Kubernetes监控与日志

时速云 – 赵宇

Agenda

- Kubernetes Summary
  - Features
  - Architecture
- Kubernetes Monitor
  - Resource Monitor
  - Service Monitor
- Kubernetes Logging
  - Fluentd
  - Elasticsearch
  - Flow
- 时速云平台实践

Features

- ✓ 资源管控
- ✓ 横向扩展
- ✓ 服务管理
- ✓ 数据持久化
- ✓ 自动化运维
- ✓ 服务发现和LB
- ✓ 安全&配置
- ✓ 工作任务

Architecture

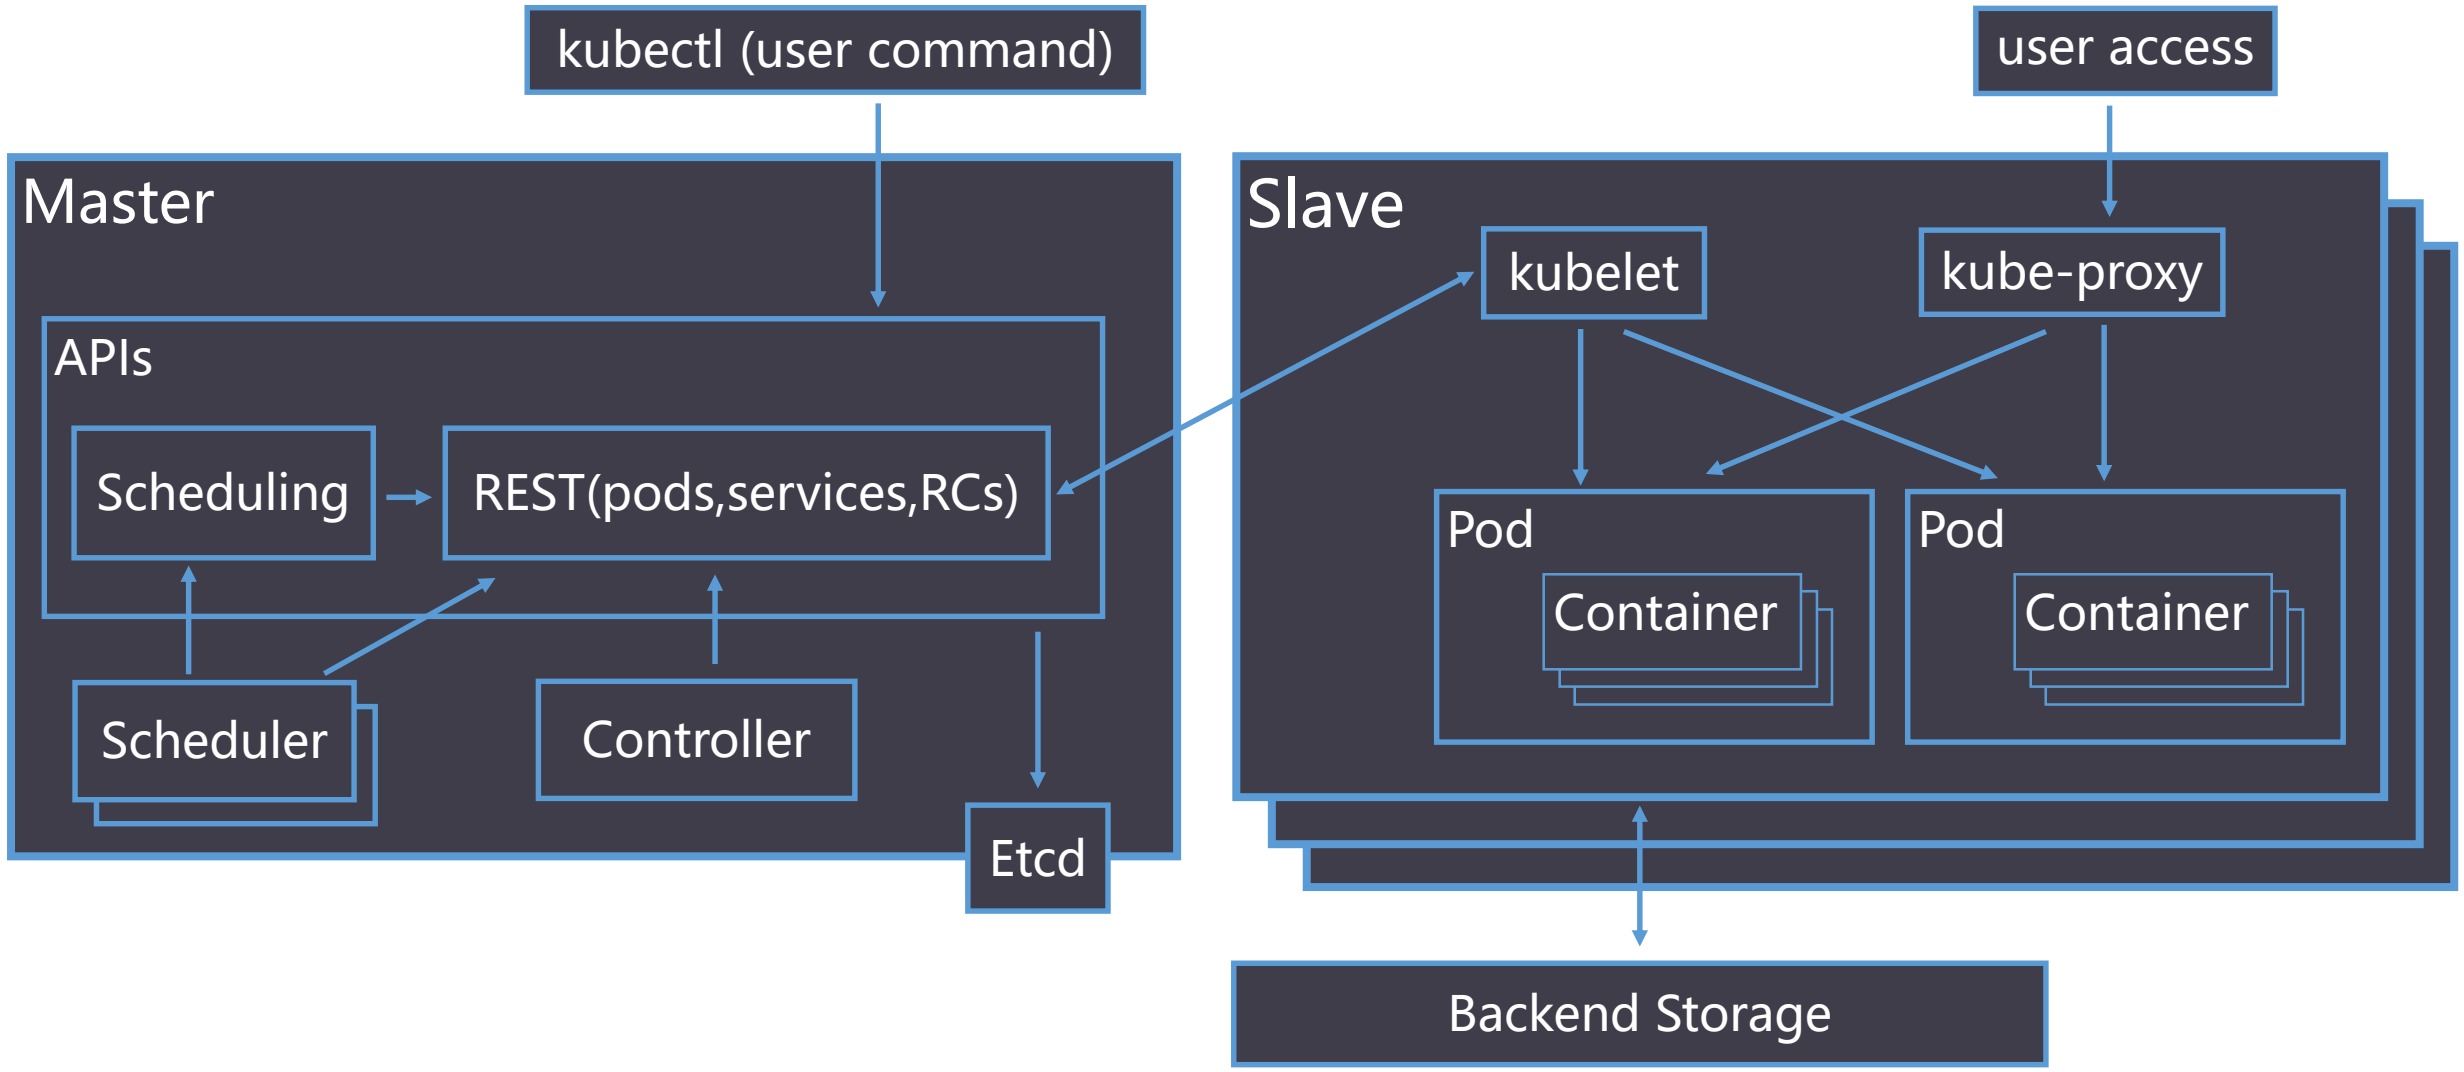

Architecture

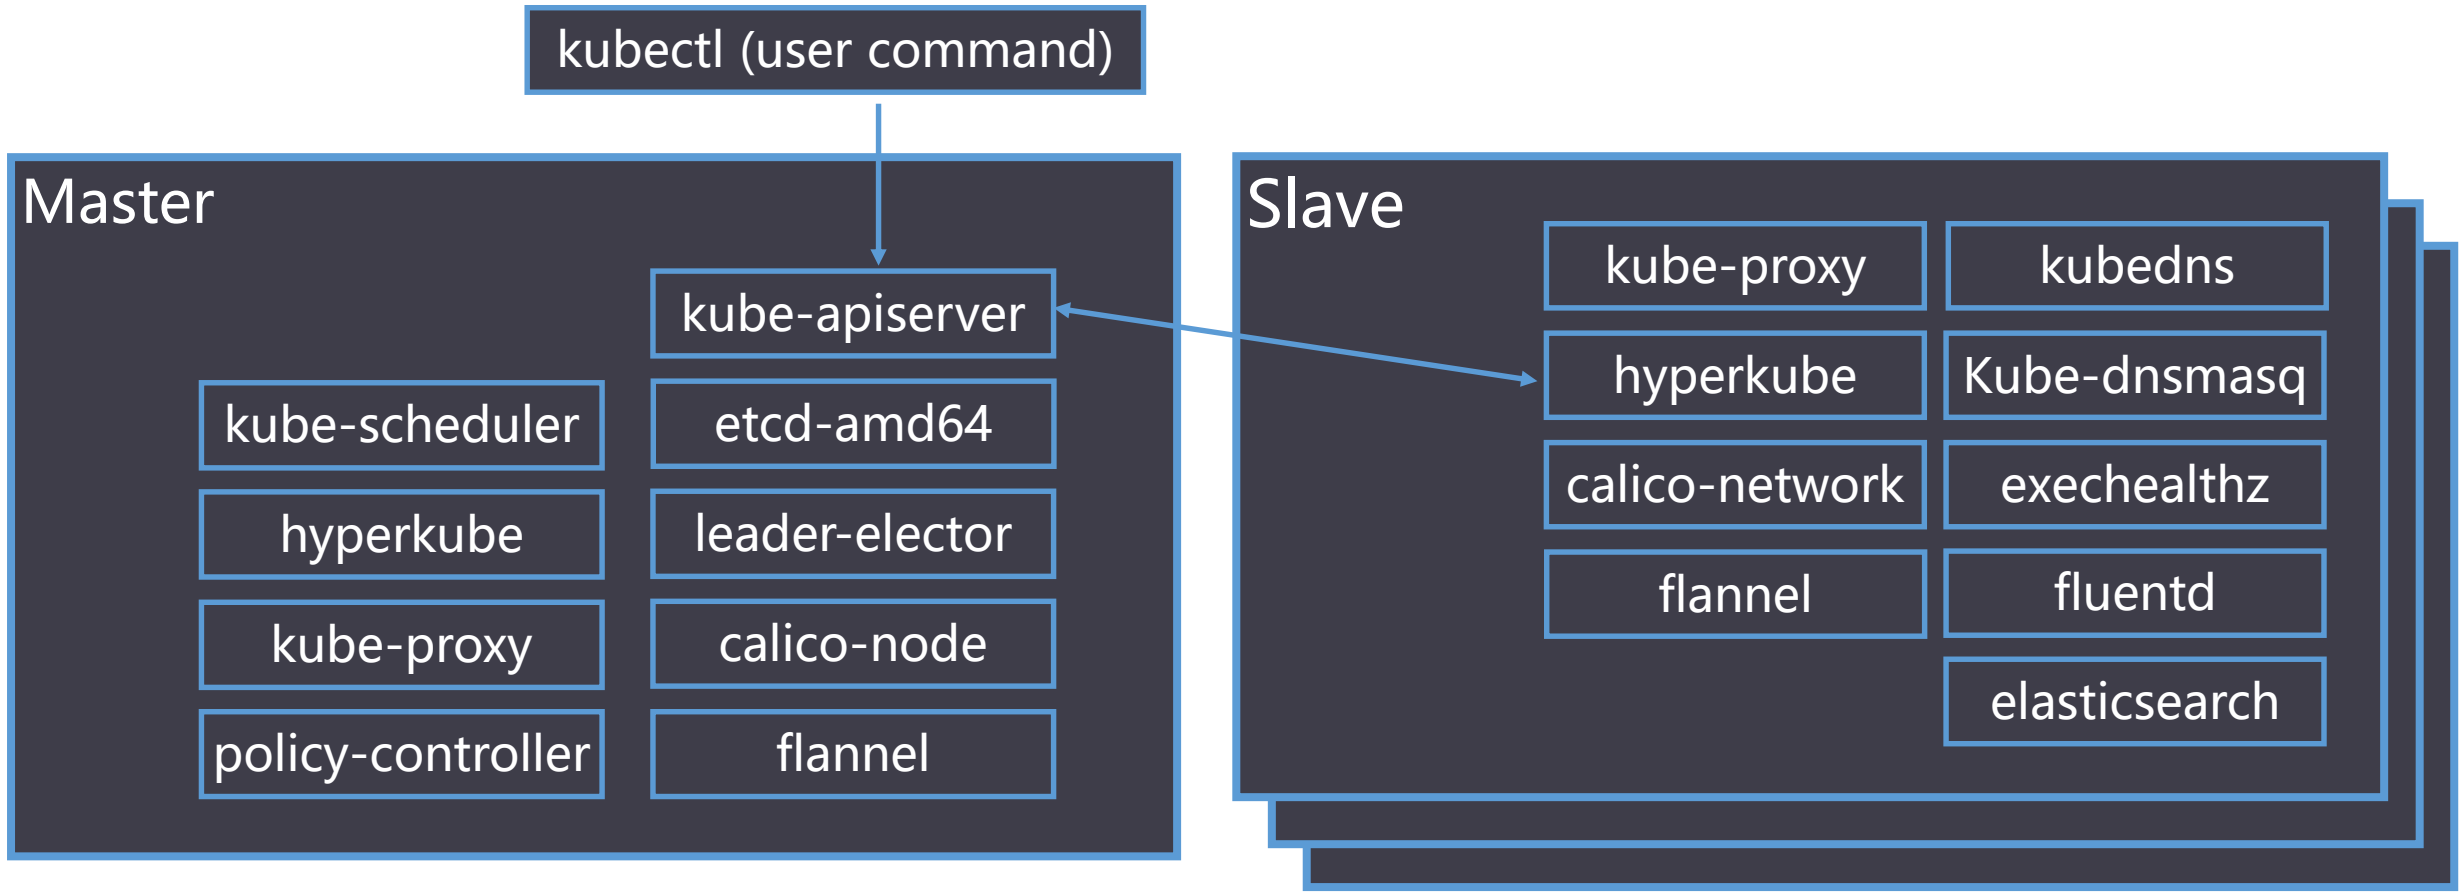

Agenda

- Kubernetes Summary
  - Features
  - Architecture
- Kubernetes Monitor
  - Resource Monitor
  - Service Monitor
- Kubernetes Logging
  - Fluentd
  - Elasticsearch
  - Flow
- 时速云平台实践

Resource Monitor

- Resource
  - ✓ Cpu：1 core = 1000m
  - ✓ Memory：measured by bytes, 128MiB
- Container：利用容器自身限制
- Pod：限制系统内创建Pod的资源
- Namespace：用户级别的资源限制，项目更多，configmaps， pods ，svc...
- 查看资源：kubectl describe po/no --namespace=**

```
Allocated resources:
  (Total limits may be over 100%, i.e., overcommitted. More info
  CPU Requests  CPU Limits      Memory Requests Memory Limits
  ------------  ----------      --------------- -------------
  4034m (50%)   4645m (58%)     15866Mi (99%)   19040Mi (118%)
```

```
Containers:
  abcd:
    Container ID:       docker:/
    Image:              index.te
    Image ID:           docker:/
    Ports:              22/TCP,
    QoS Tier:
      cpu:      Burstable
      memory:   Burstable
    Limits:
      cpu:      125m
      memory:   512Mi
    Requests:
      cpu:              109m
      memory:           448Mi
```

Resource Monitor

- 资源监控主要由kubelet，heapster和storage backends（如Influxdb）构成。Heapster可以在集群范围获取metrics和事件数据。它可以以pod的方式运行在集群中，也可以单独运行
- kubelet 从cAdvisor获取数据，heapster组织数据和推送数据到后端存储InfluxDB。cAdvisor不但可以统计节点上每个容器的资源情况，还提供整个节点的资源使用情况。

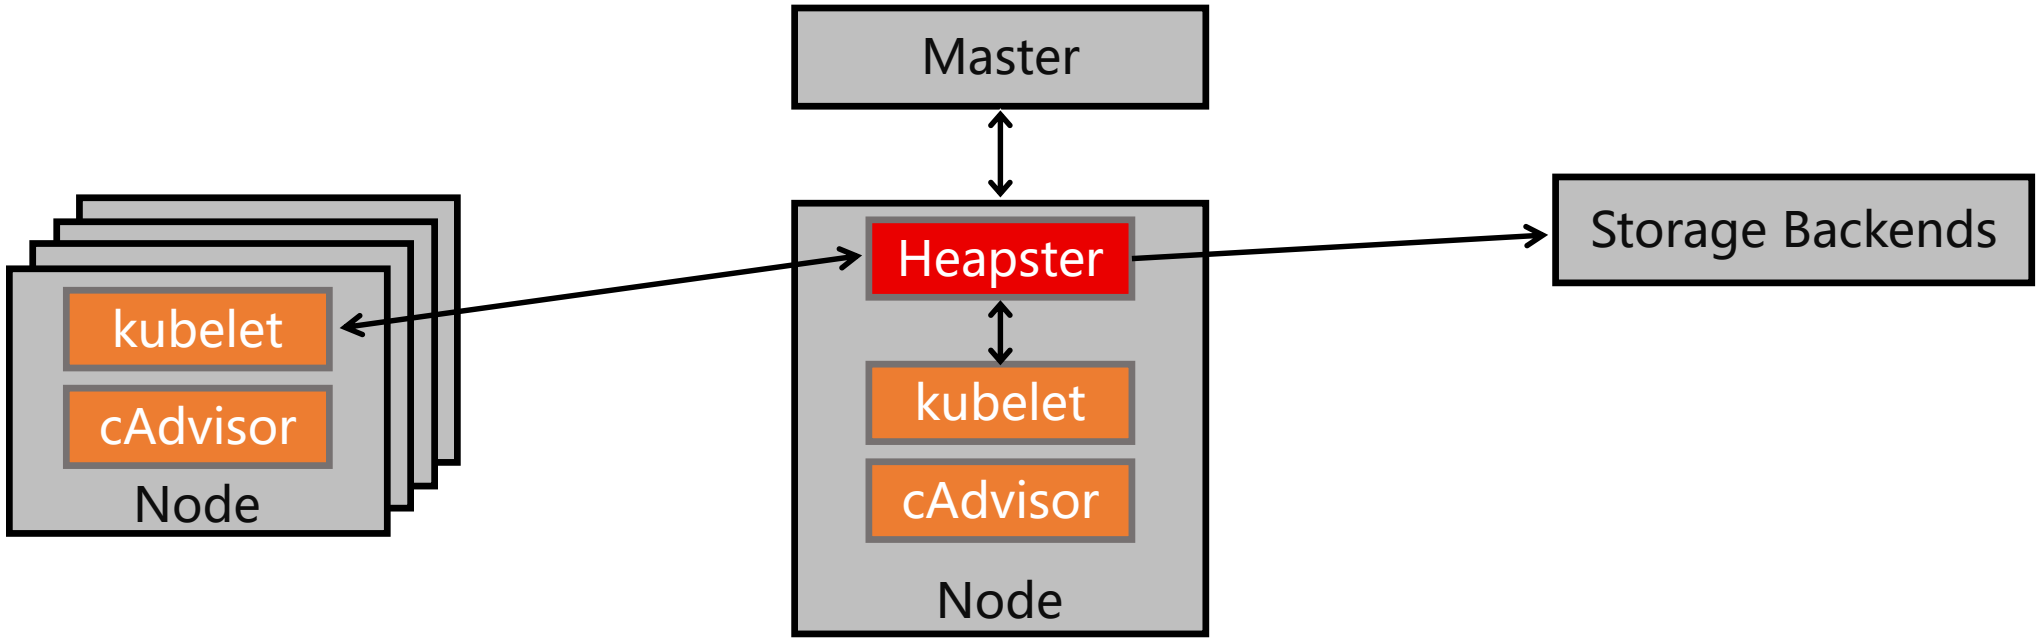

Resource Monitor

- cAdvisor（Container Advisor）是开源的专门用于监控、分析容器资源的代理程序
- cAdvisor自动发现所有容器并收集CPU、内存、磁盘和网络数据
- cAdvisor在每个节点提供简易的UI，通过端口4194访问
- 给容器打标签，并设置资源限制

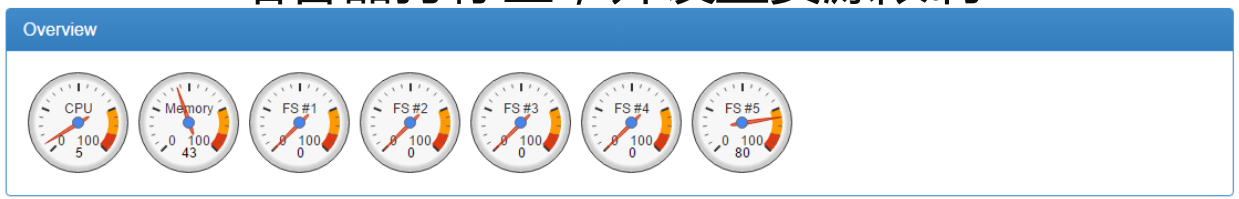

Processes

User	PID	PPID	Start Time	CPU %	MEM %	RSS	Virtual Size	Status	Running Time	Command	Container
root	6,218	5,397	Nov06	11.00	0.10	25.15 MiB	34.07 MiB	Ssl	05:10:28	kube-proxy	/docker/99d1a9e5d130402c7d
statd	8,661	5,397	Nov07	8.30	6.70	1.05 GiB	12.17 GiB	Ssl	02:26:28	java	/docker/33d5f56a8c42470569
root	63,619	5,397	Nov07	2.00	0.50	81.57 MiB	678.80 MiB	Ssl	00:35:44	hyperkube	/docker/ca1ea953ad73d65241
root	6,283	6,262	Nov06	1.70	1.10	182.85 MiB	1.22 GiB	Sl	00:49:01	td-agent	/docker/8a60c09cbf53f56581
root	5,397	1	Nov06	1.20	1.80	304.55 MiB	6.41 GiB	Ssl	00:35:36	docker	
root	2,954	5,397	Nov07	0.50	0.10	25.96 MiB	37.34 MiB	Ssl	00:10:09	etcd	/docker/7b103e22f1140266e6
root	4,556	5,397	Nov07	0.50	0.00	13.58 MiB	16.97 MiB	Ssl	00:10:20	skydns	/docker/d9130ec6a87bad6933
root	9,429	5,397	Nov07	0.20	0.30	56.48 MiB	67.13 MiB	Ssl	00:04:40	heapster	/docker/aeb4962d174db2636c
root	7	2	Nov06	0.10	0.00	0.00 B	0.00 B	S	00:04:45	rcu_sched	
root	9	2	Nov06	0.10	0.00	0.00 B	0.00 B	S	00:05:41	rcuos/0	
root	1,648	5,397	Nov07	0.10	0.20	33.72 MiB	48.16 MiB	Ssl	00:01:55	nginx-ingress-c	/docker/0ab5875d24adcd0926
root	4,141	5,397	Nov07	0.10	0.10	24.50 MiB	37.58 MiB	Ssl	00:02:23	kube2sky	/docker/6f56c304e0e3a169b7
root	1	0	Nov06	0.00	0.00	4.05 MiB	32.94 MiB	Ss	00:00:08	init	
root	2	0	Nov06	0.00	0.00	0.00 B	0.00 B	S	00:00:00	kthreadd	
root	3	2	Nov06	0.00	0.00	0.00 B	0.00 B	S	00:00:02	ksoftirqd/0	
root	8	2	Nov06	0.00	0.00	0.00 B	0.00 B	S	00:00:00	rcu_bh	

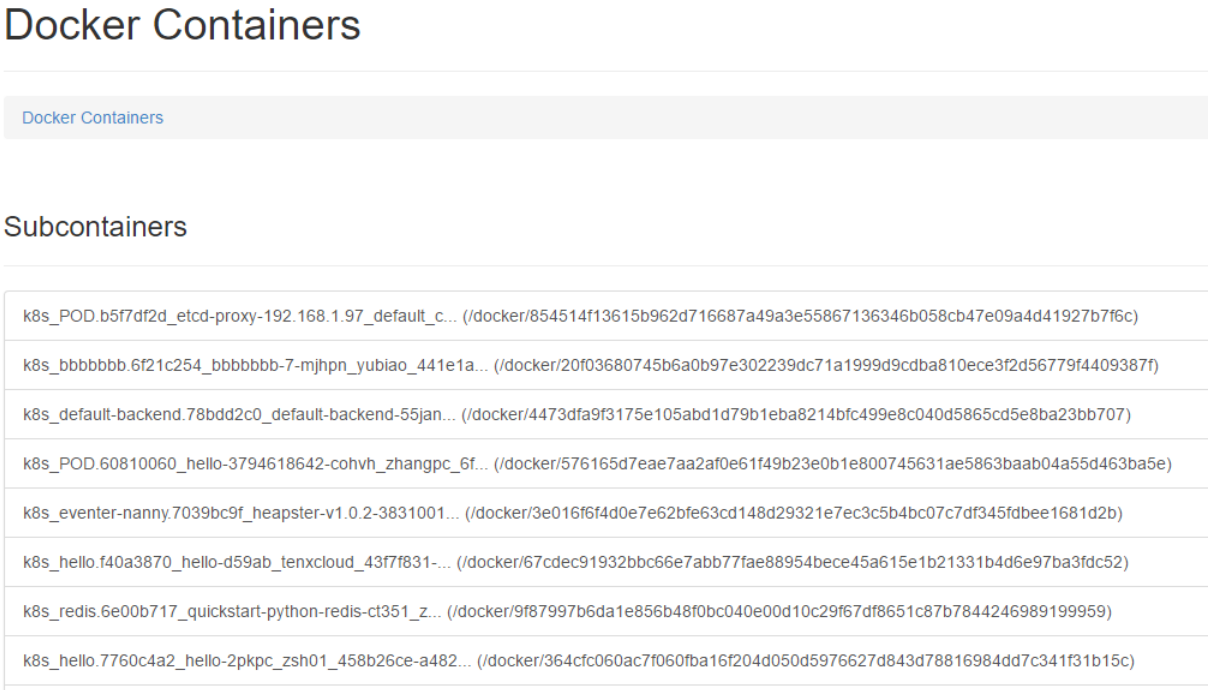

Resource Monitor

- Inside Heapster

```
[root@demo2 ~]# kubectl get svc --namespace=kube-system|grep heapster
heapster               10.0.0.142   <none>        80/TCP              24d
[root@demo2 ~]# curl 10.0.0.142/metrics
# HELP heapster_exporter_duration_microseconds Time spent exporting data to sink in microseconds
# TYPE heapster_exporter_duration_microseconds summary
heapster_exporter_duration_microseconds{exporter="InfluxDB Sink",quantile="0.5"} 3.459
heapster_exporter_duration_microseconds{exporter="InfluxDB Sink",quantile="0.9"} 3.692
heapster_exporter_duration_microseconds{exporter="InfluxDB Sink",quantile="0.99"} 3.692
heapster_exporter_duration_microseconds_sum{exporter="InfluxDB Sink"} 4858.171999999996
heapster_exporter_duration_microseconds_count{exporter="InfluxDB Sink"} 1540
heapster_exporter_duration_microseconds{exporter="Metric Sink",quantile="0.5"} 2.715
heapster_exporter_duration_microseconds{exporter="Metric Sink",quantile="0.9"} 3.735
heapster_exporter_duration_microseconds{exporter="Metric Sink",quantile="0.99"} 3.735
heapster_exporter_duration_microseconds_sum{exporter="Metric Sink"} 6019.538999999986
heapster_exporter_duration_microseconds_count{exporter="Metric Sink"} 1540
# HELP heapster_exporter_last_time_seconds Last time Heapster exported data since unix epoch in s
# TYPE heapster_exporter_last_time_seconds gauge
heapster_exporter_last_time_seconds{exporter="InfluxDB Sink"} 1.480572365e+09
heapster_exporter_last_time_seconds{exporter="Metric Sink"} 1.480572365e+09
# HELP heapster_kubelet_summary_request_duration_microseconds The Kubelet summary request latenci
# TYPE heapster_kubelet_summary_request_duration_microseconds summary
heapster_kubelet_summary_request_duration_microseconds{node="192.168.1.153",quantile="0.5"} 835
heapster_kubelet_summary_request_duration_microseconds{node="192.168.1.153",quantile="0.9"} 1180
heapster_kubelet_summary_request_duration_microseconds{node="192.168.1.153",quantile="0.99"} 118(
heapster_kubelet_summary_request_duration_microseconds_sum{node="192.168.1.153"} 1.852476e+06
heapster_kubelet_summary_request_duration_microseconds_count{node="192.168.1.153"} 1540
heapster_kubelet_summary_request_duration_microseconds{node="192.168.1.154",quantile="0.5"} 1187
```

```
[root@demo2 ~]# curl 10.0.0.142/api/v1/model/debug/allkeys
[
  "namespace:default/pod:my-nginx-1826603514-dhizz/container:my-nginx",
  "namespace:xinqiao/pod:hhhh-9vlit",
  "namespace:demo/pod:t5yw-2jrz5",
  "namespace:xinqiao",
  "namespace:default/pod:k8s-node-192.168.1.153/container:flannel",
  "node:192.168.1.154/container:kubelet",
  "node:192.168.1.154/container:docker-daemon",
  "namespace:zhongyc/pod:tomcat-nfzn6/container:tomcat",
  "namespace:xinqiao/pod:tttom-7kask/container:tttom",
  "namespace:demo",
  "namespace:kube-system/pod:fluentd-elasticsearch-192.168.1.153",
  "namespace:default/pod:curl-util",
  "namespace:zhongyc/pod:cmdb-z3rq8",
  "namespace:default",
  "namespace:kube-system/pod:fluentd-elasticsearch-192.168.1.153/container:fluentd-elasticsearch
  "namespace:default/pod:k8s-node-192.168.1.154/container:proxy",
  "namespace:demo/pod:myapp-p2xna/container:myapp",
  "namespace:kube-system/pod:heapster-23827483-91g0m/container:eventer-nanny",
  "namespace:demo/pod:myapp-ypckt",
  "namespace:kube-system/pod:heapster-23827483-91g0m",
  "namespace:default/pod:k8s-node-192.168.1.153/container:proxy",
  "namespace:default/pod:my-nginx-1826603514-4u69m/container:my-nginx",
  "node:192.168.1.153/container:kubelet",
  "namespace:xinqiao/pod:sql-nwb63",
  "namespace:default/pod:k8s-node-192.168.1.154/container:flannel",
```

Resource Monitor

- InfluxDB：开源的分布式时间序列数据库
  - ✓ 无结构，无限扩展的列
  - ✓ 单机、集群
  - ✓ 内置多样的统计函数
  - ✓ HTTP API
  - ✓ 类SQL的查询语句
  - ✓ 自带管理工具

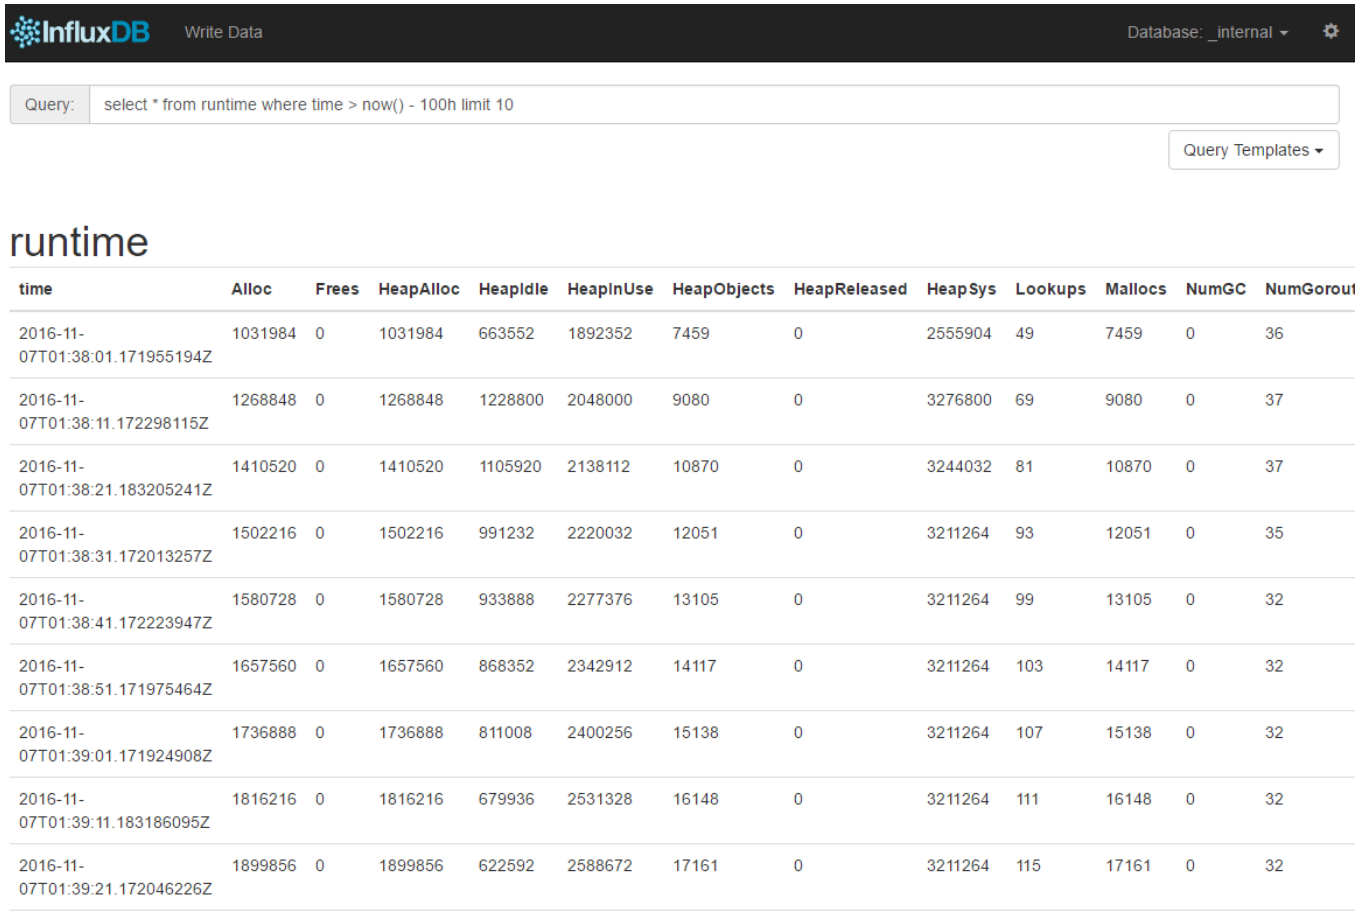

runtime

time	Alloc	Frees	HeapAlloc	HeapIdle	HeapInUse	HeapObjects	HeapReleased	HeapSys	Lookups	Mallocs	NumGC	NumGorout
2016-11-07T01:38:01.171955194Z	1031984	0	1031984	663552	1892352	7459	0	2555904	49	7459	0	36
2016-11-07T01:38:11.172298115Z	1268848	0	1268848	1228800	2048000	9080	0	3276800	69	9080	0	37
2016-11-07T01:38:21.183205241Z	1410520	0	1410520	1105920	2138112	10870	0	3244032	81	10870	0	37
2016-11-07T01:38:31.172013257Z	1502216	0	1502216	991232	2220032	12051	0	3211264	93	12051	0	35
2016-11-07T01:38:41.172223947Z	1580728	0	1580728	933888	2277376	13105	0	3211264	99	13105	0	32
2016-11-07T01:38:51.171975464Z	1657560	0	1657560	868352	2342912	14117	0	3211264	103	14117	0	32
2016-11-07T01:39:01.171924908Z	1736888	0	1736888	811008	2400256	15138	0	3211264	107	15138	0	32
2016-11-07T01:39:11.183186095Z	1816216	0	1816216	679936	2531328	16148	0	3211264	111	16148	0	32
2016-11-07T01:39:21.172046226Z	1899856	0	1899856	622592	2588672	17161	0	3211264	115	17161	0	32

Resource Monitor

- Grafana：开源的可定制、易扩展的强大丰富的数据可视化工具
  - ✓ Html/js应用
  - ✓ 灵活丰富，多种风格
  - ✓ 定义数据源
  - ✓ 配置展示页面

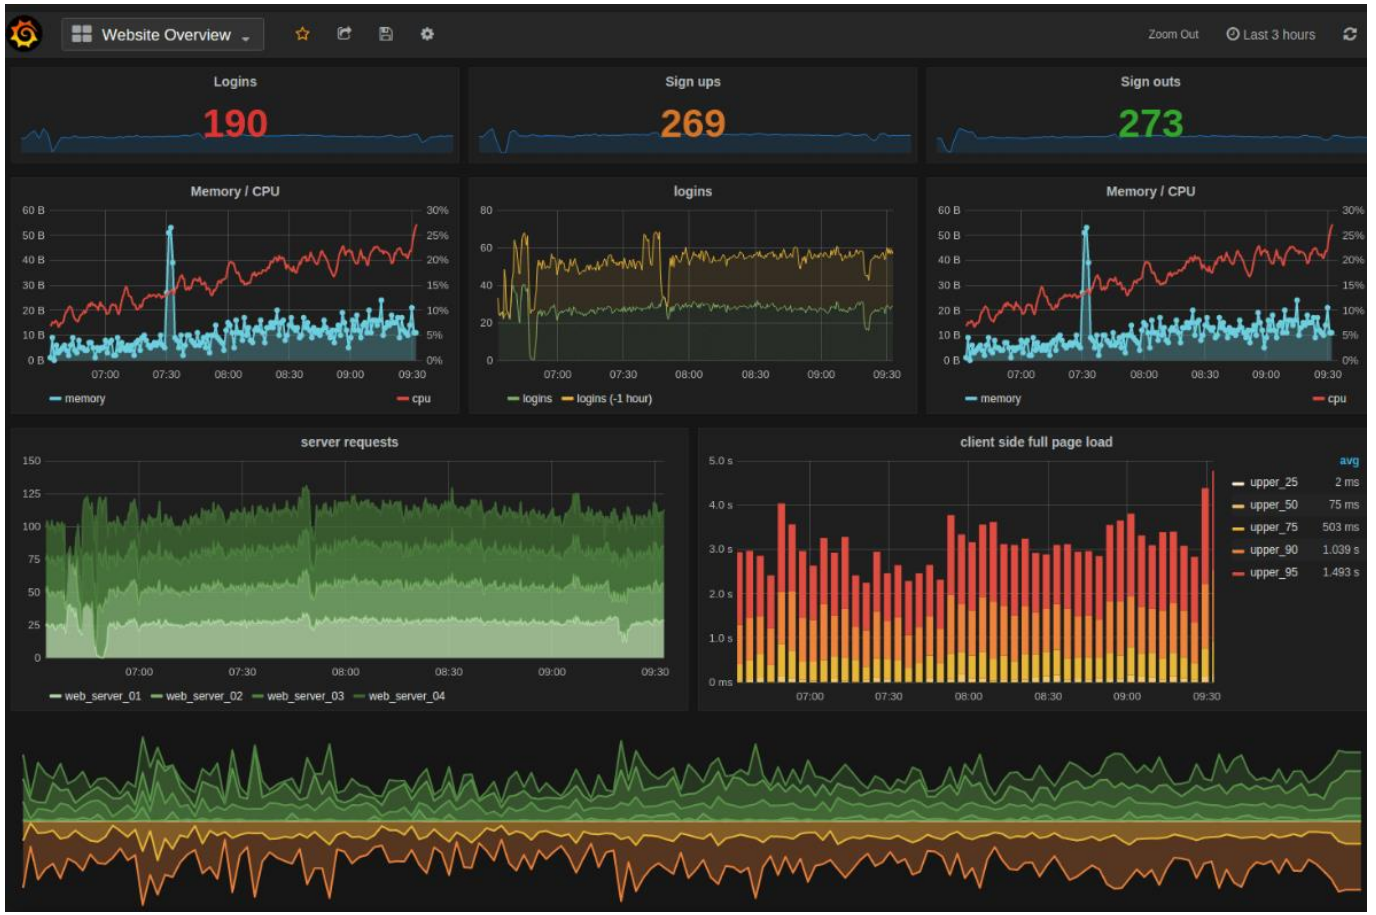

Resource Monitor

- Prometheus：开源的监控工具、易扩展的强大丰富的数据可视化工具
  - ✓ 支持k8s metrics数据格式
  - ✓ 支持k8s api进行服务发现
  - ✓ 多维度监控
  - ✓ 自定义告警

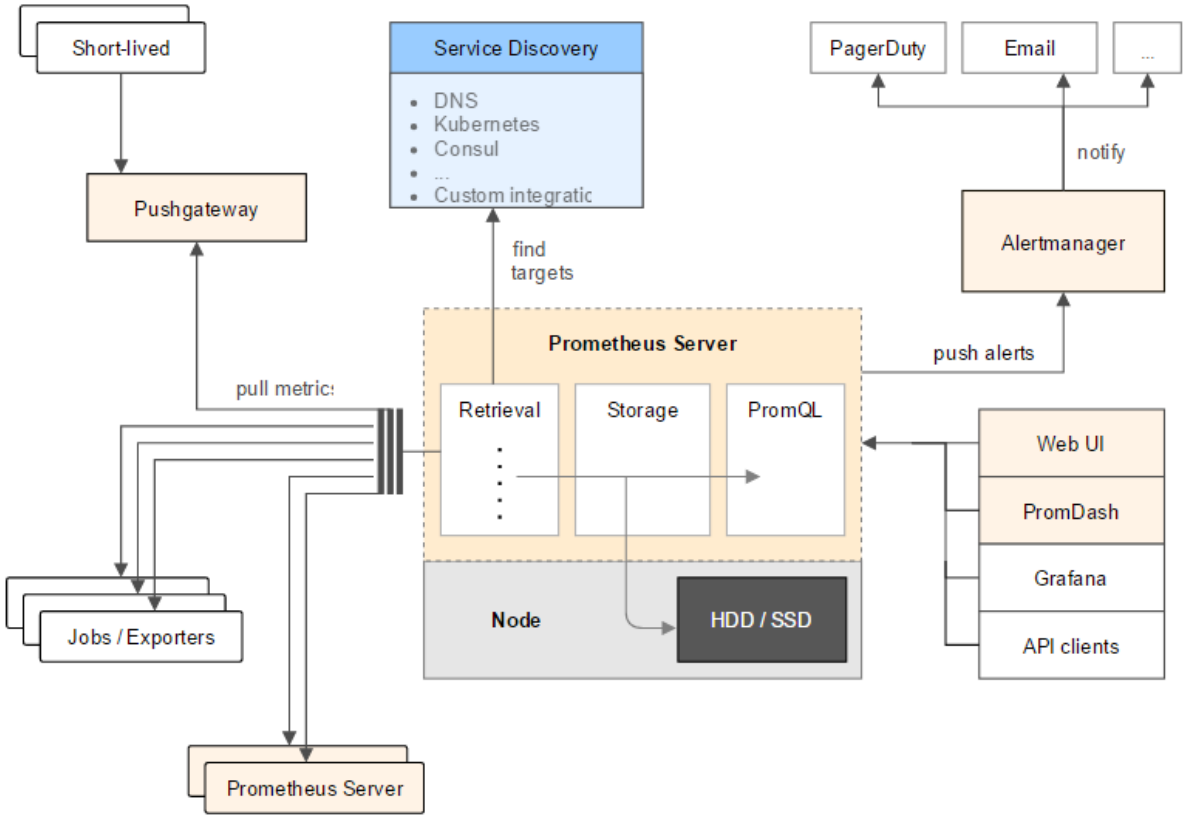

Service Monitor

- Kubernetes提供对Pod、容器运行状况的监控，通过RC保持运行数量

```
[root@demo2 ~]# kubectl get rc myapp --namespace=demo -ojson
{
    "kind": "ReplicationController",
    "apiVersion": "v1",
    "metadata": {
        "name": "myapp",
        "namespace": "demo",
        "selfLink": "/api/v1/namespaces/demo/replicationcontrollers/myapp",
        "uid": "36c1678d-b77d-11e6-be61-005056852019",
        "resourceVersion": "443757",
        "generation": 2,
        "creationTimestamp": "2016-12-01T04:18:29Z",
        "labels": {
            "name": "myapp"
        }
    },
    "spec": {
        "replicas": 3,
        "selector": {
            "name": "myapp"
        },
```

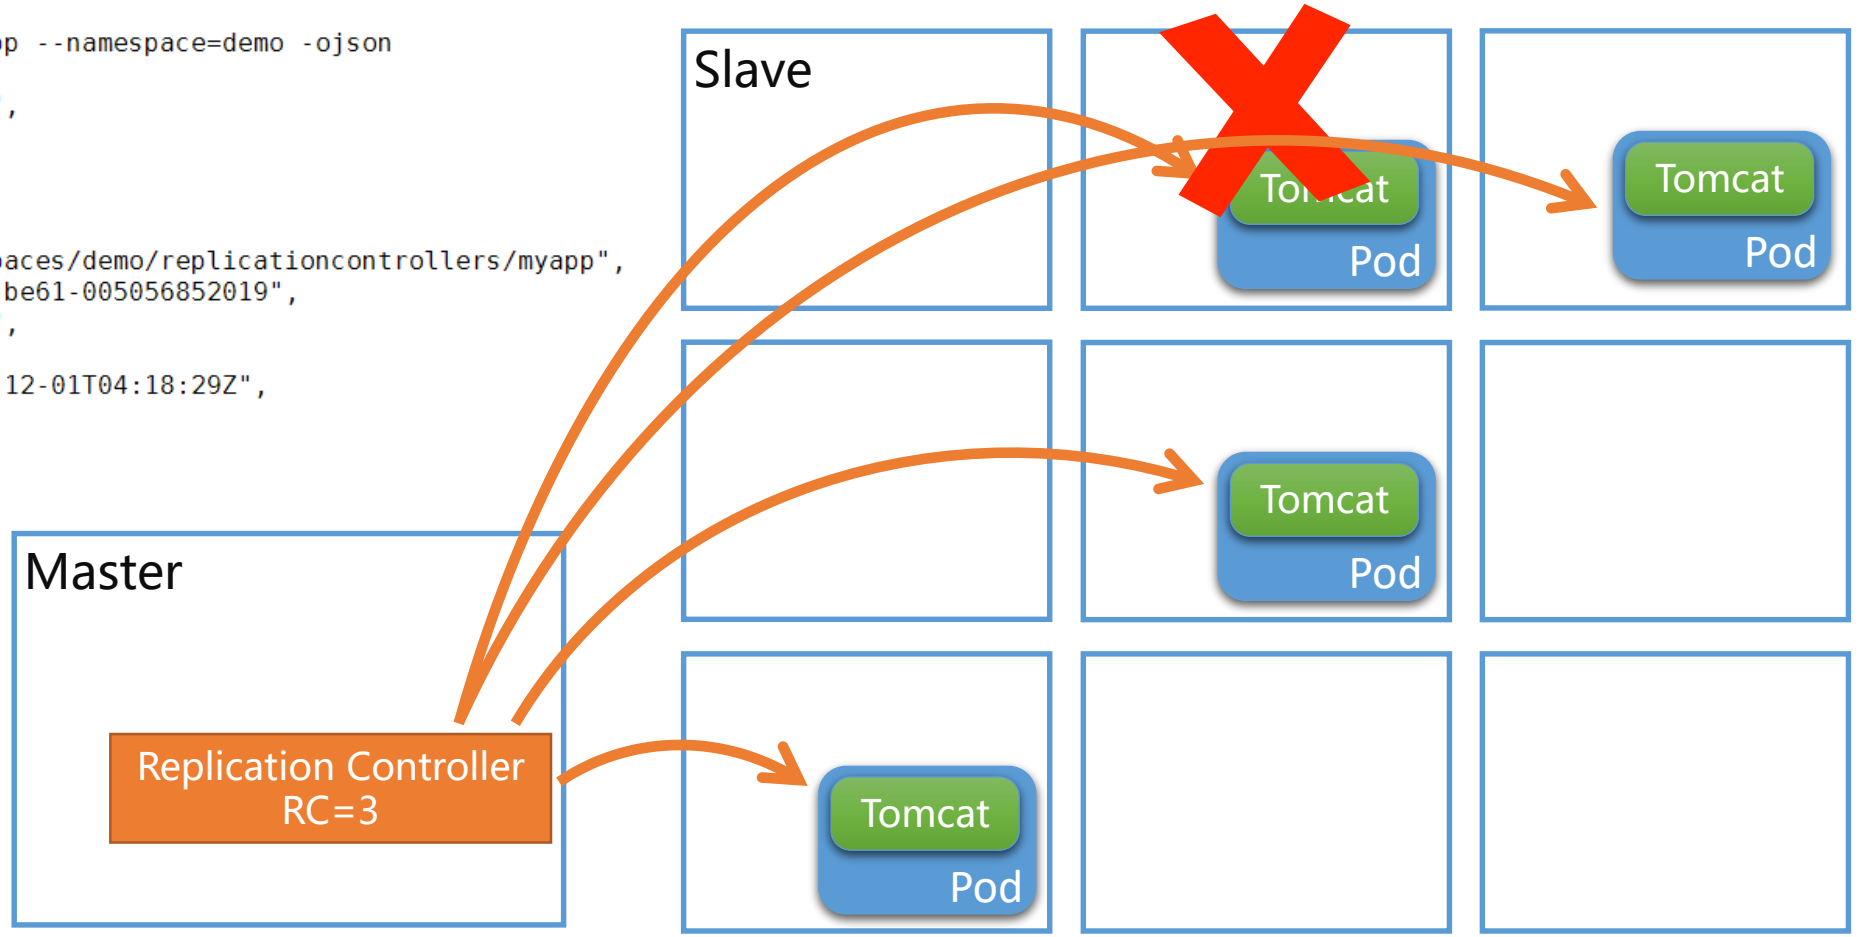

Service Monitor

- Container Probes提供HTTP、TCP的健康监控，保证应用的正常运行

```
[root@demo2 ~]# kubectl get po myapp-p2xna --namespace=demo -ojson
{
    "kind": "Pod",
    "apiVersion": "v1",
    "metadata": {
        "name": "myapp-p2xna",
        "generateName": "myapp-",
        "namespace": "demo",
        "selfLink": "/api/v1/namespaces/demo/pods/myapp-p2xna",
        "uid": "36c23bfa-b77d-11e6-be61-005056852019",
        "resourceVersion": "443485",
        "creationTimestamp": "2016-12-01T04:18:29Z",
        "labels": {
            "name": "myapp"
        },
```

```
"livenessProbe": {
    "httpGet": {
        "path": "/health",
        "port": 8080,
        "scheme": "HTTP"
    },
    "initialDelaySeconds": 10,
    "timeoutSeconds": 3,
    "periodSeconds": 5,
    "successThreshold": 1,
    "failureThreshold": 3
},
```

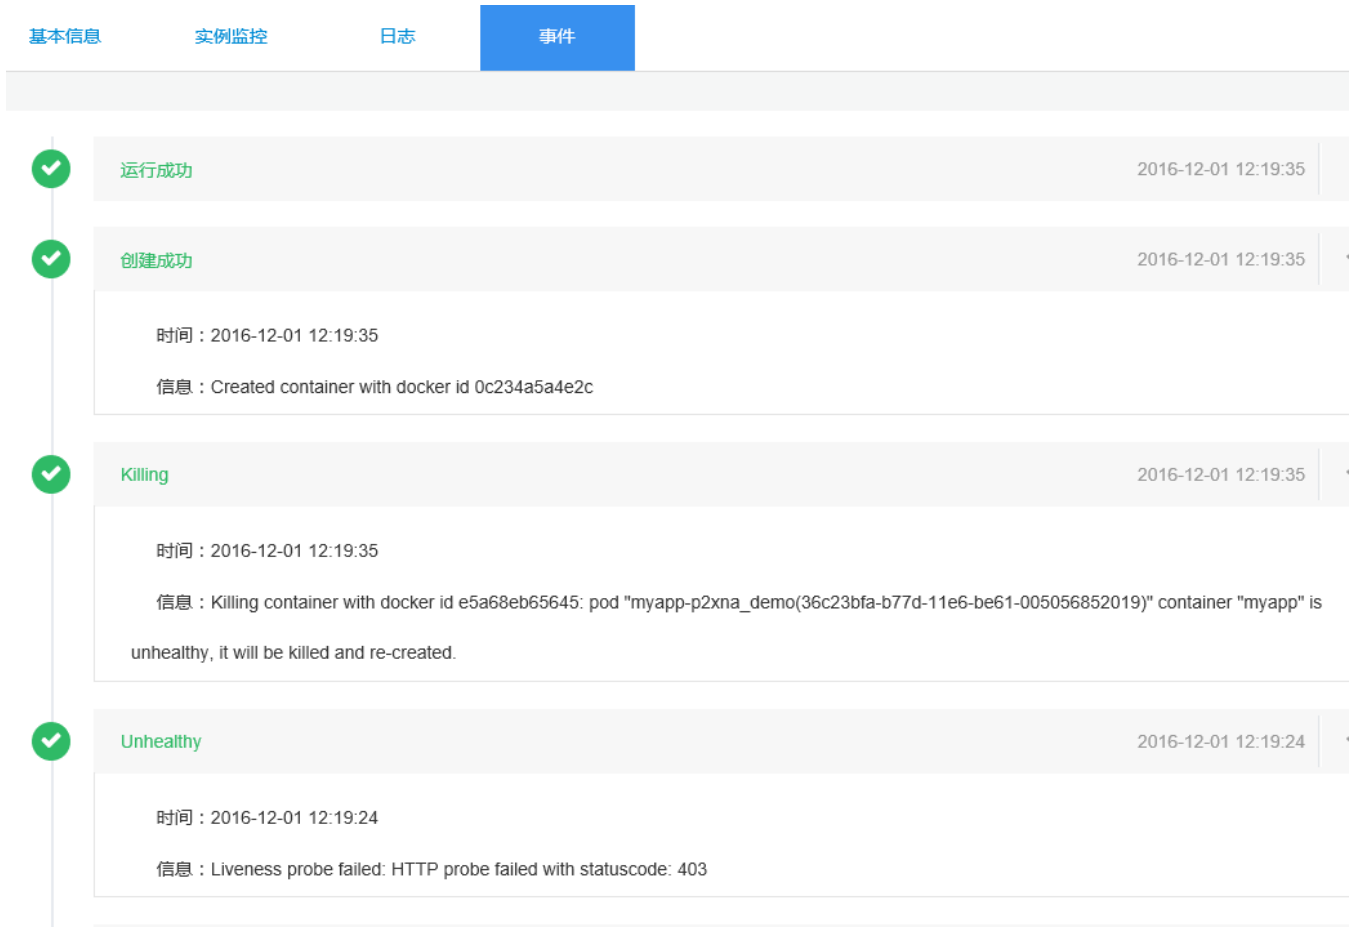

Agenda

- Kubernetes Summary
  - Features
  - Architecture
- Kubernetes Monitor
  - Resource Monitor
  - Service Monitor
- Kubernetes Logging
  - Fluentd
  - Elasticsearch
  - Flow
- 时速云平台实践

Fluentd

- Fluentd: 开源的，日志收集工具.

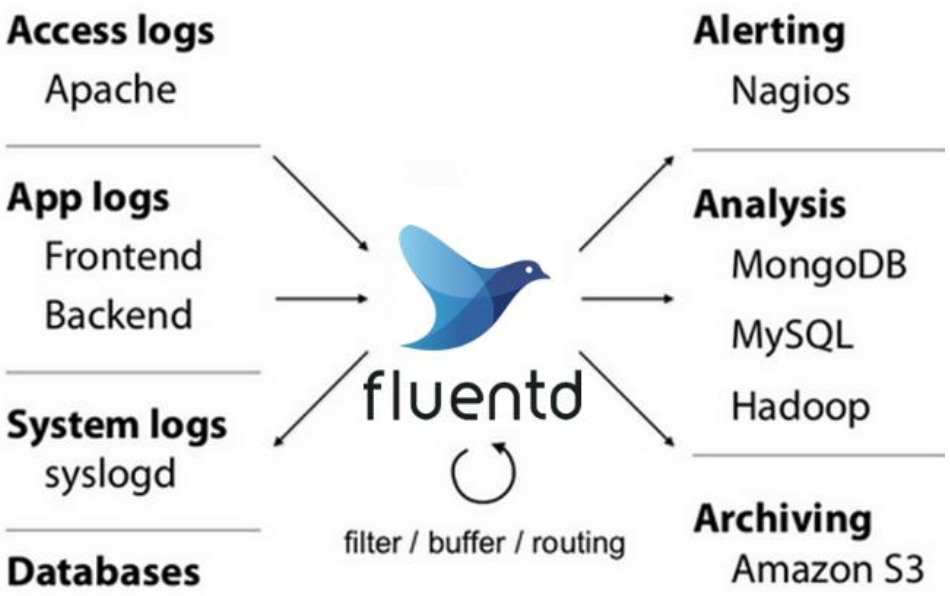

- ✓ Unified Logging Layer
  - Json 格式
  - 可靠传输
  - 易收集、易保存
- ✓ 统一接口，降低复杂度
- ✓ 插件丰富、可扩展
- ✓ 生产实践

Elasticsearch

- Elasticsearch: 开源的、实时的分布式搜索和分析引擎
  - 分布式存储
  - 分布式搜索
  - 处理PB级别数据
  - Restful API
  - 集群编排/单机 + 持久化存储 on kubernetes

Flow

- How to work in Kubernetes

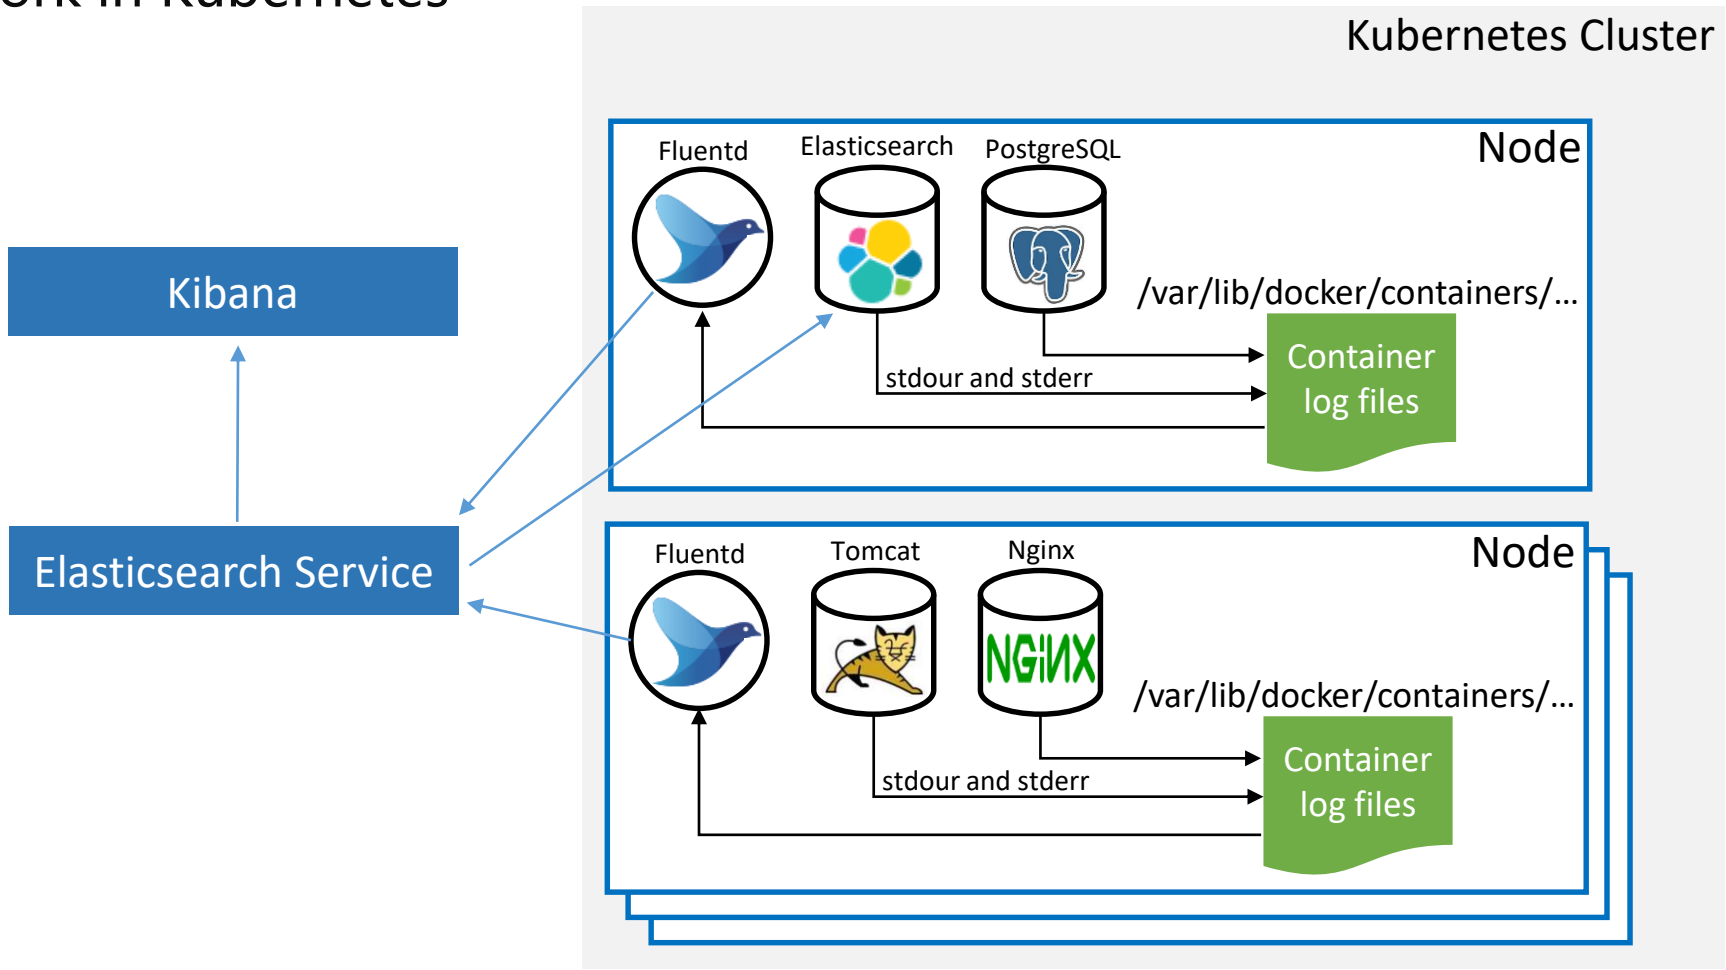

时速云平台实践

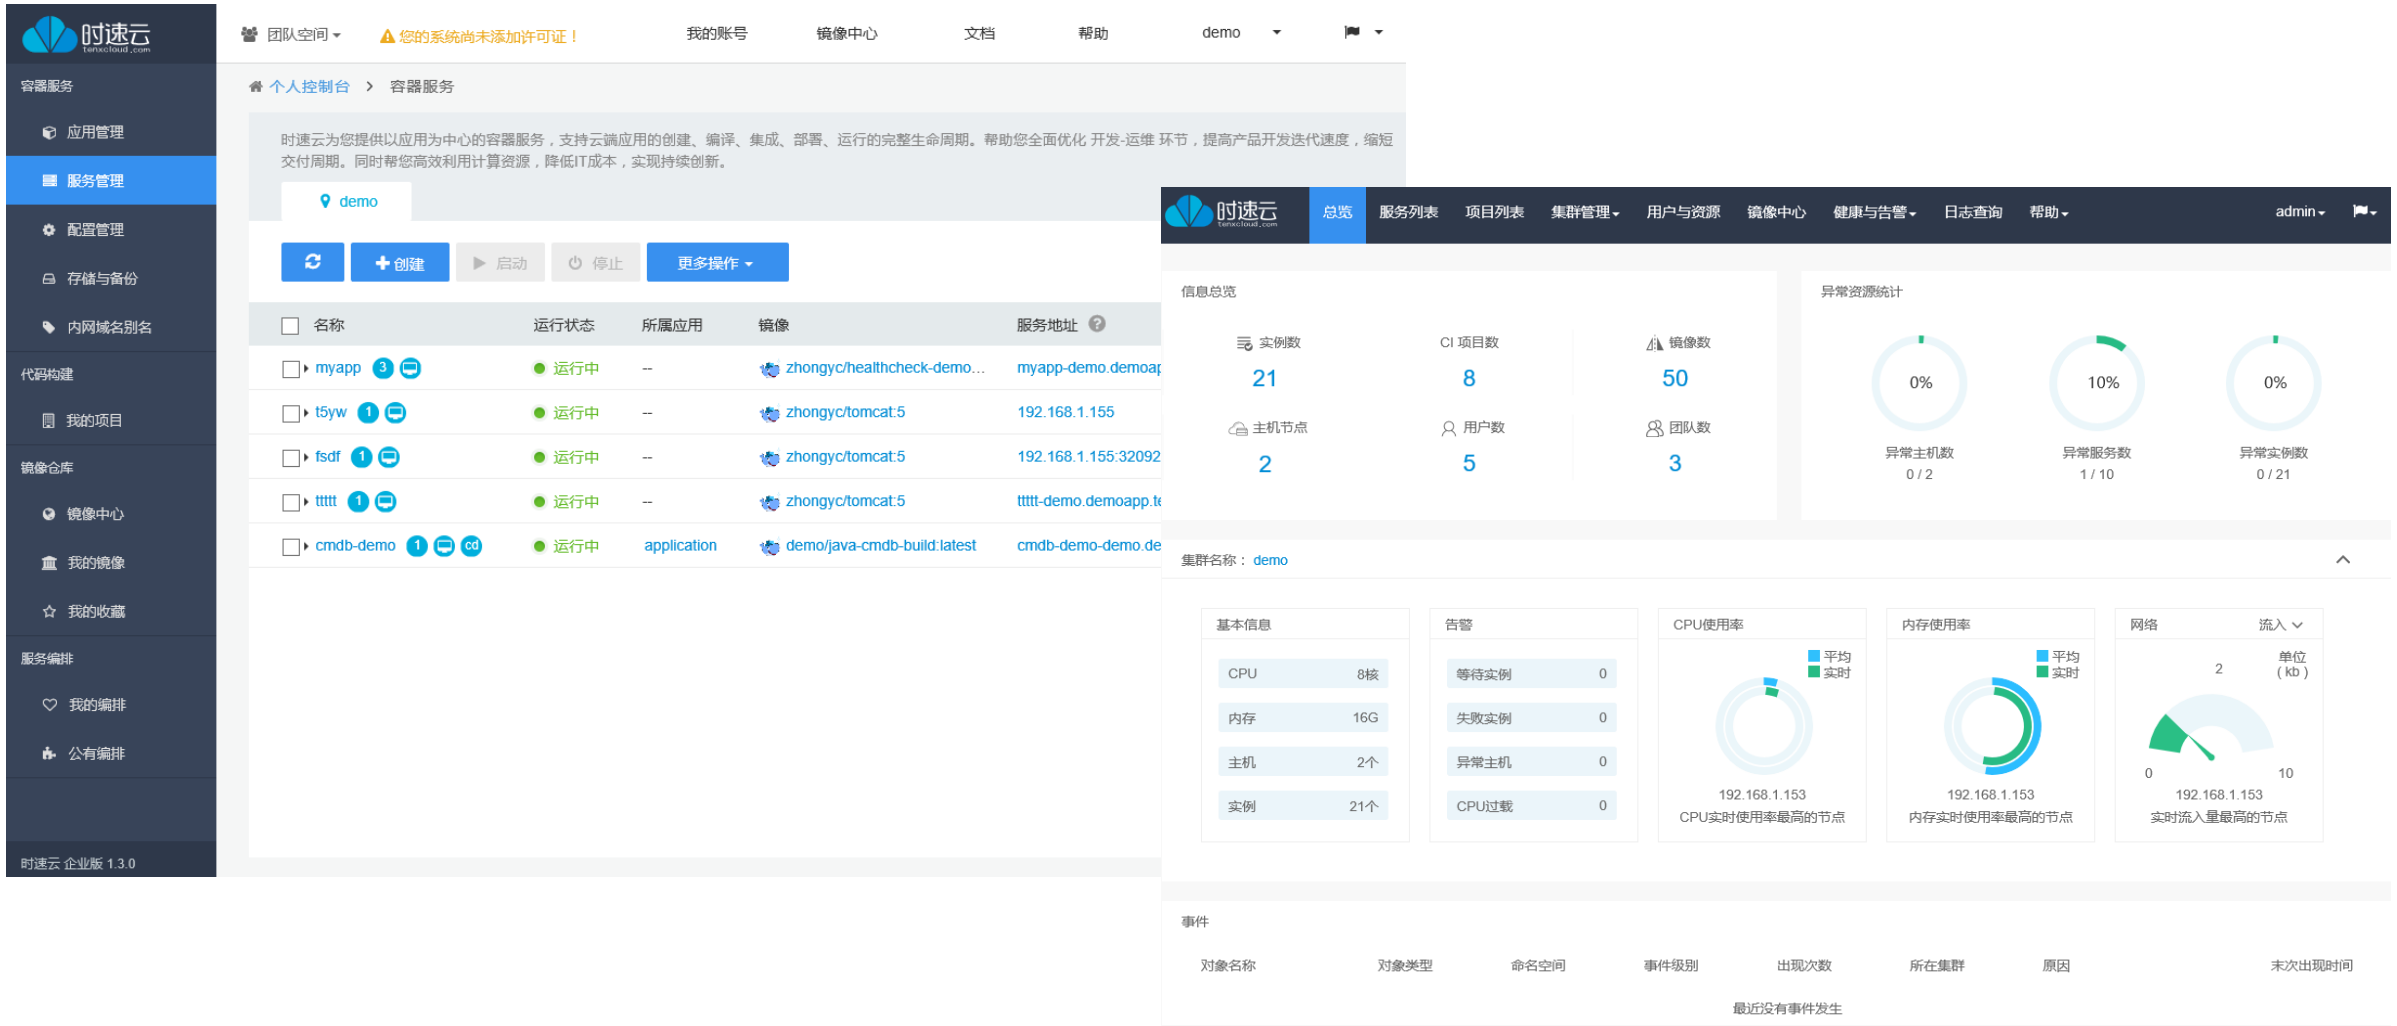
时速云
团队空间
您的系统尚未添加许可证！
我的账号
镜像中心
文档
帮助
demo
容器服务
应用管理
服务管理
配置管理
存储与备份
内网域名别名
代码构建
我的项目
镜像仓库
镜像中心
我的镜像
我的收藏
服务编排
我的编排
公有编排
时速云 企业版 1.3.0
个人控制台 > 容器服务
时速云为您提供以应用为中心的容器服务，支持云端应用的创建、编译、集成、部署、运行的完整生命周期。帮助您全面优化 开发-运维 环节，提高产品开发迭代速度，缩短交付周期。同时帮您高效利用计算资源，降低IT成本，实现持续创新。
demo
创建
启动
停止
更多操作
名称
运行状态
所属应用
镜像
服务地址
myapp
运行中
zhongyc/healthcheck-demo...
t5yw
运行中
zhongyc/tomcat:5
192.168.1.155
fsdf
运行中
zhongyc/tomcat:5
192.168.1.155:32092
ttttt
运行中
zhongyc/tomcat:5
cmdb-demo
运行中
application
demo/java-cmdb-build:latest
总览
服务列表
项目列表
集群管理
用户与资源
镜像中心
健康与告警
日志查询
帮助
admin
信息总览
实例数
21
CI 项目数
8
镜像数
50
主机节点
2
用户数
5
团队数
3
异常资源统计
0%
异常主机数
0 / 2
10%
异常服务数
1 / 10
0%
异常实例数
0 / 21
集群名称： demo
基本信息
CPU 8核
内存 16G
主机 2个
实例 21个
告警
等待实例 0
失败实例 0
异常主机 0
CPU过载 0
CPU使用率
平均
实时
192.168.1.153
CPU实时使用率最高的节点
内存使用率
平均
实时
192.168.1.153
内存实时使用率最高的节点
网络
流入
单位
(kb)
2
0
10
192.168.1.153
实时流入量最高的节点
事件
对象名称
对象类型
命名空间
事件级别
出现次数
所在集群
原因
末次出现时间
最近没有事件发生

谢谢关注

联系我们：400-626-1876
地址：北京市海淀区中关村E世界C座8层881
邮箱：service@tenxcloud.com
网址：www.tenxcloud.com