IT大咖说
知识分享平台

ELK CI/CD 部署实践

黄鑫 (@dreampuf)

Agenda

使用 SaltStack 对 ELK 维护、配置管理、持续集成和持续部署的设计与实施。

- 背景介绍
- 可用性设计
- 配置管理和持续集成
- 持续部署
- 发现的问题和未来展望

背景介绍

- 应用场景
- 多部门协作
- 日志收集为主（Filebeat）
- 历史演进（1.75 -> 5.3.2）
- 使用 SaltStack 进行状态管理（安装，更新，维护）

应用场景

- 广告投放记录
- 访问日志
- 离线计算任务追踪
- 内部服务监控

广告投放记录

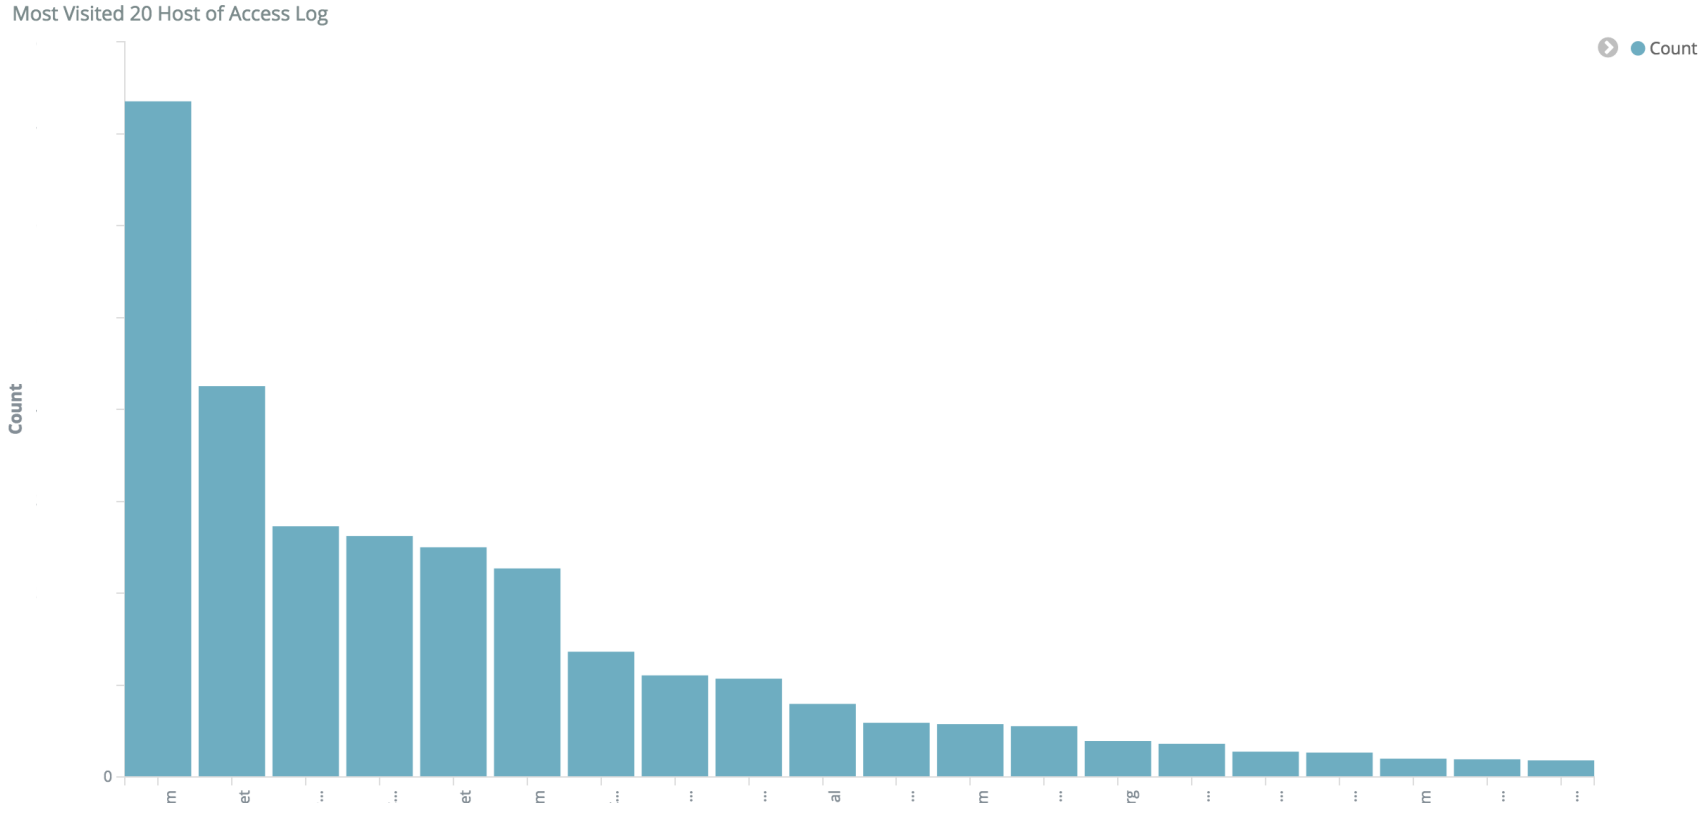

离线计算任务追踪

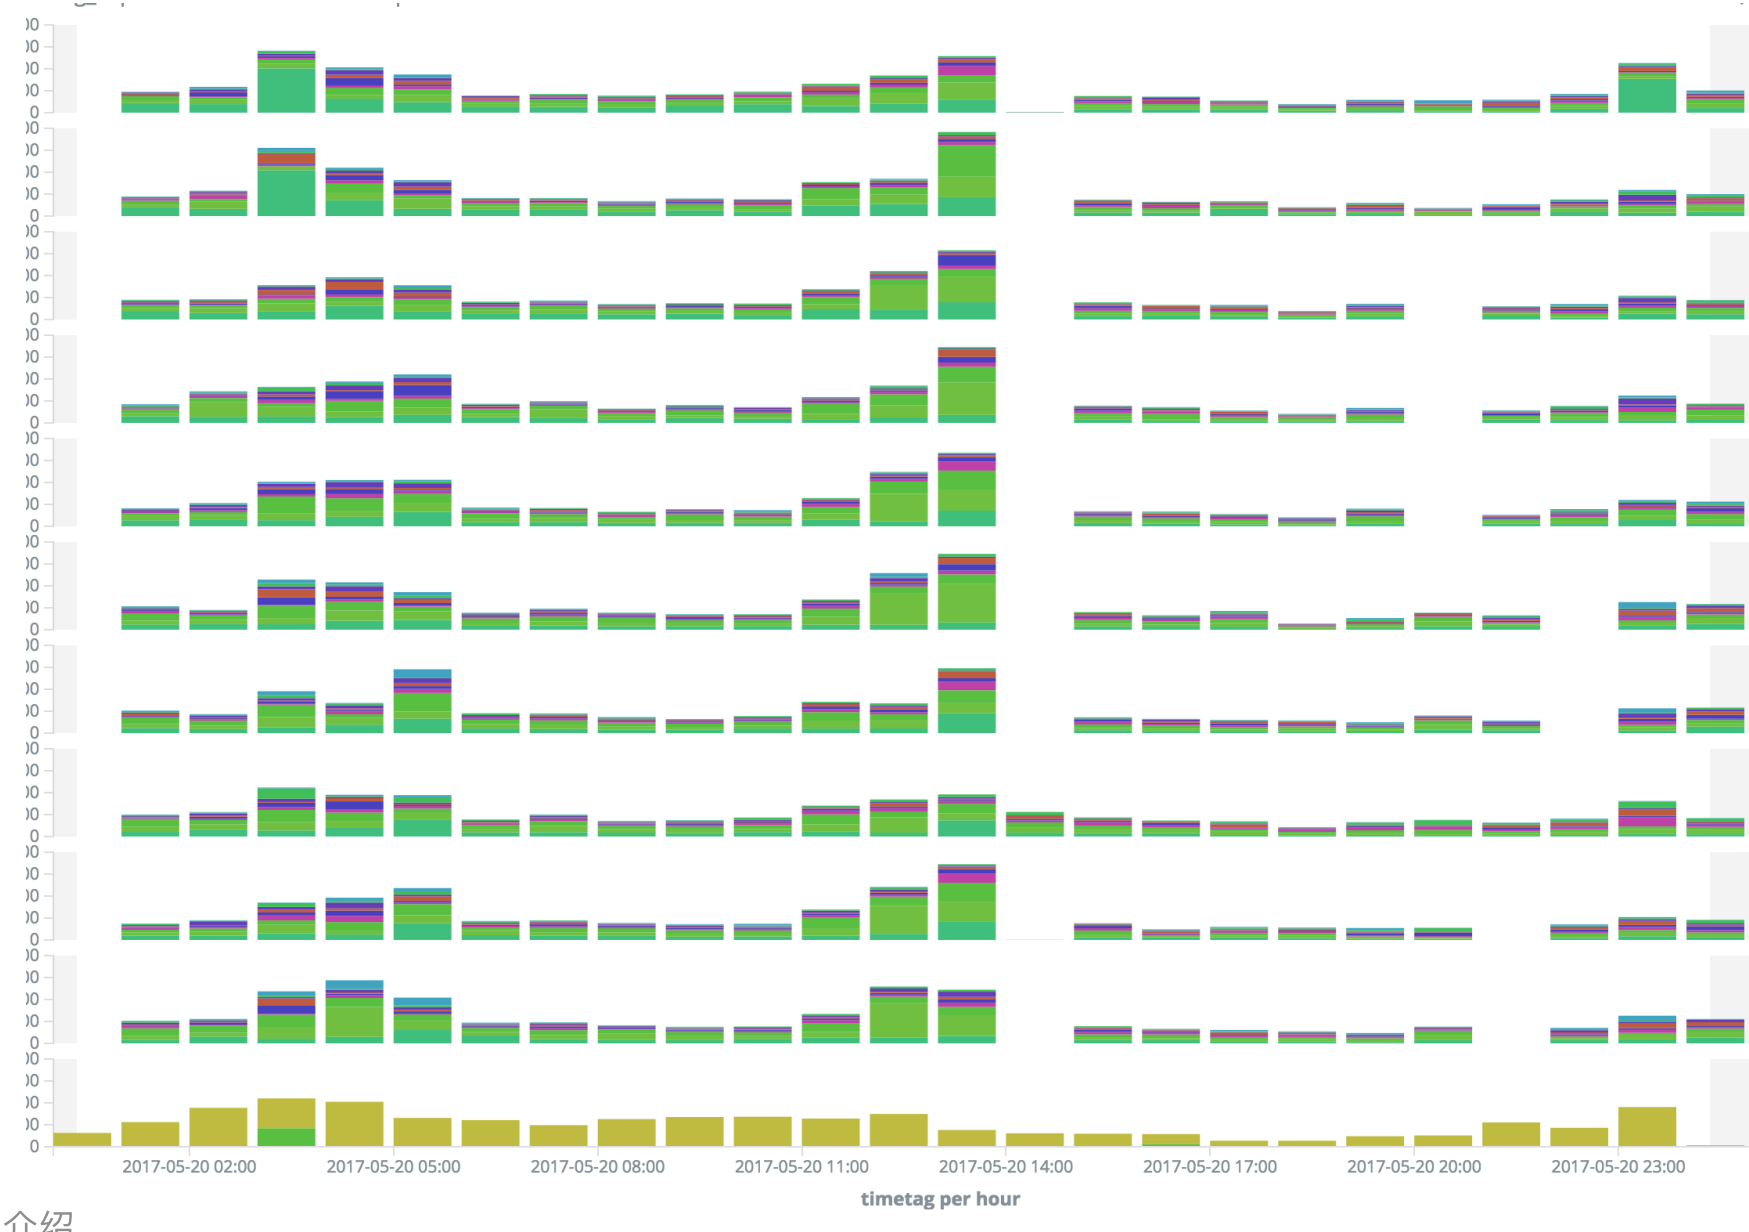
2017-05-20 02:00
2017-05-20 05:00
2017-05-20 08:00
2017-05-20 11:00
2017-05-20 14:00
2017-05-20 17:00
2017-05-20 20:00
2017-05-20 23:00
timetag per hour

多部门协作

- 前端、后端.....
- 部门之间相对隔离
- 数据存储统一，方便级联查询

数据源

- 文件日志格式收集为主（Filebeat）
- 少量网络数据（Packetbeat）
- 系统指标数据（Matricbeat/Collectd/NMon in Splunk）
- 容器日志（TBD: Kafka Driver）

架构演进

- 单机原型
- ES 独立
- 跨机房数据收集 - Redis
- ES 分离（Master/Data/Coordinating Node)

架构

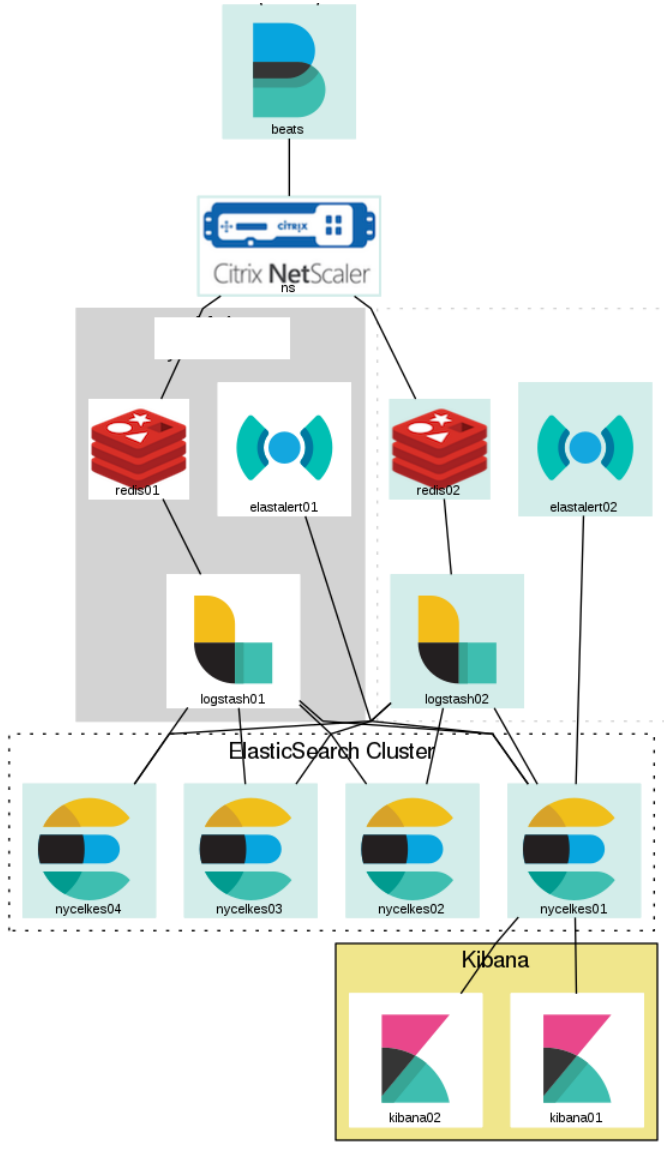
beats
Citrix NetScaler
ns
redis01
elastalert01
redis02
elastalert02
logstash01
logstash02
ElasticSearch Cluster
nycelkes04
nycelkes03
nycelkes02
nycelkes01
Kibana
kibana02
kibana01

配置管理和持续集成

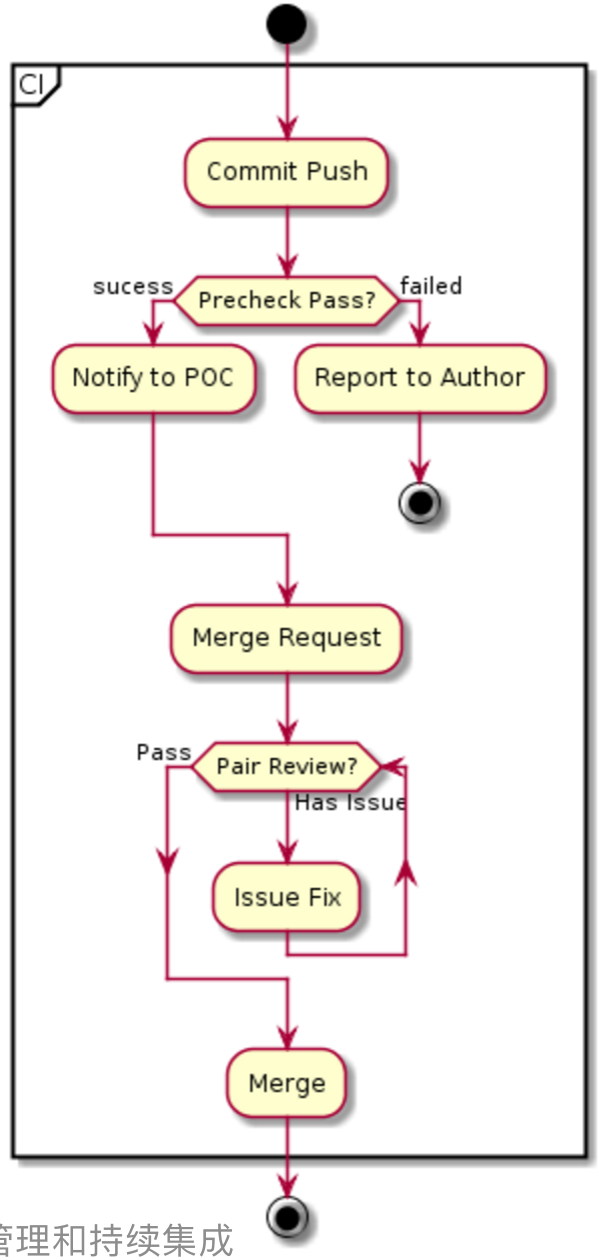
CI
Commit Push
sucess
Precheck Pass?
failed
Notify to POC
Report to Author
Merge Request
Pass
Pair Review?
Has Issue
Issue Fix
Merge

配置管理和持续集成

- 使用 Git 进行版本化
- Gitlab 作为持续集成套件

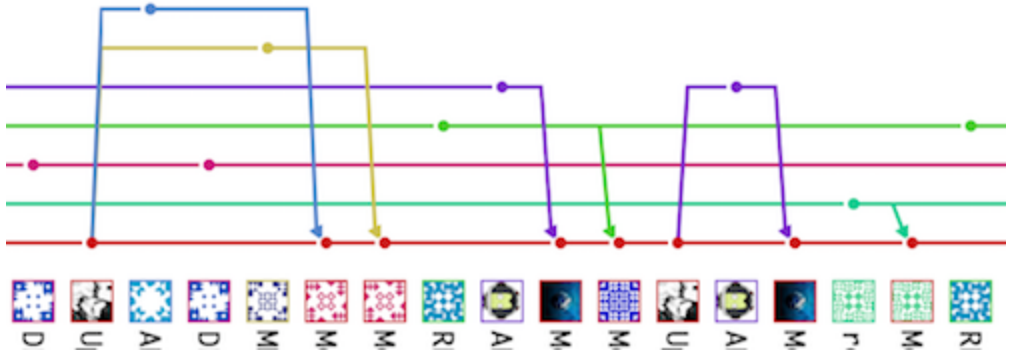

Git 进行配置版本化

```
|____config
|  |____prd
|  |  |____team_a
|  |  |  |____filebeat-a-debug-log.yml
|  |  |  |____filebeat-a-log-redis.yml
|  |  |____team_b
|  |  |  |____filebeat-c-mongo.yml
|  |  |  |____filebeat-c-redis.yml
|  |  |  |____filebeat-c-scheduler.yml
|  |  |  |____filebeat-c-mysql.yml
|  |  |  |____filebeat-c-postgresql.yml
|  |  |____team_c
|  |  |  |____filebeat-c-APP.yml
|  |  |  |____filebeat-c-APP.yml
|  |  |____team_d
|  |  |  |____filebeat-d-service1.yml
|  |  |  |____filebeat-d-service2.yml
|  |  |  |____filebeat-d-service3.yml
```

Gitlab-CI

- Feature branch workflow
- Gitlab-CI vs. Jenkins
- Precheck
  - `logstash -t`
  - `filebeat -configtest`

持续部署

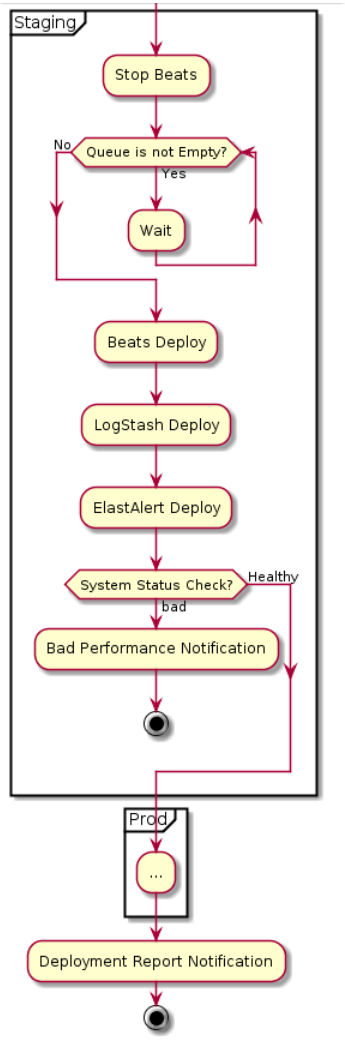
Staging
Stop Beats
No
Queue is not Empty?
Yes
Wait
Beats Deploy
LogStash Deploy
ElastAlert Deploy
System Status Check?
Healthy
bad
Bad Performance Notification
Prod
...
Deployment Report Notification

持续部署

- ELK 的状态管理
- ELK 自身监控
- 性能观察 (系统监控/x-pack monitoring)

状态管理

```
{% set appname = "metricbeat" %}

{{ appname }}:
  pkg.installed:
    - require:
      - pkgrepo: {{ repo_name }}
  file.managed:
    - name: /etc/{{ appname }}/{{ appname }}.yml
    - source: salt://fw-{{ appname }}/{{ appconfig }}
    - template: jinja
```

状态管理

```
{{ appname }}-service:
  service.running:
    - name: {{ appname }}
    - enbale: True
    - reload: True
    - require:
      - pkg: {{ appname }}
      - file: {{ appname }}-permission
      - file: {{ appname }}
    - watch:
      - file: {{ appname }}
```

状态管理

```
supervisor:
  pkg.installed: []
  file.managed:
    - name: /etc/supervisord.d/logstash.ini
    - contents: |+
        {% for path in paths %}{% set pname = get_filenam
        [program:{{ pname }}]
        command=logstash --path.settings /opt/logstash/et
                  -f {{ path }} --path.data /opt/logstash/
                  -n {{ grains['fqdn'] }}:{{ pname }}
        directory=/opt/logstash/
        user=logstash
        environment=LS_HEAP_SIZE=256m
        {% endfor %}
    - require:
      - file: logstash
```

状态管理

```
elasticsearch:
  file.managed:
    - name: /etc/elasticsearch/elasticsearch.yml
    - source: salt://es/elasticsearch-5.y/elasticsearch.yr
    - template: jinja
  cmd.run:
    - name: elasticsearch-plugin install x-pack
    - unless: elasticsearch-plugin list | grep x-pack
  service.running:
    - enbale: True
    - reload: True
    - require:
      - pkg: elasticsearch-pkg
      - file: elasticsearch
      - file: elasticsearch-link
    - watch:
      - file: elasticsearch
```

自身监控

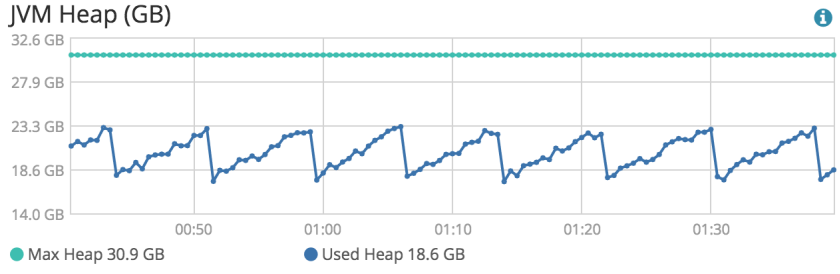
JVM Heap (GB)
32.6 GB
27.9 GB
23.3 GB
18.6 GB
14.0 GB
00:50
01:00
01:10
01:20
01:30
Max Heap 30.9 GB
Used Heap 18.6 GB

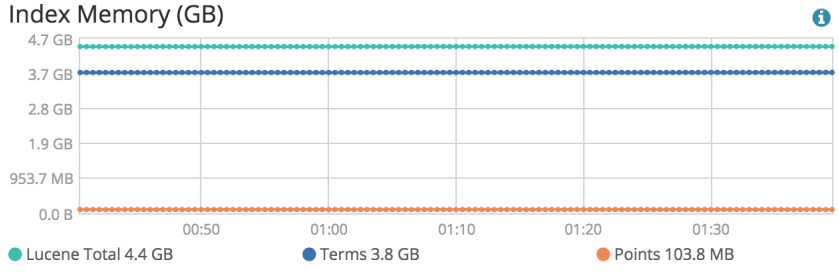
Index Memory (GB)
4.7 GB
3.7 GB
2.8 GB
1.9 GB
953.7 MB
0.0 B
00:50
01:00
01:10
01:20
01:30
Lucene Total 4.4 GB
Terms 3.8 GB
Points 103.8 MB

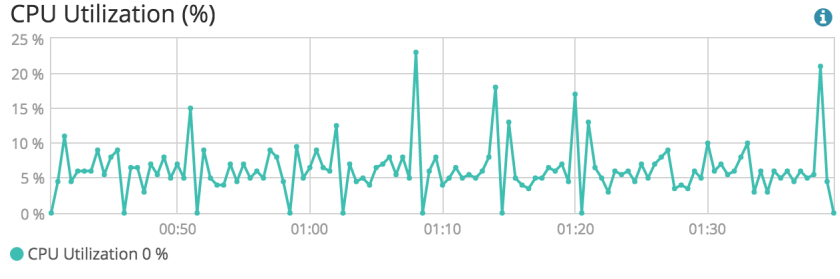
CPU Utilization (%)
25 %
20 %
15 %
10 %
5 %
0 %
00:50
01:00
01:10
01:20
01:30
CPU Utilization 0 %

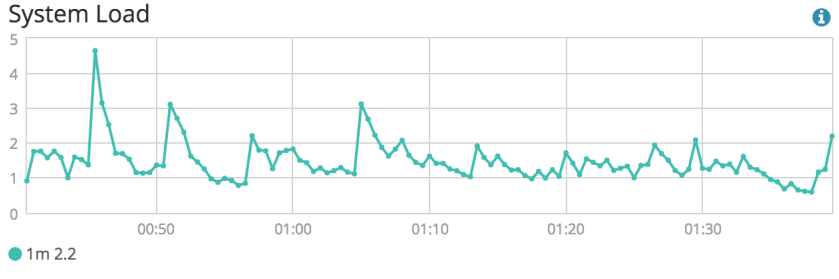
System Load
5
4
3
2
1
0
00:50
01:00
01:10
01:20
01:30
1m 2.2

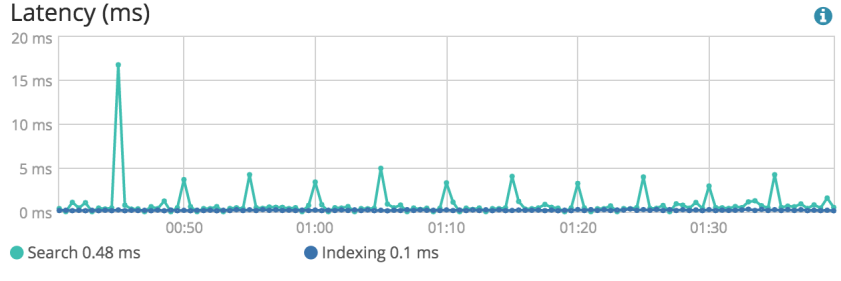
Latency (ms)
20 ms
15 ms
10 ms
5 ms
0 ms
00:50
01:00
01:10
01:20
01:30
Search 0.48 ms
Indexing 0.1 ms

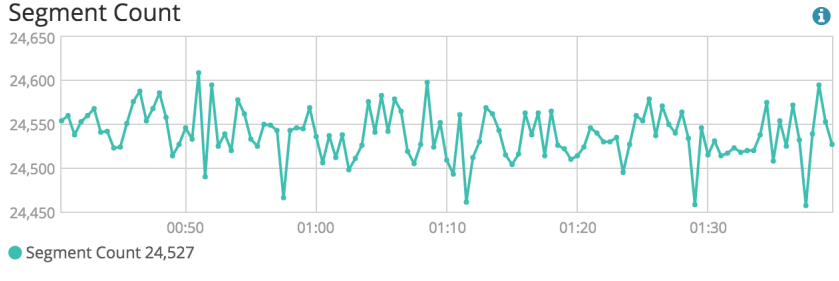
Segment Count
24,650
24,600
24,550
24,500
24,450
00:50
01:00
01:10
01:20
01:30
Segment Count 24,527

发现的问题

- Salt-API timeout
- 新 LogStash 配置没有启动
- ElastAlert
- Filebeat 对 RHEL5 的支持

Filebeat 对 RHEL5 的支持

```
diff -urpN output_orig/etc/init.d/filebeat output/etc/ini
--- output_orig/etc/init.d/filebeat
+++ output/etc/init.d/filebeat
@@ -27,8 +27,13 @@ pidfile=${PIDFILE-/var/run/filebeat.pi
-wrapper="/usr/share/filebeat/bin/filebeat-god"
-wrapperopts="-r / -n -p $pidfile"
+if grep -q -i "release 7" /etc/redhat-release; then
+    wrapper="/usr/share/filebeat/bin/filebeat-god"
+else
+    wrapper="/usr/share/filebeat/bin/filebeat-god-el5"
+fi
+beat_user="monitor"
+wrapperopts="-r / -n -p $pidfile -u $beat_user"
 RETVAL=0

 test() {
-       $agent $args $test_args
+       runuser -s /bin/bash $beat_user -c "$agent $args
 }
```

未来展望

- Kafka
- MetricBeat
- Container Monitoring
- ELK - Dockerize

Q&A