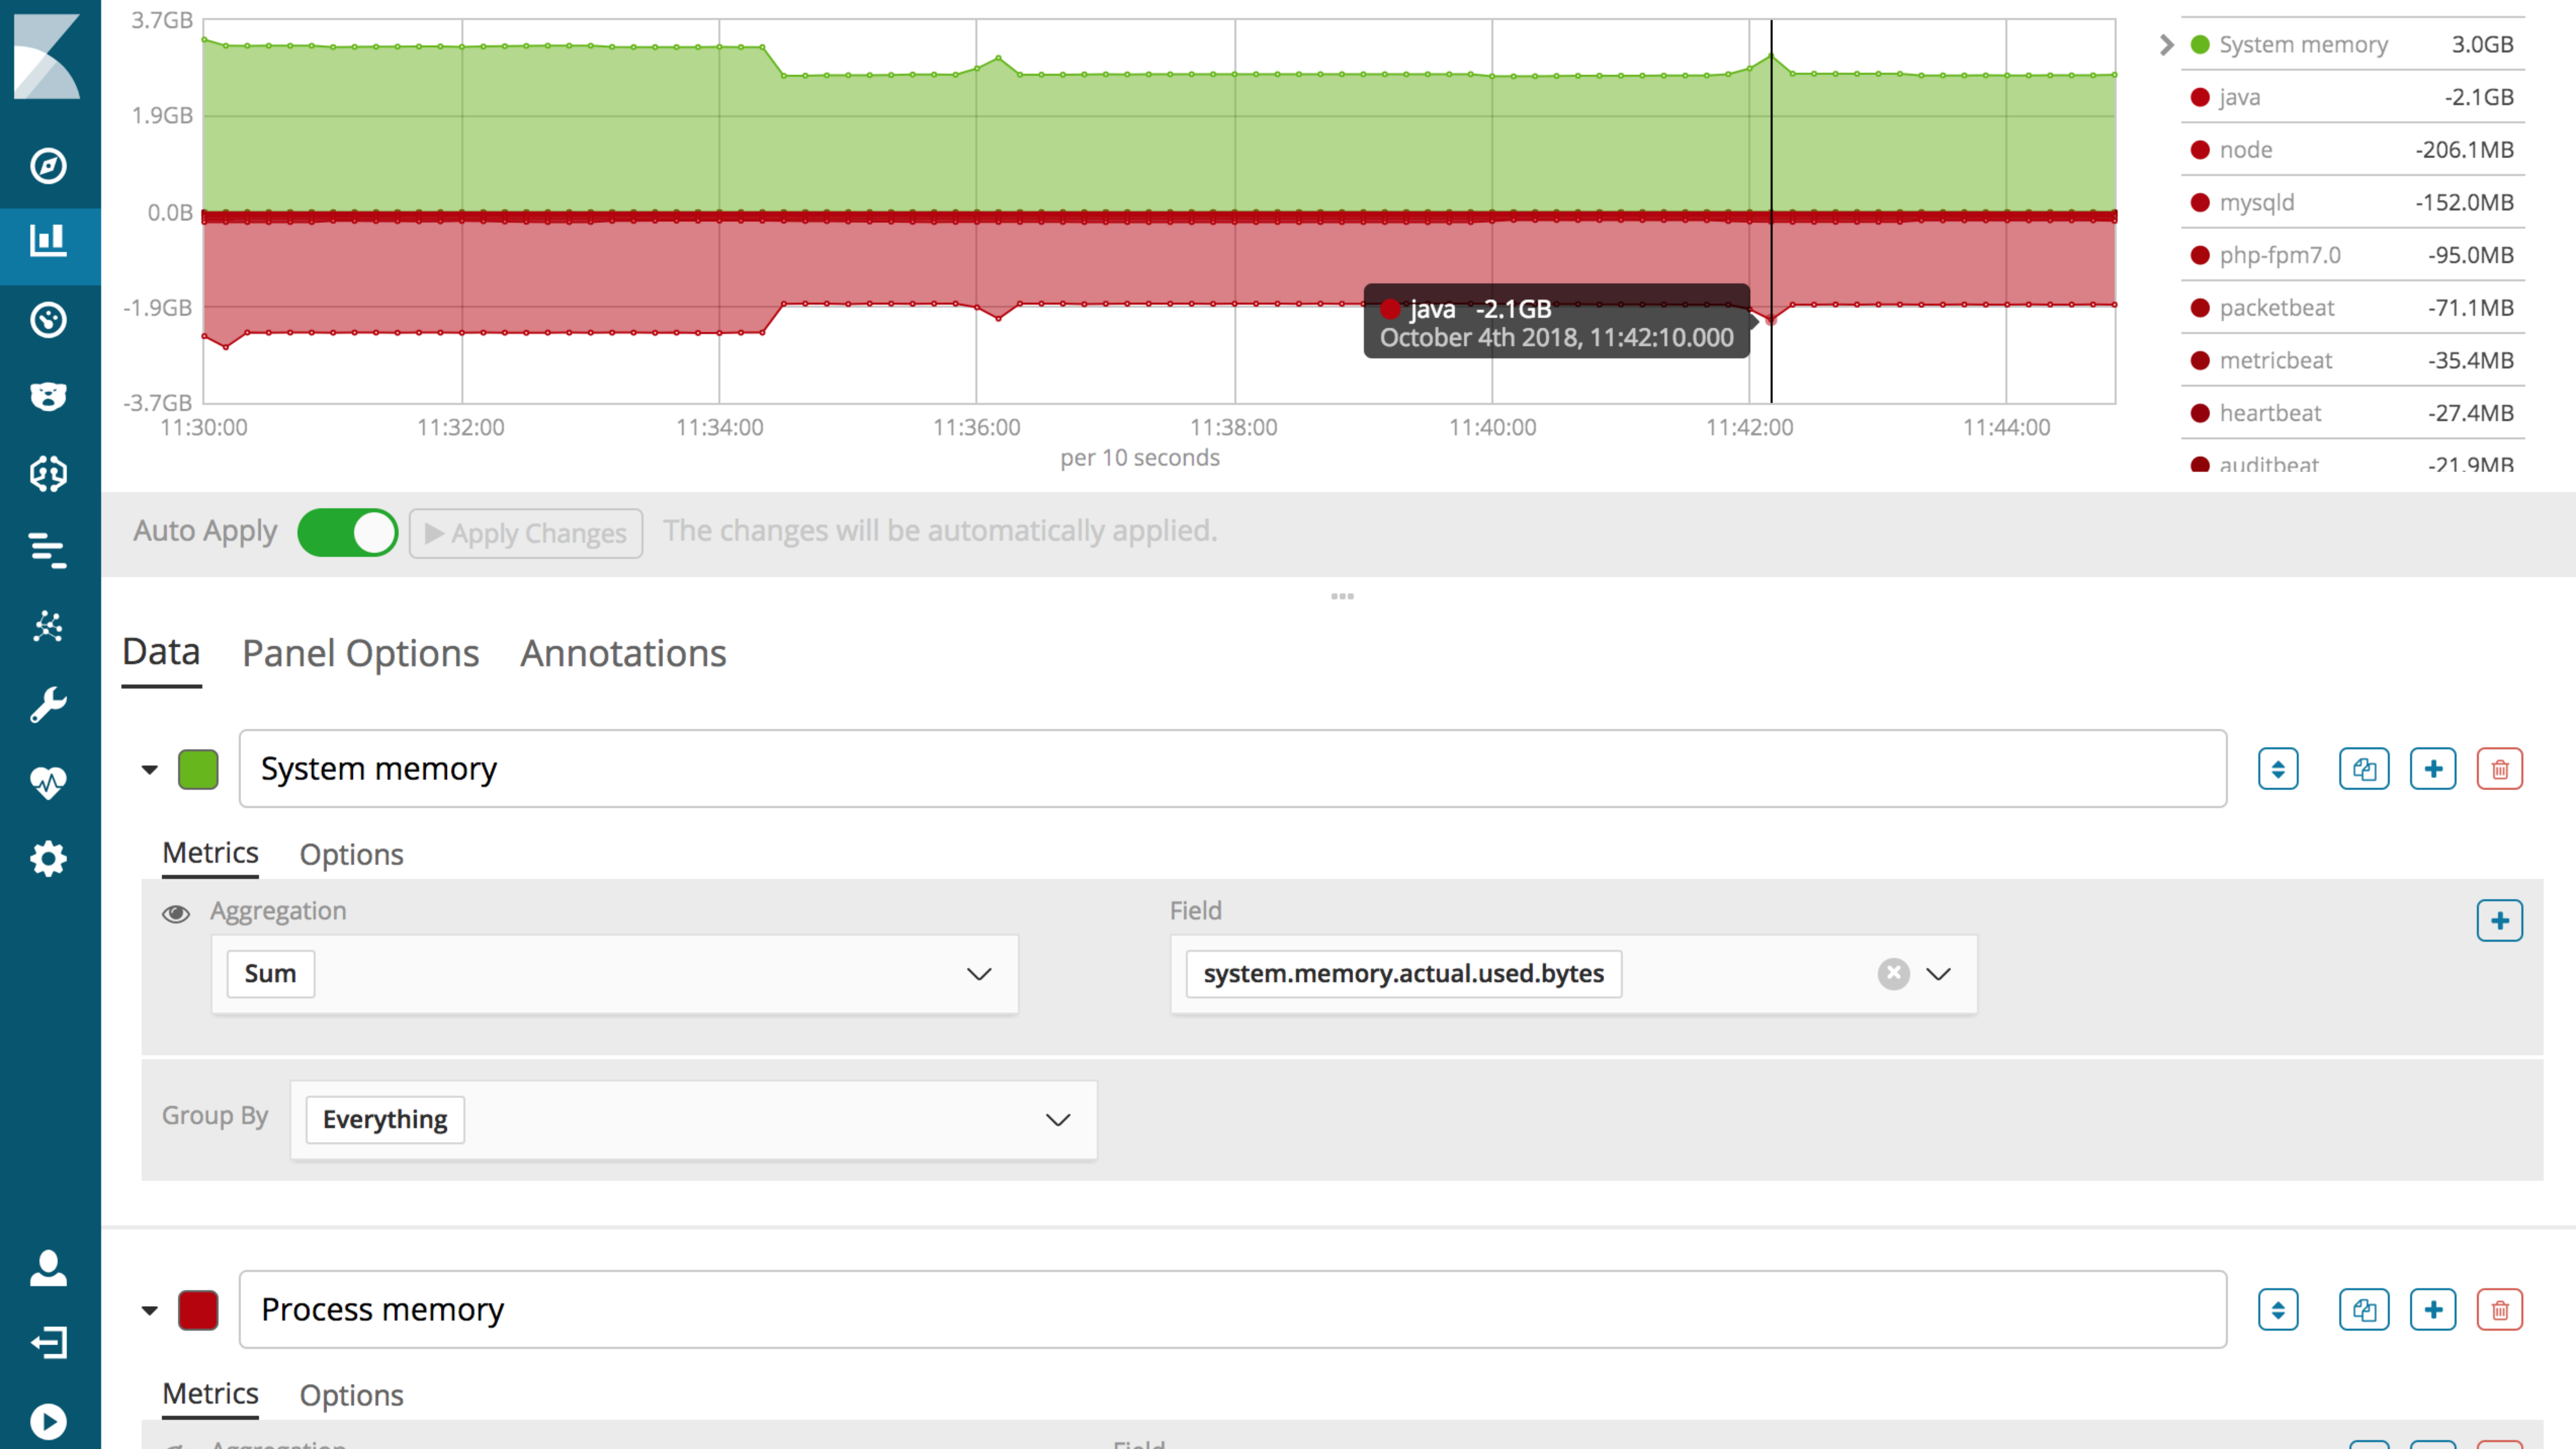

🟩 System memory	3.0GB
🟥 java	-2.1GB
🟥 node	-206.1MB
🟥 mysqld	-152.0MB
🟥 php-fpm7.0	-95.0MB
🟥 packetbeat	-71.1MB
🟥 metricbeat	-35.4MB
🟥 heartbeat	-27.4MB
🟥 auditbeat	-21.9MB

Auto Apply ● ▶ Apply Changes The changes will be automatically applied.

Data | Panel Options | Annotations

🟩 System memory

Metrics | Options

Aggregation: Sum

Field: system.memory.actual.used.bytes

Group By: Everything

🟥 Process memory

Metrics | Options

Aggregation

Field

Packetbeat

Call /, /good, /bad, and /foobar

[Packetbeat] Overview, [Packetbeat] Flows, [Packetbeat] HTTP, and [Packetbeat] DNS Tunneling dashboards

Packetbeat

Raw events in Discover

Process enrichment for nginx, Java, and the APM server

Filebeat Modules

[Filebeat Nginx] Access and error logs, [Filebeat System] Syslog dashboard, and [Osquery Result] Compliance pack dashboards

Custom Log Files

Elastic Common Schema

https://github.com/elastic/ecs

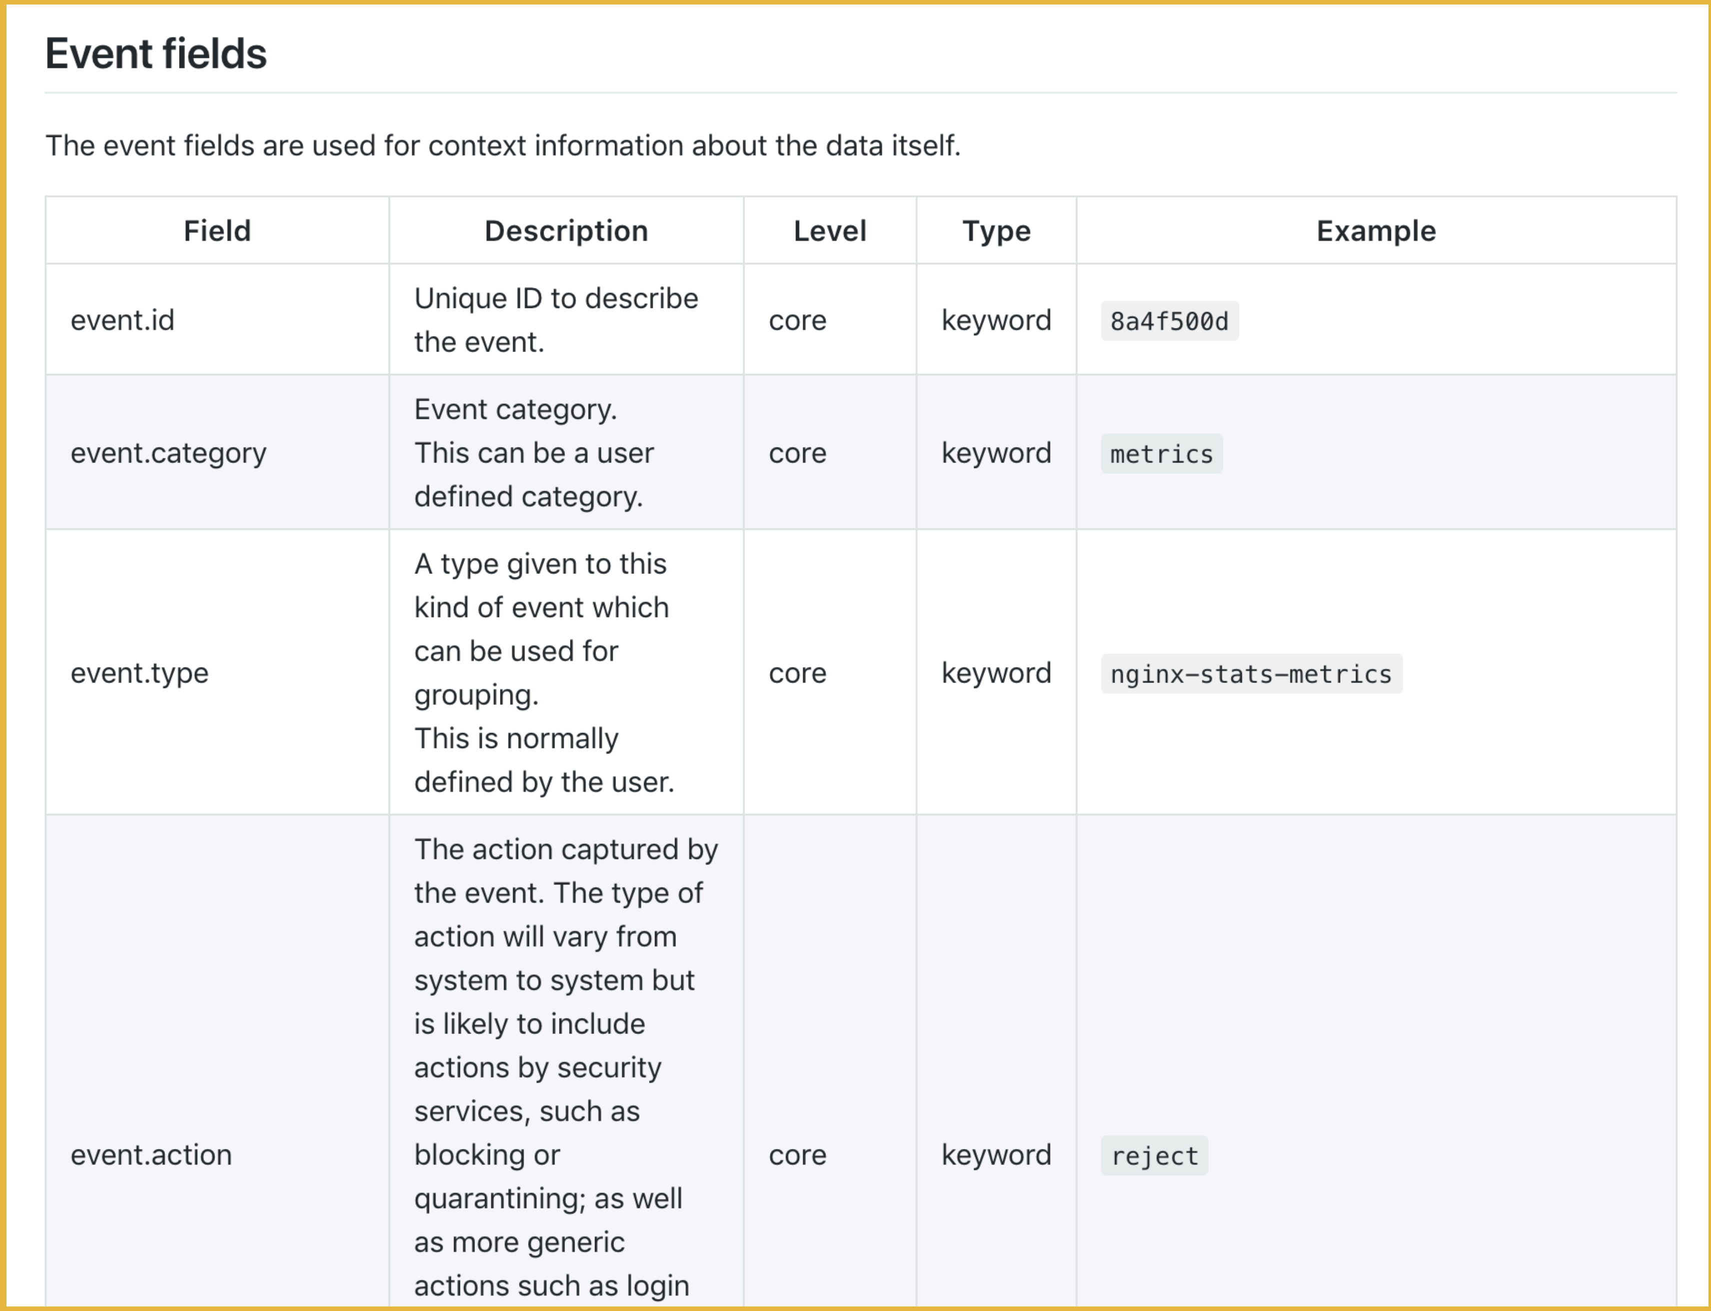

Event fields

The event fields are used for context information about the data itself.

Field	Description	Level	Type	Example
event.id	Unique ID to describe the event.	core	keyword	`8a4f500d`
event.category	Event category. This can be a user defined category.	core	keyword	`metrics`
event.type	A type given to this kind of event which can be used for grouping. This is normally defined by the user.	core	keyword	`nginx-stats-metrics`
event.action	The action captured by the event. The type of action will vary from system to system but is likely to include actions by security services, such as blocking or quarantining; as well as more generic actions such as login	core	keyword	`reject`

elastic

Setting for `negate`	Setting for `match`	Result	Example `pattern: ^b`
`false`	`after`	Consecutive lines that match the pattern are appended to the previous line that doesn't match.	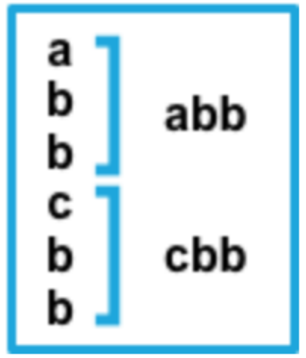
`false`	`before`	Consecutive lines that match the pattern are prepended to the next line that doesn't match.	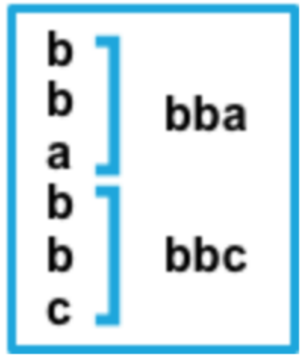
`true`	`after`	Consecutive lines that don't match the pattern are appended to the previous line that does match.	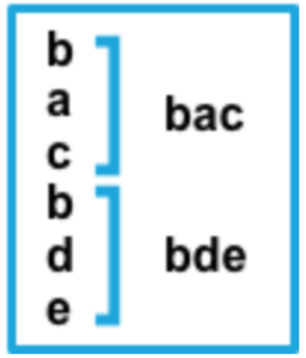
`true`	`before`	Consecutive lines that don't match the pattern are prepended to the next line that does match.	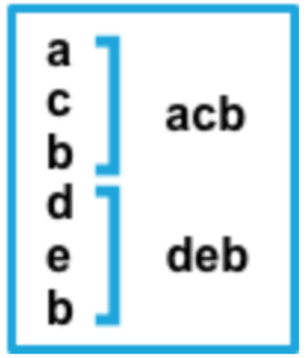

Dev Tools

Grok Debugger

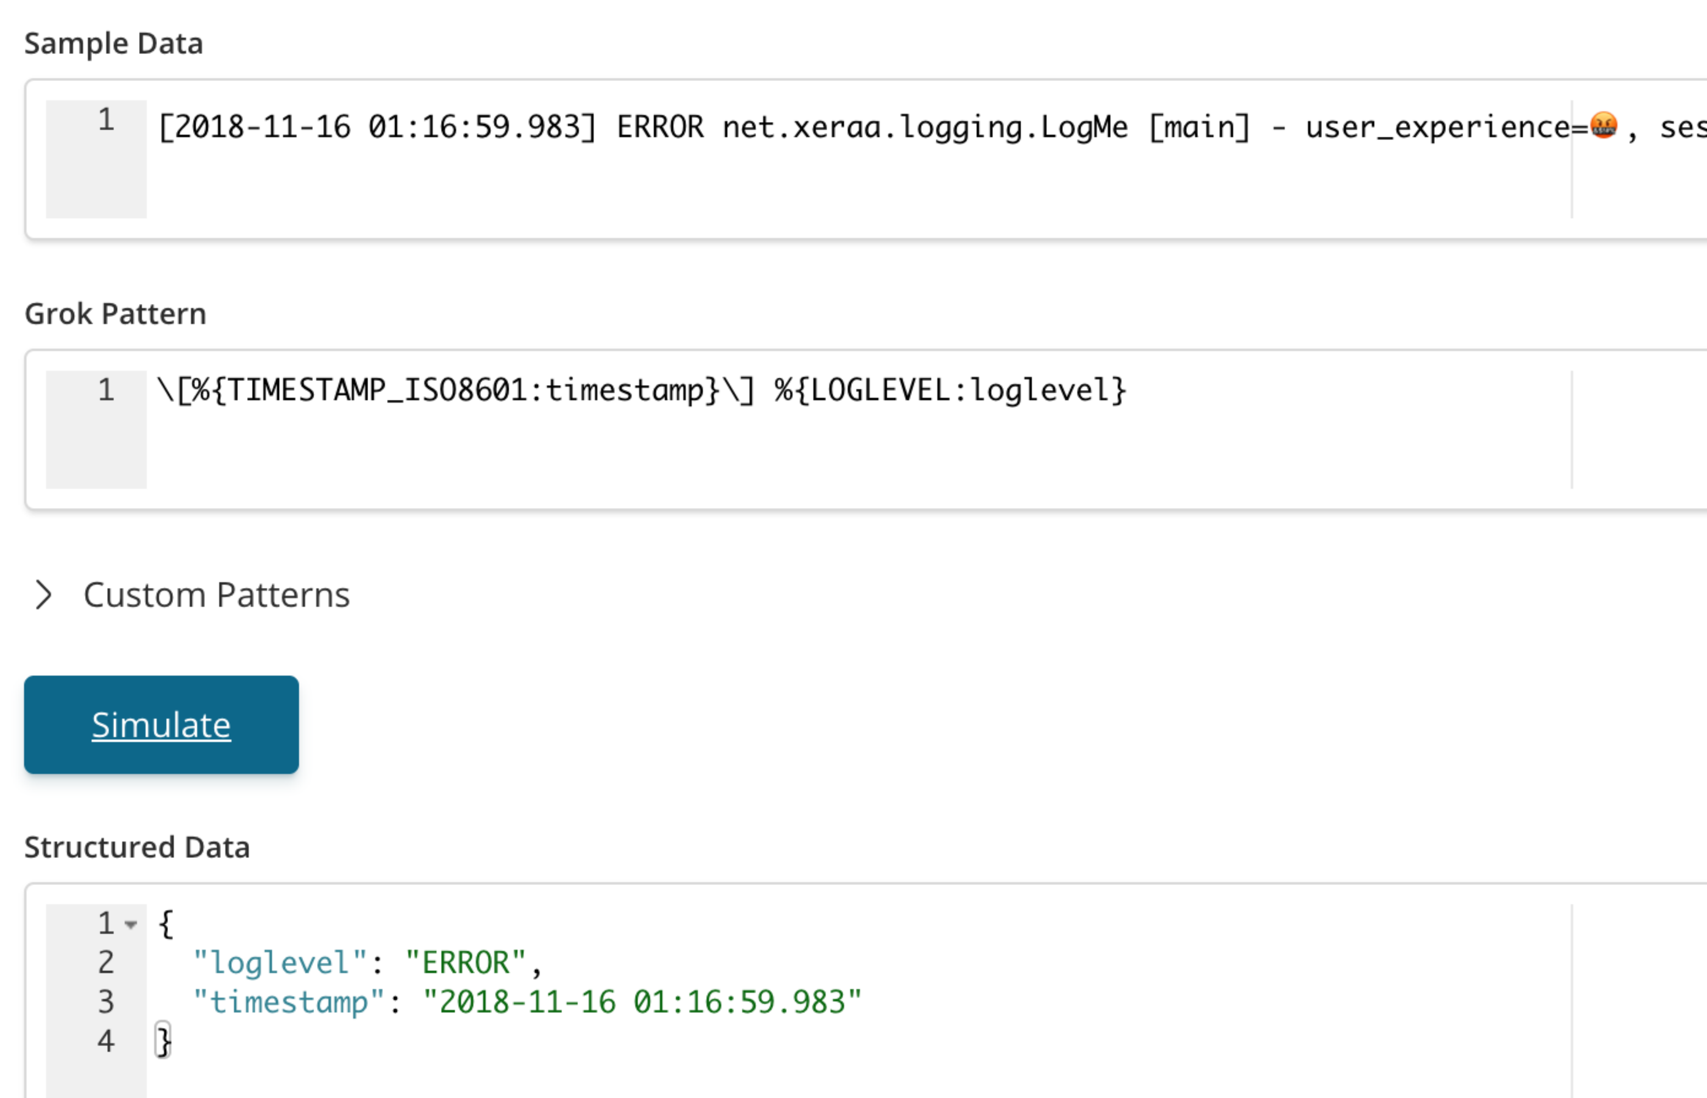

Sample Data

```
[2018-11-16 01:16:59.983] ERROR net.xeraa.logging.LogMe [main] - user_experience=🤬 , ses
```

Grok Pattern

```
\[%{TIMESTAMP_ISO8601:timestamp}\] %{LOGLEVEL:loglevel}
```

Structured Data

```
{
  "loglevel": "ERROR",
  "timestamp": "2018-11-16 01:16:59.983"
}
```

elastic

Machine Learning

Data Visualizer

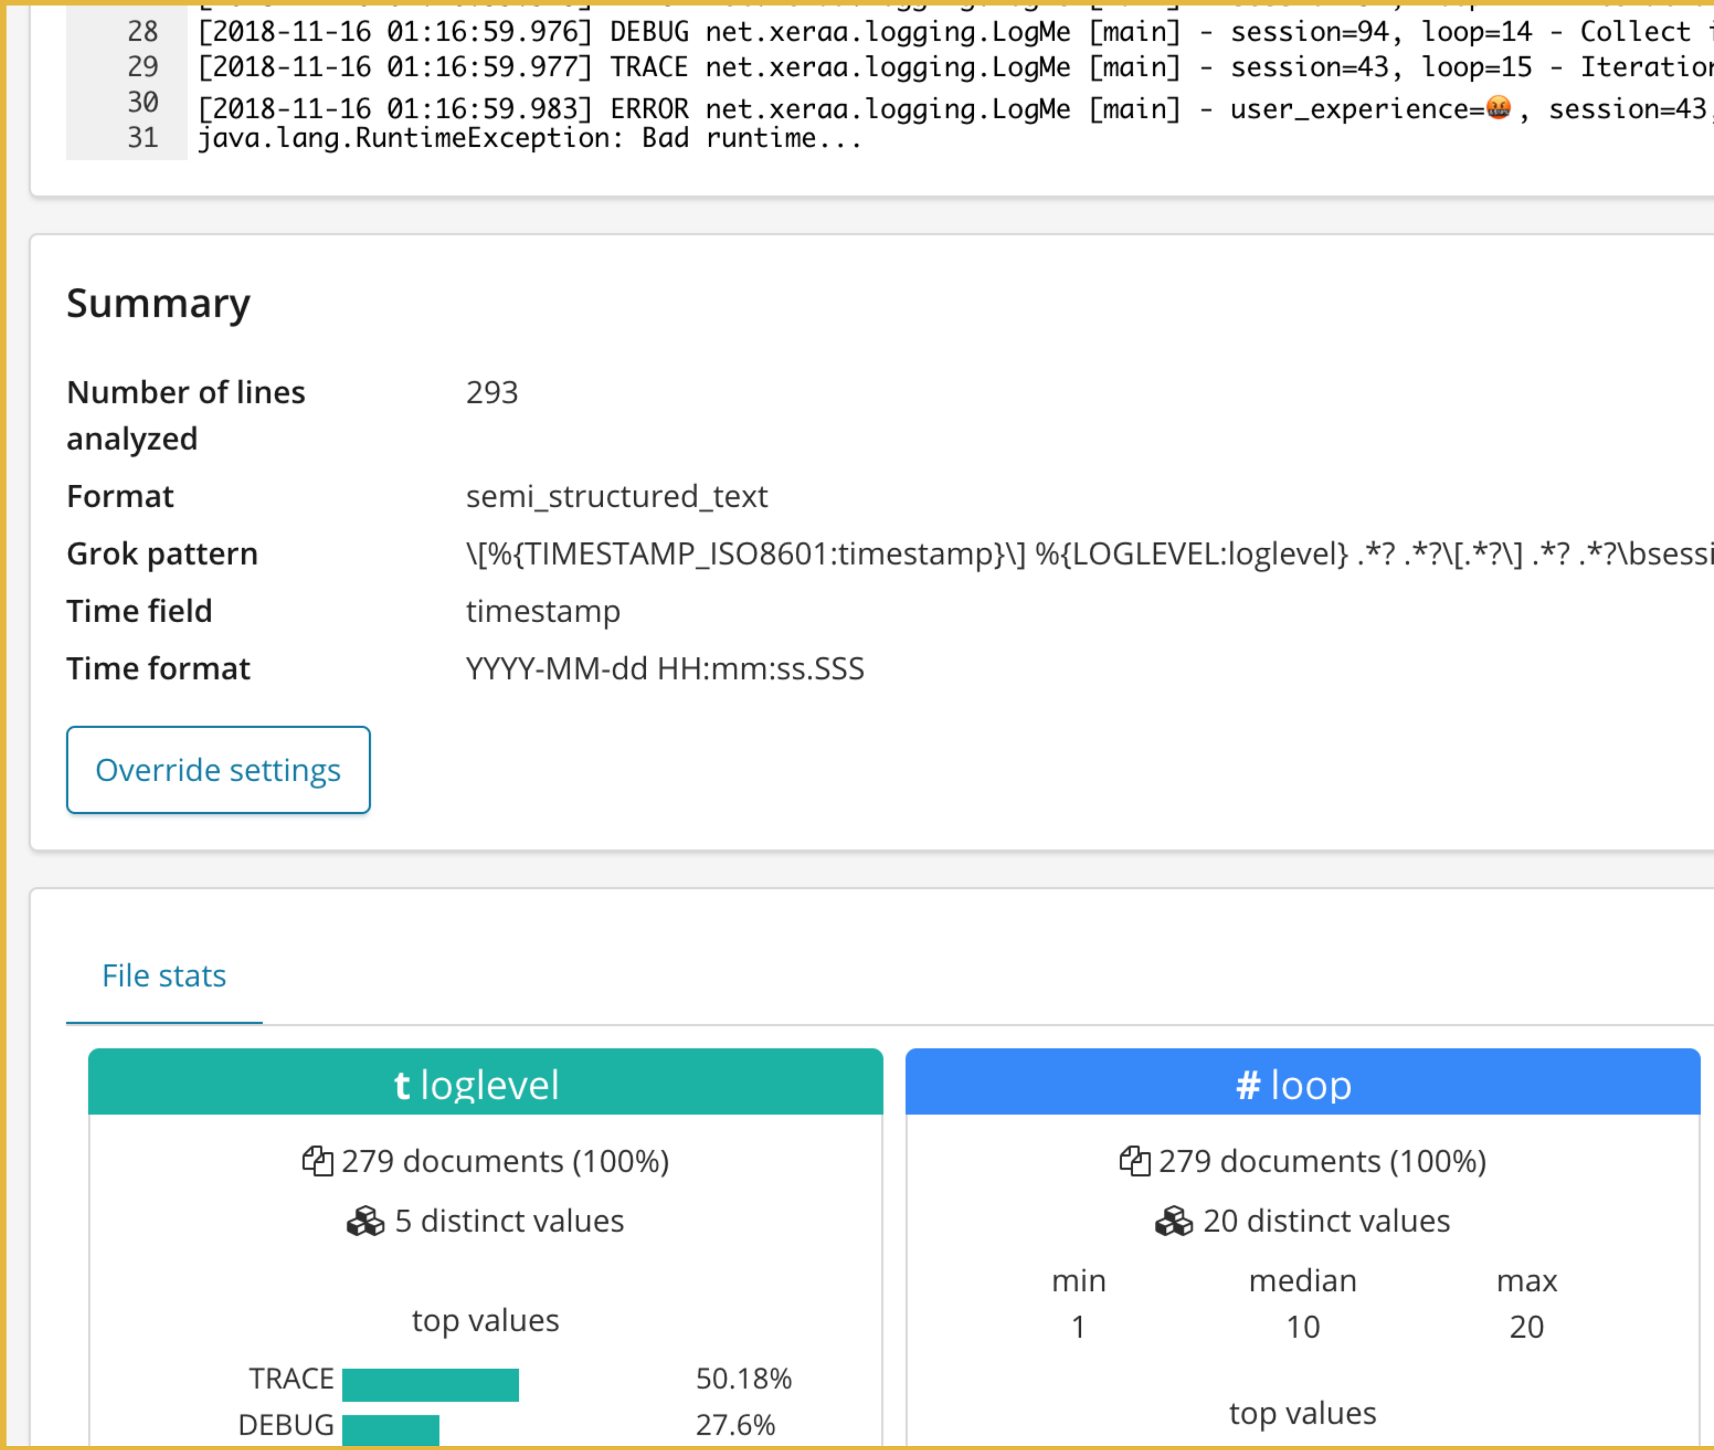

```
[2018-11-16 01:16:59.976] DEBUG net.xeraa.logging.LogMe [main] - session=94, loop=14 - Collect
[2018-11-16 01:16:59.977] TRACE net.xeraa.logging.LogMe [main] - session=43, loop=15 - Iteration
[2018-11-16 01:16:59.983] ERROR net.xeraa.logging.LogMe [main] - user_experience=🤬, session=43
java.lang.RuntimeException: Bad runtime...
```

Summary

Number of lines analyzed	293
Format	semi_structured_text
Grok pattern	\[%{TIMESTAMP_ISO8601:timestamp}\] %{LOGLEVEL:loglevel} .*? .*?\[.*?\] .*? .*?\bsessi
Time field	timestamp
Time format	YYYY-MM-dd HH:mm:ss.SSS

t loglevel

279 documents (100%)

5 distinct values

top values

TRACE 50.18%

DEBUG 27.6%

# loop

279 documents (100%)

20 distinct values

min	median	max
1	10	20

top values

elastic

Log UI

Infra UI

elastic

Filebeat

Raw events in Discover

/good: MDC logging under json.name and the context view for one log message

meta.* and host.* information

Filebeat

/bad and /null: Stacktraces by filtering down on application:java and json.severity:ERROR

Visualize json.stack_hash

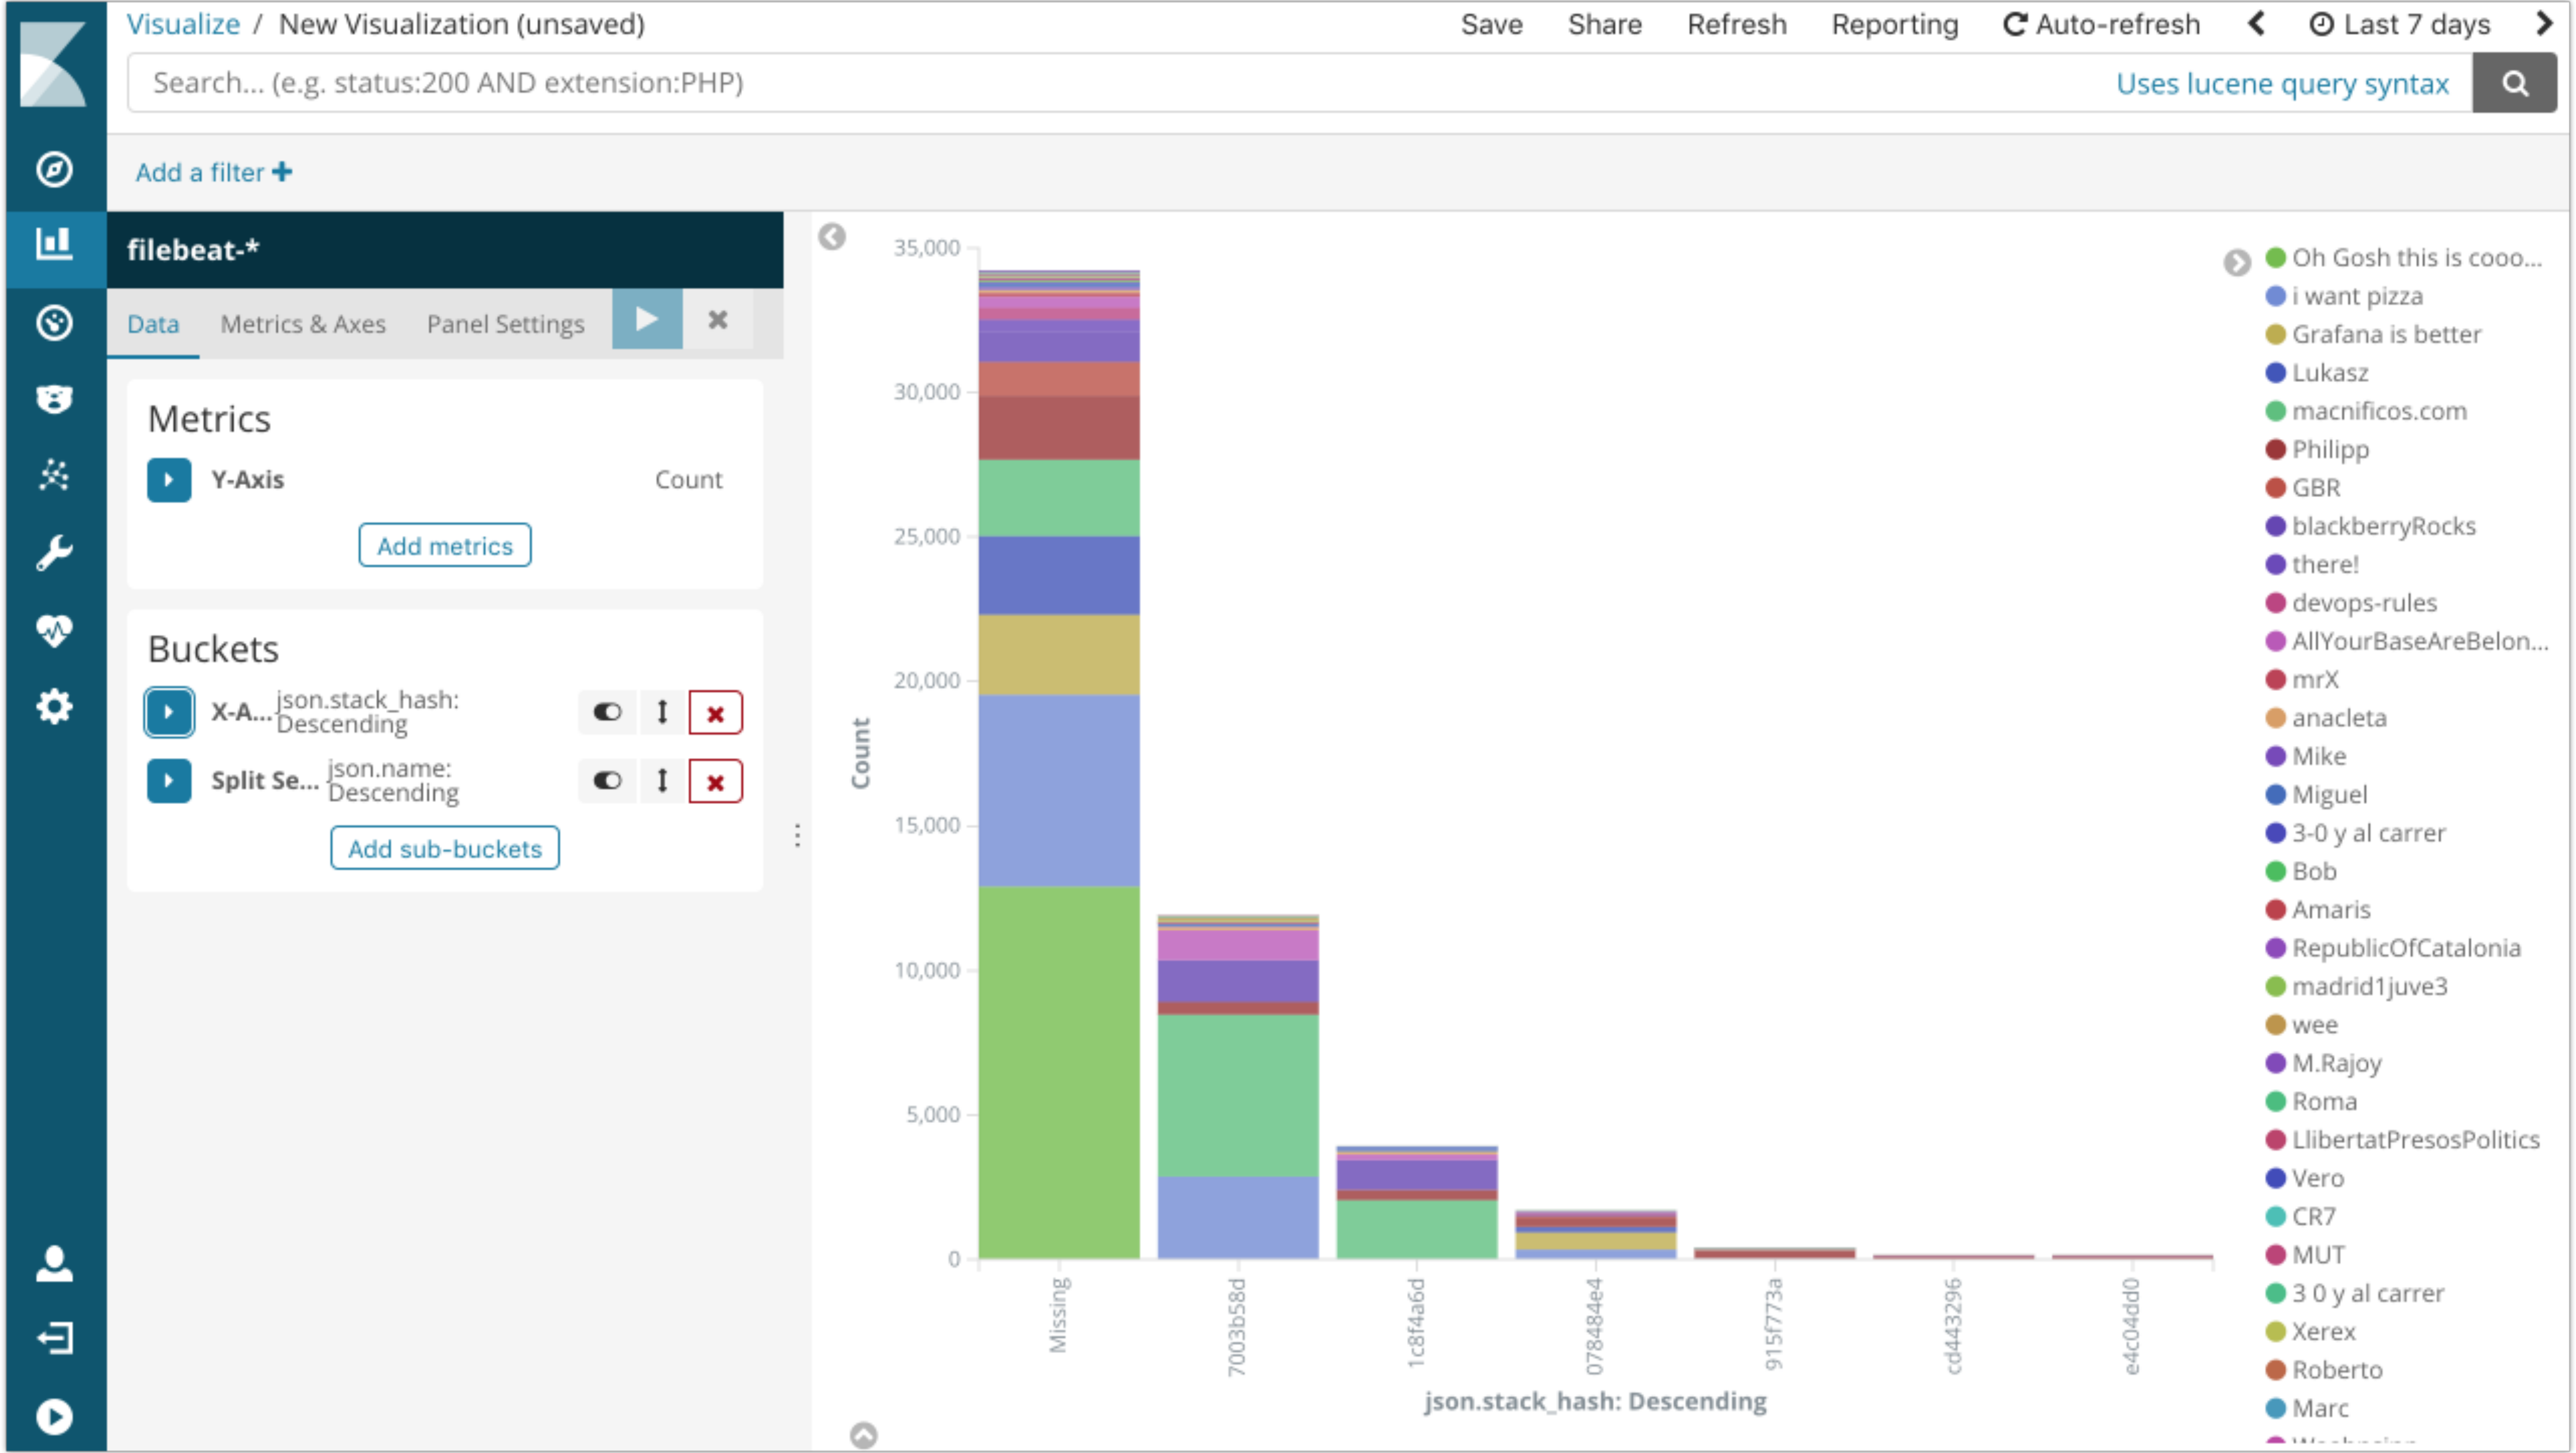

Visualize / New Visualization (unsaved)

Search... (e.g. status:200 AND extension:PHP)

filebeat-*

Metrics

Y-Axis Count

Buckets

X-A... json.stack_hash: Descending

Split Se... json.name: Descending

HEARTBEAT

Heartbeat HTTP monitoring dashboard

Stop and start the frontend application while auto refreshing

elastic

Metricbeat nginx

[Metricbeat Nginx] Overview dashboard

elastic

Metricbeat HTTP

/health and /metrics endpoints

Collected information in Discover

Metricbeat JMX

Same data

Visualize the heap usage: jolokia.metrics.memory.heap_usage.used divided by the max of jolokia.metrics.memory.heap_usage.max

Annotations

Add changes from the events index

elastic

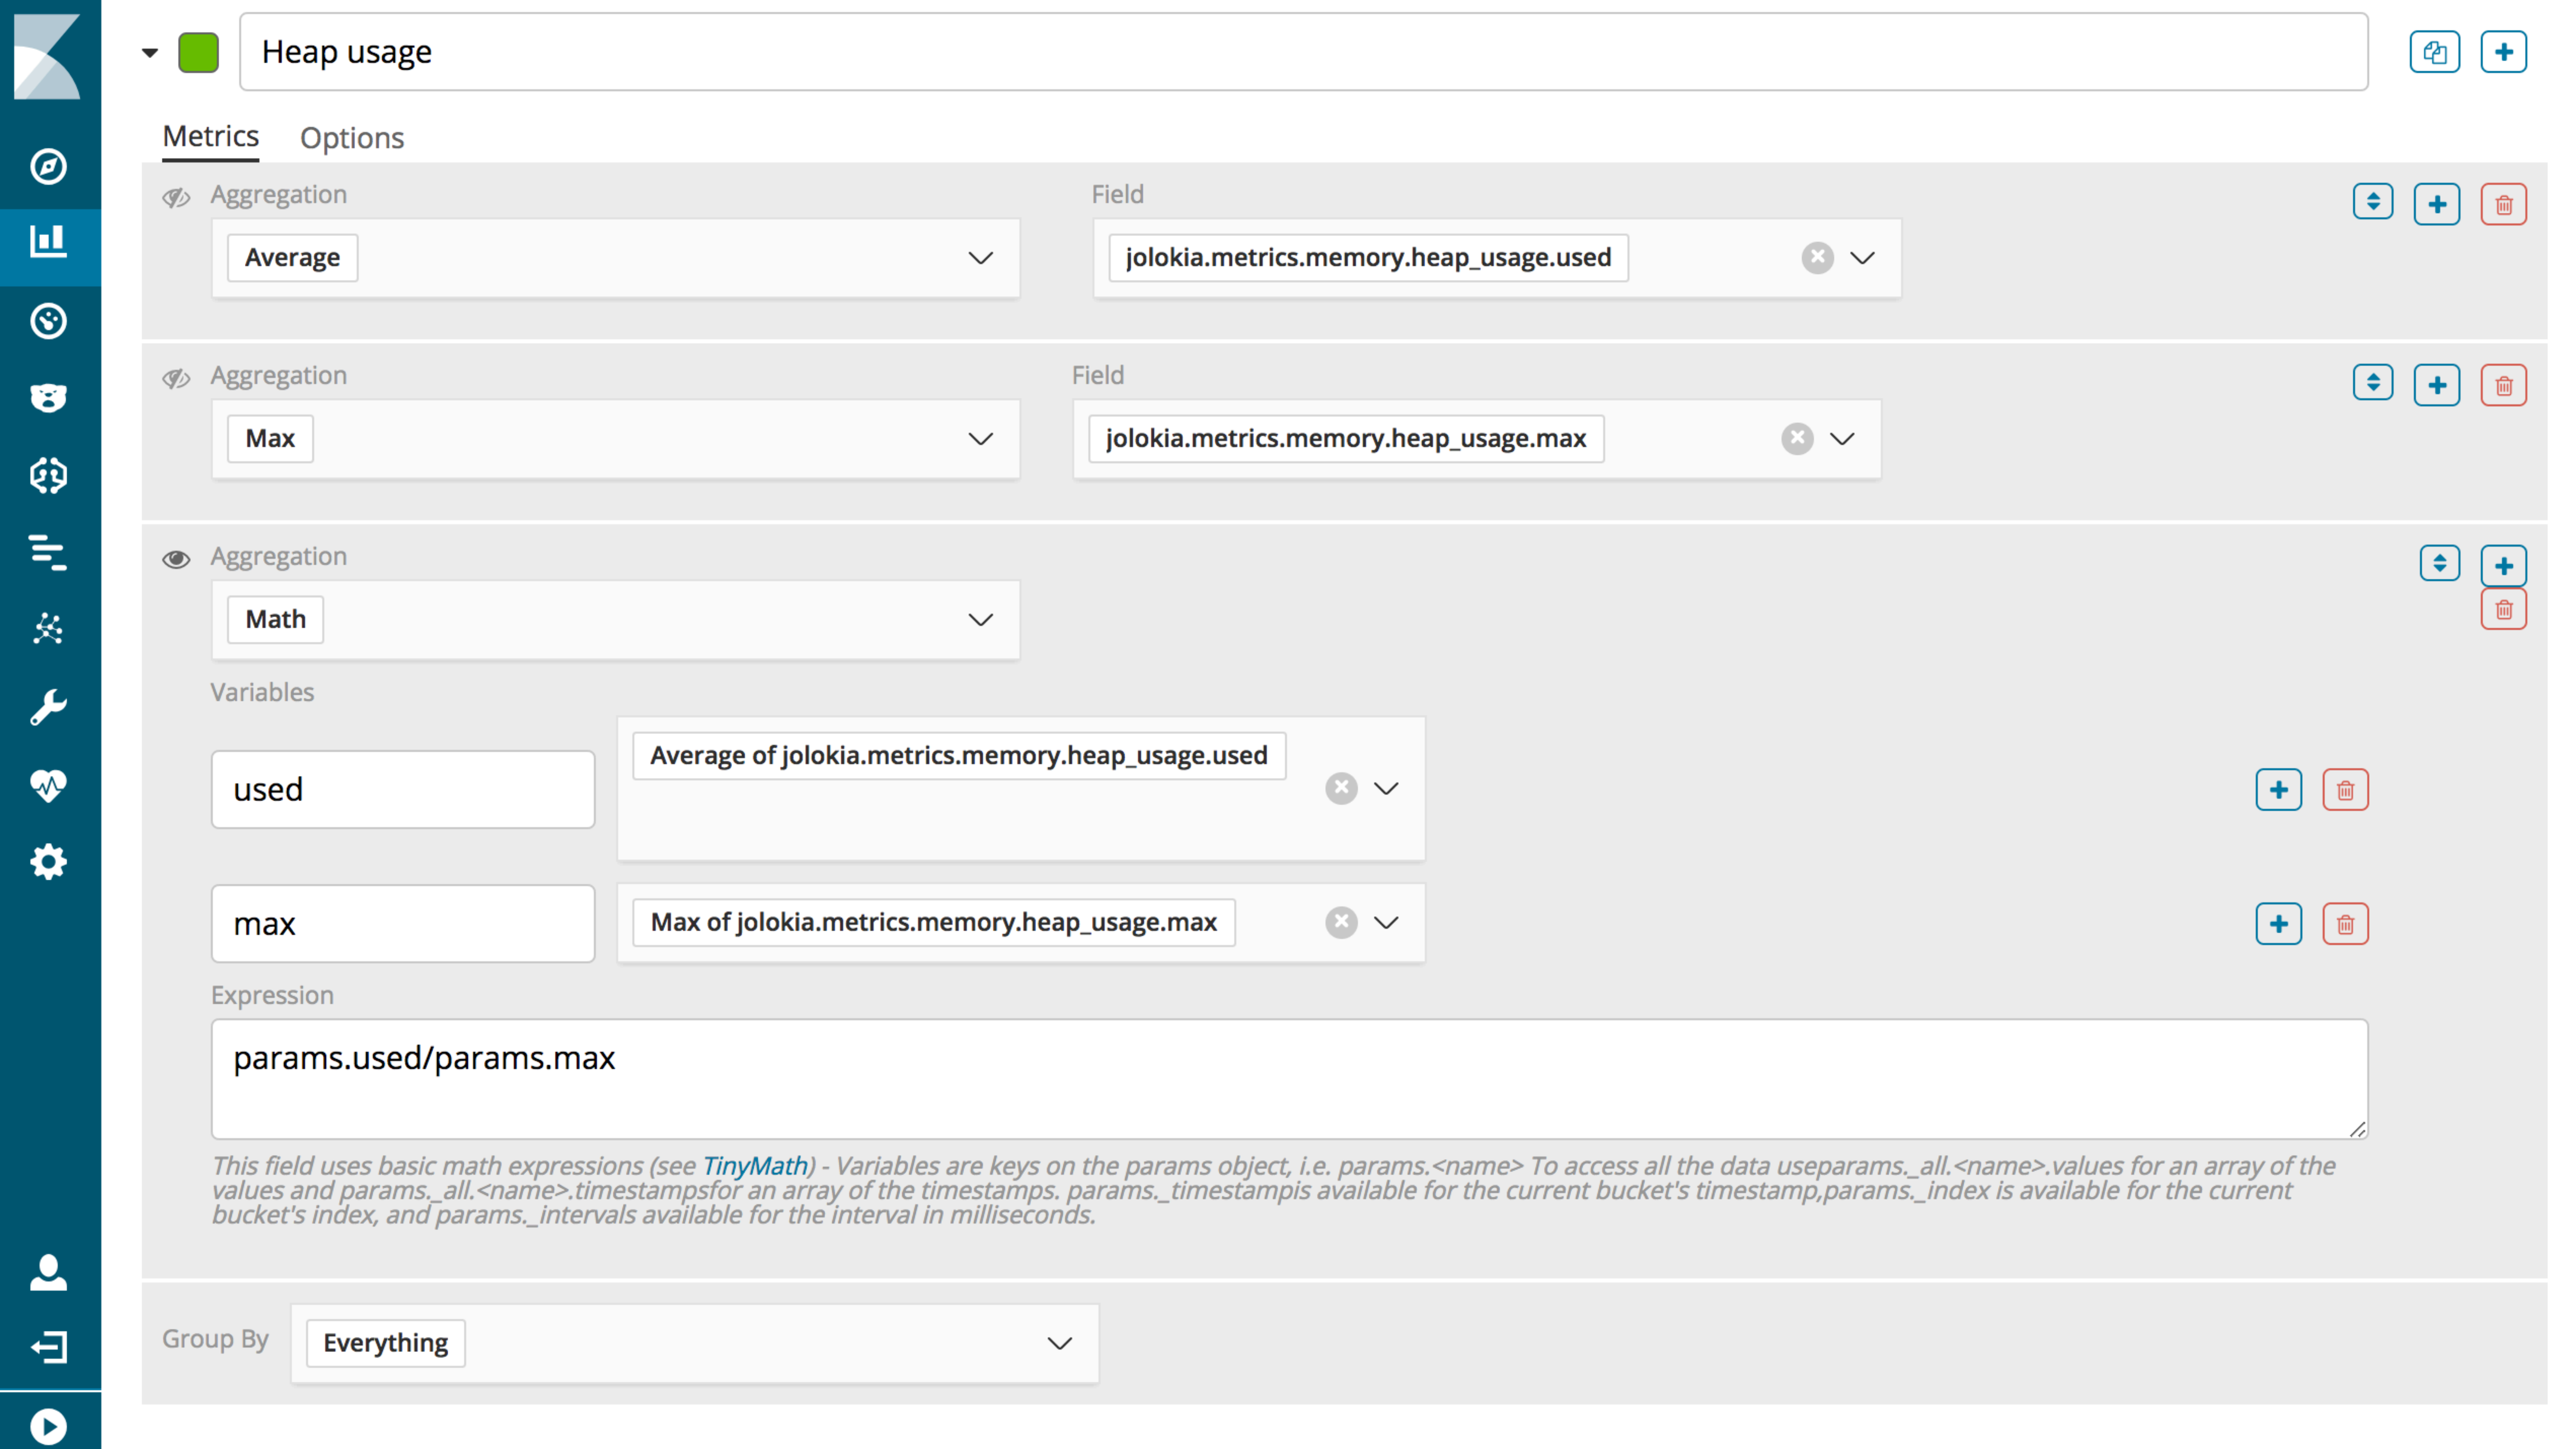

Heap usage

Aggregation: Average
Field: jolokia.metrics.memory.heap_usage.used

Aggregation: Max
Field: jolokia.metrics.memory.heap_usage.max

Aggregation: Math

Variables

used: Average of jolokia.metrics.memory.heap_usage.used

max: Max of jolokia.metrics.memory.heap_usage.max

Expression

```
params.used/params.max
```

This field uses basic math expressions (see TinyMath) - Variables are keys on the params object, i.e. params.<name> To access all the data useparams._all.<name>.values for an array of the values and params._all.<name>.timestampsfor an array of the timestamps. params._timestampis available for the current bucket's timestamp,params._index is available for the current bucket's index, and params._intervals available for the interval in milliseconds.

Group By: Everything

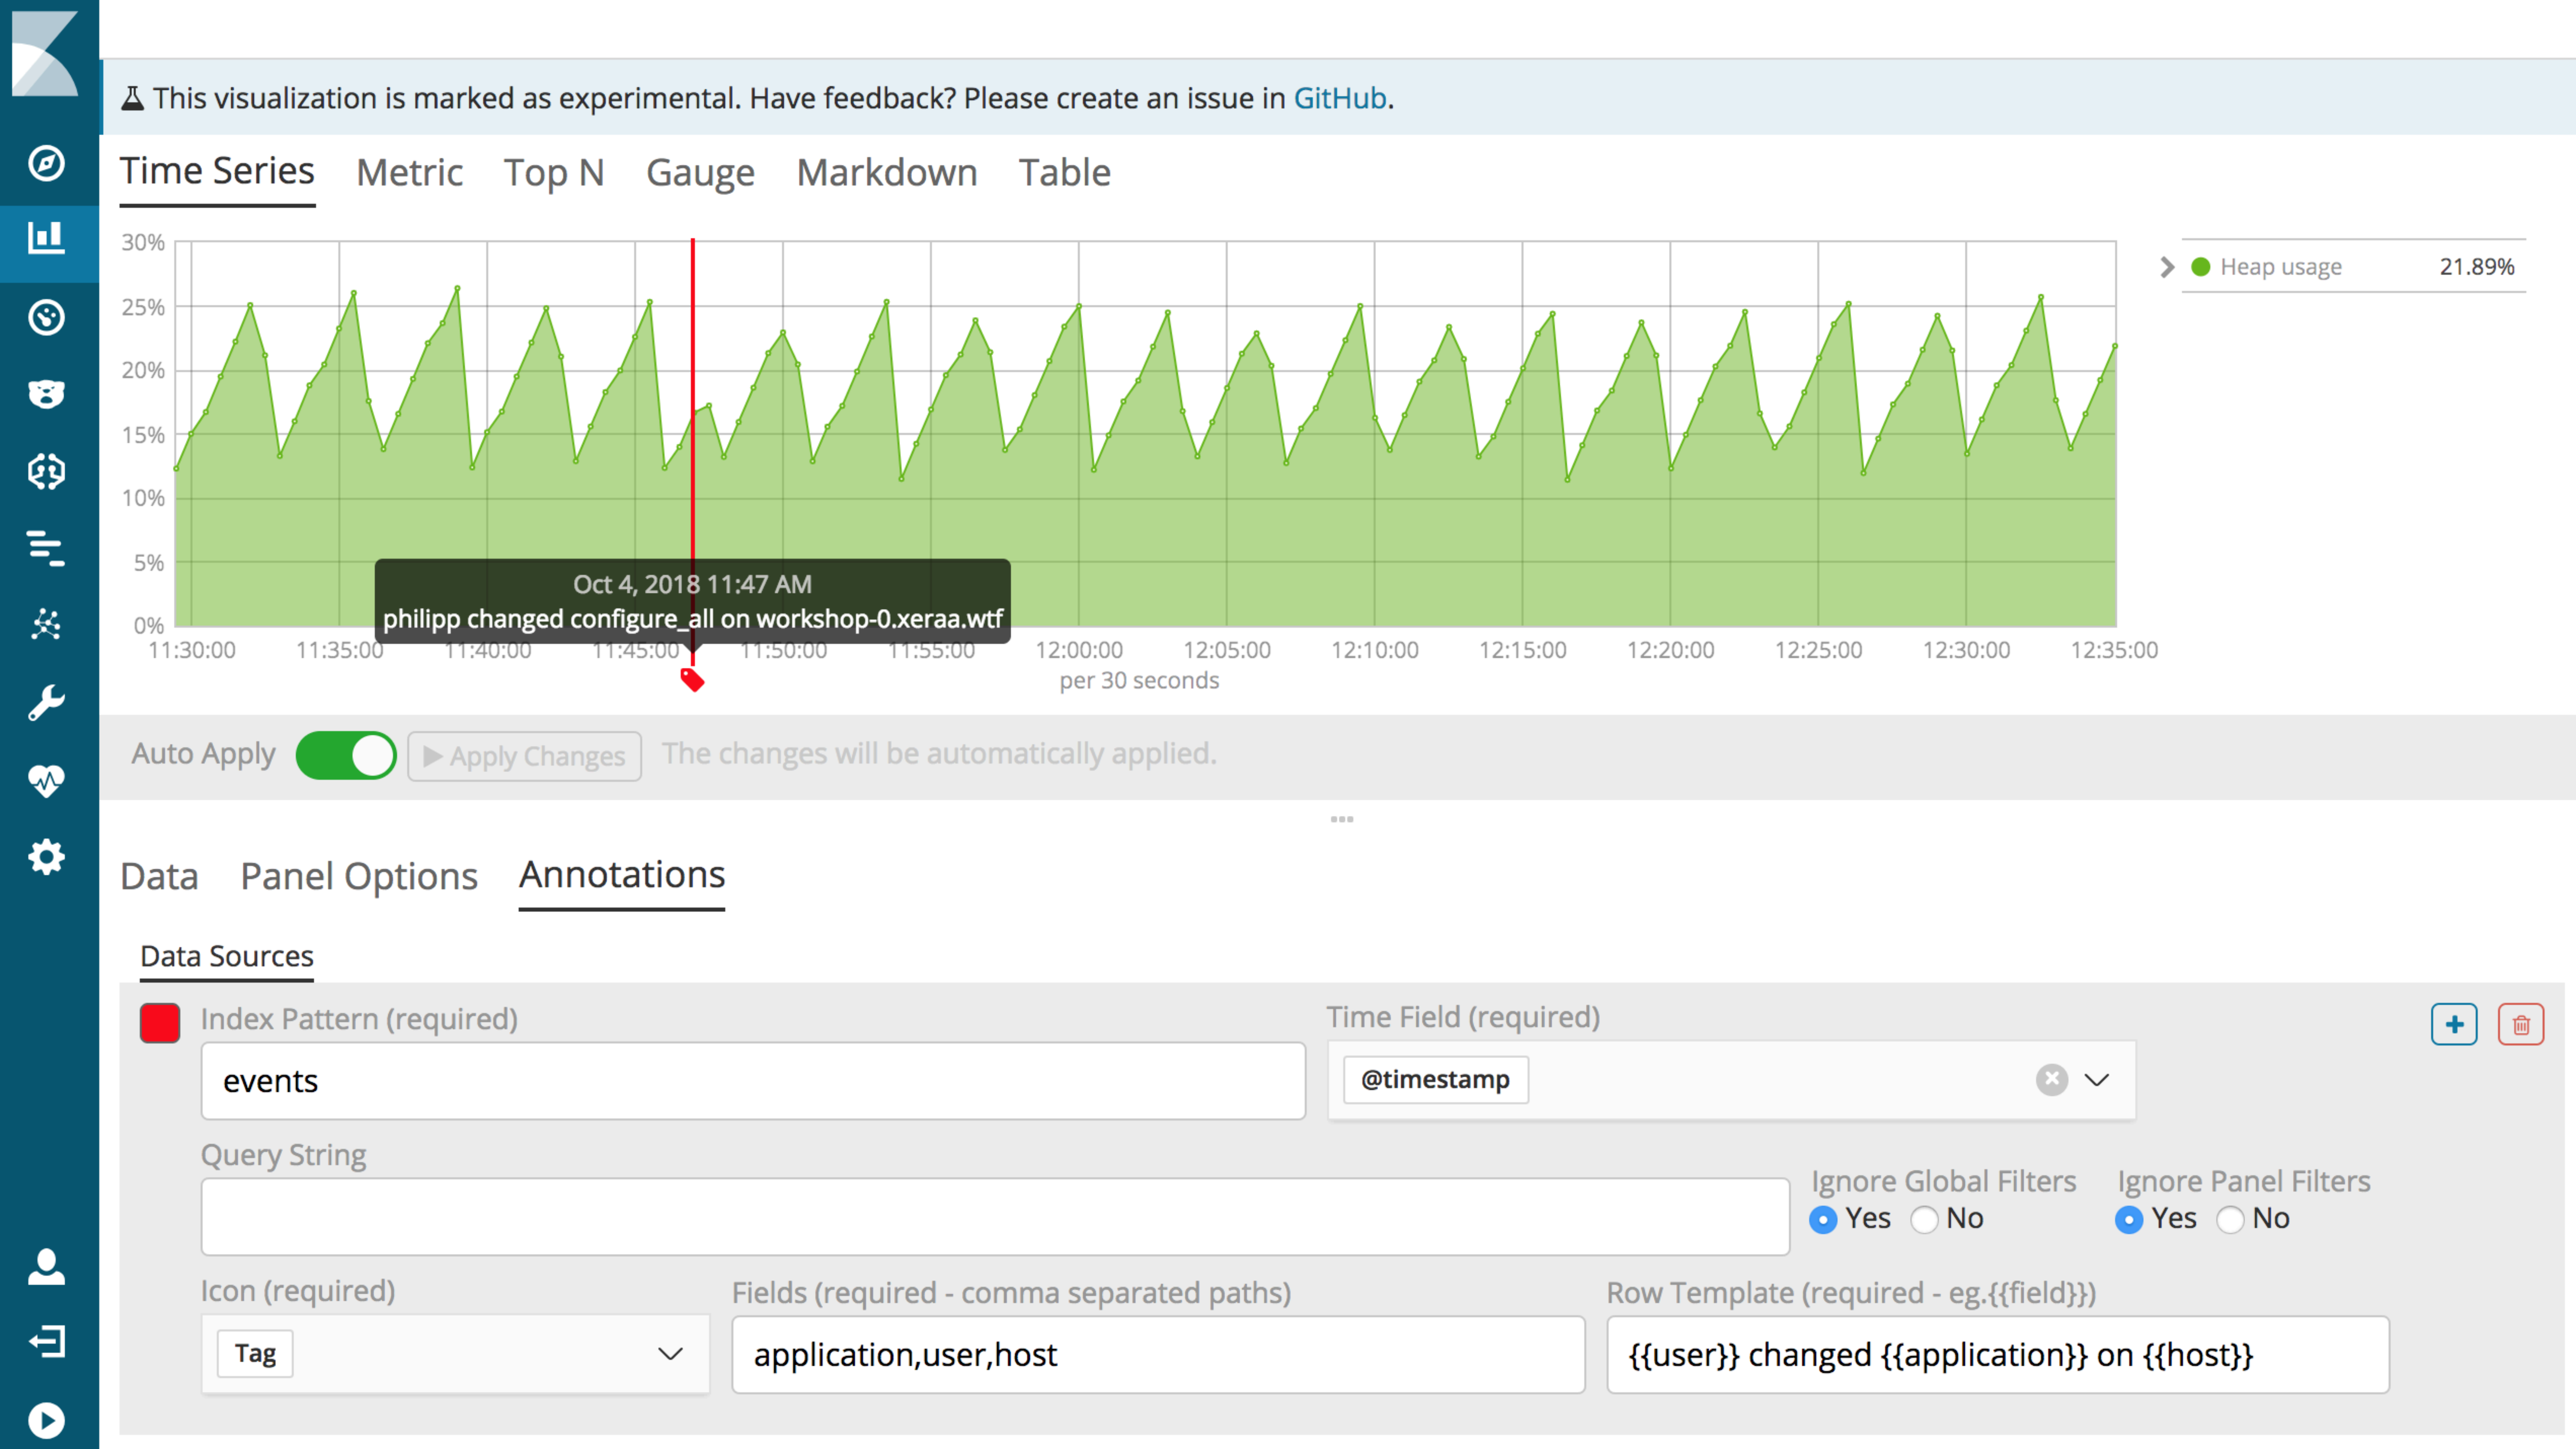

This visualization is marked as experimental. Have feedback? Please create an issue in GitHub.

Oct 4, 2018 11:47 AM
philipp changed configure_all on workshop-0.xeraa.wtf

The changes will be automatically applied.

Index Pattern (required): events

Time Field (required): @timestamp

Query String:

Ignore Global Filters: ● Yes ○ No

Ignore Panel Filters: ● Yes ○ No

Icon (required): Tag

Fields (required - comma separated paths): application,user,host

Row Template (required - eg.{{field}}): {{user}} changed {{application}} on {{host}}

APM

Distributed Tracing

More Features

elastic cloud

Index Lifecycle Management

Currently

https://github.com/elastic/curator

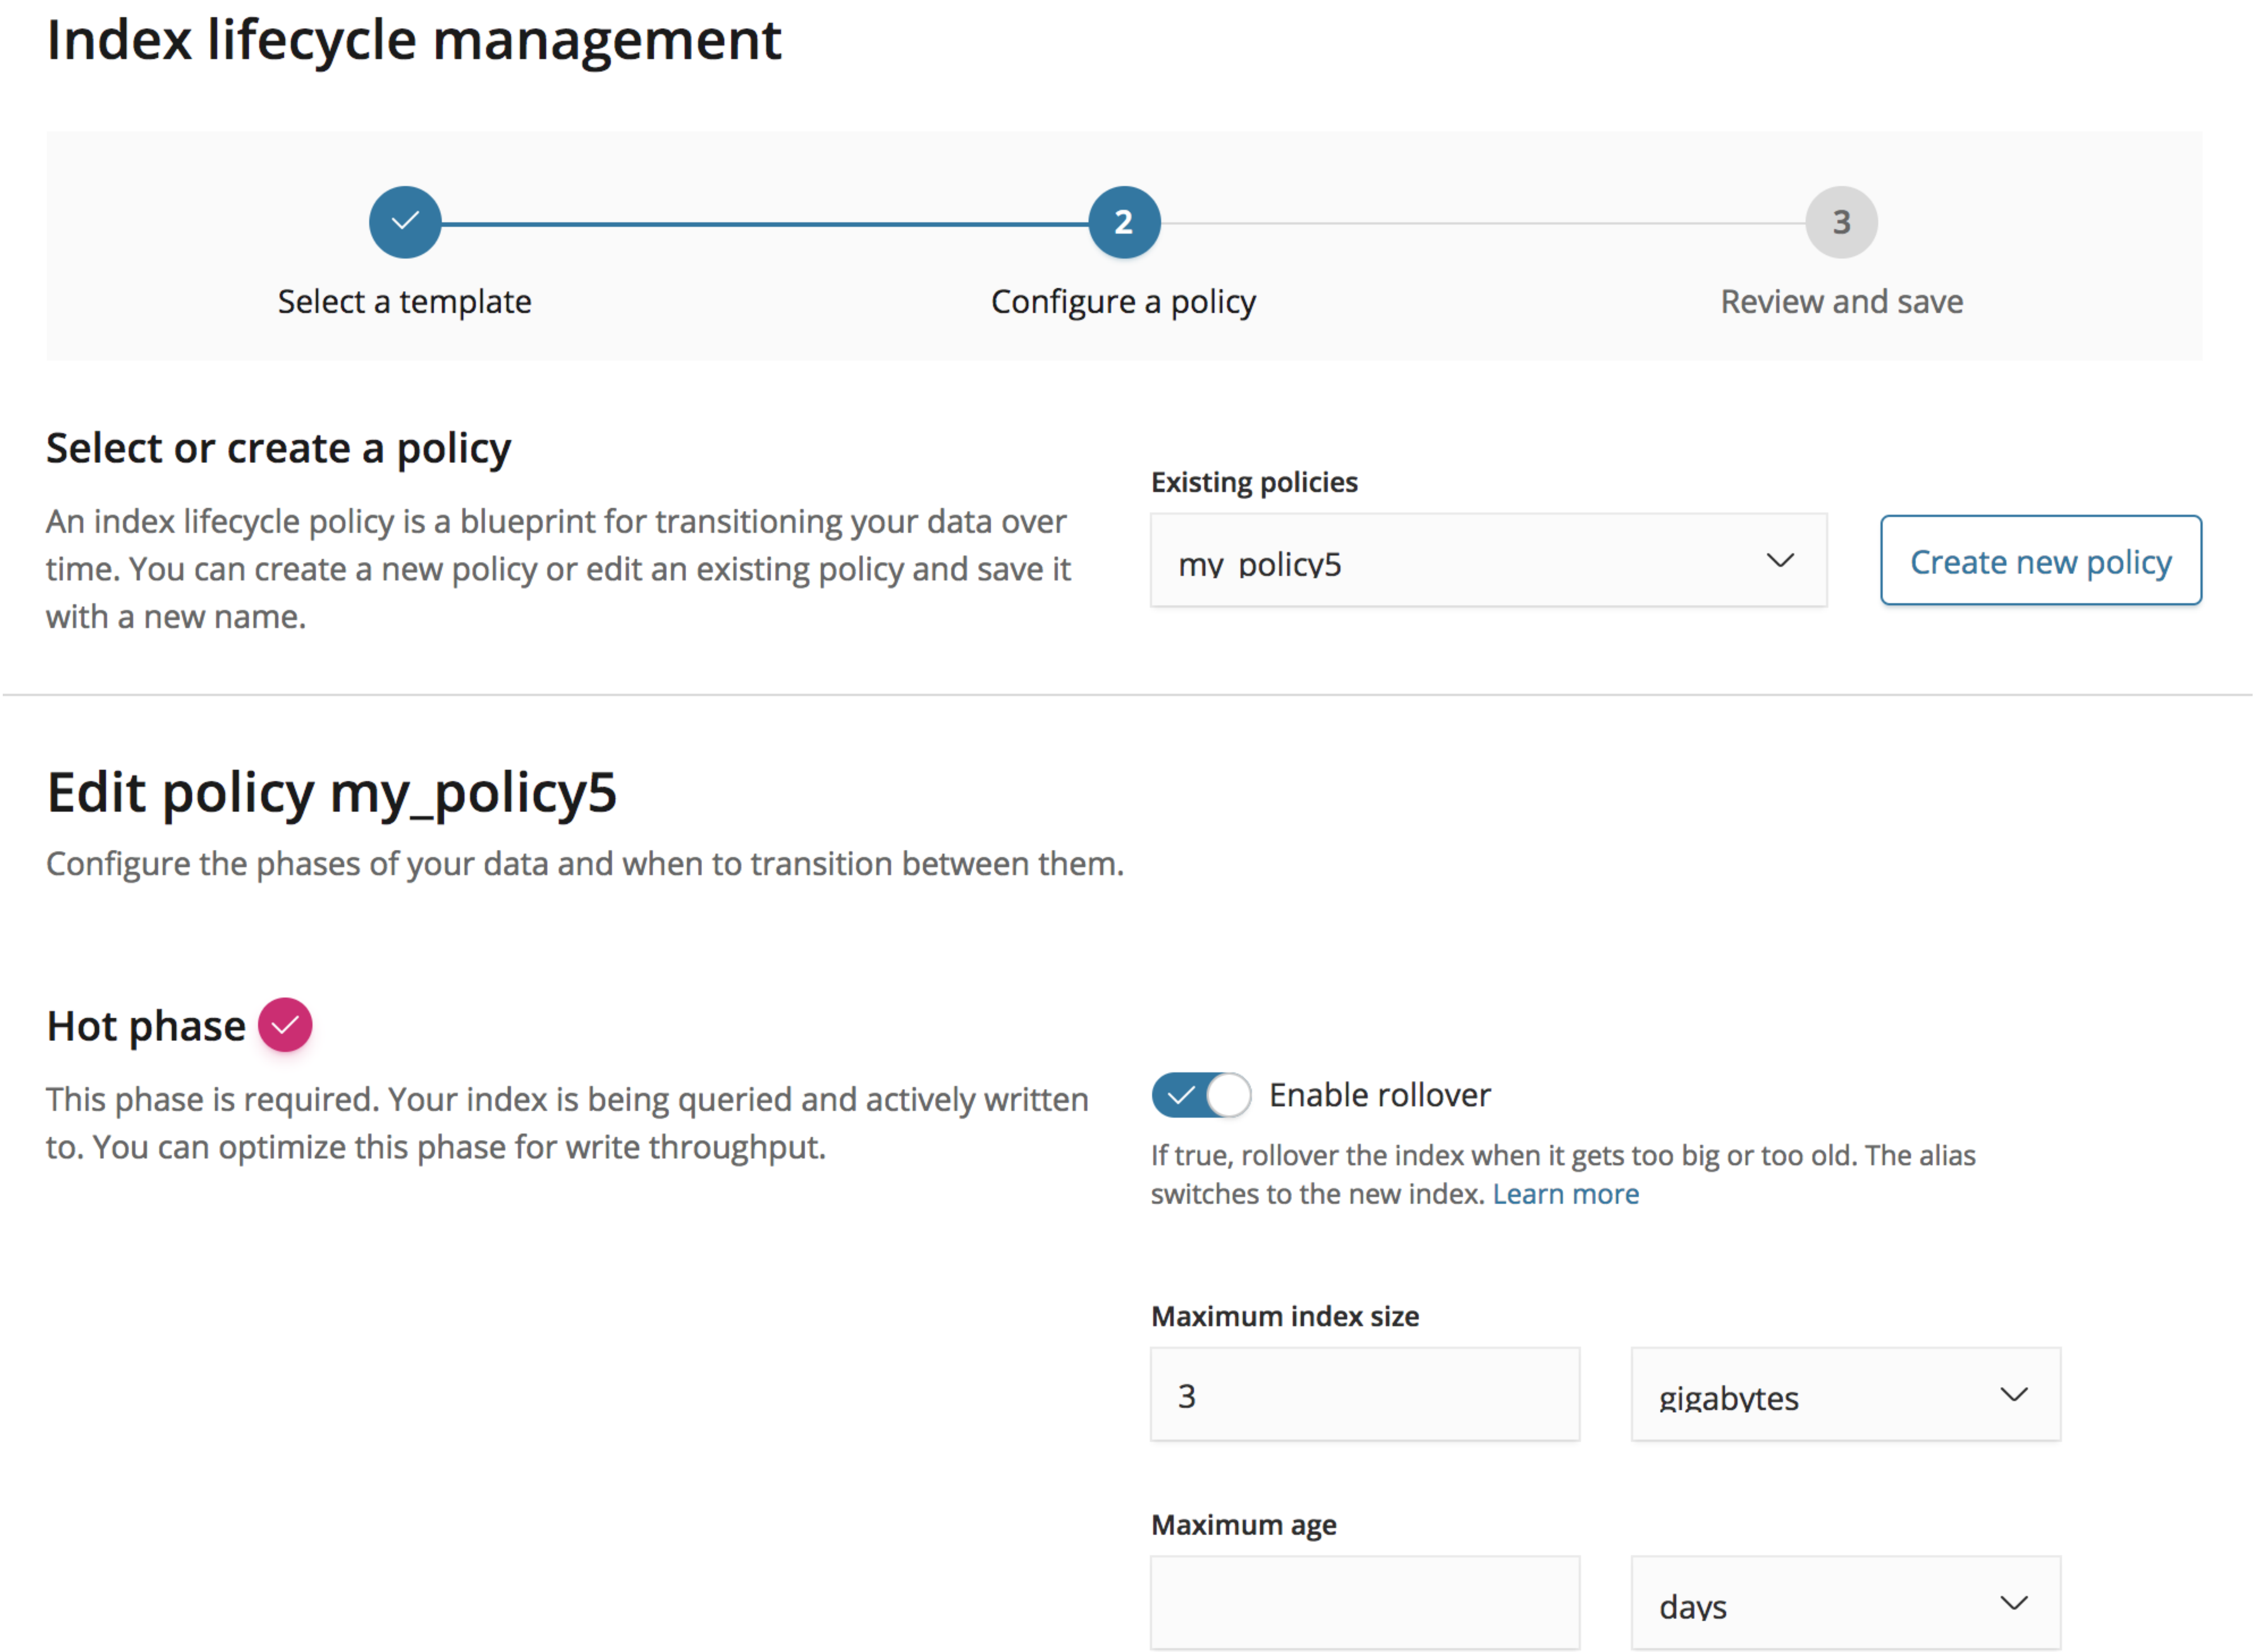

Index lifecycle management

Select or create a policy

An index lifecycle policy is a blueprint for transitioning your data over time. You can create a new policy or edit an existing policy and save it with a new name.

Existing policies: my_policy5

Edit policy my_policy5

Configure the phases of your data and when to transition between them.

Hot phase

This phase is required. Your index is being queried and actively written to. You can optimize this phase for write throughput.

Enable rollover: ✓

If true, rollover the index when it gets too big or too old. The alias switches to the new index. Learn more

Maximum index size: 3 gigabytes

Maximum age: days

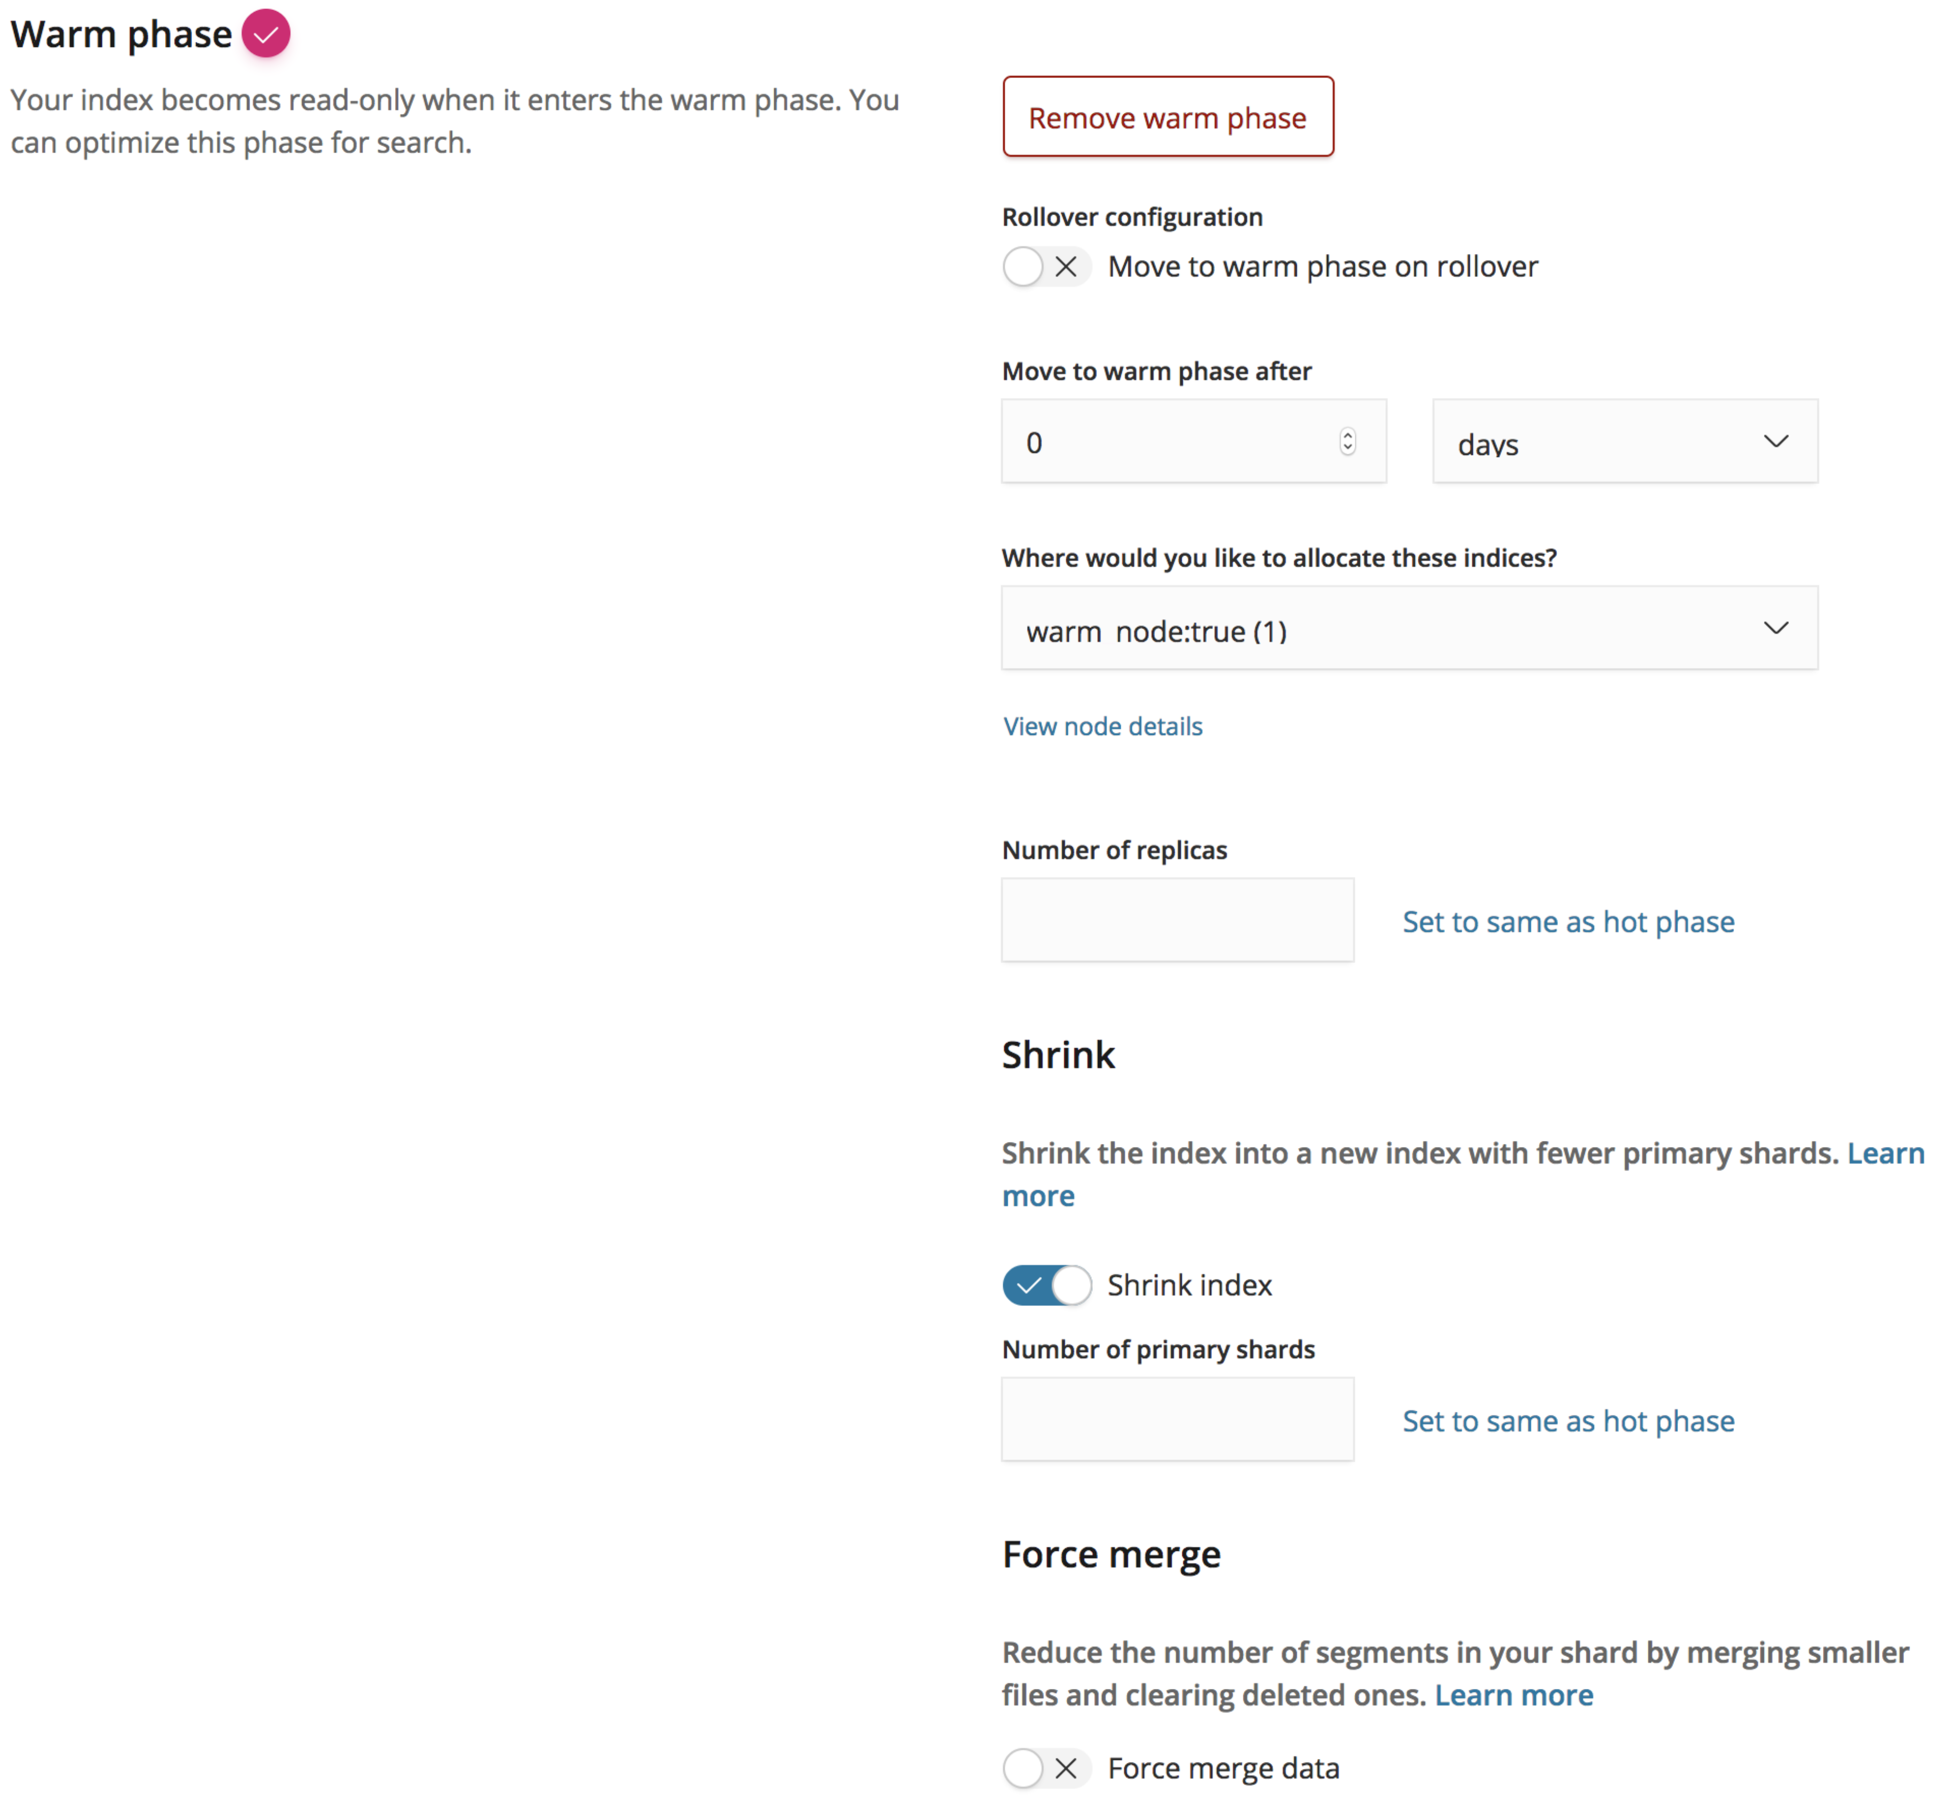

Warm phase ✓

Your index becomes read-only when it enters the warm phase. You can optimize this phase for search.

Remove warm phase

Rollover configuration

☒ Move to warm phase on rollover

Move to warm phase after

0 | days

Where would you like to allocate these indices?

warm node:true (1)

View node details

Number of replicas

Set to same as hot phase

Shrink

Shrink the index into a new index with fewer primary shards. Learn more

☑ Shrink index

Number of primary shards

Set to same as hot phase

Force merge

Reduce the number of segments in your shard by merging smaller files and clearing deleted ones. Learn more

☒ Force merge data

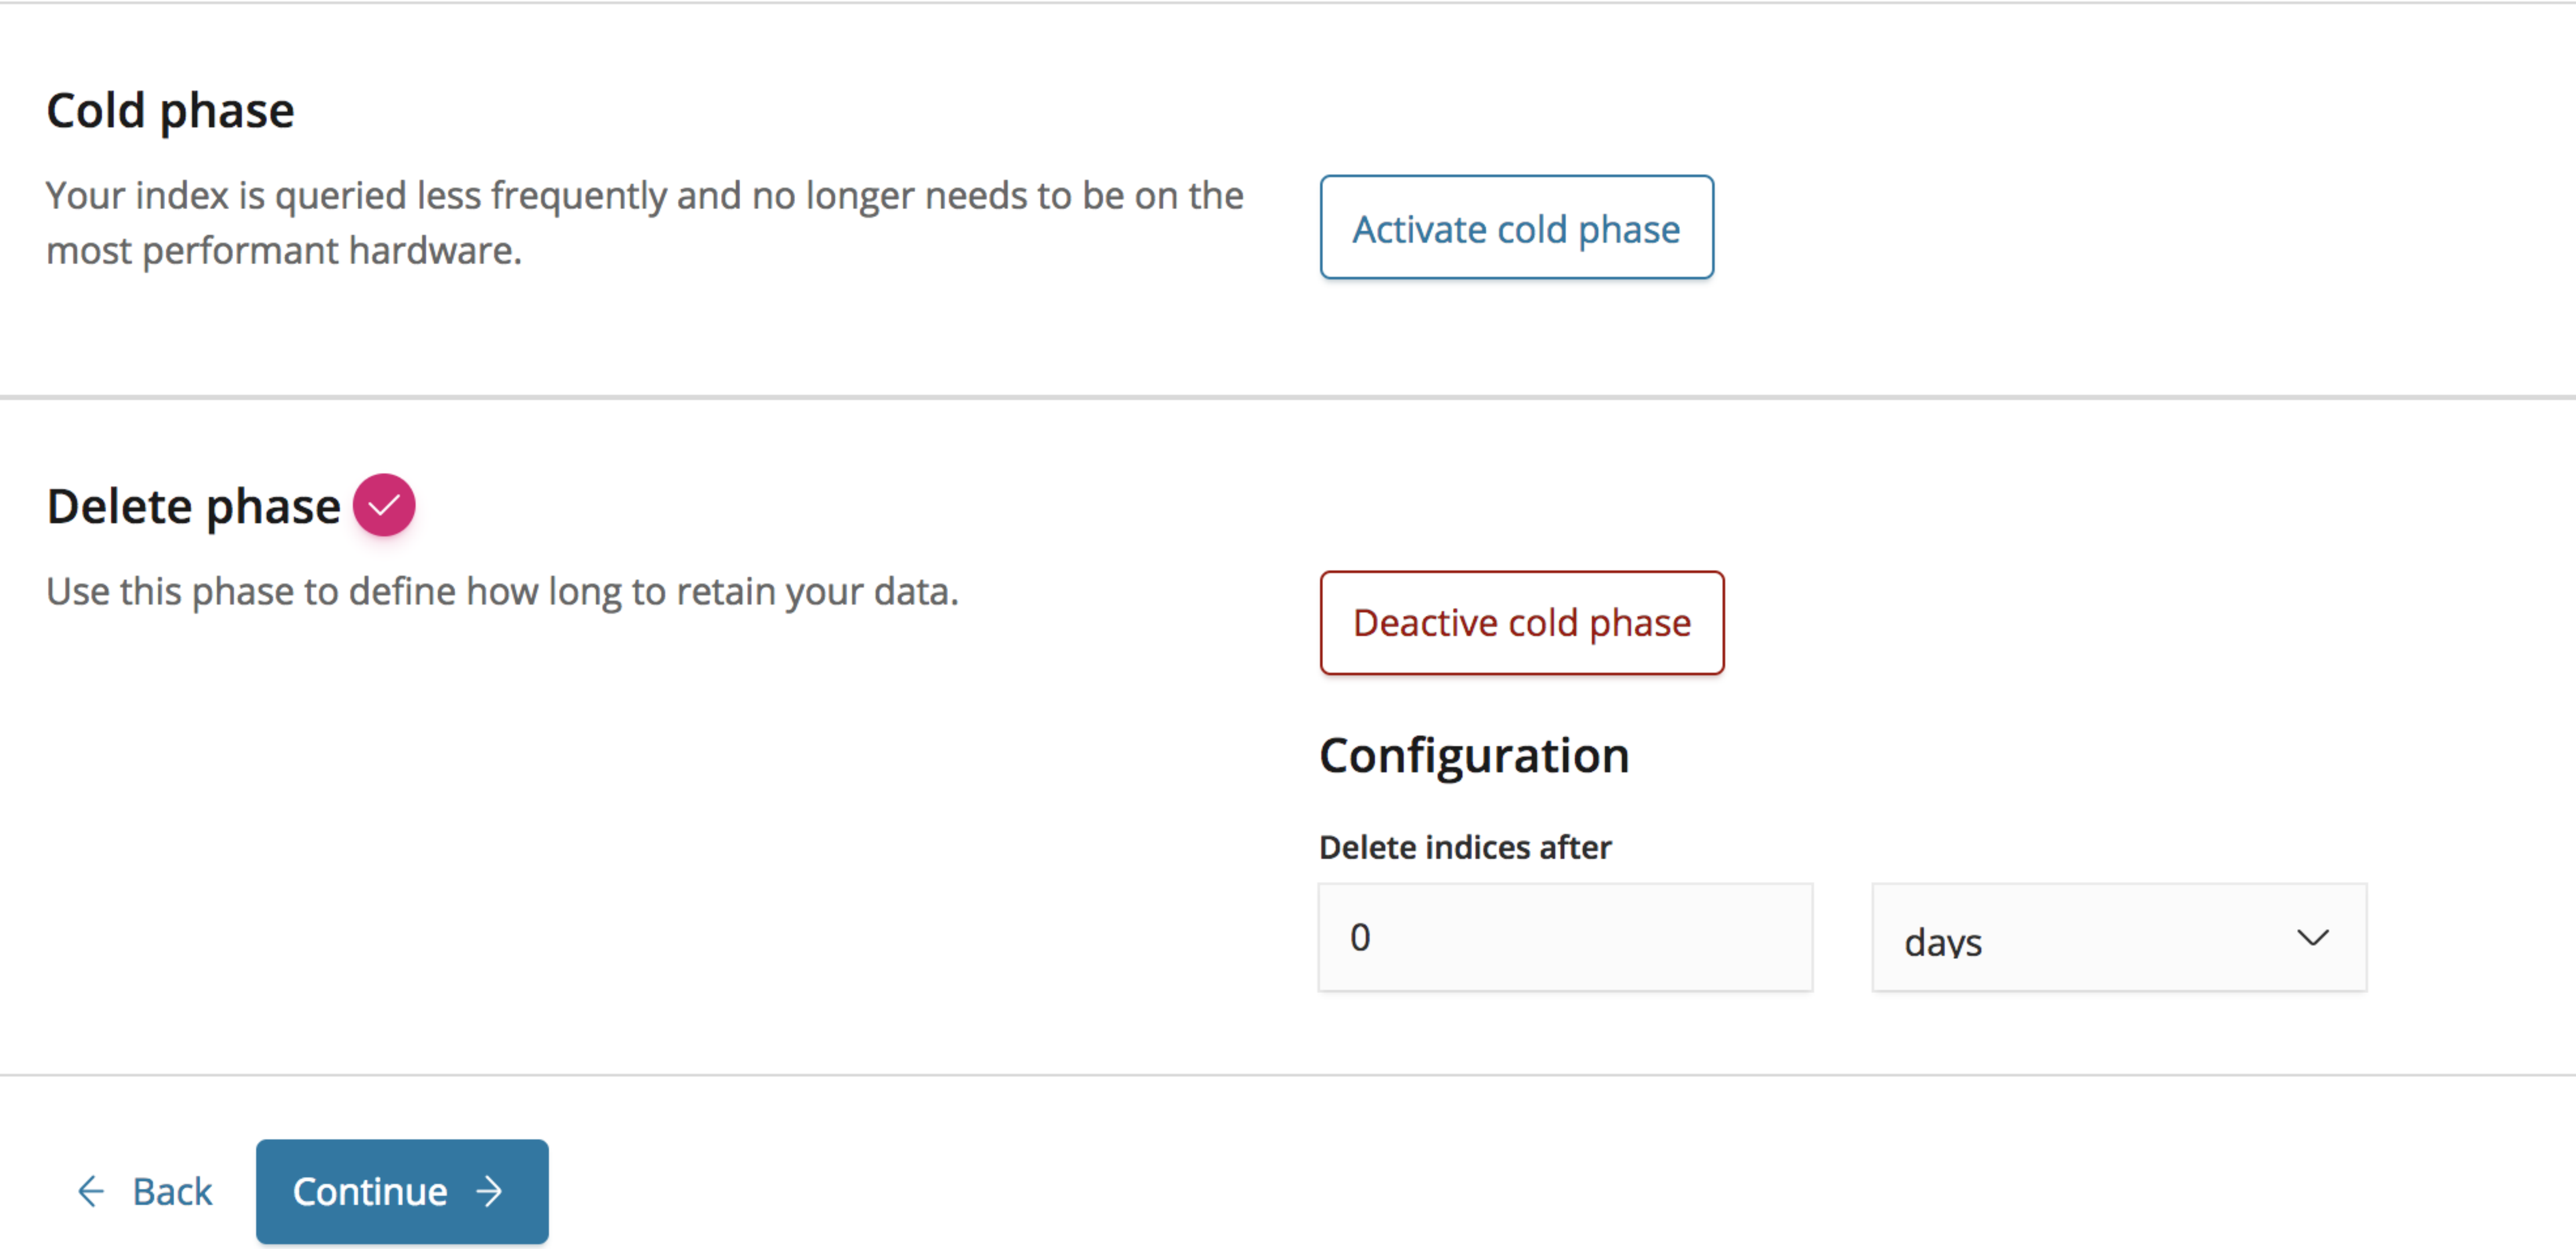

Cold phase

Your index is queried less frequently and no longer needs to be on the most performant hardware.

Delete phase ✓

Use this phase to define how long to retain your data.

Configuration

Delete indices after

0 days

Alerting[a]

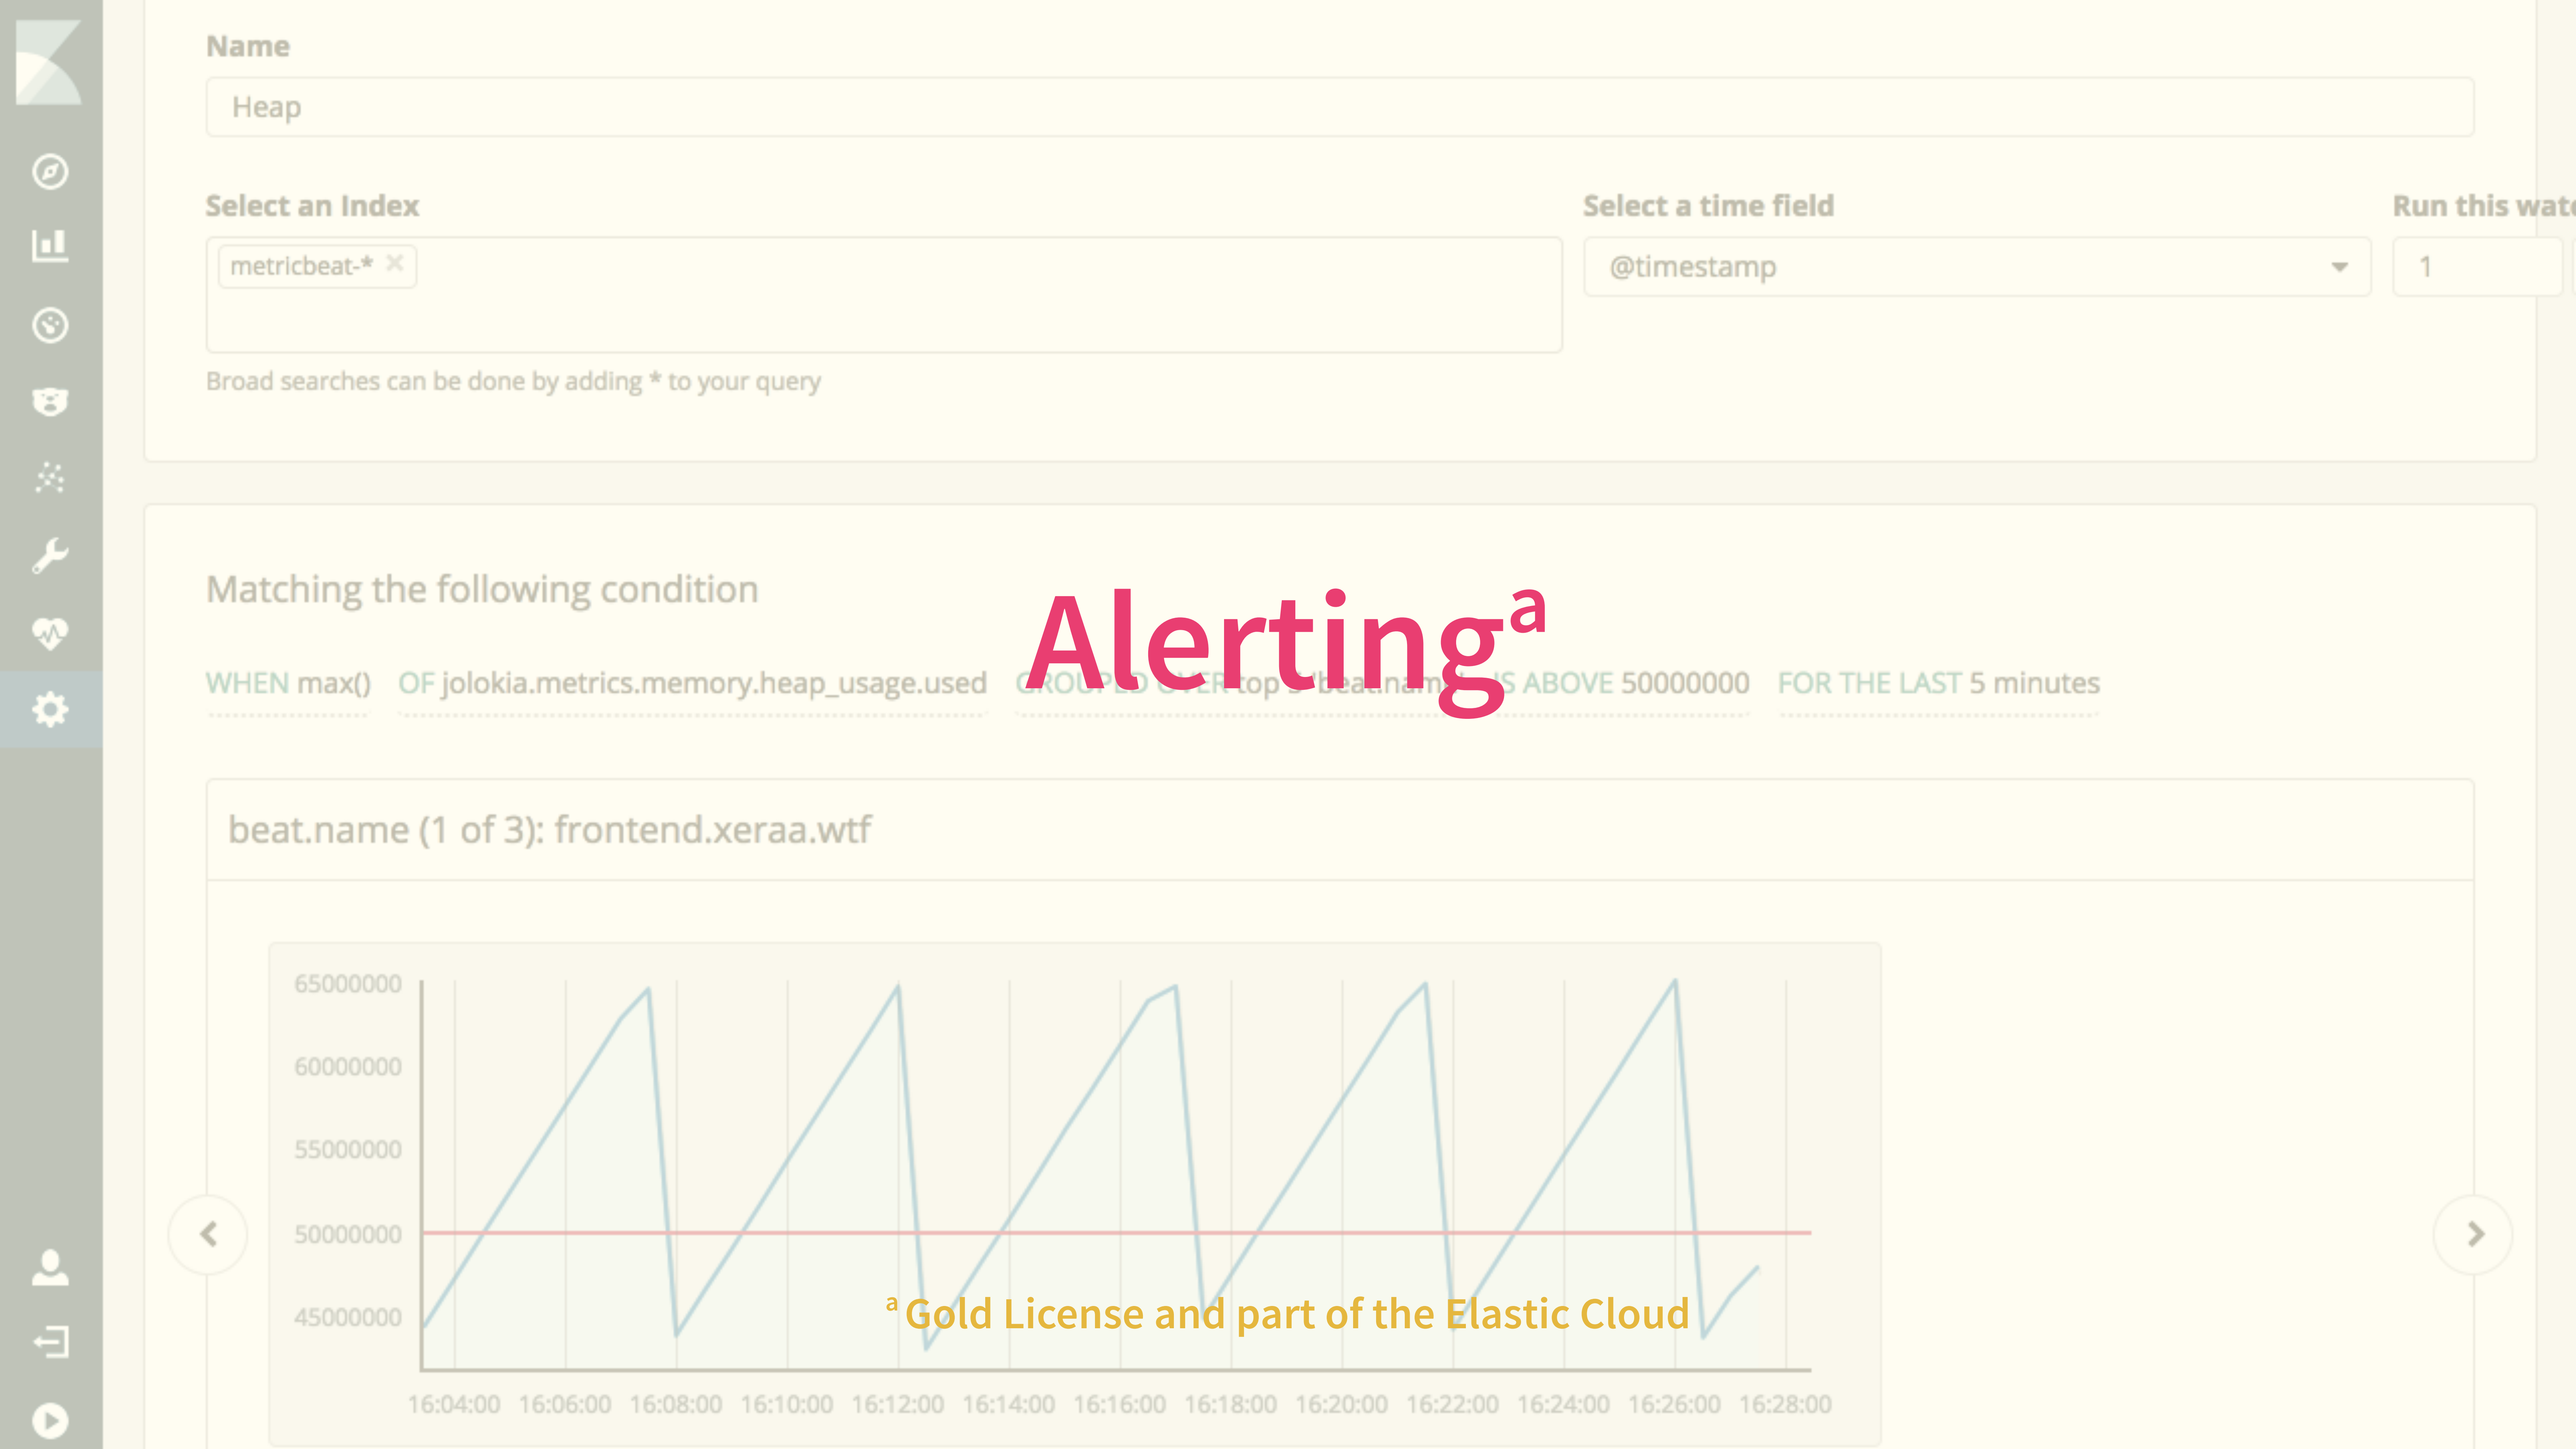

Name

Heap

Select an Index

metricbeat-* ✕

Broad searches can be done by adding * to your query

Select a time field

@timestamp

Run this wat

1

Matching the following condition

WHEN max() OF jolokia.metrics.memory.heap_usage.used GROUPED OVER top 3 'beat.name' IS ABOVE 50000000 FOR THE LAST 5 minutes

beat.name (1 of 3): frontend.xeraa.wtf

[a] Gold License and part of the Elastic Cloud

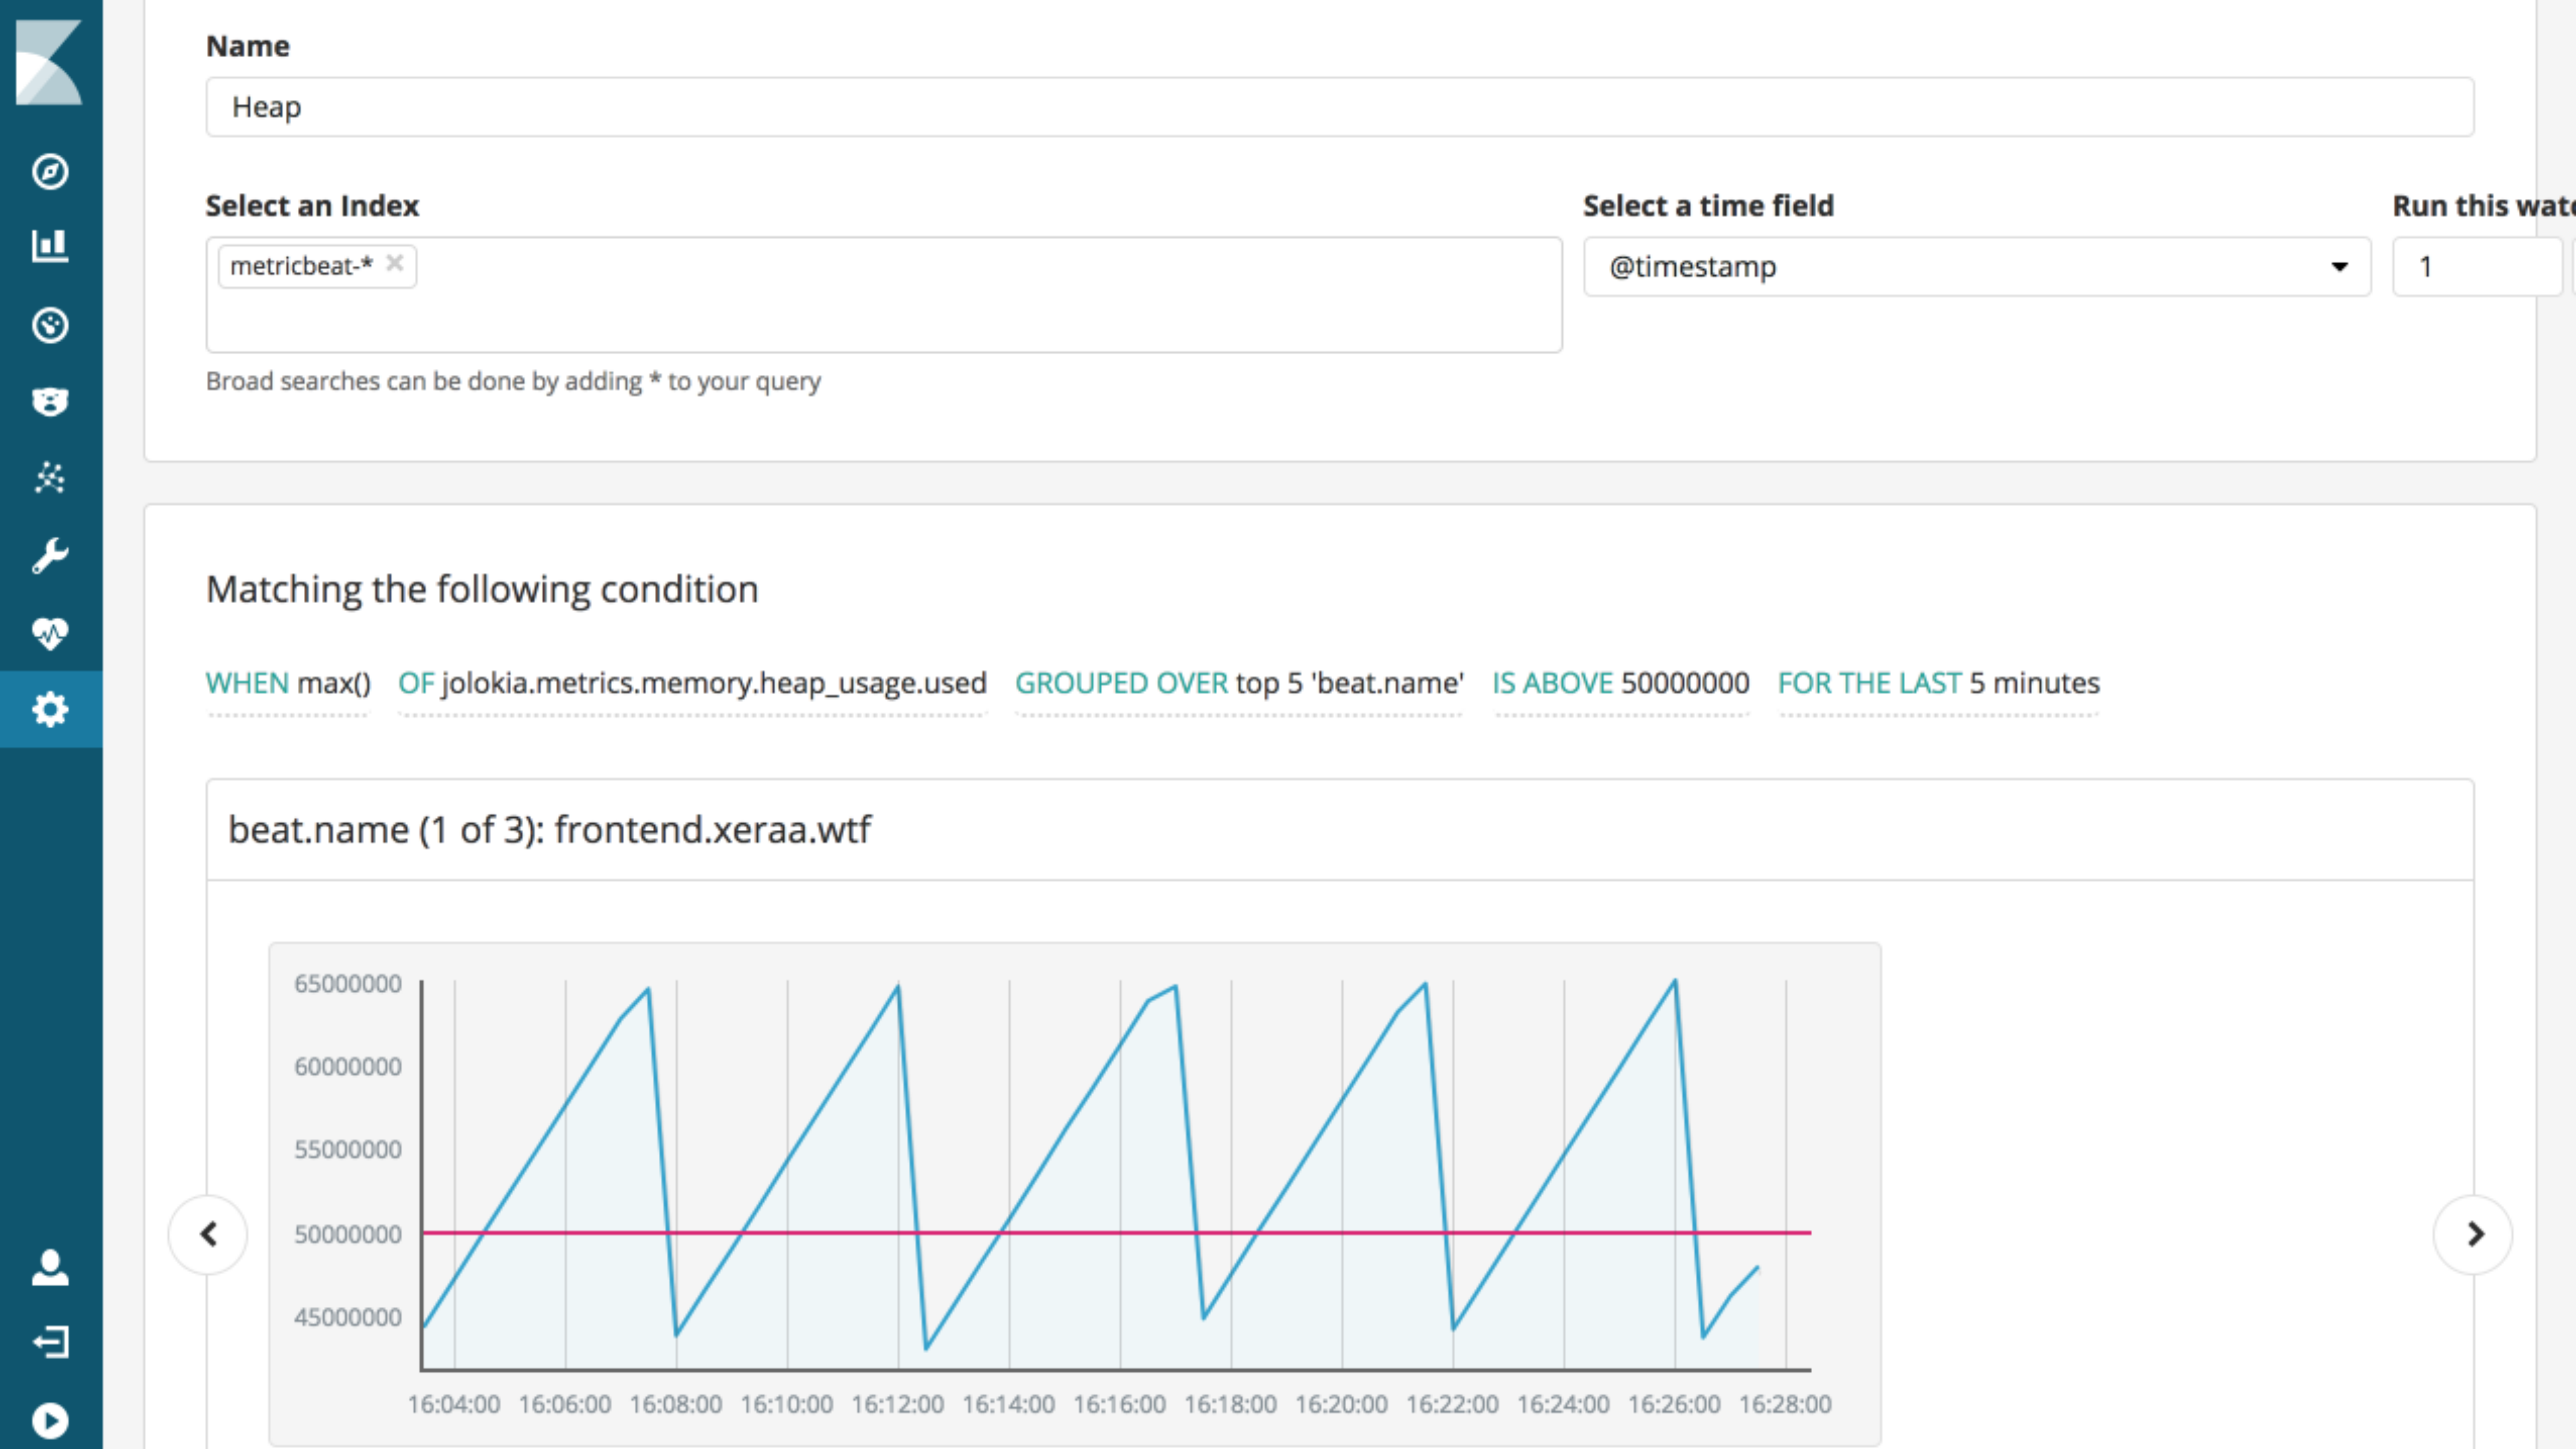

Name

Heap

Select an Index

metricbeat-* ✕

Broad searches can be done by adding * to your query

Select a time field

@timestamp

Run this wat

1

Matching the following condition

WHEN max() OF jolokia.metrics.memory.heap_usage.used GROUPED OVER top 5 'beat.name' IS ABOVE 50000000 FOR THE LAST 5 minutes

beat.name (1 of 3): frontend.xeraa.wtf

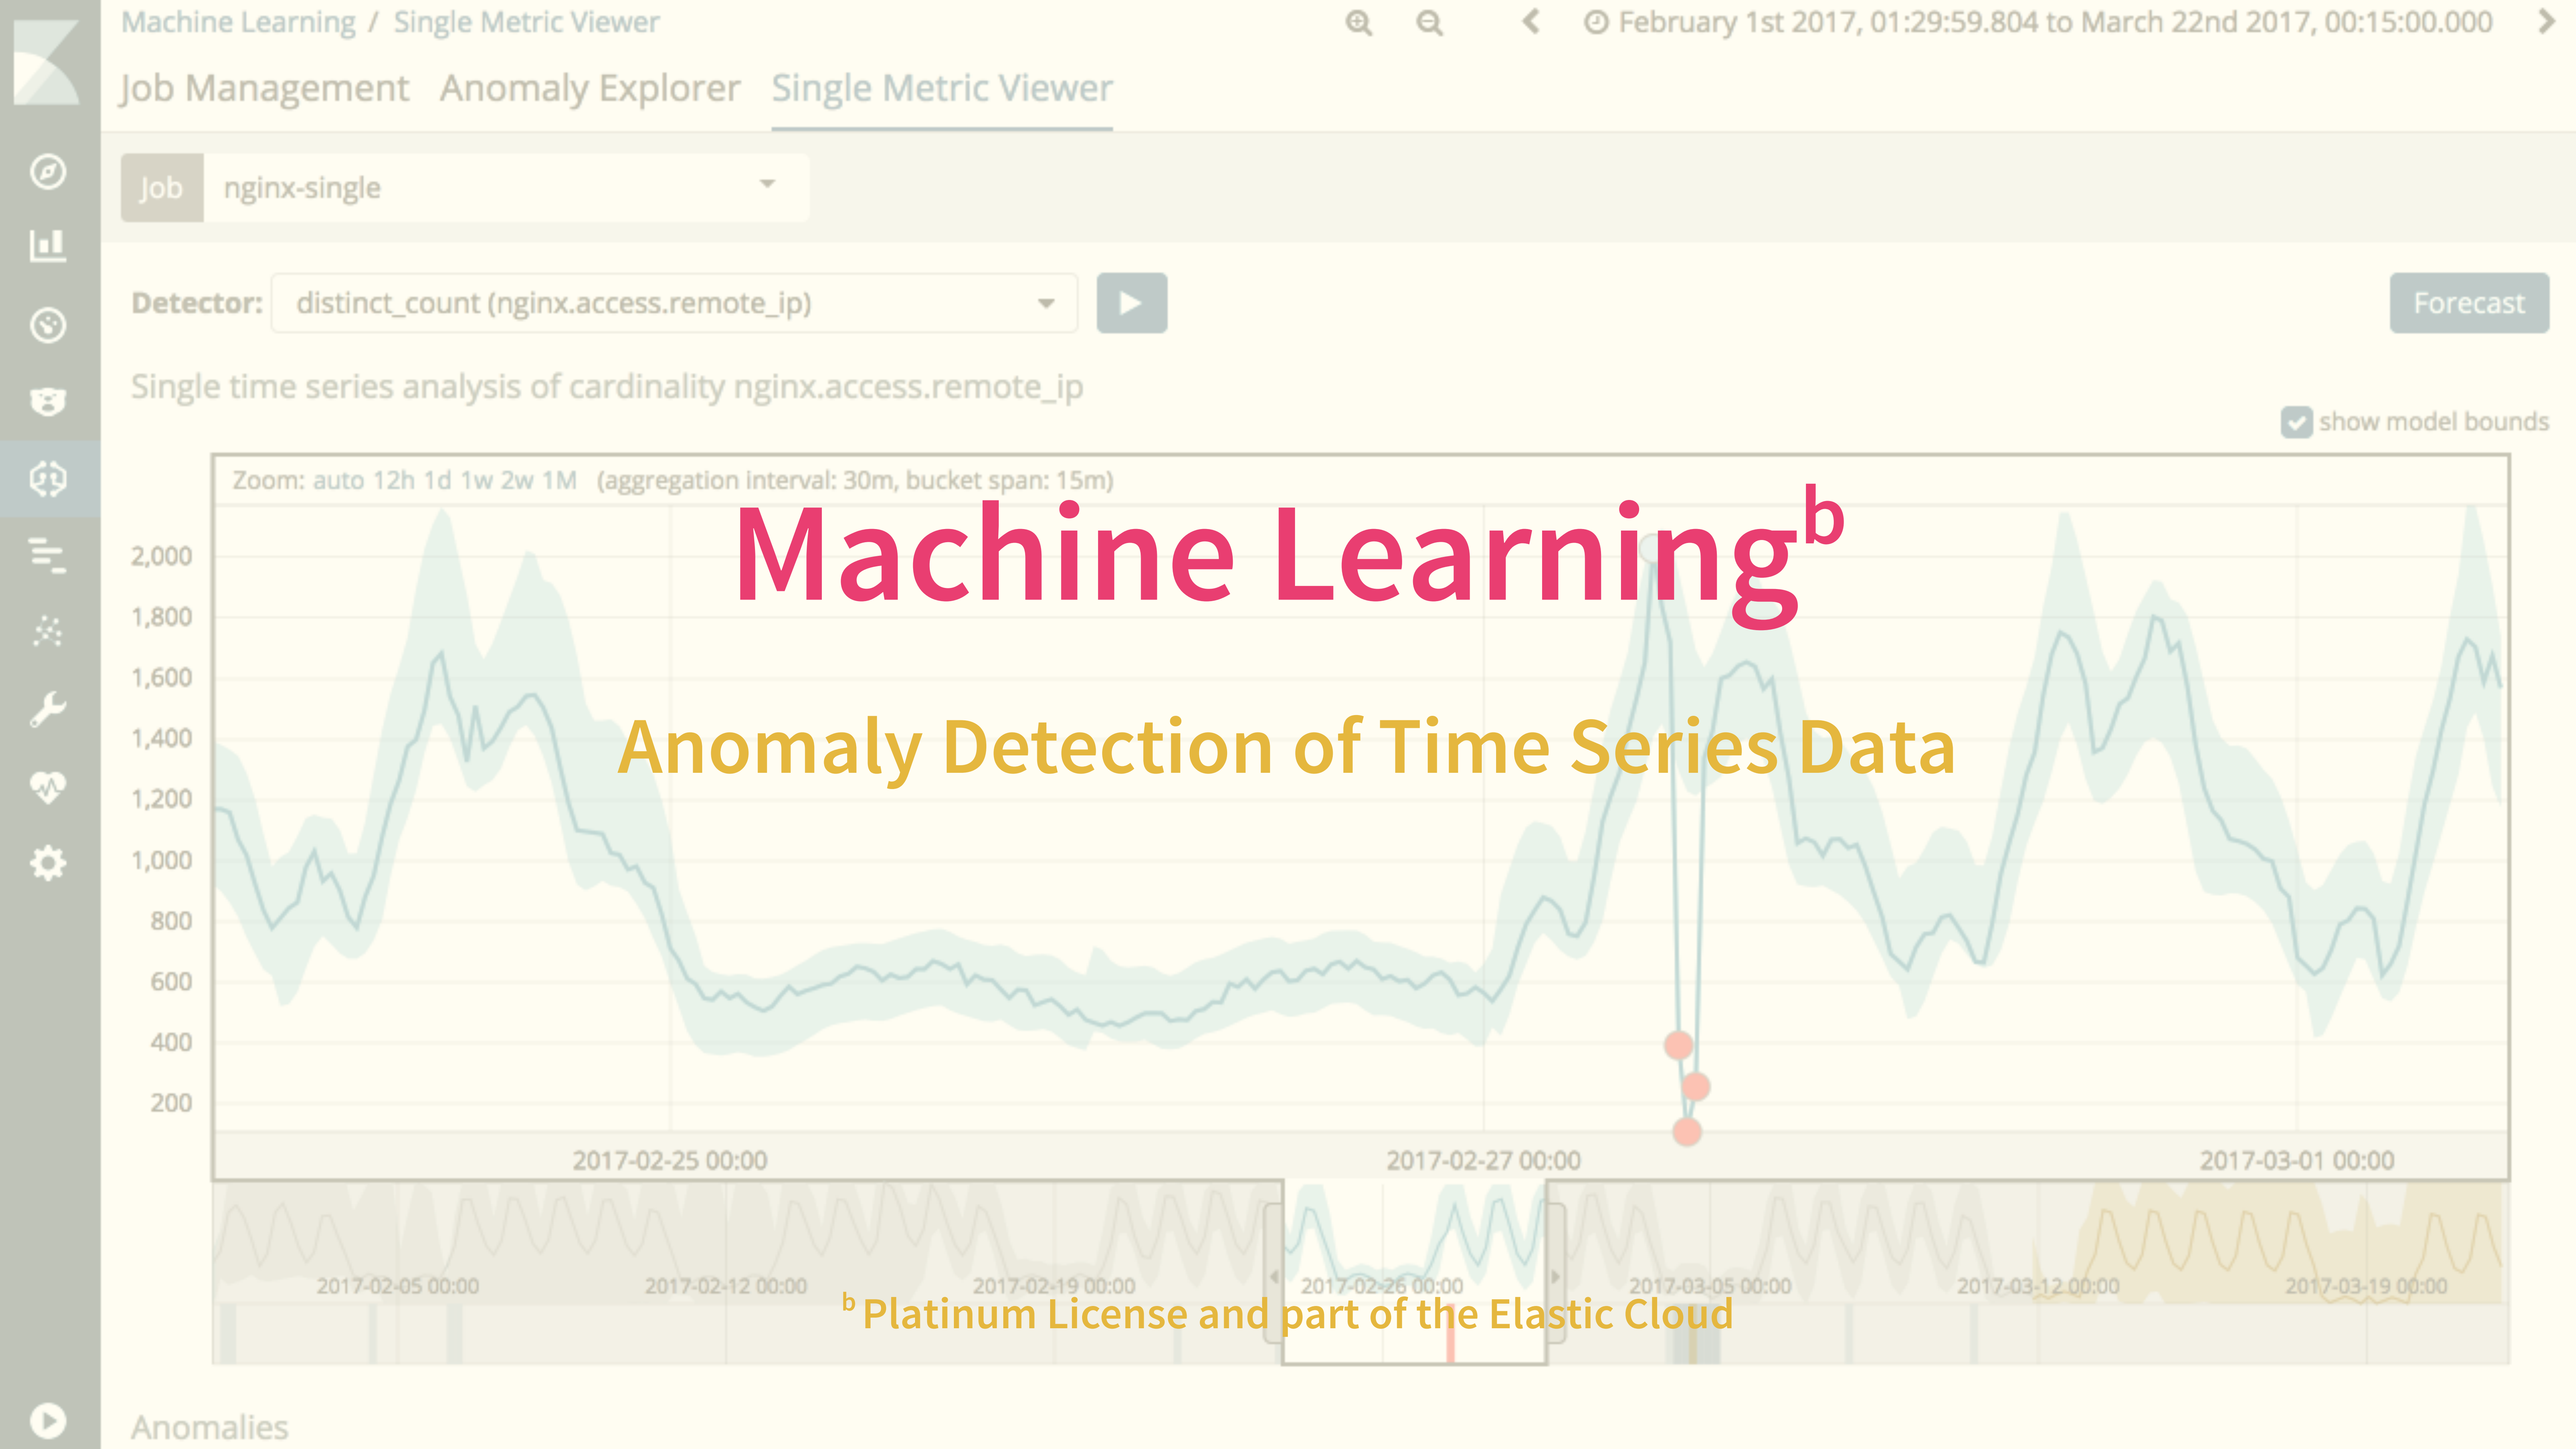

February 1st 2017, 01:29:59.804 to March 22nd 2017, 00:15:00.000

Job: nginx-single

Detector: distinct_count (nginx.access.remote_ip)

Single time series analysis of cardinality nginx.access.remote_ip

☑ show model bounds

(aggregation interval: 30m, bucket span: 15m)

Machine Learning[b]

Anomaly Detection of Time Series Data

[b] Platinum License and part of the Elastic Cloud

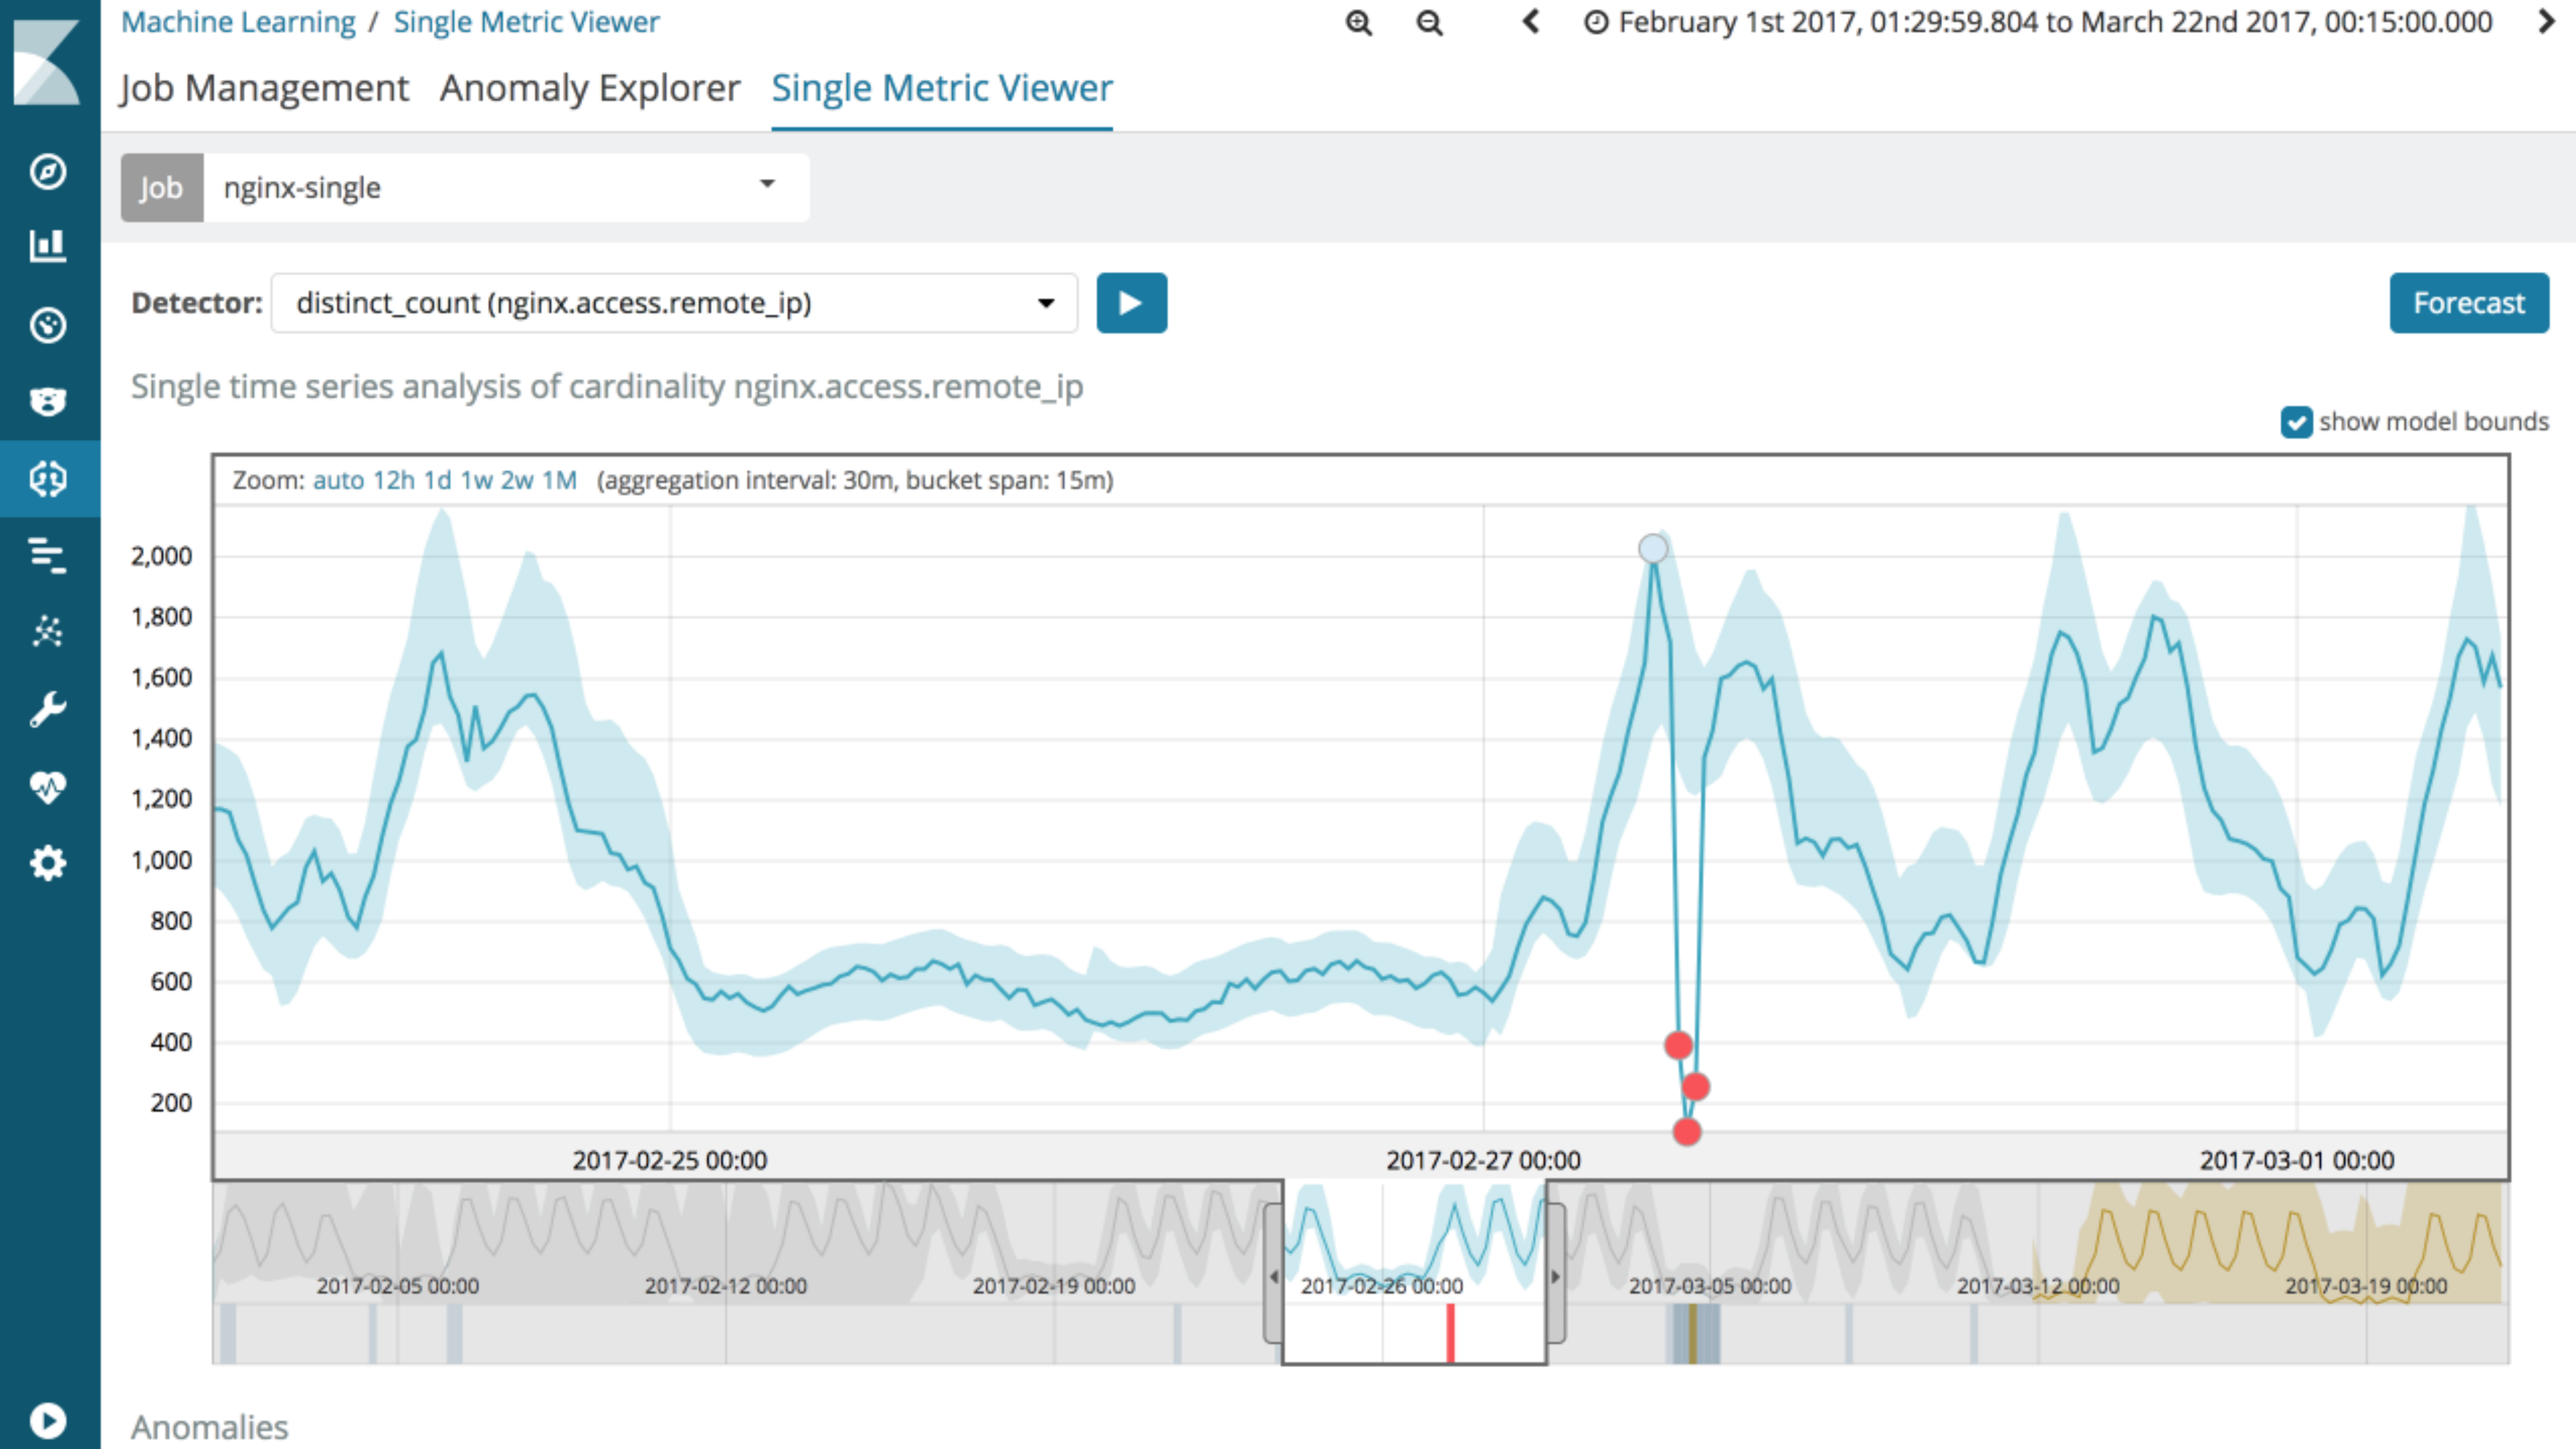

February 1st 2017, 01:29:59.804 to March 22nd 2017, 00:15:00.000

Job nginx-single

Detector: distinct_count (nginx.access.remote_ip)

Single time series analysis of cardinality nginx.access.remote_ip

Zoom: auto 12h 1d 1w 2w 1M (aggregation interval: 30m, bucket span: 15m)

Security[c]

[c] Gold / Platinum License and part of the Elastic Cloud

elastic

Q&A + YOUR APPS

Conclusion

System metrics & network

Filebeat modules & Auditbeat

Application logs

Uptime

Application metrics

Request tracing

Code

https://github.com/xeraa/microservice-monitoring

THANK YOU

Philipp Krenn @xeraa

PS: Sticker