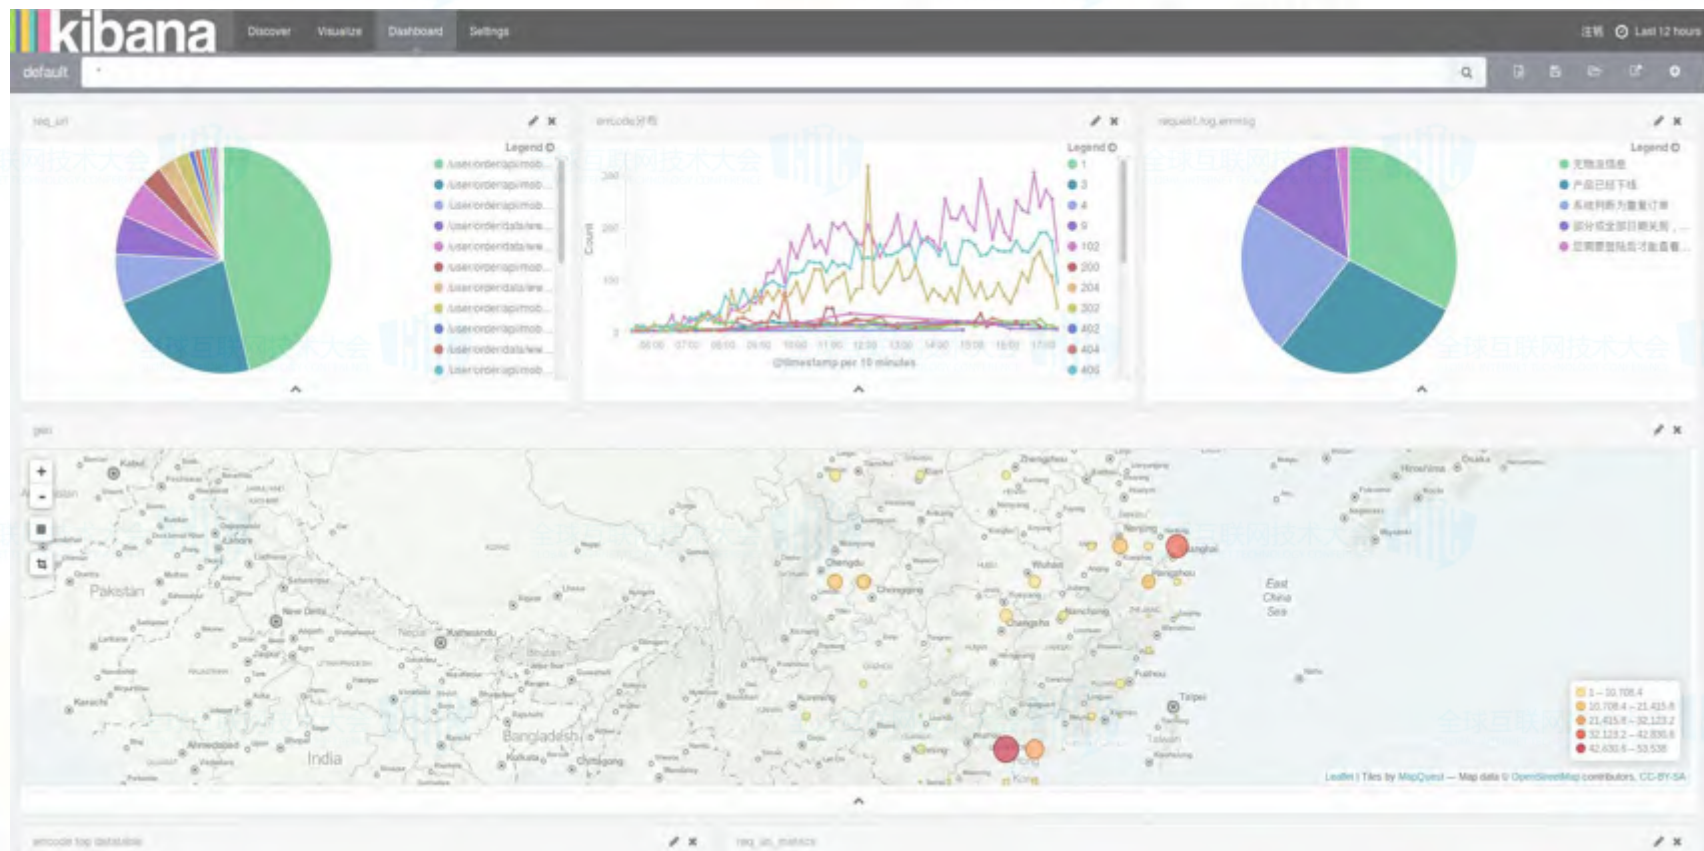


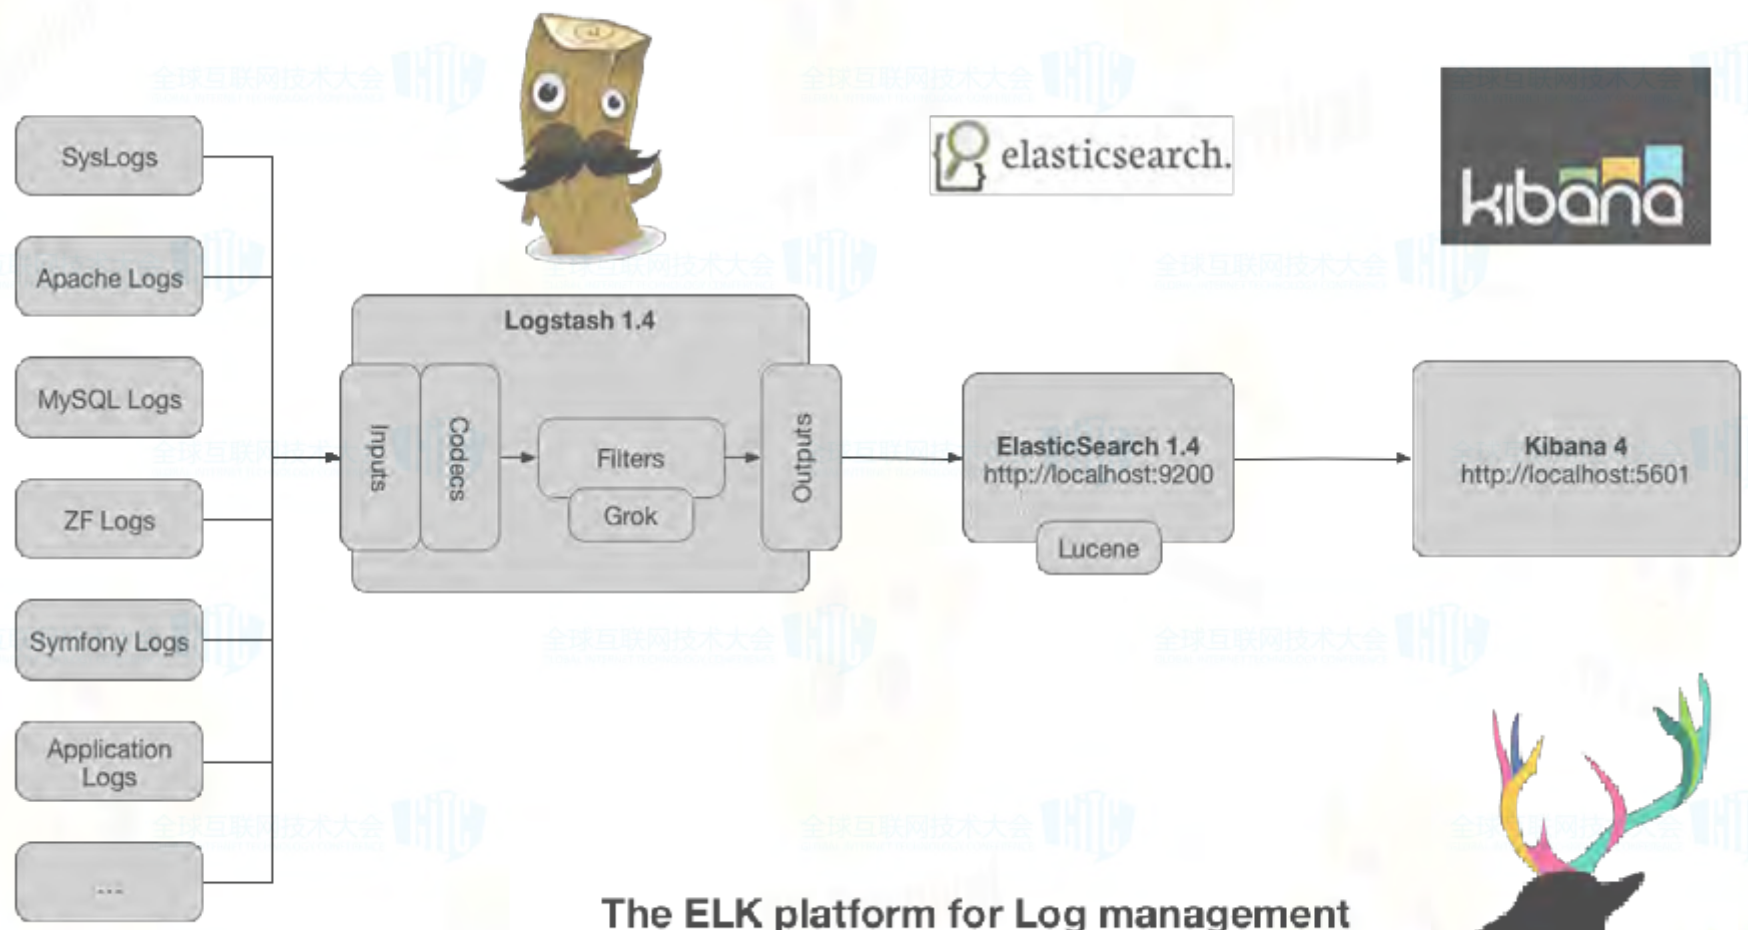


微信扫一扫

Mesos在Qunar的应用

2015年Q2——ELK





The ELK platform for Log management





解决问题定位的瓶颈





面临的问题

- 快速构建业务
 - 自动化
- 快速扩容
 - Scale up
 - Scale out
- 业务隔离
 - 数据隔离
 - 访问隔离
 - 资源隔离



方案



docker



MARATHON



MESOS



2015年Q4——流处理


Spark



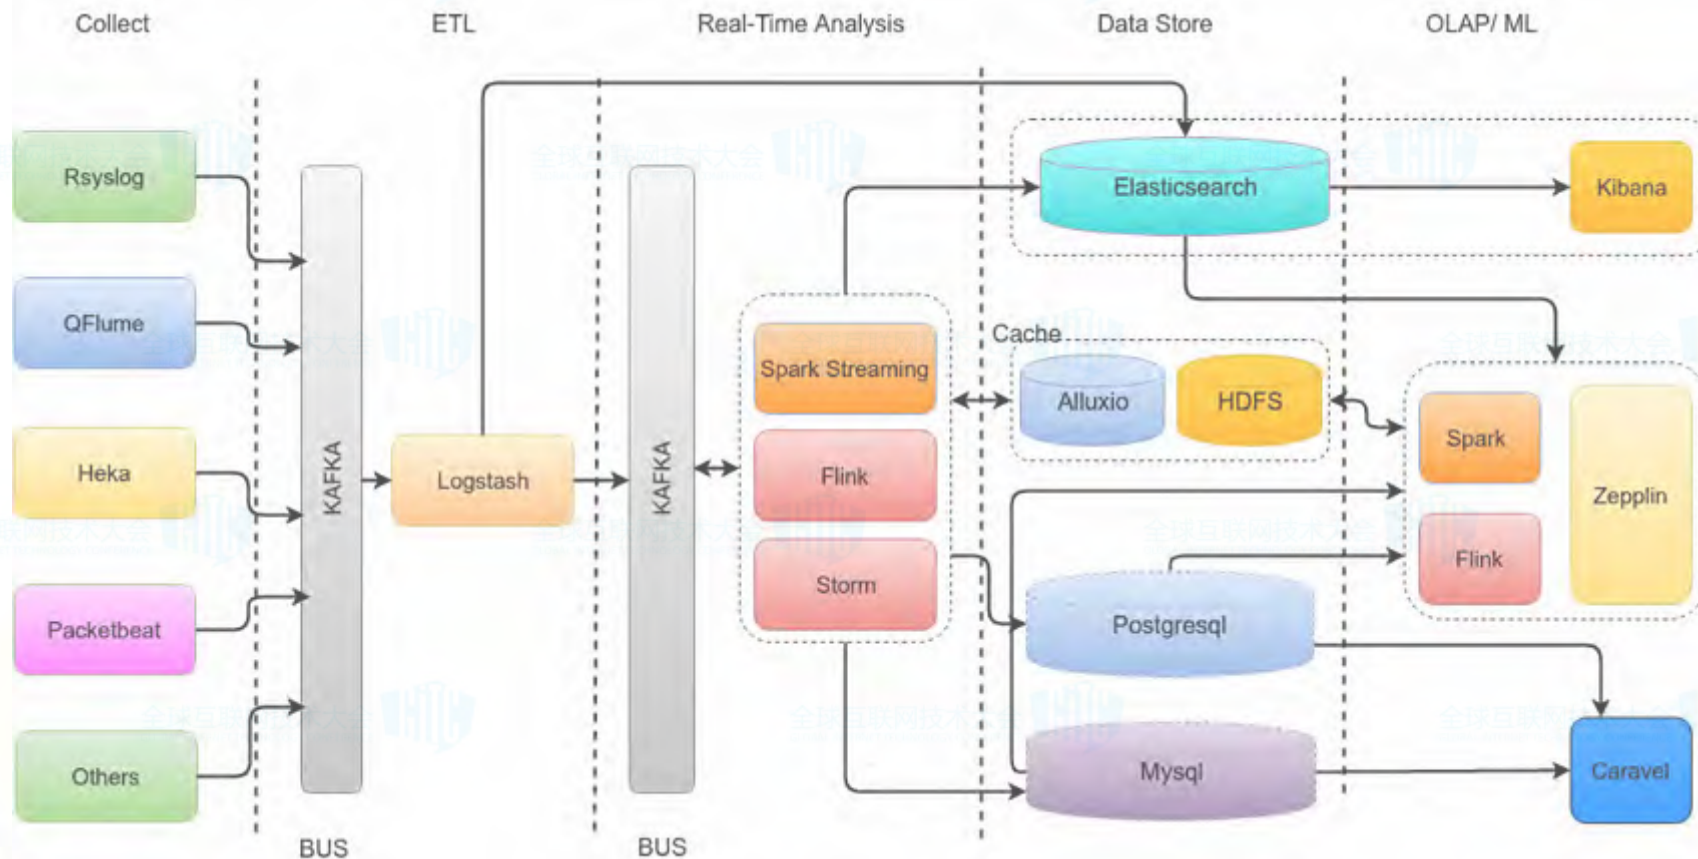
STORM



Flink



实时流系统





微信扫一扫

stateless -> stateful

有状态服务才是更广泛的需求



Mesos的新特性

- Mesos > 0.27
 - Persistent Volume
 - Dynamic Reservation
 - Quota
 - Oversubscription
- Marathon ~ 1.0+



2016年——ES as a Service

The screenshot displays the Esas console interface for a cluster named 'ops_event_log_qes'. The status is 'Health'. The configuration table is as follows:

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

Additional details shown include '计费' (Billing) with zero usage for Time, Memory, Storage, SSD, High availability, Hourly rate, and Monthly rate. Node types 'masternode' and 'datanode' are listed with redacted IP addresses. A 'Dashboards' section at the bottom shows thumbnails for 'kib' and 'logstash'.



为什么选择ES?

- Eat your own dog food
 - 我们的实时平台内部就有两套ES集群
- 业务线大量使用ES支撑系统
 - 维护麻烦
 - 扩容麻烦
 - 资源利用率
 - 分散精力，没办法专注于业务

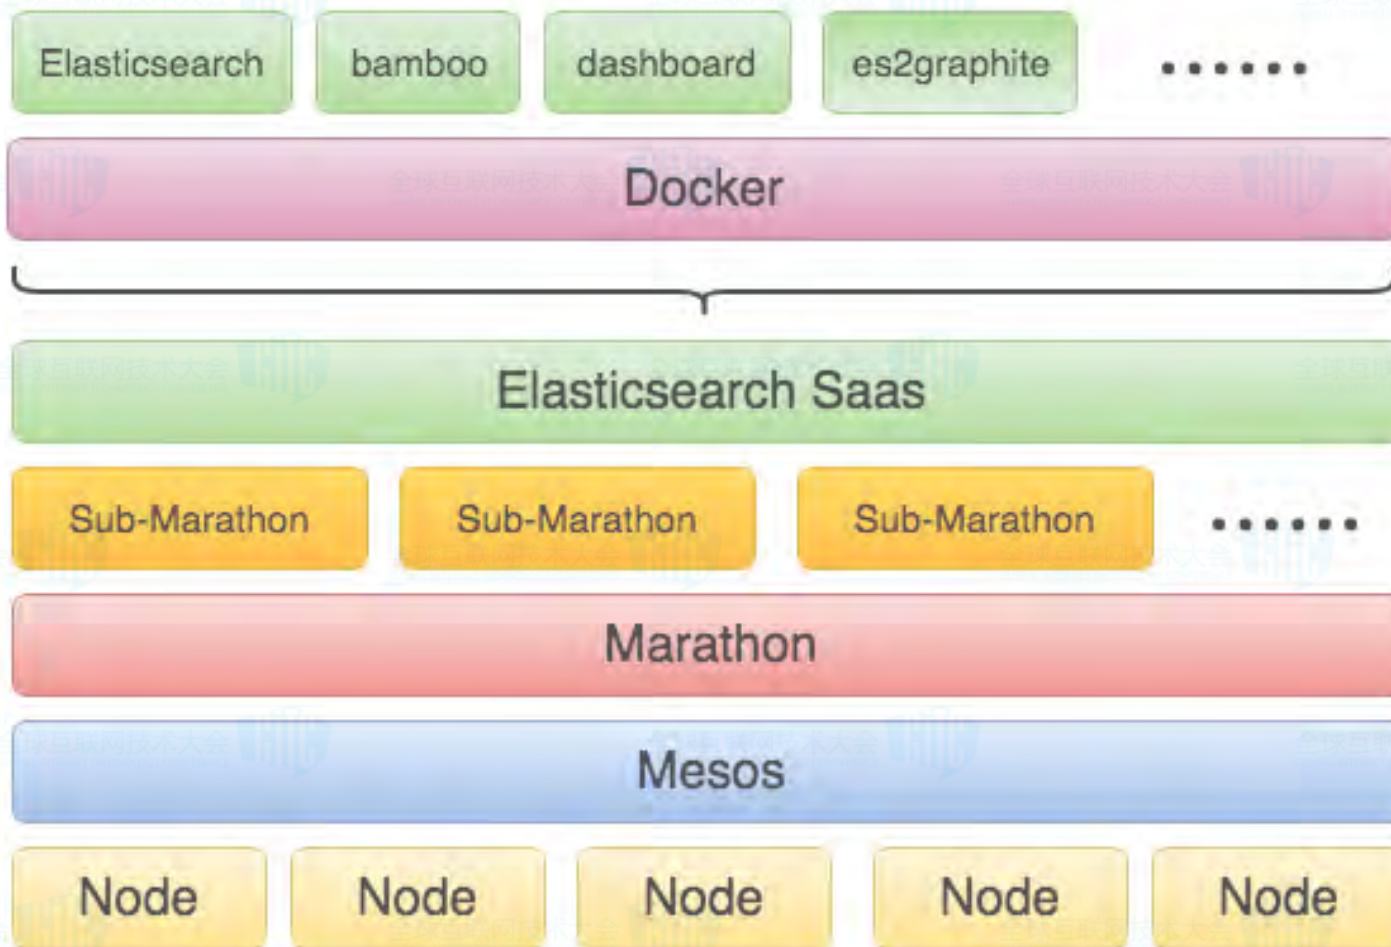


选型

- Elasticsearch on Mesos (Cisco)
- Elastic Cloud (官方)
- AWS Elasticsearch
- Elasticsearch on Mesos via Marathon

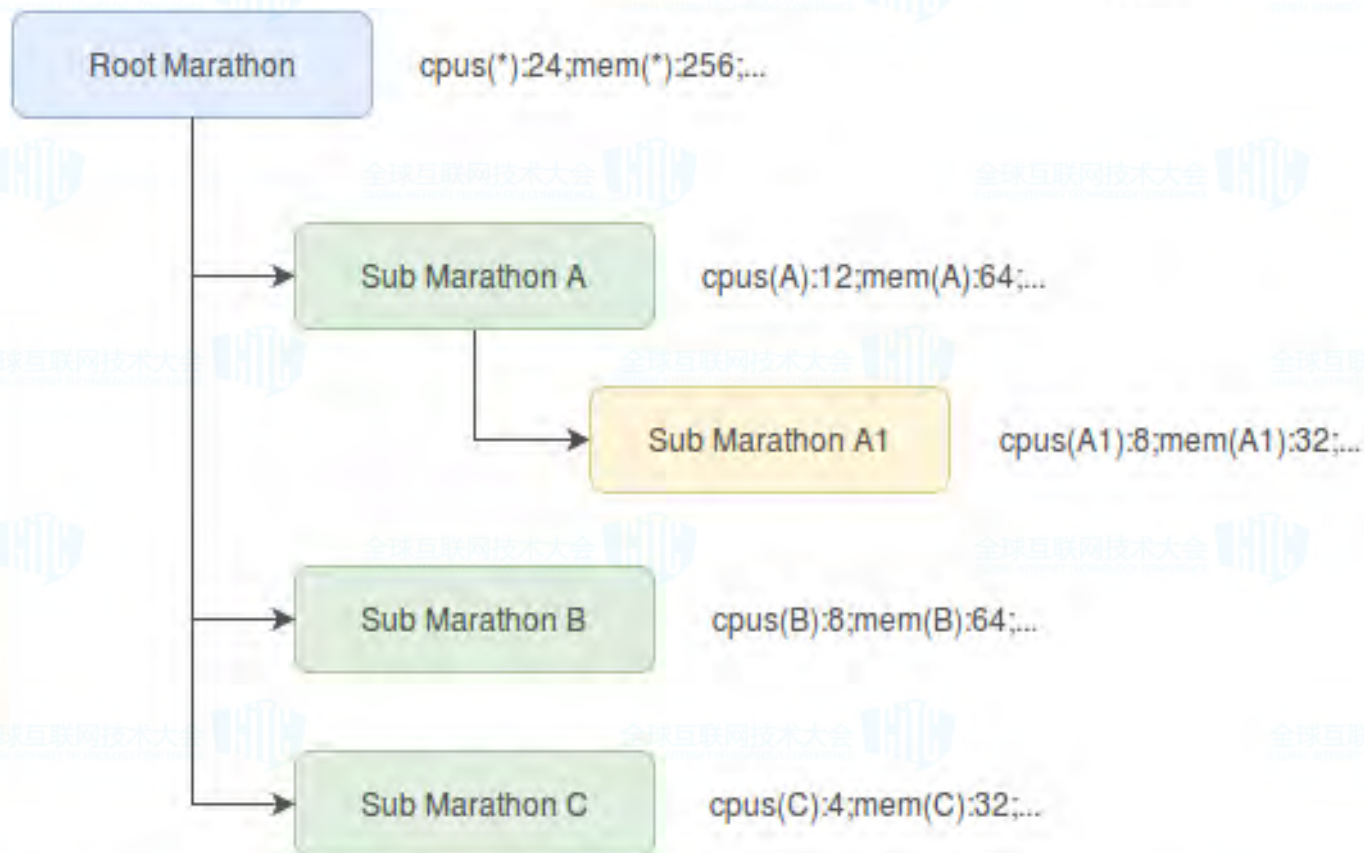


Esaas结构





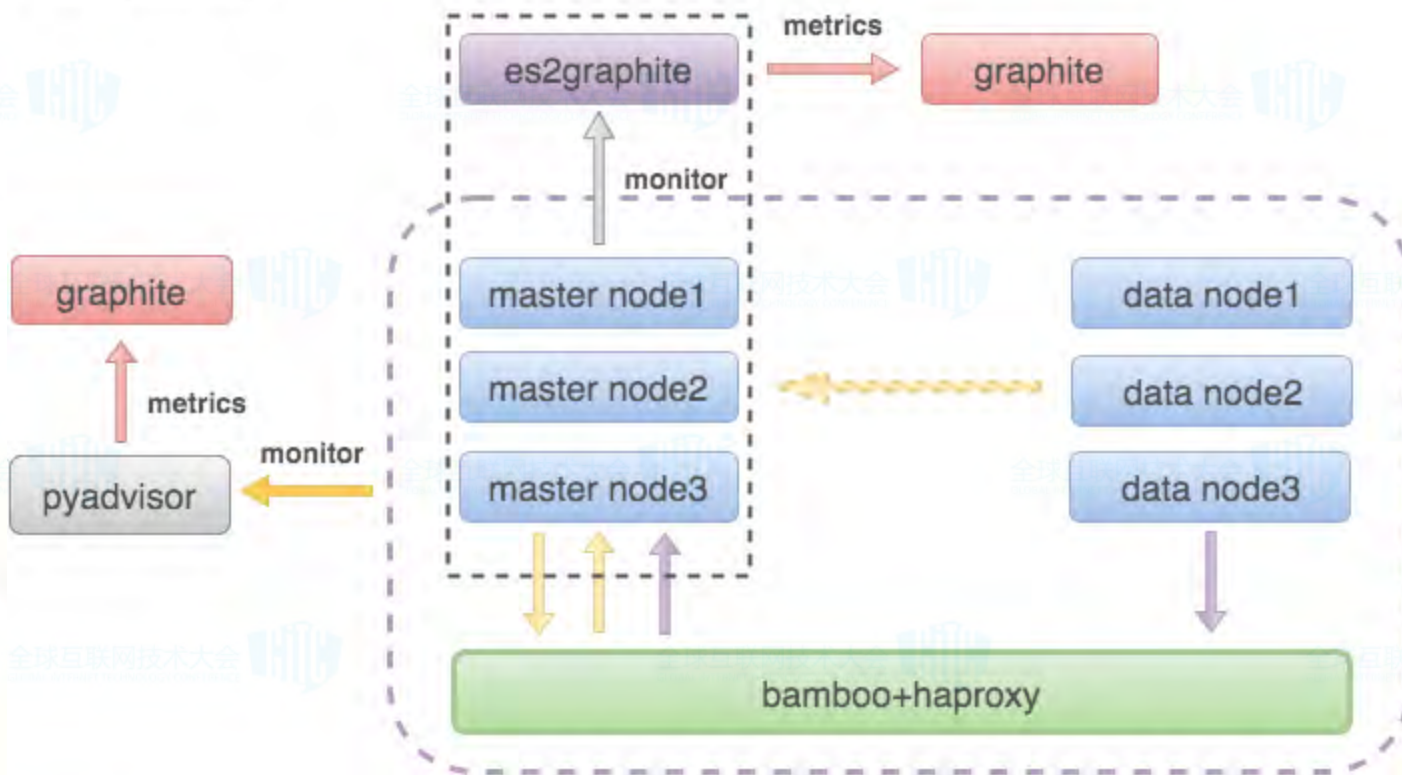
Quota





ES内部服务发现

Sub-Marathon





运行时维护

Esaas

wen.ma

ops_event_log_qes

状态: Health

配置 打开 关闭

集群概况

集群配置

操作日志

ops_watcher_events_qes

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- }
    
```



运行时维护

The screenshot displays the Marathon web interface. At the top, there's a search bar for applications and buttons for 'Create Group' and 'Create Application'. The main content area is divided into 'STATUS' and 'HEALTH' sections. The 'STATUS' section shows a list of applications with columns for Name, CPU, Memory, Status, Running Instances, and Health. The 'HEALTH' section shows a summary of application health, including Healthy, Unhealthy, and Unknown counts.

STATUS	Count	Name	CPU	Memory	Status	Running Instances	Health
Running	5						
Deploying							
Suspended							
Delayed							
Waiting							
HEALTH							
Healthy	4						
Unhealthy							
Unknown	1						
RESOURCES							
Volumes	3						



OPS Trace

Esaas

wen.ma

ops_event_log_ops

状态: Health

系统概况

系统配置

操作日志

ops_watchr_events_ops

操作日志

Q 请输入搜索内容

时间	用户	变更url	变更内容
2016-10-09 12:28:34	wen.ma	/login	login
2016-10-09 12:15:04	wen.ma	/login	login
2016-09-27 10:51:08	songkuan.zheng	/login	login
2016-09-23 10:37:06	wen.ma	/install_plugin	install error: install kibana error: ["message":"An app with id [ops-event-log-queskibana] already exists."]
2016-09-21 11:35:39	tiantian.lian	/login	login
2016-09-18 19:49:48	xulei.xu	/login	login
2016-09-12 19:17:18	congxiue.zhang	/login	login
2016-09-06 16:00:43	yun.huang	/login	login
2016-09-06 15:41:05	wen.ma	/edit_switch	edit switch status: True
2016-09-02 14:27:32	yun.huang	/login	login
2016-09-02 12:23:16	wen.ma	/login	login
2016-08-22 16:28:13	wen.ma	/edit_switch	edit switch status: False
2016-08-22 16:25:09	xuxiao.xu	/login	login
2016-08-22 16:25:06	xueyan.li	/login	login
2016-08-22 16:24:20	wen.ma	/edit_switch	edit switch status: True
2016-08-22 08:16:02	wen.ma	/edit_switch	edit switch status: True
2016-08-21 20:41:05	wen.ma	/edit_switch	edit switch status: False
2016-08-21 20:38:23	wen.ma	/modify_replica	modify success, replicas=1
2016-08-21 20:36:53	wen.ma	/cluster_setting	setting success, settings={\"transient\": {}, \"persistent\": {\"discovery.zen.minimum_master_nodes\": 2}}

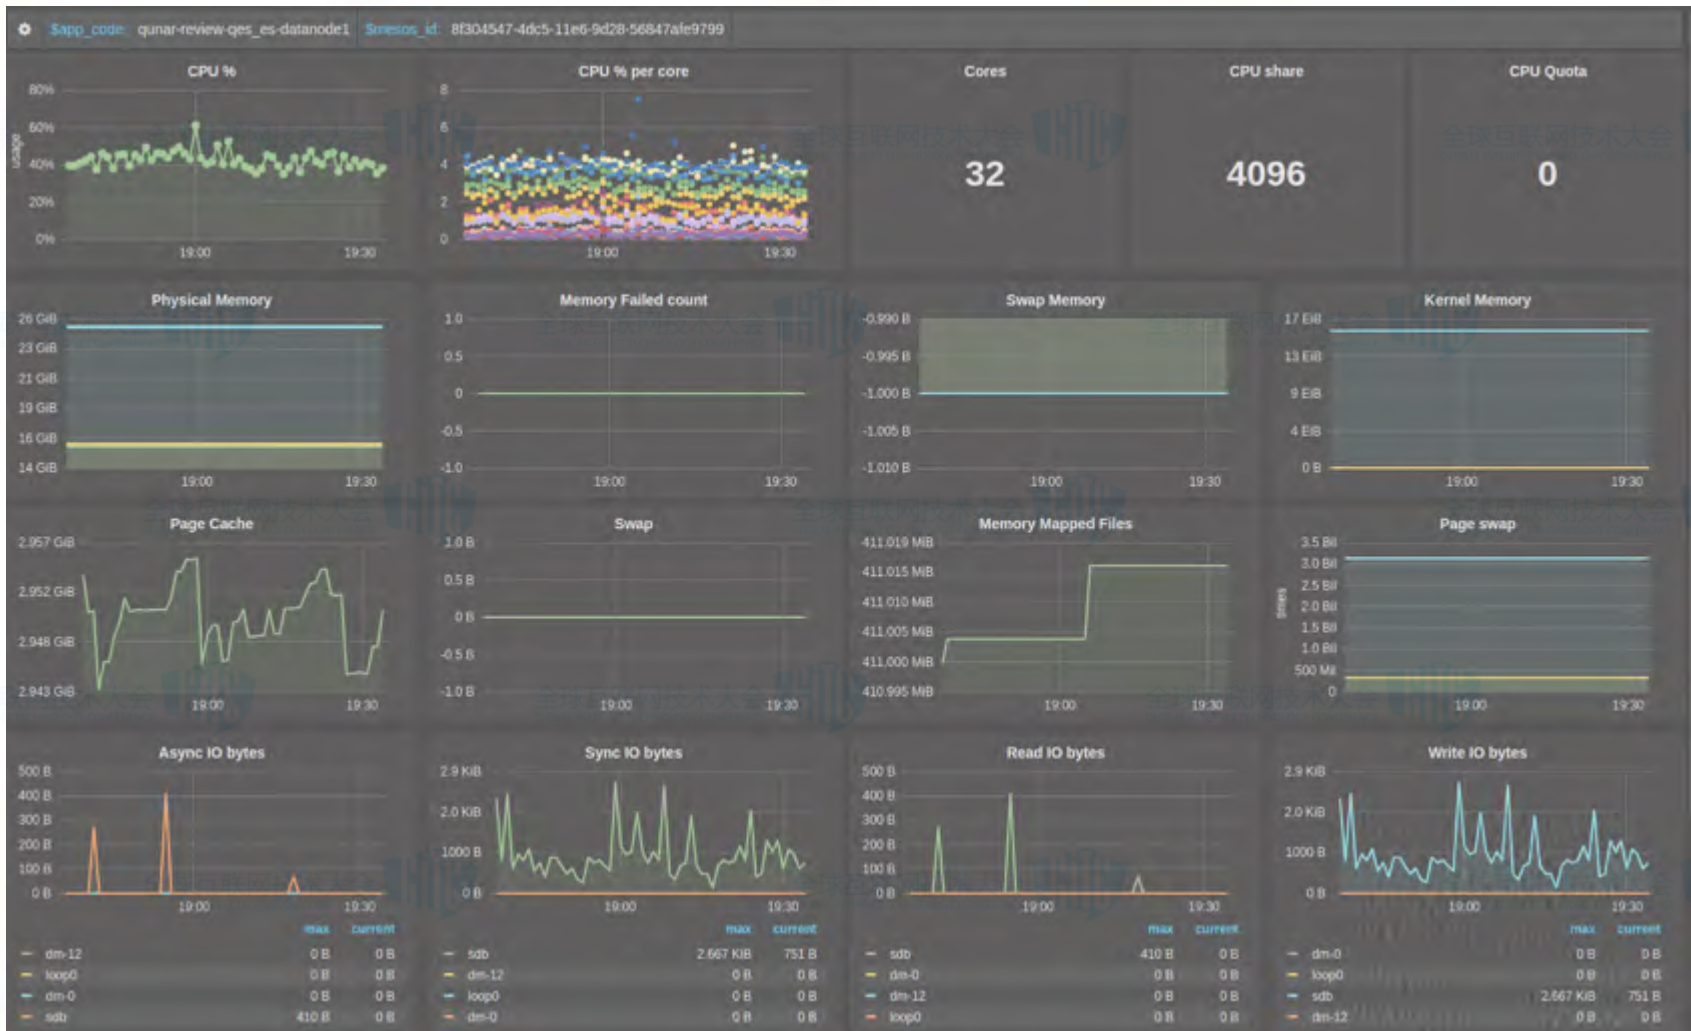
每页20条,共19条



监控/报警

- 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







规模

- **ESAAS集群机器数量：77台**
- **数据节点机器数量：66台**
- **当前托管的集群数量：44个**
- **当前数据存储总量量：120TB+**
- **当前覆盖业务线：30个**



微信扫一扫

谢谢

