Monitor Your Applications

with Logs, Metrics, Pings, and Traces

Philipp Krenn @xeraa

elastic

Honest Status Page
@honest_update

Follow

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

RETWEETS	LIKES
2,882	2,256

4:10 PM - 7 Oct 2015

18 2.9K 2.3K


THIS IS FINE.


TH-
THIS IS NOT FINE!!
OH MY GOD EVERY THINGS ON FIRE
WHAT THE HELL IS MY PROBLEM


ALL THE THINGS!

How?

 vs 

elastic

elastic

Developer 🥑

Disclaimer

I build highly monitored Hello World apps

Disclaimer

This is not a training

https://training.elastic.co

elasticsearch.

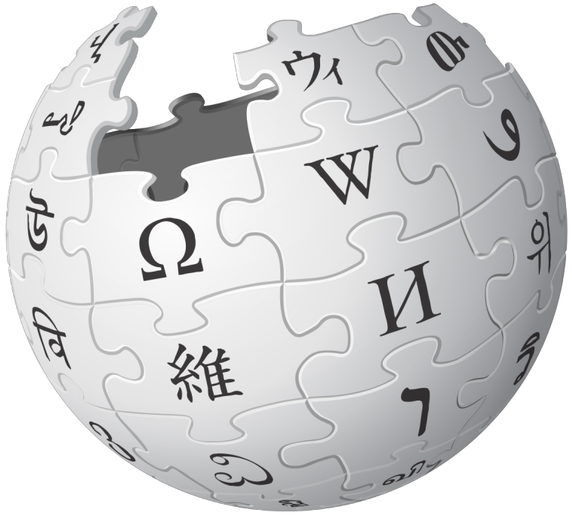

logstash


lyft


slack


fitbit

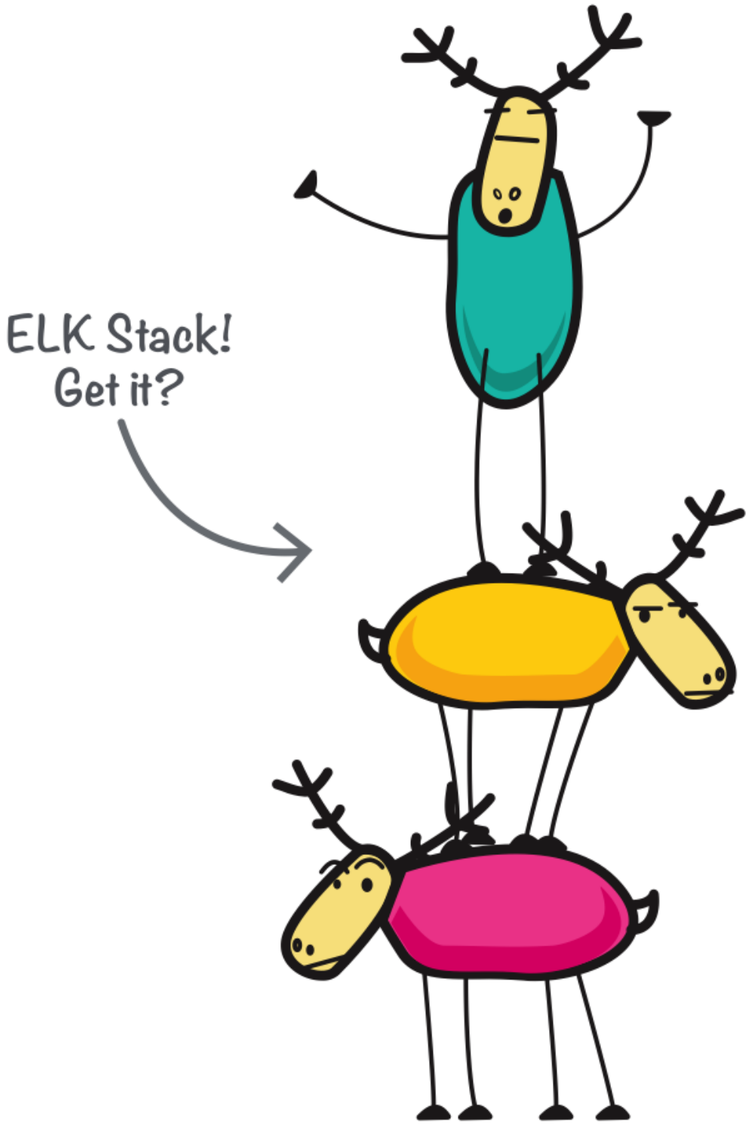

E Elasticsearch

L Logstash

K Kibana

Apparently, I'm an ELKB personality.

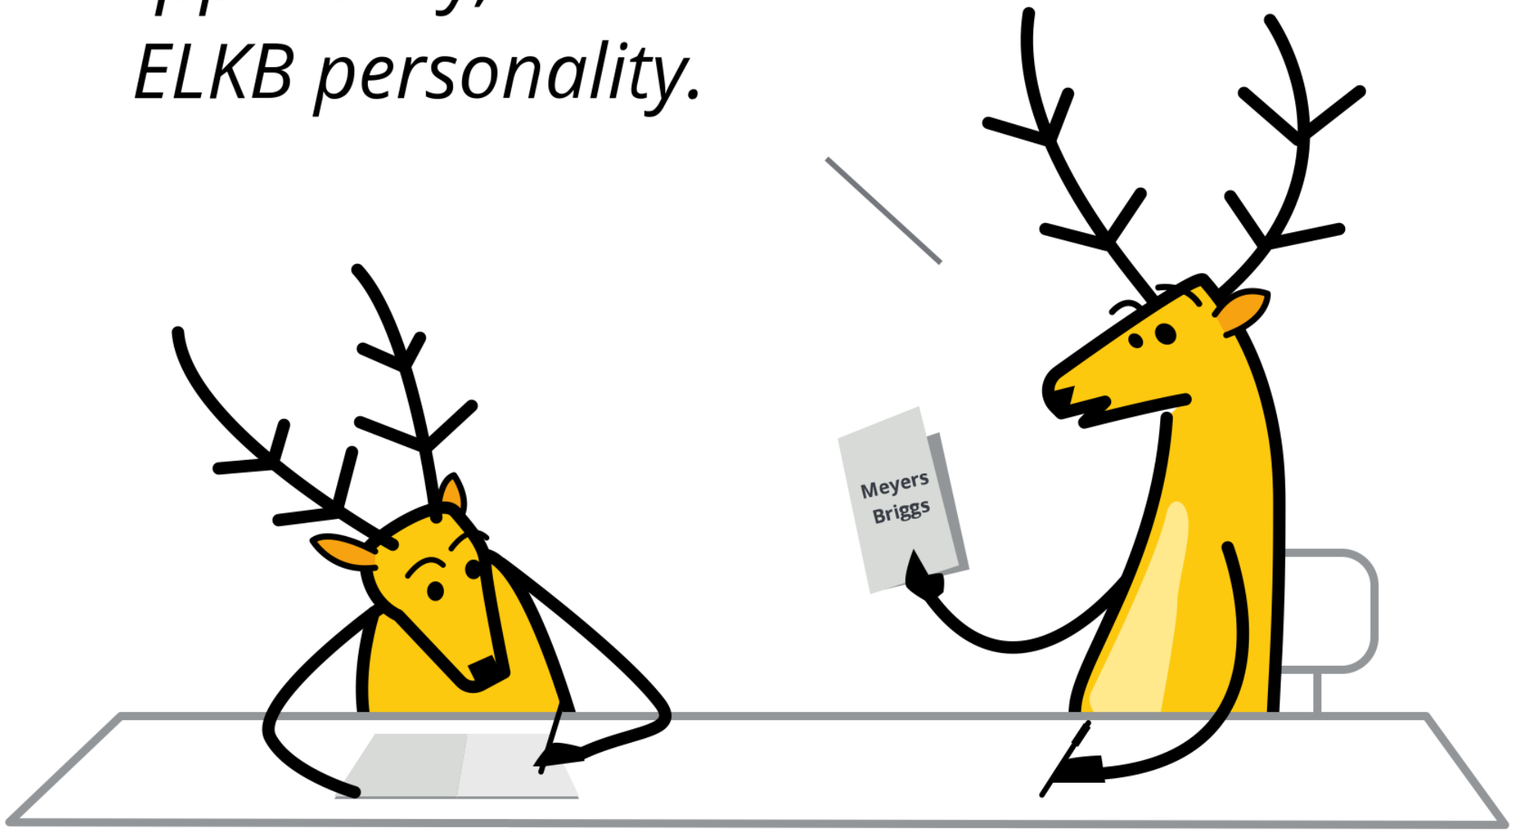

elastic

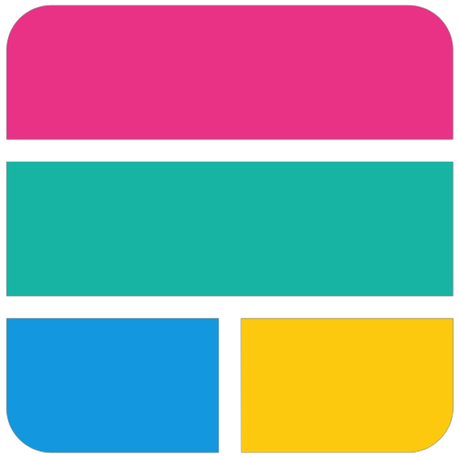

elastic stack


BLIZZARD
ENTERTAINMENT


CAR
2GO


ebay

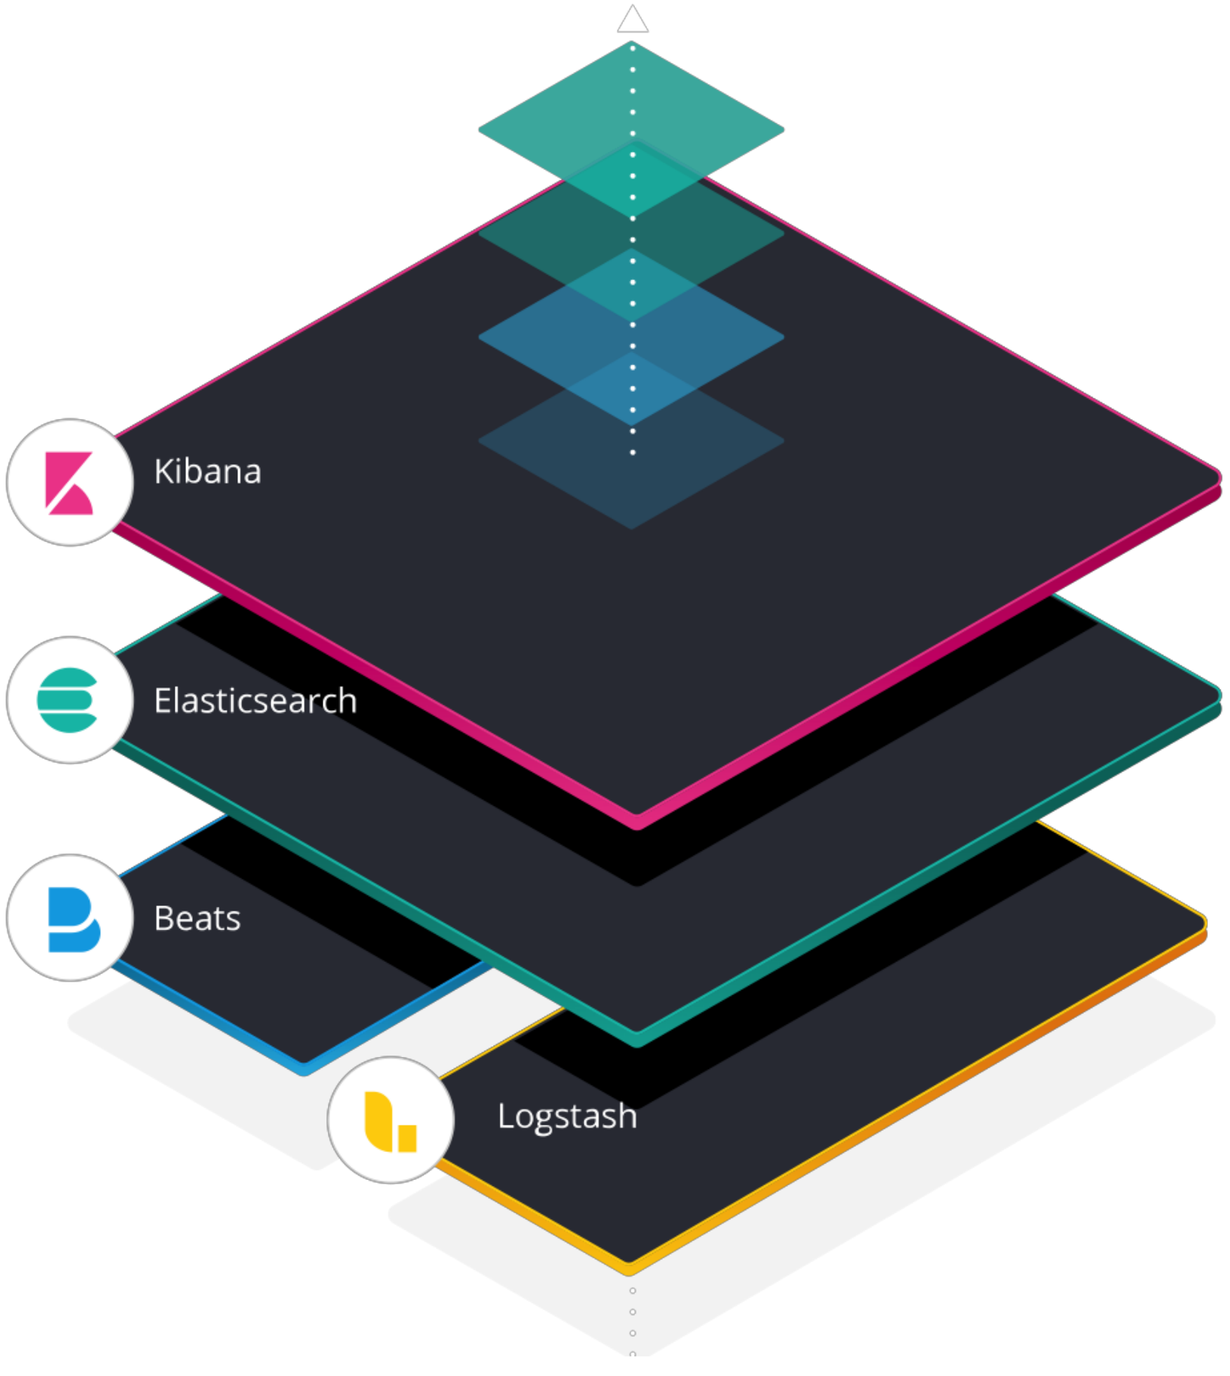
Kibana
Elasticsearch
Beats
Logstash

Licensing

Open Source Apache-2.0

Basic free

Commercial 

elastic

Agenda

9:00	Intro + Basics
10:00	Monitor Java
13:30	Some Security
14:30	Monitor PHP
16:00	Q&A + Your Apps

elastic

Intro + Basics

Intro

Experience Goals

elastic

Elasticsearch

Search Engine

REST

Horizontal Scalability

elastic

Cluster, Node, Index, Shard

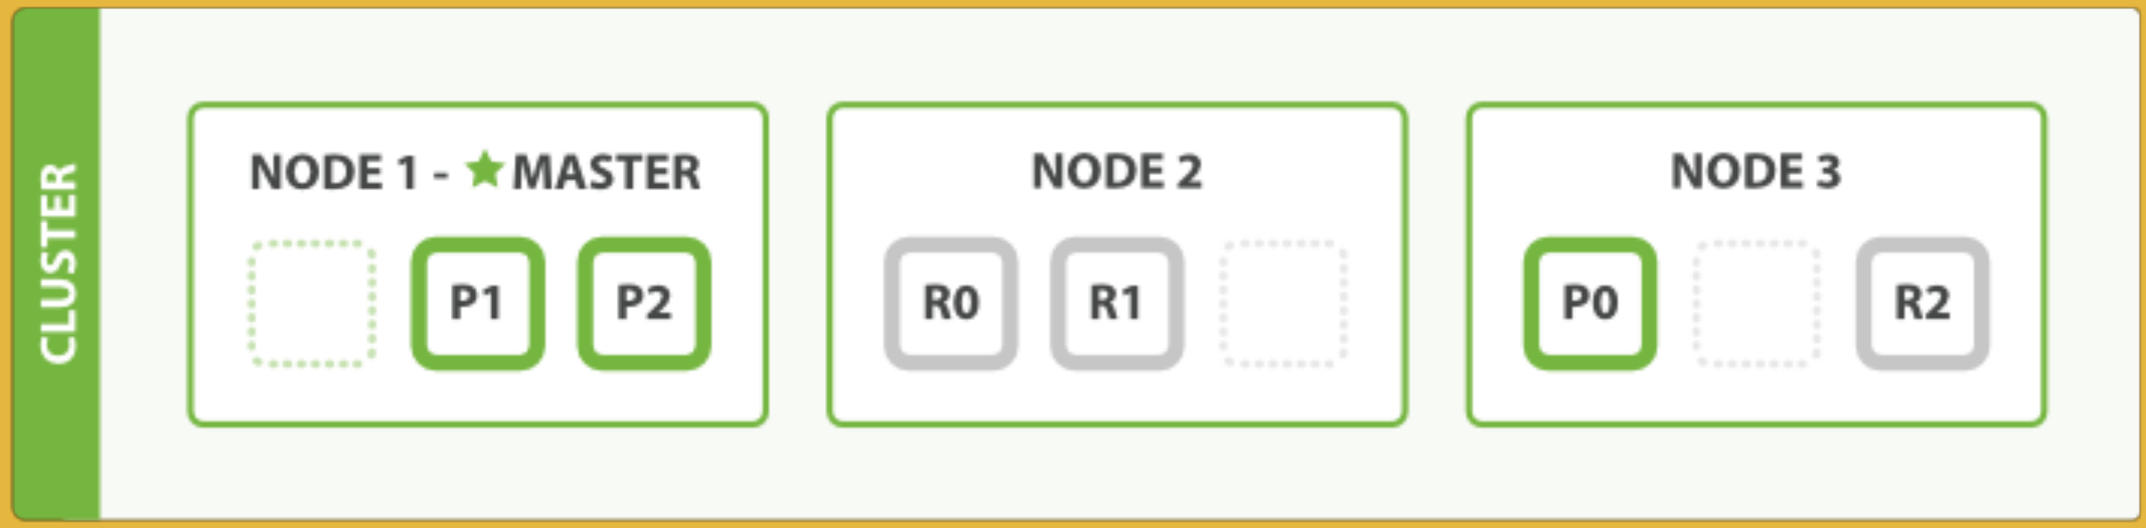

Write

Coordinating Node, ID, Hash, Primary, Replica(s)

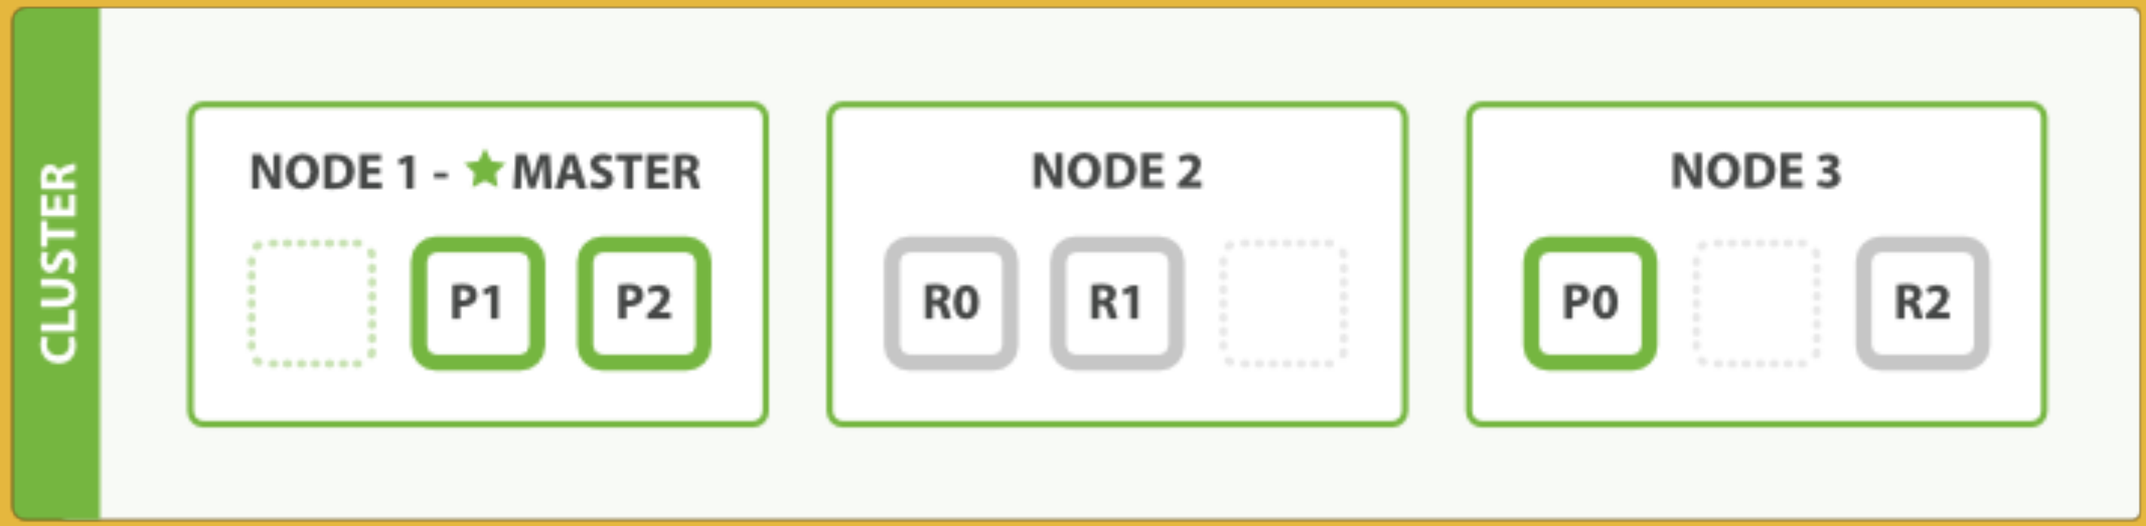

GET

ID, Coordinating Node, Hash, Shard

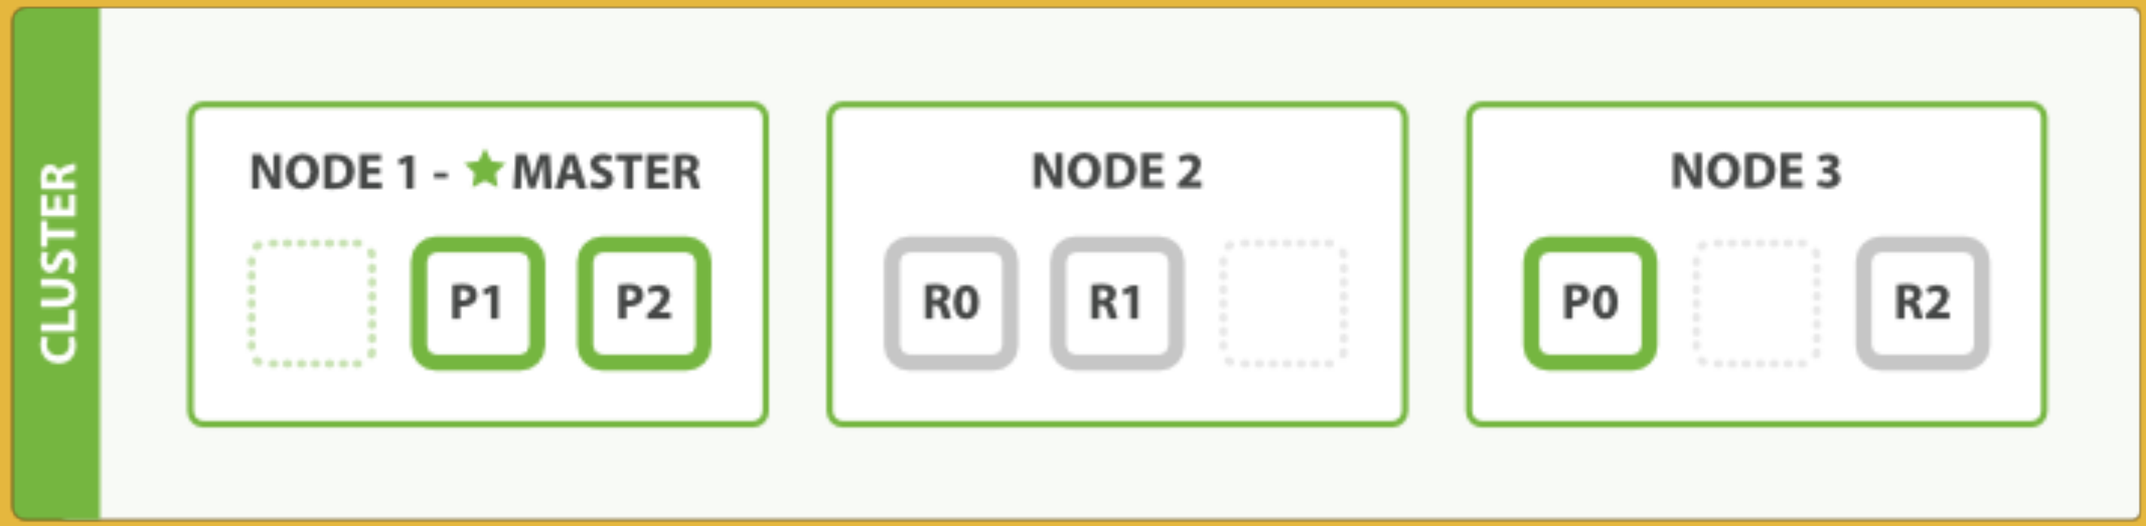

Search

Coordinating Node, Query then Fetch

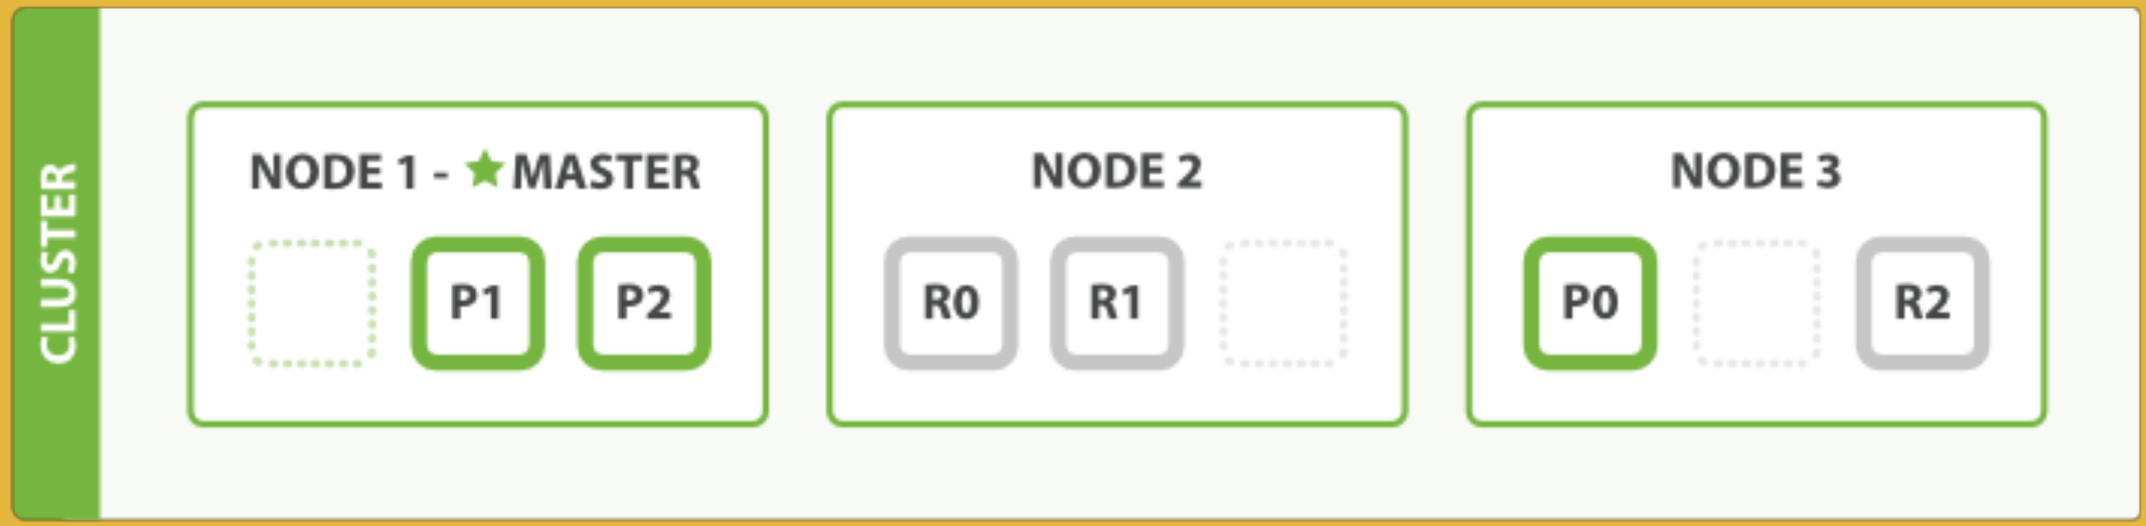

Time Based Indices

```
"metricbeat-%{[beat.version]}-%{+yyyy.MM.dd}"
```

Nodes

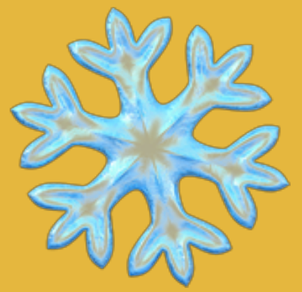

elastic

```
$ bin/elasticsearch
  -Enode.attr.rack=rack1
  -Enode.attr.size=hot
```

```
PUT /metricbeat/_settings
{
  "index.routing.allocation.include.size": "hot"
}
```

Monitor Java

Code

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

elastic

Simple

No discovery, load-balancing,...

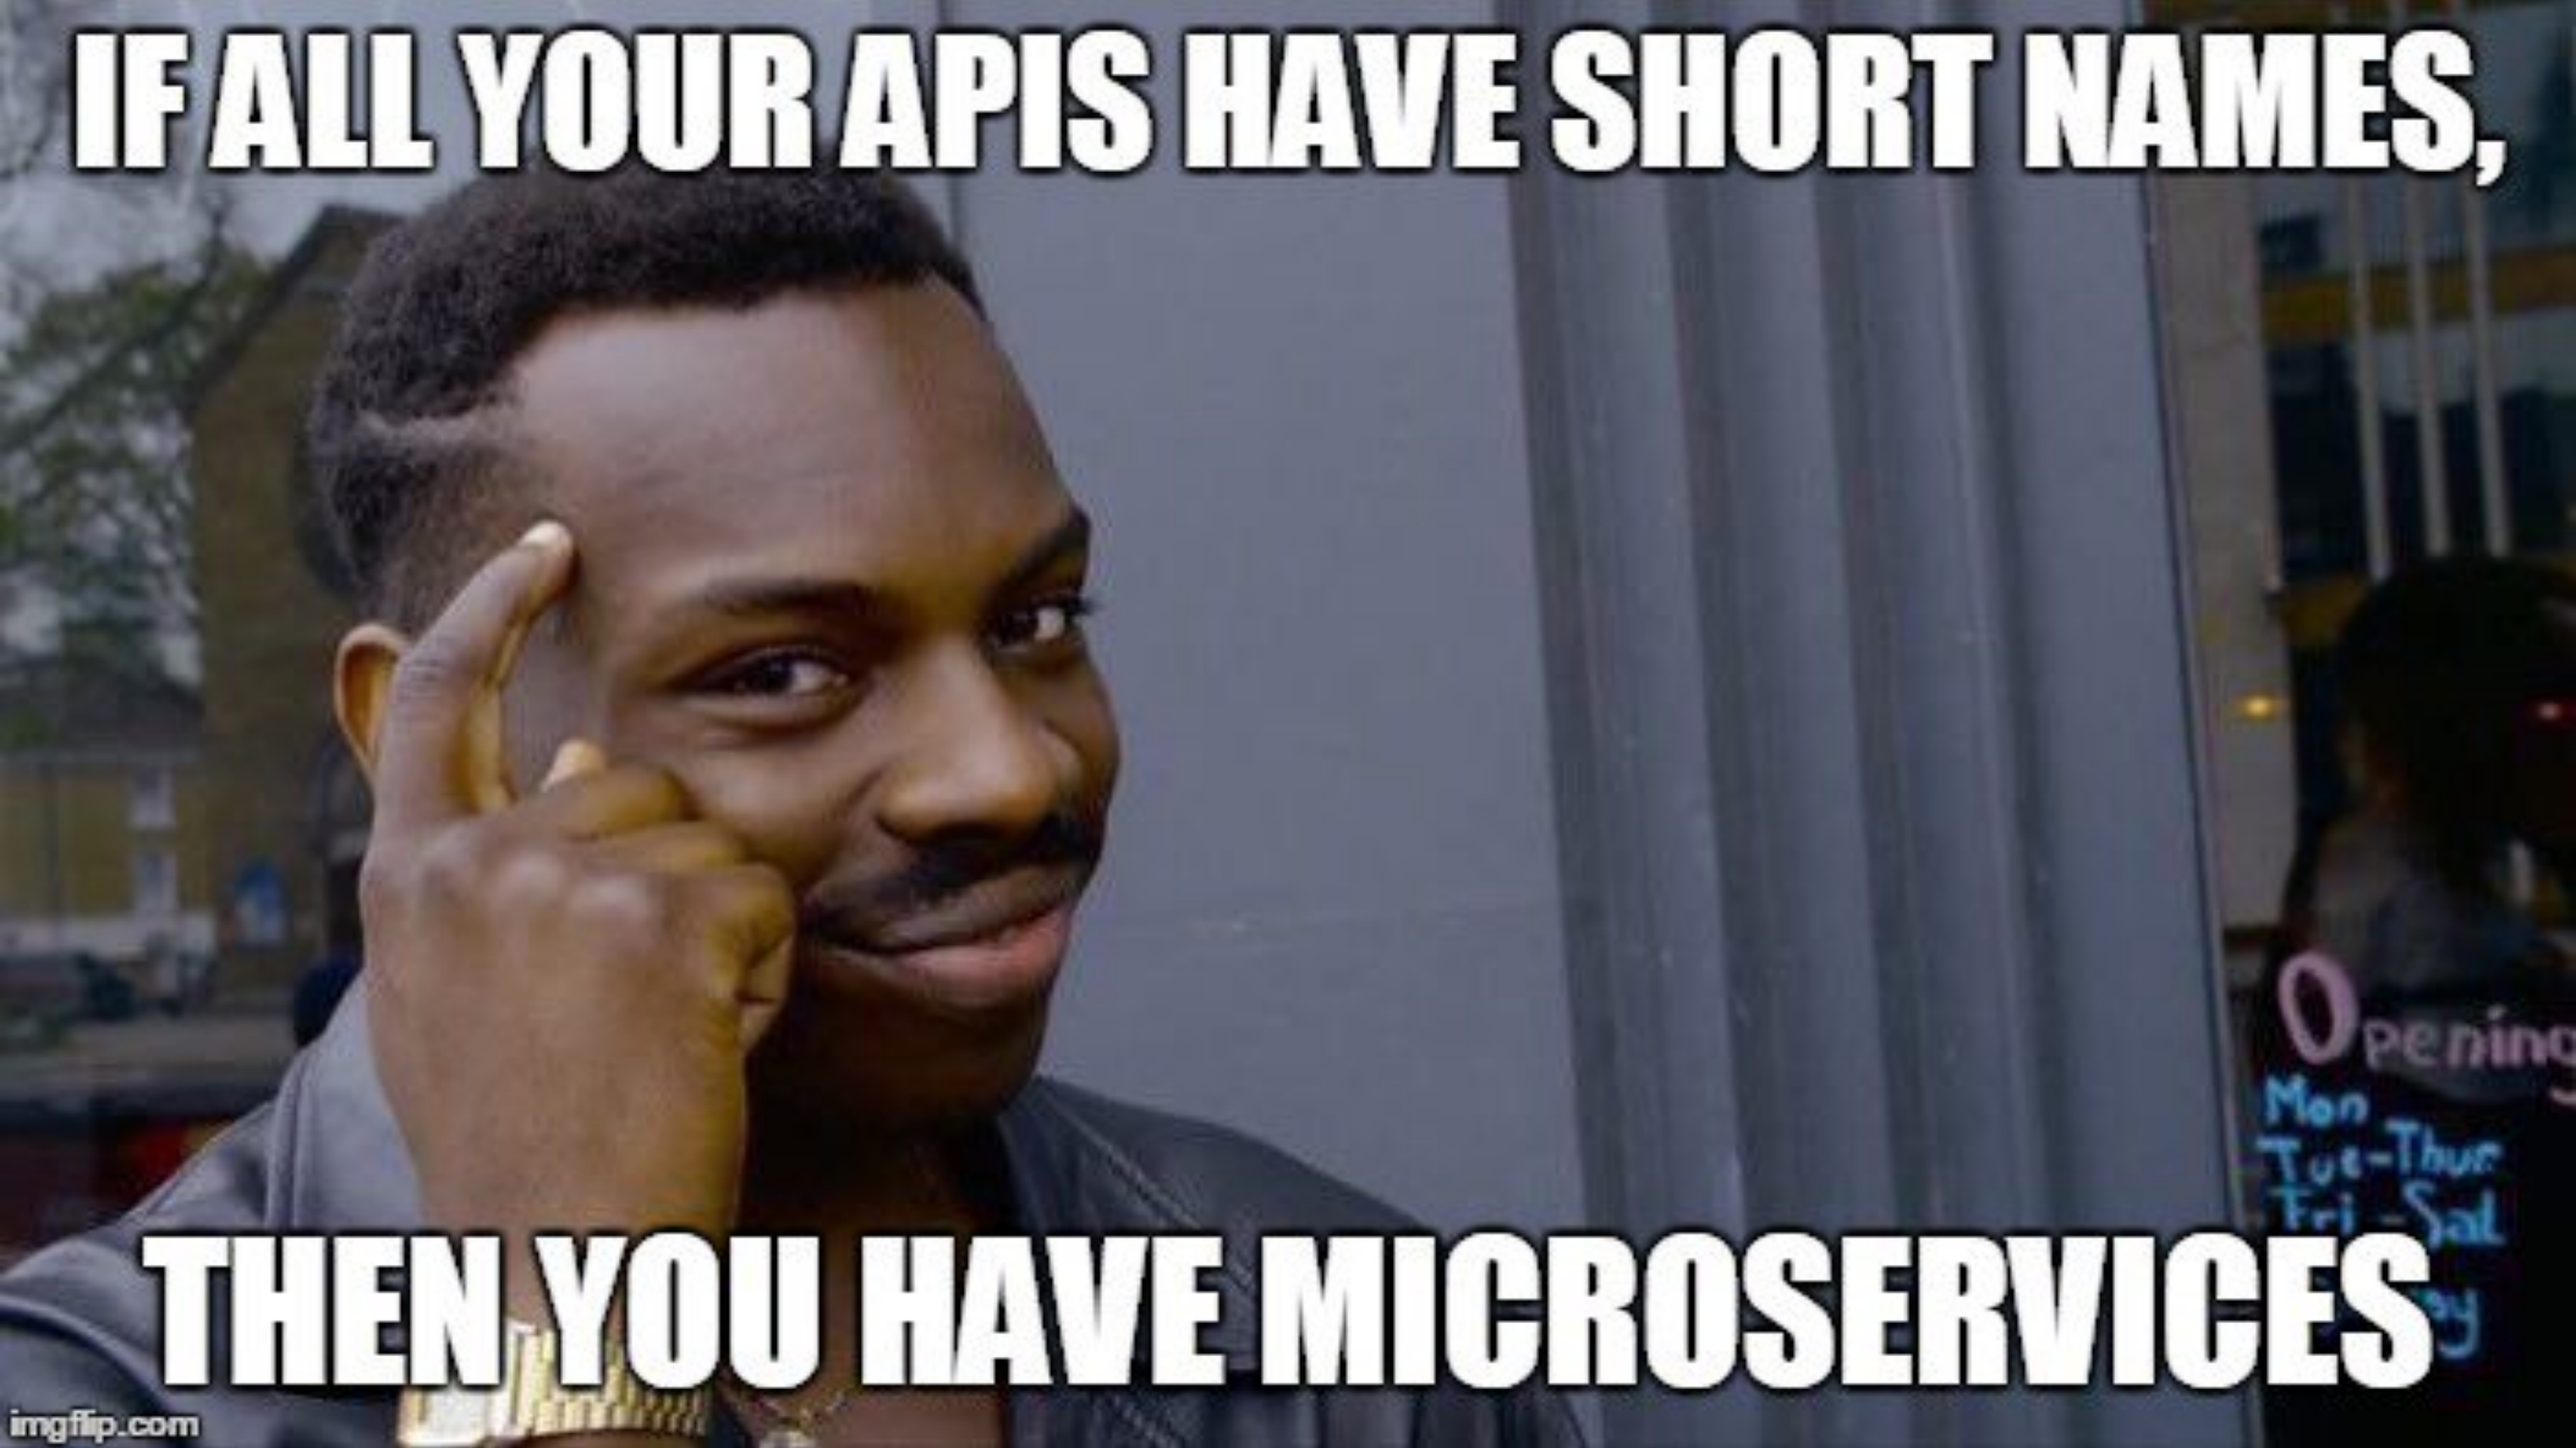
IF ALL YOUR APIS HAVE SHORT NAMES,
THEN YOU HAVE MICROSERVICES
imgflip.com


Auditbeat
Filebeat
Heartbeat
Metricbeat
Packetbeat

elastic

Workshop

SSH: `ssh elastic-admin@workshop-<#>.xeraa.wtf`
elastic-admin / secret

Elasticsearch: `http://localhost:9200`
admin / secret

Kibana: `http://workshop-<#>.xeraa.wtf:5601`
admin / secret

Java Application: `http://workshop-<#>.xeraa.wtf`

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

elastic

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

elastic

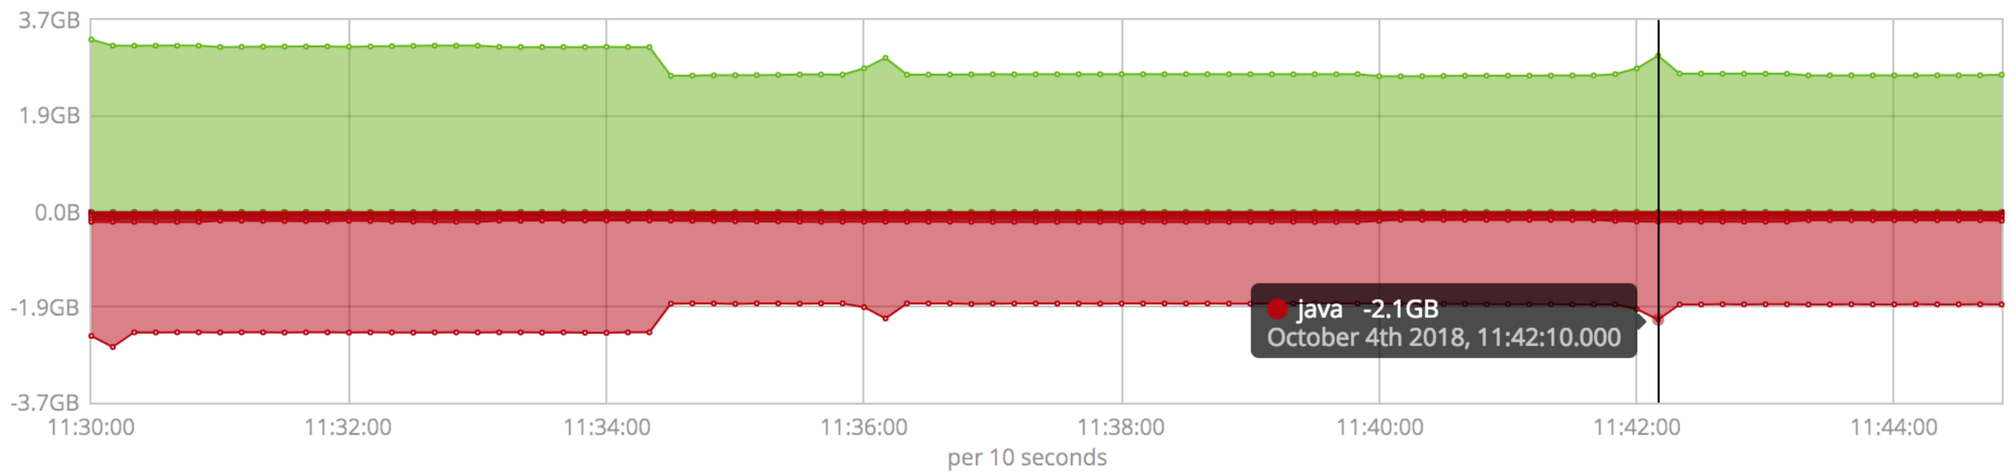

Auto Apply ▶ Apply Changes The changes will be automatically applied.

Data Panel Options Annotations

System memory

Metrics Options

Aggregation: Sum

Field: system.memory.actual.used.bytes

Group By: Everything

Process memory

Metrics Options

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.	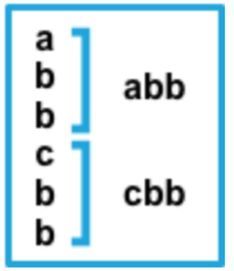
`false`	`before`	Consecutive lines that match the pattern are prepended to the next line that doesn't match.	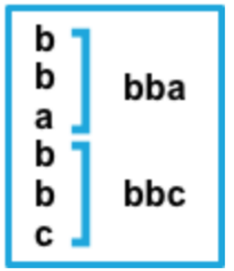
`true`	`after`	Consecutive lines that don't match the pattern are appended to the previous line that does match.	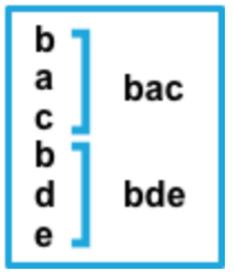
`true`	`before`	Consecutive lines that don't match the pattern are prepended to the next line that does match.	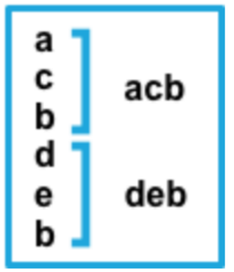

Dev Tools

Grok Debugger

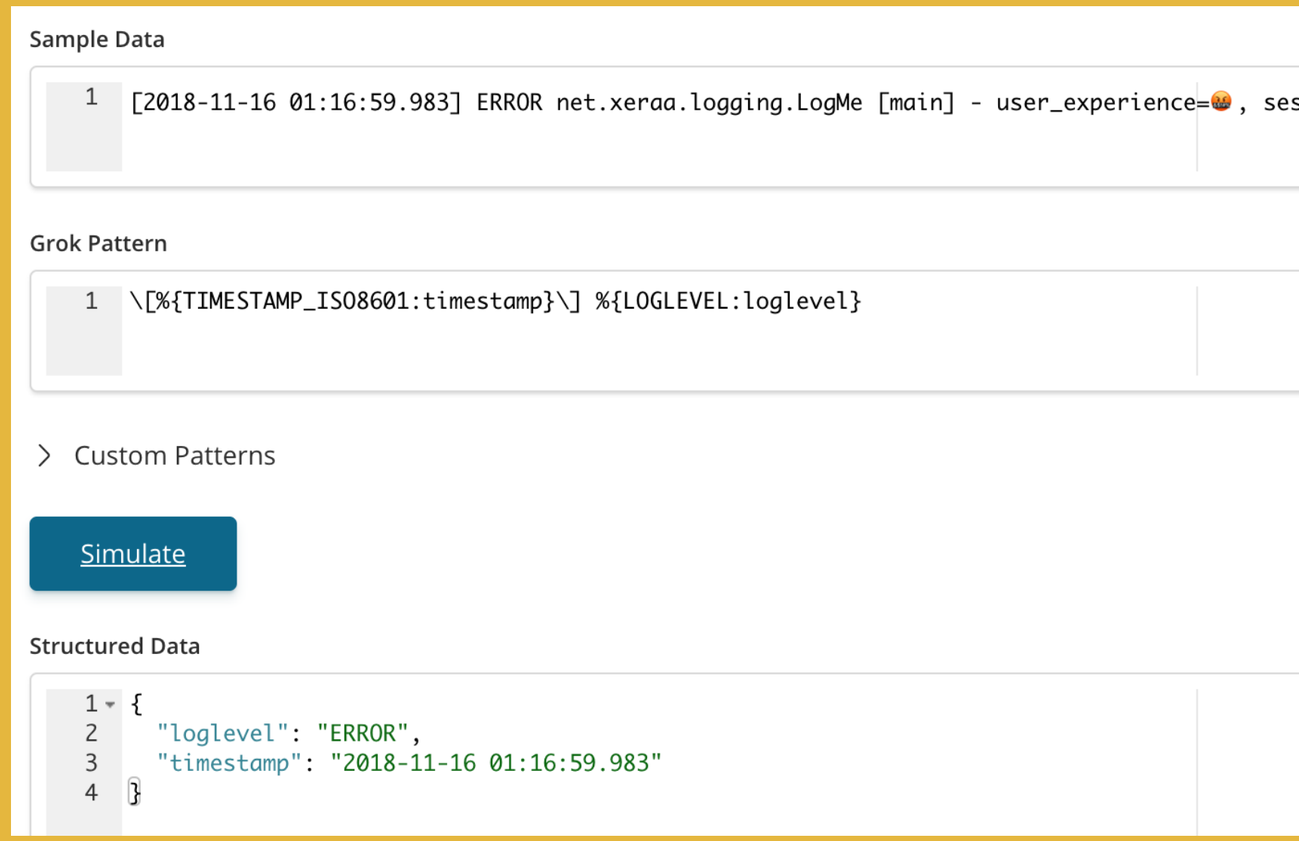

elastic

Machine Learning

Data Visualizer

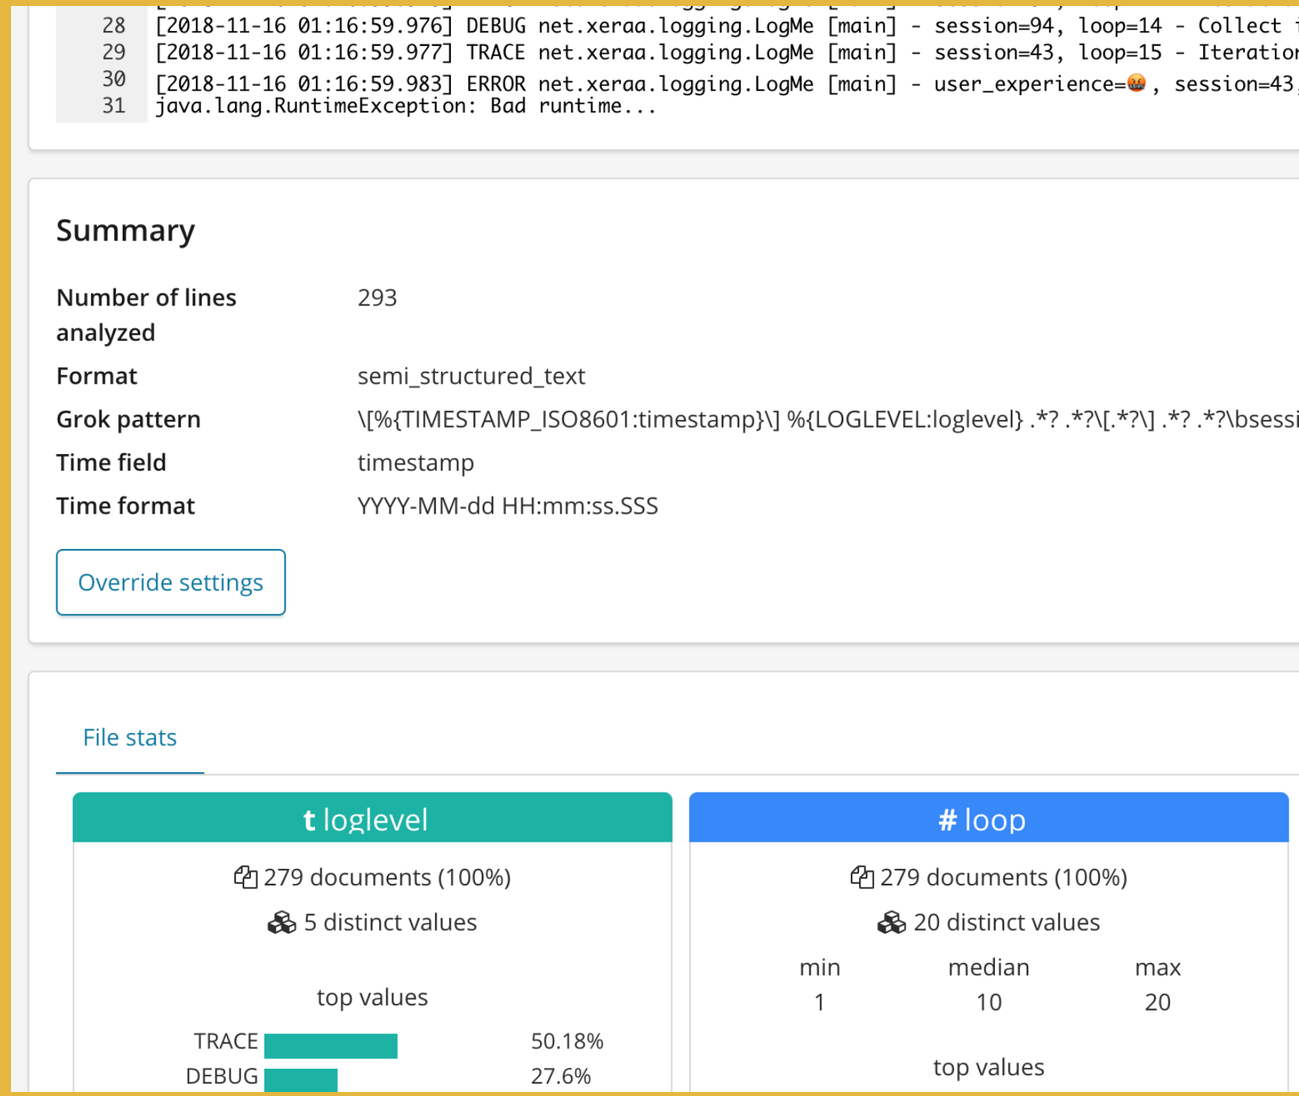

elastic

LOG UI

elastic

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

Search... (e.g. status:200 AND extension:PHP) Uses lucene query syntax

Add a filter +

filebeat-*

Data | Metrics & Axes | Panel Settings

Metrics

Y-Axis Count

Add metrics

Buckets

X-A... json.stack_hash: Descending

Split Se... json.name: Descending

Add sub-buckets

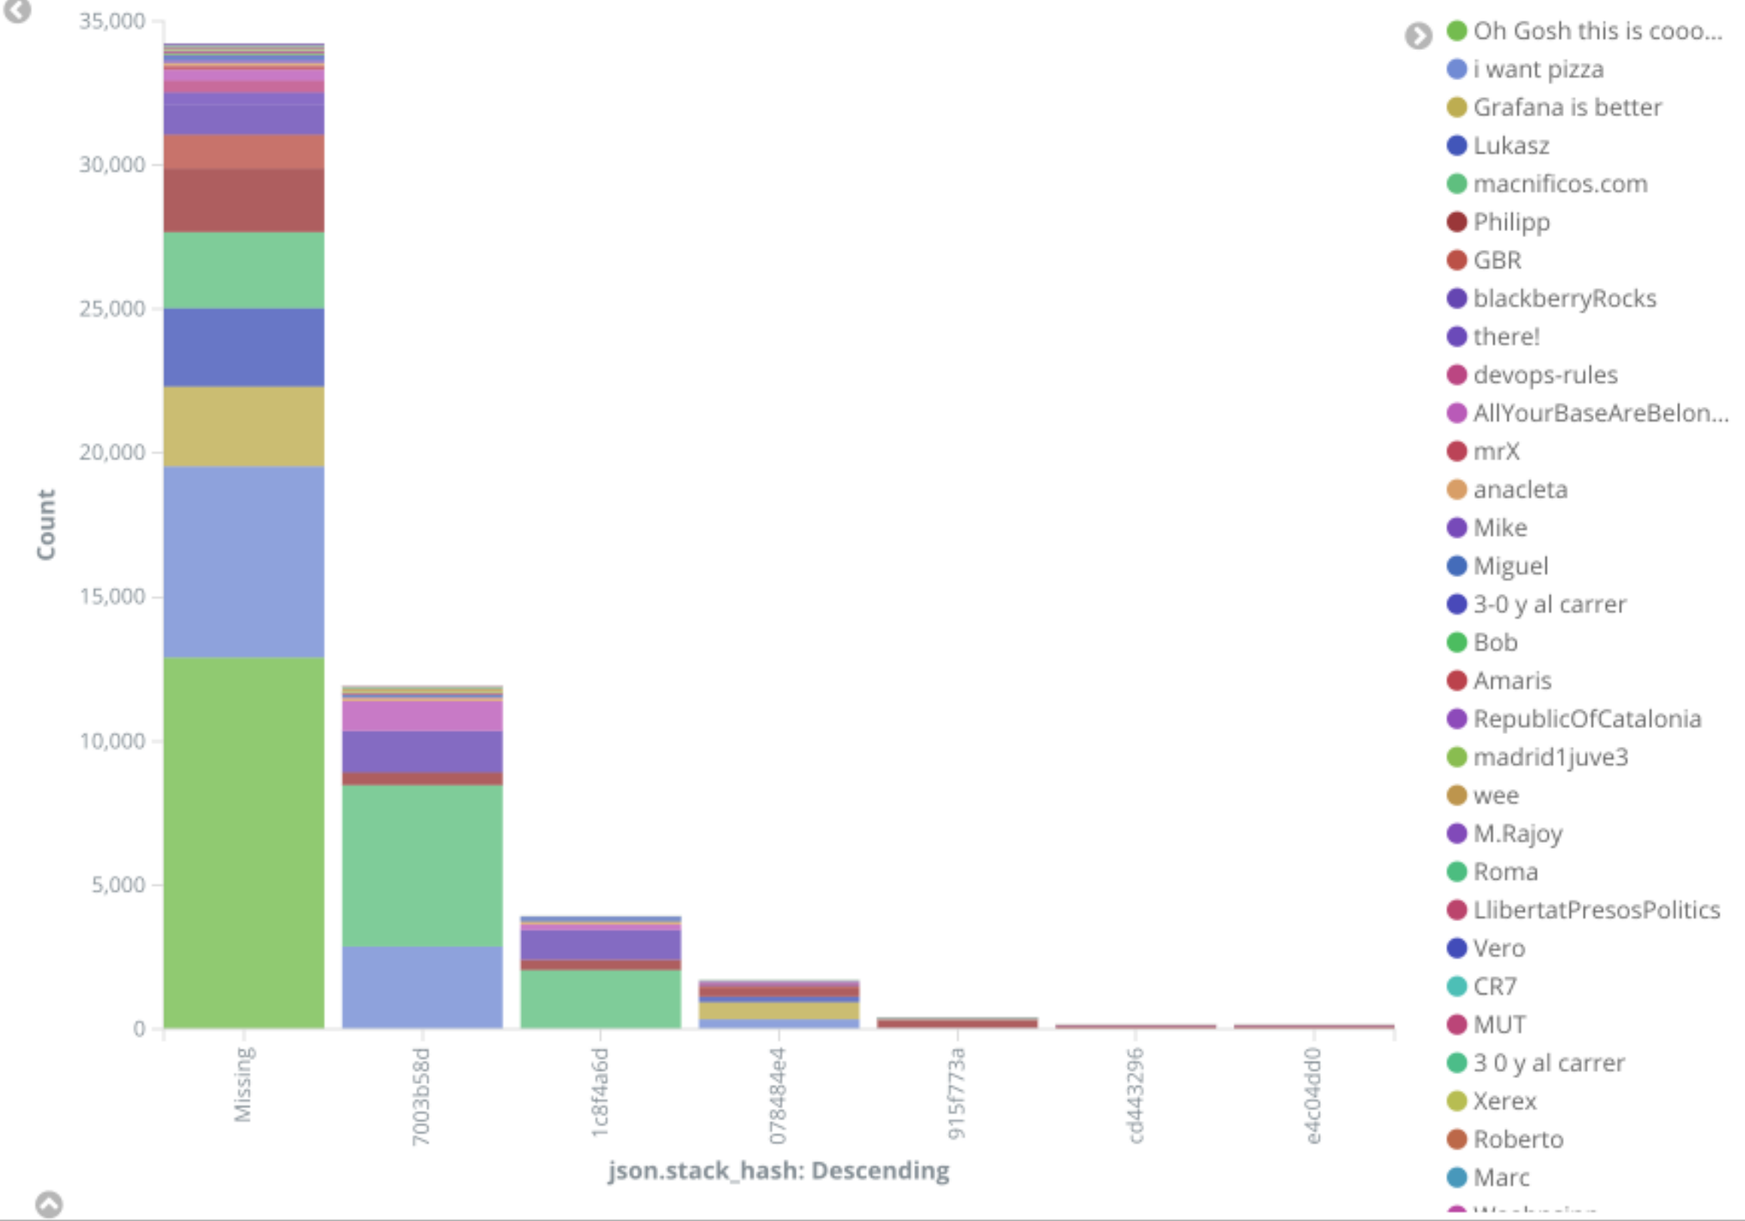

Heartbeat

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

Heap usage

Metrics | Options

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.

Time Series | Metric | Top N | Gauge | Markdown | Table

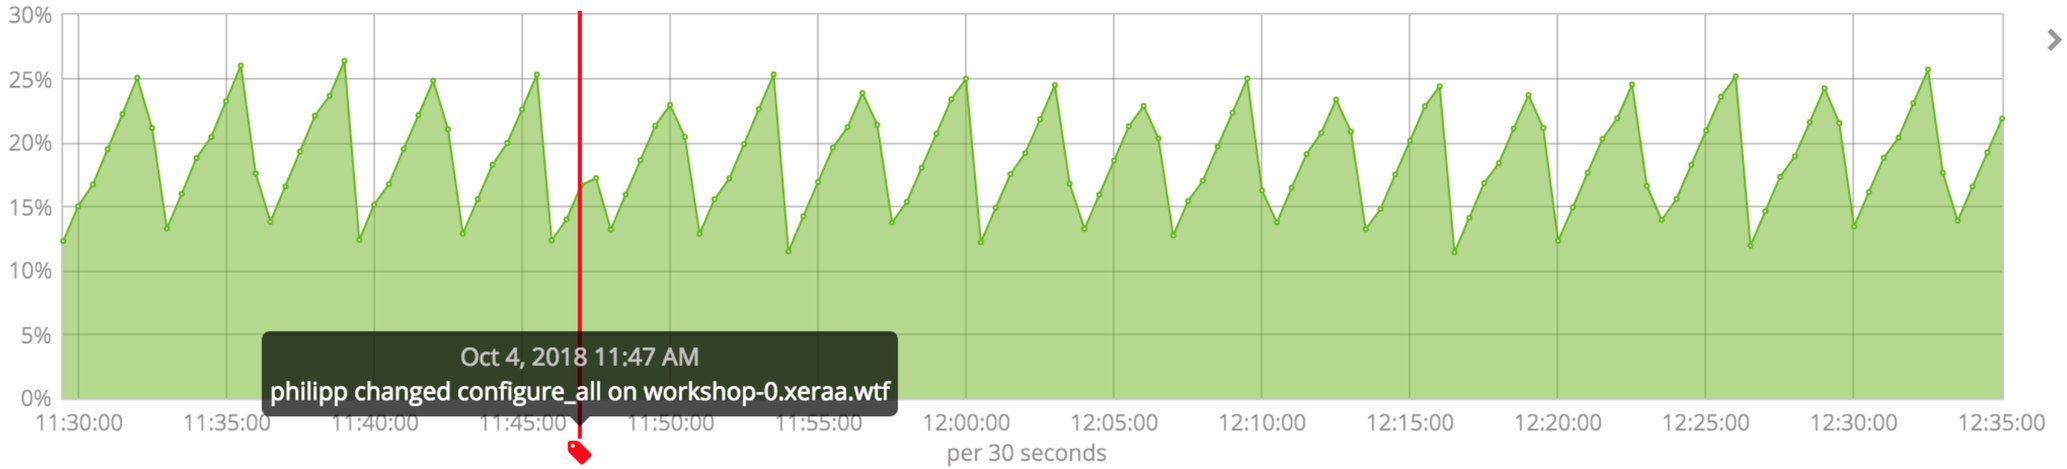

Auto Apply | ▶ Apply Changes | The changes will be automatically applied.

Data | Panel Options | Annotations

Data Sources

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

elastic

Some Security

elastic

Filebeat Modules

[Filebeat Auditd] Audit Events, [Filebeat System] New users and groups, and [Filebeat System] Sudo commands dashboards

https://github.com/linux-audit

"auditd is the userspace component to the Linux Auditing System. It's responsible for writing audit records to the disk. Viewing the logs is done with the ausearch or aureport utilities."

Auditd Monitors

File and network access

System calls

Commands run by a user

Security events

elastic

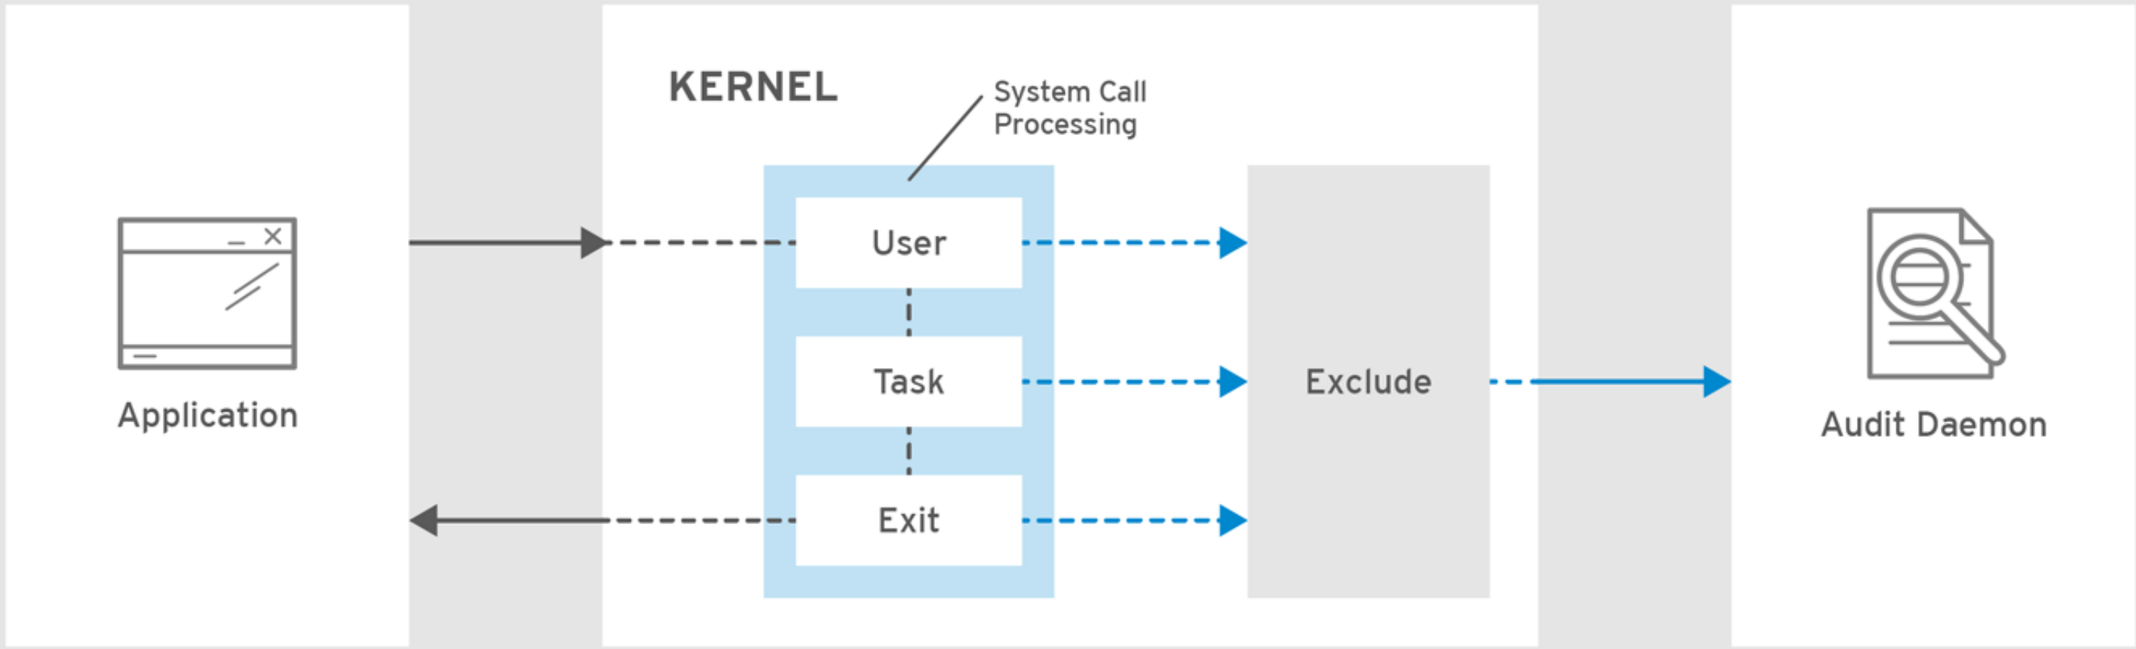

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/security_guide/chap-system_auditing

Understanding Logs

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/security_guide/sec-understanding_audit_log_files

Auditbeat

[Auditbeat Auditd] Overview dashboard

elastic

FAIL SSH

ssh elastic-user@xeraa.wtf
with a bad password

[Filebeat System] SSH login attempts
dashboard

elastic

Success

ssh elastic-user@xeraa.wtf with a good password

Run service nginx restart and pick the elastic-admin user

Audit Event

[Auditbeat Auditd] Executions dashboard filter elastic-user

Audit Event

cat /etc/passwd

Filter for tags is developers-passwd-read in Discover

Power Abuse

ssh elastic-admin@xeraa.wtf

sudo cat /home/elastic-user/secret.txt

Tag power-abuse in Discover

File Integrity

Change something in
/var/www/html/index.html

[Auditbeat File Integrity] Overview
dashboard

elastic

MONITOR PHP

elastic

Heartbeat

Add HTTP on port 88

elastic

Packetbeat

Add HTTP on port 88

elastic

Metricbeat

php-fpm

```
- module: php_fpm
  metricsets: ["pool"]
  period: 10s
  status_path: "/status"
  hosts: ["http://localhost:88"]
```

Filebeat

Collect /var/www/html/silverstripe/logs/*.json

More Features

elastic cloud

Index Lifecycle Management

Currently

https://github.com/elastic/curator

elastic

Index lifecycle management

Select a template — 2 Configure a policy — 3 Review and save

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

Create new policy

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.

Activate cold phase

Delete phase

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

Deactive cold phase

Configuration

Delete indices after

0

days

← Back

Continue →

Alerting[a]

Name

Heap

Select an Index

metricbeat-*

Select a time field

@timestamp

Run this wat

1

Broad searches can be done by adding * to your query

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

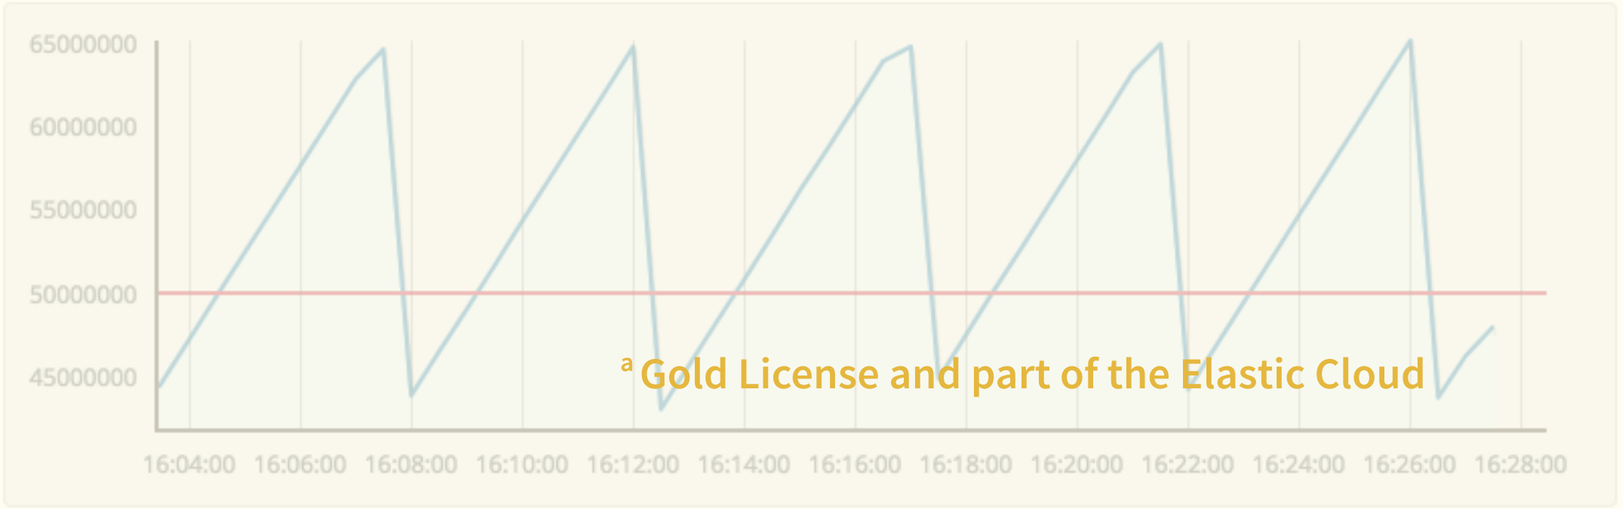

[a] Gold License and part of the Elastic Cloud

Name

Heap

Select an Index

metricbeat-* ×

Select a time field

@timestamp

Run this wat

1

Broad searches can be done by adding * to your query

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

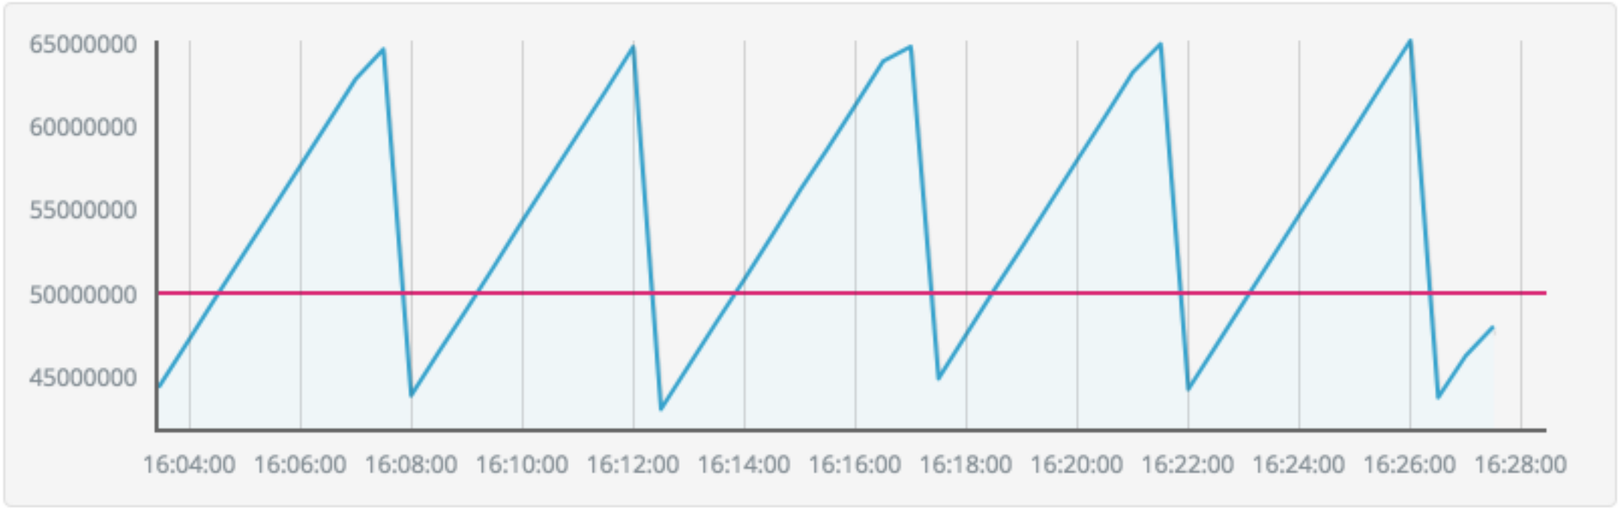

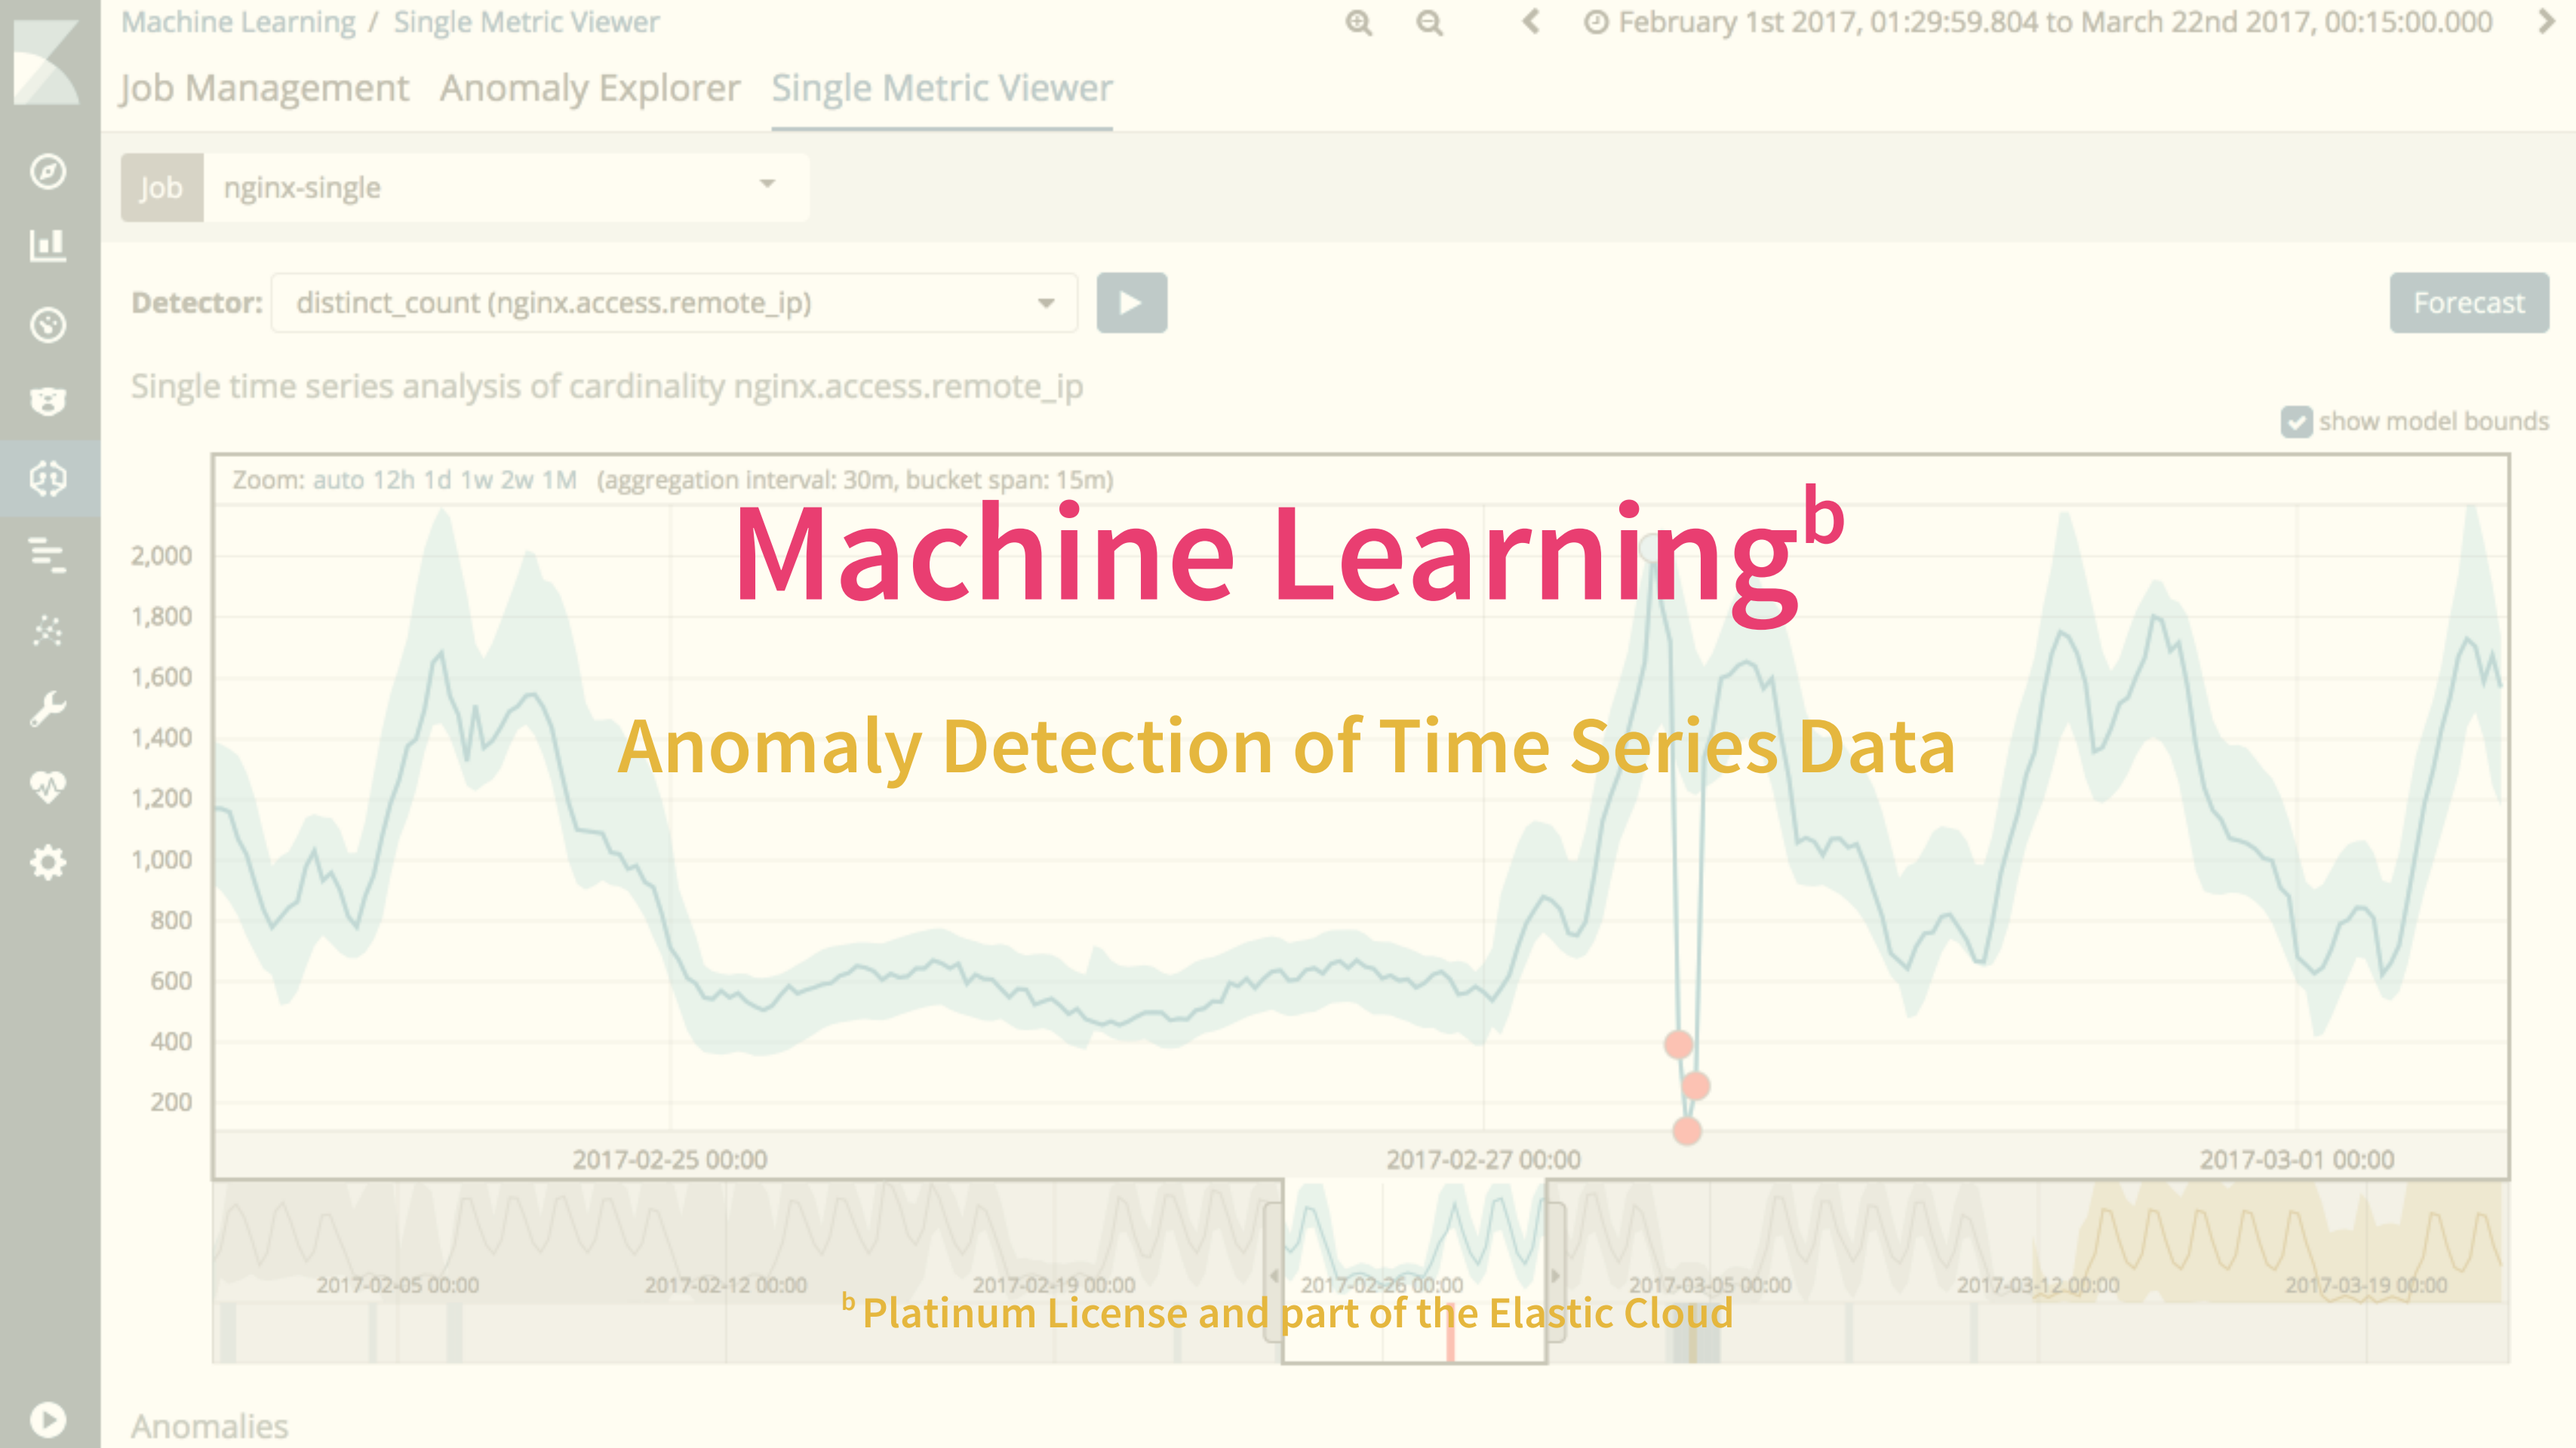

Machine Learning[b]

Anomaly Detection of Time Series Data

[b] Platinum License and part of the Elastic Cloud

Job nginx-single

Detector: distinct_count (nginx.access.remote_ip)

Forecast

Single time series analysis of cardinality nginx.access.remote_ip

show model bounds

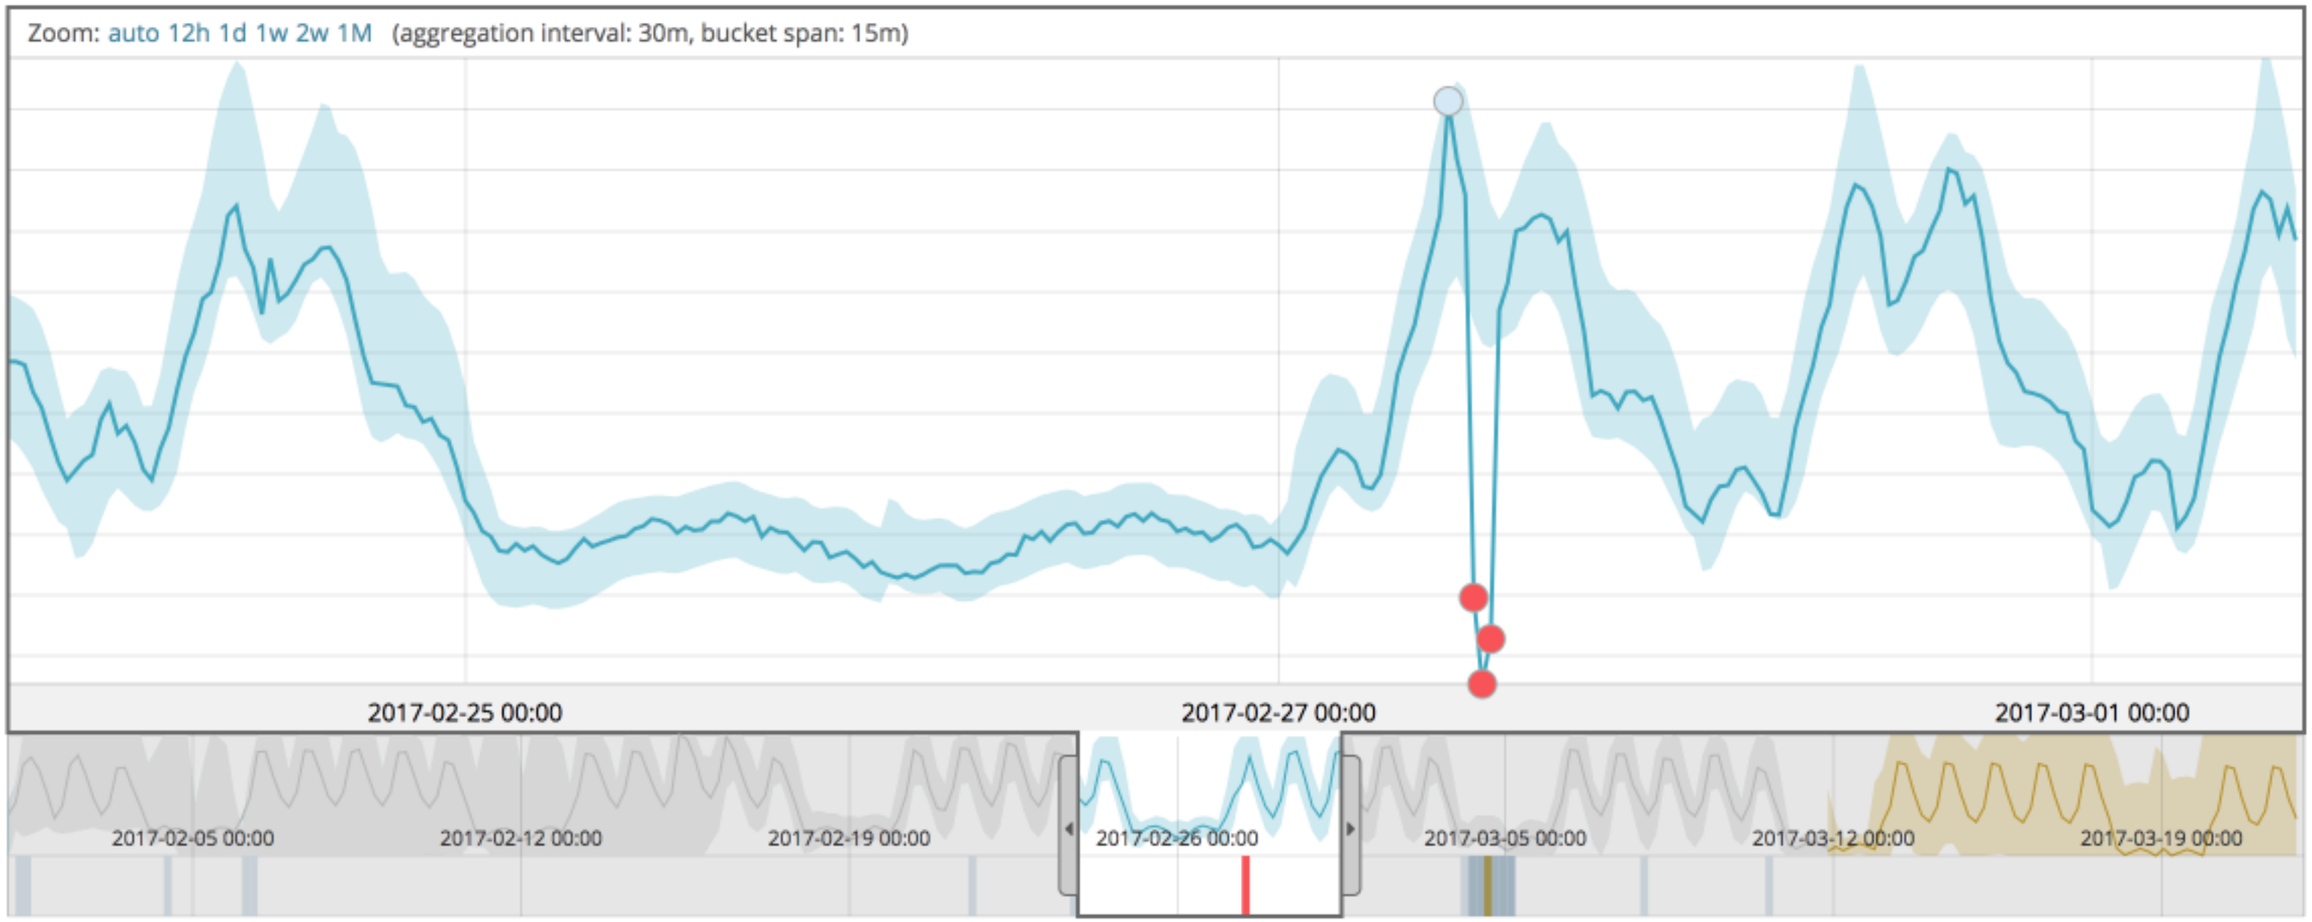

Anomalies

SECURITY[C]

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

elastic

Q&A + YOUR APPS

elastic

Conclusion

System metrics & network

Filebeat modules & Auditbeat

Application logs

Uptime

Application metrics

Request tracing

Benchmarks

Fair

Reproducible

Close to Production

elastic

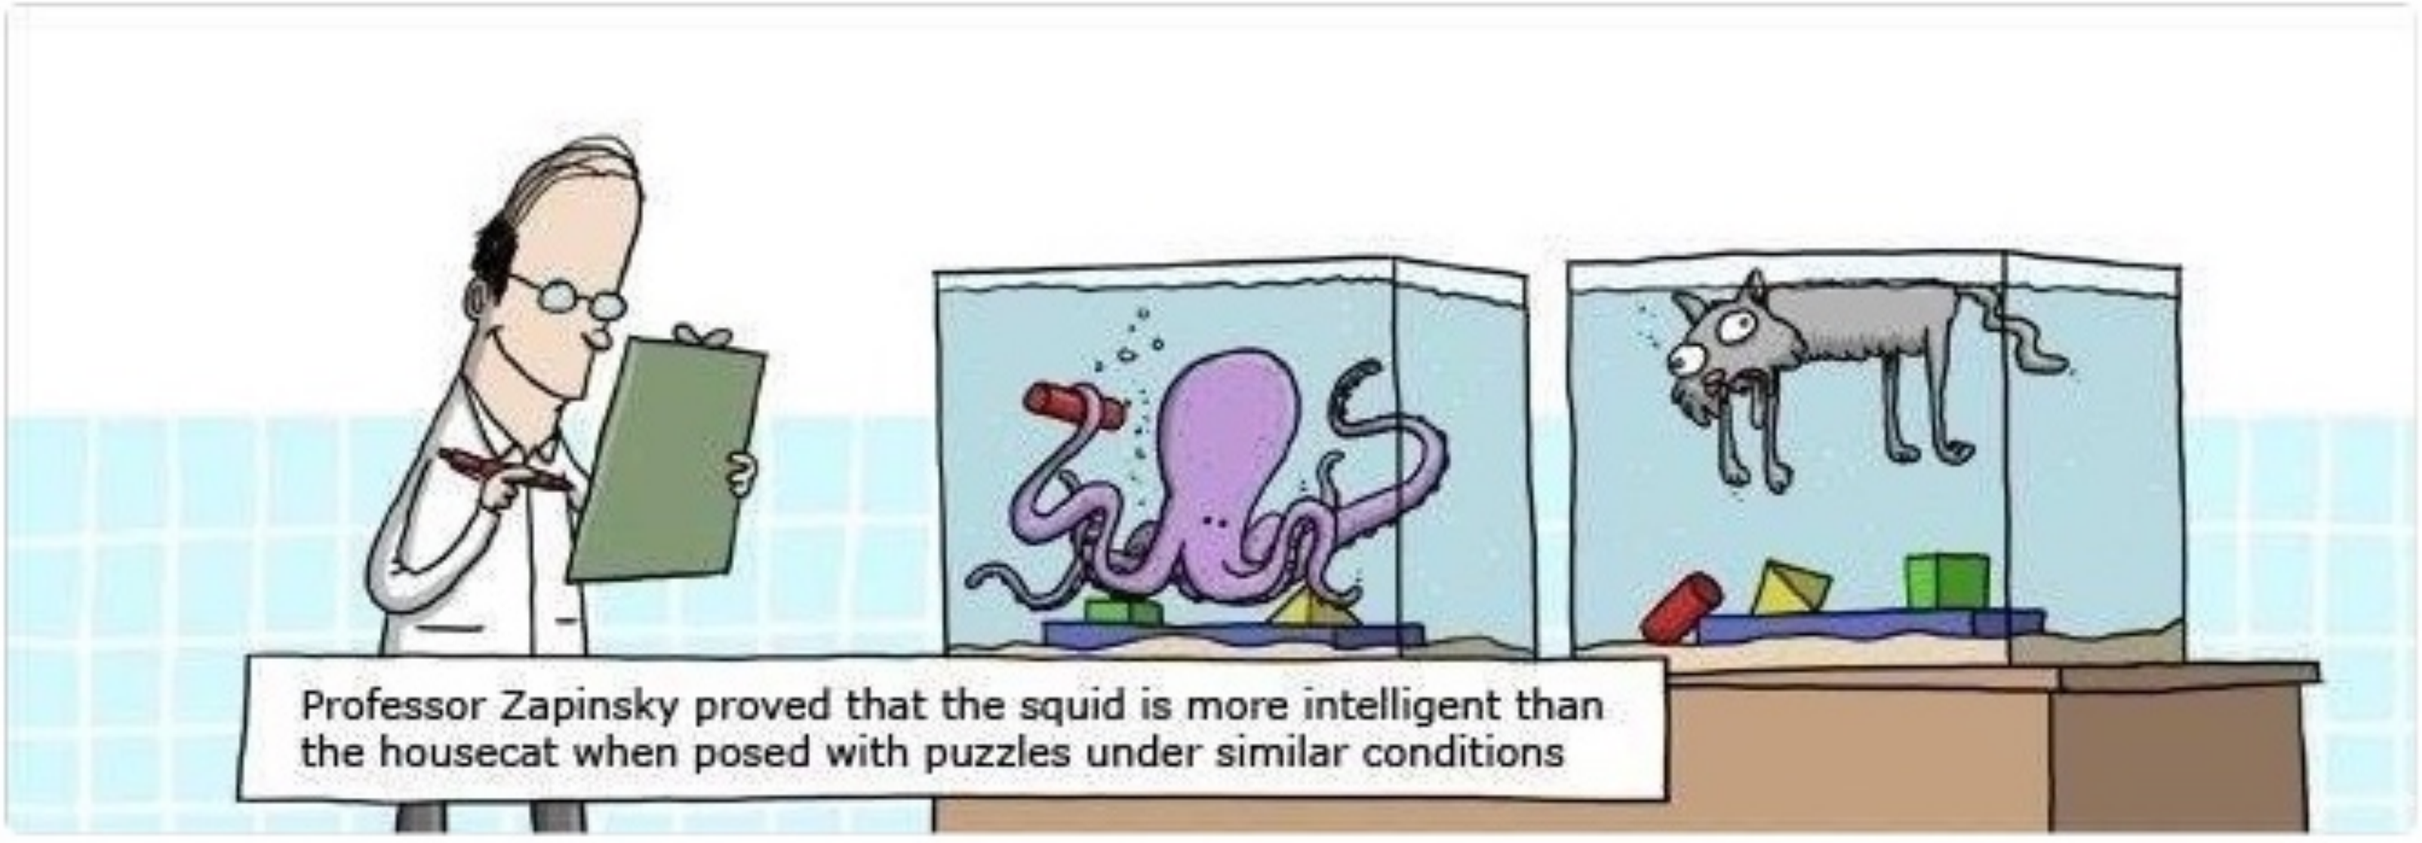
Professor Zapinsky proved that the squid is more intelligent than the housecat when posed with puzzles under similar conditions

Code

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

elastic

Thank You

Philipp Krenn @xeraa

PS: Sticker