阿里云 MVP

ELK在运维工作中应用两三事

上海安畅运维专家 韩军辉

阿里云 MVP

场景一：数据中心流量分析

阿里云 MVP

安畅网络背景介绍

阿里云 MVP

面对的挑战和思考

挑战

单数据中心每秒最高10G+流量

单数据中心每月高达3000次DDOS

思考

流量数据的分析与展现

快速检查出DDOS攻击及异常流量

快速识别出攻击的线路、IP地址、对应客户信息等

阿里云 MVP

数据中心流量全景

北艾机房流量全景视图

*

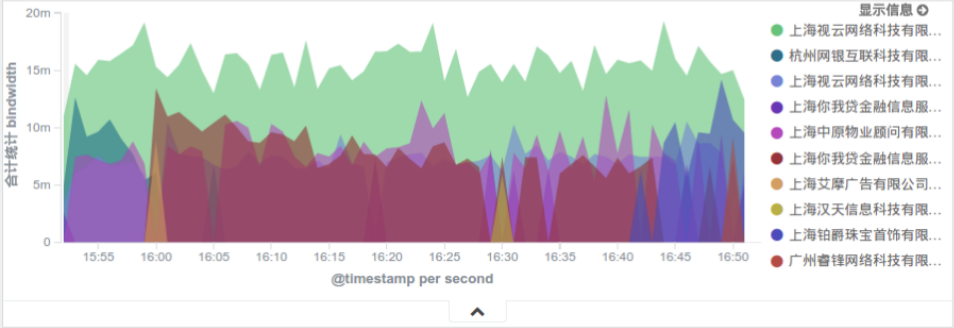

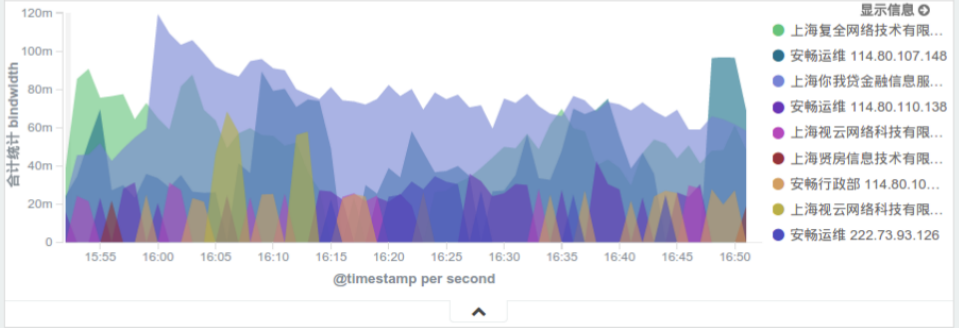

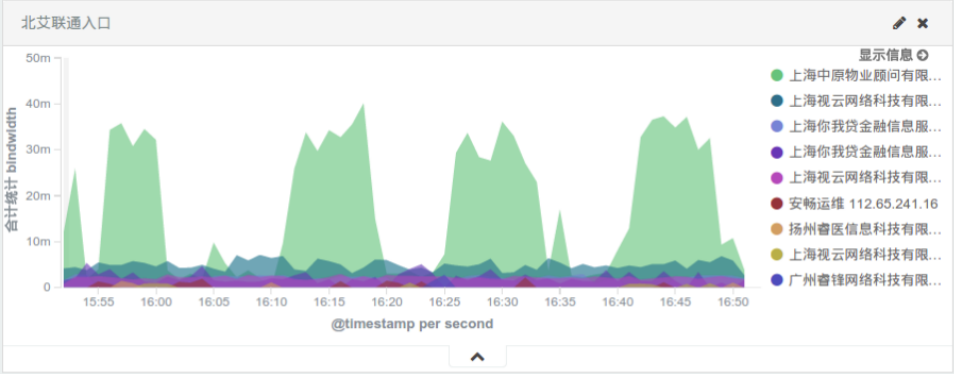

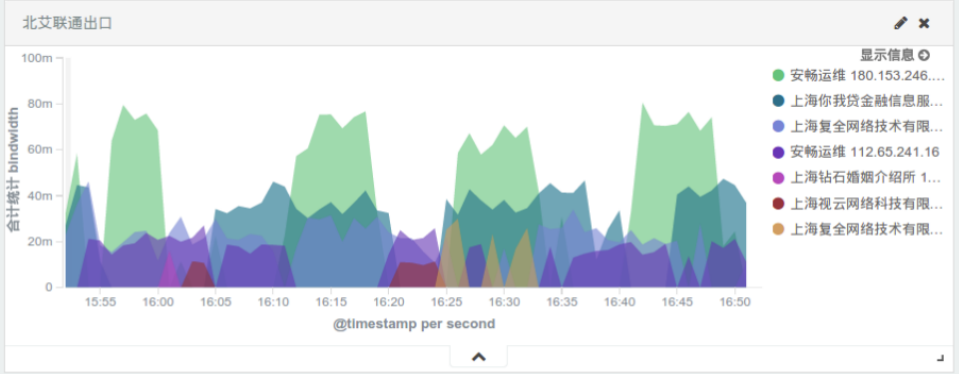

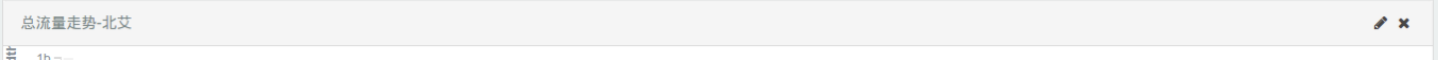

阿里云 MVP

安畅网络业务平台流量全景

流量全景图 *

入口流量走势

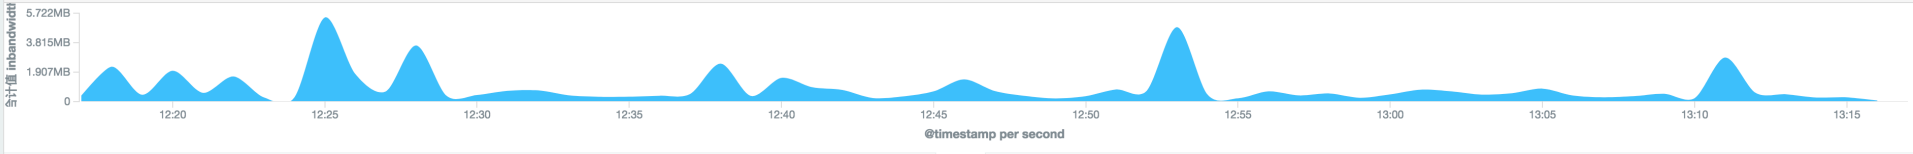

应用层协议分布

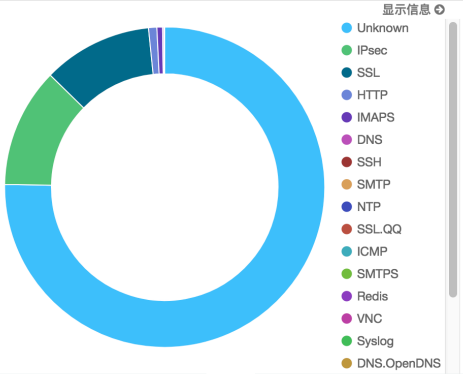

访问流量最多的地区

传输层协议分布

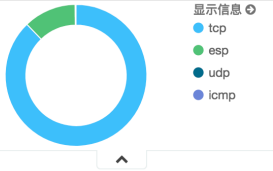

入口运营商统计

入口总流量

395.391MB

合计值 inbytes

进入流量最多的来源IP

Top 10 IPV4_SRC_ADDR	合计值 inbytes
118.193.144.4	84.162MB
221.228.82.34	54.36MB
101.227.66.81	48.933MB
117.121.25.6	36.8MB

访客来源地热点图-flow

阿里云 MVP

DDOS攻击告警

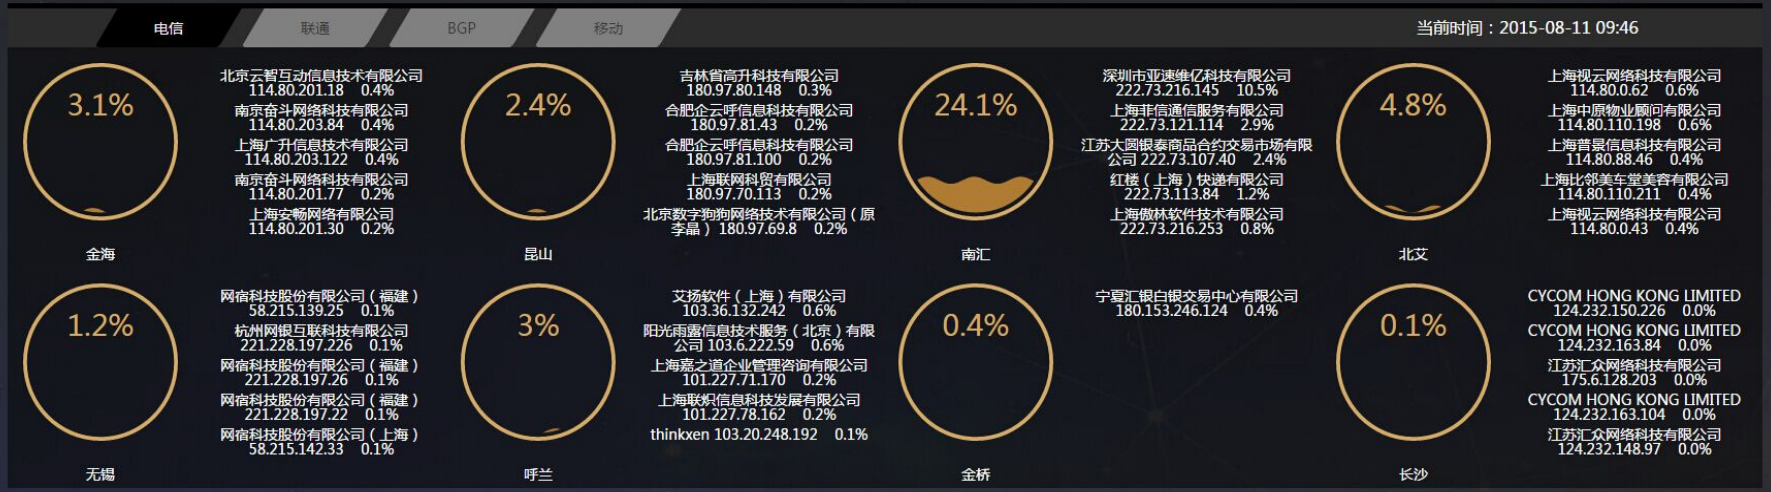

TIME	DATA CENTER	AGENT IP	PERCENT	CUSTOMER
2015-08-18 16:11:24	长沙	124.232.137.153	16.81%	CYCOM HONG KONG LIMITED 124.232.163.120
2015-08-18 16:11:24	无锡	58.215.171.6	10.33%	吉林省高升科技有限公司 103.21.119.5
2015-08-18 16:11:24	怒江	117.135.137.81	10.35%	CYCOM HONG KONG LIMITED 221.181.64.40
2015-08-18 16:11:24	昆山	180.97.70.3	10.27%	合肥企云呼信息科技有限公司 180.97.81.100
2015-08-18 16:11:24	鲁谷	124.202.141.9	12.36%	恒拓开源（天津）信息科技股份有限公司 211.155.89.196
2015-08-18 16:11:24	鲁谷	124.202.141.9	21.42%	上海视云网络科技有限公司 124.202.157.13
2015-08-18 16:11:24	南汇	222.73.124.237	13.62%	上海菲信通信服务有限公司 114.141.132.130
2015-08-18 16:11:24	长沙	124.232.137.153	11.41%	江苏汇众网络科技有限公司 175.6.128.202
2015-08-18 16:11:24	北艾	222.73.184.128	14.50%	上海中原物业顾问有限公司 114.80.110.198
2015-08-18 16:11:24	鲁谷	124.202.141.9	15.57%	上海五帝科技有限公司 124.202.157.35
2015-08-18 16:11:24	无锡	58.215.171.6	11.69%	吉林省高升科技有限公司 103.21.119.7

阿里云 MVP

技术架构

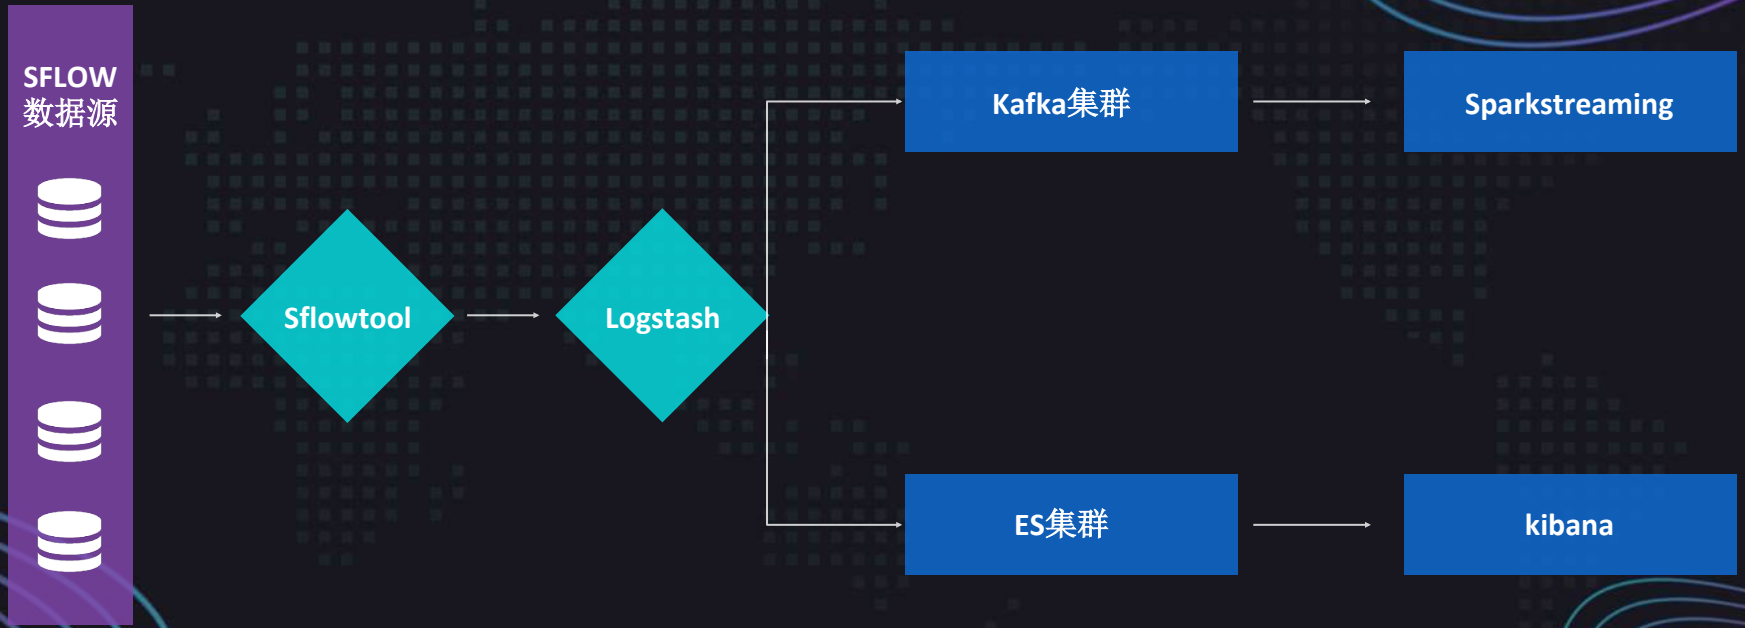
SFLOW
数据源
Sflowtool
Logstash
Kafka集群
Sparkstreaming
ES集群
kibana

阿里云 MVP

关键配置

```
sFlow Version 5 Information:
------------------------------------------------------------------------------
Agent Information:

     IP Address: 103.6.222.166
 Address family: IPV4
   Vpn-instance: NA
------------------------------------------------------------------------------
Collector Information:

   Collector ID: 1
     IP Address: 58.[illegible]2
 Address family: IPV4
   Vpn-instance: NA
           Port: 2065
  Datagram size: 1400
       Time out: NA
    Description: Sflow

   Collector ID: 2
     IP Address: 221.[illegible]
 Address family: IPV4
   Vpn-instance: NA
           Port: 6352
  Datagram size: 1400
       Time out: NA
    Description: JiDanTu
------------------------------------------------------------------------------
Port on slot 1/5 Information:
```

阿里云 MVP

关键配置

```
Interface: XGE1/6/0/3
 Flow-sample collector: NA              Counter-sample collector   : 1,2
 Flow-sample rate(1/x): 512             Counter-sample interval(s): 2
 Flow-sample maxheader: 128
 Flow-sample direction: IN,OUT
------------------------------------------------------------------------------
Port on slot 2/5 Information:
------------------------------------------------------------------------------
Port on slot 2/6 Information:

Interface: XGE2/6/0/5
 Flow-sample collector: NA              Counter-sample collector   : 1,2
 Flow-sample rate(1/x): 512             Counter-sample interval(s): 2
 Flow-sample maxheader: 128
 Flow-sample direction: IN,OUT
```

"114.80.200.1": 安畅运维
"114.80.200.2": [redacted]限公司
"114.80.200.3": [redacted]技有限公司
"114.80.200.4": [redacted]限公司
"114.80.200.5": [redacted]科技有限公司
"114.80.200.6": [redacted]息技术有限公司
"114.80.200.7": [redacted]科技有限公司
"114.80.200.9": 安畅运维
"114.80.200.10": [redacted]技有限公司
"114.80.200.11": 安畅运维

阿里云 MVP

关键配置

```
input {
        pipe {
                type => "sflow"
                command => "/var/scripts/sflowtool-wrapper.sh -l -p 6343"
        }
        pipe {
                type => "sflow"
                command => "/var/scripts/sflowtool-wrapper.sh -l -p 6344"
        }
        pipe {
                type => "sflow"
                command => "/var/scripts/sflowtool-wrapper.sh -l -p 6345"
        }
```

阿里云 MVP

关键配置

```
filter {
## Sflow Monitor
        if "sflow" == [type] {
                grok {
                        match => [ "message","%{WORD},%{IP:Agent_adress},%{NUMBER:In_Port},%{NUMBER:Out_Port},%{DATA:Src_MAC},%{DATA:Dst_MAC},%{DATA:ethernet_type},%{NUMBER:in_vlan},%{NUMBER:out_vlan},%{IP:Src_IP},%{IP:Dst_IP},%{NUMBER:IP_protocol},%{DATA:ip_tos},%{NUMBER:ip_ttl},%{NUMBER:Src_Port},%{NUMBER:Dst_Port},%{DATA:Tcp_flags},%{INT:packetsize:int},%{NUMBER:IP_size},%{NUMBER:sampling_rate}" ]
                }
                translate {
                        field => "IP_protocol"
                        destination => "protocols"
                        dictionary_path => "/home/translate/protocol"
                        refresh_interval => 86400
                }
                translate {
                        field => "Src_IP"
                        destination => "customer"
                        dictionary_path => "/home/translate/customerip"
                        refresh_interval => 86400
                        add_field => { "sfully" => "%{customer} %{Src_IP}" }
                }
                translate {
                        field => "Dst_IP"
                        destination => "customer"
                        dictionary_path => "/home/translate/customerip"
                        refresh_interval => 86400
```

阿里云 MVP

场景二：日志分析和性能监控

阿里云 MVP

为什么需要日志收集

- 掌握操作系统的运行状况
- 掌握应用运行中的详情、运行异常和业务日志输出
- 监控当前web服务的运行情况，根据关键字查询日志详情
- 等级保护的需要
- ....

阿里云 MVP

过去的困境

- 开发人员不能登录线上服务器查看详细日志
- 各个系统都有日志，日志数据分散，查找困难并且效率低下
- 磁盘空间不足删除了日志
- 没有实时监控、分析
- 无法提供全文搜索
- ....

阿里云 MVP

微服务其中一个服务日志

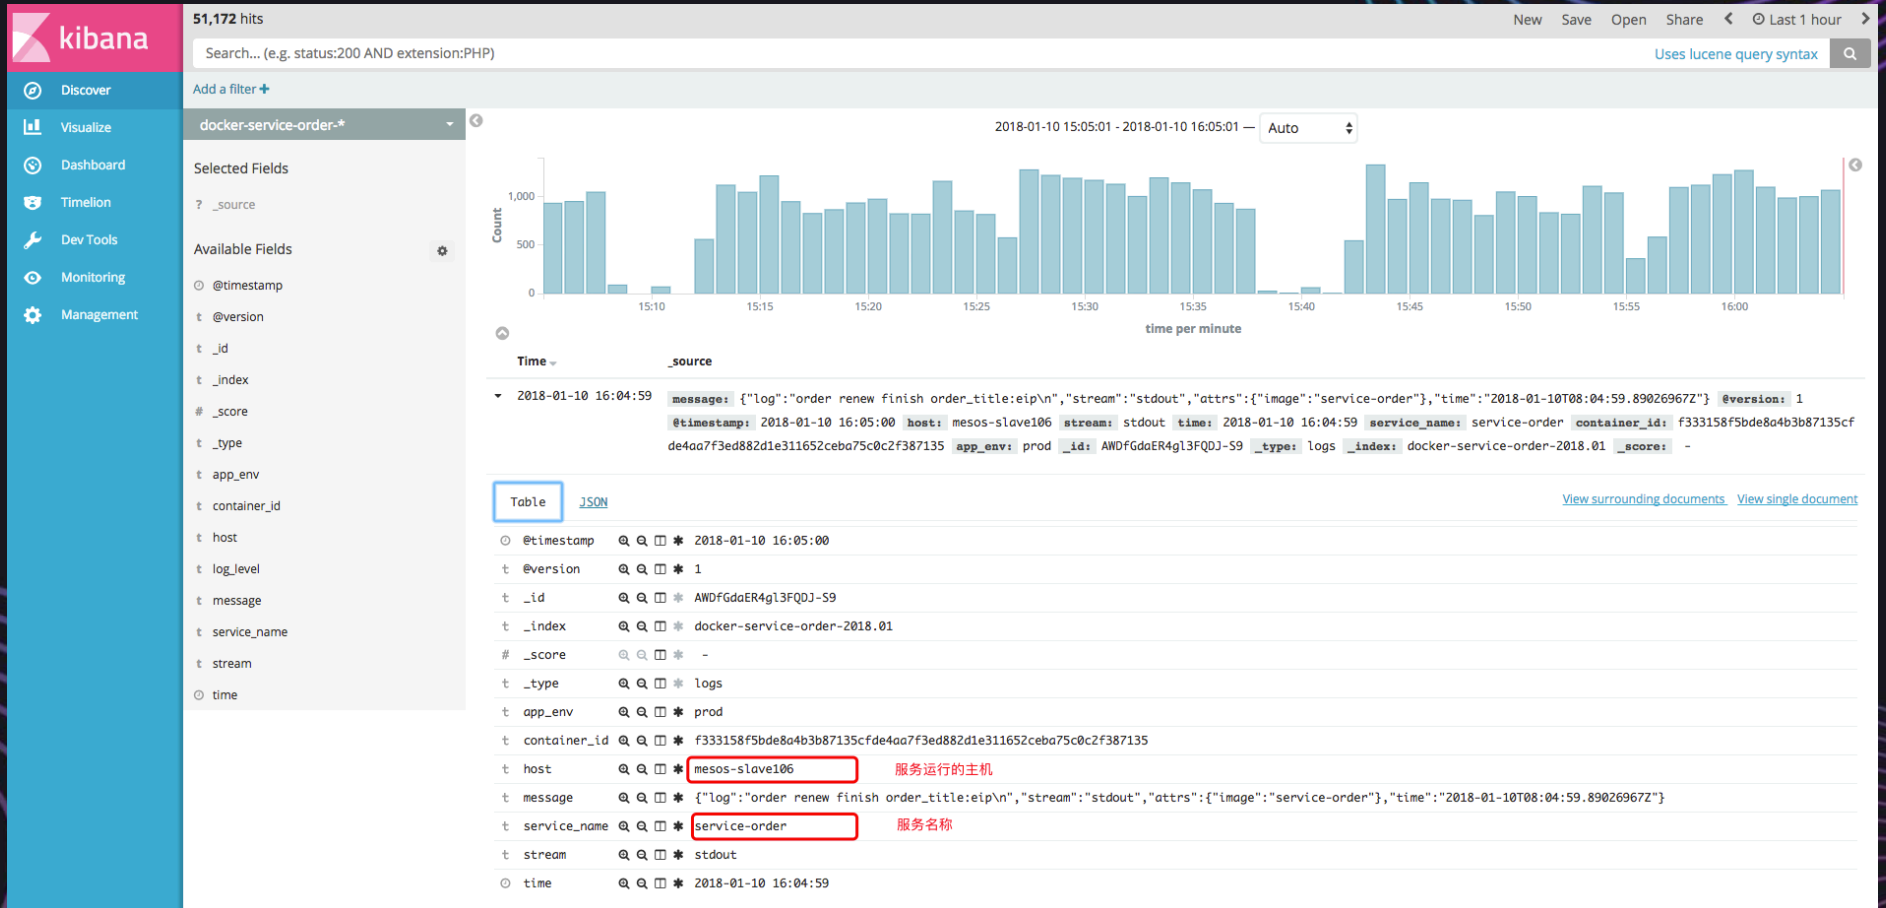

阿里云 MVP

Packetbeat收集mysql日志

kibana

- Discover
- Visualize
- Dashboard
- Timelion
- Dev Tools
- Monitoring
- Management
- Collapse

t port
t proc
t query
responsetime
t server
t type

Table JSON

View surrounding documents View single document

@timestamp	2017-12-01 23:59:59	
t _id	AWASzlY-M5YQLkgNxbo3	
t _index	packetbeat-mysql-2017.12	
# _score	-	
t _type	mysql	
t beat.hostname	mysql-master	
t beat.name	mysql-master	
t beat.version	5.4.0	
# bytes_in	472	
# bytes_out	14	
t client_ip	192.168.9.8	执行客户端
t client_port	50387	
t client_proc		
t client_server		
t direction	in	
t ip	192.168.9.32	MySql主机ip
t method	INSERT	
# mysql.affected_rows	1	
# mysql.error_code	0	
t mysql.error_message		
t mysql.insert_id	605501	
mysql.iserror	false	
t mysql.num_fields	0	
t mysql.num_rows	0	
t path		
t port	3306	
t proc		执行语句
t query	insert into Res_SysEmail(Title,EmType,Content,TemplateId,addResSee,BccEmail,RegularlySentTS,SentId,CreateUserId,CreateTs,isSend,CcEmail,OrderId,CustomerId,FailedSendem) values ('机柜分配状态异常（机柜是未分配状态但有分配记录','noreply','Id:1048机柜编号：I04 Id:8718机柜编号：I08 Id:13666机柜编号：J08 ',null,'wangcq@51idc.com;chengll@51idc.com',null,'2017-12-2 0:0:0',null,1,'2017-12-2 0:0:0',1,null,null,null,null)	
# responsetime	0	

阿里云 MVP

Nginx日志Dashboard

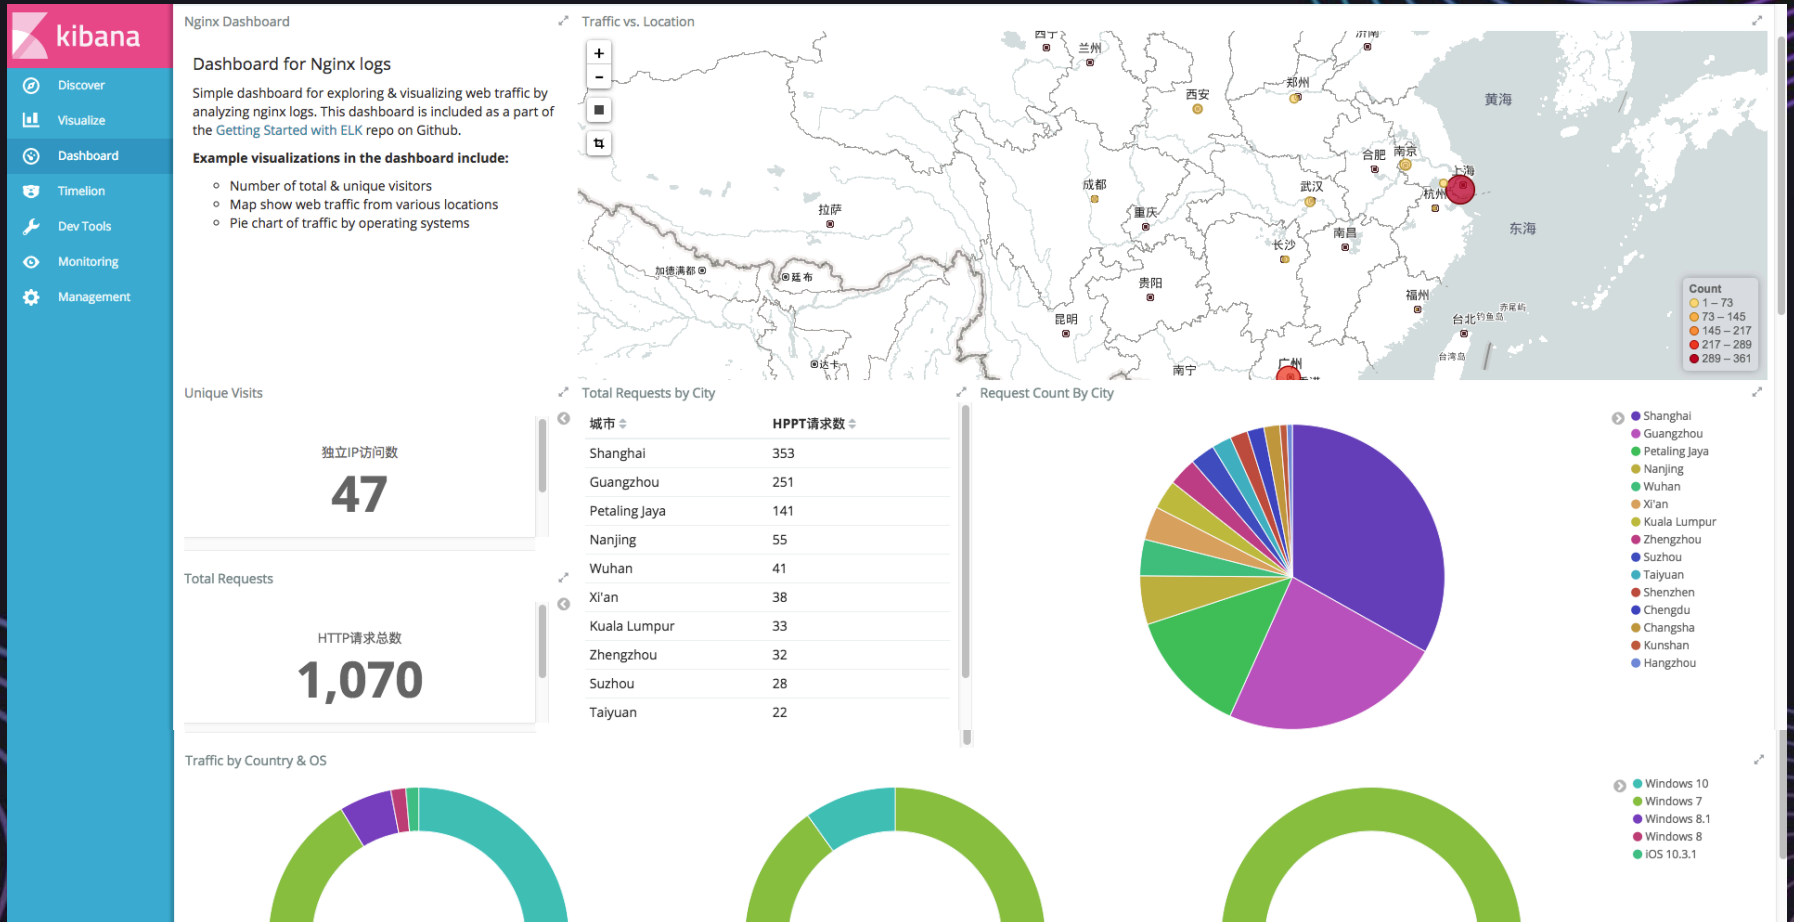

阿里云 MVP

X-Pack监控

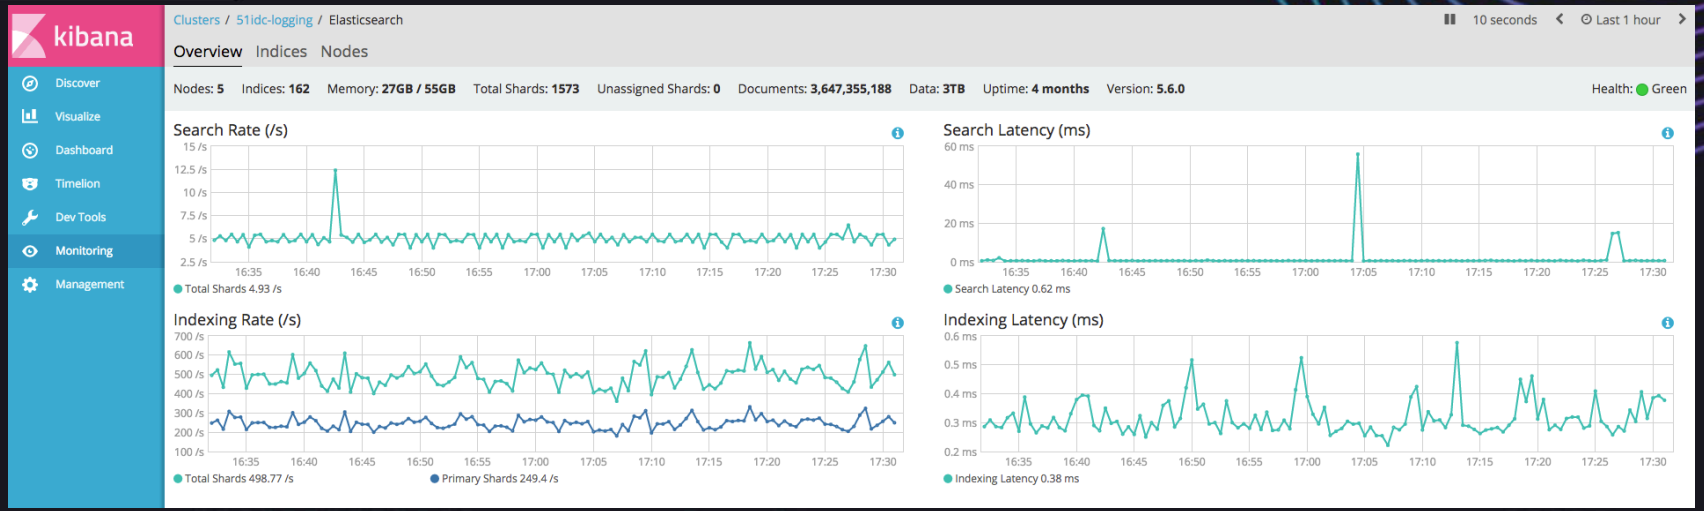

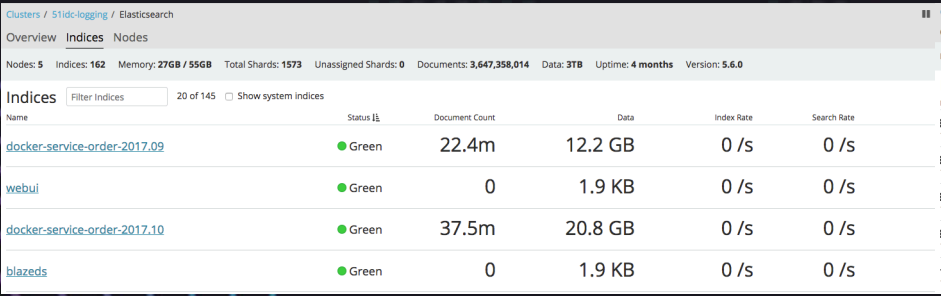

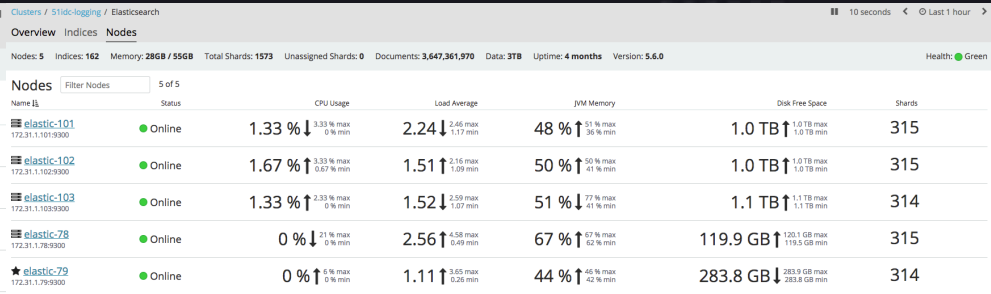

阿里云 MVP

技术方案

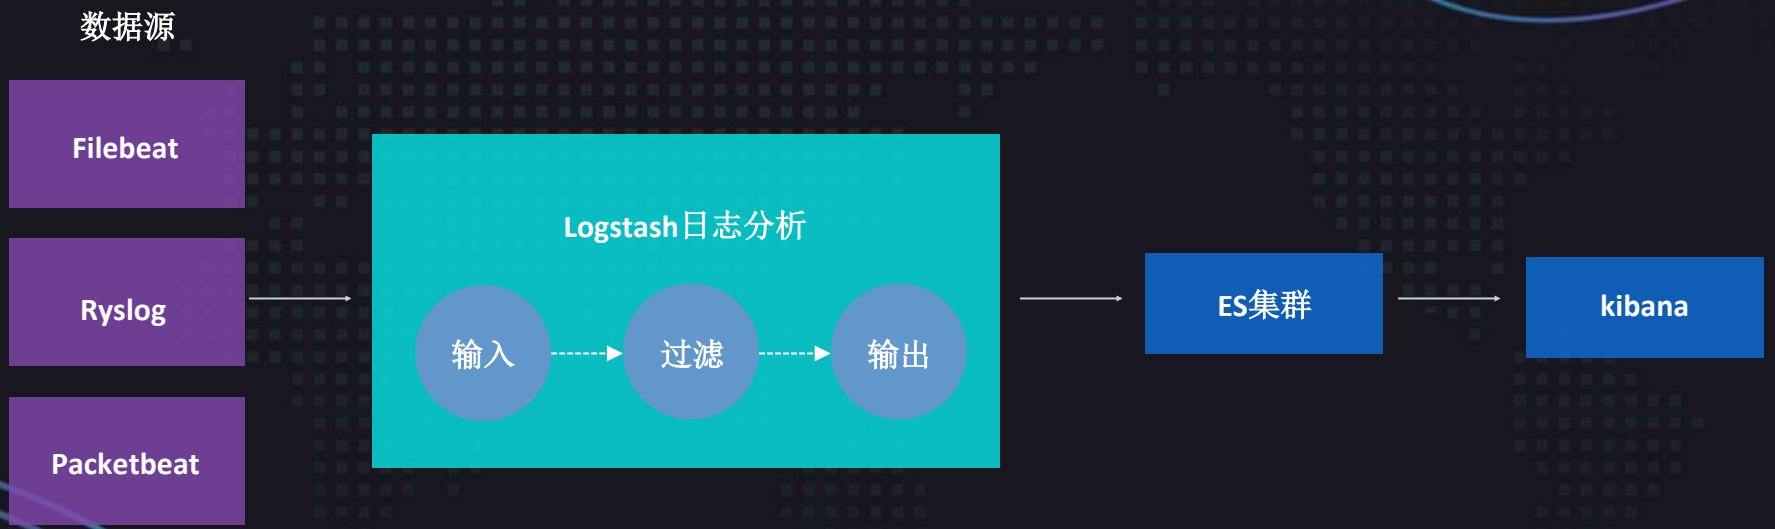
数据源
Filebeat
Ryslog
Packetbeat
Logstash日志分析
输入
过滤
输出
ES集群
kibana

阿里云 MVP

关键配置

```
input {
  file {
    path => ["/var/lib/docker/containers/*/*-json.log"]
    sincedb_path => "/root/.docker_logstash_sincedb"
    start_position => "beginning"
  }
}
filter {
  json{
    source => message
    add_field => { "service_name" => "%{[attrs][image]}" }
    remove_field => ["attrs"]
  }
  grok {
    match => ["log","\[(?<log_time>.*?)\]\s(?<log_level>.*?)\b\s(?<content>.*)"]
    match => ["path", "/(?<container_id>\w+)-json.log"]
    remove_field => ["path"]
    remove_field => ["log_time"]
  }
  mutate {
    remove_field => ["log"]
    rename => { "content" => "message" }
    add_field => { "app_env" => "prod" }
  }
}
output {
 if [service_name] != "%{[attrs][image]}"{
      elasticsearch {
        hosts => ["192.168.9.101:9200","192.168.9.102:9200","192.168.9.103:9200"]
        index => "docker-%{[service_name]}-%{+YYYY.MM}"
        workers => 20
        template_overwrite => true
      }
  }
}
```

```
filter {
    grok {
        match => {"message" => '%{IPORHOST:remote_ip} - %{DATA:user_name} \
[%{HTTPDATE:time}\] "%{WORD:request_action} %{DATA:request} HTTP/%{NUMBER:h
ttp_version}" %{NUMBER:response} %{NUMBER:bytes} "%{DATA:referrer}" "%{DATA
:agent}"'}
    }
    date {
        match => ["time", "dd/MMM/yyyy:HH:mm:ss Z"]
        target => "time"
    }
    geoip {
        source => "remote_ip"
        target => "geoip"
        database => "/opt/logstash-5.1.2/config/GeoLite2-City.mmdb"
    }
    useragent {
        source => "agent"
        target => "user_agent"
    }
    mutate {
        convert => { "bytes" => "integer" }
    }
}

output {
    elasticsearch {
        hosts => ["192.168.9.101:9200","192.168.9.102:9200","192.168.9.103:
9200","192.168.9.78:9200","192.168.9.79:9200"]
        index => "nginx-log"
        document_type => "log"
        workers => 5
        template => "/etc/logstash/nginx_template.json"
        template_name => "nginx-access"
        template_overwrite => true
     }
}
```

阿里云 MVP

为什么用ELK做监控

- 节约资源，监控、日志在一个平台搞定
- 降低学习成本，无需额外的学习另外的监控平台
- ELK社区比较活跃，看好ELK的前景
- ….

阿里云 MVP

Docker容器监控Dashboard

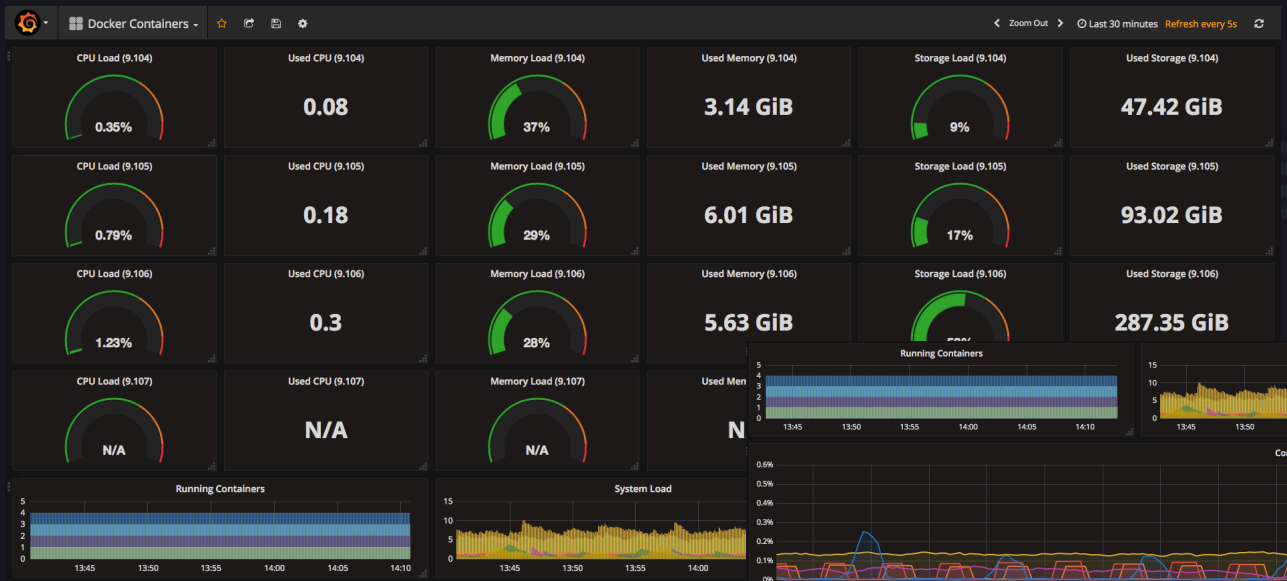

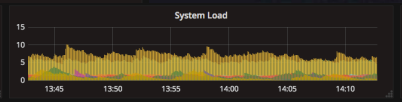

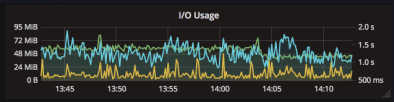

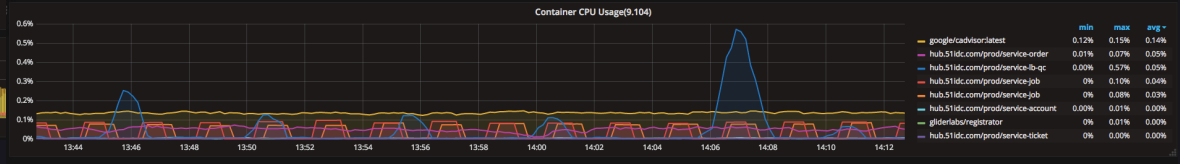

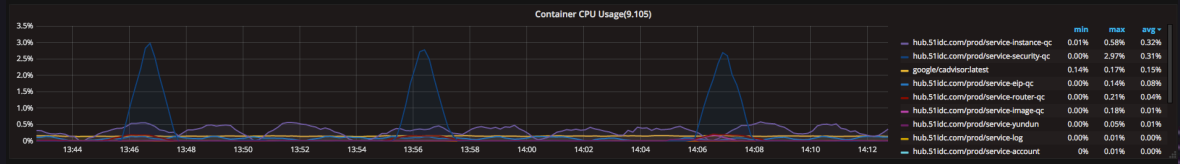

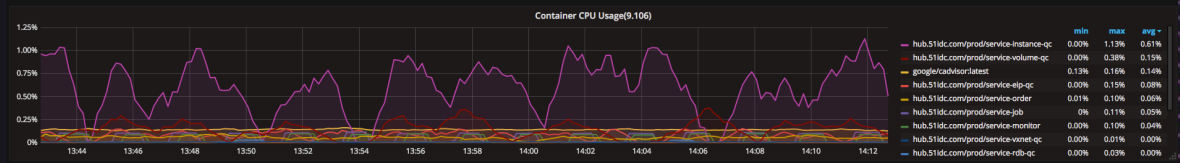

阿里云 MVP

主机监控Dashboard

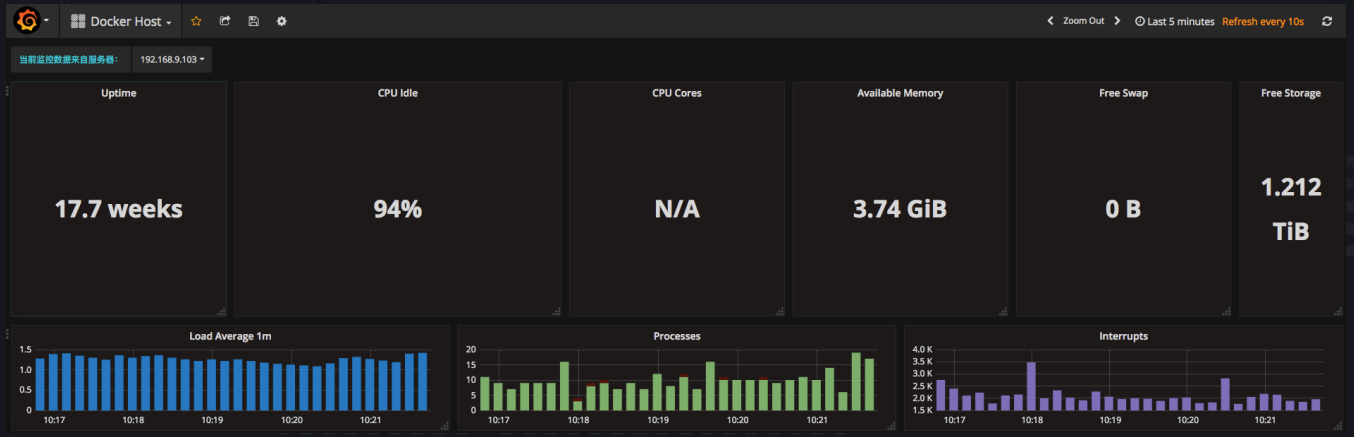

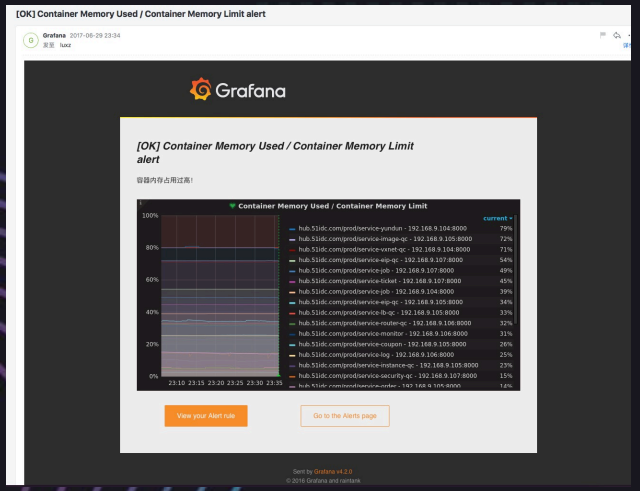

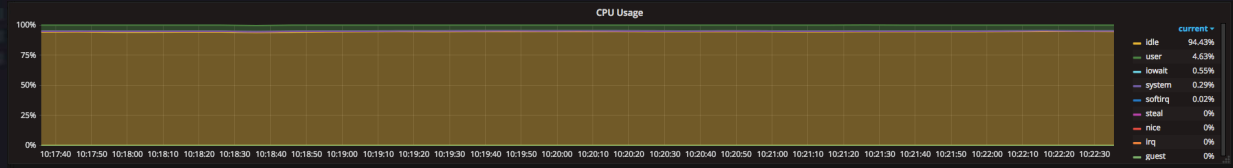

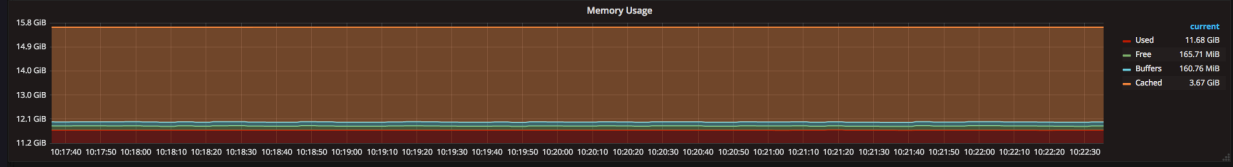

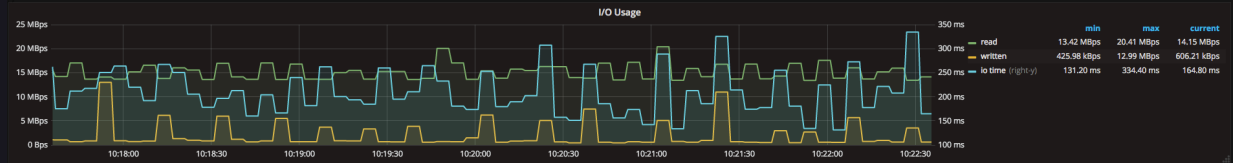

阿里云 MVP

性能指标采集工具

- Topbeat
- Metricbeat
- Collectd
- statsd
- ….

阿里云 MVP

技术方案

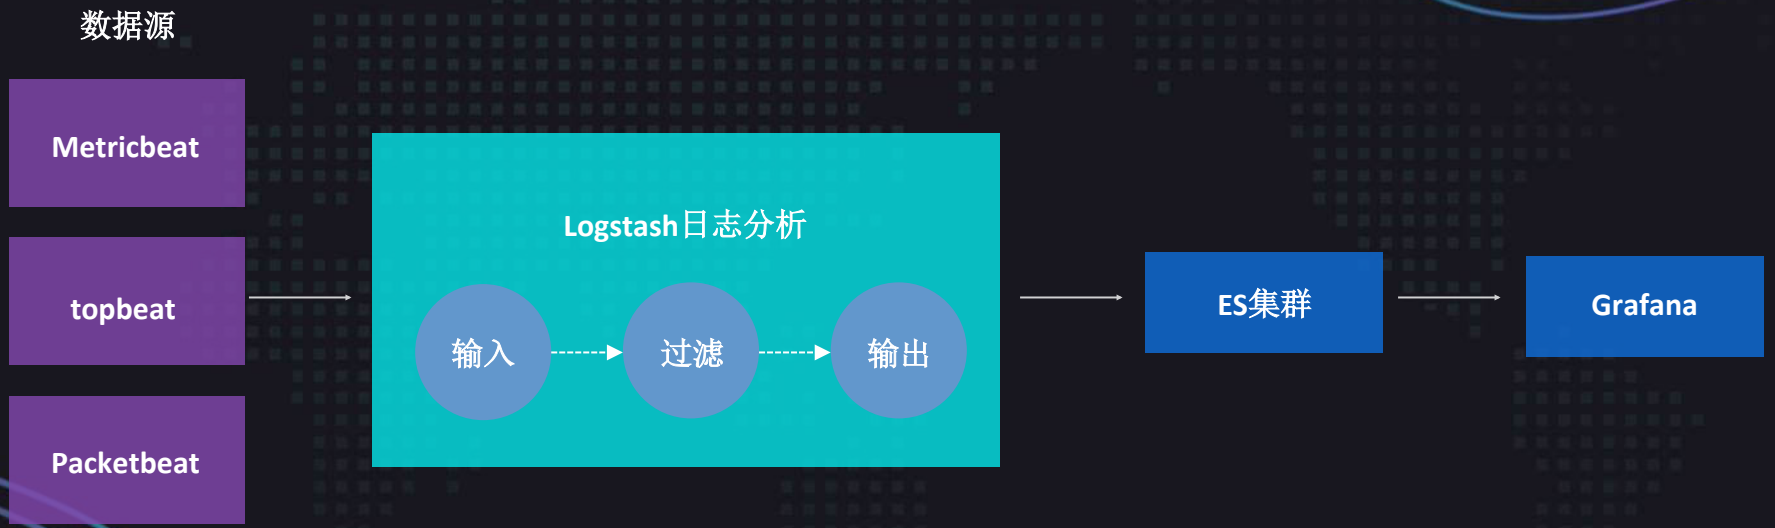
数据源
Metricbeat
topbeat
Packetbeat
Logstash日志分析
输入
过滤
输出
ES集群
Grafana

阿里云 MVP

ELK优势

- 开源、免费、上手容易
- 实时全文搜索
- 横向可扩展性
- 强大的灵活性
- 。。。

阿里云 MVP

ELK使用过程中遇到的问题

- 开源套件没有权限管理
- ES集群的监控告警
- 数据传输没有加密
- 日志审核
- 版本升级

Thanks!