使用Elastic Stack快速搭建SIEM体系

21st/Jul/2018

吴斌 - Elastic 架构师

什么是SIEM?

- 由Gartner的Mark Nicolett和Amrit Williams在2005年提出了Security Information Event Management (SIEM) 的概念

- Describes the product capabilities of gathering, analyzing and presenting information from network and security devices; identity and access management applications; vulnerability management and policy compliance tools; operating system, database and application logs; and external threat data.

核心目标

- 识别
  - 潜在的威胁
  - 已经造成的破坏
- 审计
  - 收集安全以及合规日志
- 分析
  - 调查
  - 关联性

安全事件工作流

NIST Cybersecurity Framework
https://www.nist.gov/cyberframework

我们致力于帮助客户解决所有实时数据处理问题

- **100,000+** 社区参与者
- **225M+** 下载 (持续增长中)
- 政府, 金融保险, 电信, 高科技
- **900+** 员工
- **5,000+** 订阅客户

NETFLIX hp verizon UBER COMCAST CISCO SOCIETE GENERALE SwissLife ZURICH
Scottrade HSBC E*TRADE AXA Microsoft salesforce ebay Audi GROUPON elastic

Elastic 产品选择

elastic

Elastic Stack

100% 开源
无商业/企业版本
全新6.x版本

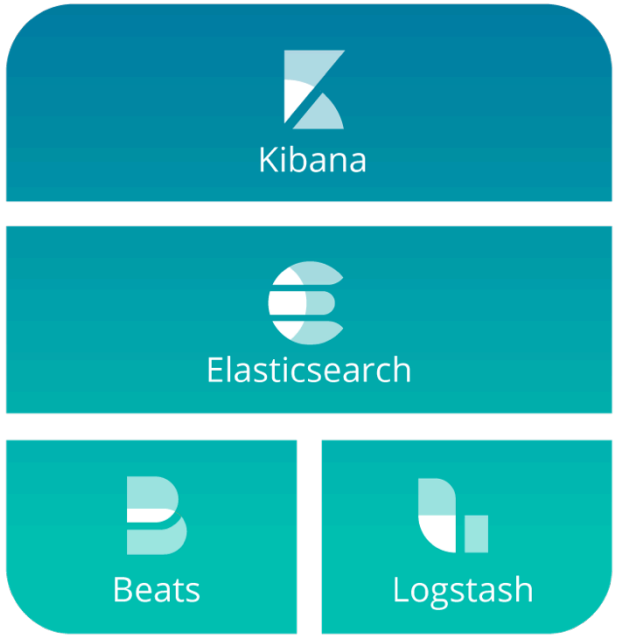

elastic

ELK to ELKB (Elastic Stack)

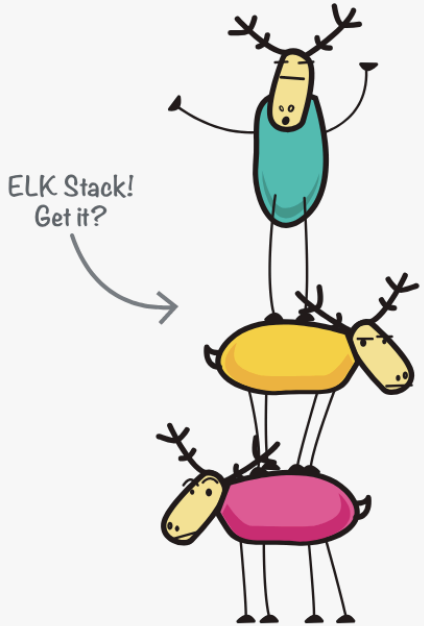

E Elasticsearch

L Logstash

K Kibana

elastic

X-Pack

6.3+ 默认整合
Elastic Stack的企业级扩展

数据安全

告警

监控

报表

图查询

机器学习

elastic

Beats: 轻量级数据搬运工

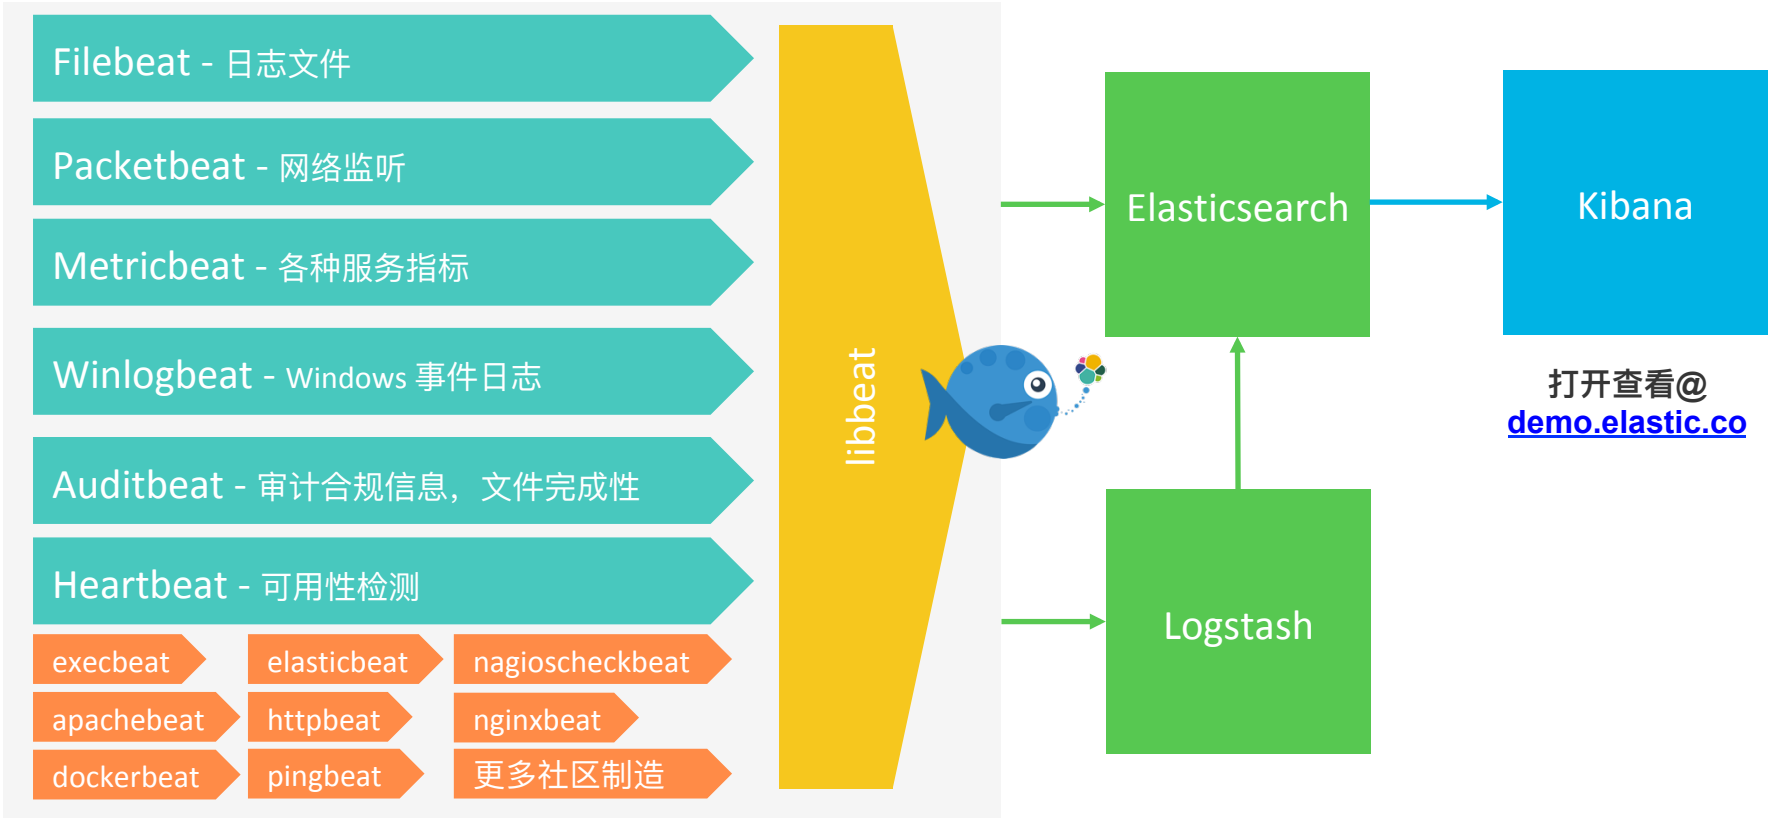

elastic

Beats 预置模块

简化收集、解析和可视化常见日志格式

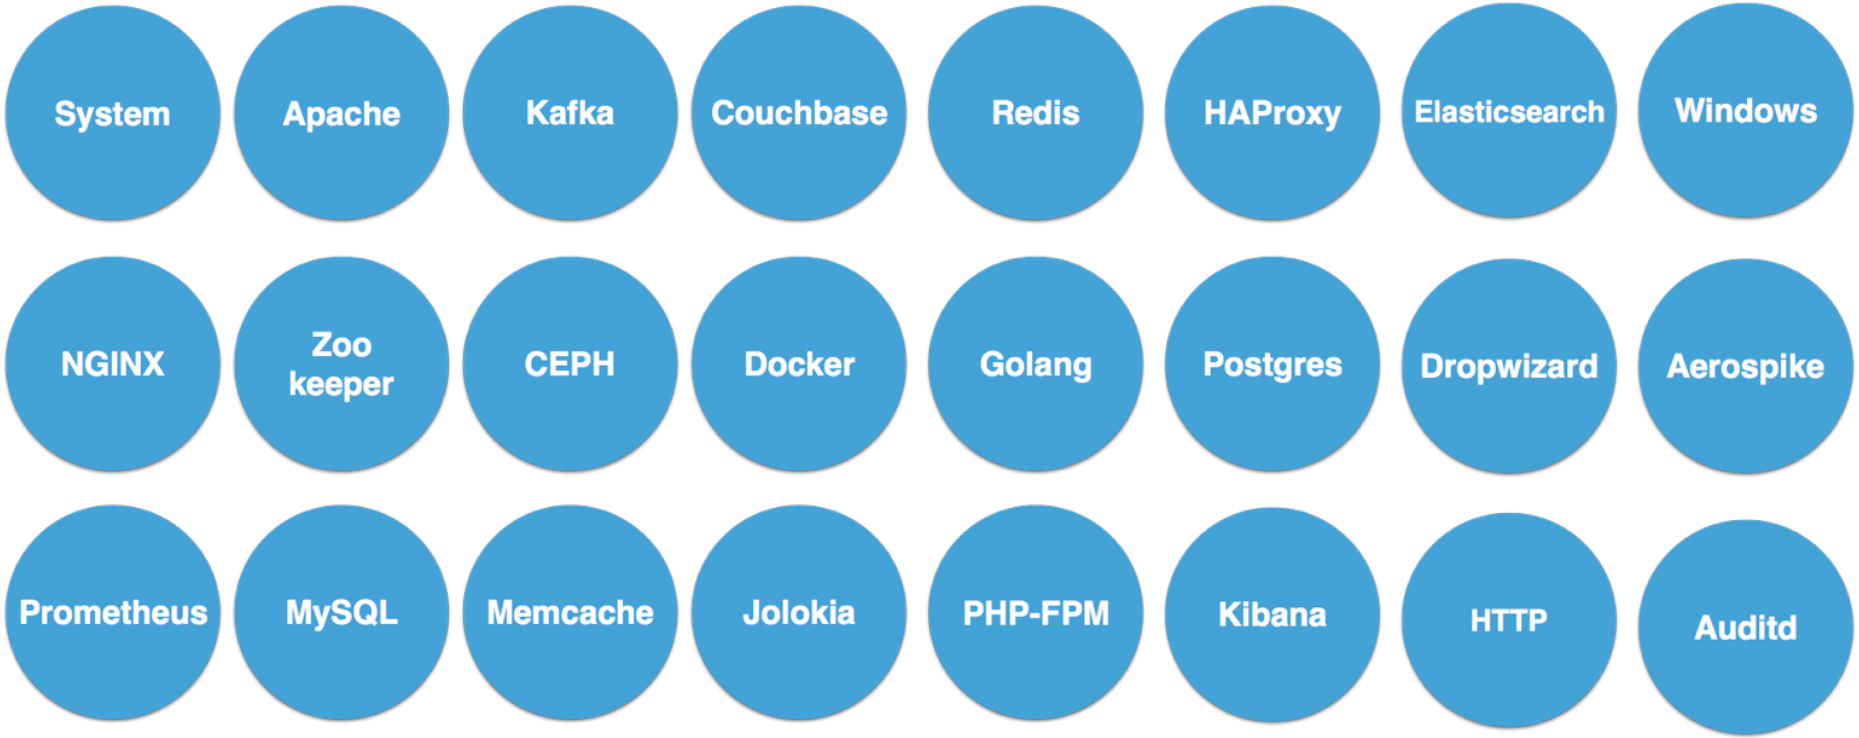

elastic

Beats预置面板

elastic

X-Pack

机器学习

时序型数据的非监督、增量学习

- 异常检测自动化
- 加速归因分析

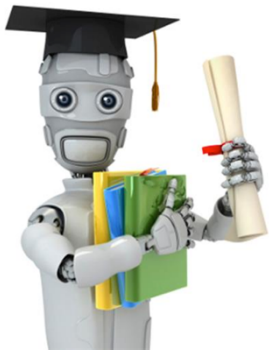

Elastic

X-Pack

elastic

Elastic Stack满足SIEM场景一些基本诉求

安全分析 | 日志分析 | 指标分析 | 商业分析 | 搜索

保护你的数据

数据告警

异常检测

监测管理
Elastic Stack

找到数据关联

共享你的报表

elastic

保护你的 Elastic Stack

elastic

场景

限制数据访问

Elastic 认证:
- 原生认证（用户名/密码）
- LDAP, Active Directory
- 定制 (like Kerberos, SSO)
- SAML 6.2

Elastic 授权:
- 索引, 文档, 字段
- APIs

IP 过滤规则

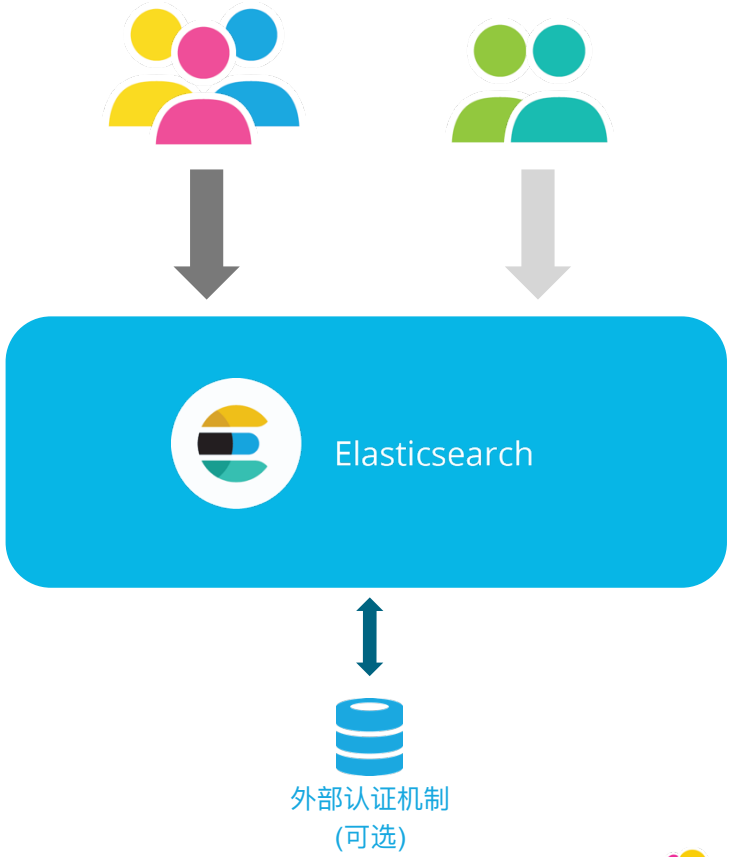

elastic

数据加密

通讯:

- 和Elasticsearch集群的交互
- Elasticsearch节点间通讯

磁盘数据:

- 支持 *dm-crypt* 加密

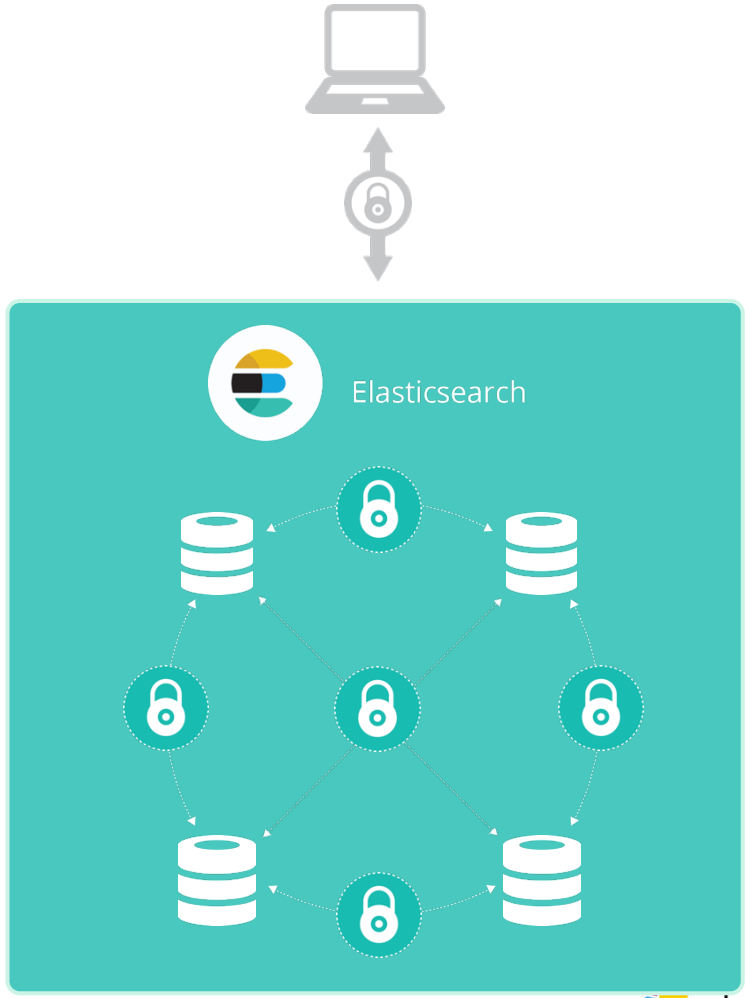

审计/合规

使用安全审计功能:

- 洞悉正常和非常规行为
- 对非常规行为进行告警
- 构建真对访问控制的面板

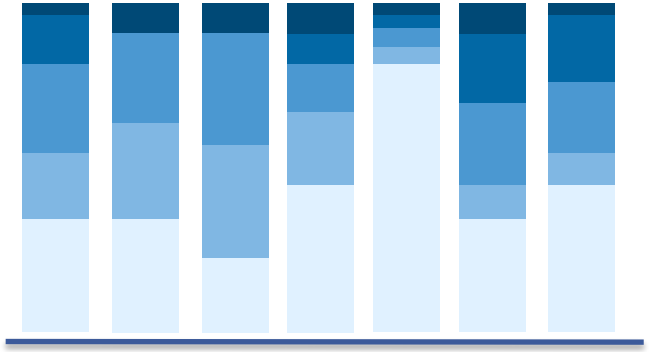

Access vs. time by users

elastic

对异常数据/行为进行告警

elastic

IT 运营分析

使用规则

在下列情况发生时通知:

- 磁盘剩余空间不足 5%
- Elasticsearch 集群健康状态红色预警

使用异常检测

在下列情况发生时通知:

- 突发的高错误率
- 未知的高CPU使用率

安全分析

使用规则

在下列情况发生时通知:

- 过去5分钟内单个机器上超过5次登录失败
- 进程 X 在任意服务器上启动

使用异常检测

在下列情况发生时通知:

- 在非常规的时间/地点登陆
- 向系统外可疑的数据传输

追踪记录所有告警

所有告警记录可追溯:

- SLAs是否达标?
- 安全事件的趋势?
- 那台服务器故障问题最多?
- 哪些事件与异常数据存在相关性?

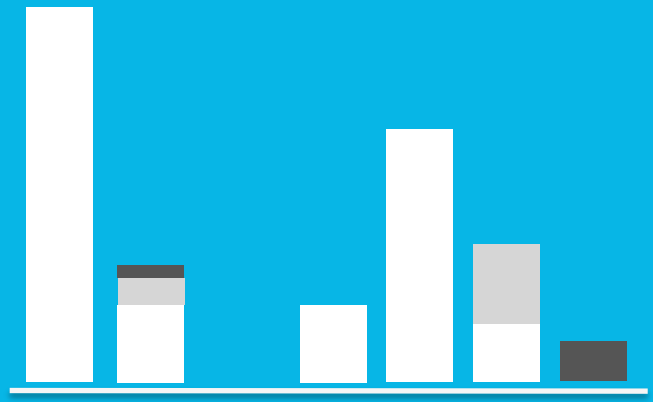

内存空间告警

-- 服务器 1 -- 服务器 2 -- 服务器 3

机器学习

Image Classification **Recommendations**

Autonomous cars Voice Recognition **Predictive Medicine**

Fraud detection **异常检测**

Learn to Rank Speech Recognition

Language Translation **Entity Resolution**

elastic

异常检测是规则引擎很好的补充

- 规则并不是定义正常/异常最好的方法
- 定义再好的规则，也不会随着数据改变而持续演化、改进

如何设置伐值?

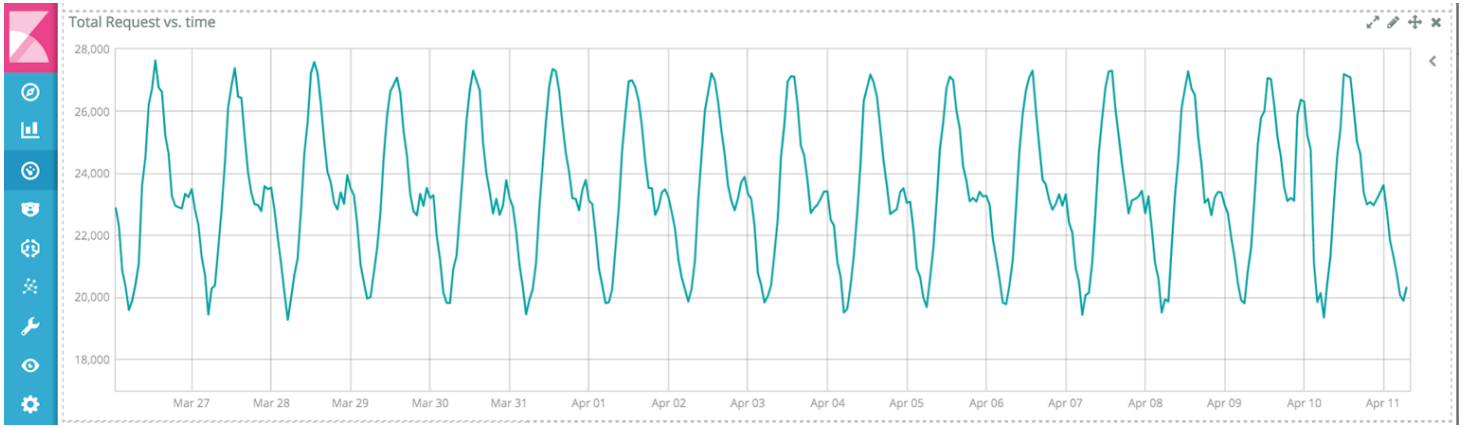

elastic

非监督机器学习

全自动察觉数据正常范围，随数据变化而持续改进、更新

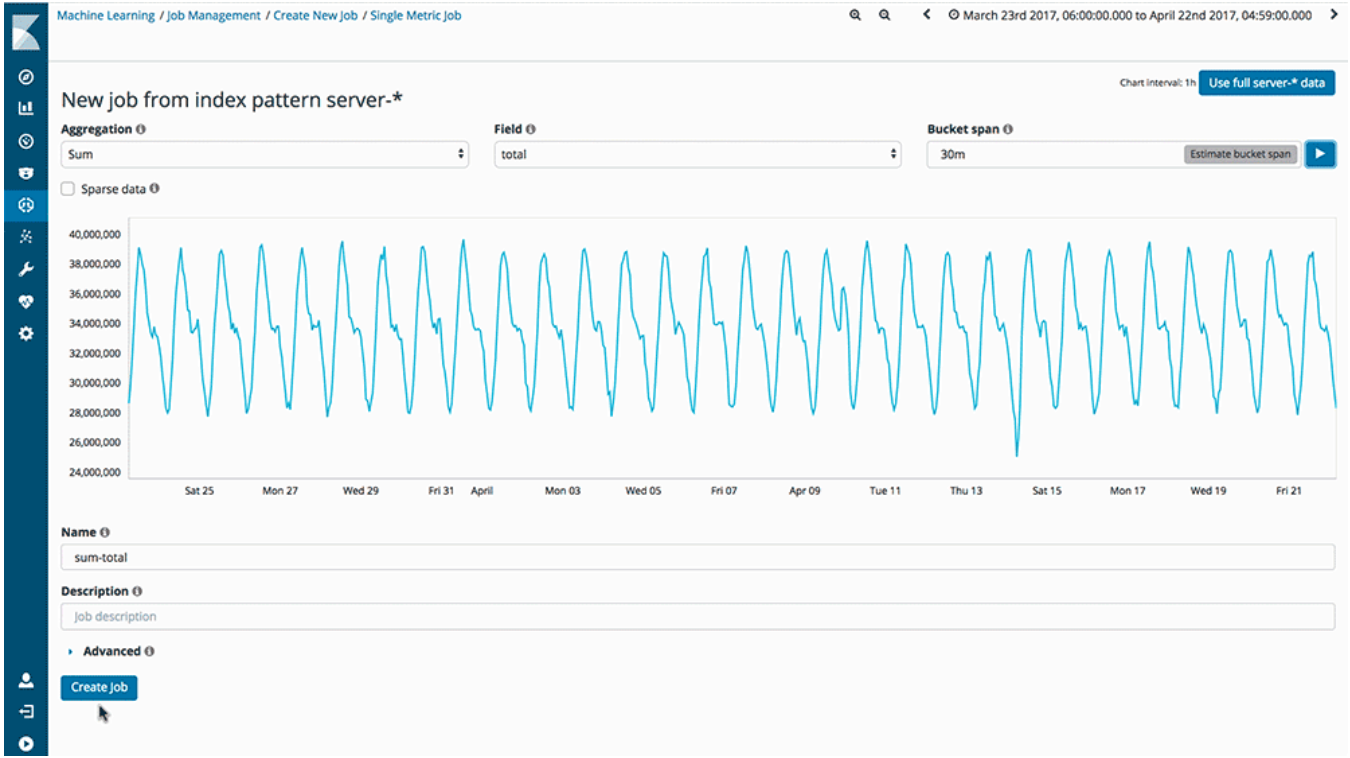

elastic

加速归因分析，锁定影响因子

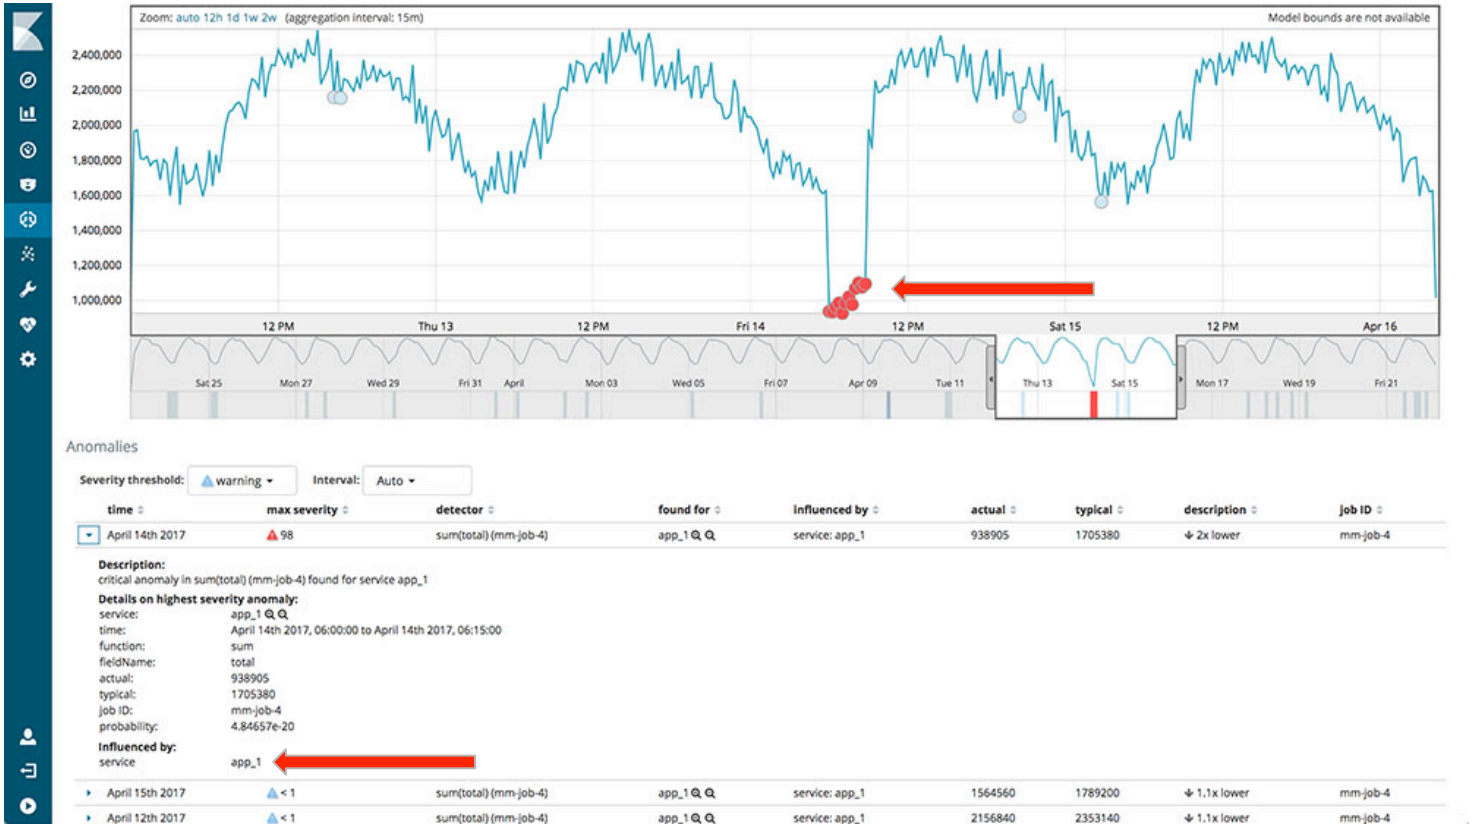

elastic

预测、洞悉未来

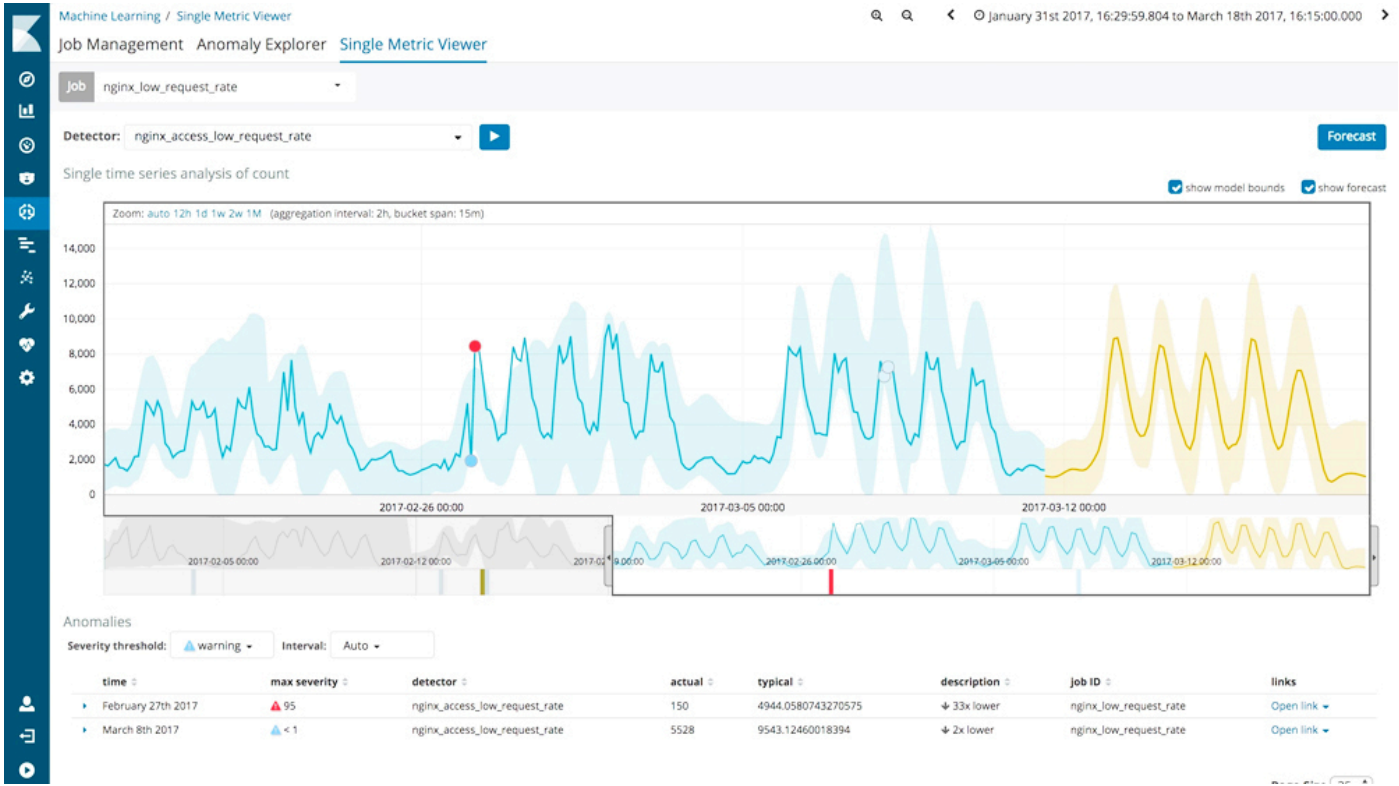

elastic

监控 Elastic Stack

elastic

实时监控

- 集群整体的工作状态、压力
- 是否有压力暴增的单个节点?
- 在灾难发生前洞察问题并着手解决

统一的监控入口，掌控集群整体状态

elastic

监控数据也会被索引到 Elasticsearch

数据可以留作日后对集群状态的分析和调优结果评估

Table JSON


```
{
  "_index": ".marvel-es-1-2016.06.06",
  "_type": "node_stats",
  "_id": "AVUnC_T4qQKi2Snrly9N",
  "_score": null,
  "_source": {
    "cluster_uuid": "ReVr5CXFTi2fucBt2LbLNA",
    "timestamp": "2016-06-06T18:49:02.968Z",
    "source_node": {
      "uuid": "OxWM08tpSn-ijbx-kY17VA",
      "host": "10.137.126.111",
      "transport_address": "10.137.126.111:19168",
      "ip": "10.137.126.111",
      "name": "tiebreaker-0000000060",
      "attributes": {
        "logical_availability_zone": "tiebreaker",
        "availability_zone": "us-east-1b",
        "data": "false",
        "max_local_storage_nodes": "1",
        "region": "us-east-1",
        "master": "true"
      }
    },
    "node_stats": {
      "node_id": "OxWM08tpSn-ijbx-kY17VA",
      "node_master": false,
      "mlockall": false,
      "disk_threshold_enabled": false,
      "disk_threshold_watermark_high": 90,
      "indices": {
```


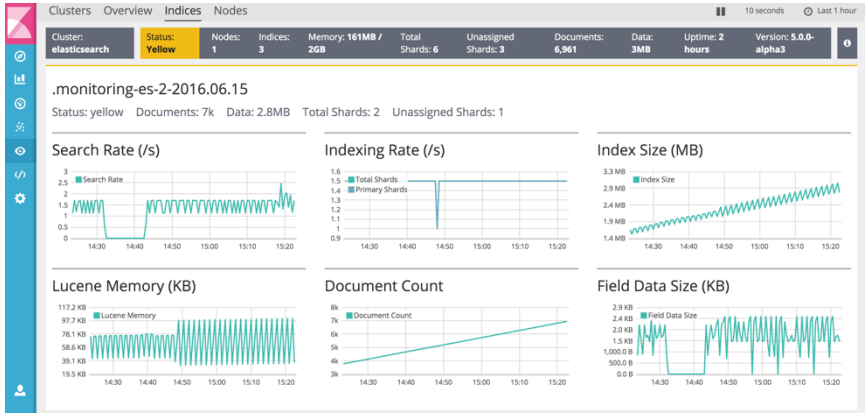

JVM 堆内存使用过高*!*

elastic

找到数据之间的关联

elastic

行为分析

一个可疑的IP地址还访问了哪些endpoint?
还有哪些IP跟他可能有潜在关联?

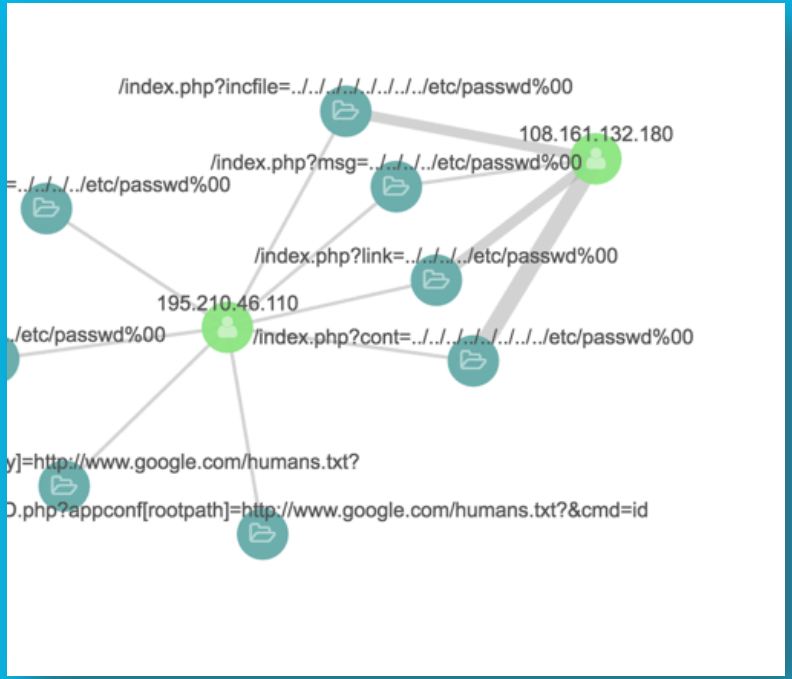

洞悉数据的共享

elastic

按需或者按计划生成报表

- 把面板导出为 PDF 文档
- 给 Elastic Stack 用户共享
- 提供离线数据访问

- 定时发送报表到指定邮箱
  –节省重复工作时间

生成报表来共享面板数据

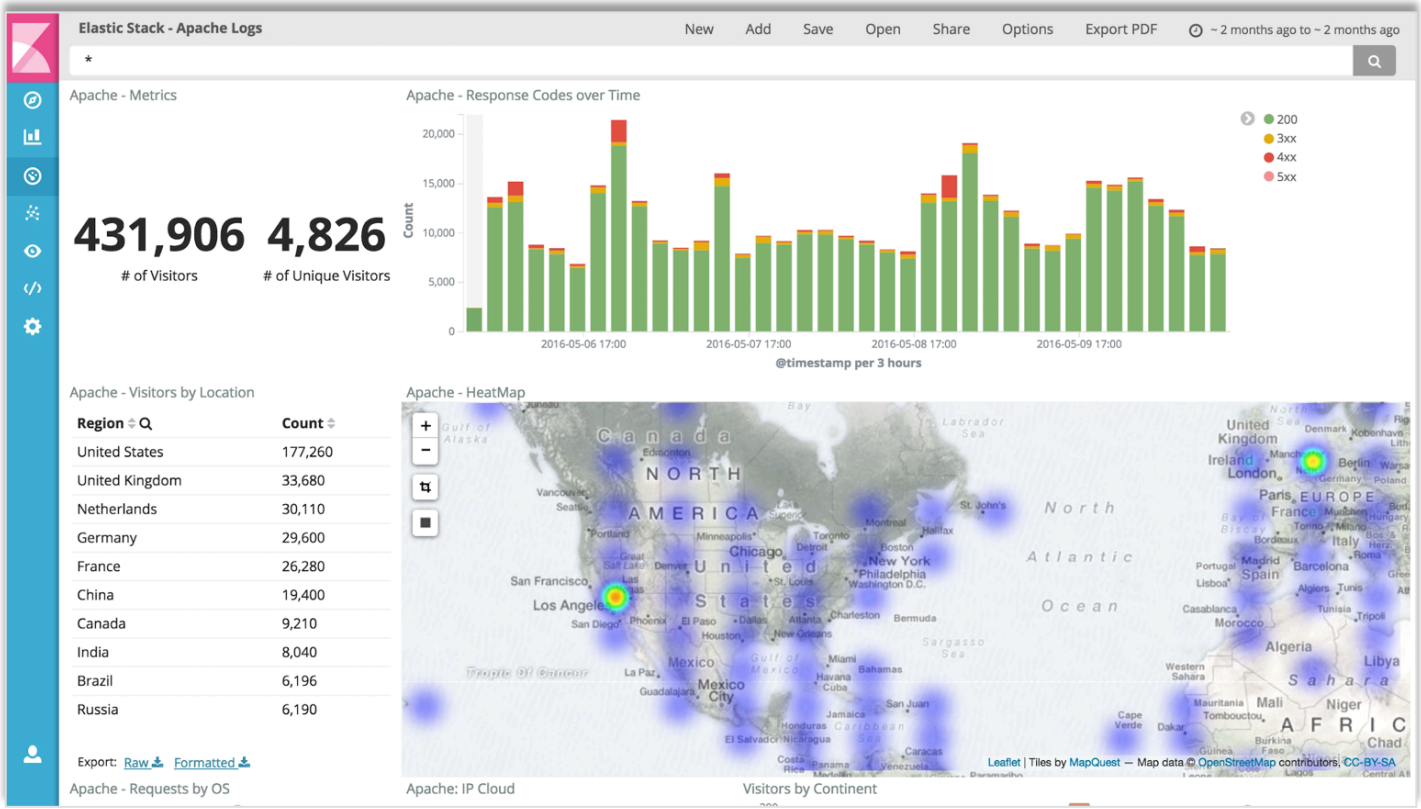

x-pack

在SIEM应用中的价值

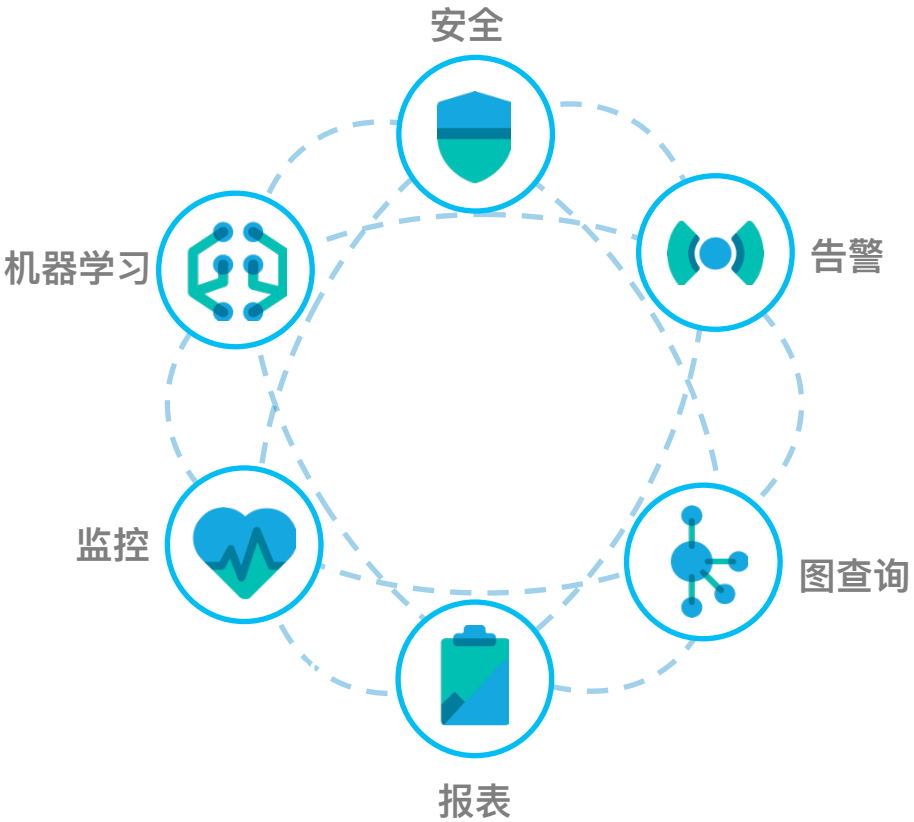

elastic

基于Elastic Stack的SIEM

部署快速上手

elastic

数据处理管道逻辑图

Beats

- 轻量级数据搬运工
- 文件、系统指标、网络数据、审计事件
- 可扩展框架
- 丰富的预置模块

Logstash

- 过滤、丰富数据
- 标准化数据格式，方便统一分析
- 集中的配置和管理
- 持久化队列

Elasticsearch

- 高性能搜索和分析
- 高可用、容错性
- 专用节点角色分担相应负载
- REST API

X-Pack

Kibana

- 探索数据
- 丰富的可视化工具
- 开发和管理工具

X-Pack

Machine Learning

Alerting

Graph

elastic

快速开始

Beats, Elasticsearch and Kibana

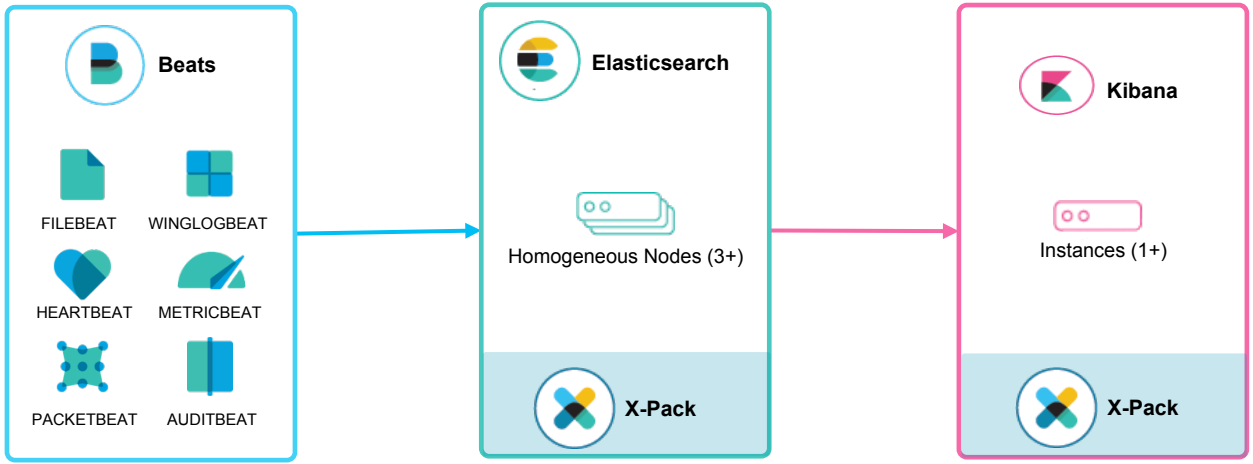

elastic

更加完善和容错更好的数据管道

添加了Logstash和按照角色分配Elasticsearch的任务

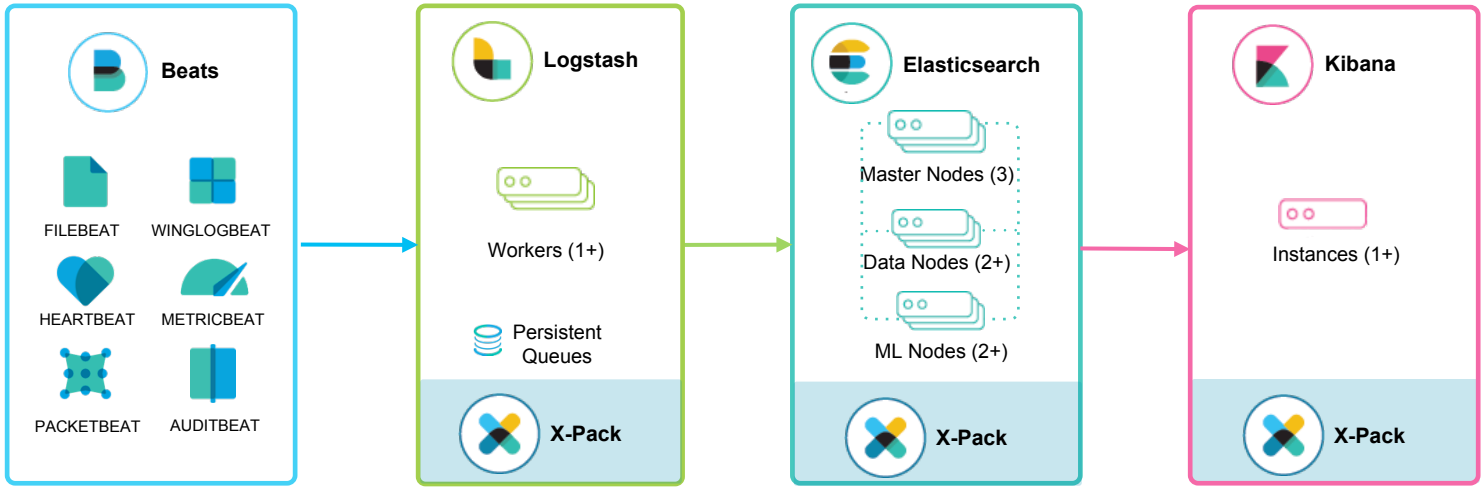

elastic

企业级部署

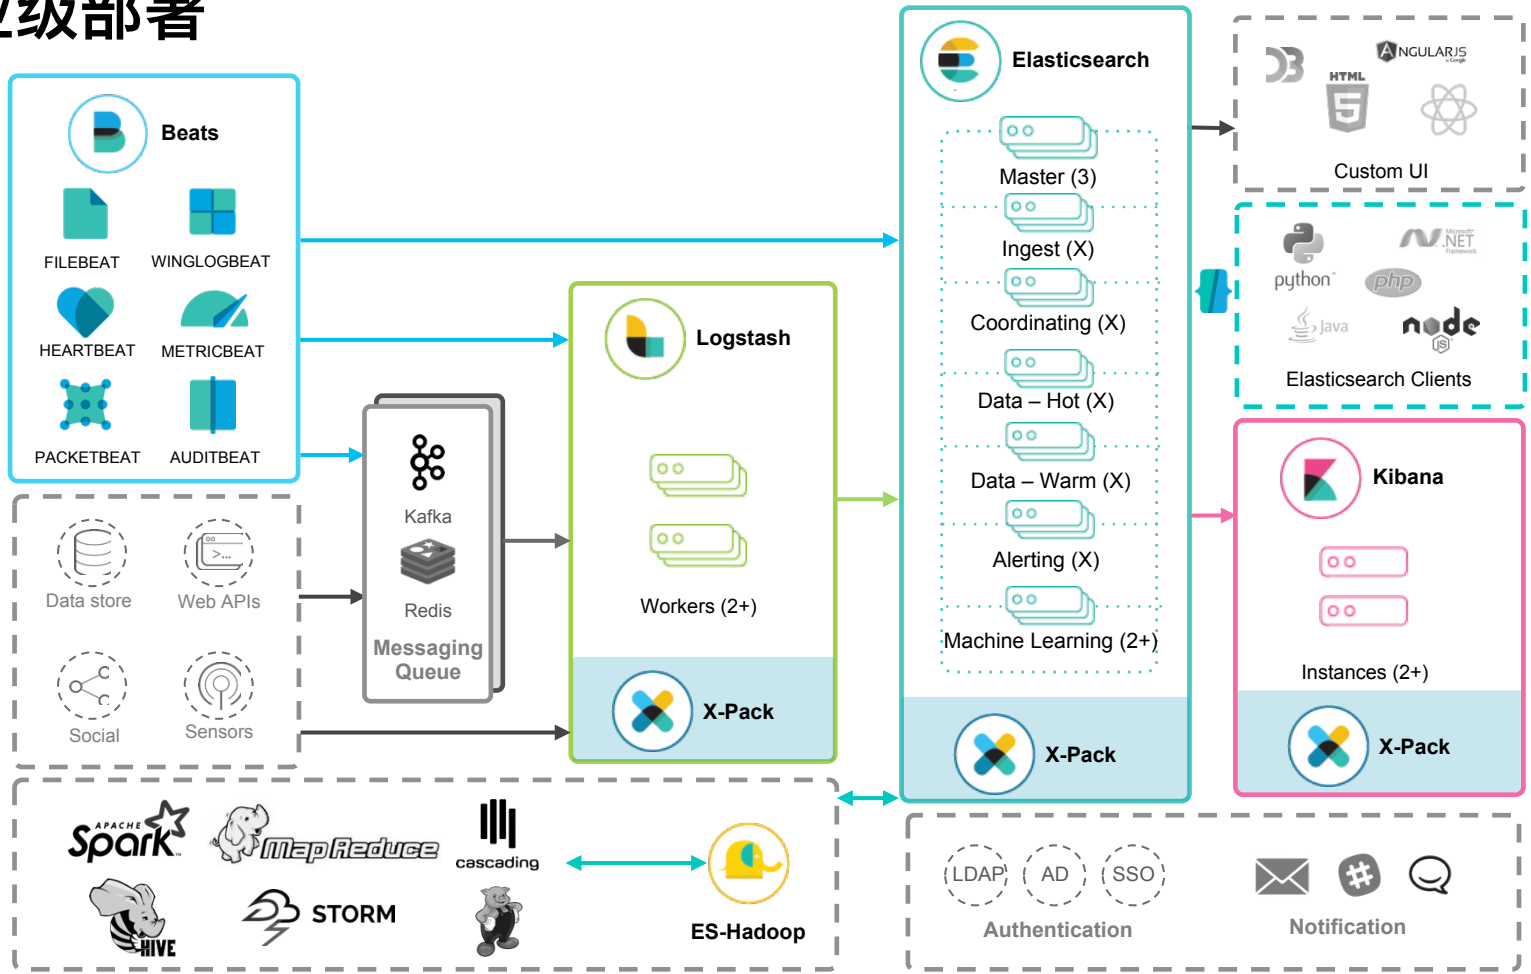

跨数据中心冗余数据

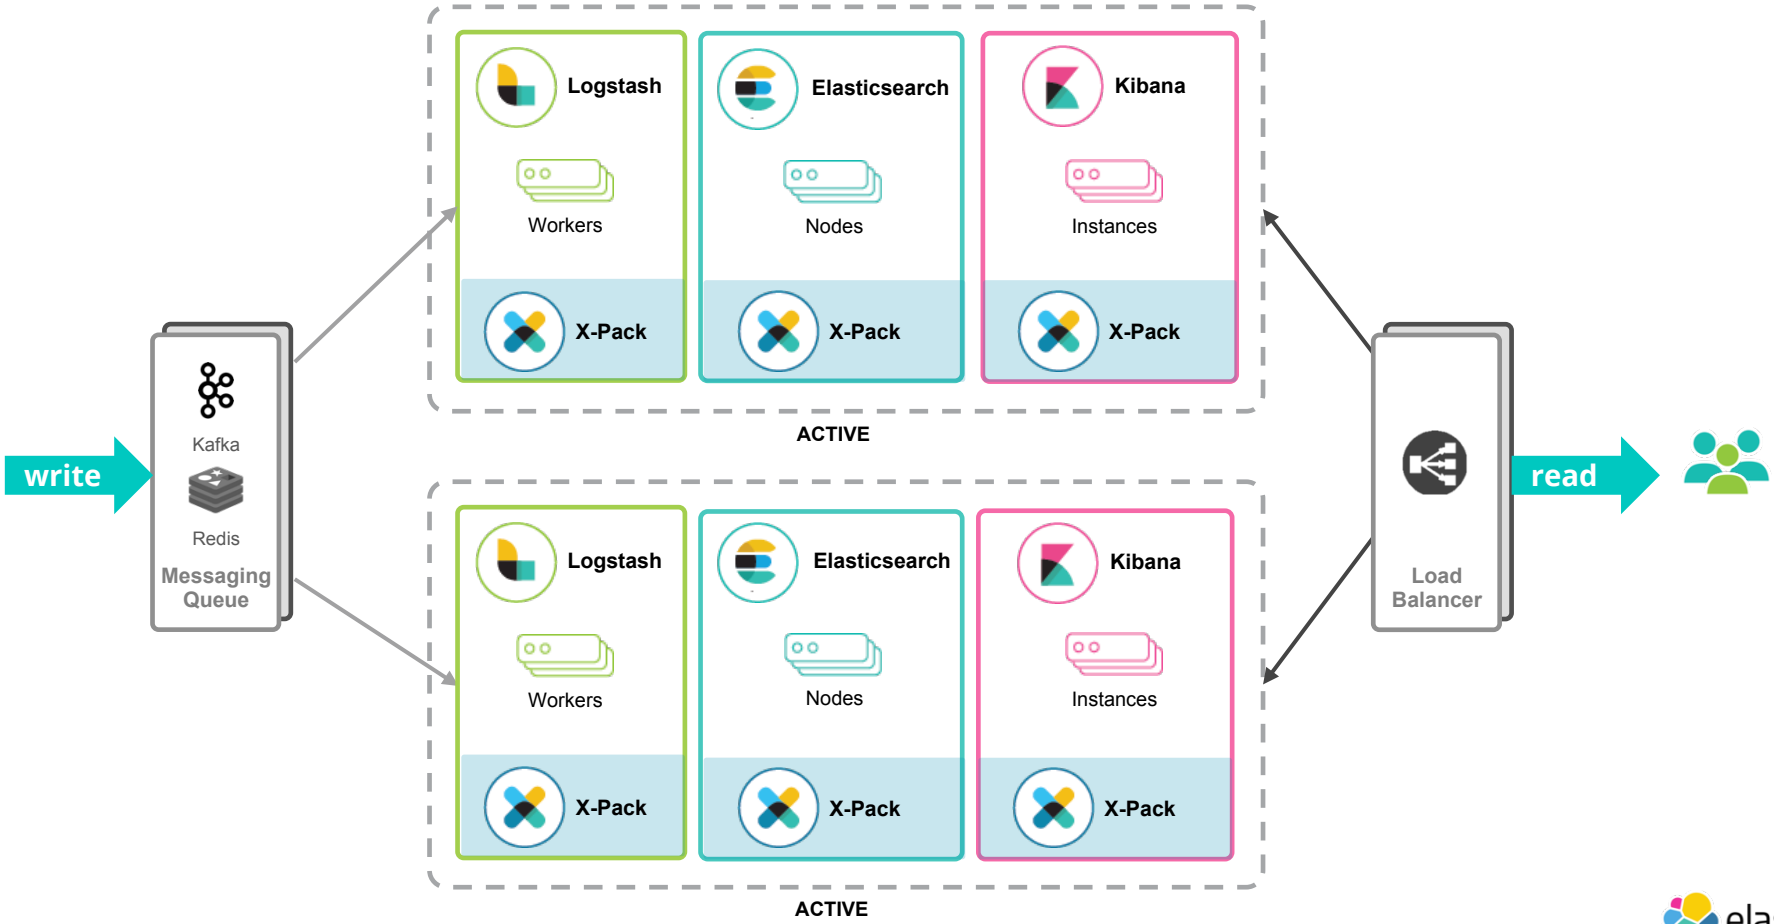

elastic

多数据中心非冗余数据夸区查询

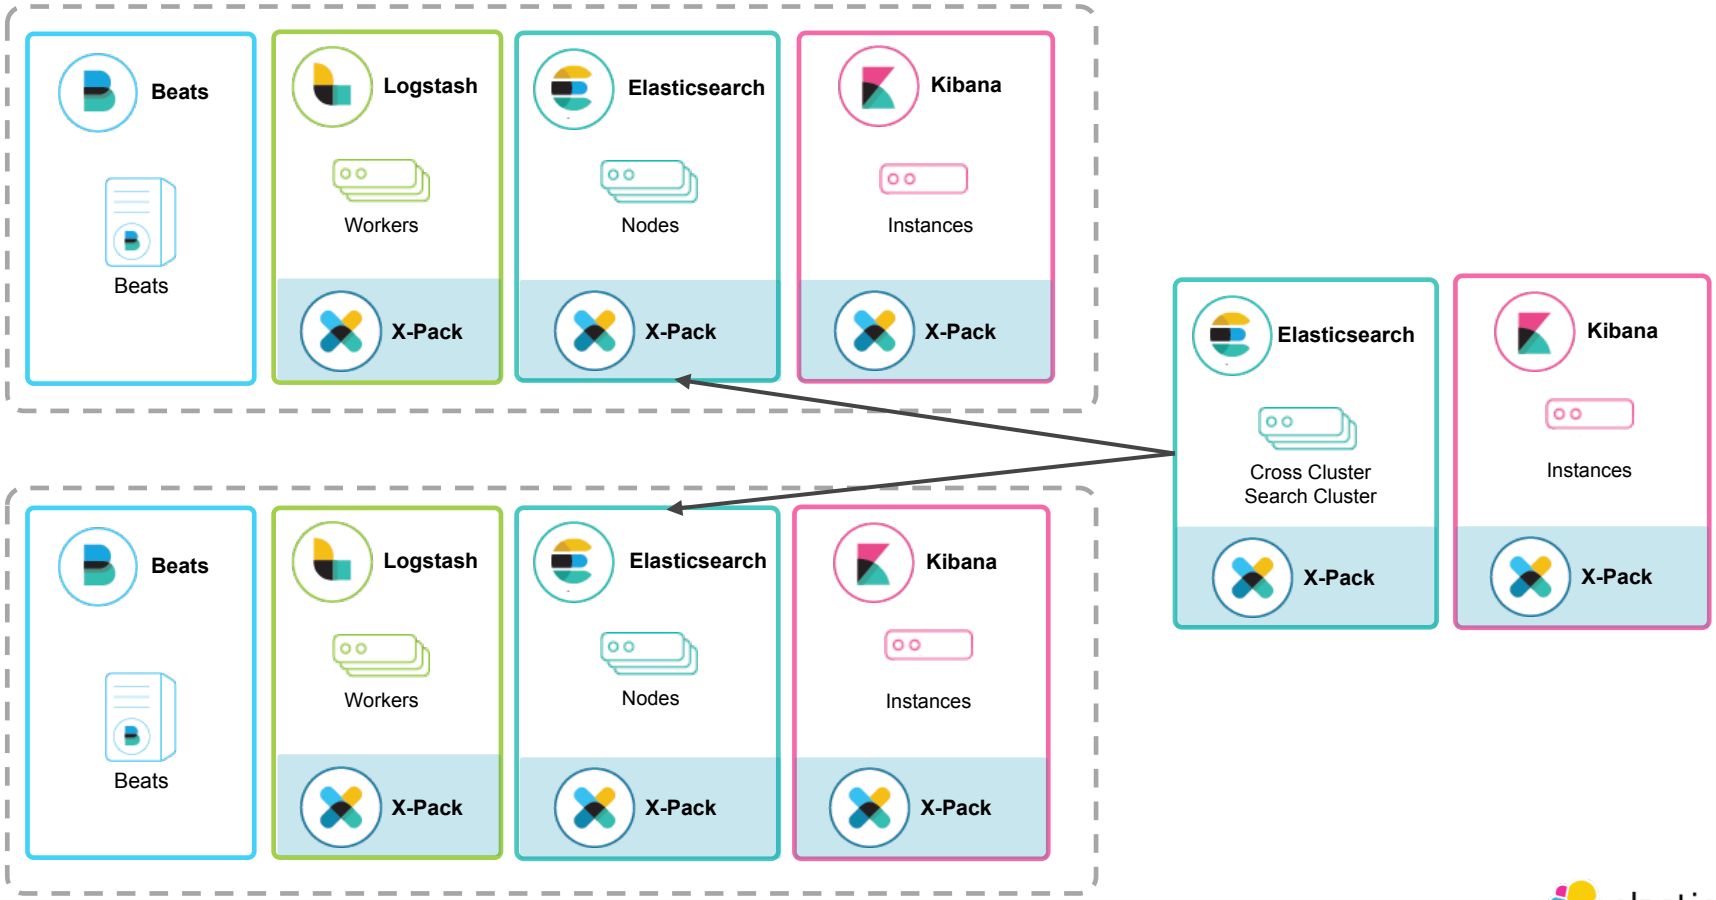

elastic

基于Elastic Stack的SIEM

业务场景构建

elastic

SSH 登陆情况面板

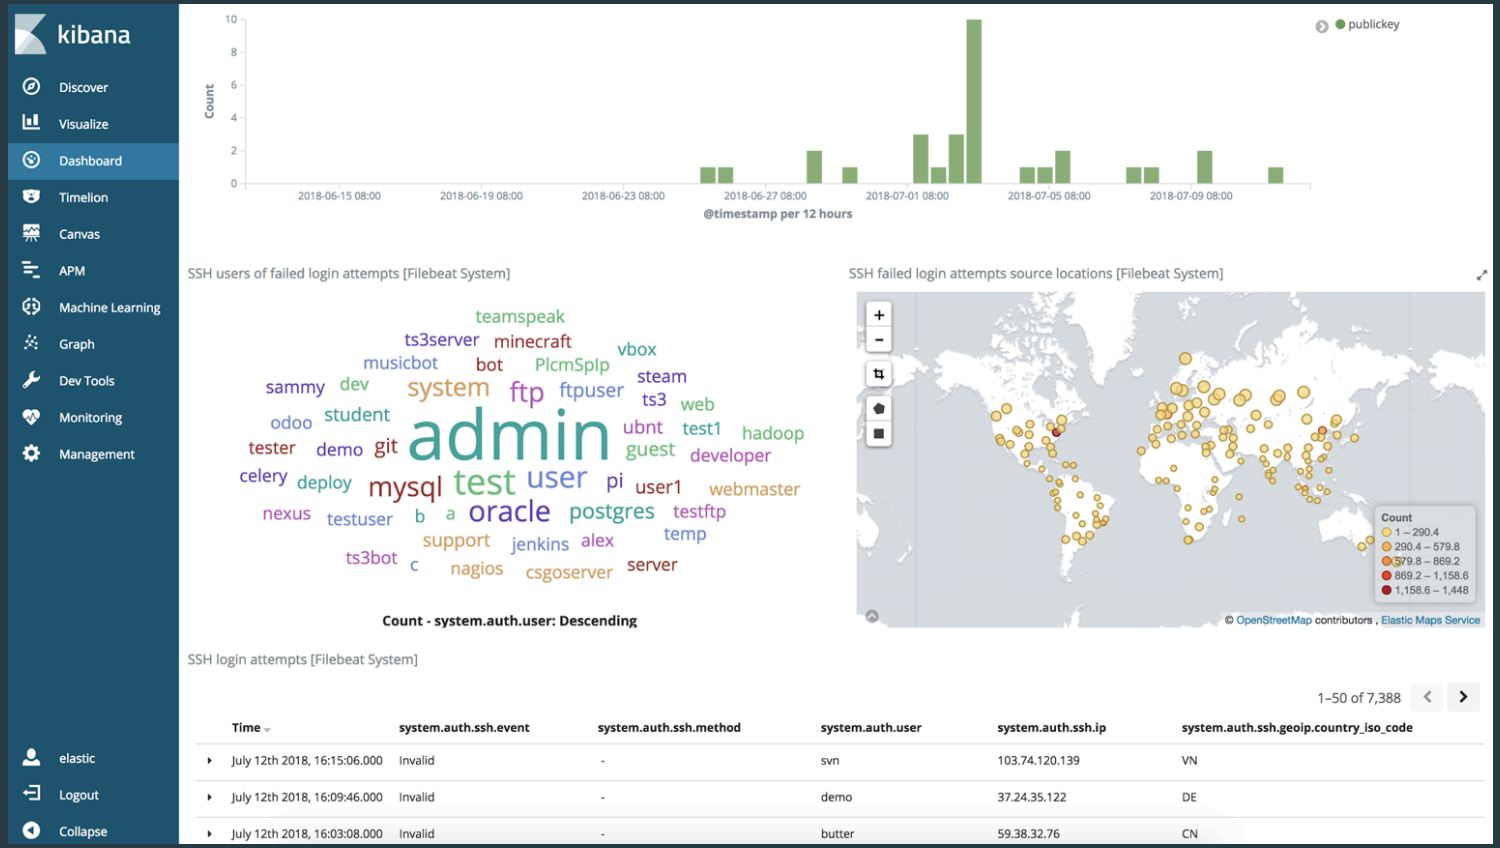

SSH 登陆场景

业务实体

- 用户名（谁）
- 时间
- IP（地点）

登陆成功

- 异常时间
- 异常地点

如何确定异常?

登陆失败

- 暴力登陆
  - 多少次算暴力登陆呢?

SSH 登陆信息探索

```
      },
      "filter": {
        "range": {
          "@timestamp": {
            "gte": "2018-07-01T00:00:00"
          }
        }
      }
    }
  }
}

GET filebeat-*/_search
{
  "size": 0,
  "aggs": {
    "NAME": {
      "cardinality": {
        "field": "system.auth.ssh.event"
      }
    }
  }
}

GET filebeat-*/_search
{
  "size": 0,
  "aggs": {
    "NAME": {
      "terms": {
        "field": "system.auth.ssh.event",
        "size": 10
      }
```

```
{
  "took": 3,
  "timed_out": false,
  "_shards": {
    "total": 1,
    "successful": 1,
    "skipped": 0,
    "failed": 0
  },
  "hits": {
    "total": 93549,
    "max_score": 0,
    "hits": []
  },
  "aggregations": {
    "NAME": {
      "value": 3
    }
  }
}
```

```
  }
}

GET filebeat-*/_search
{
  "size": 0,
  "aggs": {
    "NAME": {
      "cardinality": {
        "field": "system.auth.ssh.event"
      }
    }
  }
}

GET filebeat-*/_search
{
  "size": 0,
  "aggs": {
    "NAME": {
      "terms": {
        "field": "system.auth.ssh.event",
        "size": 10
      }
    }
  }
}

GET /_settings/index.version.created

GET _xpack/license

GET server-metrics/_search

PUT _ingest/pipeline/plus_one_year
{
  "processors": [
    {
      "script": {
        "lang": "painless",
        "source": """
```

```
{
  "took": 8,
  "timed_out": false,
  "_shards": {
    "total": 1,
    "successful": 1,
    "skipped": 0,
    "failed": 0
  },
  "hits": {
    "total": 93549,
    "max_score": 0,
    "hits": []
  },
  "aggregations": {
    "NAME": {
      "doc_count_error_upper_bound": 0,
      "sum_other_doc_count": 0,
      "buckets": [
        {
          "key": "Invalid",
          "doc_count": 10625
        },
        {
          "key": "error:",
          "doc_count": 45
        },
        {
          "key": "Accepted",
          "doc_count": 38
        }
      ]
    }
  }
}
```

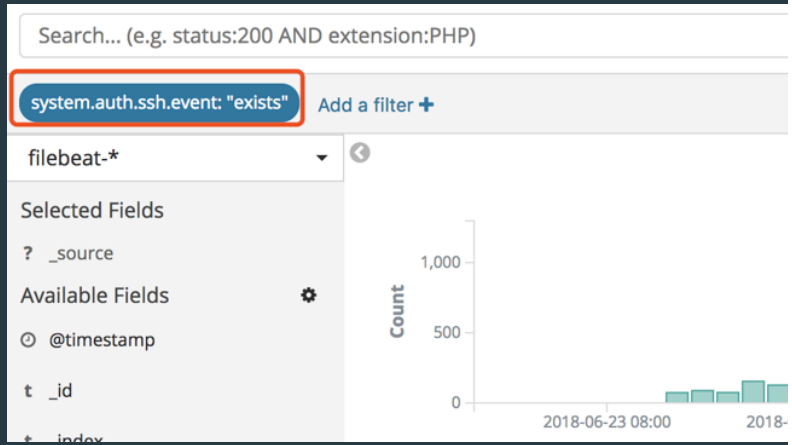


```
"@timestamp": "2018-07-20T08:39:07.000Z",
"system": {
  "auth": {
    "hostname": "ip-172-31-13-50",
    "ssh": {
      "geoip": {
        "continent_name": "Asia",
        "city_name": "Shenzhen",
        "country_iso_code": "CN",
        "region_name": "Guangdong",
        "location": {
          "lon": 114.1333,
          "lat": 22.5333
        }
      },
      "ip": "36.36.201.21",
      "event": "Invalid"
    },
    "pid": "25062",
    "user": "nexus",
    "timestamp": "Jul 20 08:39:07"
  }
},
"beat": {
  "hostname": "ip-172-31-13-50",
  "name": "ip-172-31-13-50",
  "version": "6.3.0"
},
"host": {
  "name": "ip-172-31-13-50"
}
},
```

SSH暴力破解异常检测

elastic

选择学习数据

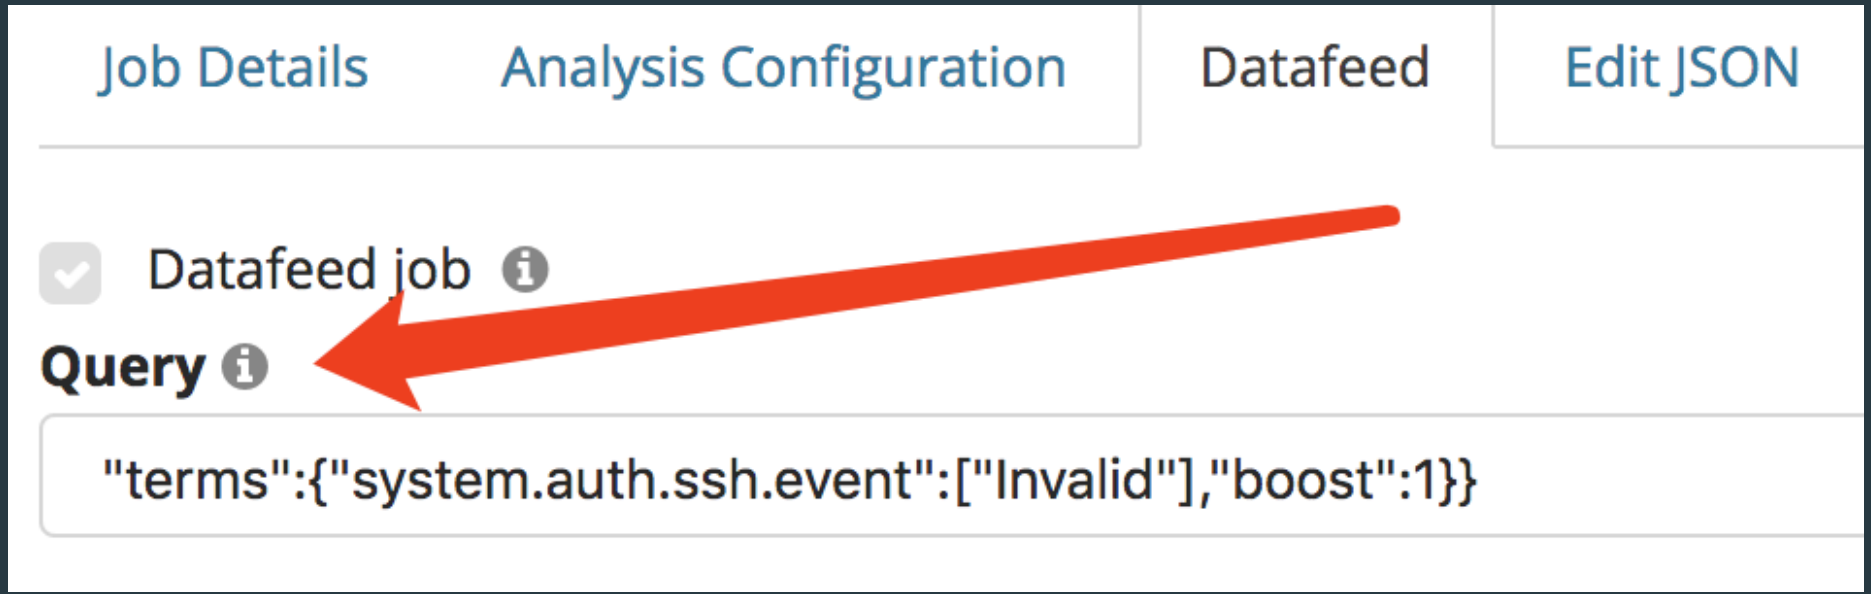

指定观测指标

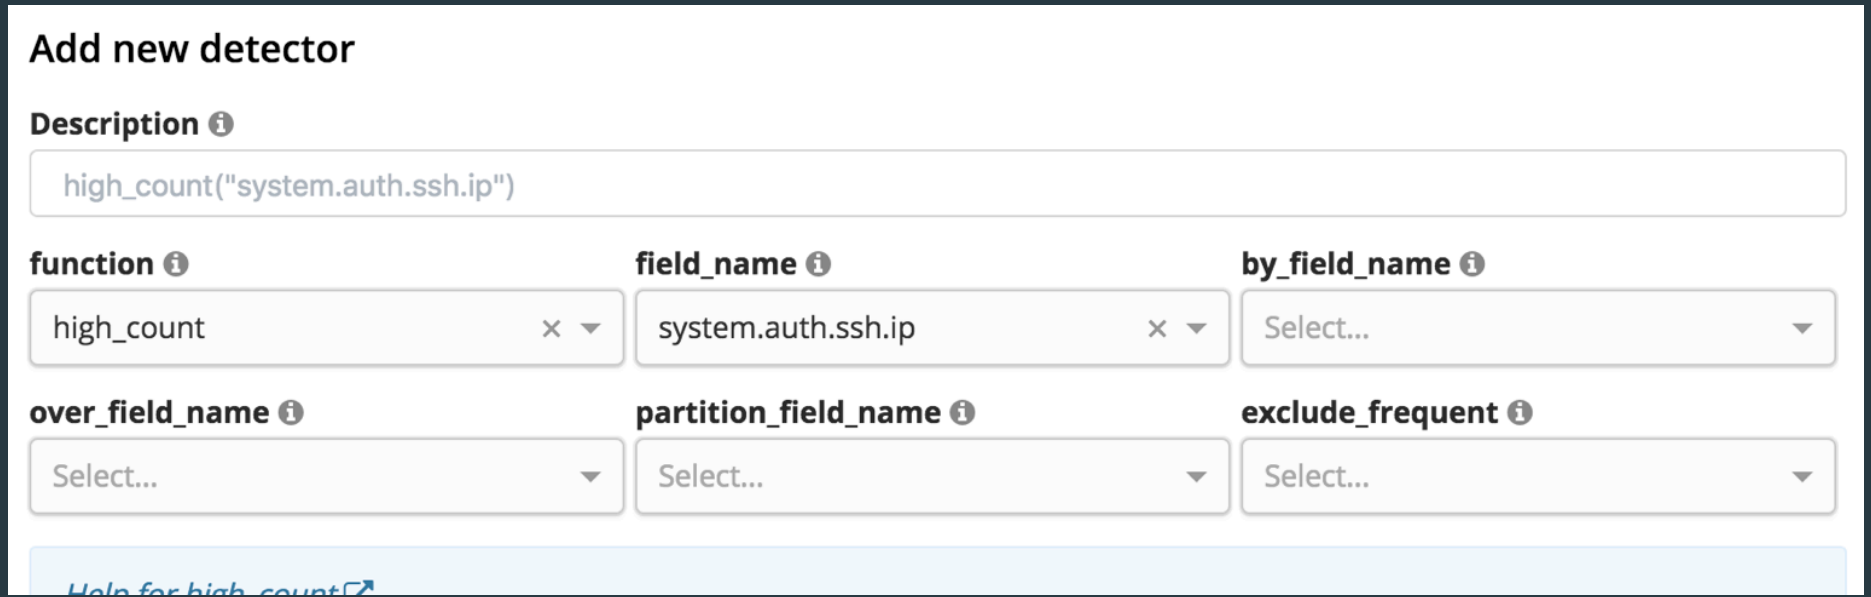

Influencers

- [ ] system.auth.ssh.geoip.country_iso_code
- [ ] system.auth.ssh.geoip.region_name
- [x] system.auth.ssh.ip
- [ ] system.auth.ssh.method
- [ ] system.auth.ssh.signature
- [ ] system.auth.sudo.command
- [ ] system.auth.sudo.error
- [ ] system.auth.sudo.pwd
- [ ] system.auth.sudo.tty
- [ ] system.auth.sudo.user
- [ ] system.auth.timestamp
- [x] system.auth.user
- [ ] system.auth.useradd.home

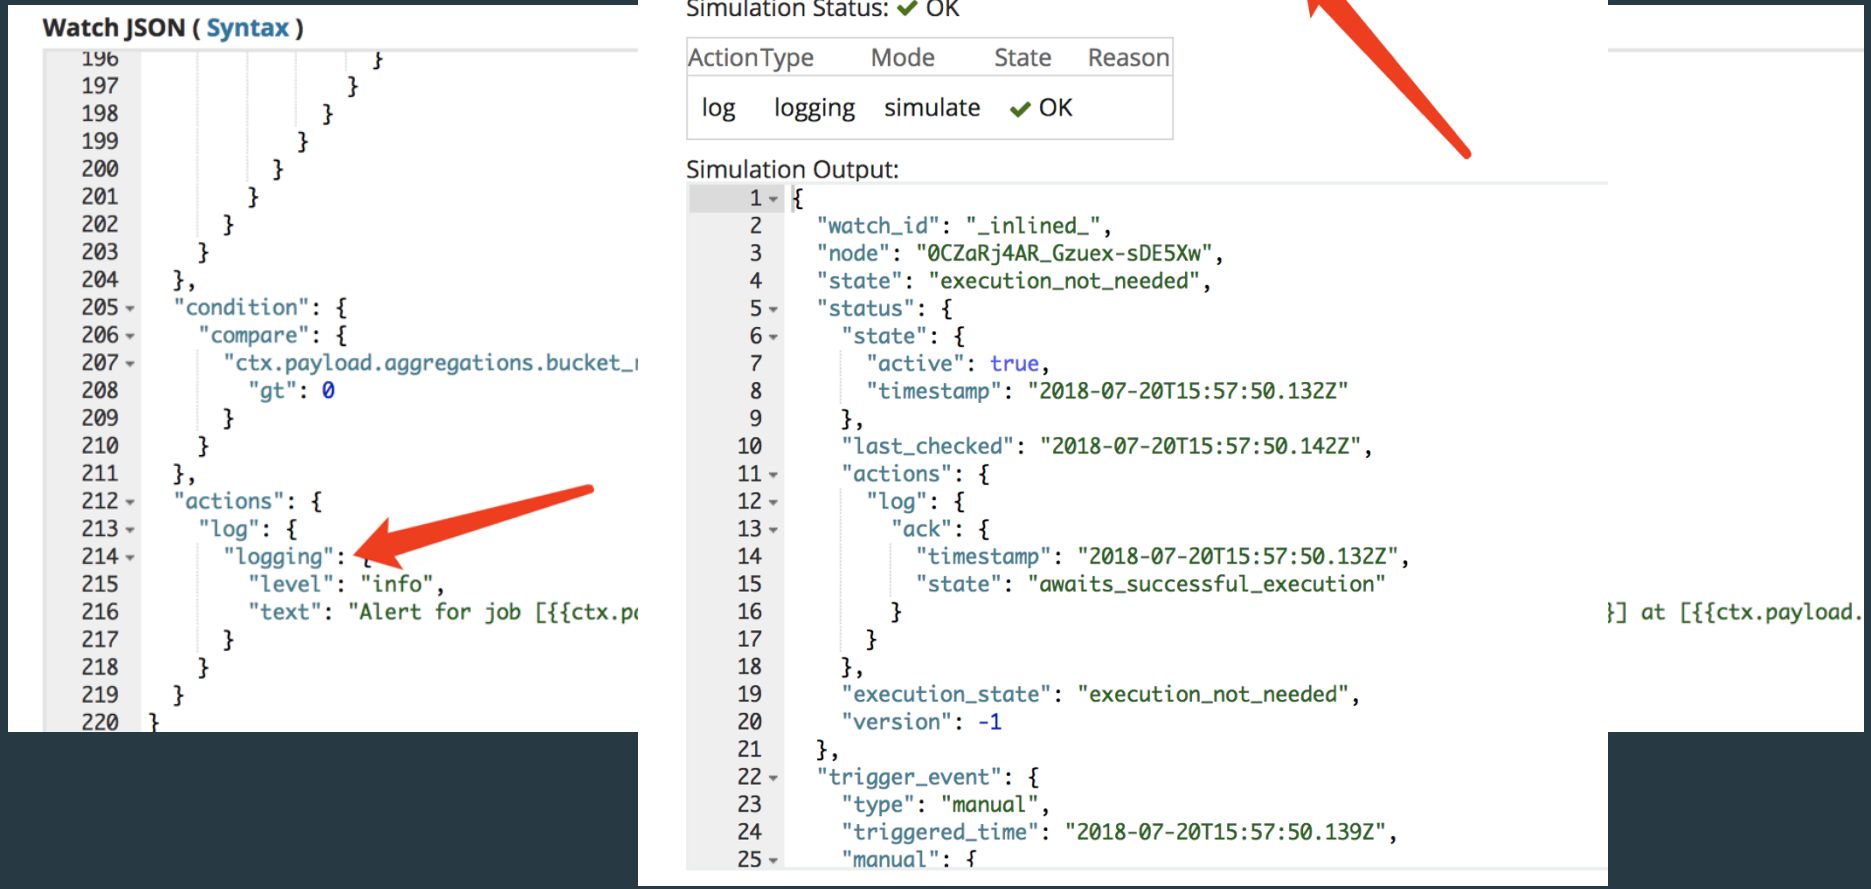

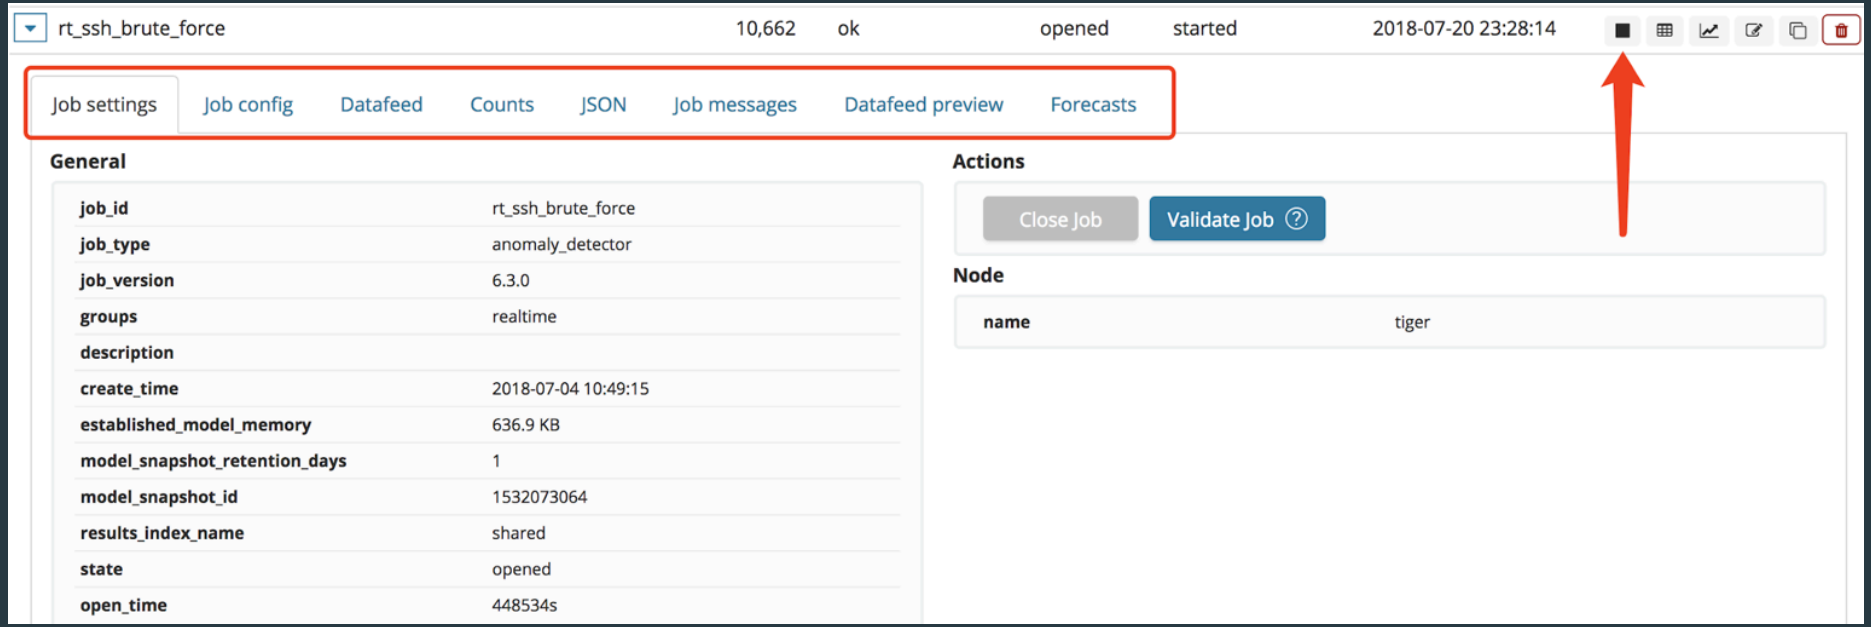

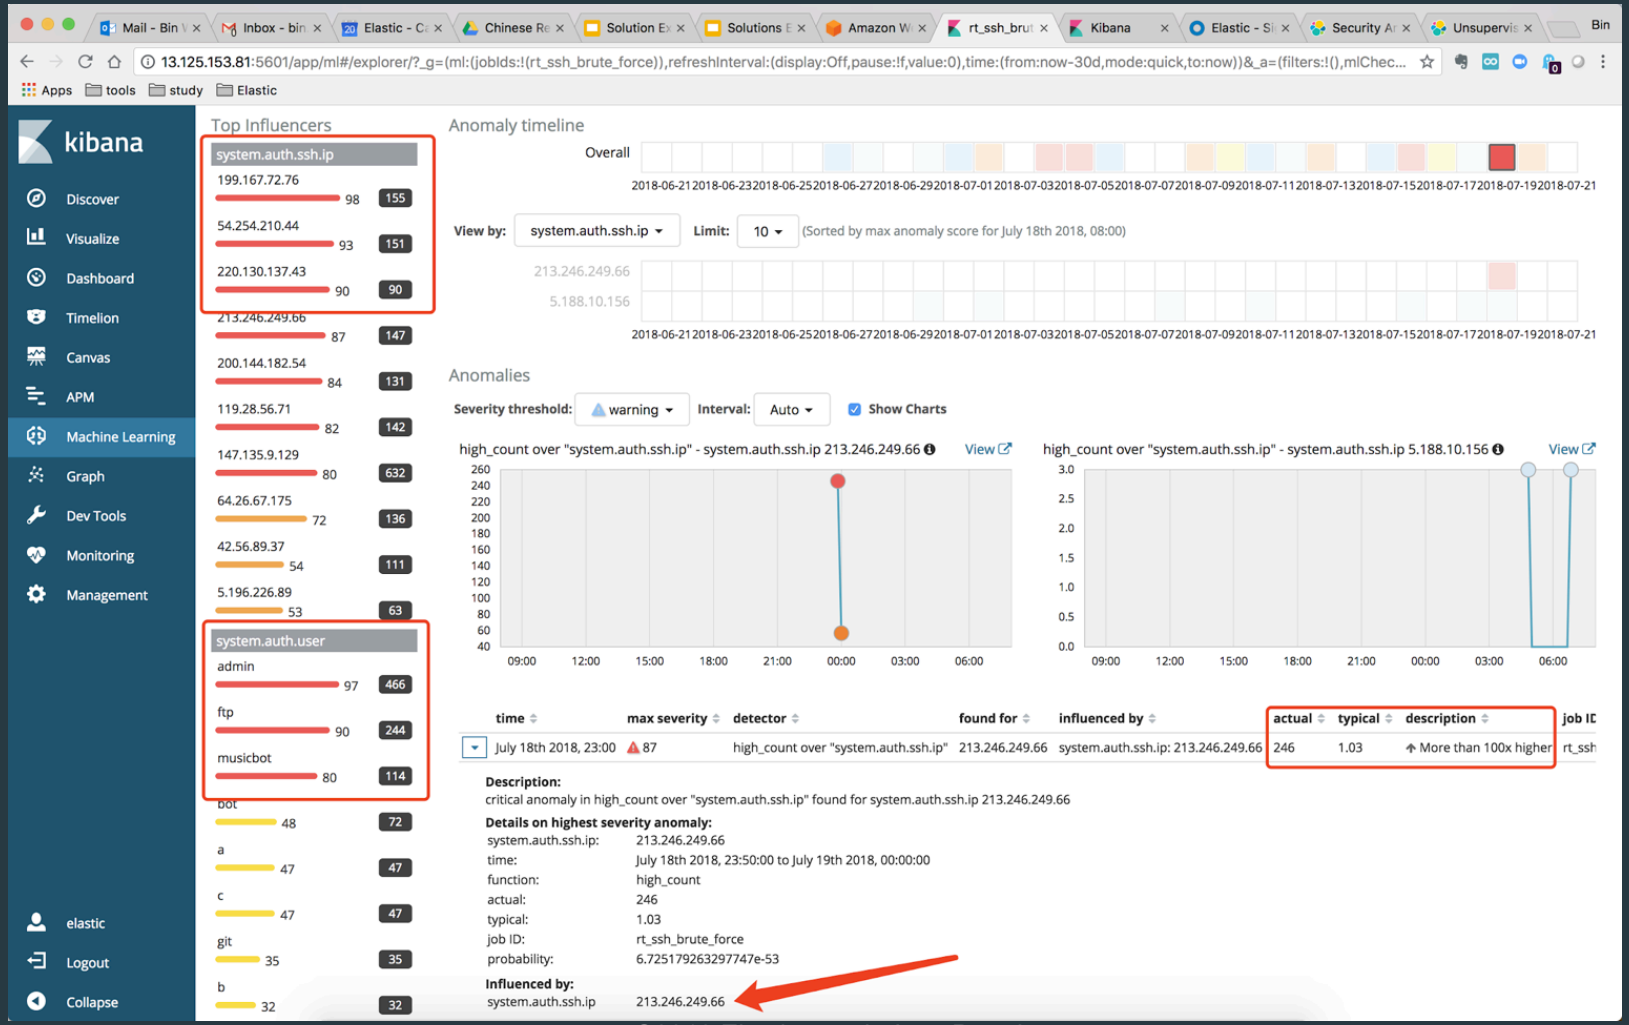

图关联查询

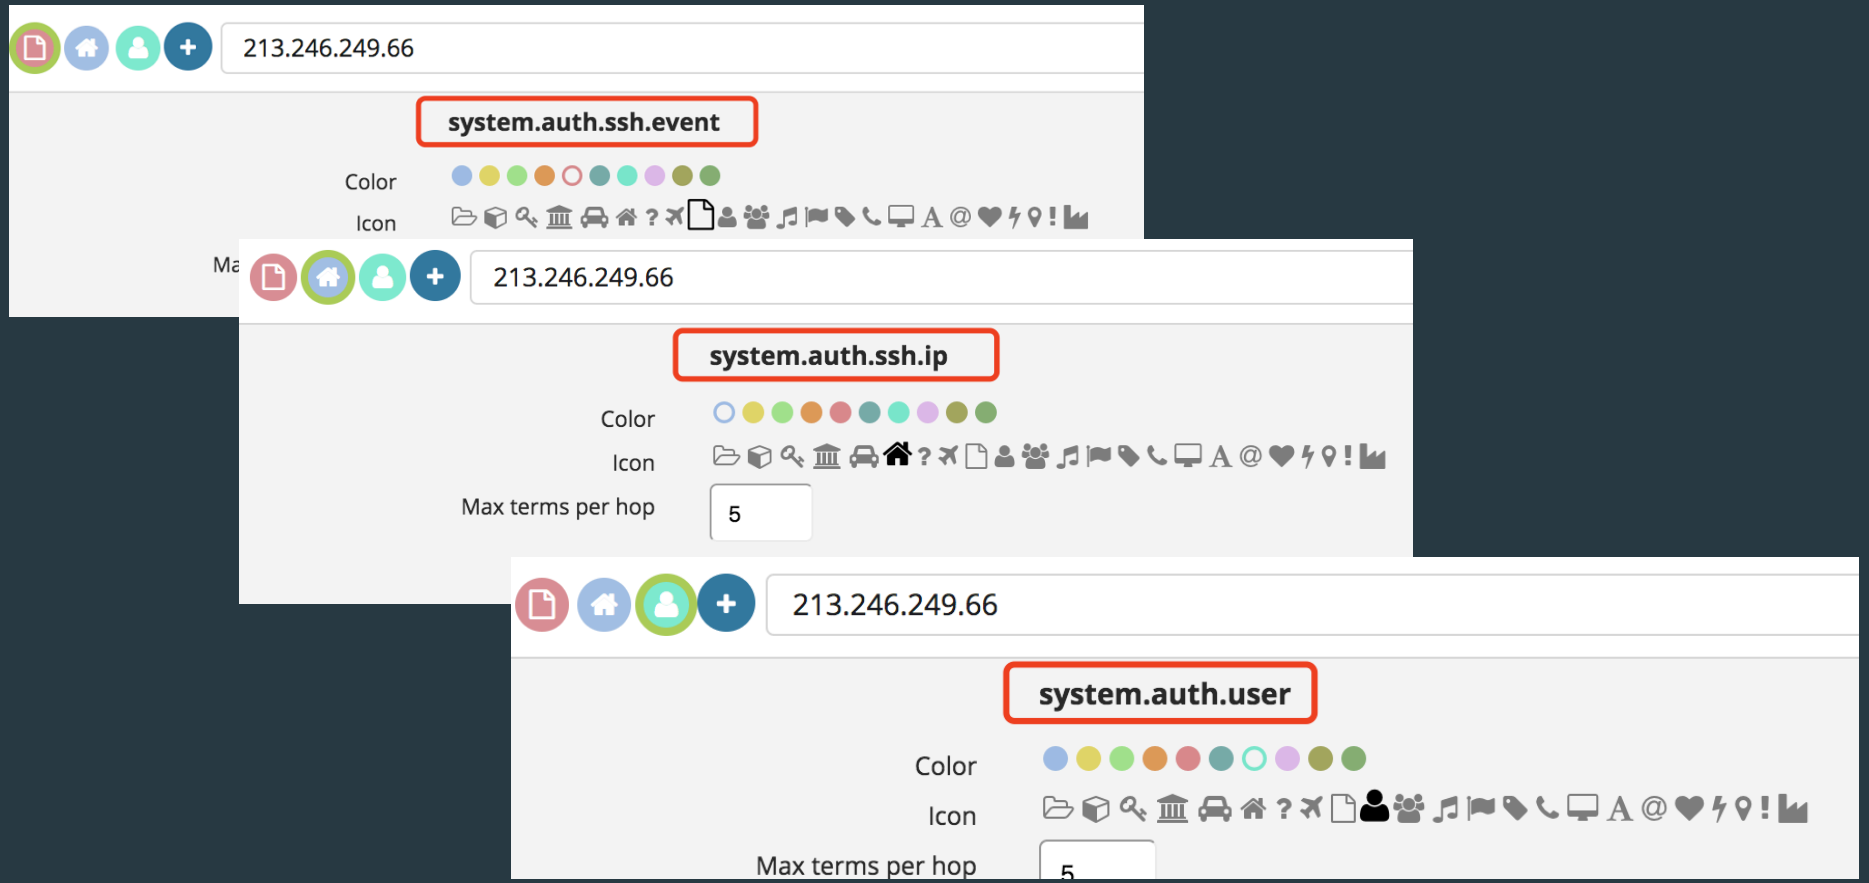

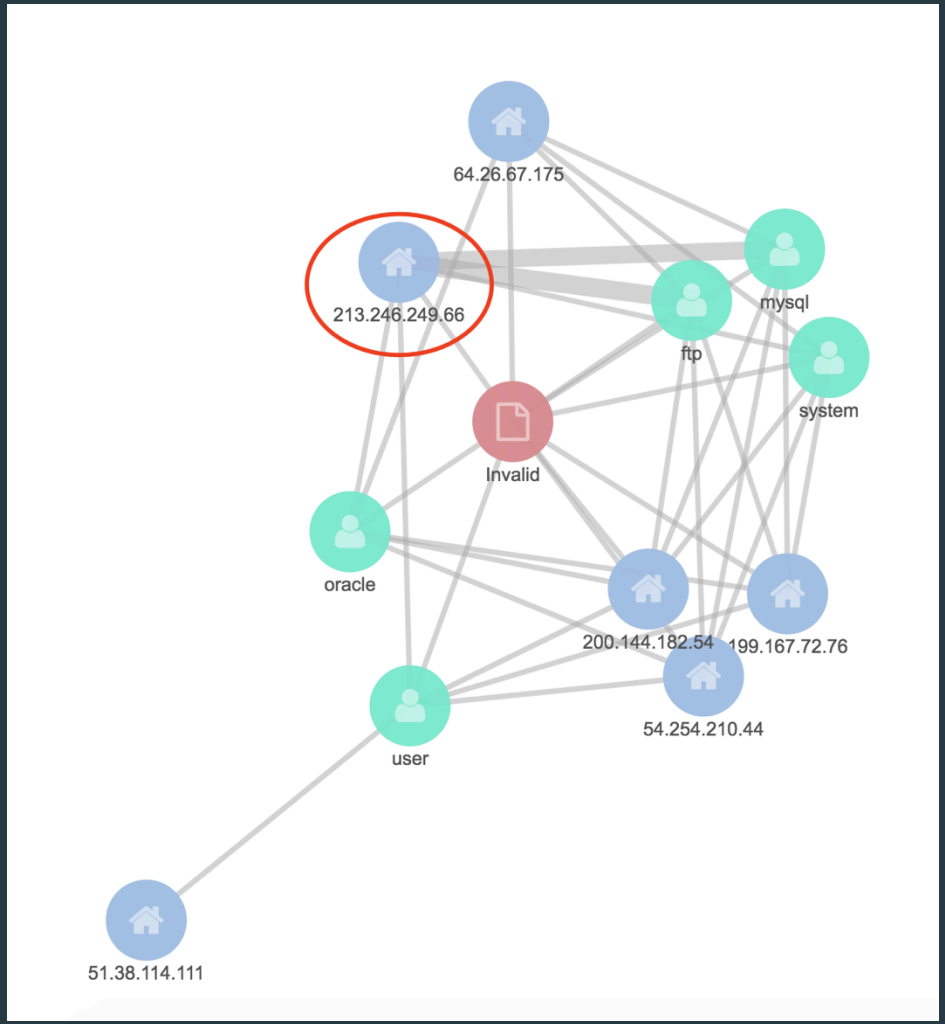

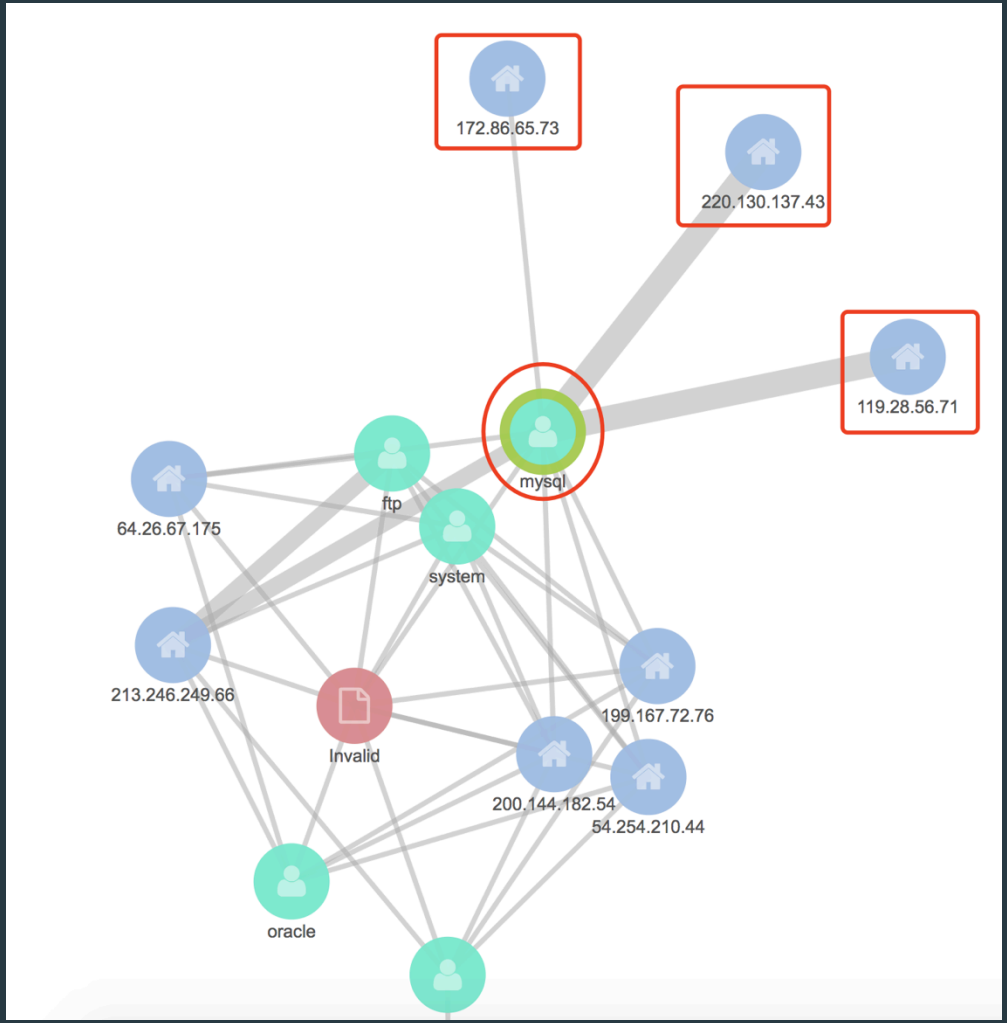

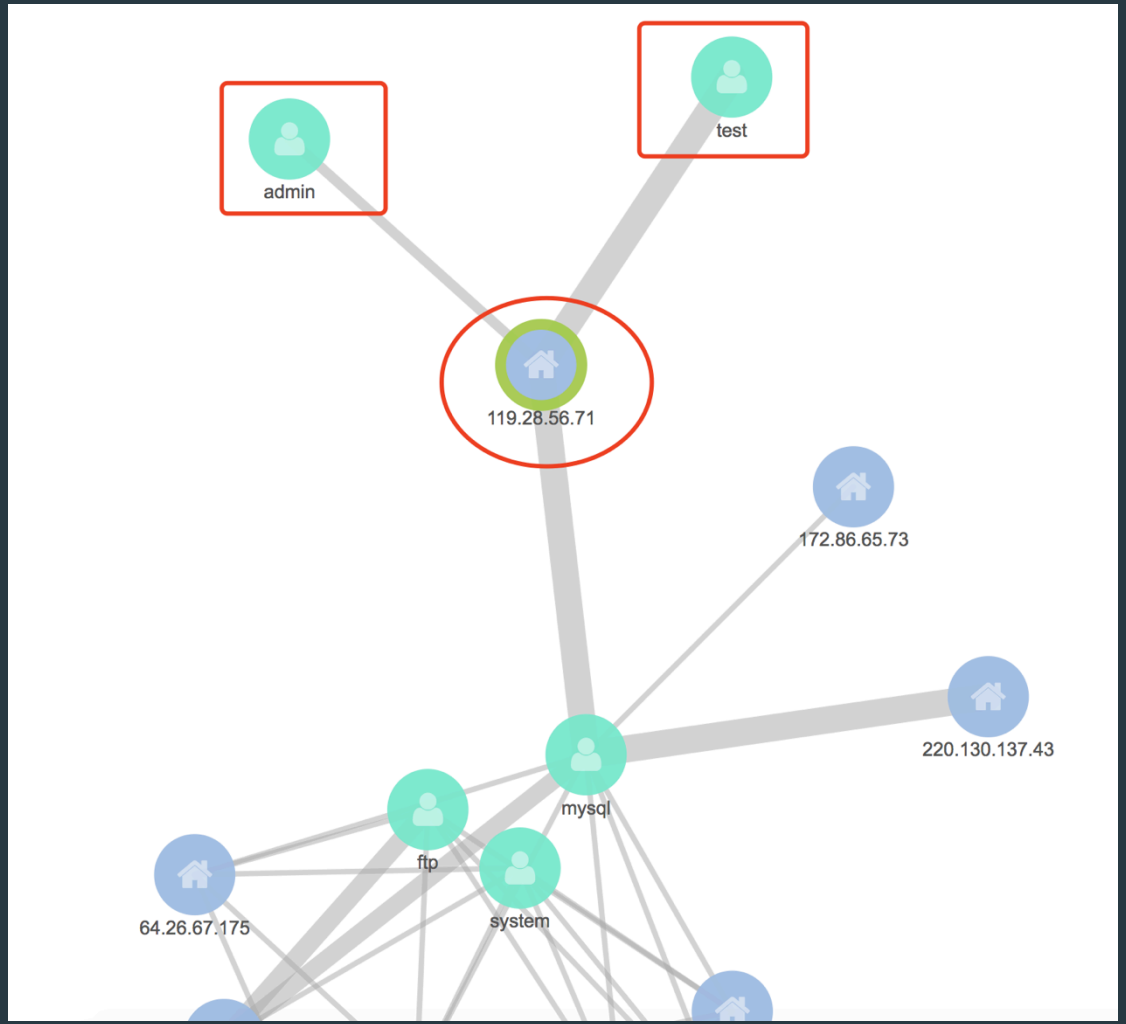

异常检测任务组

rt_port_scan	3,642,361	ok	opened	started	2018-07-21 00:09:50	
rt_ssh_brute_force	10,665	ok	opened	started	2018-07-21 00:00:10	
rt_unusual_login_location	36	ok	opened	started	2018-07-19 08:37:18	
rt_unusual_login_time	36	ok	opened	started	2018-07-19 08:37:18	
rt_unusual_process	7,495,861	ok	opened	started	2018-07-21 00:04:59	

Detectors

high_distinct_count("dest.port")

Influencers

dest.ip

Detectors

rare by "system.auth.ssh.geoip.city_name"

Influencers

system.auth.ssh.ip, system.auth.user, system.auth.ssh.geoip.country_iso_code, system.auth.ssh.geoip.continent_name, system.auth.ssh.geoip.city_name

Detectors

time_of_day by "system.auth.user"

Influencers

system.auth.ssh.ip, system.auth.user

Detectors

rare by "system.process.name"

Influencers

system.process.cmdline, system.process.cwd, system.process.username, system.process.state

https://www.elastic.co/solutions/security-analytics

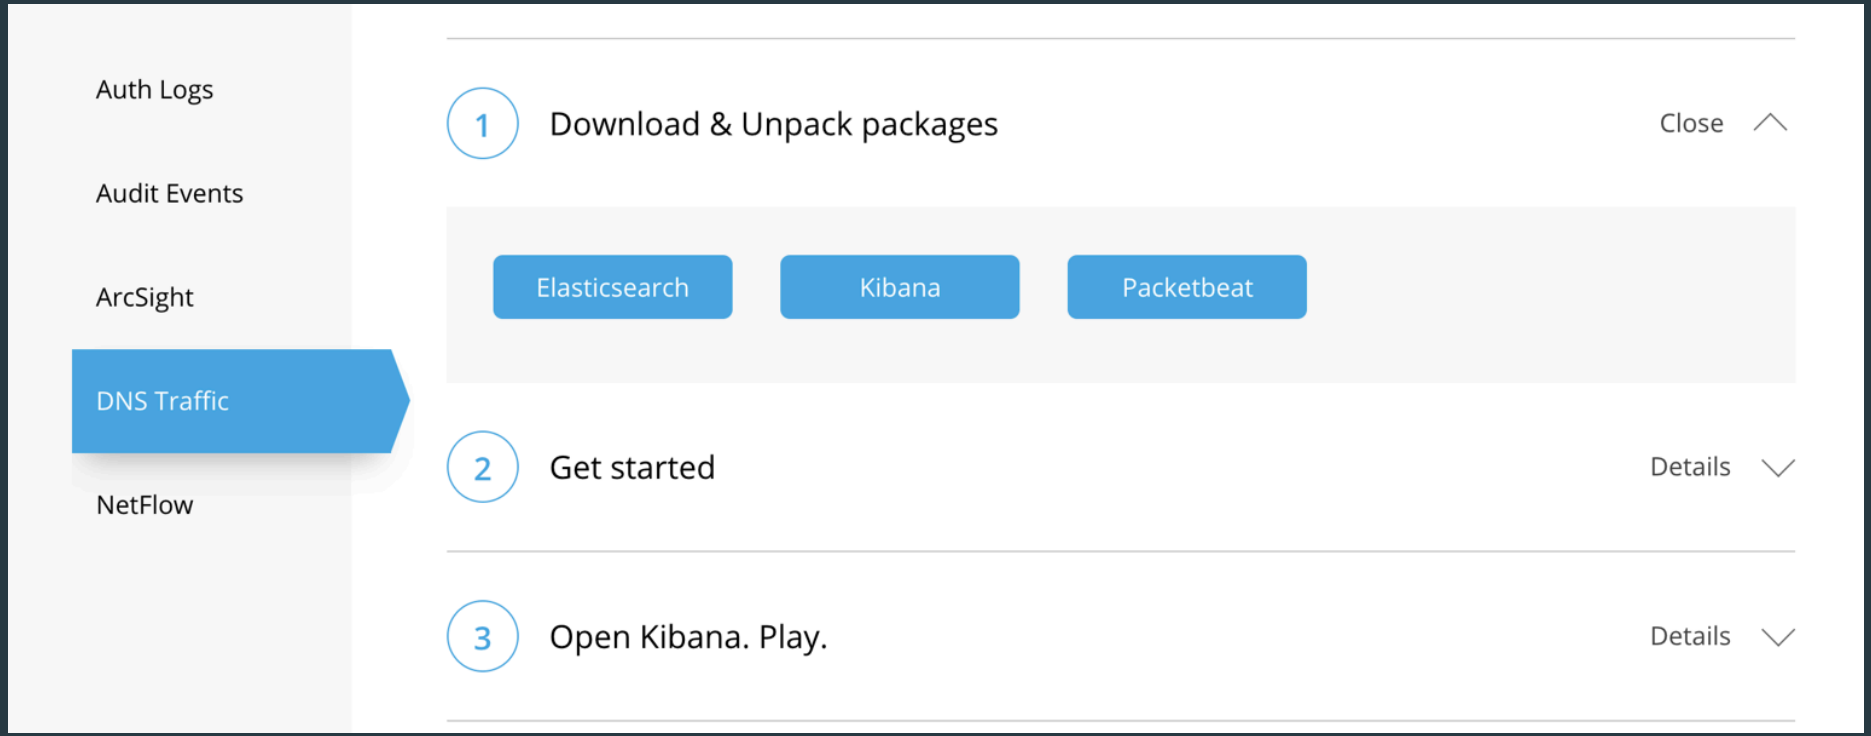

https://www.elastic.co/products/x-pack/machine-learning/recipes

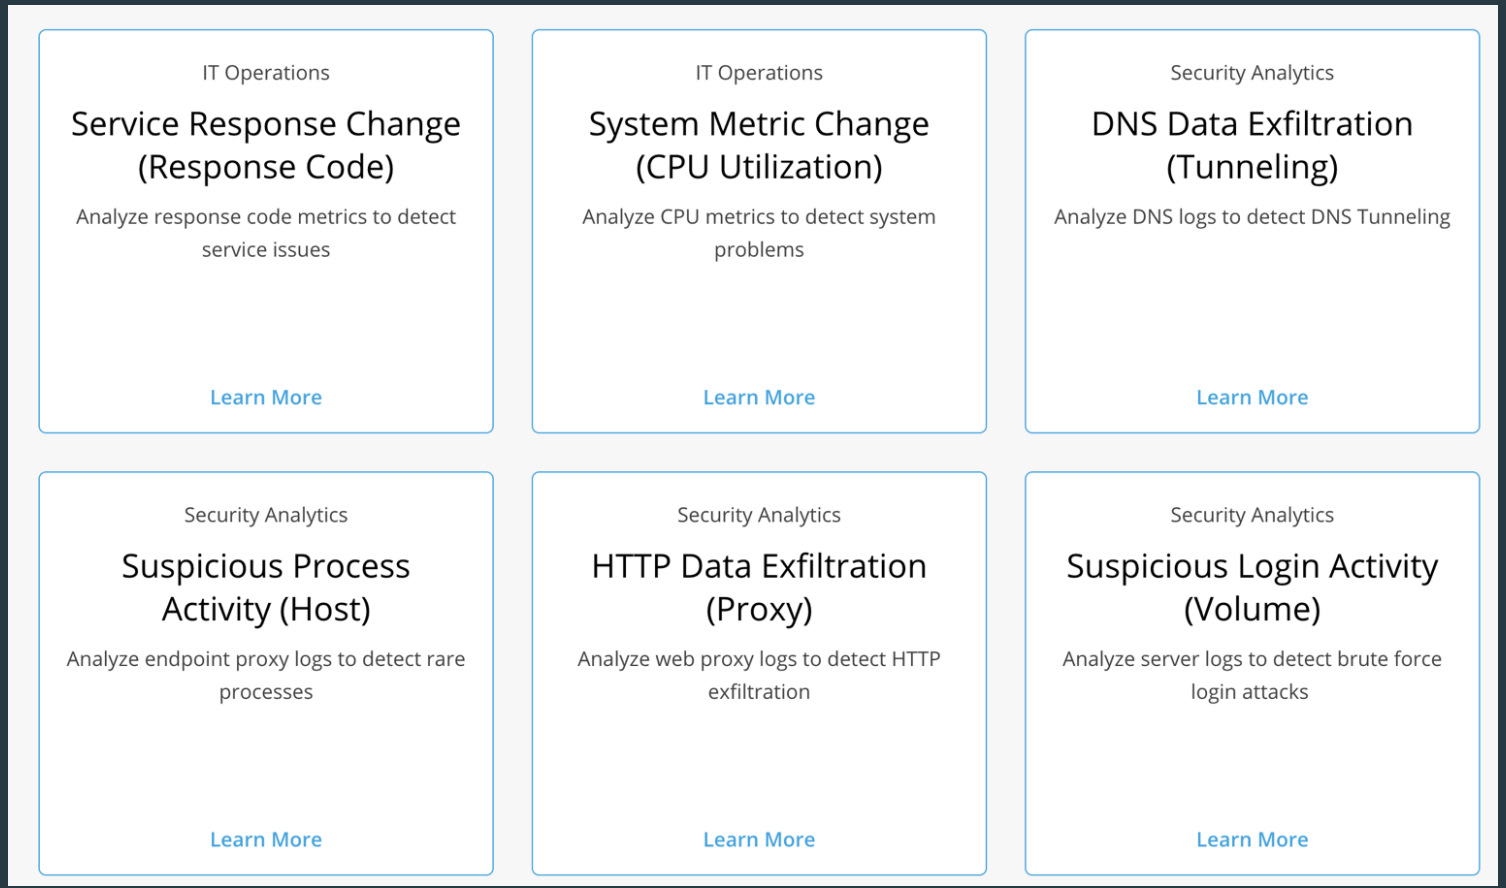

elastic
中文社区

专业、垂直、纯粹的 Elastic 开源技术交流社区
https://elasticsearch.cn/