



OpsWorld 运维世界 大会·深圳站

基于Mesos和Docker 构建企业级SaaS服务


徐磊

大纲


1. Apache Mesos简介
2. Mesos在Qunar的应用
3. Stateful service
4. Elasticsearch as a Service

Apache Mesos

A distributed systems kernel




Apache Mesos



Apache Mesos

- 强大的二次开发能力
- REST API，支持所有语言



Apache Mesos


- Mesos通过DockerExecutor调用docker命令实现对容器的操作
- Framework通过编程方式指定使用DockerExecutor




Mesos在Qunar的应用

- 日志处理
 - 问题定位与回溯
 - 快速生成数据报表
- 流处理
 - 风控、推荐和ETL等

Mesos在Qunar的应用



Mesos在Qunar的应用



Mesos在Qunar的应用

- 有状态应用的尝试
 - 手工管理数据目录生命周期
 - 资源耗尽导致应用无法在含有数据节点rebuild
 - “约束”硬编码

Stateful service

stateless -> stateful

有状态服务才是更广泛的需求

Stateful service

- Mesos version < 0.23
- 状态如何落地？
 - Local filesystem + volume
 - Distribute filesystem
- 资源如何预留？
 - Mesos目前仅支持抢占式调度
- 动态划分资源？
 - Static Reservation

Stateful service

- Mesos version ≥ 0.23
- Persistent Volume
- Dynamic Reservation
- Quota
- Oversubscription

Stateful service

- 选择Elasticsearch作为stateful改造的切入点
 - Eat your own dog food
 - 业务线大量使用ES支撑系统
- 目标
 - 即申请即用
 - 整合流处理系统，结果自动写入业务线的Elasticsearch
 - 统一管理、规范运维

Elasticsearch as a Service

Esaas

WENJIA

ops_event_log_qes

集群概况

集群配置

操作日志

ops_watcher_events_qes

状态: Health

概况

名称	ops_event_log_qes
版本	v2.3.3
http端口号	12101
transport端口	12100
node数量	4
shards总量	62
访问地址	
配置地址	
dashboard	

计费

Time	0
Memory	0
Storage	0
SSD	0
High availability	0
Hourly rate	0
Monthly rate	0

masternode

更多masternode

datanode

更多datanode

Dashboards

kopf

head

bigdesk

Elasticsearch as a Service

Esaas wen.ma

ops_event_log_qes
集群概况
集群配置
操作日志
ops_watcher_events_qes

状态: Health

配置 打开 关闭

插件安装

delete-by-query 安装

Kibana

4.5.4 安装

Replica

该设置为全局设置，集群所有的index都会被修改，谨慎操作！线上备份数线上设置应该大于等于1(replicas>=1)

数量 1 配置

自助配置

/_cluster/settings 配置


方法: PUT

```
1- {  
2-   "transient": {},  
3-   "persistent": {  
4-     "discovery.zen.minimum_master_nodes": 2  
5-   }  
6- }
```

Elasticsearch as a Service


- 调研了一些产品：
 - Elasticsearch on Mesos（Cisco）
 - Elastic Cloud（官方）
 - AWS Elasticsearch
- 结合自身情况：
 - Elasticsearch on Mesos via Marathon

Elasticsearch as a Service




Elasticsearch as a Service

- 按照业务线划分资源
 - Multiple role based framework(MESOS-1763)
 - Sub-marathon



Elasticsearch as a Service




The screenshot displays the Marathon web interface. The top navigation bar includes the Marathon logo, tabs for 'Applications' and 'Deployments', a search bar labeled 'Search all applications', and buttons for 'Create Group' and 'Create Application'. The left sidebar contains filters for 'STATUS' (Running: 5, Deploying, Suspended, Delayed, Waiting) and 'HEALTH' (Healthy: 4, Unhealthy, Unknown: 1). The main content area shows a list of applications under the breadcrumb 'Applications > sec-log-qes'. The table lists applications with columns for Name, CPU, Memory, Status, Running Instances, and Health. All listed applications are in a 'Running' state with green health bars.

Name	CPU	Memory	Status	Running Instances	Health
bamboo-haproxy	0.5	256 MiB	Running	1 of 1	Healthy
es-dashboard	0.3	100 MiB	Running	1 of 1	Healthy
es-datanode	12.0	36 GiB	Running	3 of 3	Healthy
es-master	3.0	6 GiB	Running	3 of 3	Healthy
es2graphite	0.5	50 MiB	Running	1 of 1	Healthy

Elasticsearch as a Service

- Elasticsearch内部服务发现
 - 利用HAproxy做了一个hack
 - 仅在启动集群时生效




Elasticsearch as a Service

- Elasticsearch data/plugin dir
 - mount --bind (Mesos Persistent Volume)
- Elasticsearch log dir
 - Mesos sandbox + logrotate module
- Elasticsearch部署策略
 - host unique + rack unique
- Metrics
 - pyadvisor + es2graphite

Elasticsearch as a Service



Elasticsearch as a Service



Elasticsearch as a Service

- 延迟
 - Kafka Topic Lag
- 流量
 - Search/Docs Count
 - Kafka Message Count
- 错误
 - Reject
- 饱和度
 - IO使用率、CPU使用率

谢谢