与「十」俱进，码出未来

2018源创会年终盛典

主办方 开源中国 oschina.net

基于 Elastic Stack 的 Logging、Metrics 和 APM 分析

曾勇

Dec 2018

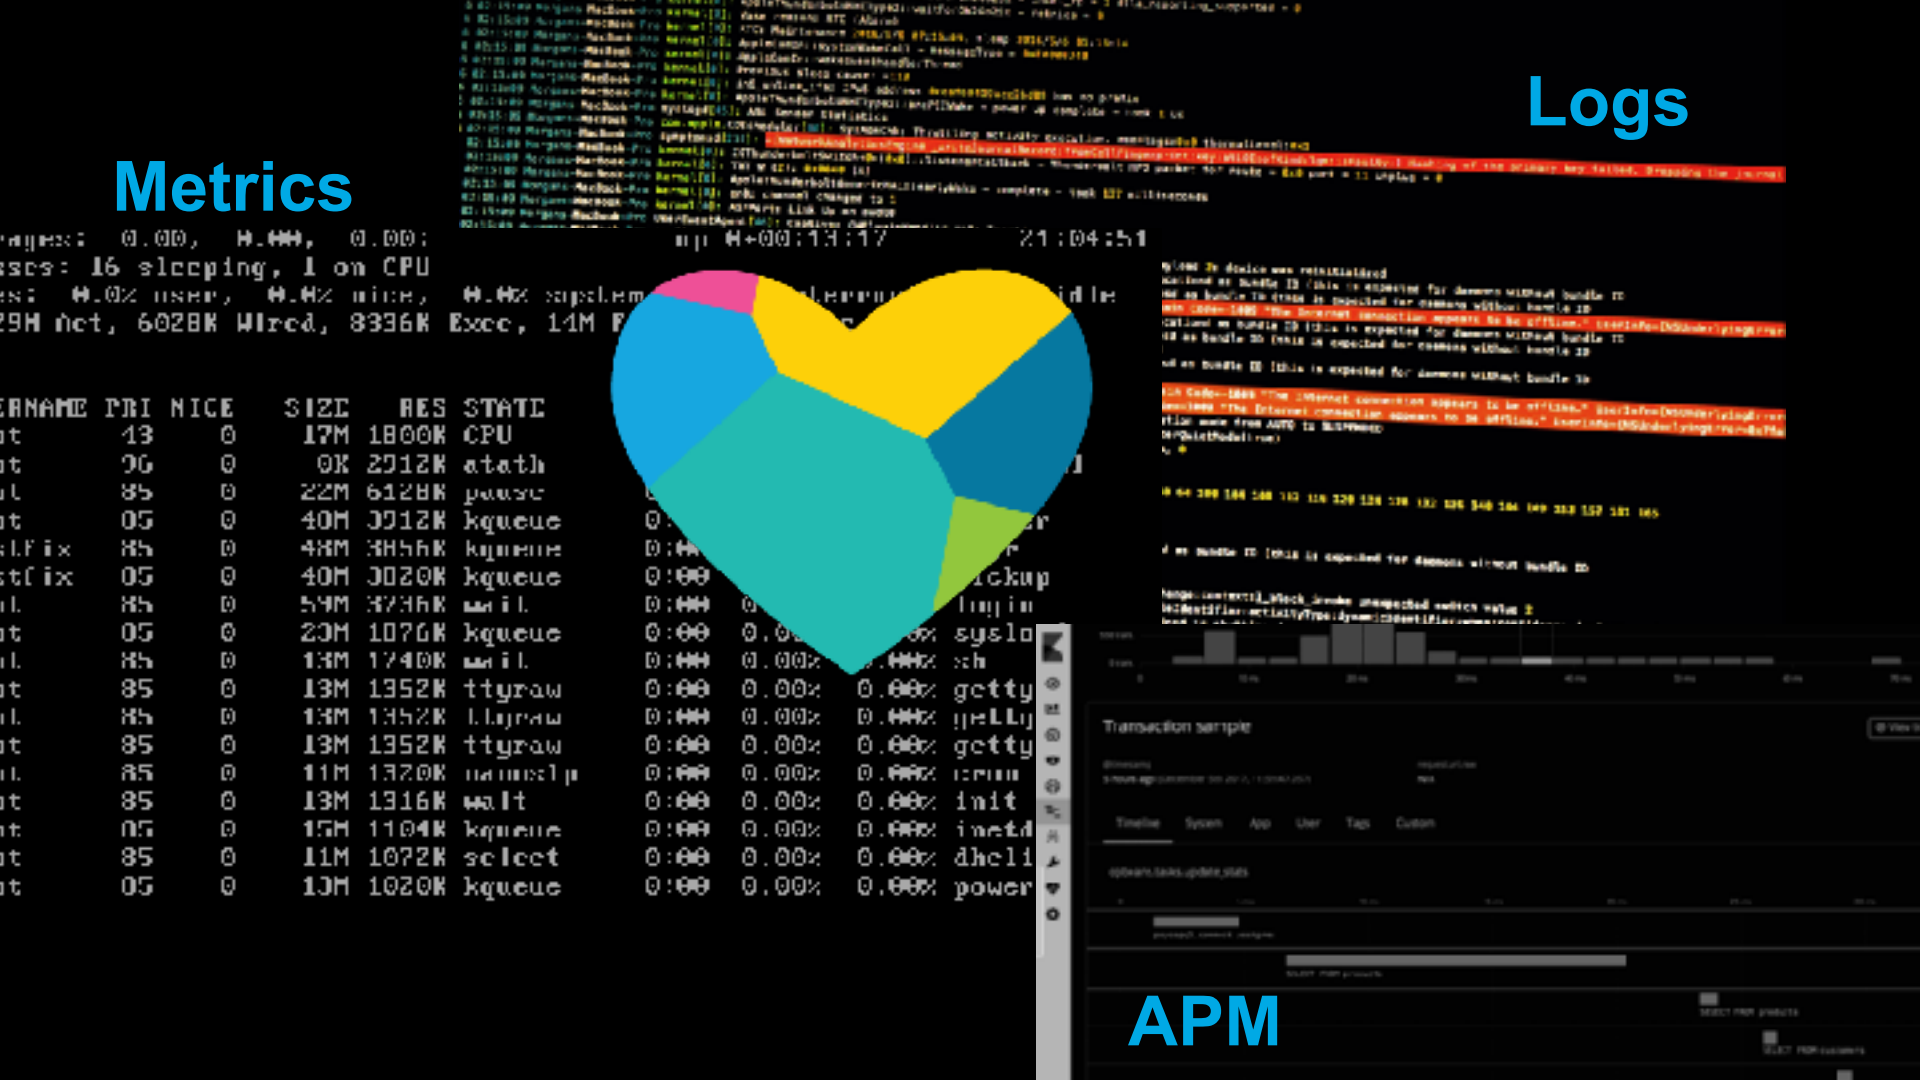
Metrics
Logs
APM
Transaction sample

Logs、Metrics、APM
都在 Elastic Stack 的好处

统一的报表

相同的 UI 界面，KPI 摘要或是根因分析

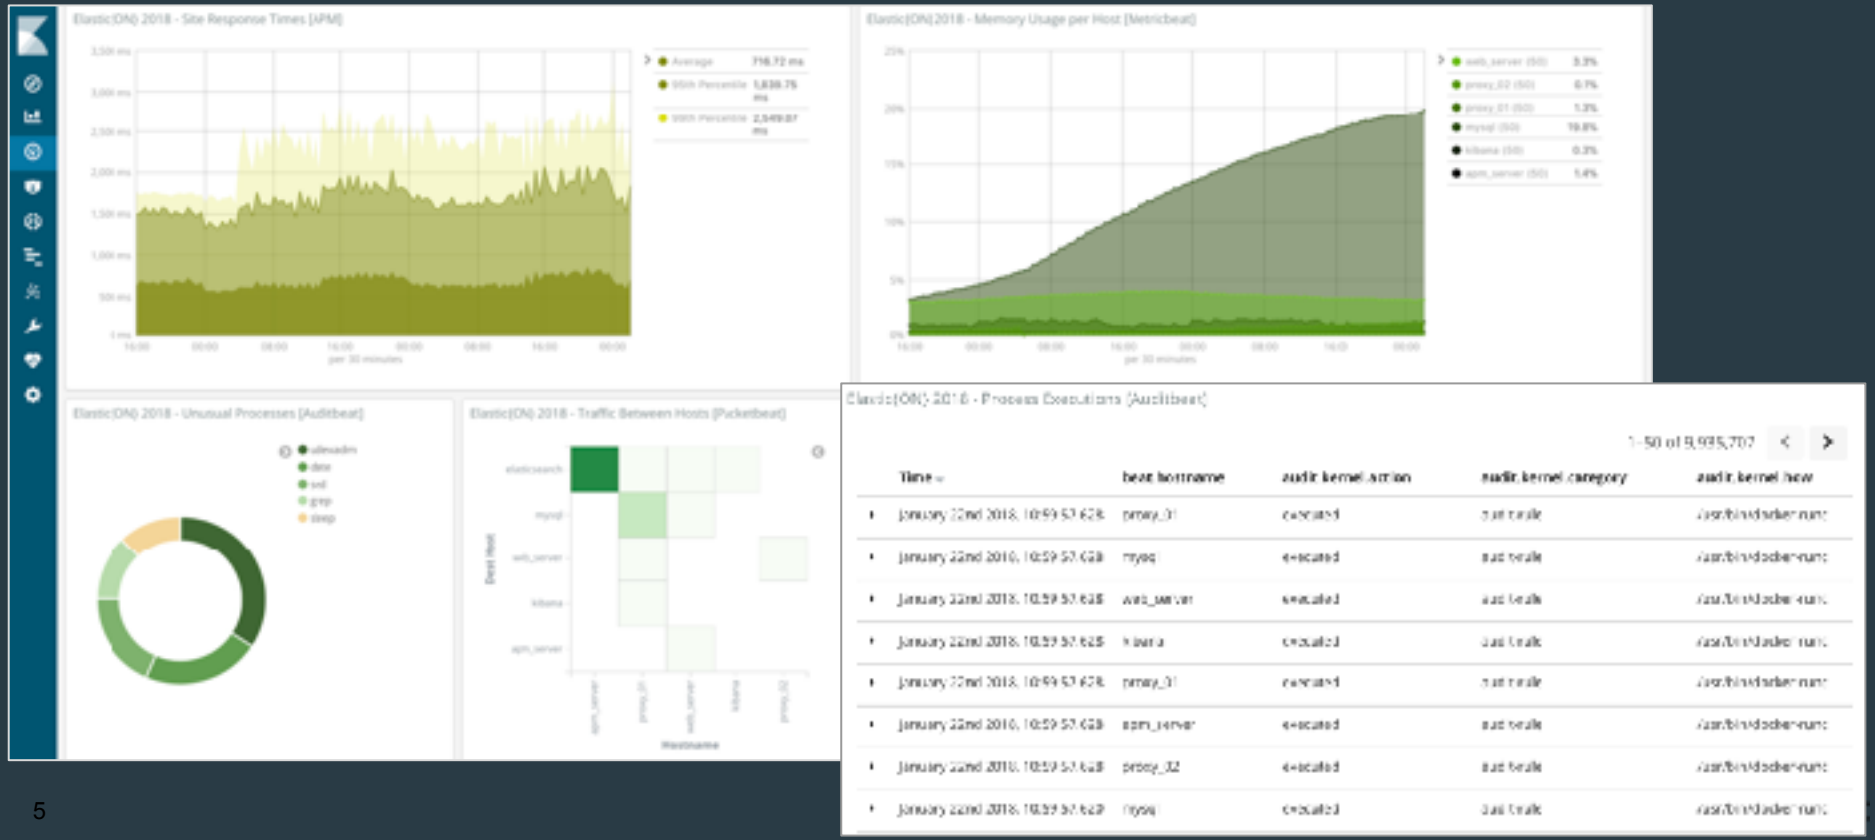

统一的告警

从任何数据的触发统一的告警

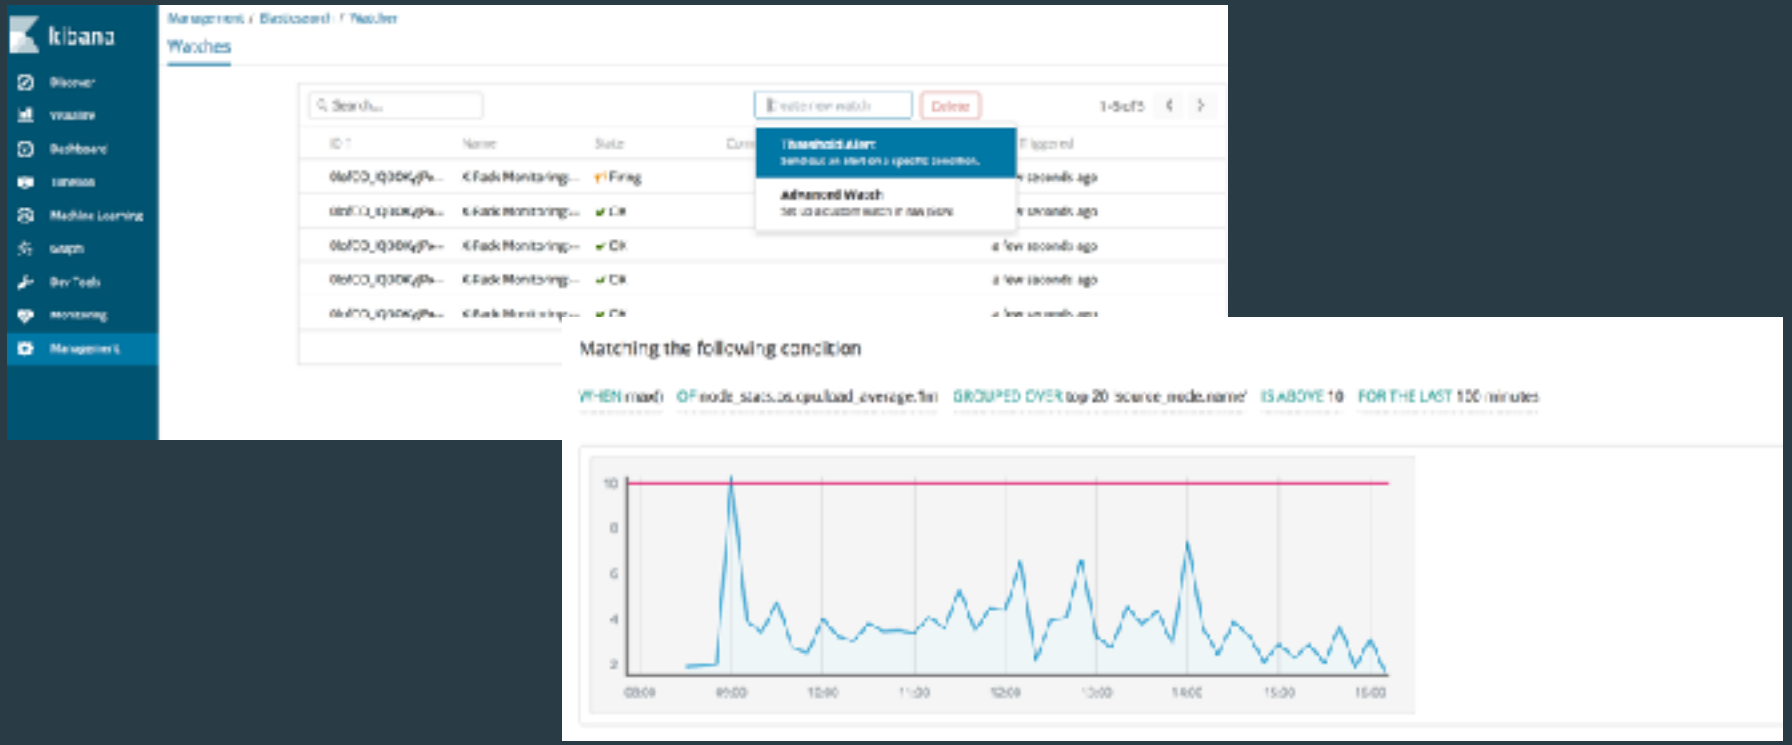

统一的机器学习

关联多个数据源，实现更智能的异常检测

elastic

实施成本

一站式技术解决方案节省运营成本

Elastic Stack for logs

elastic

Metrics vs Logs

日志是事件在时间上的顺序记录

64.242.88.10 - - [07/Mar/2017:16:10:02 -0800] "GET /mailman/listinfo/hsdivision HTTP/1.1" 200 6291

64.242.88.10 - - [07/Mar/2017:16:11:58 -0800] "POST /twiki/bin/view/TWiki/WikiSyntax HTTP/1.1" 404 7352

64.242.88.10 - - [07/Mar/2017:16:20:55 -0800] "GET /twiki/bin/view/Main/DCCAndPostFix HTTP/1.1" 200 5253

每个事件，输出当前正在发生的事情。

预置常用模块，开箱即用

- 常用数据类型，开箱即用
- 从数据到报表，一键直达
- 全自动的日志解析和处理
- 内置报表、告警规则和机器学习任务

elastic

Logging 模块

AUDITBEAT FILEBEAT WINLOGBEAT



Infrastructure

System
- Linux / MacOS
- Windows Events

Containers
- Docker
- Kubernetes

Applications

Databases
- MySQL
- PostgreSQL

Queues
- Kafka
- Redis

Web servers
- Apache
- Nginx

Audit data
- Filesystem
- System calls

elastic

Log 文件导入

自动的日志结构识别

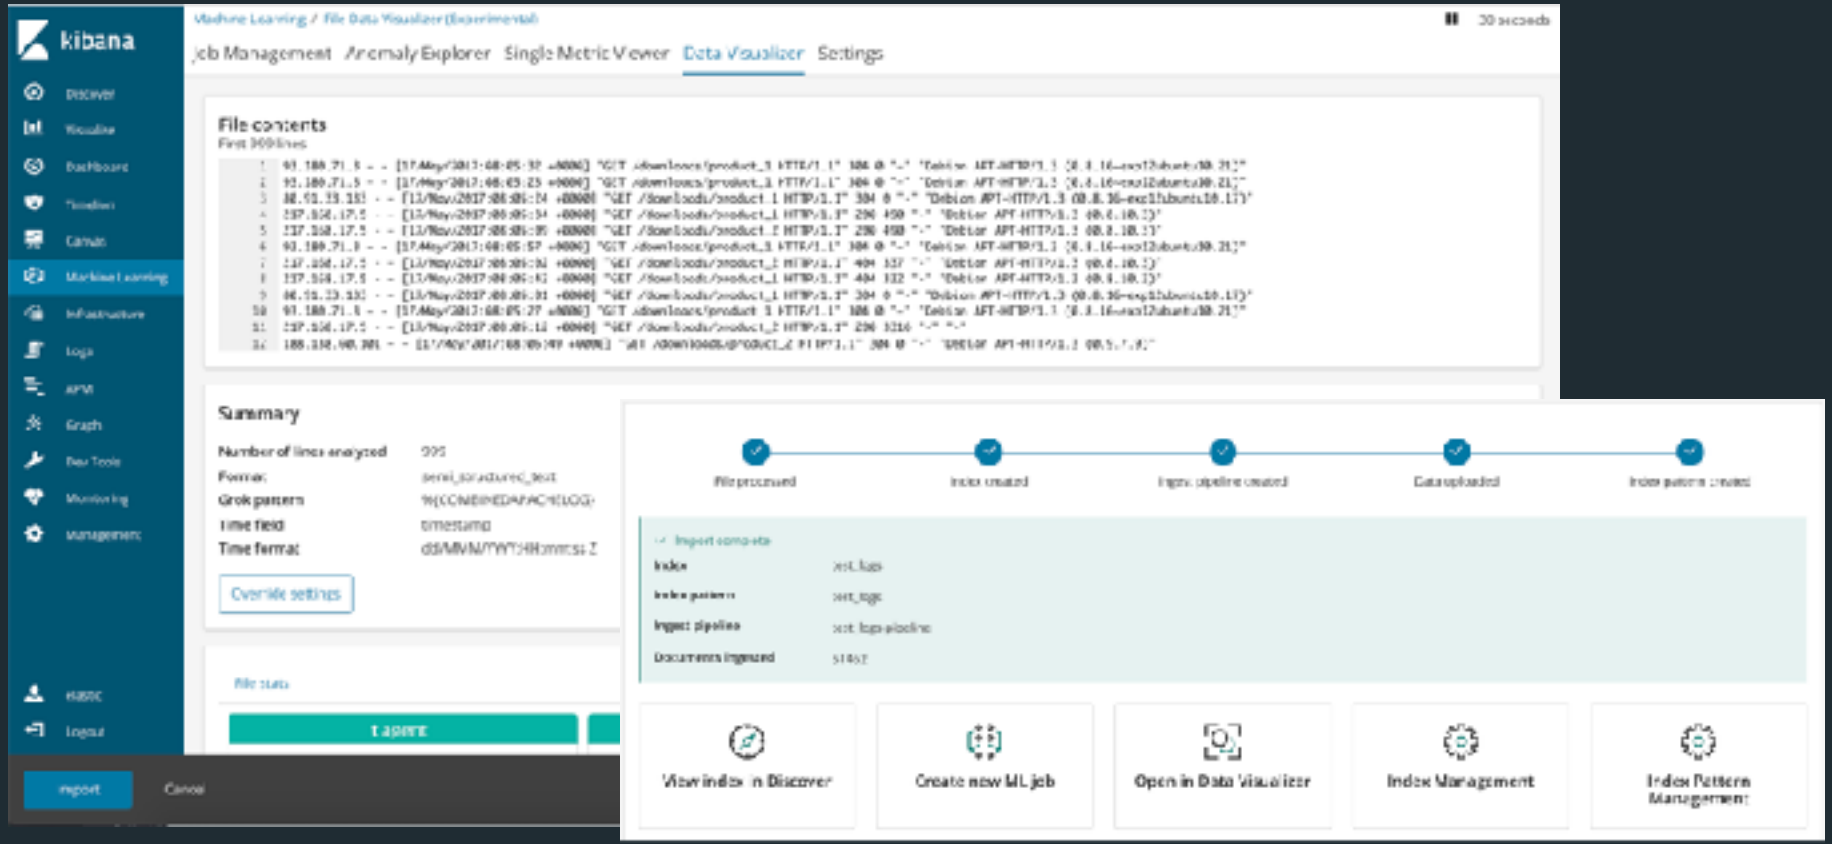

即时的日志搜索、可视化分析

Kibana Discover, Visualize, Dashboard

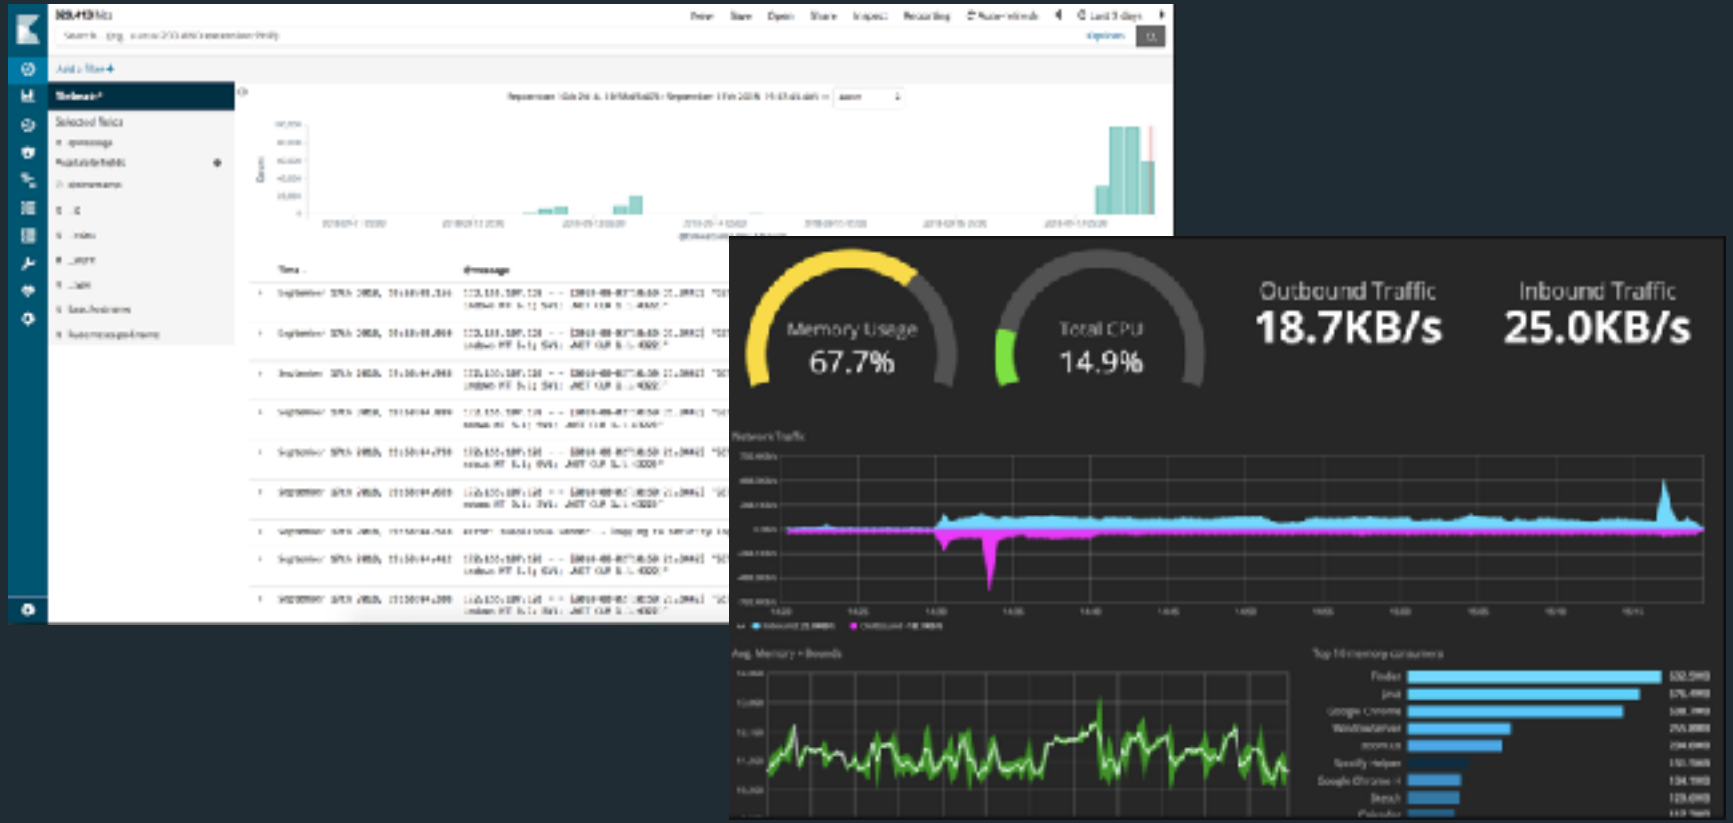

支持 Hot/Warm 架构部署

- 一键 **hot-warm** 部署
- 可大大优化数据存储

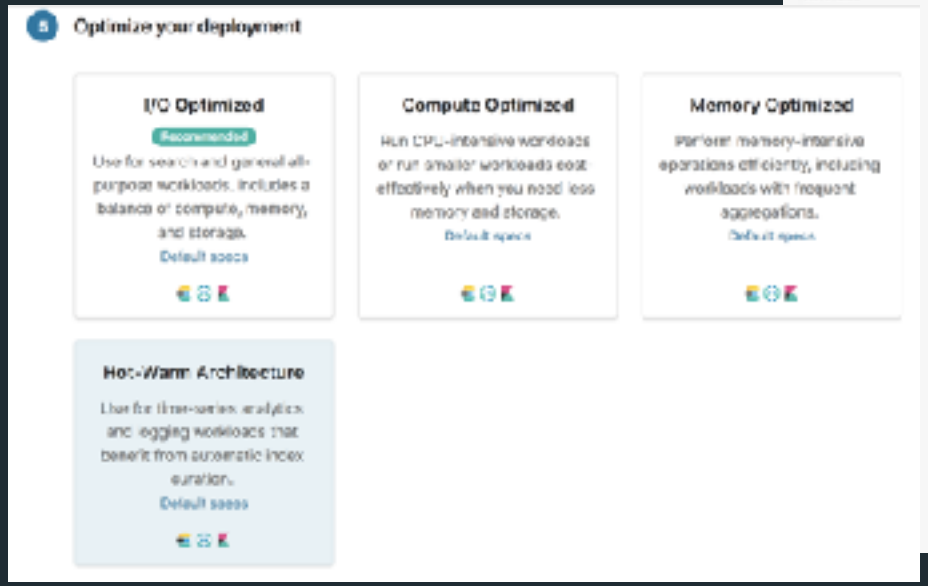

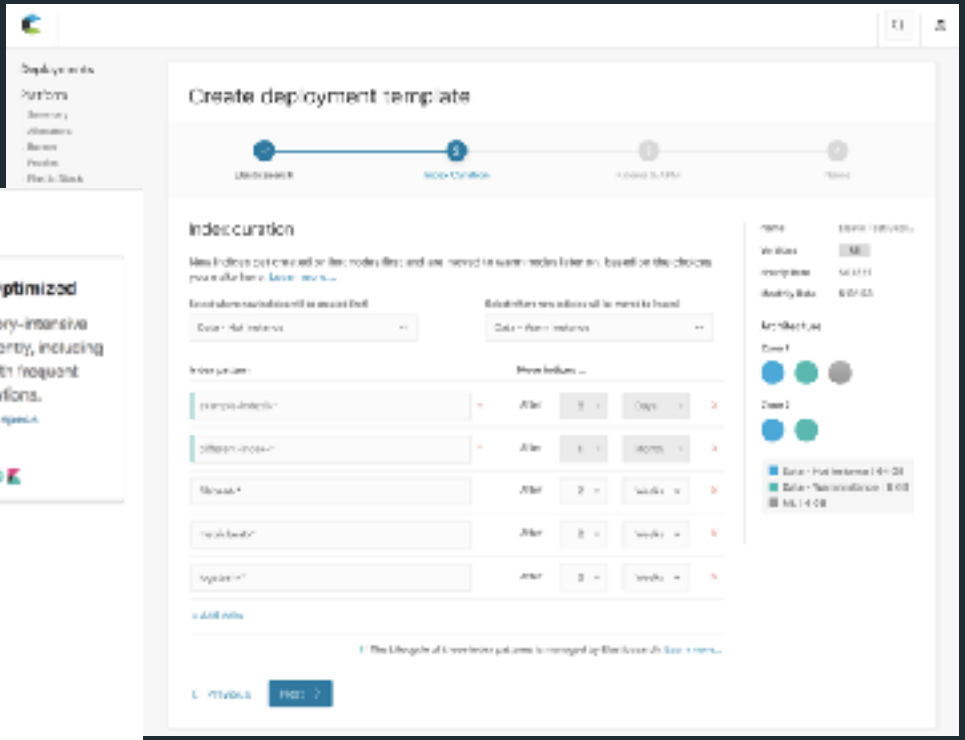

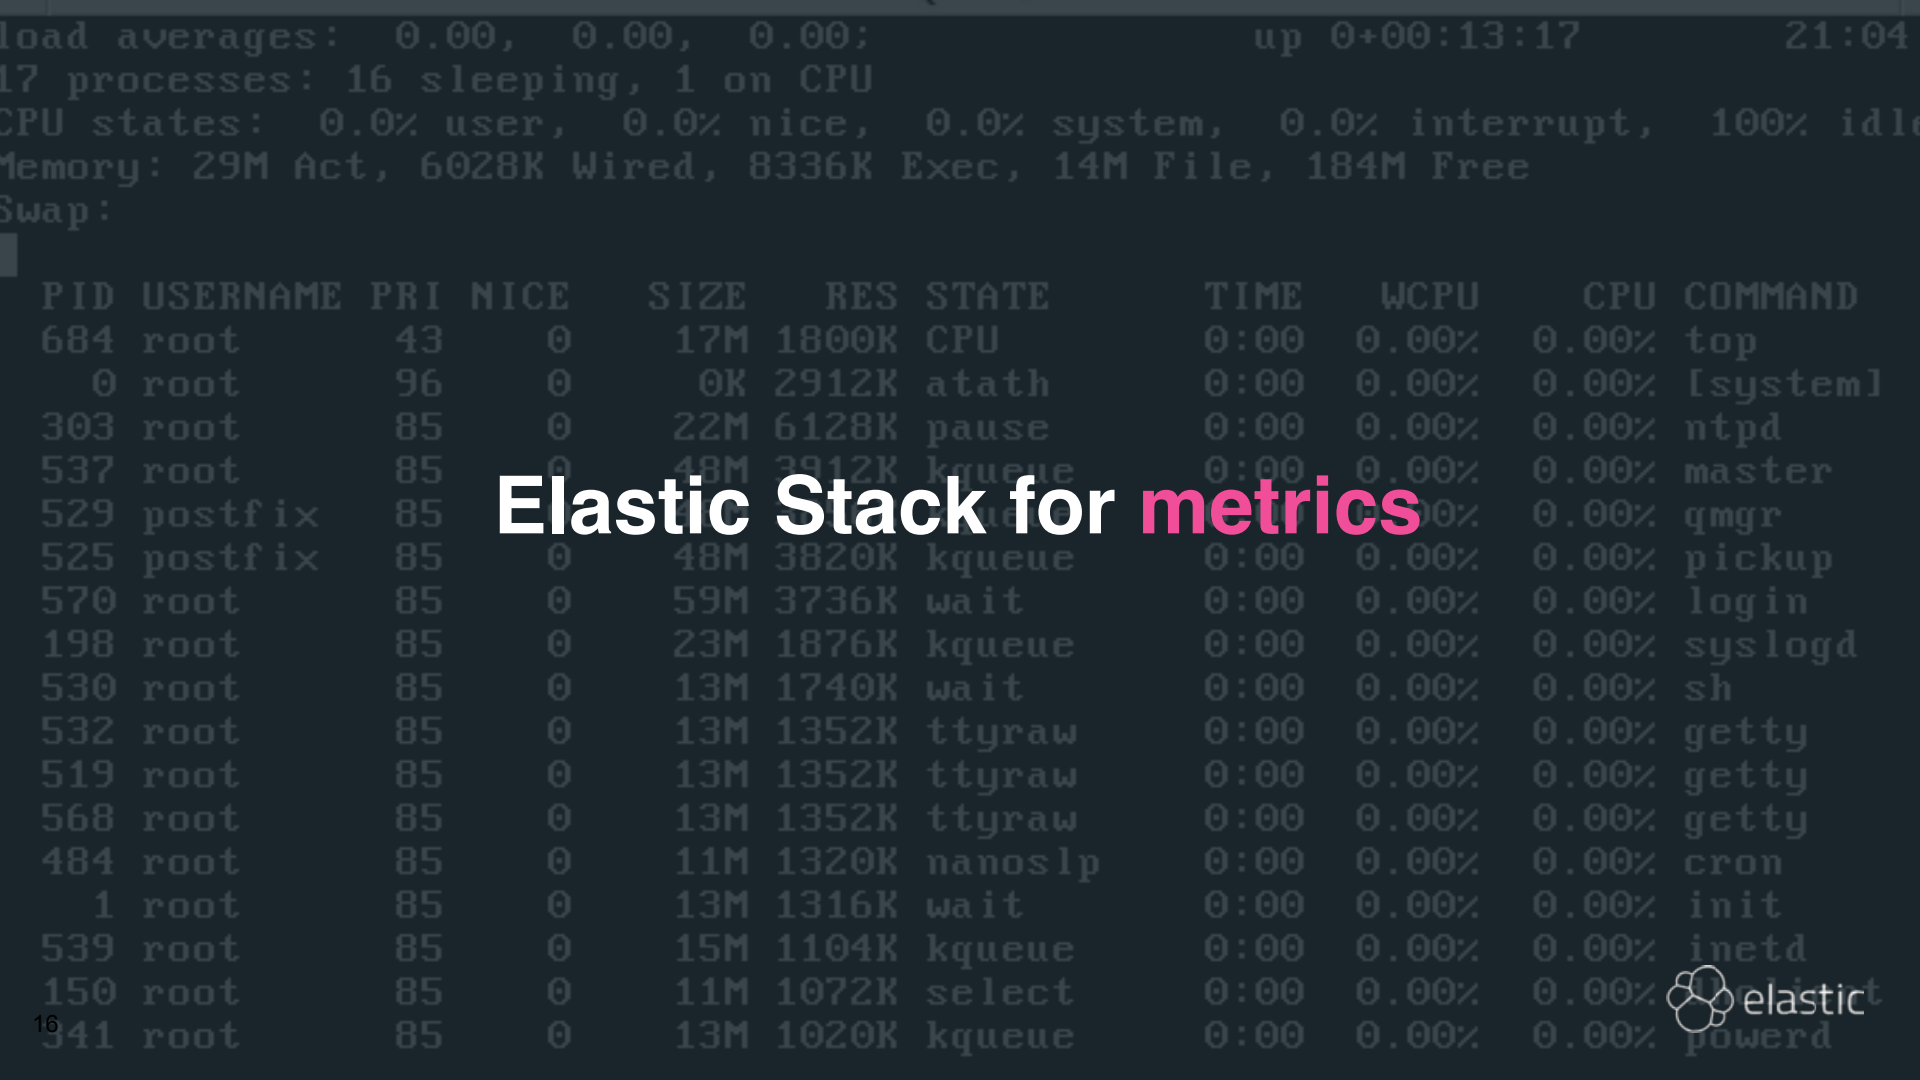

Elastic Stack for metrics

elastic

Metrics vs Logs

日志是事件在时间上的顺序记录

64.242.88.10 - - [07/Mar/2017:16:10:02 -0800] "GET /mailman/listinfo/hsdivision HTTP/1.1" 200 6291

64.242.88.10 - - [07/Mar/2017:16:11:58 -0800] "POST /twiki/bin/view/TWiki/WikiSyntax HTTP/1.1" 404 7352

64.242.88.10 - - [07/Mar/2017:16:20:55 -0800] "GET /twiki/bin/view/Main/DCCAndPostFix HTTP/1.1" 200 5253

每个事件，输出当前正在发生的事情。

Metrics 是一系列 KPI 的周期度量

07/Mar/2017 16:10:00	all	2.58	0.00	0.70	1.12	0.05	95.55	server1	containerX	regionA
07/Mar/2017 16:20:00	all	2.56	0.00	0.69	1.05	0.04	95.66	server2	containerY	regionB
07/Mar/2017 16:30:00	all	2.64	0.00	0.65	1.15	0.05	95.50	server2	containerZ	regionC

每隔x分钟，测量CPU负载并将其打印出来，并用元数据进行注释。

elastic

Elasticsearch 在 Metrics 存储上的演进

Elasticsearch 搜索和数值型的分析

Inverted Index for full-text search

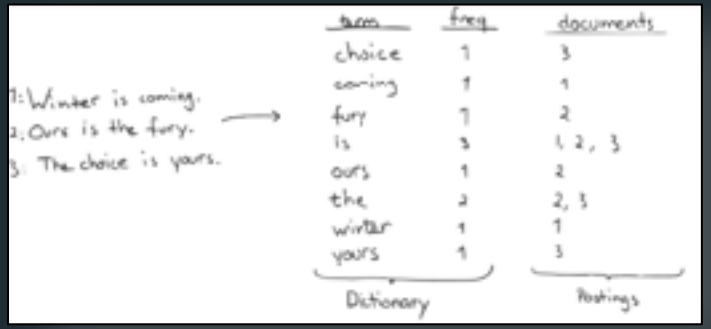

Columnar store for structured data

usernid	first	middle	last	city	state
johr123	John	James	Smith	Alamo	California
jrice	Jill	Amy	Rice		
mt123	Jeff		Twain	Toledo	Ohio
sadams	Sue		Adams		
adoe	Amy		Doe	Miami	Florida

BKD Trees for numerical operations

Rollups

elastic

Elasticsearch 早期

2010 年

- Elasticsearch 主要用于 **应用搜索**
- Lucene 数据结构: **倒排索引**

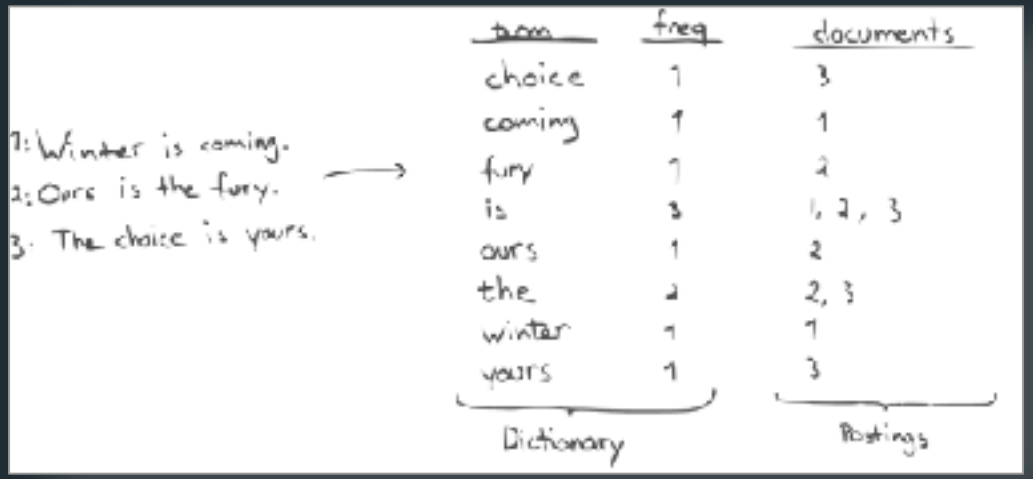

elastic

Elasticsearch 进化到对数据分析的支持

~ 2010 到 2014

- Elasticsearch 1.0 开始支持 **列式存储** (基于 Lucene **“doc values”**)
- 结构化的文本和数值型数据能够更快的取出来并用于分析

usernid	first	middle	last	city	state
john123	John	James	Smith	Alamo	California
jrice	Jill	Amy	Rice		
mt123	Jeff		Twain	Toledo	Ohio
sadams	Sue		Adams		
adoe	Amy		Doe	Miami	Florida

https://www.elastic.co/blog/elasticsearch-as-a-column-store

Elasticsearch 改善存储效率

2016

- Elasticsearch 5.0 添加更多数据结构来高效存储数值型数据 (**BKD Trees**)
- 数值型数据以及地理位置数据默认都采用这种数据结构来进行存储

1-Dimension

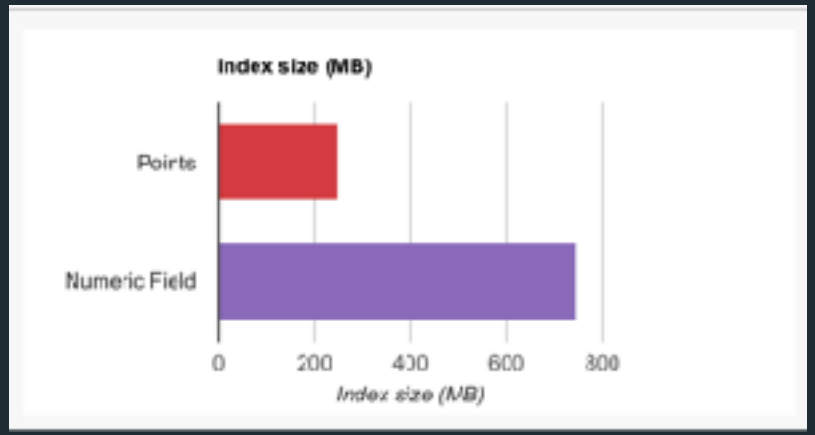

2-Dimensions

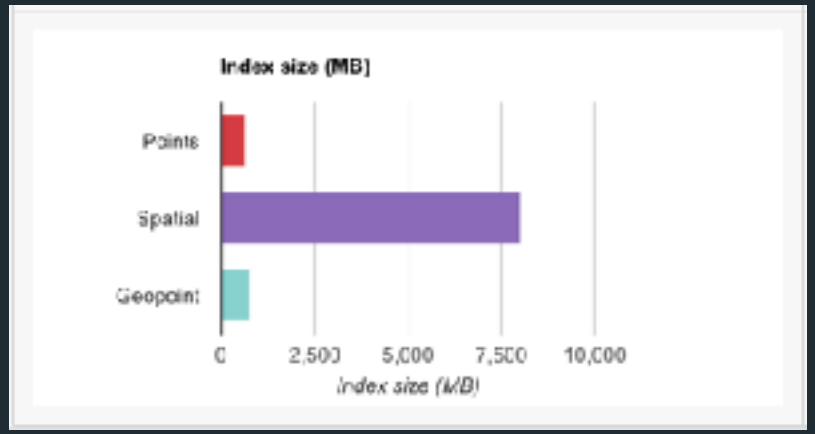

https://www.elastic.co/blog/searching-numb3rs-in-5.0

elastic

Elasticsearch 改善存储效率

2017

- Elasticsearch 6.0 改进 **sparse values** 的存储效率 (Metricbeat 节省 59.5% 的索引存储)

usernid	first	middle	last	city	state
john123	John	James	Smith	Alamo	California
jrice	Jill	Amy	Rice		
mt123	Jeff		Twain	Toledo	Ohio
sadams	Sue		Adams		
adoe	Amy		Doe	Miami	Florida

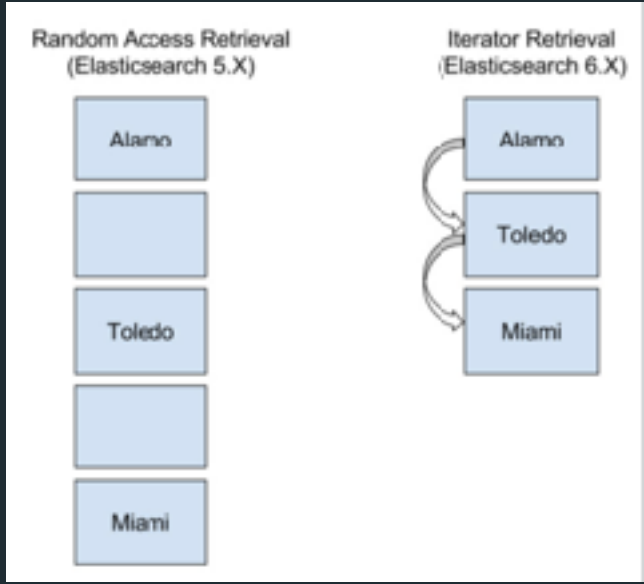

https://www.elastic.co/blog/minimize-index-storage-size-elasticsearch-6-0

elastic

Rollup 用于长期数据存储

Elasticsearch 6.3 新增

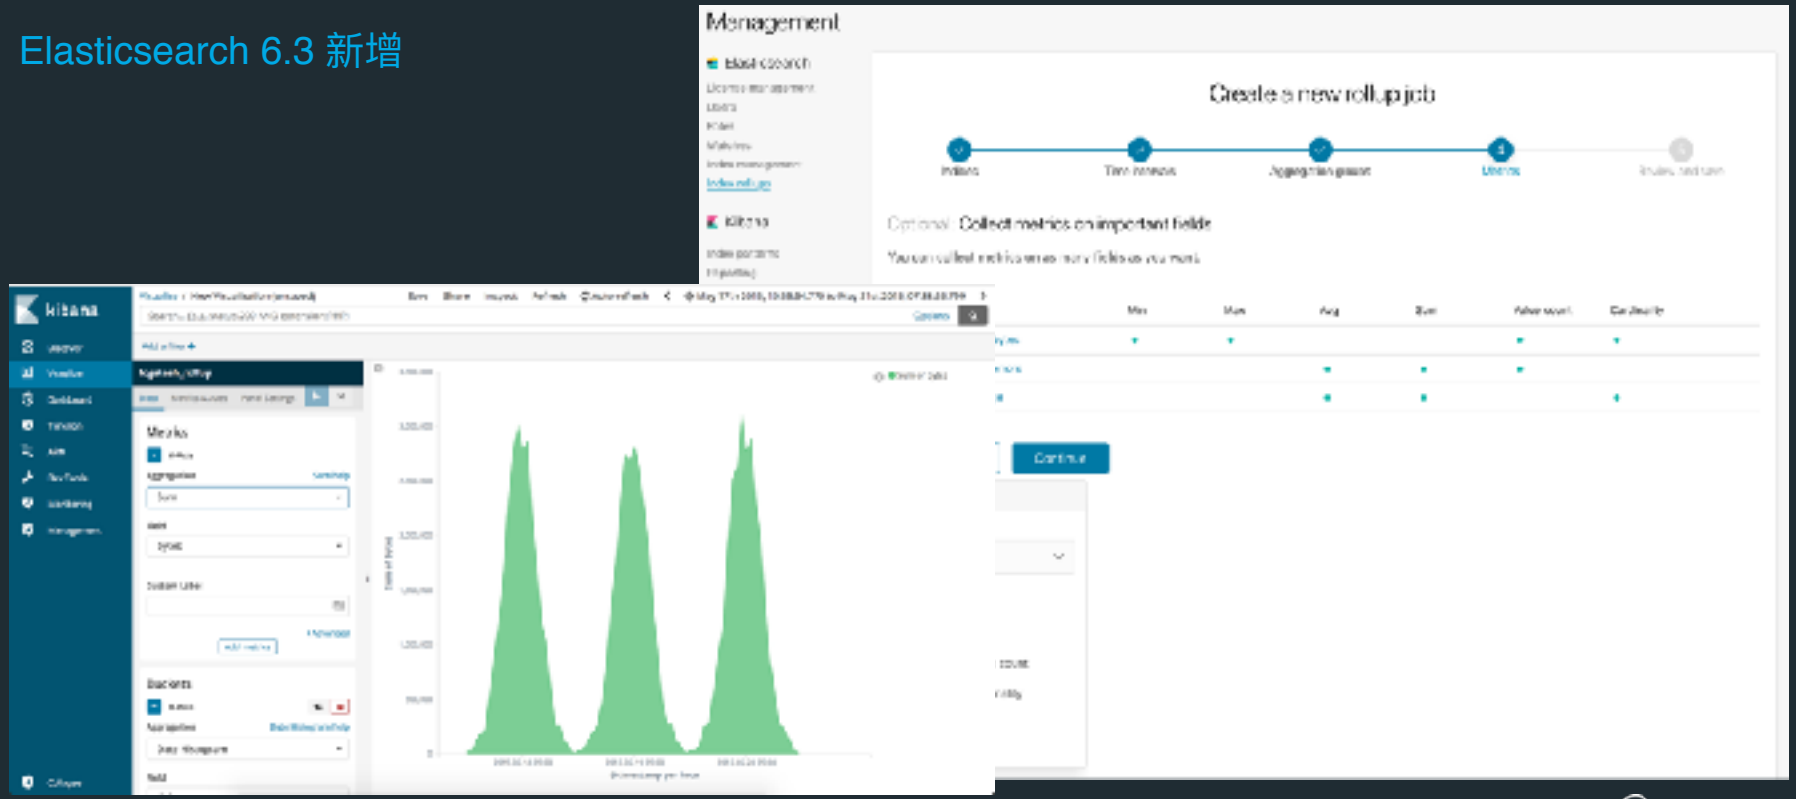

https://www.elastic.co/blog/data-rollups-in-elasticsearch-you-know-for-saving-space

elastic

Elastic Stack 作为 Metrics 解决方案

Metrics 模块

HEARTBEAT 

METRICBEAT 

PACKETBEAT 

Infrastructure

System
- Linux
- MacOS
- Windows
- Perfmon

Containers
- Docker
- Kubernetes

Virtualization
- vSphere

Cloud
- AWS
- GCP
- Azure
- DigitalOcean
- Alibaba

Network
- Netflow
- Packets
- TLS Envelope

Storage
- Ceph

elastic

Metrics 模块

HEARTBEAT METRICBEAT PACKETBEAT



Applications

Datastores
- MySQL
- PostgreSQL
- MongoDB
- Couchbase
- Aerospike
- Graphite

Queues
- Kafka
- Redis
- RabbitMQ

Caches
- Memcached

Uptime
- Heartbeat

Custom apps
- JMX/Jolokia
- PHP-FPM
- Golang

Web servers
- Apache
- Nginx

Other
- HAProxy
- Zookeeper

elastic

Heartbeat: 心跳监测

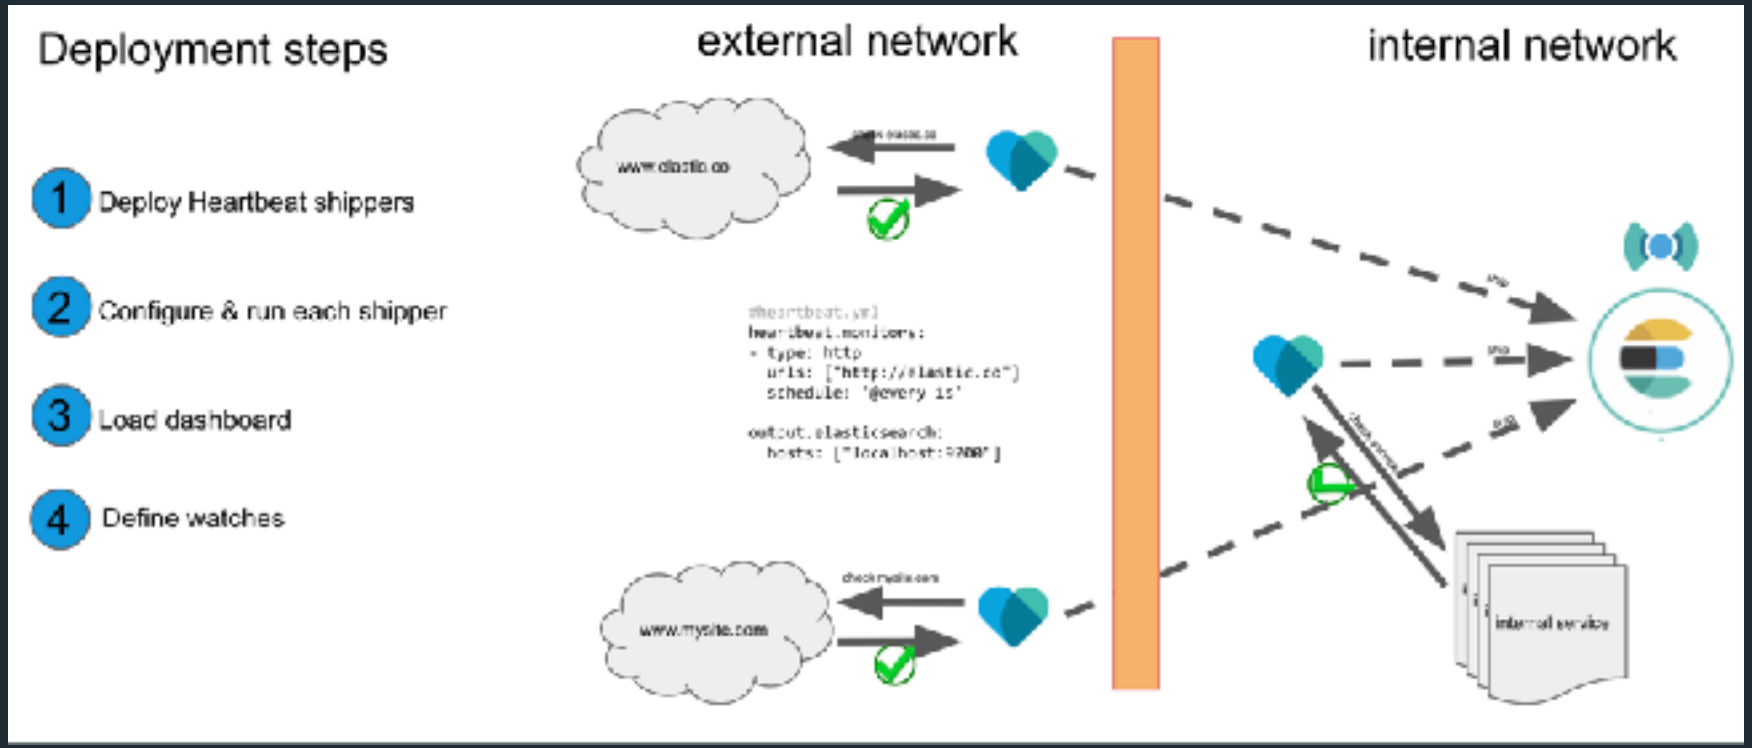

elastic

Heartbeat: 心跳监测

elastic

Functionbeat: Serverless 日志采集器

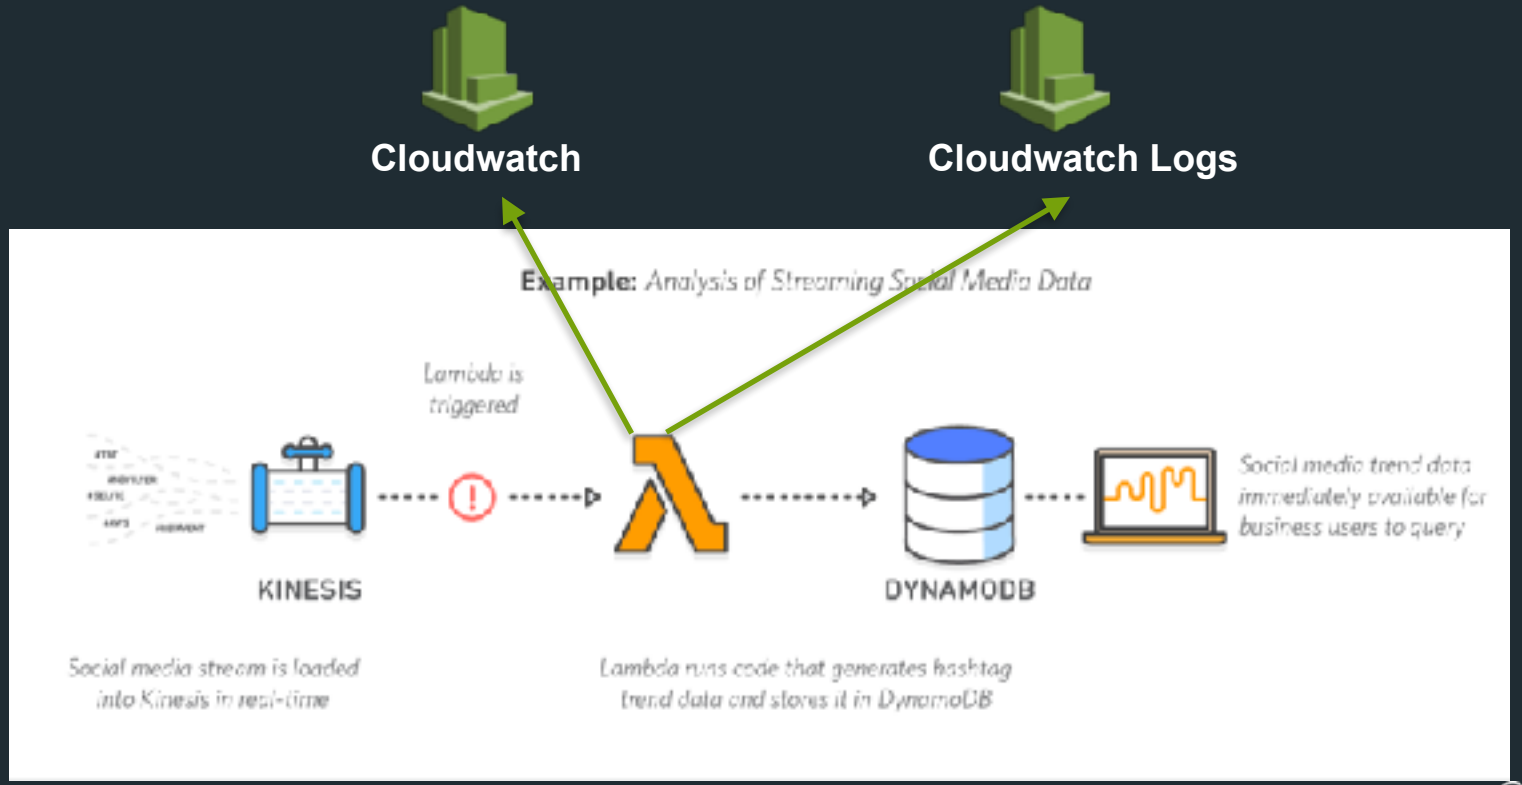

elastic

Functionbeat: Serverless 日志采集器

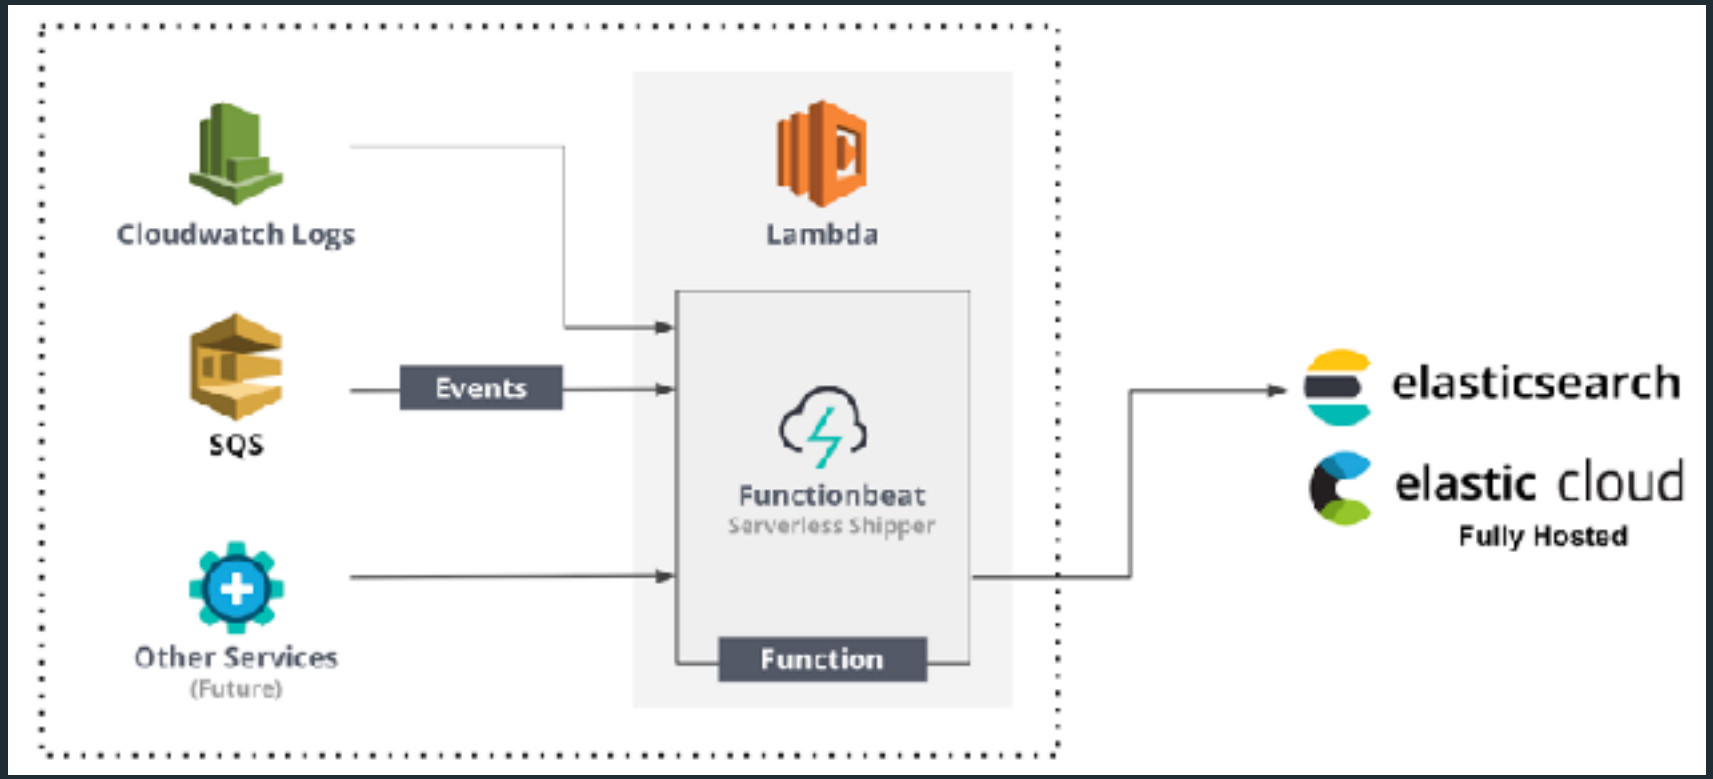

elastic

将 Logs, metrics, APM 关联起来

Elastic Common Schema

Benefits

- Correlate data from different sources
- Ability to re-use analysis content
- Ability to re-use Elastic-provided content

Status

- Beta published: github.com/elastic/ecs
- Integrating into Elastic products in progress
- Community feedback welcome!

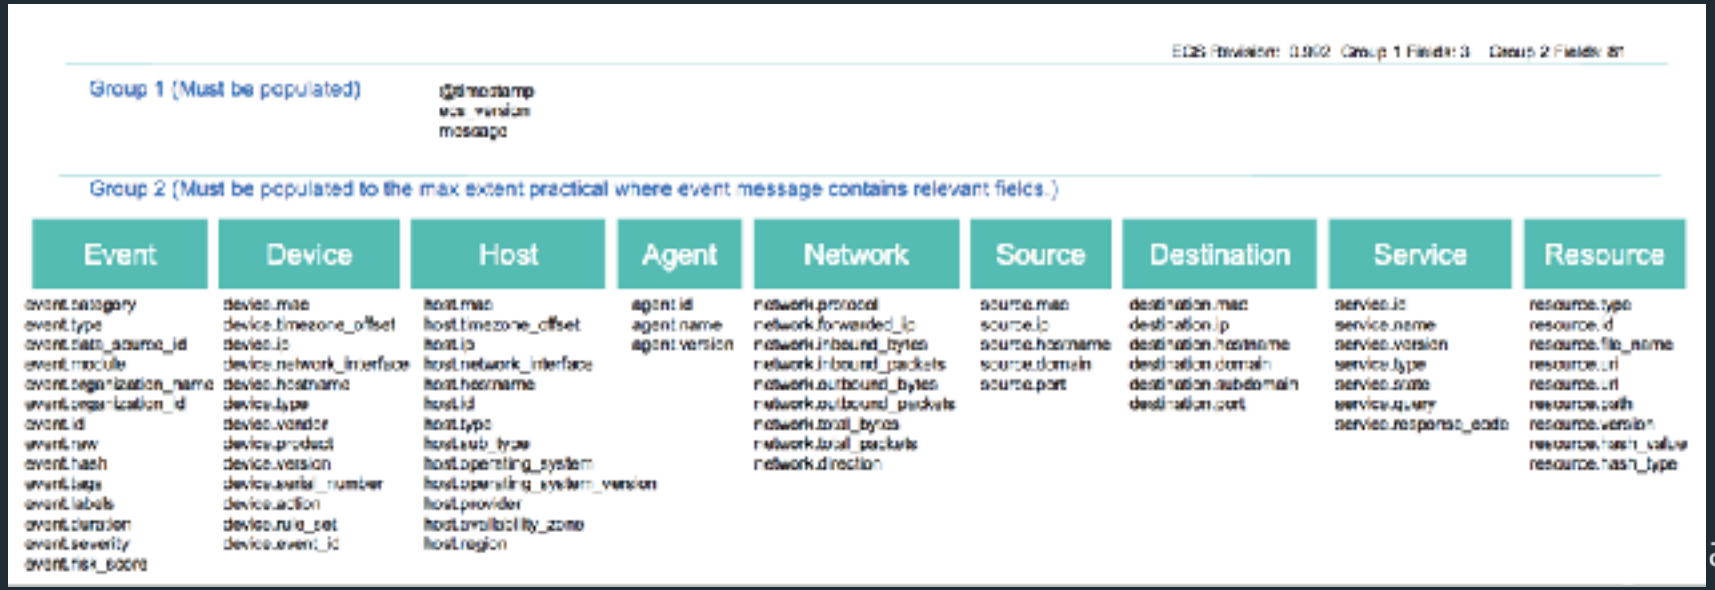

时序型数据的可视化

Time Series Visual Builder

elastic

时序型数据的可视化

Annotations

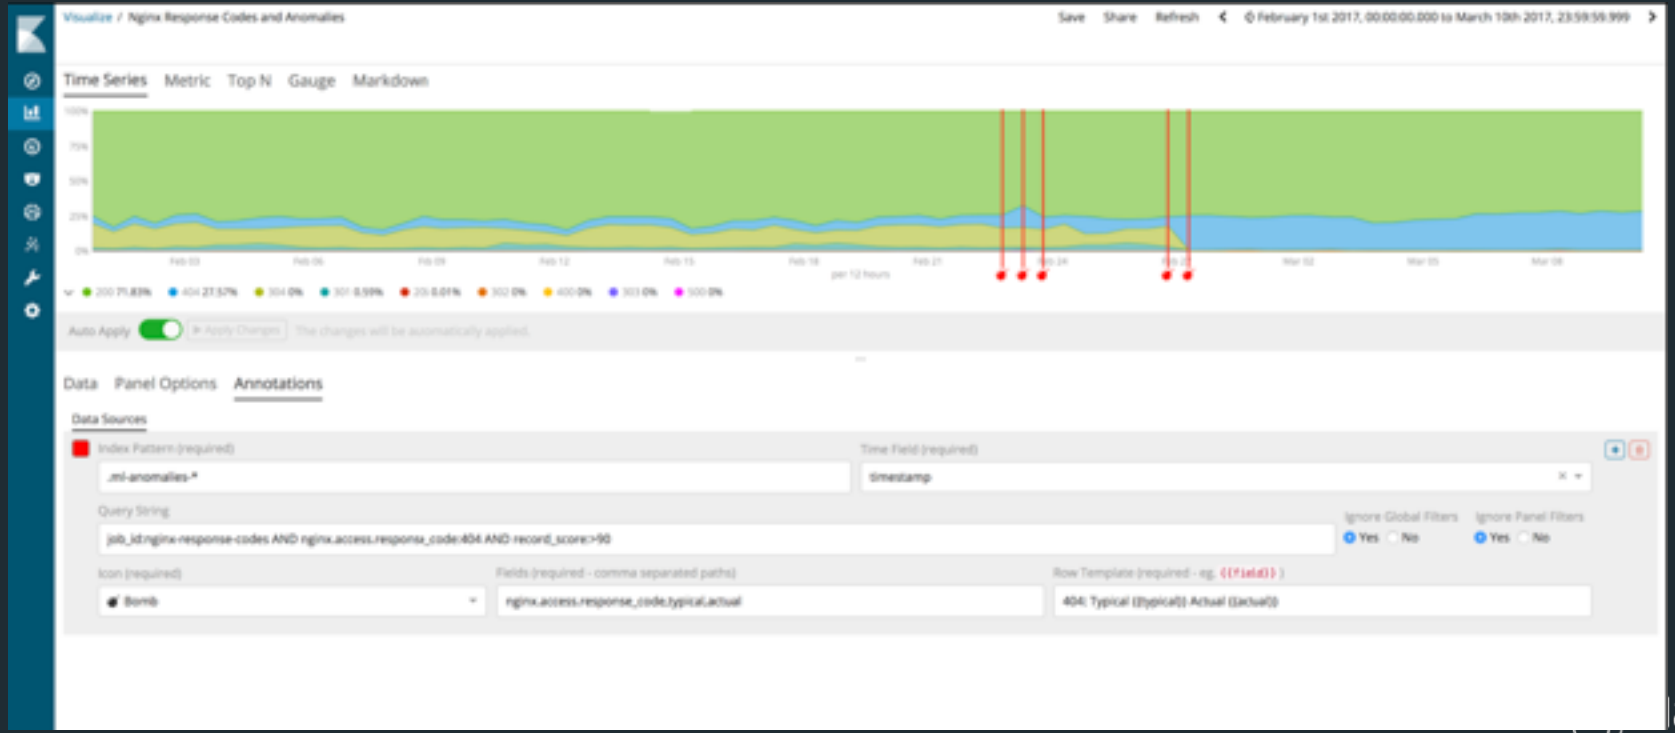

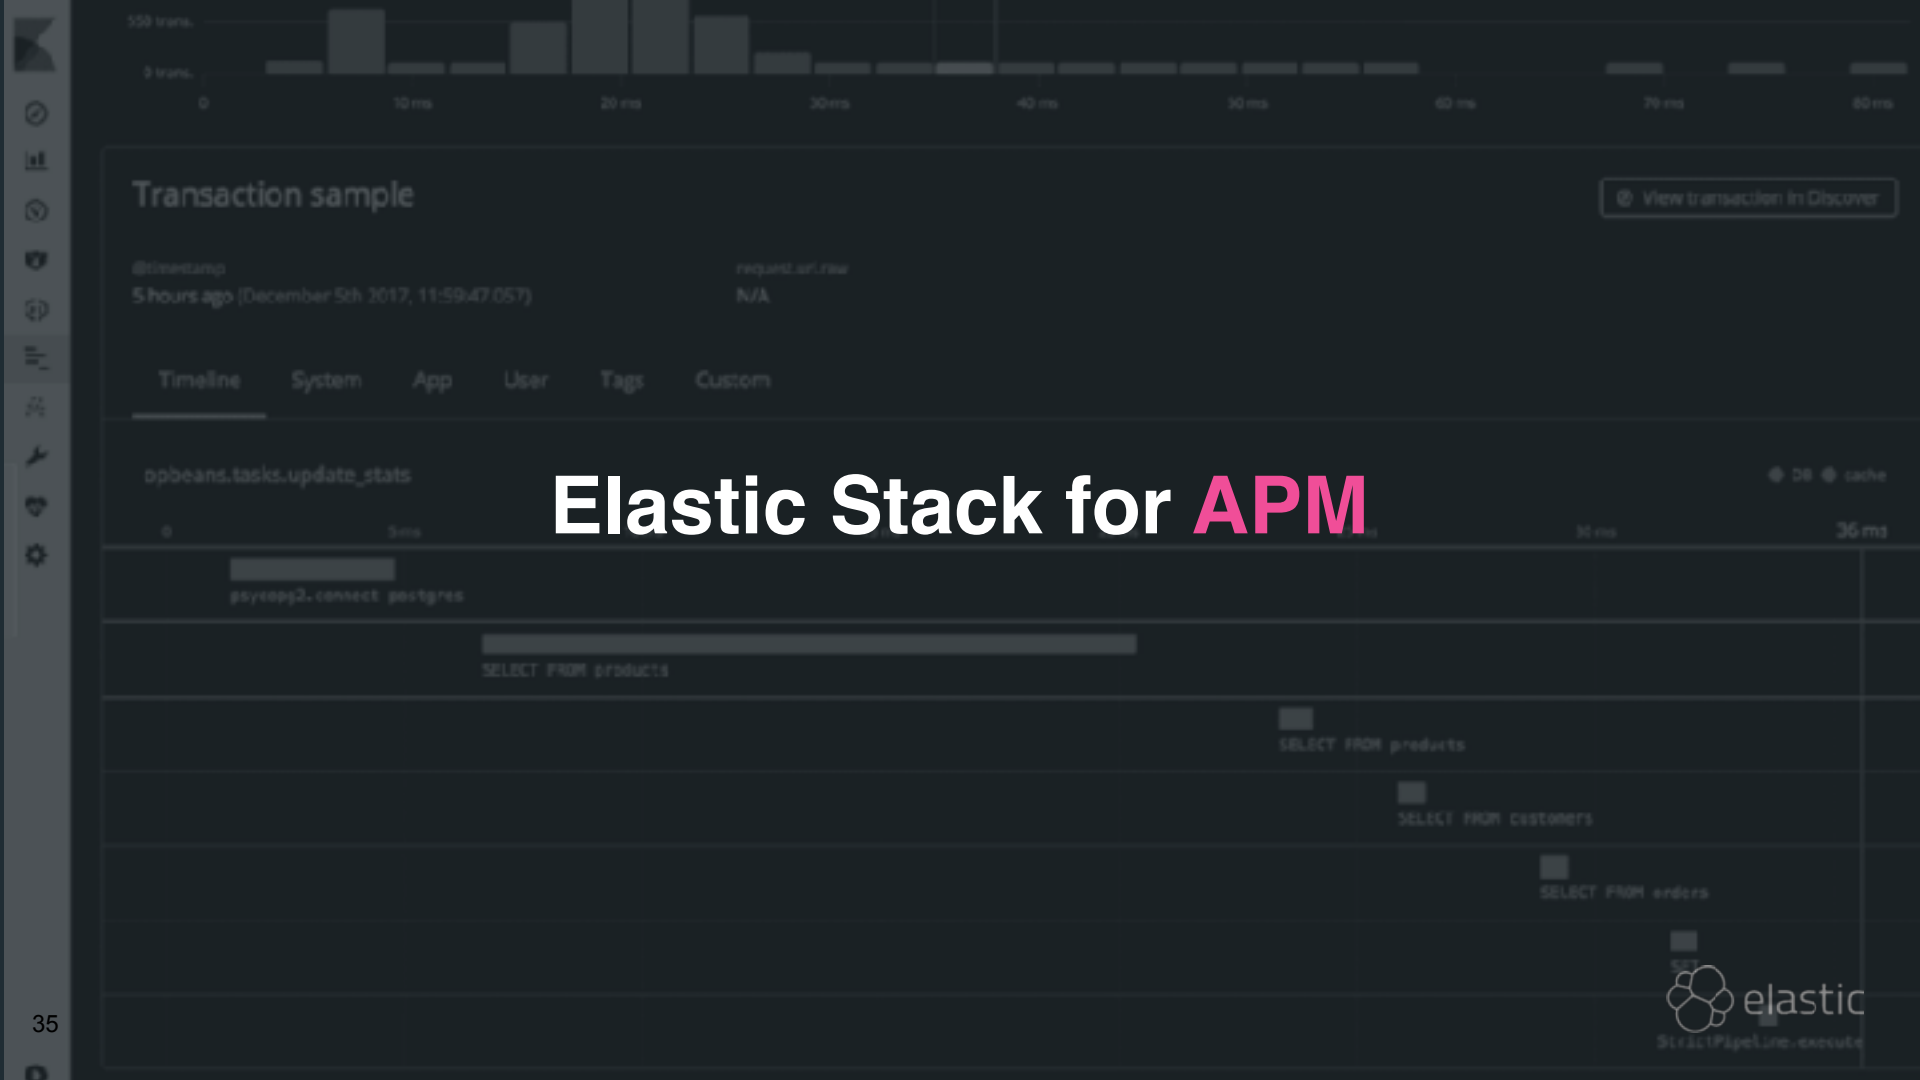

Elastic Stack for APM

elastic

什么是 APM?

Example

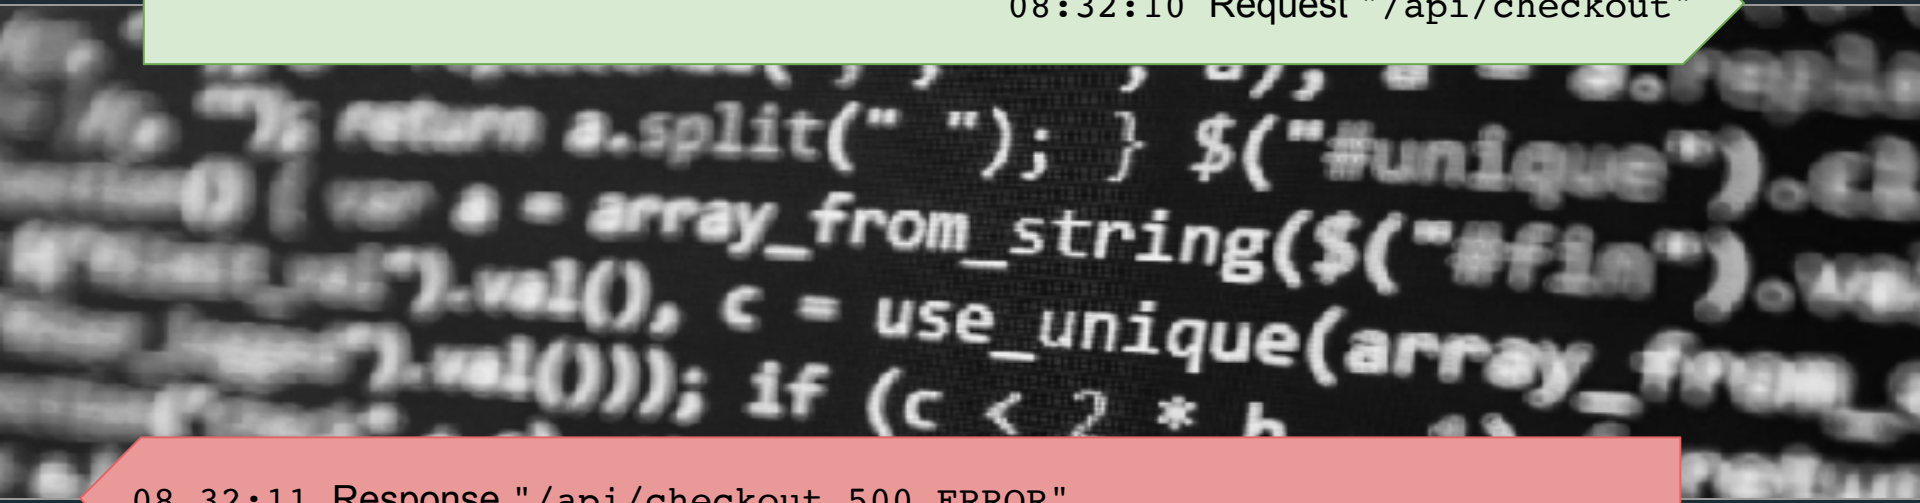

elastic

什么是 APM?

Example

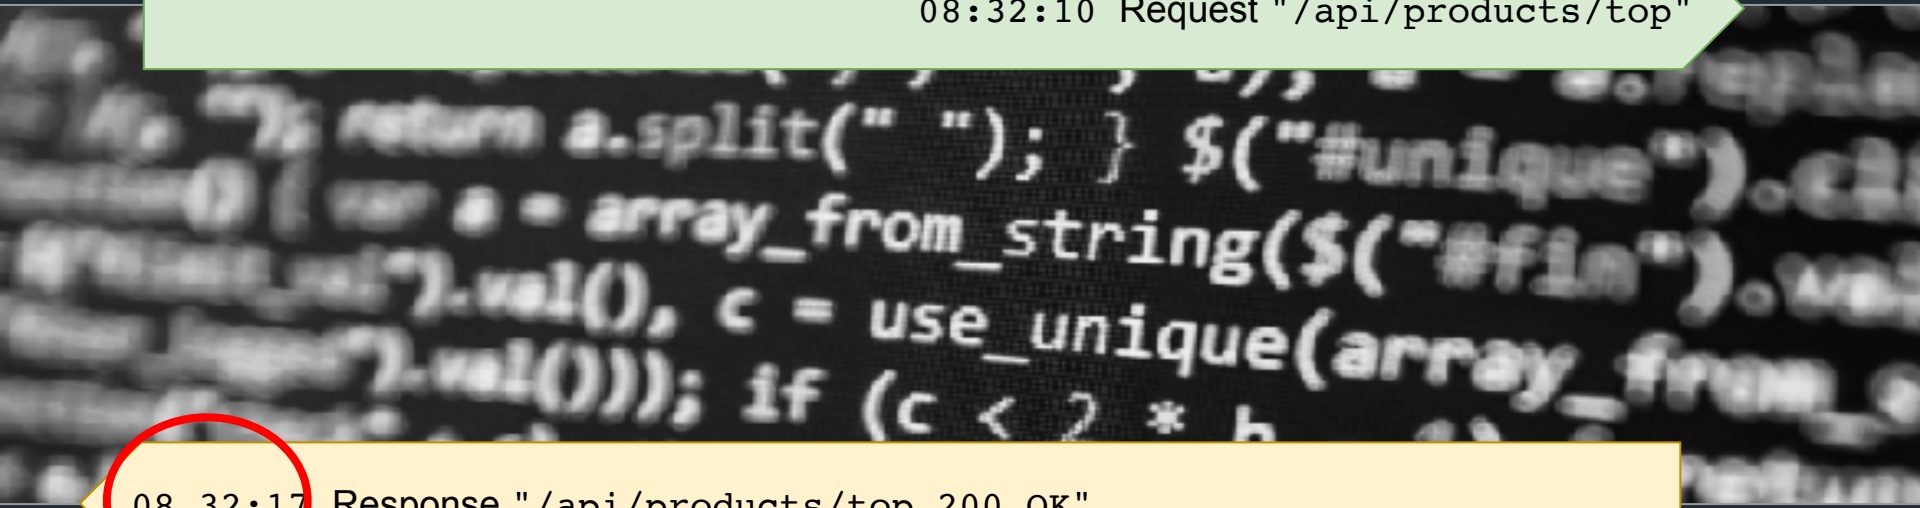

7 seconds - zZzzZZz

elastic

Elastic APM 如何工作?

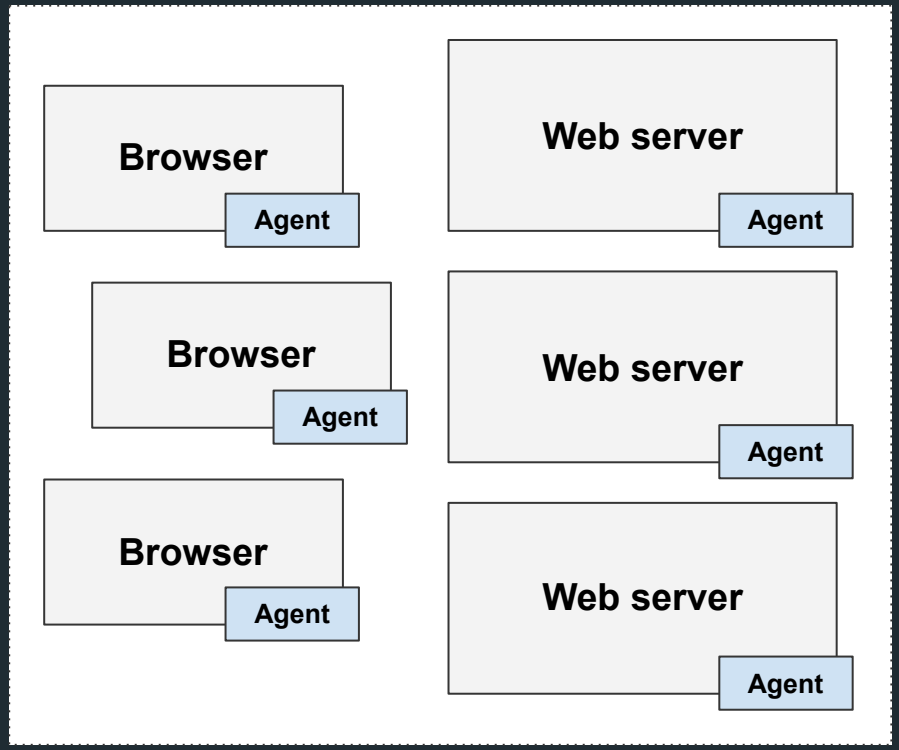

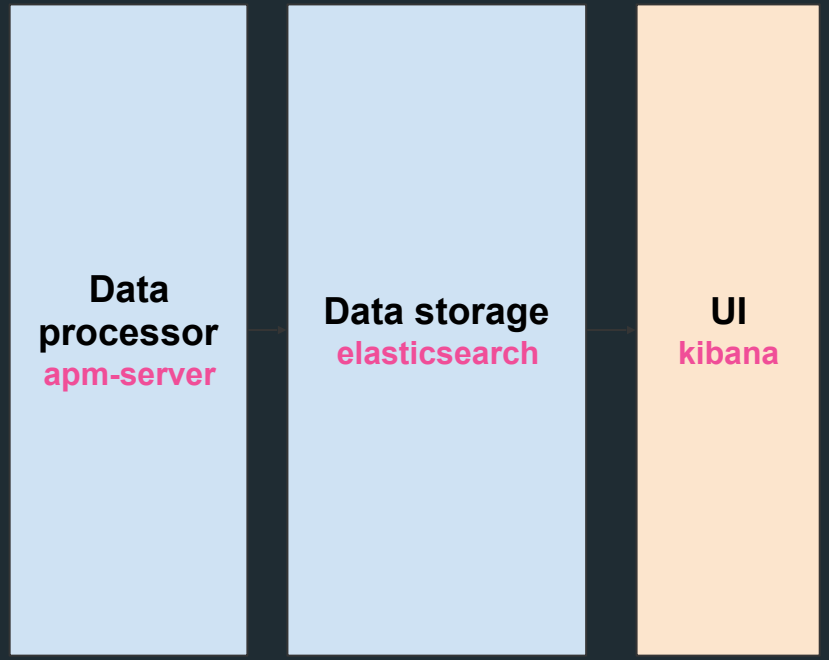

elastic

Elastic APM

APM adds end-user experience and application-level monitoring to the stack

- Focuses on search experience on top of APM data
- ‘Just another index’ in Elastic Stack

Language support

- Python
- Node.js
- Ruby
- RUM
- Java
- Go
- .NET (in dev)

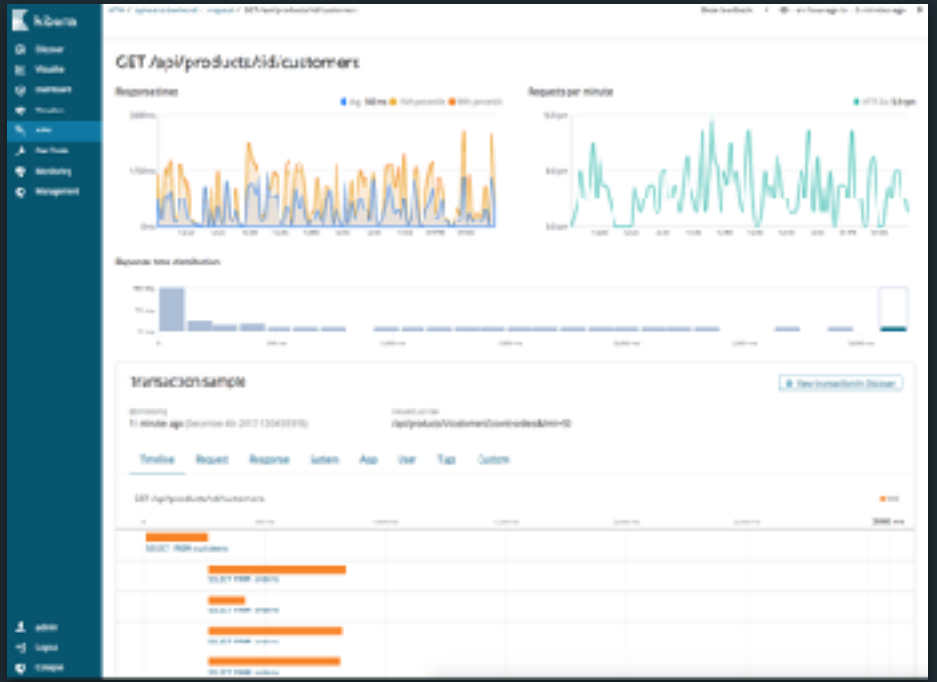

elastic

完善的 APM 分析 UI

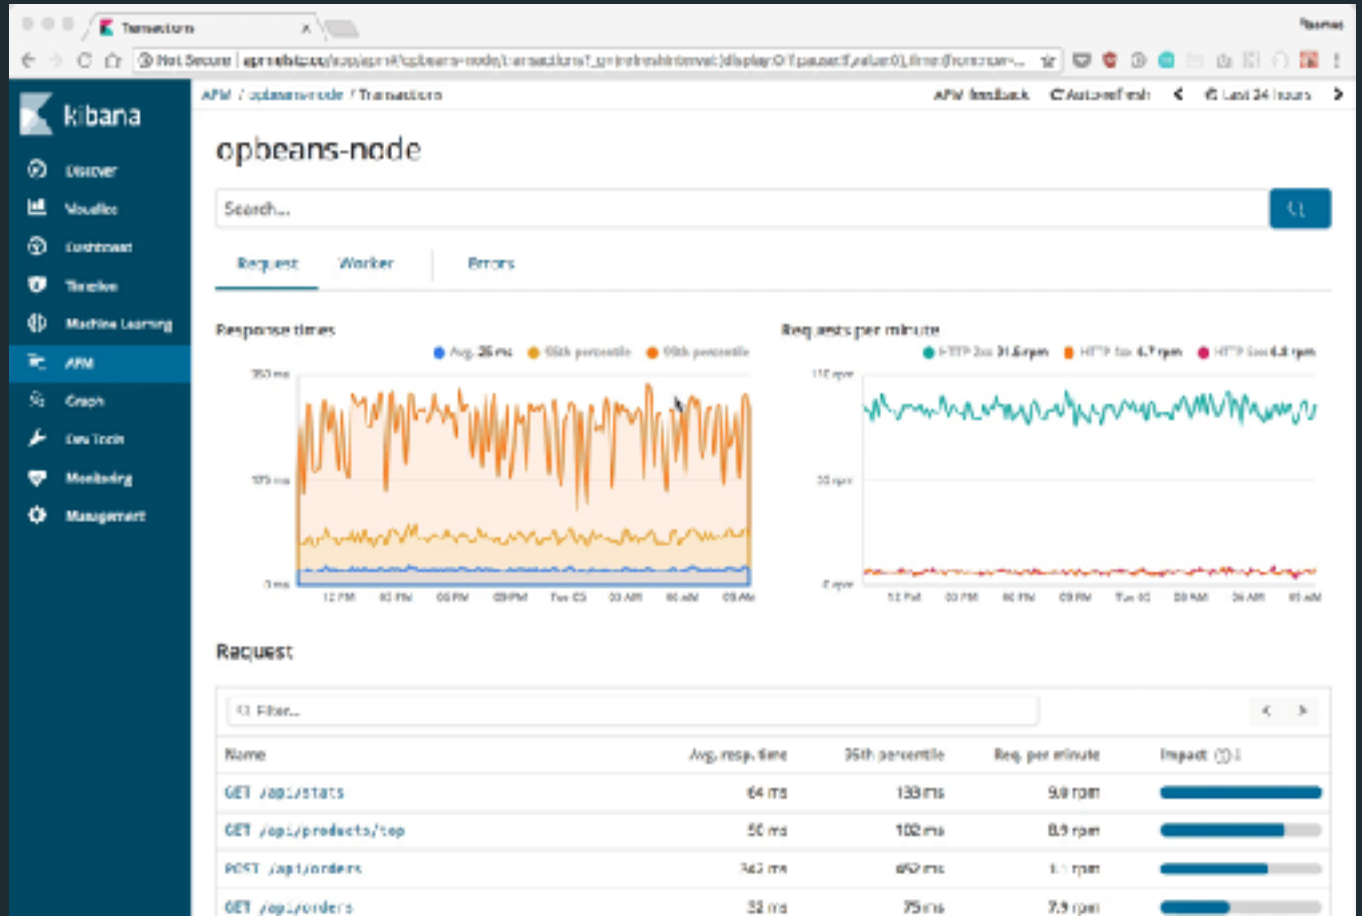

Combine custom workflow with freedom of search

elastic

Distributed Tracing

Single transaction

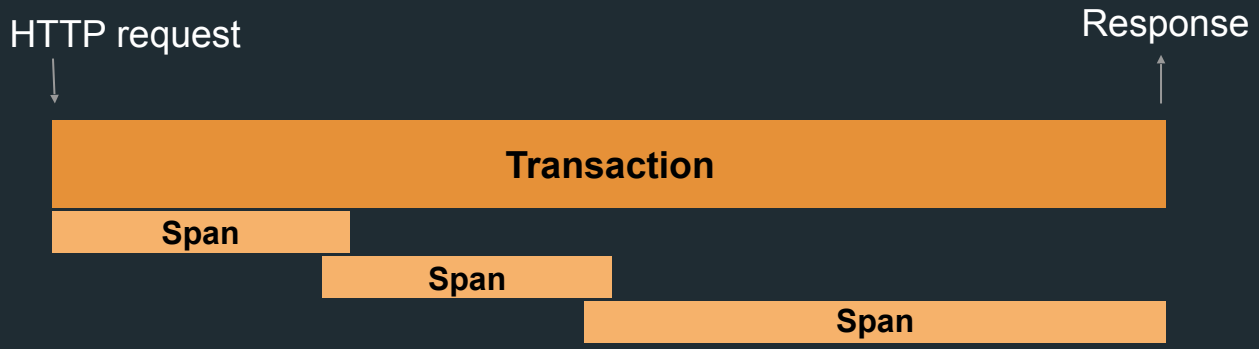

elastic

Distributed Tracing

Distributed tracing example

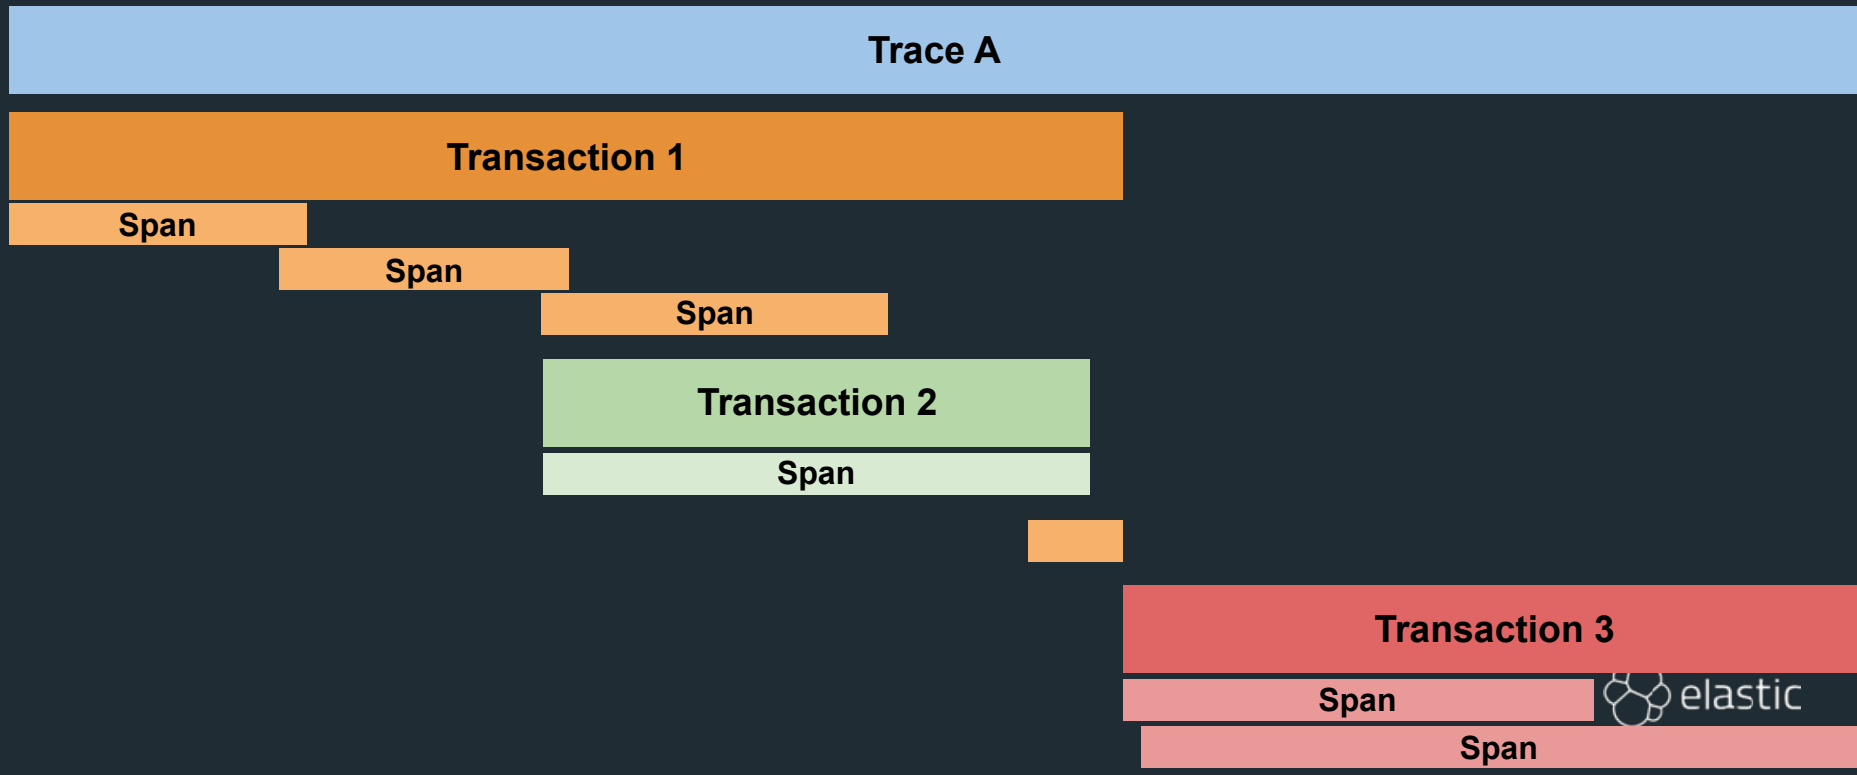

Distributed Tracing

Trace and map across multiple services

- 查看端到端视图并导航到各个事务
- 基于 Trace ID 的分布式跨服追踪
- 兼容 OpenTracing API 和 W3C trace context 规范

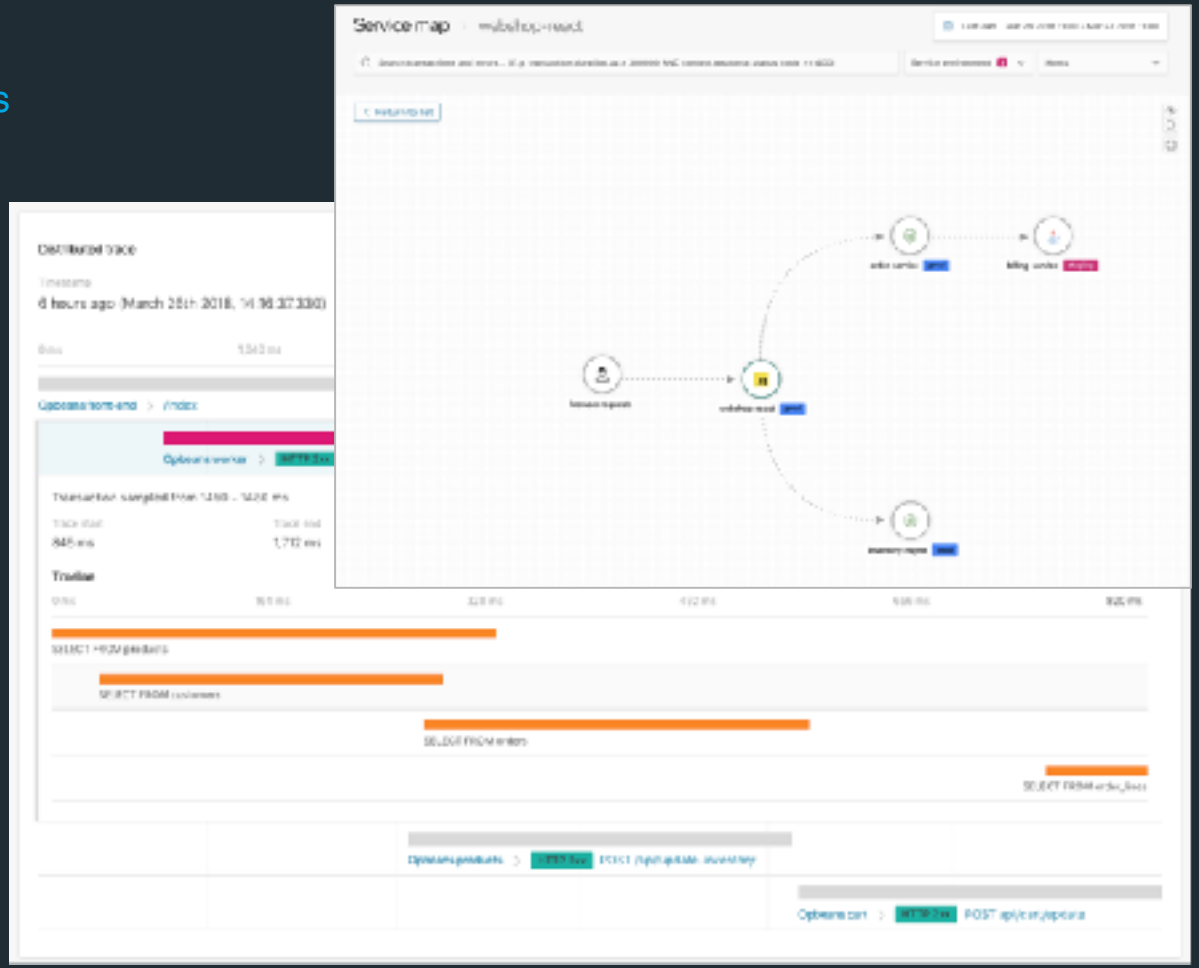

APM is another index in Elasticsearch

Need another visualization? Build a dashboard, no need to wait for your vendor

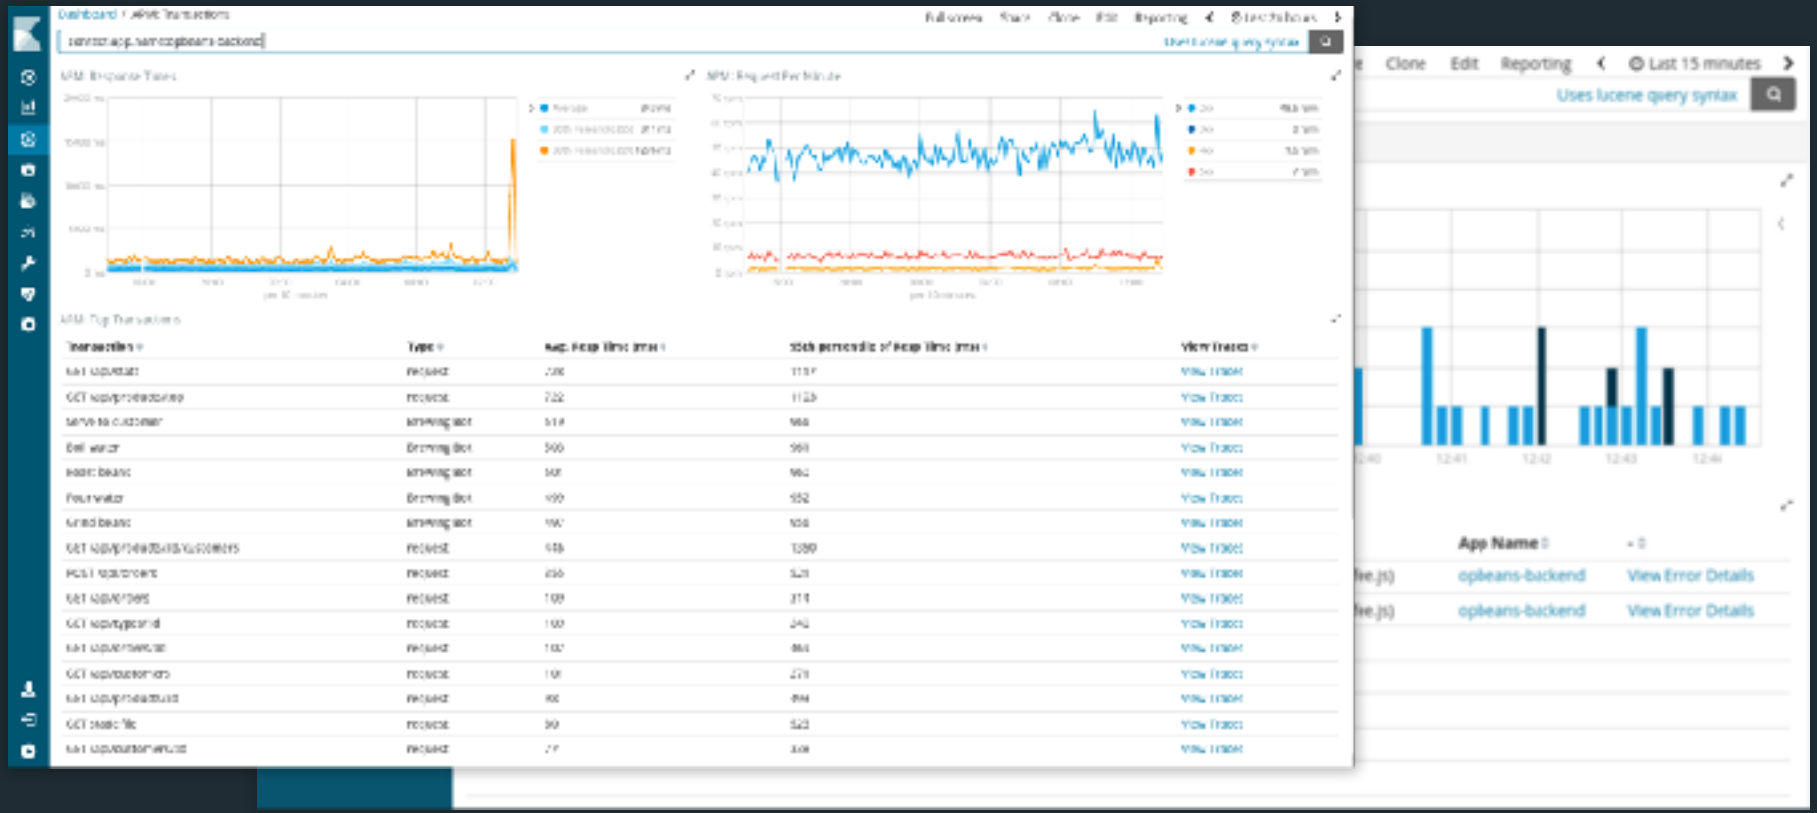

DEMO

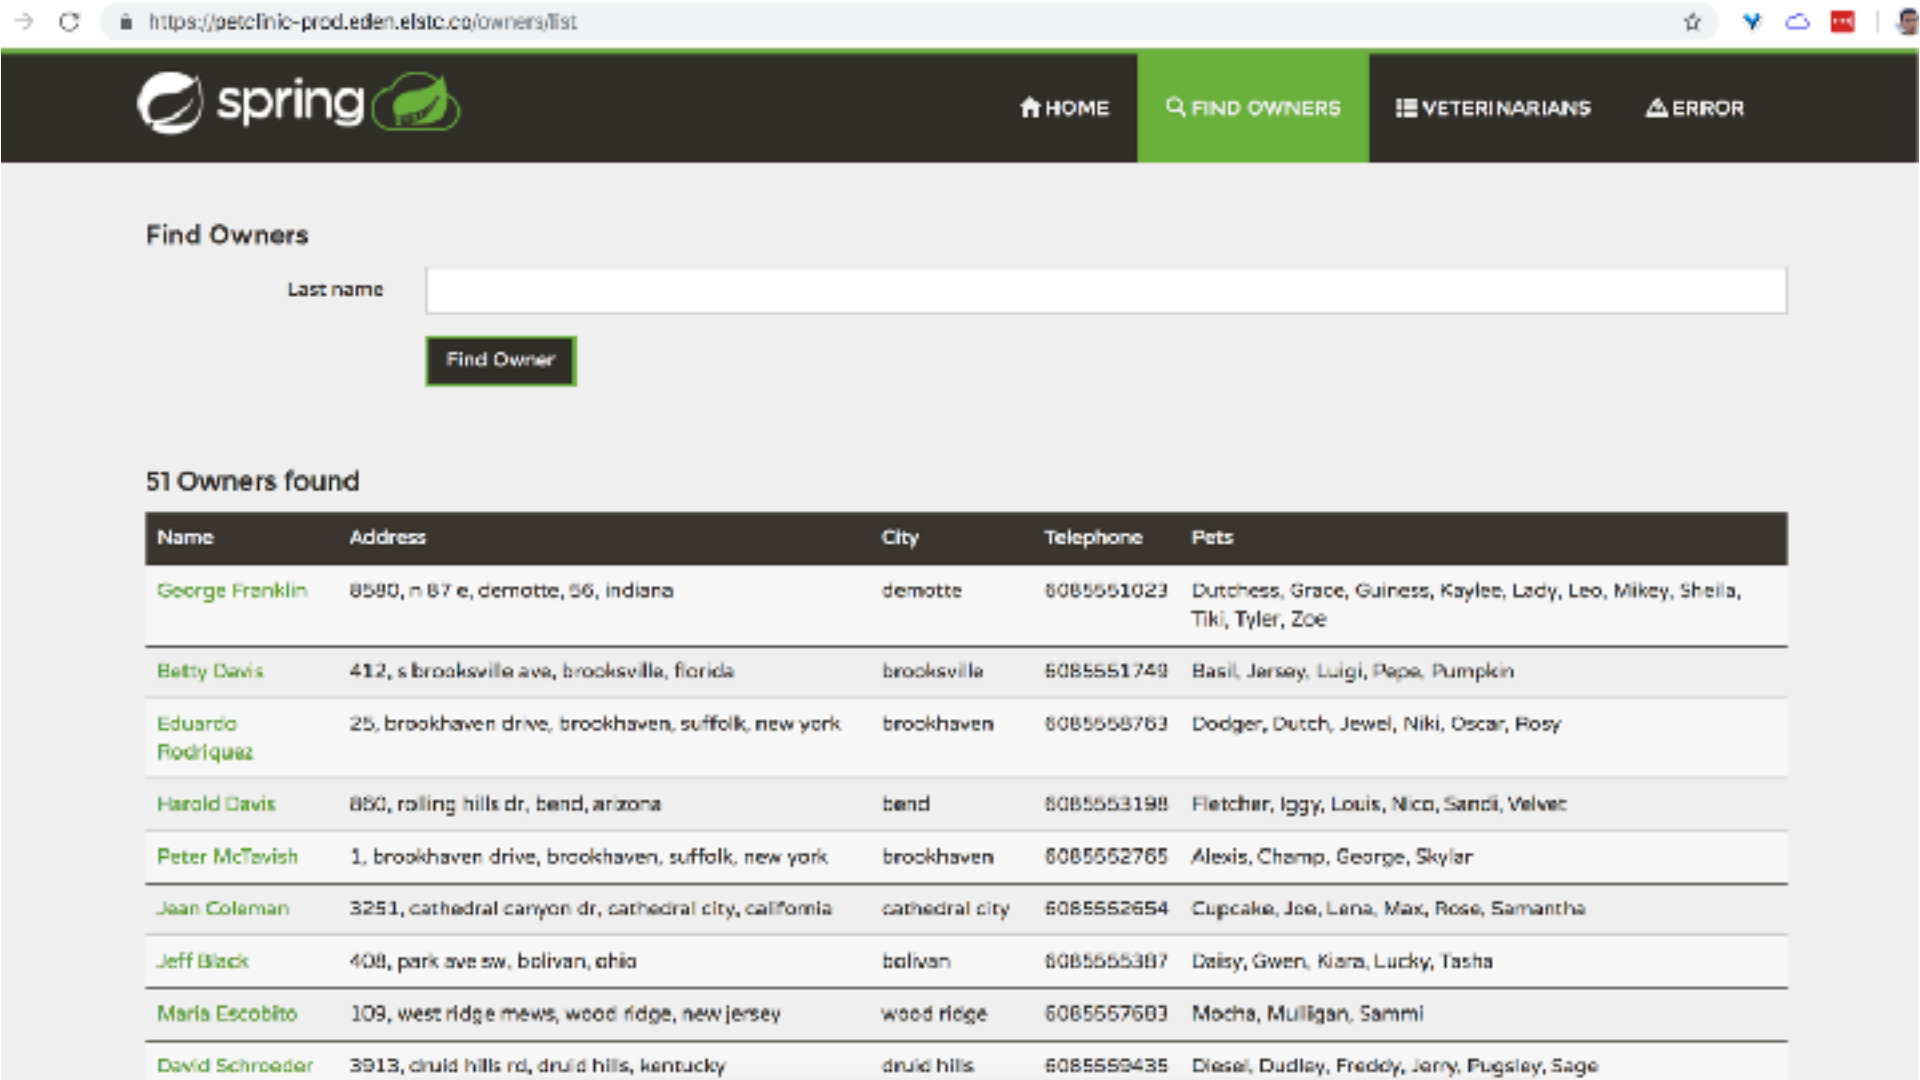
https://petclinic-prod.eden.elstc.co/owners/list
spring
HOME
FIND OWNERS
VETERINARIANS
ERROR
Find Owners
Last name
Find Owner
51 Owners found
Name
Address
City
Telephone
Pets
George Franklin
8580, n 87 e, demotte, 56, indiana
demotte
6085551023
Dutchess, Grace, Guiness, Kaylee, Lady, Leo, Mikey, Sheila, Tiki, Tyler, Zoe
Betty Davis
412, s brooksville ave, brooksville, florida
brooksville
6085551749
Basil, Jersey, Luigi, Pepe, Pumpkin
Eduardo Rodriquez
25, brookhaven drive, brookhaven, suffolk, new york
brookhaven
6085558763
Dodger, Dutch, Jewel, Niki, Oscar, Rosy
Harold Davis
850, rolling hills dr, bend, arizona
bend
6085553198
Fletcher, Iggy, Louis, Nico, Sandi, Velvet
Peter McTavish
1, brookhaven drive, brookhaven, suffolk, new york
brookhaven
6085552765
Alexis, Champ, George, Skylar
Jean Coleman
3251, cathedral canyon dr, cathedral city, california
cathedral city
6085552654
Cupcake, Joe, Lena, Max, Rose, Samantha
Jeff Black
408, park ave sw, bolivan, ohio
bolivan
6085555387
Daisy, Gwen, Kiara, Lucky, Tasha
Maria Escobito
109, west ridge mews, wood ridge, new jersey
wood ridge
6085557603
Mocha, Mulligan, Sammi
David Schroeder
3913, druid hills rd, druid hills, kentucky
druid hills
6085559435
Diesel, Dudley, Freddy, Jerry, Pugsley, Sage

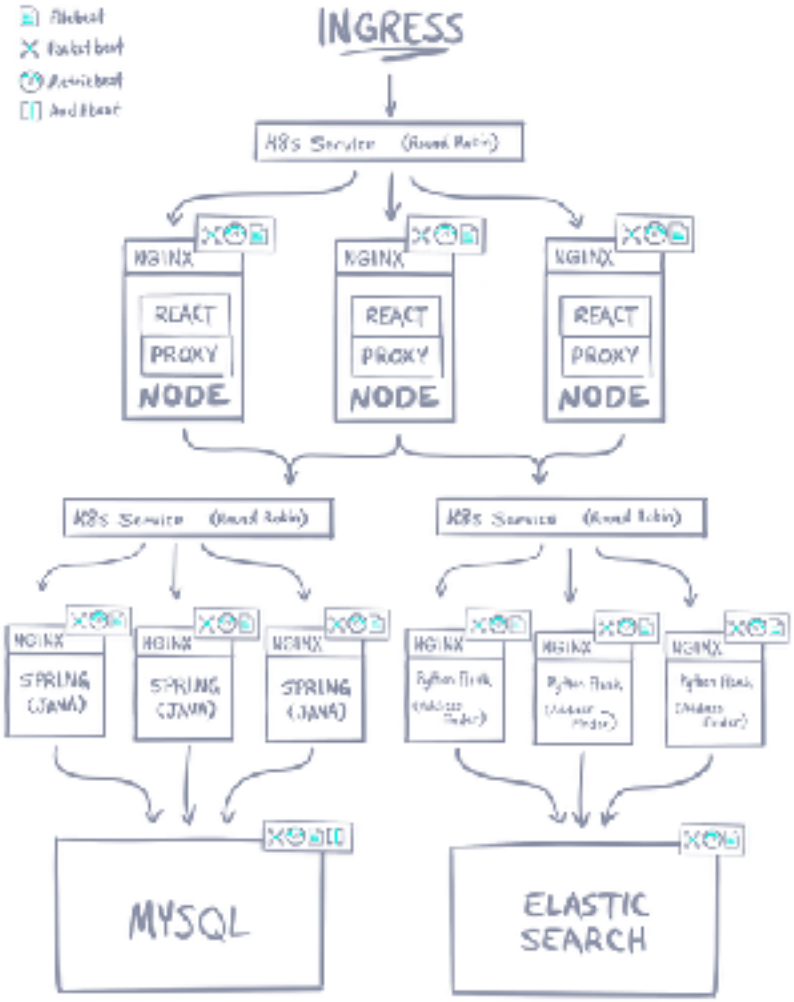
INGRESS
NGINX
REACT
PROXY
NODE
NGINX
REACT
PROXY
NODE
NGINX
REACT
PROXY
NODE
NGINX
SPRING
(JAVA)
NGINX
SPRING
(JAVA)
NGINX
SPRING
(JAVA)
NGINX
NGINX
NGINX
MYSQL
ELASTIC
SEARCH

怎样?

要不要下载试试?

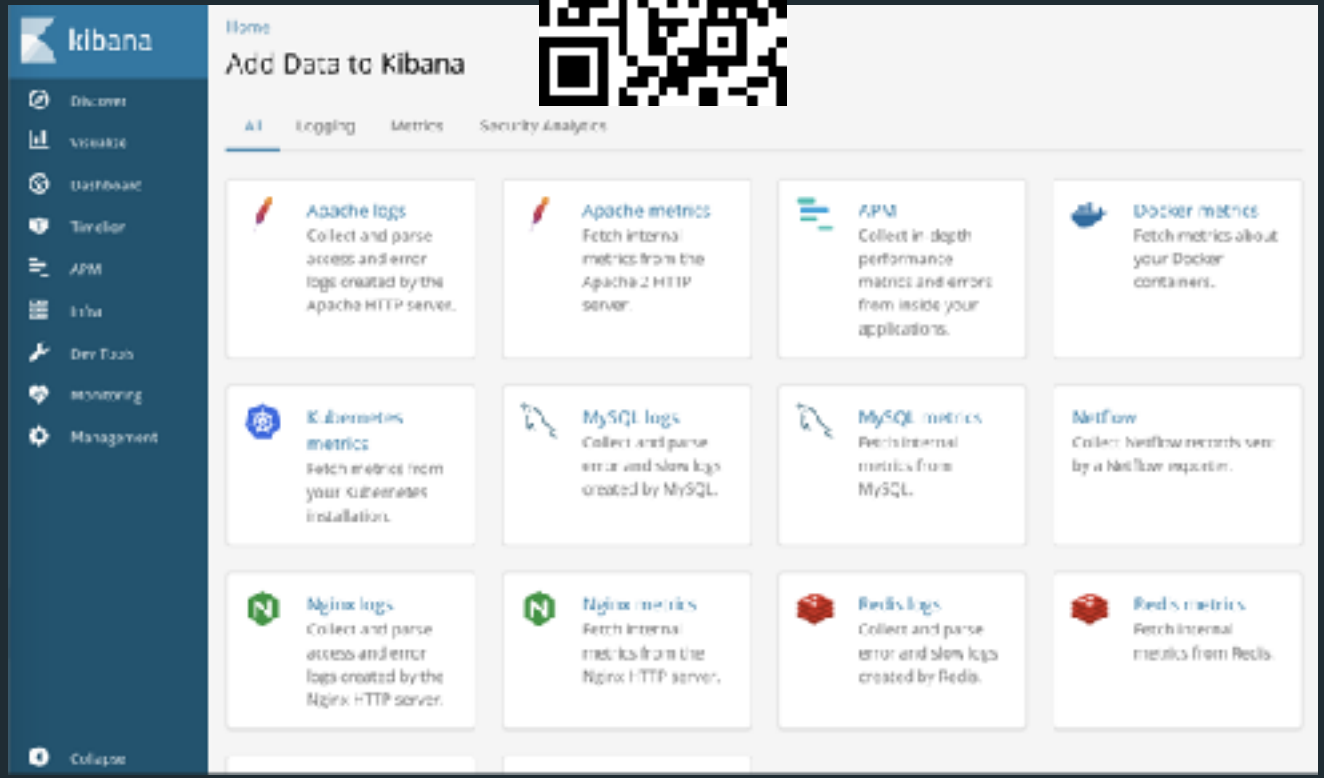

elastic

Questions?

更多 Demo: http://demo.elastic.co
官方网站：http://elastic.co/cn