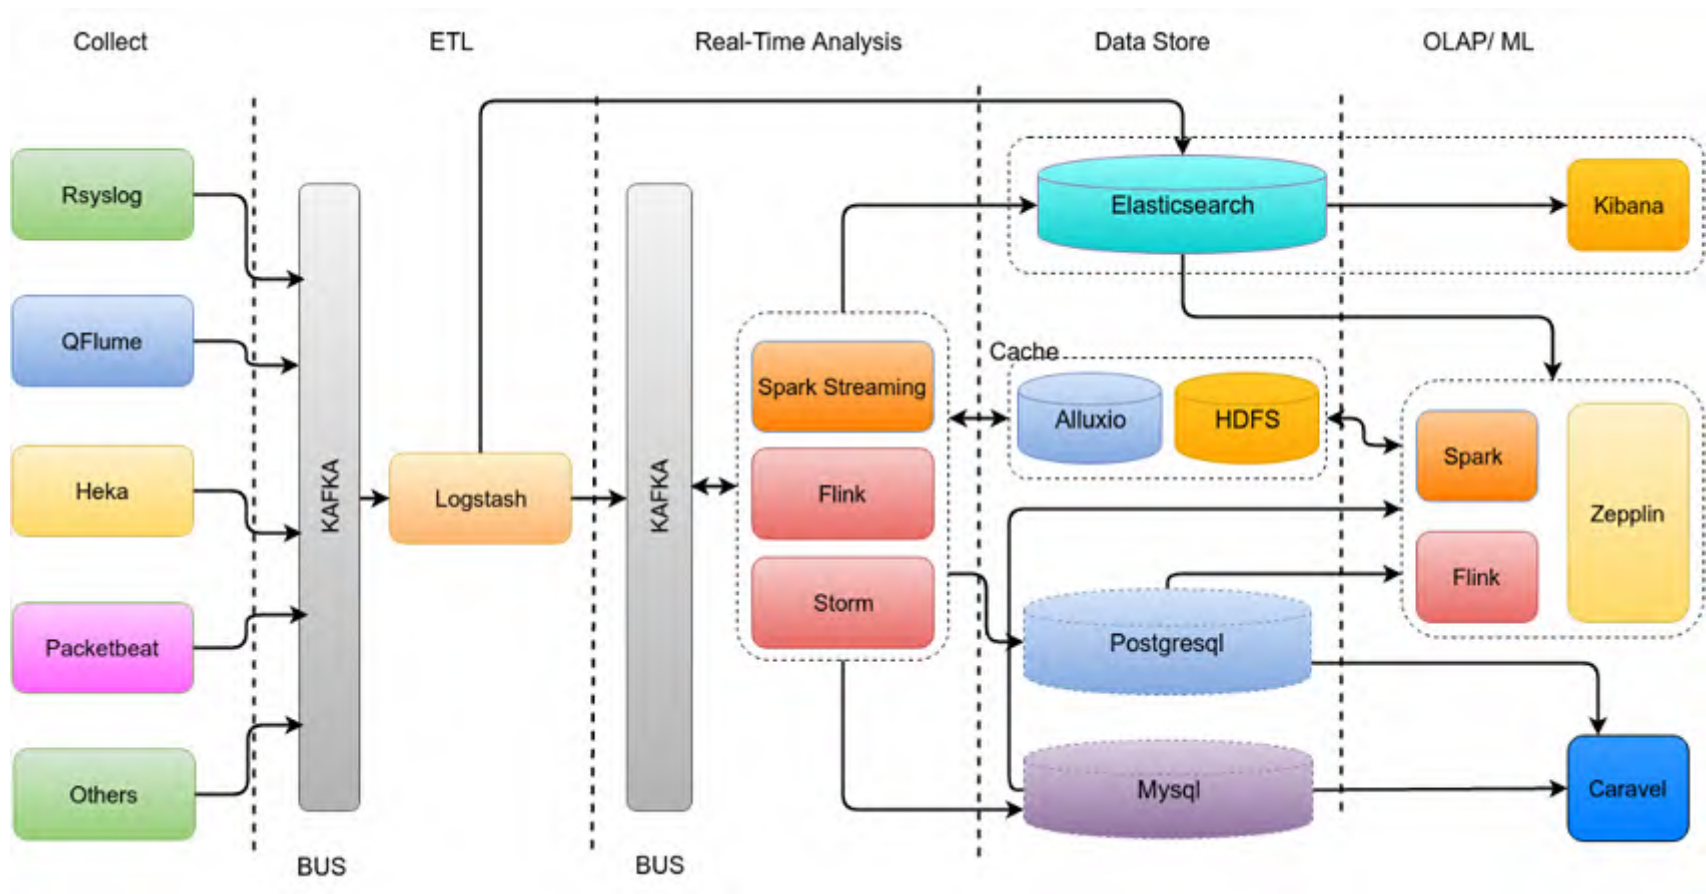


Qunar 实时数据系统实践

吕晓旭 去哪儿网实时数据平台

- 我们的实时数据平台-Prism
- 从这里开始
- 演进
- 集群规划
- 监控
- 规模

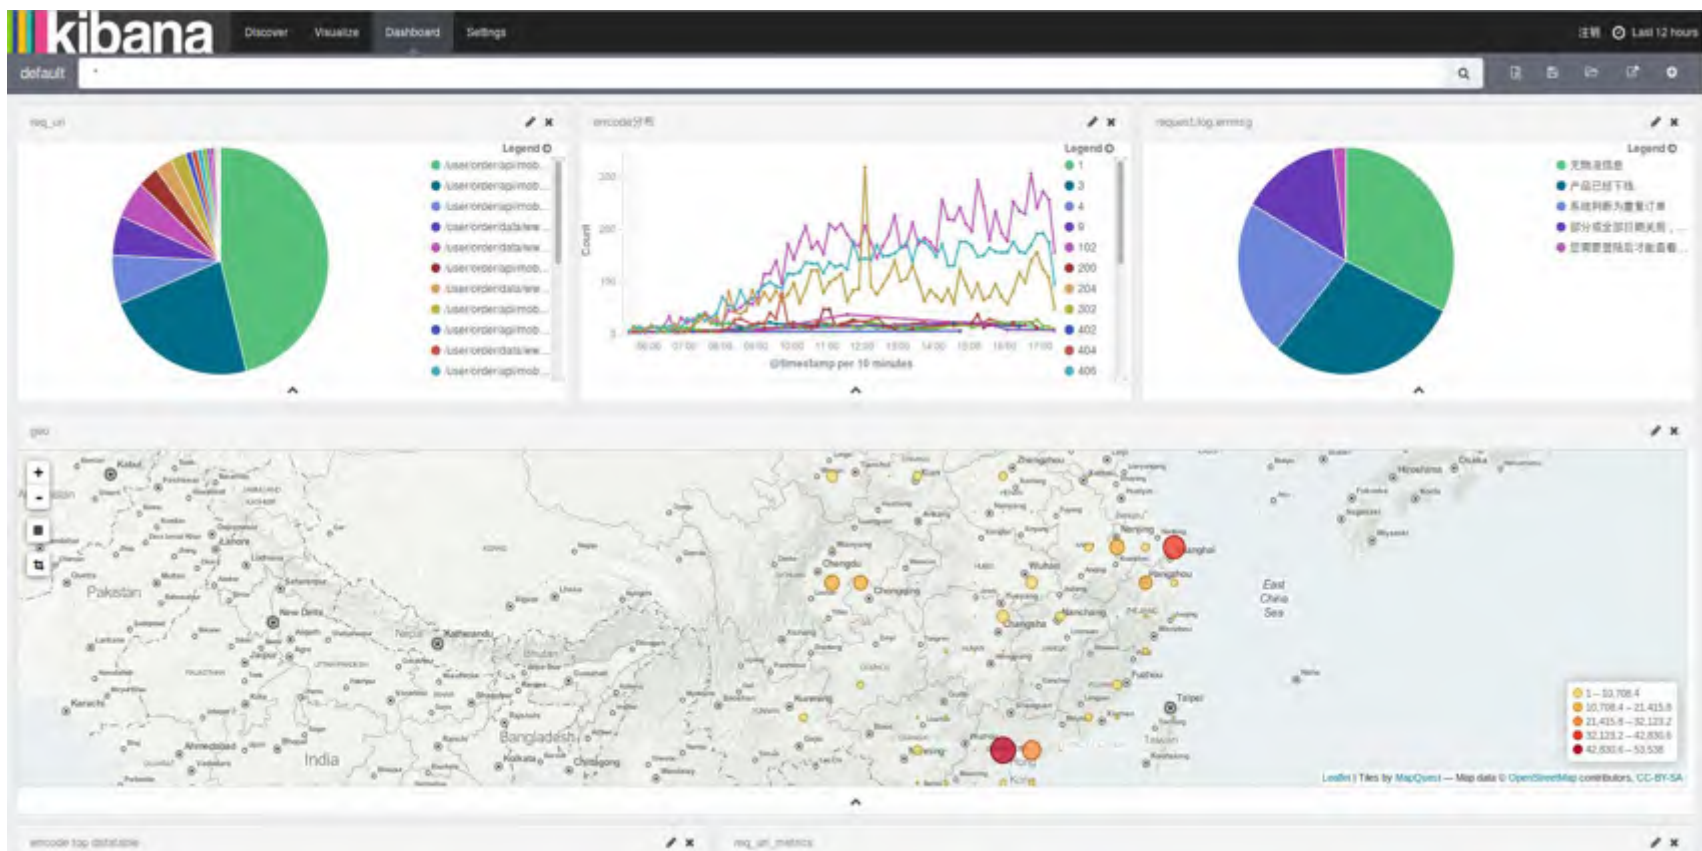
- 宗旨
 - 以数据可视化作为出发点
 - 以降低数据和数据分析软件获取成本为己任
 - 的实时数据平台
- 提供哪些服务
 - 日志实时监控 - ELK
 - 数据总线 - Kafka
 - 数据实时分析 - Spark Streaming/Storm/Flink
 - 数据存储 - Elasticsearch as a Service
 - OLAP/试验平台 - Zeppelin+Spark/Flink

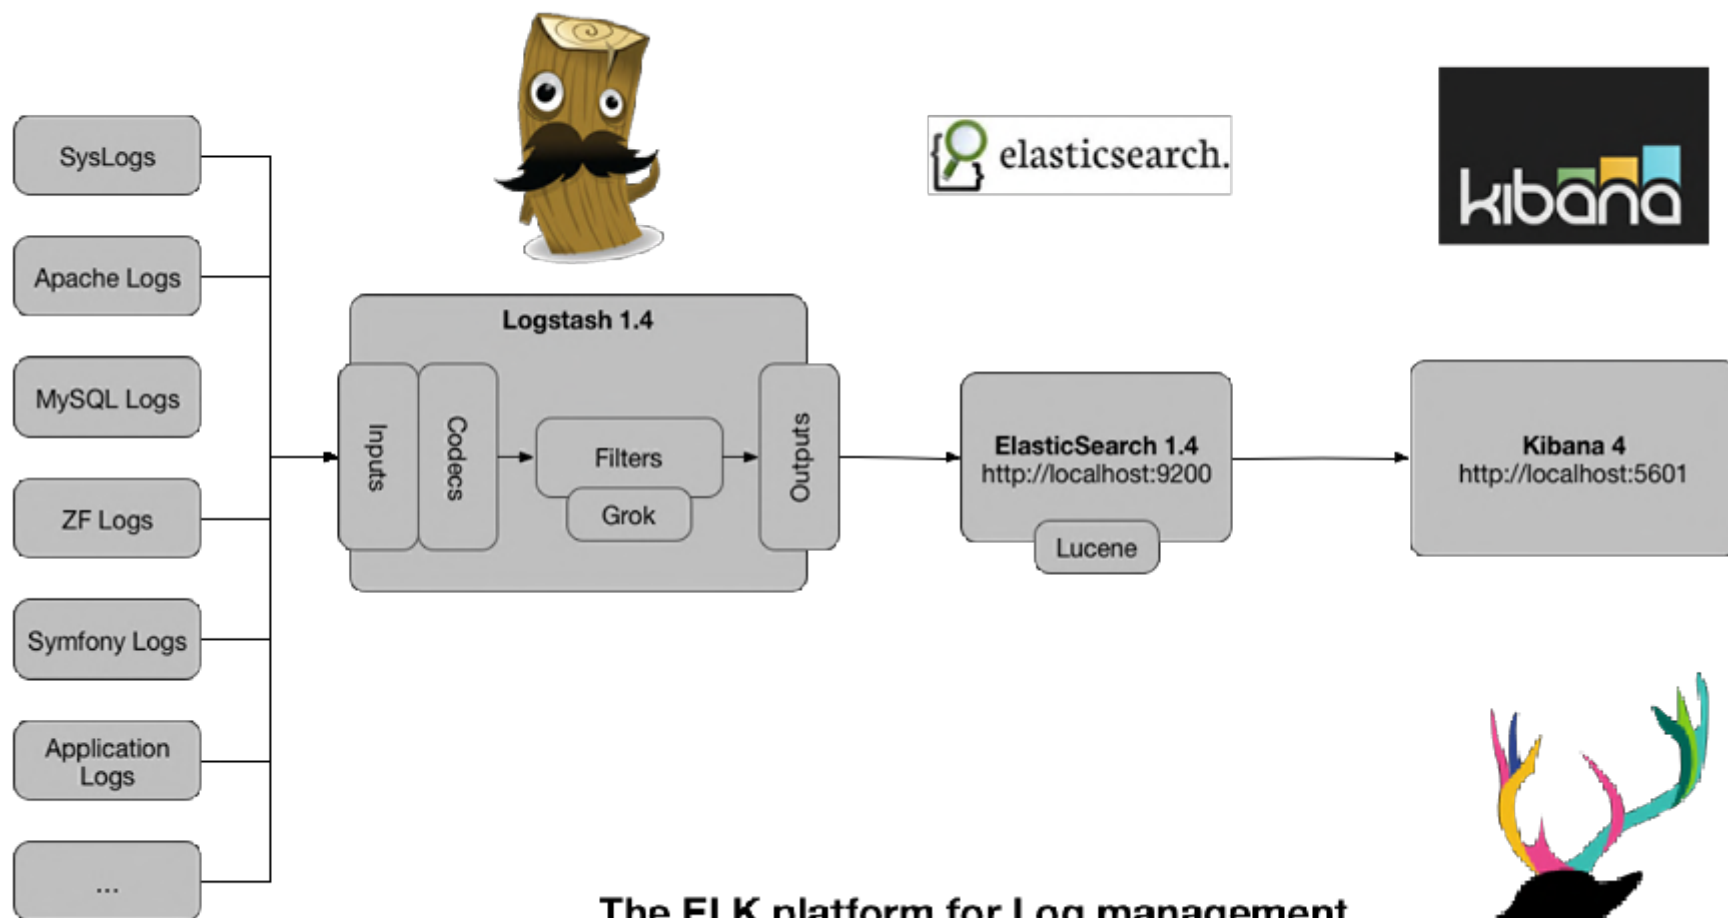


这里开始

Dev成为问题定位的瓶颈







大受欢迎



- 部署方式
 - 申请虚拟机/添加账号
 - 使用salt部署
- 面临的问题
 - 无法快速构建业务流
 - 无法快速增减容量

怎么办？



docker

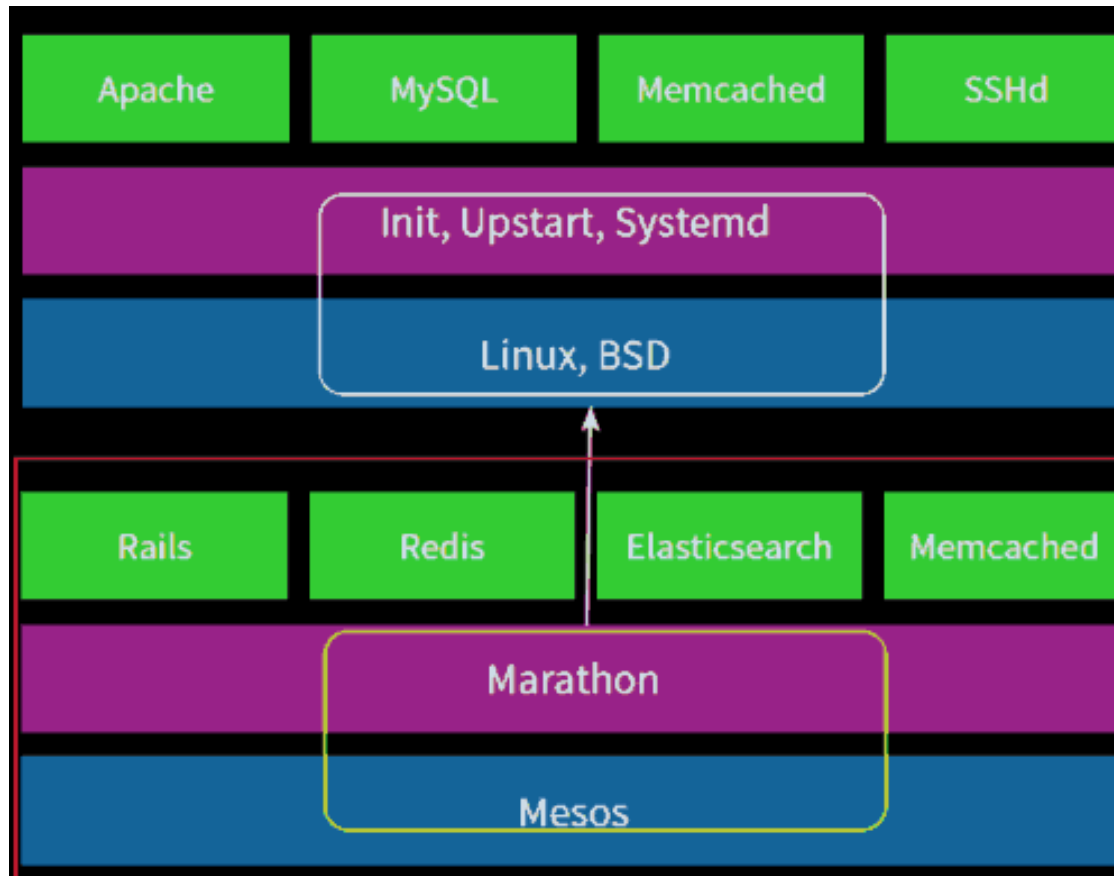



MARATHON



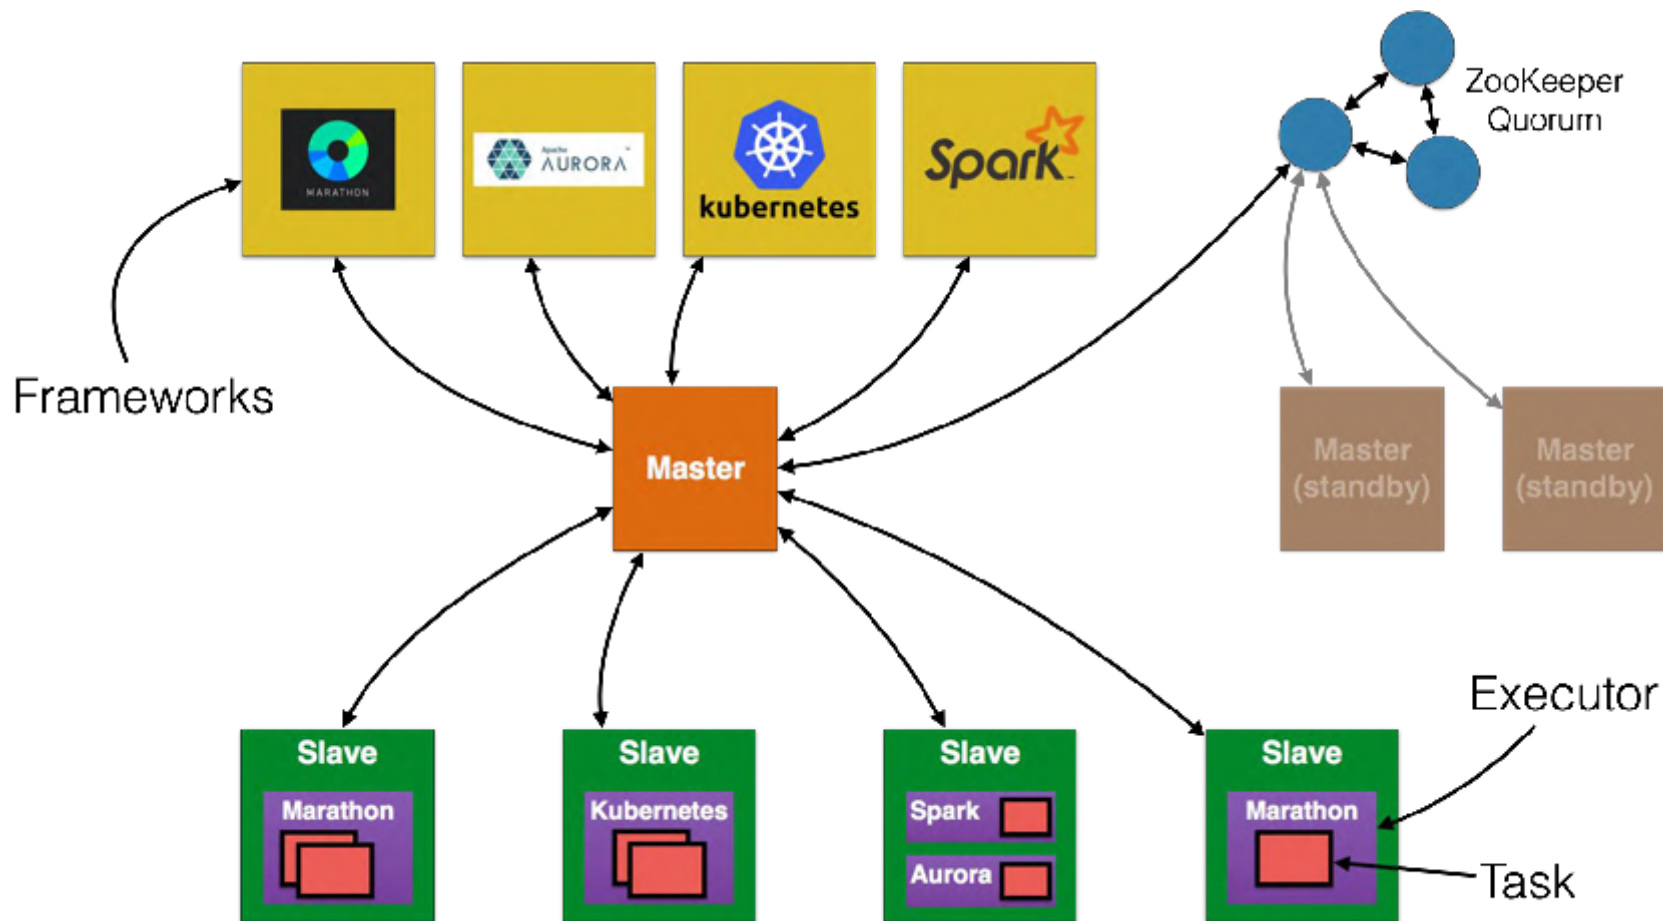
MESOS

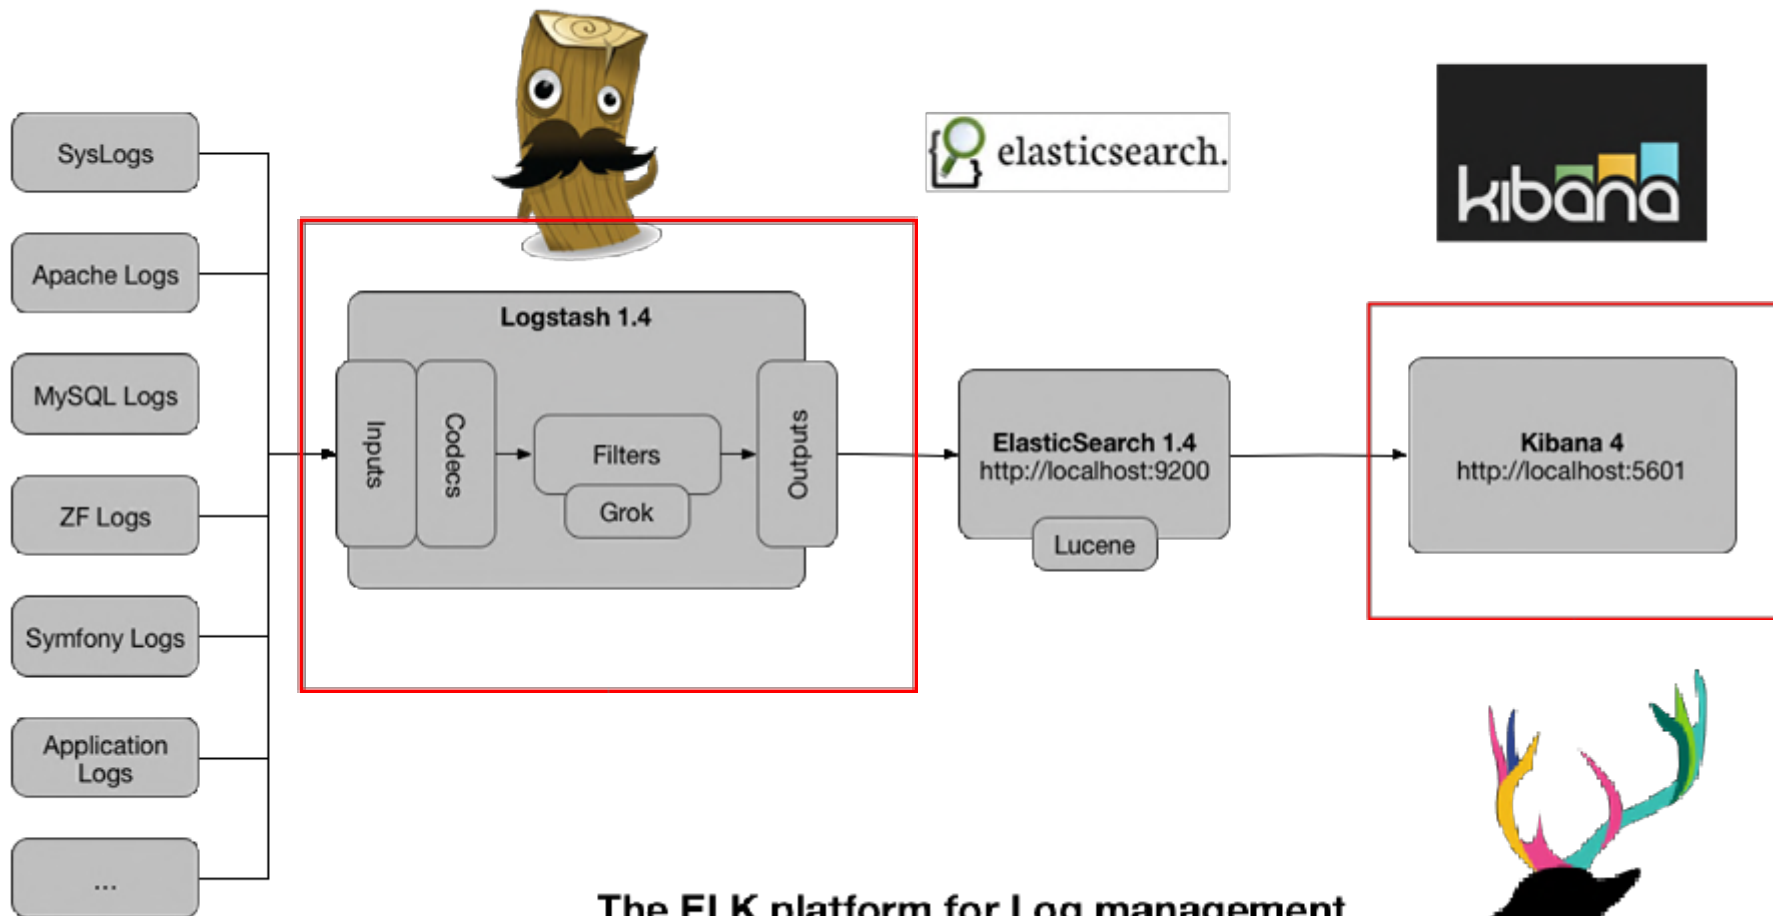
- 快速增减容量
- 新工具快速支持
- 提高硬件资源利用率
- 降低数据软件的使用成本



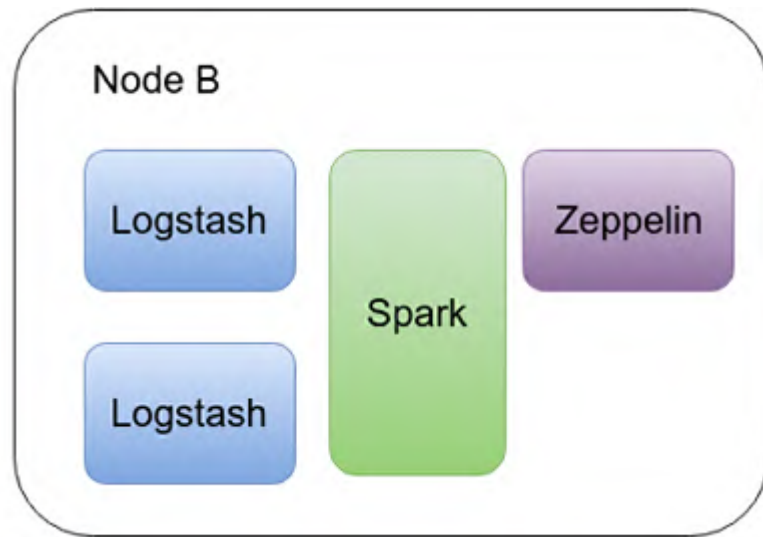
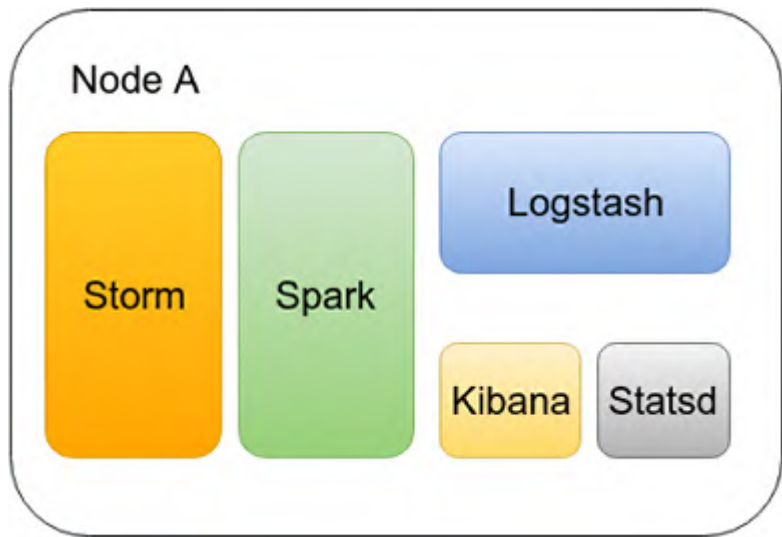


Marathon和Spark的位置

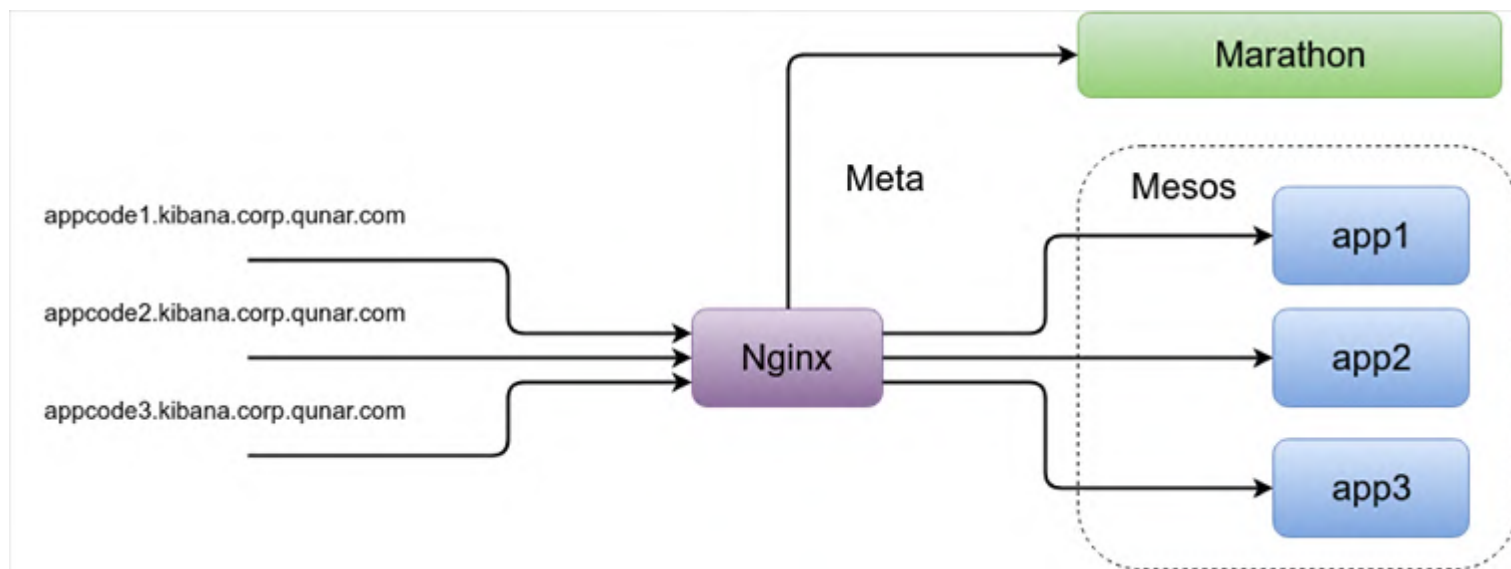


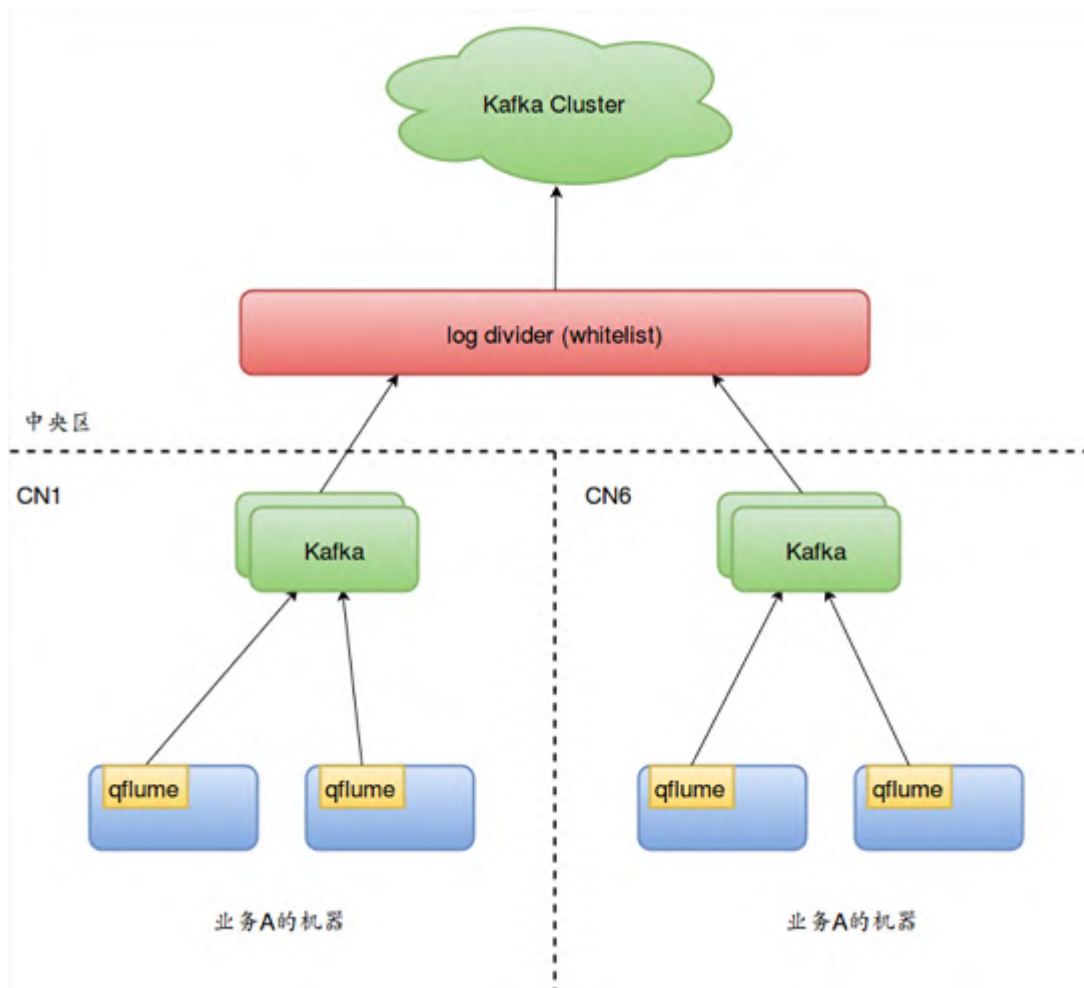


The ELK platform for Log management



- 网络方案
 - --net=host
 - Calico
 - CNI (Mesos version \geq 1.0)
- 请求路由/服务发现 (HTTP)





The screenshot displays the Logstash debugger interface with the following sections:

- Header:** Logstash debugger | Life's too short for bad tools | GELab | Wiki | xiaoxu.lv | 注册
- Config 277:** Filter by name. A list of configurations including:
 - etl.dujia.m_innovation_pitcher_poster
 - etl.hotel.h_qdstat_qreport
 - etl.train.t_train_mwhotdog
 - etl.train.t_web_kylin
 - etl.ucenter.u_hotdog
 - etl.wireless.m_innovation_pitcher_poster
 - etl.wireless.m_invocation_flight
 - etl.wireless.m_invocation_flight_snk
 - etl.wireless.m_invocation_hotdog
 - etl.wireless.m_invocation_kylin
 - logger.beta.car
 - logger.beta.flight
 - logger.beta.hotel
 - logger.beta.piao
 - logger.beta.sp
 - logger.beta.ucenter
 - logger.beta.wireless
 - logger.beta
 - logger.car.m_car_awardsms
- Pattern:** logs.logger.flight_av_data | Gmsab | Washer | Kibana. The pattern is:

```
HTTPDATE %{MONTHDAY}/%{MONTH}/%{YEAR}:%{TIME} \+%{INT}
LOGTIME %{YEAR}\-%{MONTHNUM}\-%{MONTHDAY} %{HOUR}:%{MINUTE}:%{SECOND}
ERROR_KEYWORD \b(ERROR!.*?[eE]xception!.*?[tT]imeout!Trying to connect to!WARN)
CLIENT ((unknown!%{IPV4}),[ ]?)*%{IPV4}!-
```
- Input:** 自定义input字段. Input field contains "1".
- Filter:**

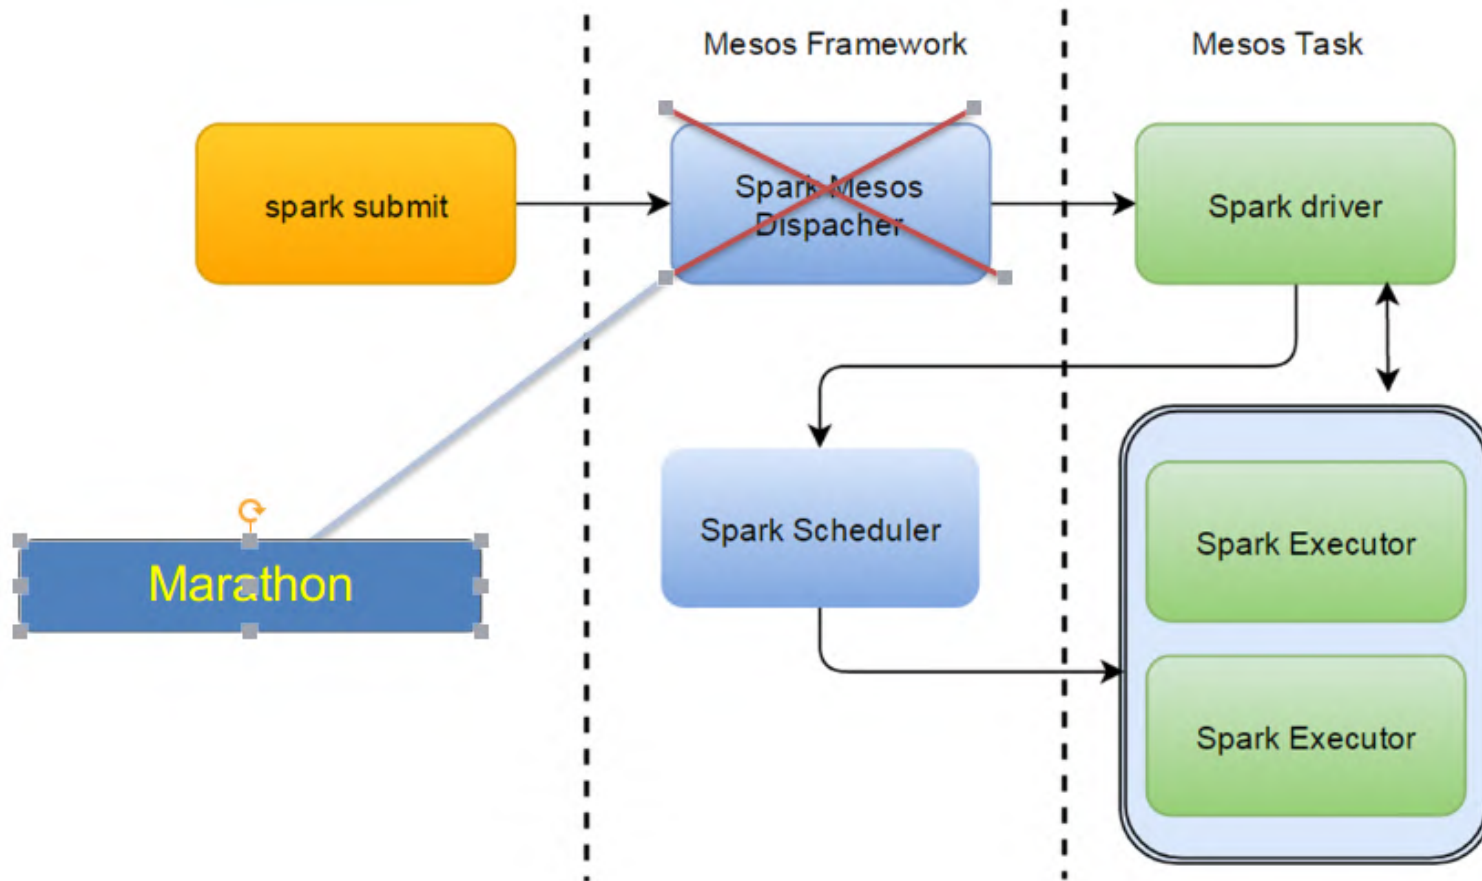
```
1 filter [
2 #   ruby [
3 #     code => "event[log_duration] = (Time.parse(event[@timestamp].to_s).to_f * 1000).to_i - event[send_time].to_i"
4 #   ]
5 #}
```
- Buttons:** Debug (blue), Commit (orange), Deployment (red).
- Terminal:**

```
scrat -s /tmp/s.png
+ - scrat -s /tmp/s.png
+ - scrat -s /tmp/s.png
```

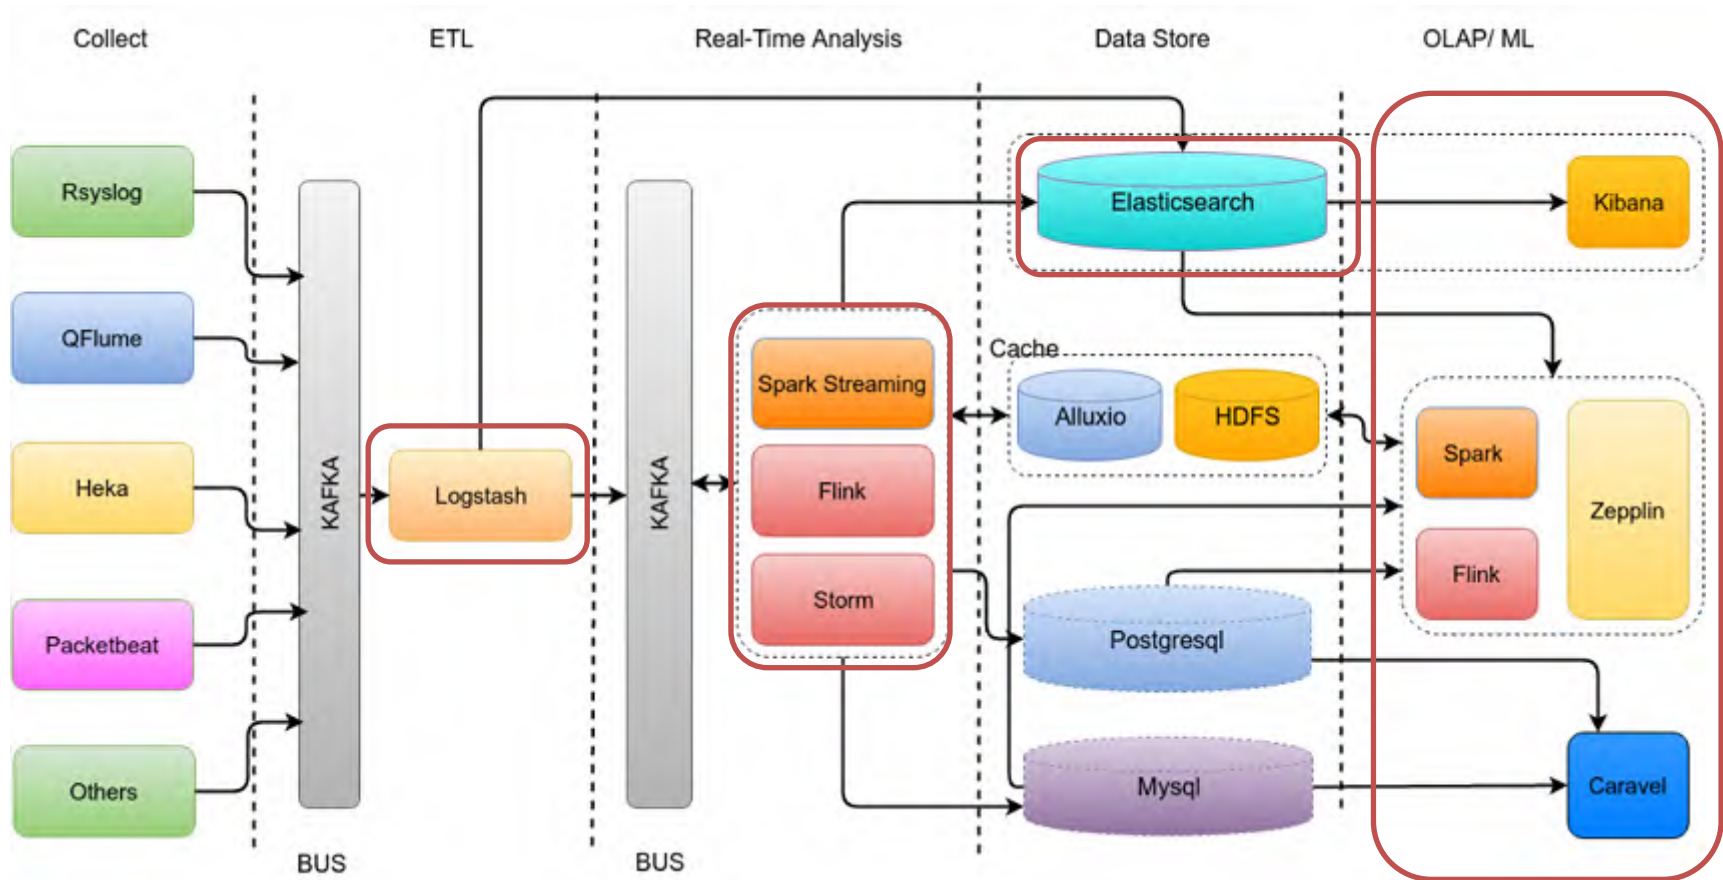
新需求

- 实时推荐
- 多数据源实时JOIN
- Logstash能力不足以支撑新的需求


引入Spark on Mesos



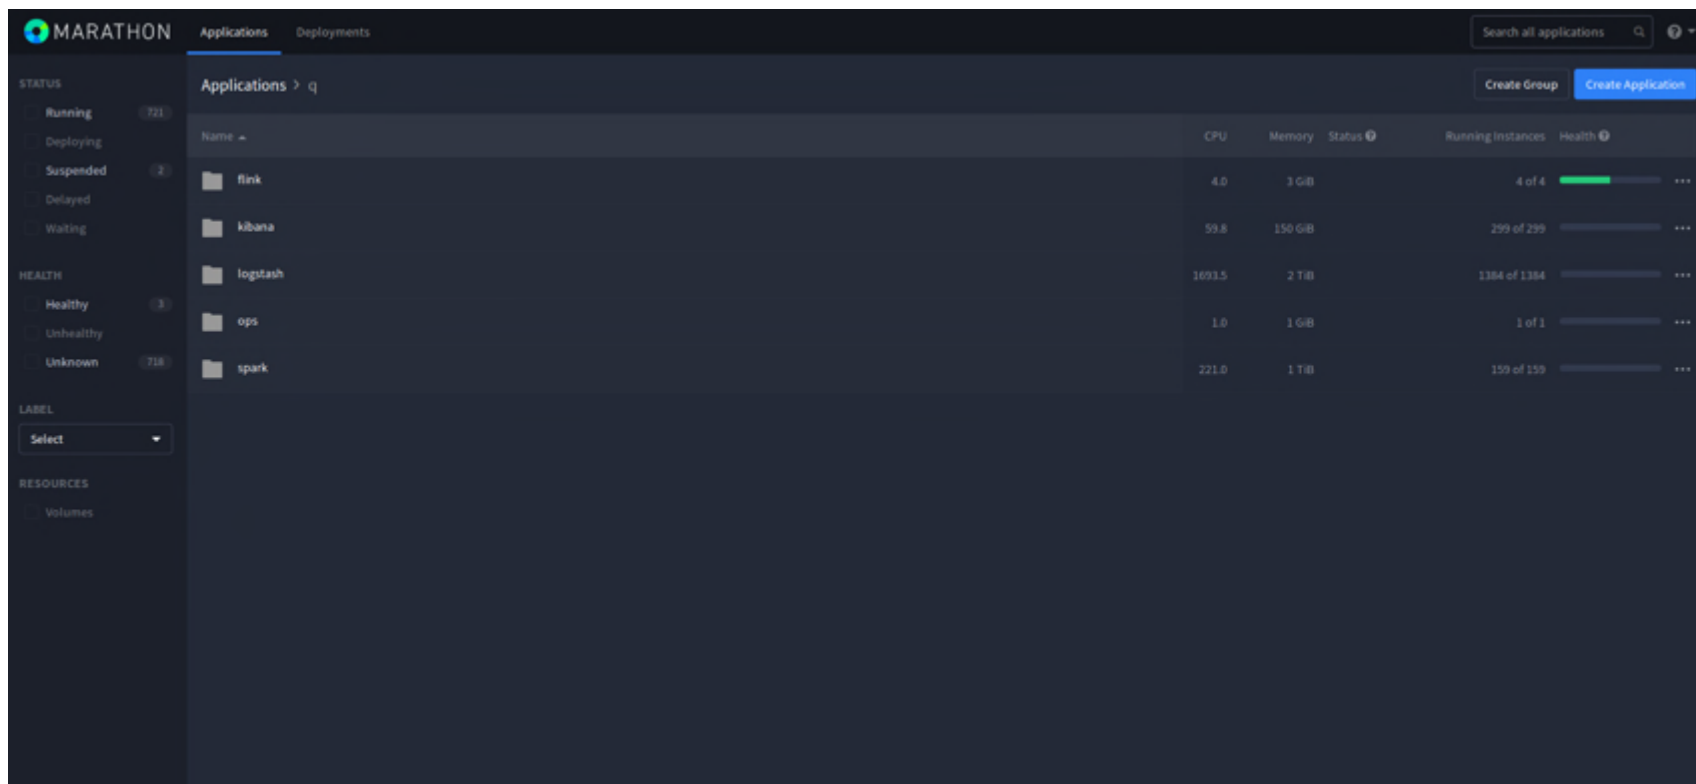
Software on Mesos



logs_wireless_m_pub_web_hotdog的网络



- 实时日志分析：300模块，kafka峰值带宽6G
- Spark Streaming任务：50个
- Storm集群：5个
- Flink集群：2个

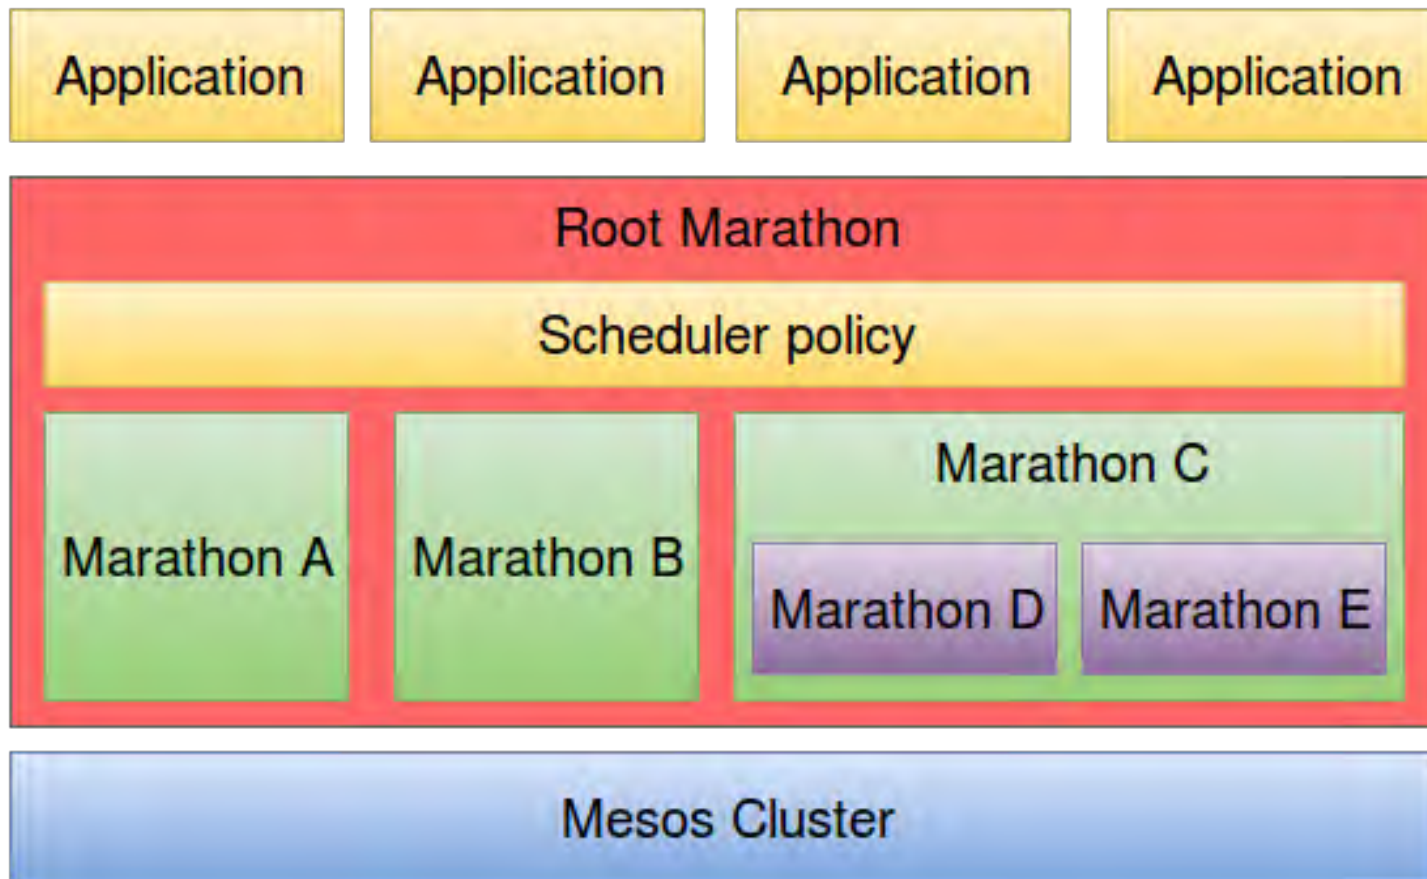


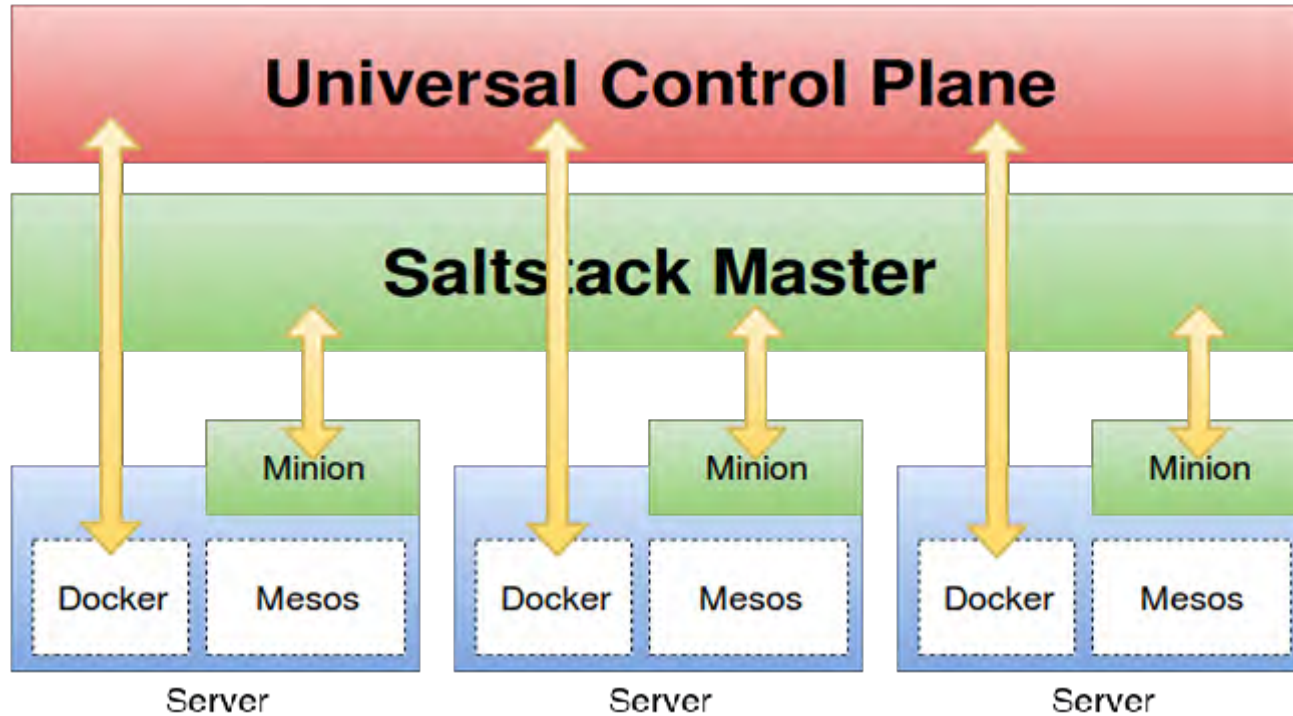
The screenshot shows the Apache Marathon web interface. The main content area displays a table of applications with the following columns: Name, CPU, Memory, Status, Running Instances, and Health. The table lists five applications: flink, kibana, logstash, ops, and spark. The flink application is highlighted with a green bar, indicating it is healthy. The kibana application has 299 of 299 running instances. The logstash application has 1384 of 1384 running instances. The ops application has 1 of 1 running instances. The spark application has 159 of 159 running instances.

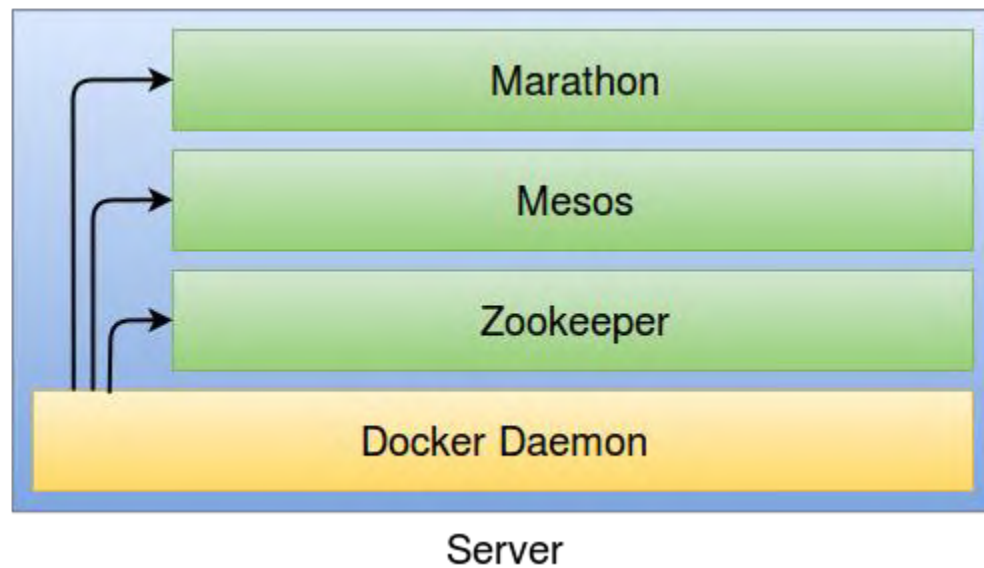
Name	CPU	Memory	Status	Running Instances	Health
flink	4.0	1 GB	Running	4 of 4	Healthy
kibana	59.8	150 GB	Running	299 of 299	Healthy
logstash	1693.5	2 TB	Running	1384 of 1384	Healthy
ops	1.0	1 GB	Running	1 of 1	Healthy
spark	221.0	1 TB	Running	159 of 159	Healthy

集群规划

- 基础设施这一层技术收敛
- Fail Over
- 多租户资源隔离



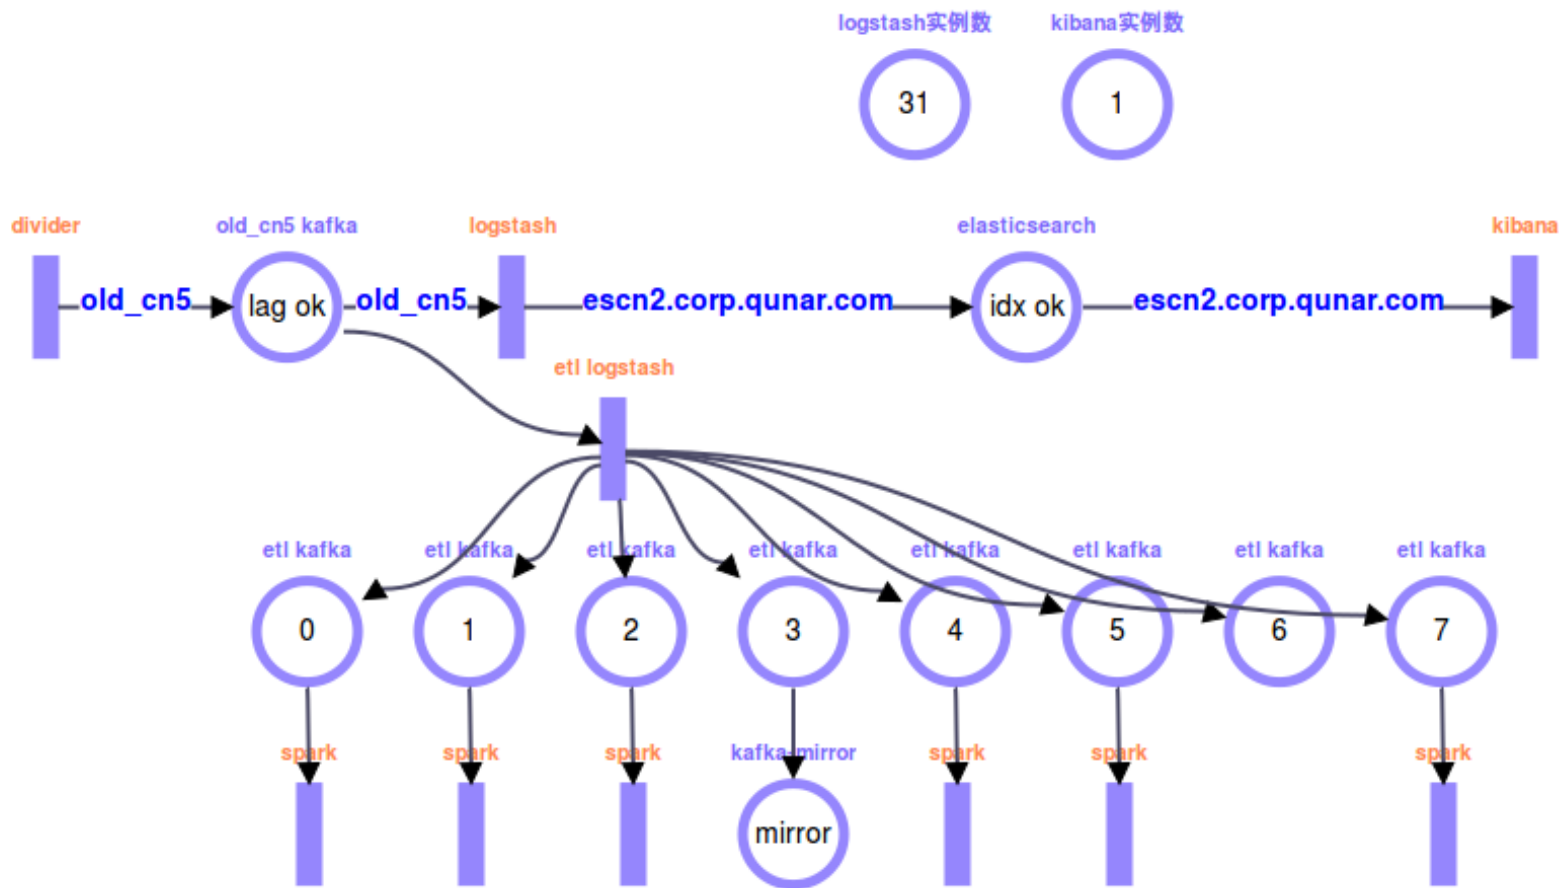




监控

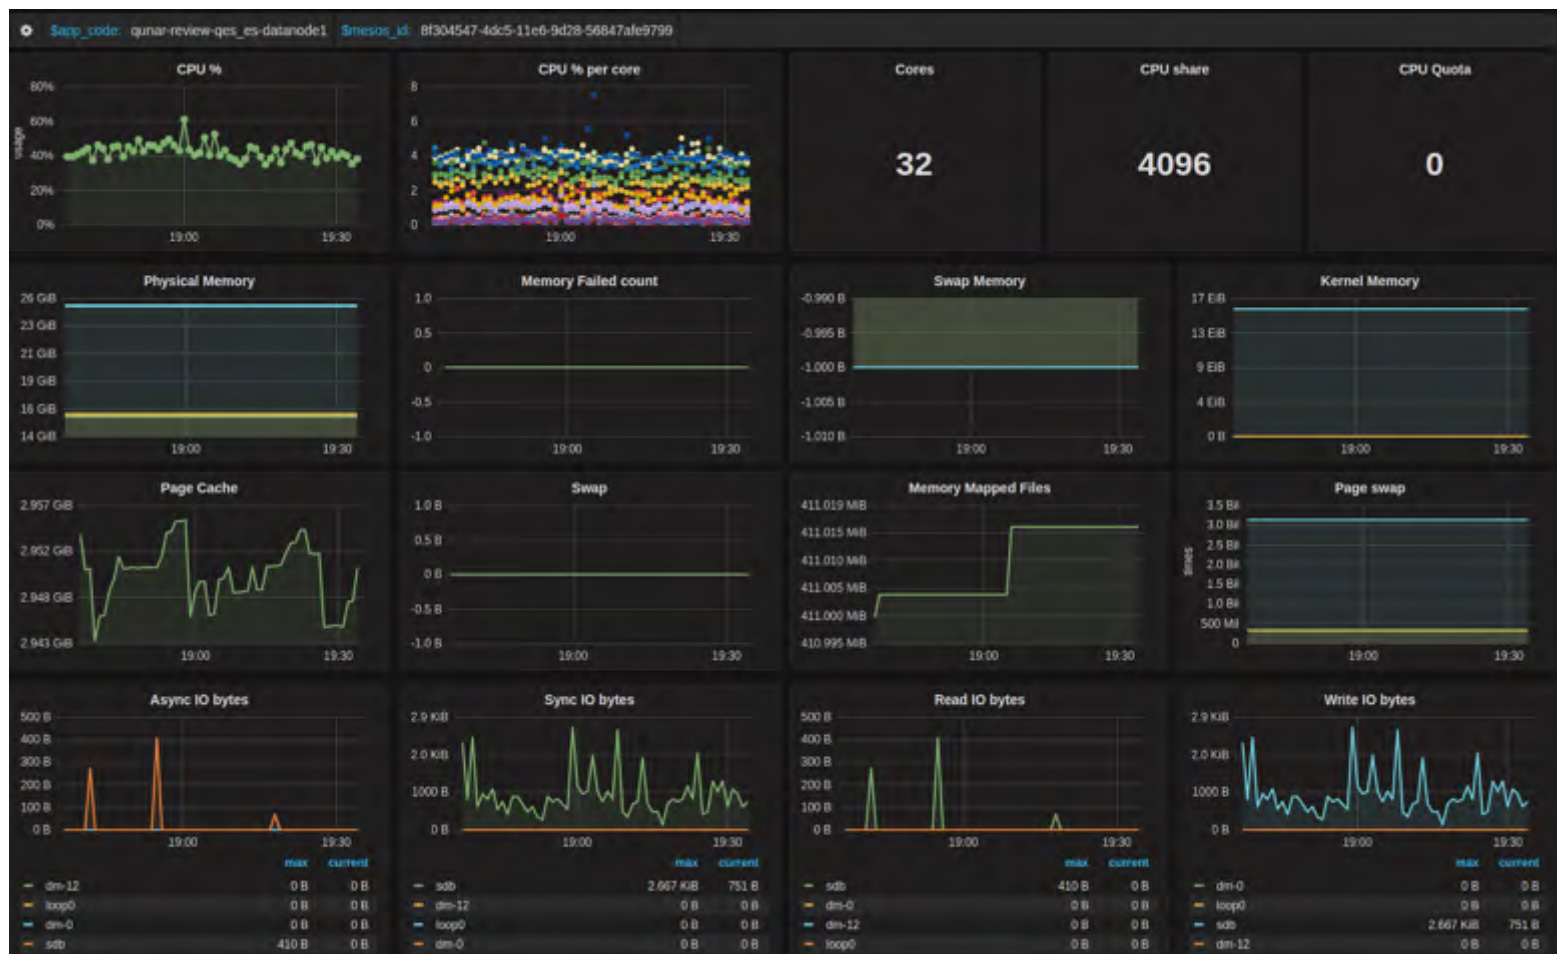
- 数据处理模块拓扑监控
- 业务统计监控
 - 队列堆积：Kafka Topic Lag
 - 流量：Search Count/Message Count
 - 错误：Reject/Exception
- 基础监控/容量监控
 - IO使用率
 - CPU使用率
 - 内存使用率
 - JVM/GC等
 - 集群资源使用量

logs_wireless_m_pub_web_hotdog的网络











- <https://github.com/QunarOPS/pyadvisor>

- 计算集群120+ ; 2600+ 容器
- ES中存储的日志160T , 4万Shards
- Esaas 50+; 47 ES集群 ; 600+ 容器

- 我们做的事儿
 - 实时数据治理
 - 解决数据软件的部署的门槛
 - 解决Mesos环境部署的门槛
- 仍存在的问题
 - 负载不均匀
 - 数据异常定位速度慢
- 下一步计划
 - 解决以存在的问题
 - 接入新的软件
 - GPU计算平台建设



BDTC 2016中国大数据技术大会
Big Data Technology Conference 2016



BDTC 2016中国大数据技术大会
Big Data Technology Conference 2016