Cloud Native

- 容器生态
- DevOps
- 持续交付
- 微服务

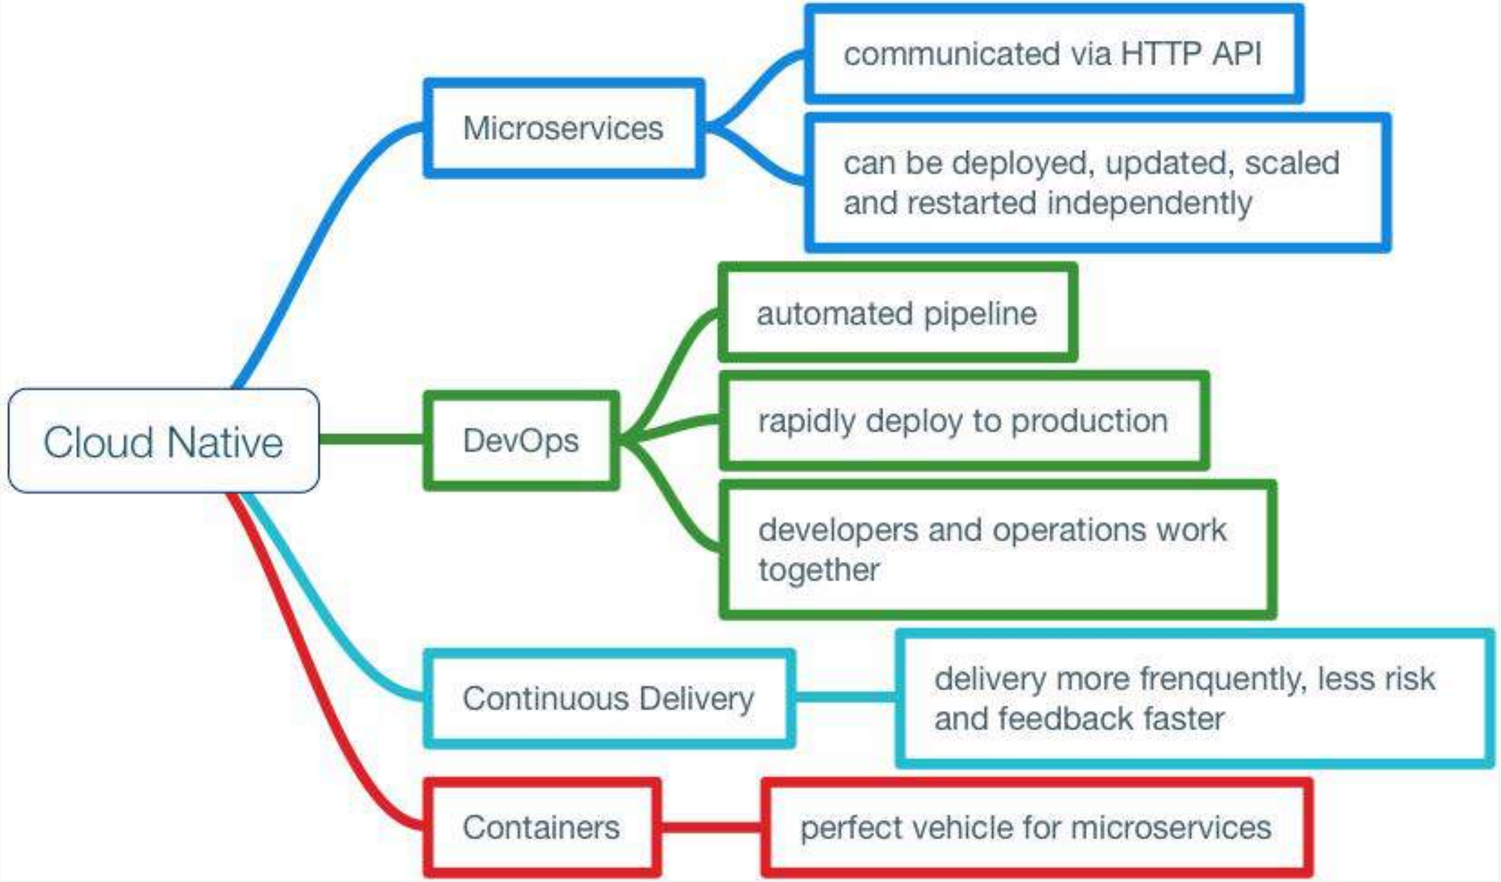

Cloud Native Features

Container packaged

Running applications and processes in software containers as an isolated unit of application deployment, and as a mechanism to achieve high levels of resource isolation.

Dynamically managed

Actively scheduled and actively managed by a central orchestrating process.

Microservices oriented

Loosely coupled with dependencies explicitly described (e.g., through service endpoints).

Cloud Native hosted projects

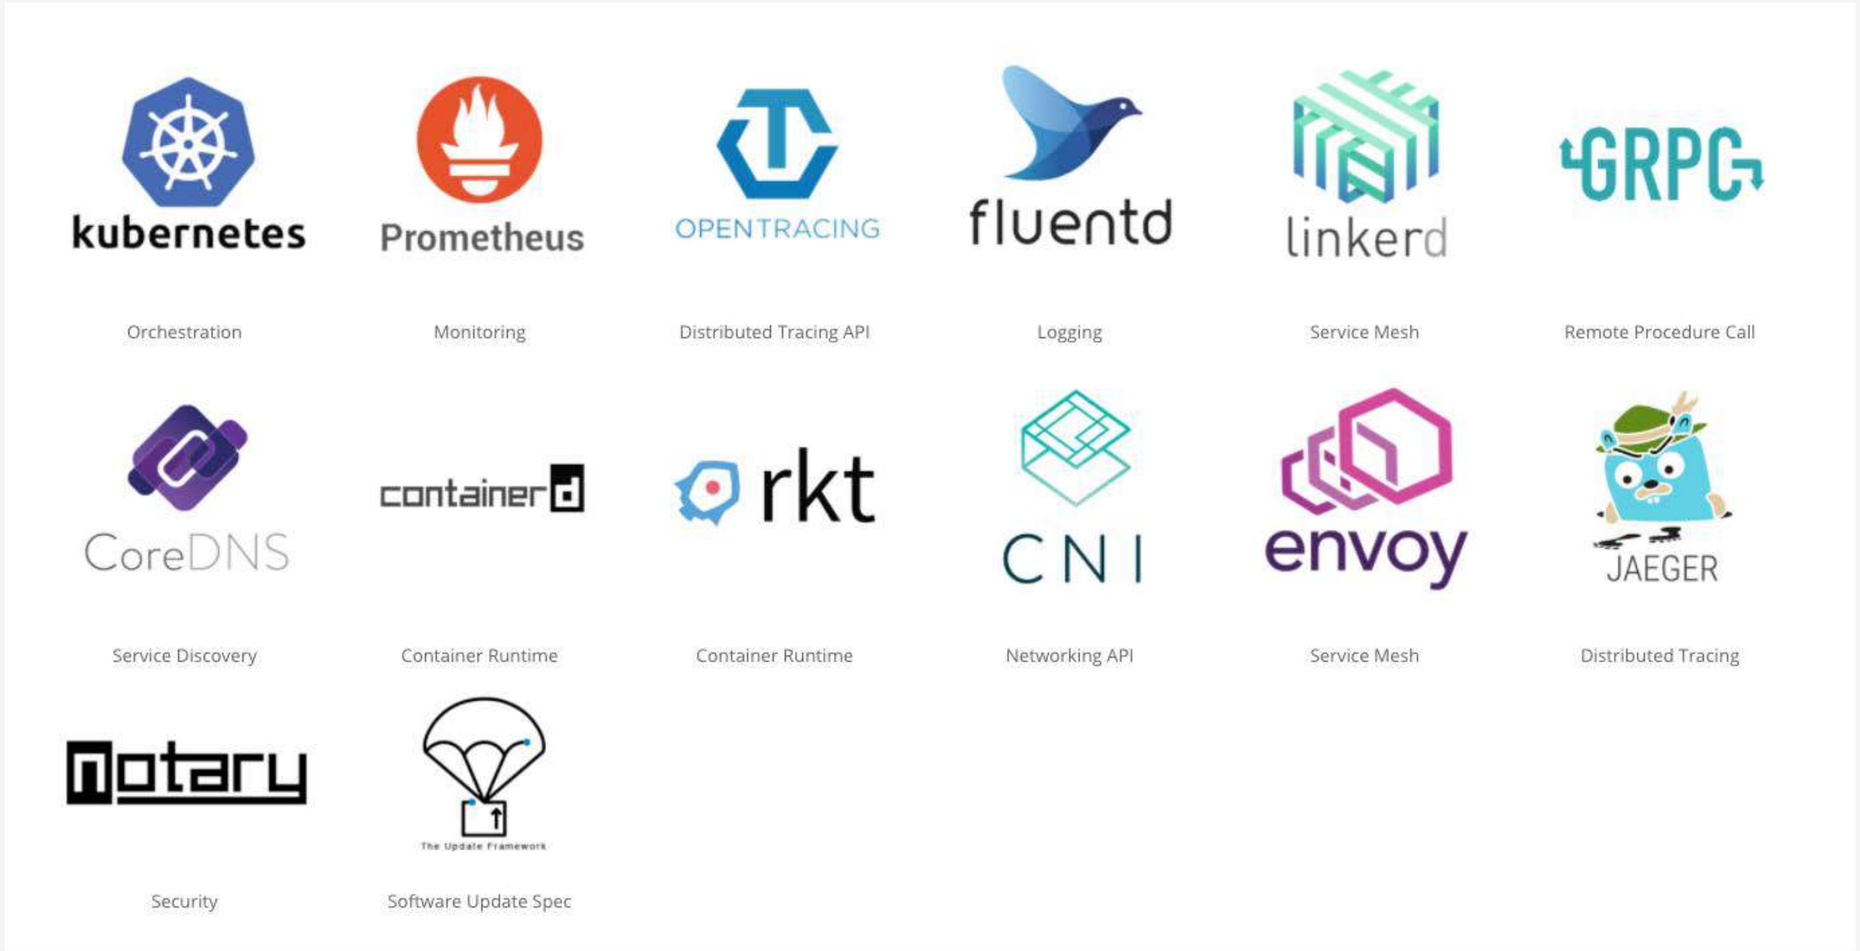

source https://cncf.io

Building Cloud Native Architecture with Kubernetes

source https://jimmysong.io

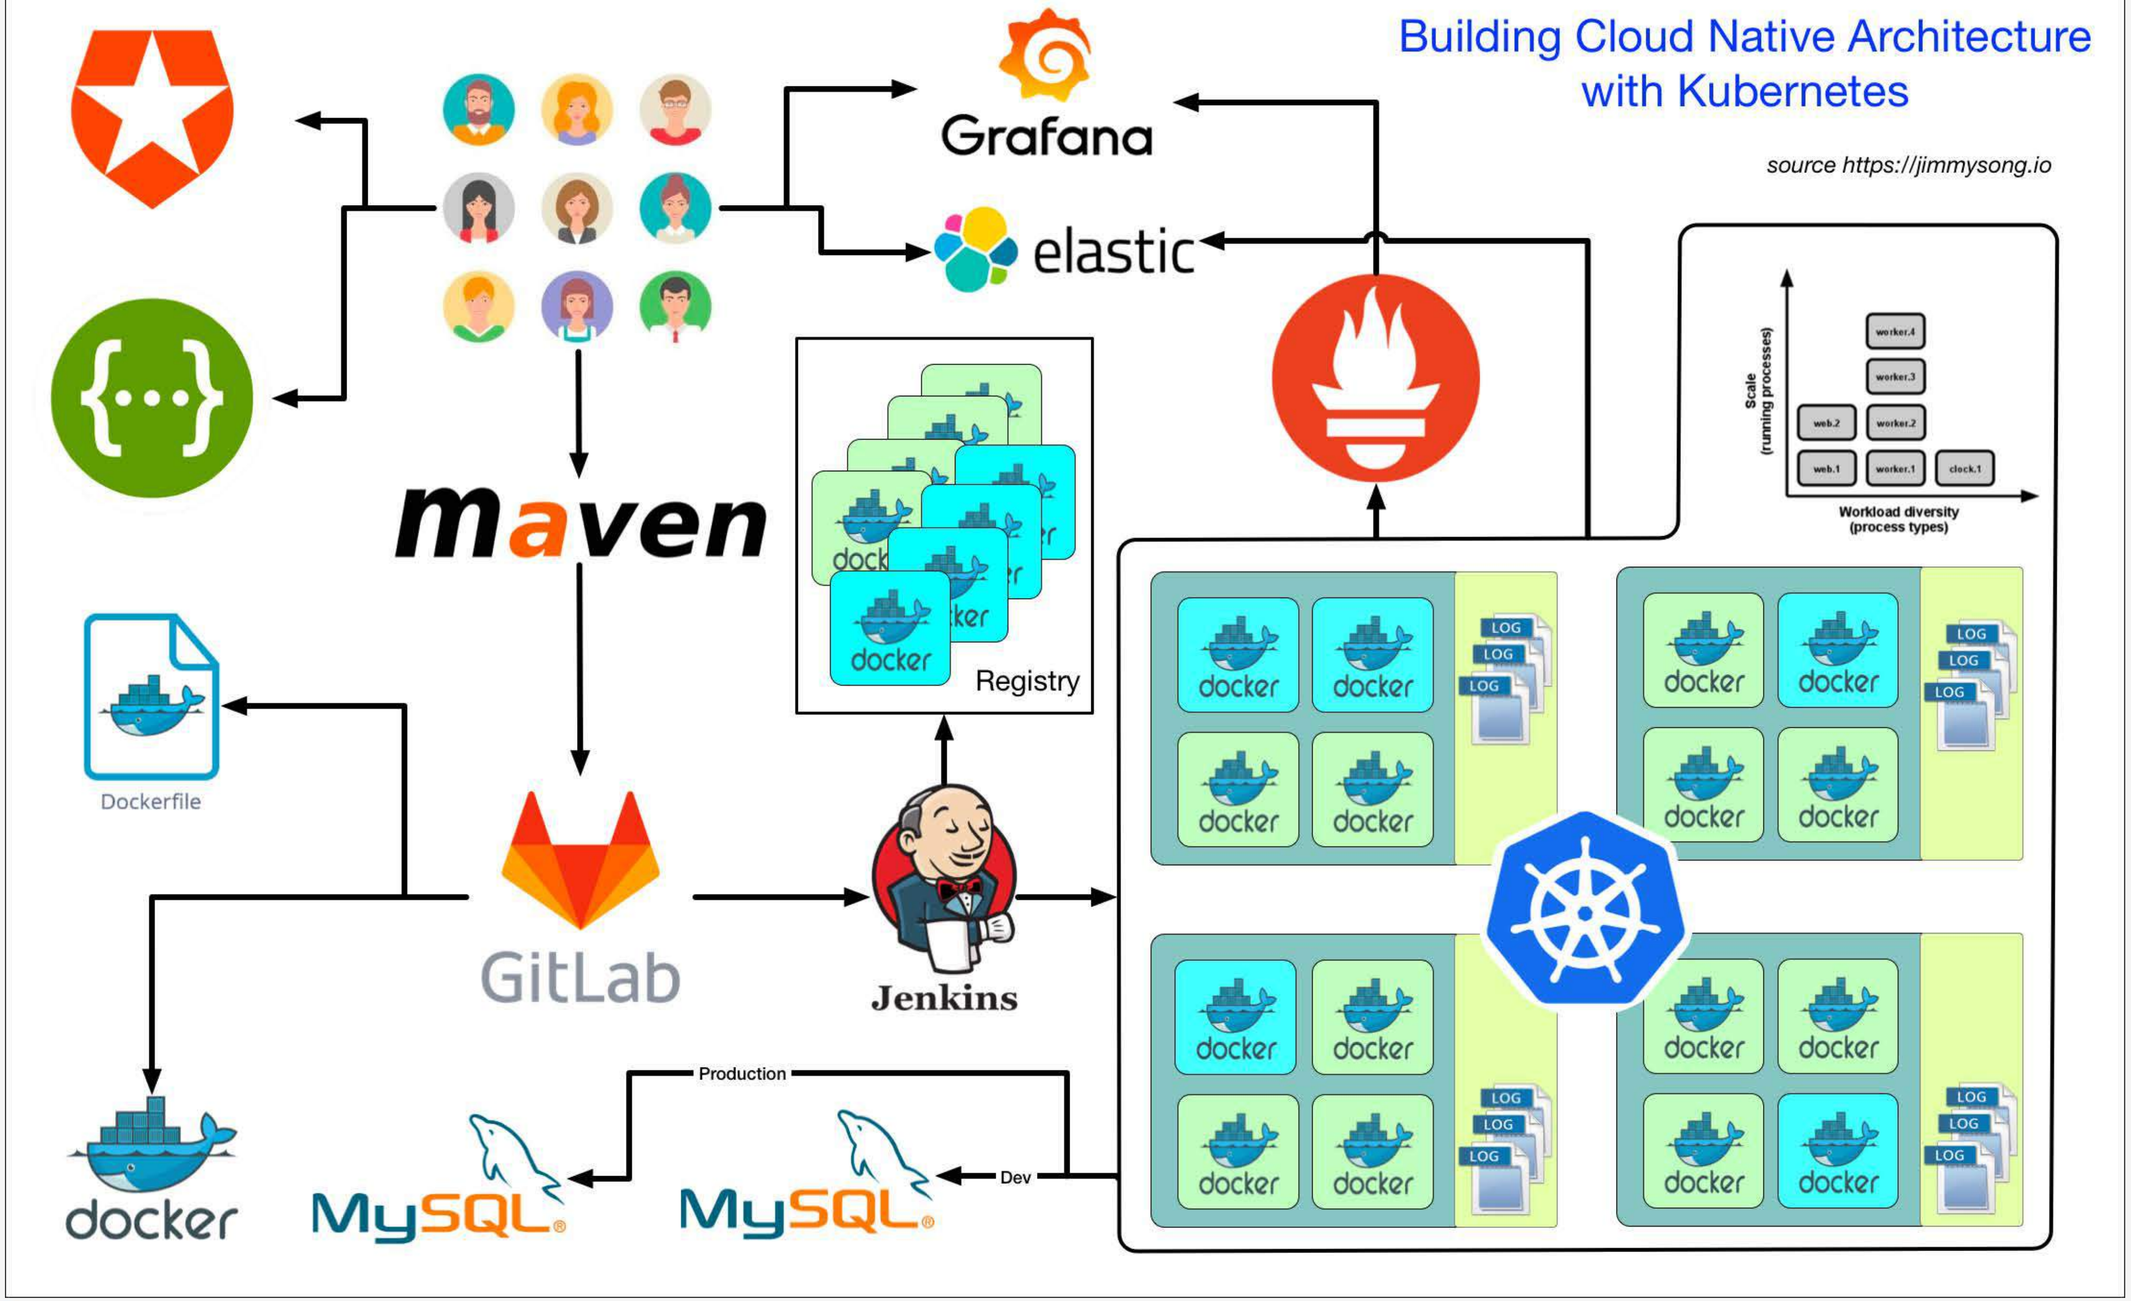

12 factors

1. 基准代码
2. 依赖管理
3. 配置
4. 后端服务
5. 构建，发布，运行
6. 无状态进程
7. 端口绑定
8. 并发
9. 易处理
10. 开发环境与线上环境等价
11. 日志作为事件流
12. 管理进程

Additional

- API声明管理
- 认证和授权
- 监控与告警

source https://12factor.net/

How to create a Cloud Native app

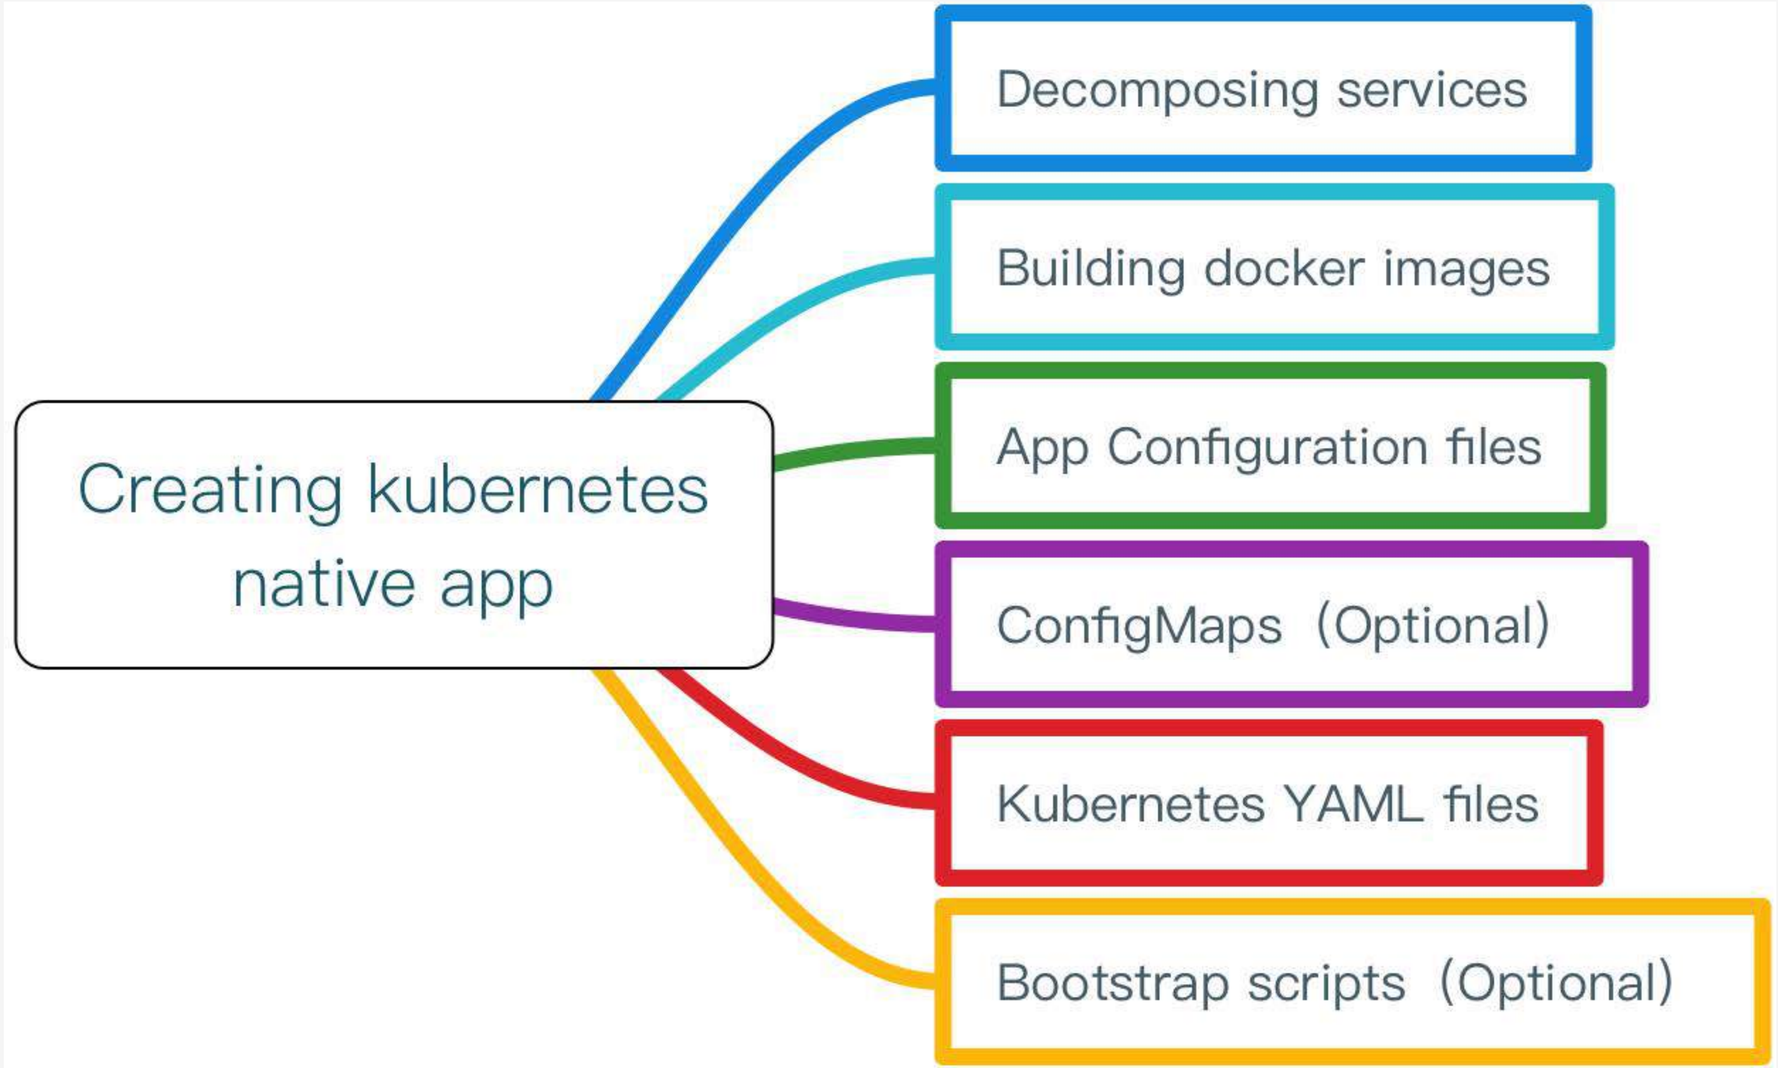

Move to Cloud

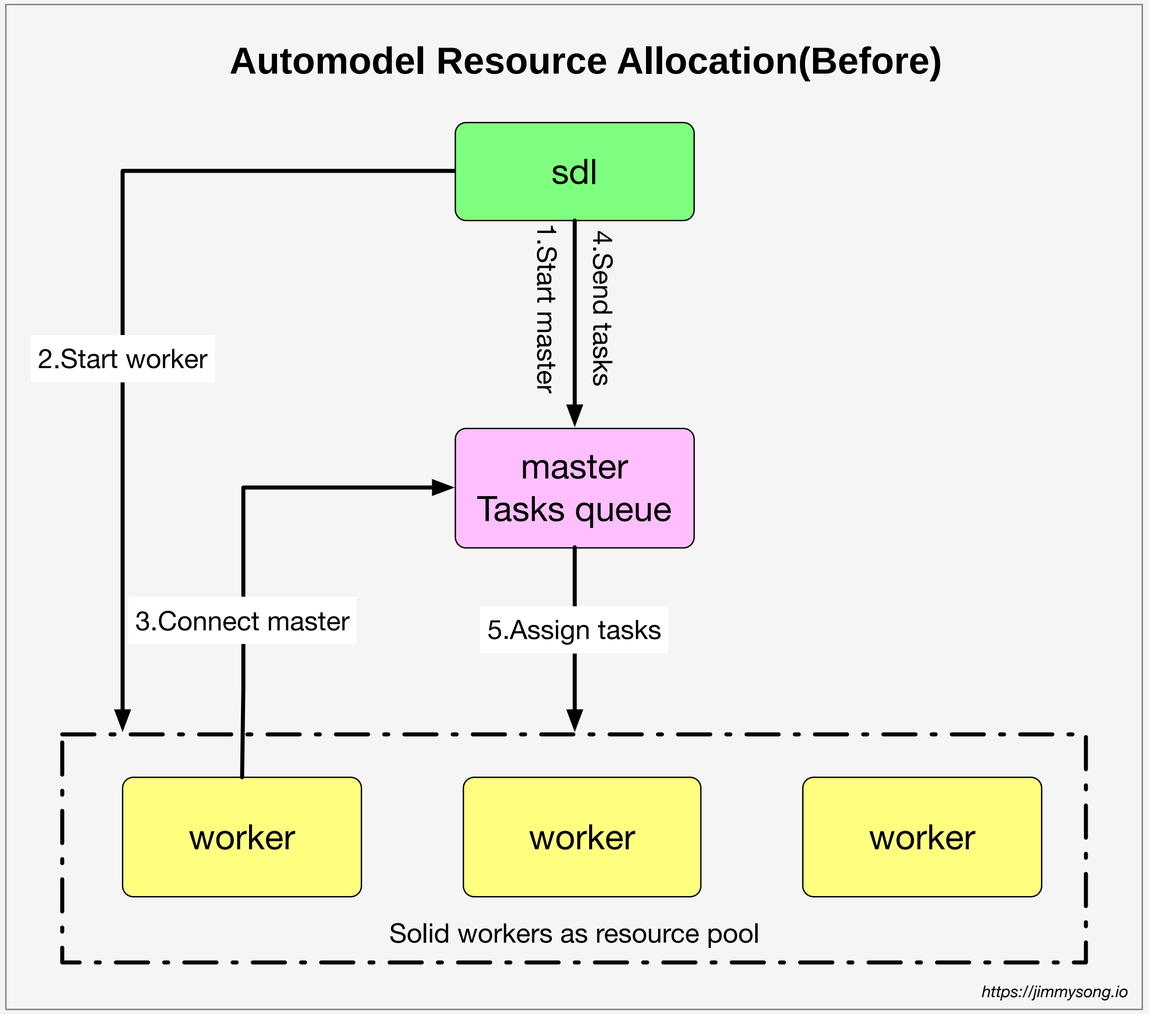

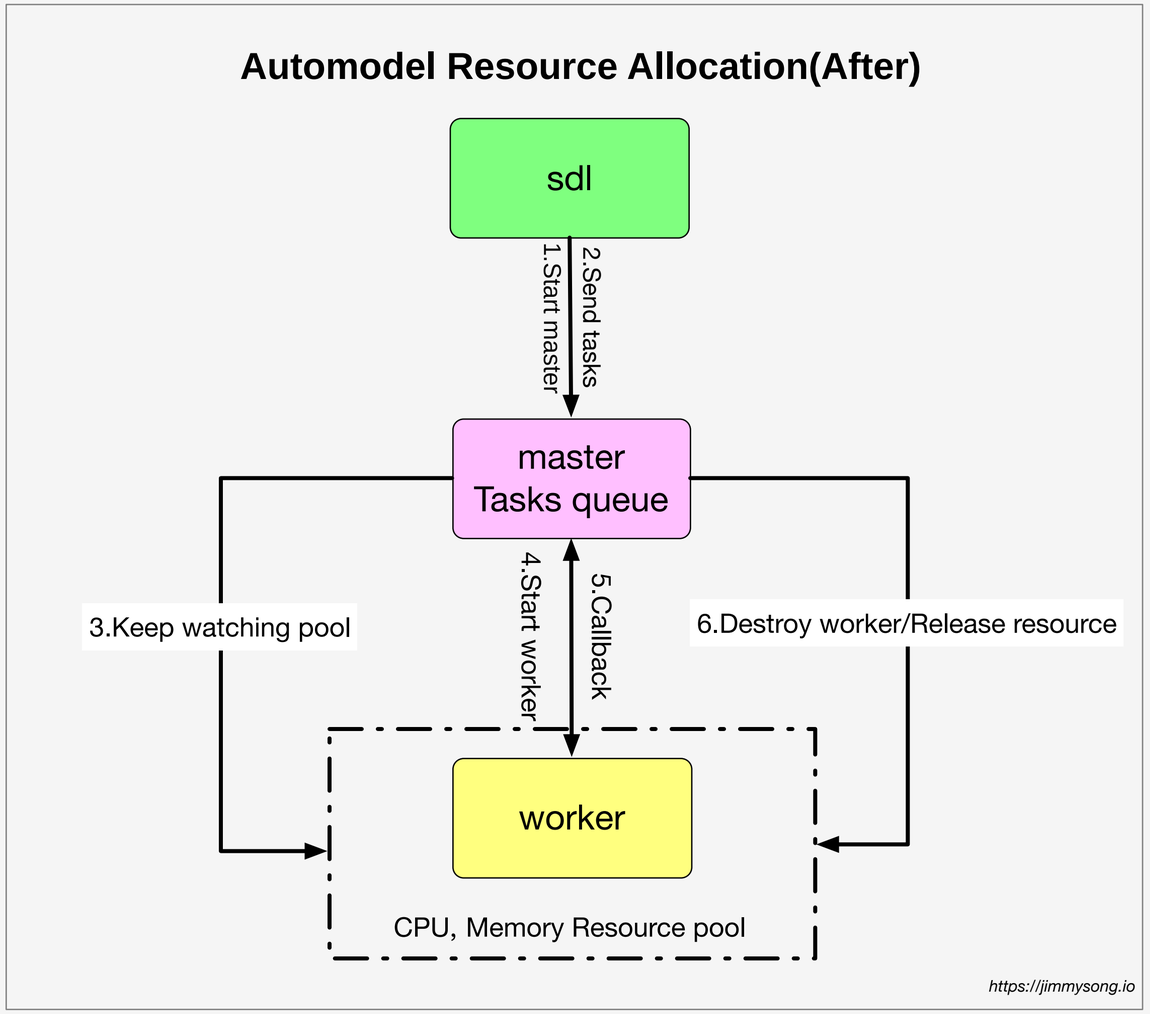

CI/CD

- Jenkins
- Harbor
- PowerDNS
- EFK

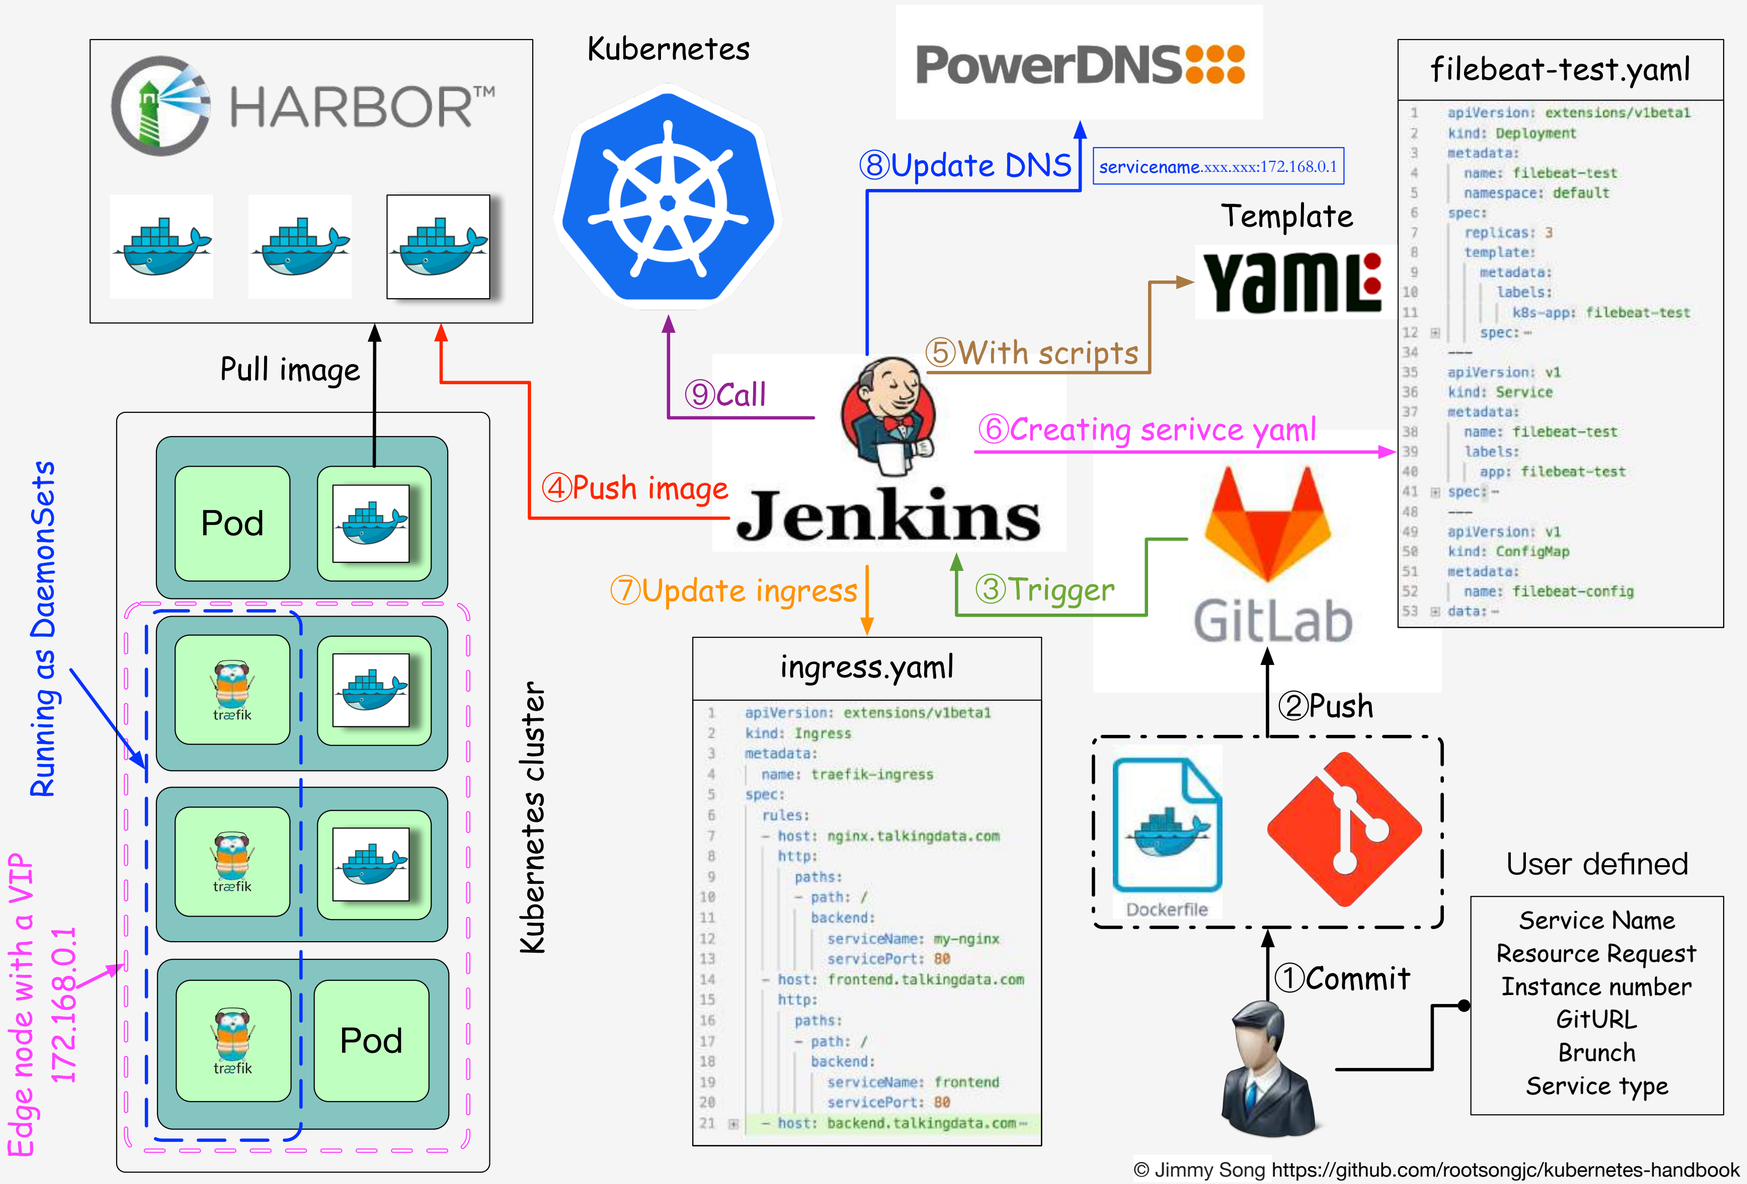

Edge node

- 边缘节点
- 流量出入口
- host + path

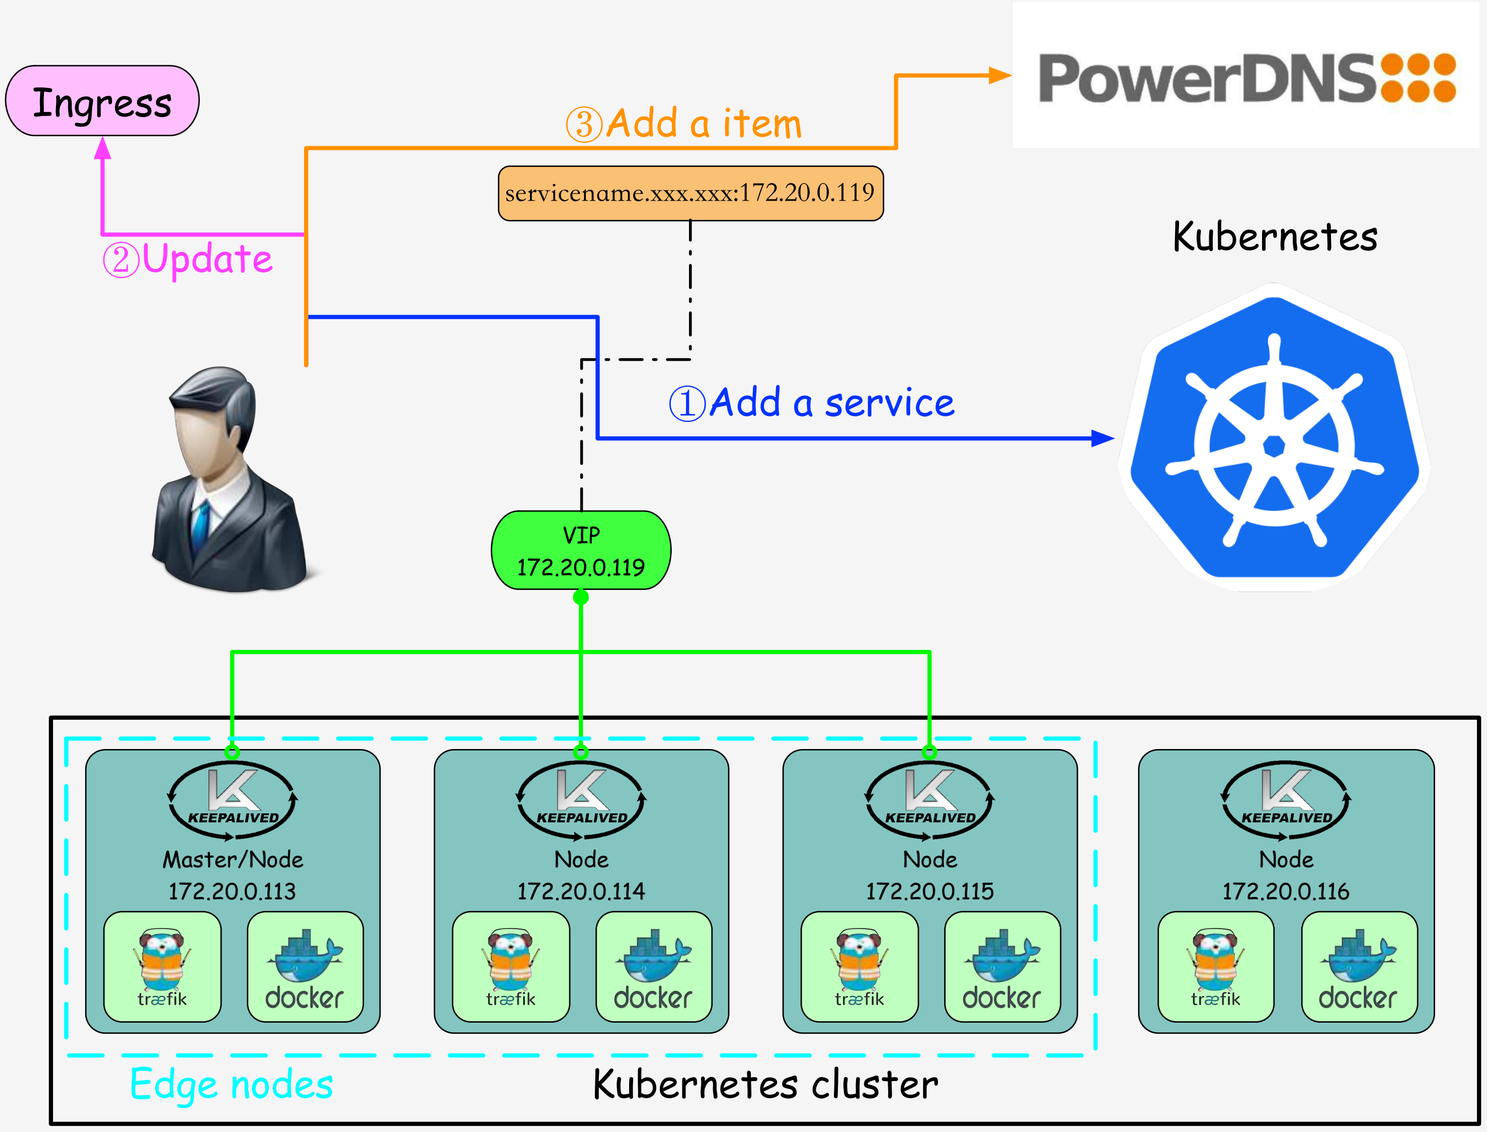

Monitoring

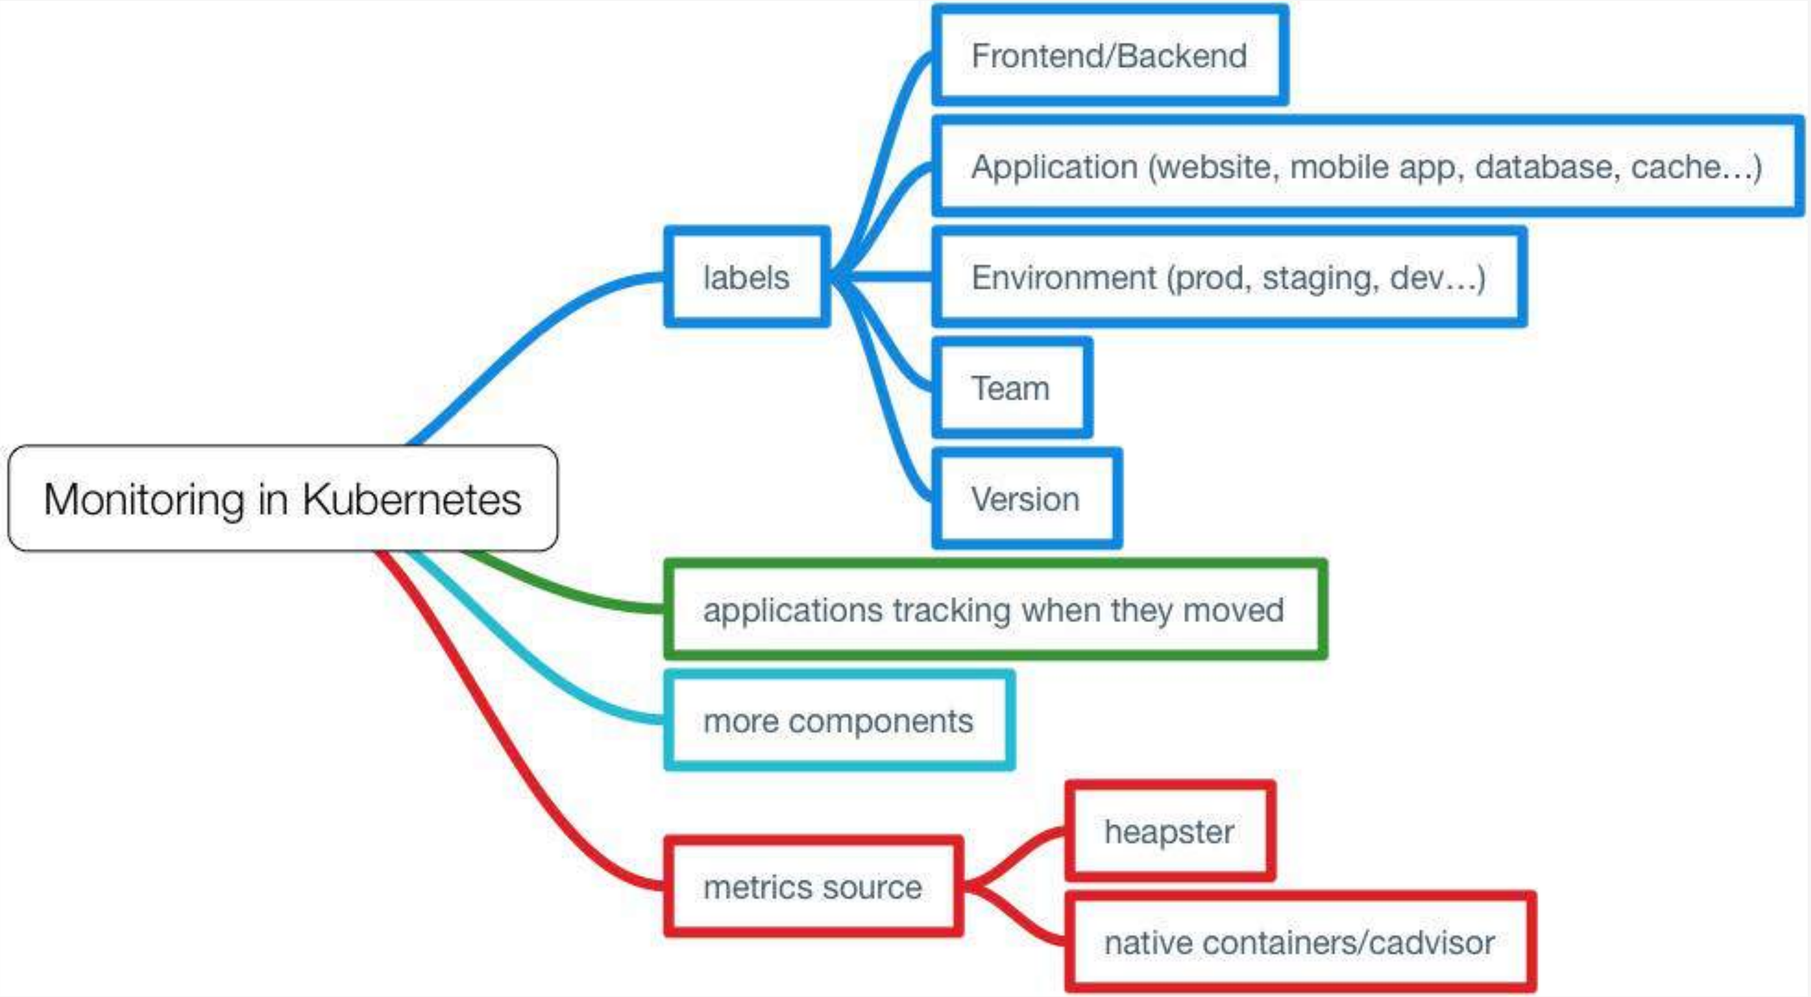
Monitoring in Kubernetes
labels
Frontend/Backend
Application (website, mobile app, database, cache…)
Environment (prod, staging, dev…)
Team
Version
applications tracking when they moved
more components
metrics source
heapster
native containers/cadvisor

Outline

- **Docker in TalkingData**
- **Containers**
- **Why Kubernetes?**
- **Microservices**
- **Cloud Native**
- **Service Mesh**
- **Use Cases**
- **Open Source**

Service Mesh

Microservices Government on Kubernetes

Istio

Linkerd

- Traffic Management：API网关
- Observability：服务调用和性能分析
- Policy Enforcment：控制服务访问策略
- Service Identity and Security：安全保护

- 专用的基础设施层
- 轻量级高性能网络代理
- 提供安全的、快速的、可靠地服务间通讯
- 扩展kubernetes的应用负载均衡机制，实现灰度发布
- 完全解耦于应用，应用可以无感知，加速应用的微服务和云原生转型

CONDUIT

https://github.com/runconduit/conduit

Service Mesh

Istio VS Linkerd

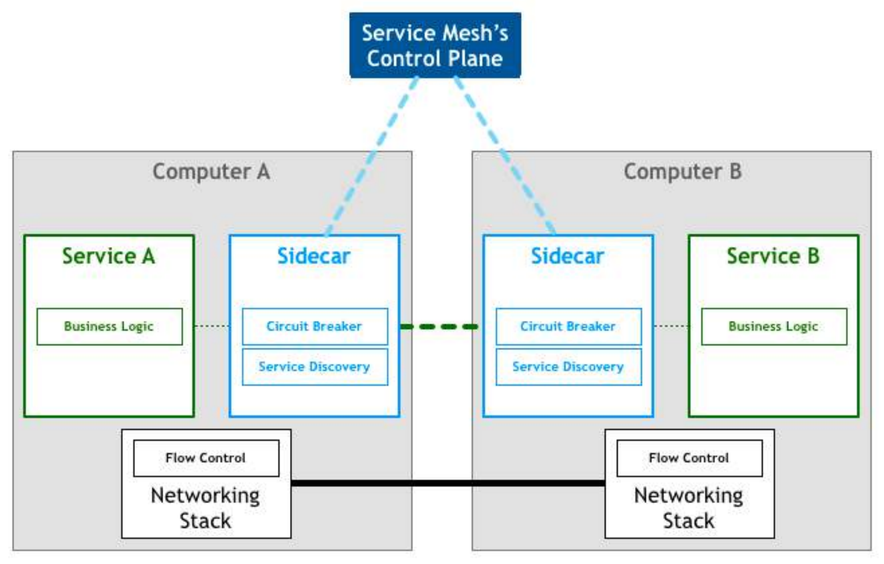

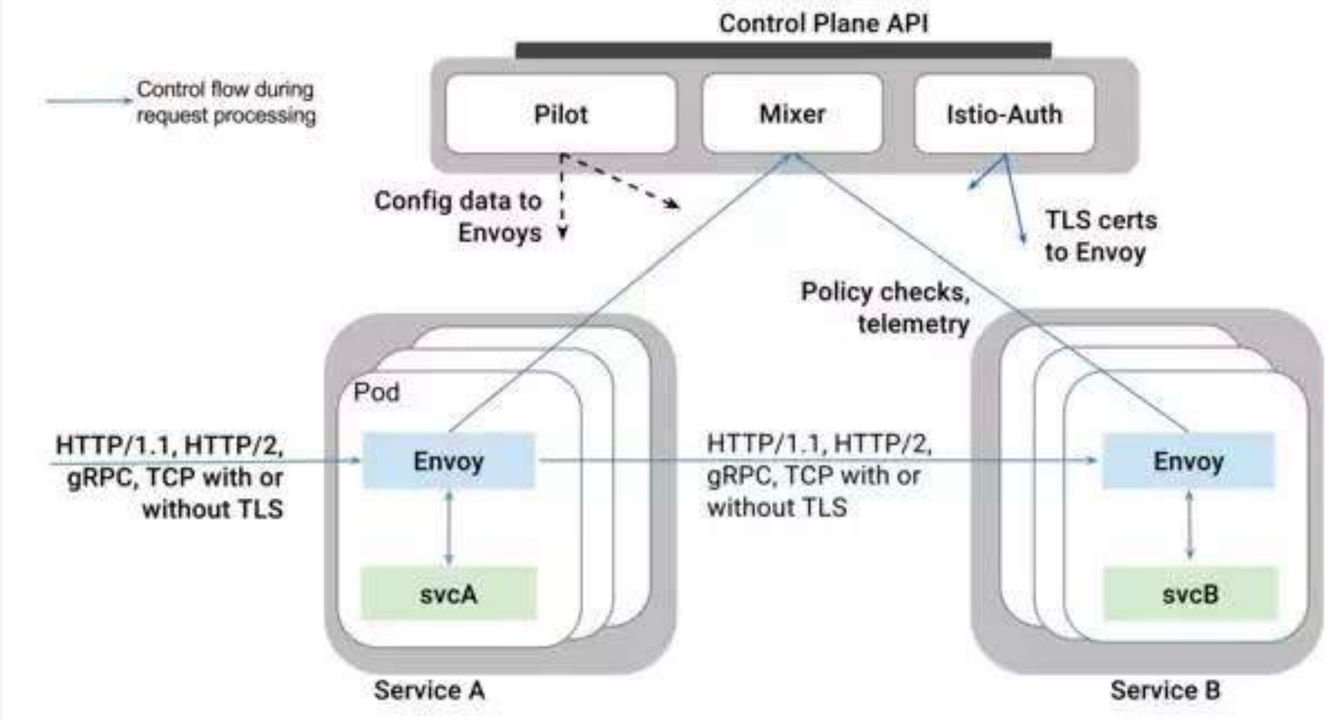

Service Mesh中文网：http://www.servicemesh.cn/

Istio vs Linkerd

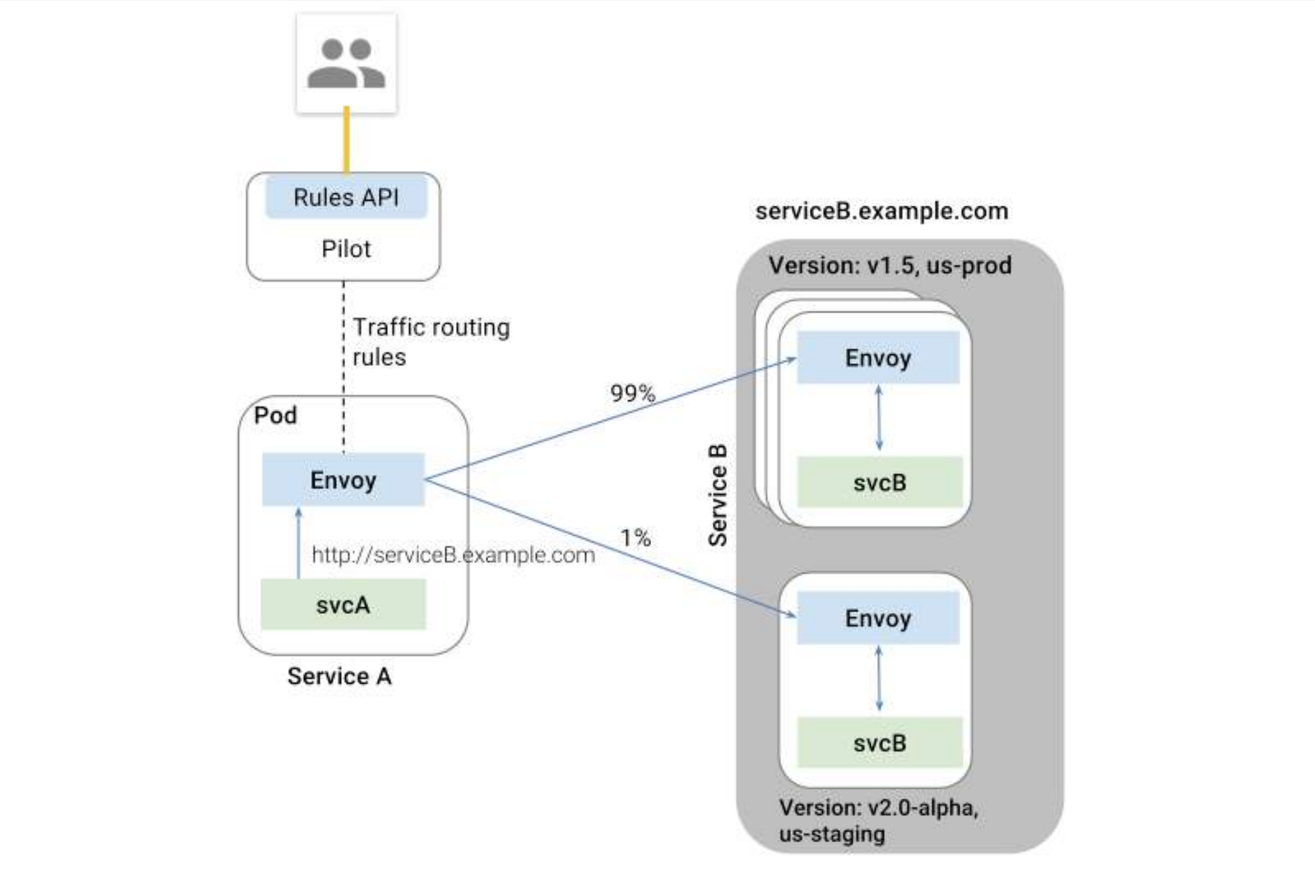

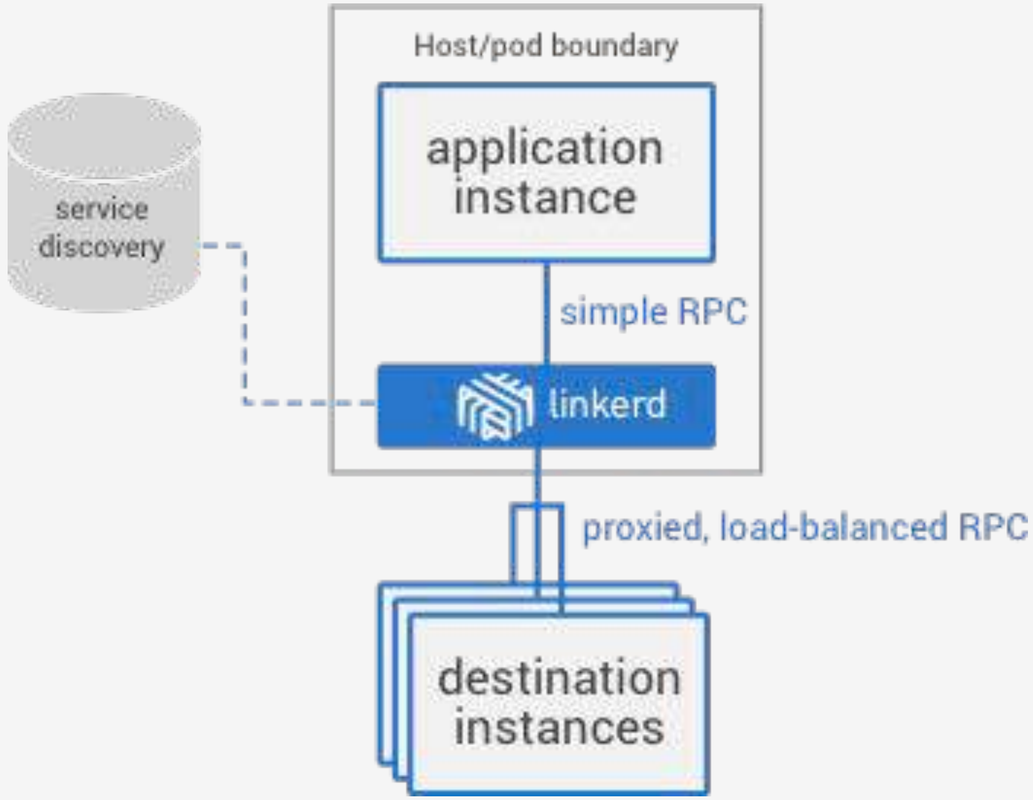

Istio vs Linkerd

Feature	Istio	Linkerd
Deployment	Envoy/ Sidecar	DaemonSets
Easy to use	complicated	easy
Platform	kuberentes	kubernetes/ mesos/Istio/local
Version	0.3.0	1.3.3
Production	No	Yes

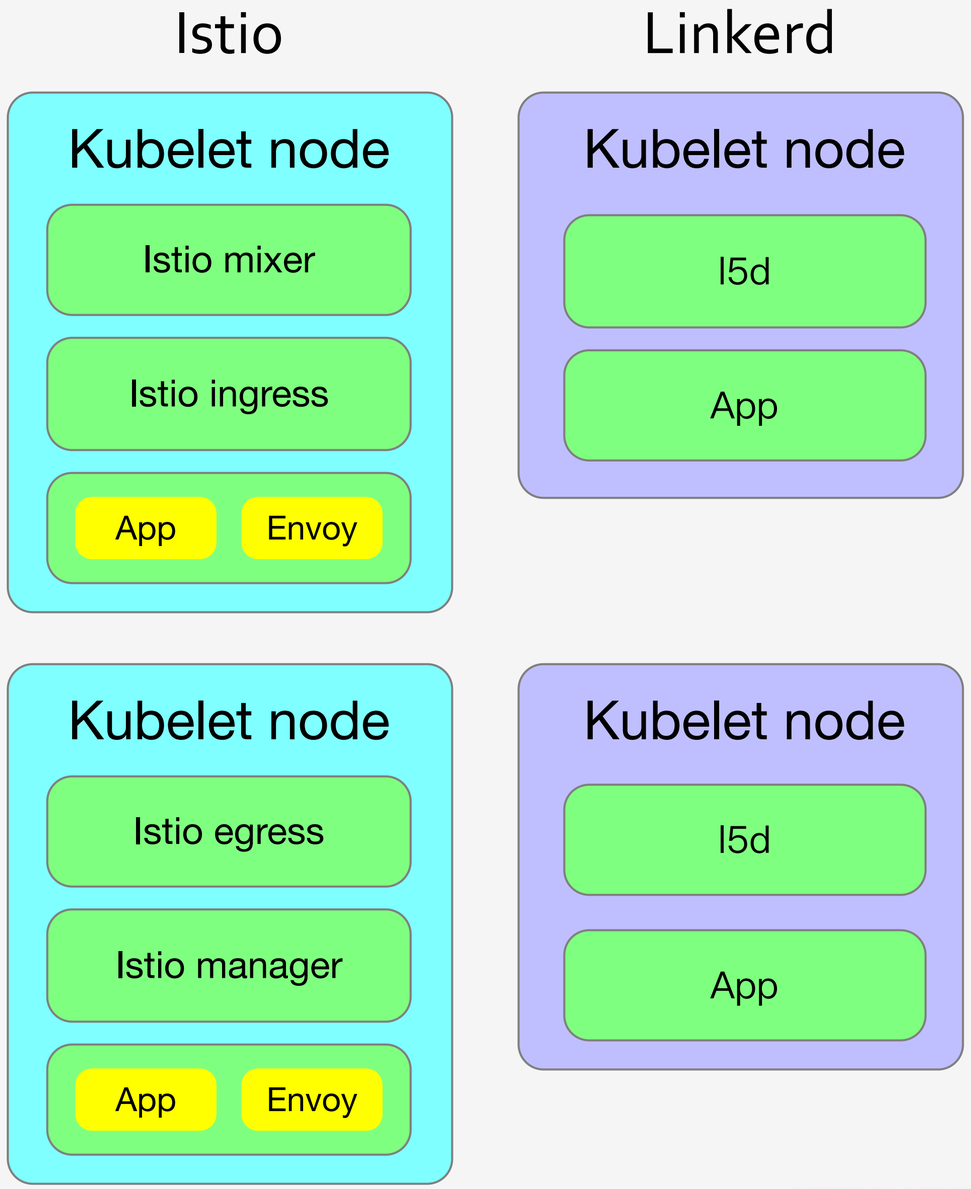

Grafana

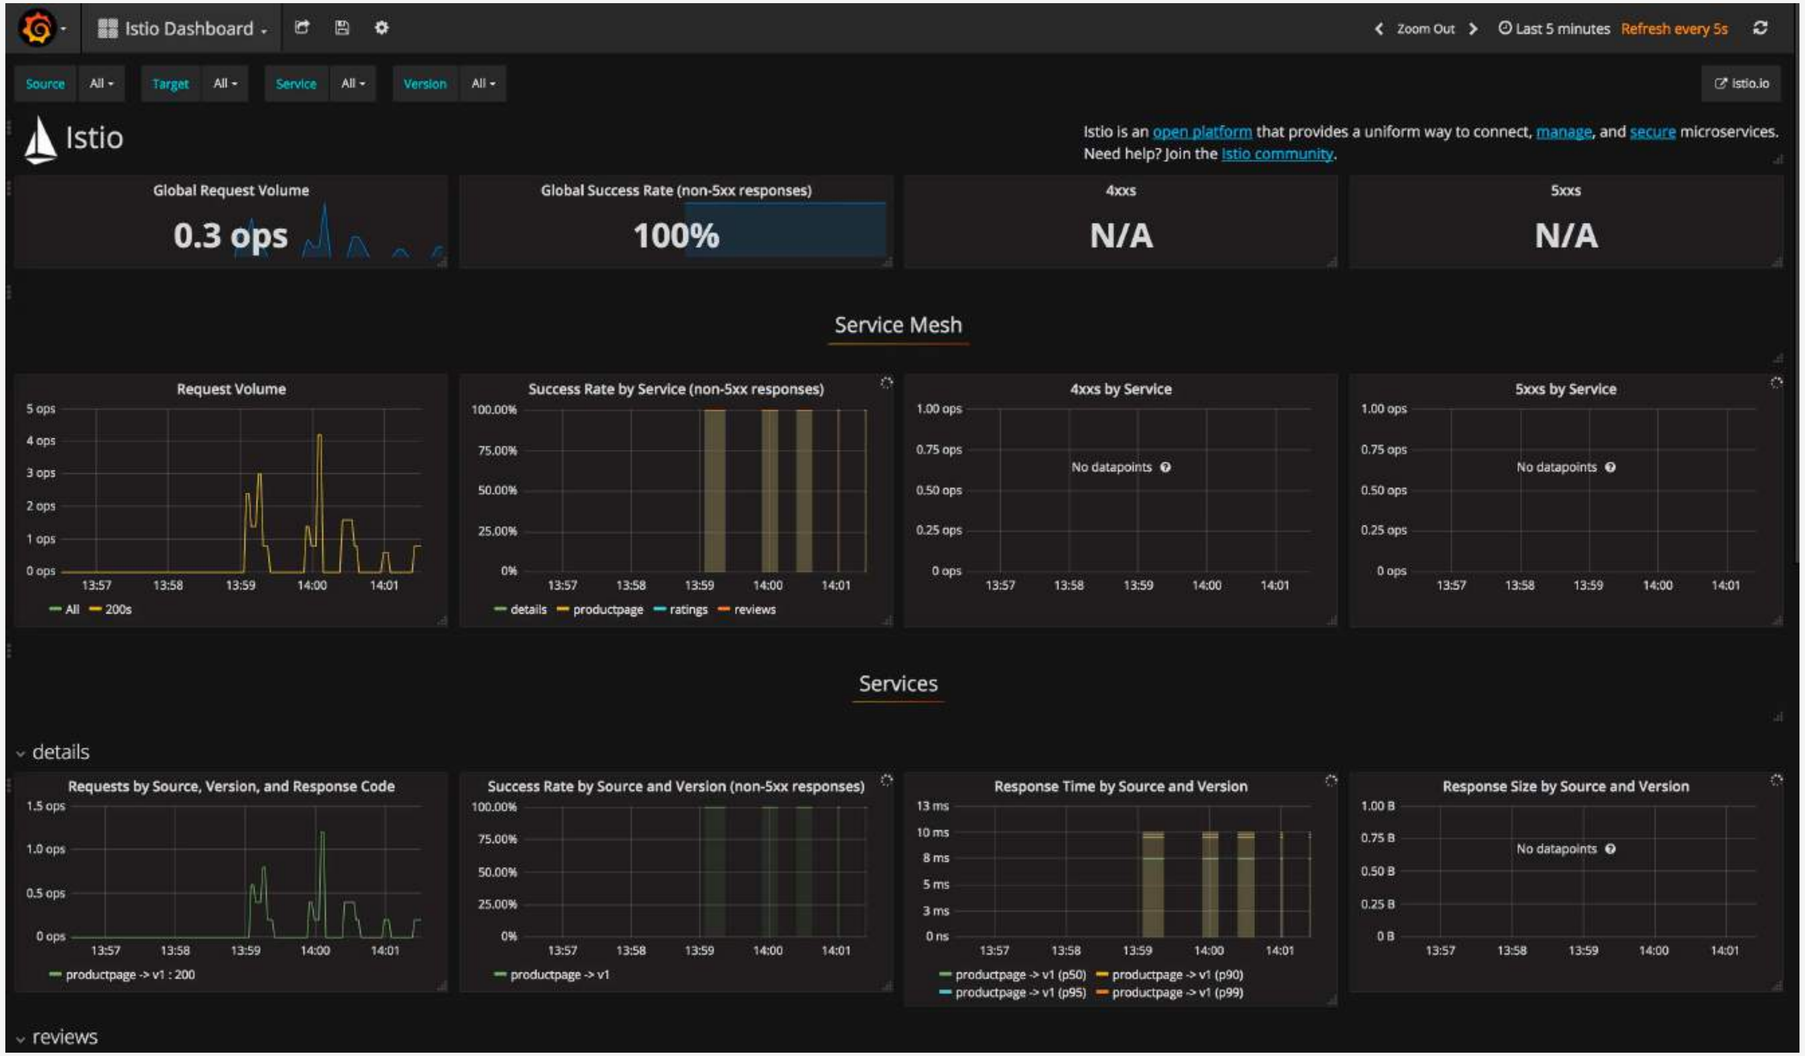
Istio Dashboard
Zoom Out
Last 5 minutes
Refresh every 5s
Source
All
Target
All
Service
All
Version
All
Istio.io
Istio
Istio is an open platform that provides a uniform way to connect, manage, and secure microservices.
Need help? Join the Istio community.
Global Request Volume
0.3 ops
Global Success Rate (non-5xx responses)
100%
4xxs
N/A
5xxs
N/A
Service Mesh
Request Volume
5 ops
4 ops
3 ops
2 ops
1 ops
0 ops
13:57
13:58
13:59
14:00
14:01
All
200s
Success Rate by Service (non-5xx responses)
100.00%
75.00%
50.00%
25.00%
0%
details
productpage
ratings
reviews
4xxs by Service
1.00 ops
0.75 ops
0.50 ops
0.25 ops
0 ops
No datapoints
5xxs by Service
Services
details
Requests by Source, Version, and Response Code
1.5 ops
1.0 ops
0.5 ops
productpage -> v1 : 200
Success Rate by Source and Version (non-5xx responses)
productpage -> v1
Response Time by Source and Version
13 ms
10 ms
8 ms
5 ms
3 ms
0 ns
productpage -> v1 (p50)
productpage -> v1 (p90)
productpage -> v1 (p95)
productpage -> v1 (p99)
Response Size by Source and Version
1.00 B
0.75 B
0.50 B
0.25 B
0 B
reviews

Zipkin

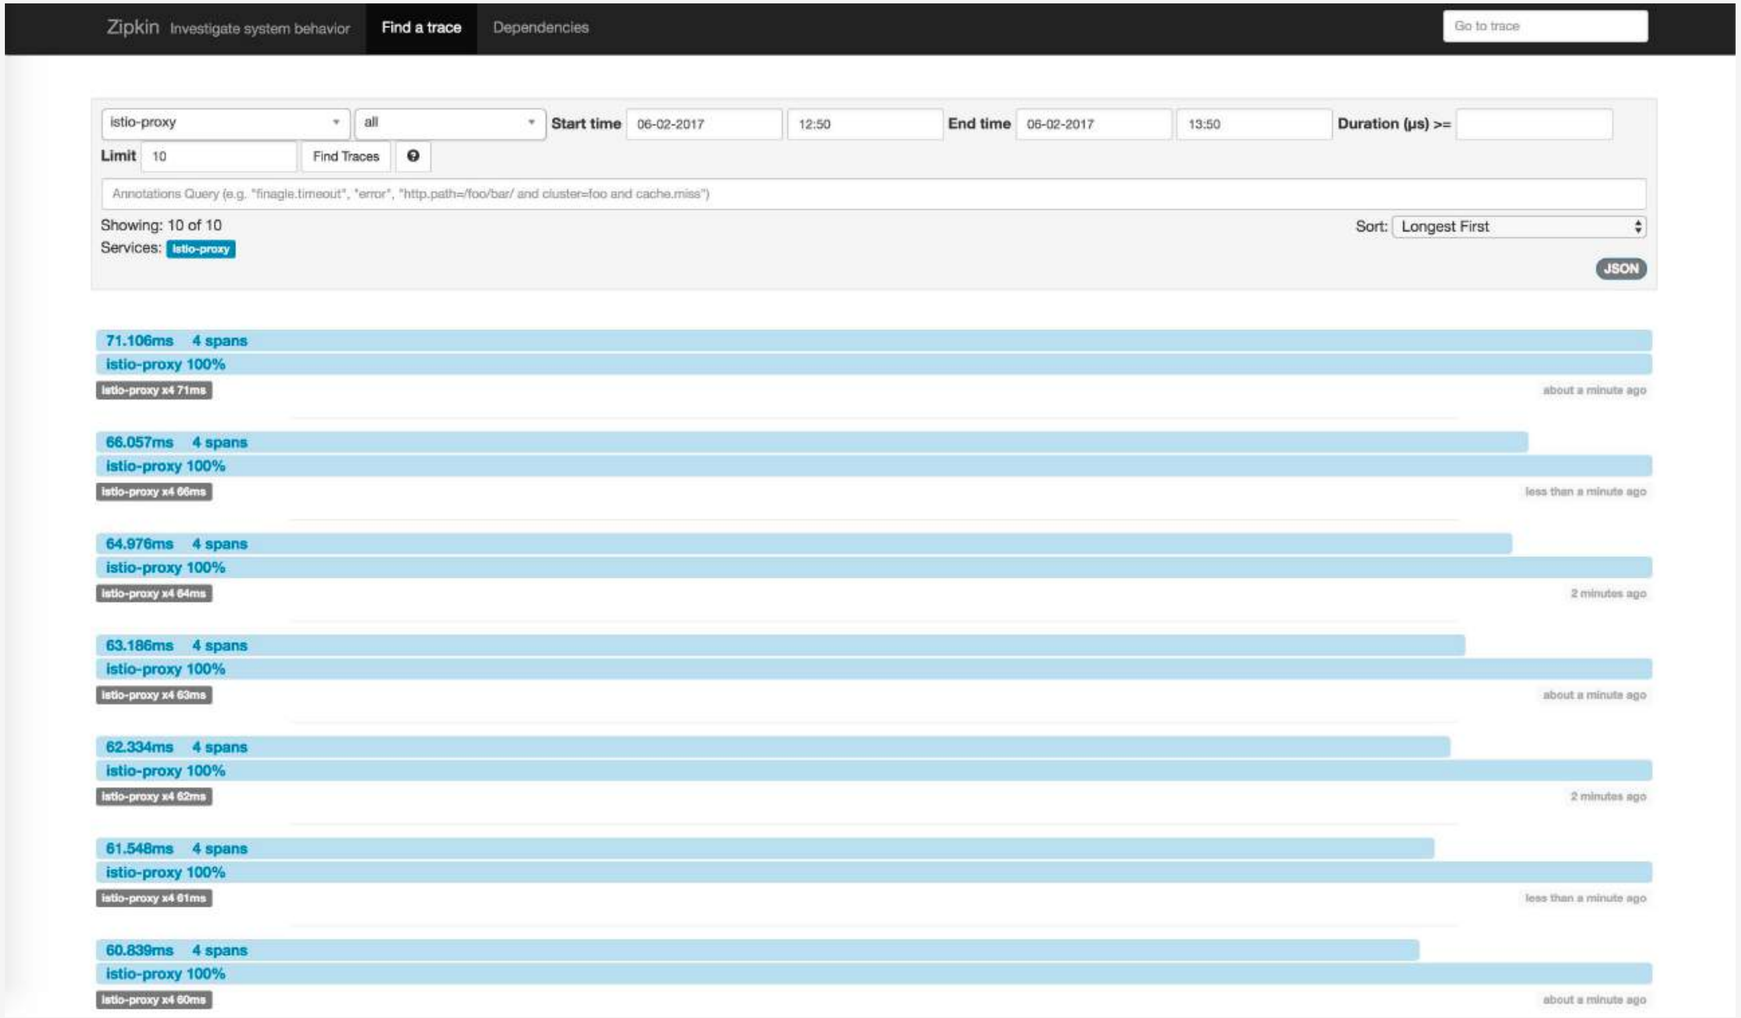
Zipkin Investigate system behavior
Find a trace
Dependencies
Go to trace
istio-proxy
all
Start time 06-02-2017 12:50
End time 06-02-2017 13:50
Duration (µs) >=
Limit 10
Find Traces
Annotations Query (e.g. "finagle.timeout", "error", "http.path=/foo/bar/ and cluster=foo and cache.miss")
Showing: 10 of 10
Sort: Longest First
Services: istio-proxy
JSON
71.106ms 4 spans
istio-proxy 100%
istio-proxy x4 71ms
about a minute ago
66.057ms 4 spans
istio-proxy 100%
istio-proxy x4 66ms
less than a minute ago
64.976ms 4 spans
istio-proxy 100%
istio-proxy x4 64ms
2 minutes ago
63.186ms 4 spans
istio-proxy 100%
istio-proxy x4 63ms
about a minute ago
62.334ms 4 spans
istio-proxy 100%
istio-proxy x4 62ms
2 minutes ago
61.548ms 4 spans
istio-proxy 100%
istio-proxy x4 61ms
less than a minute ago
60.839ms 4 spans
istio-proxy 100%
istio-proxy x4 60ms
about a minute ago

Weave scope

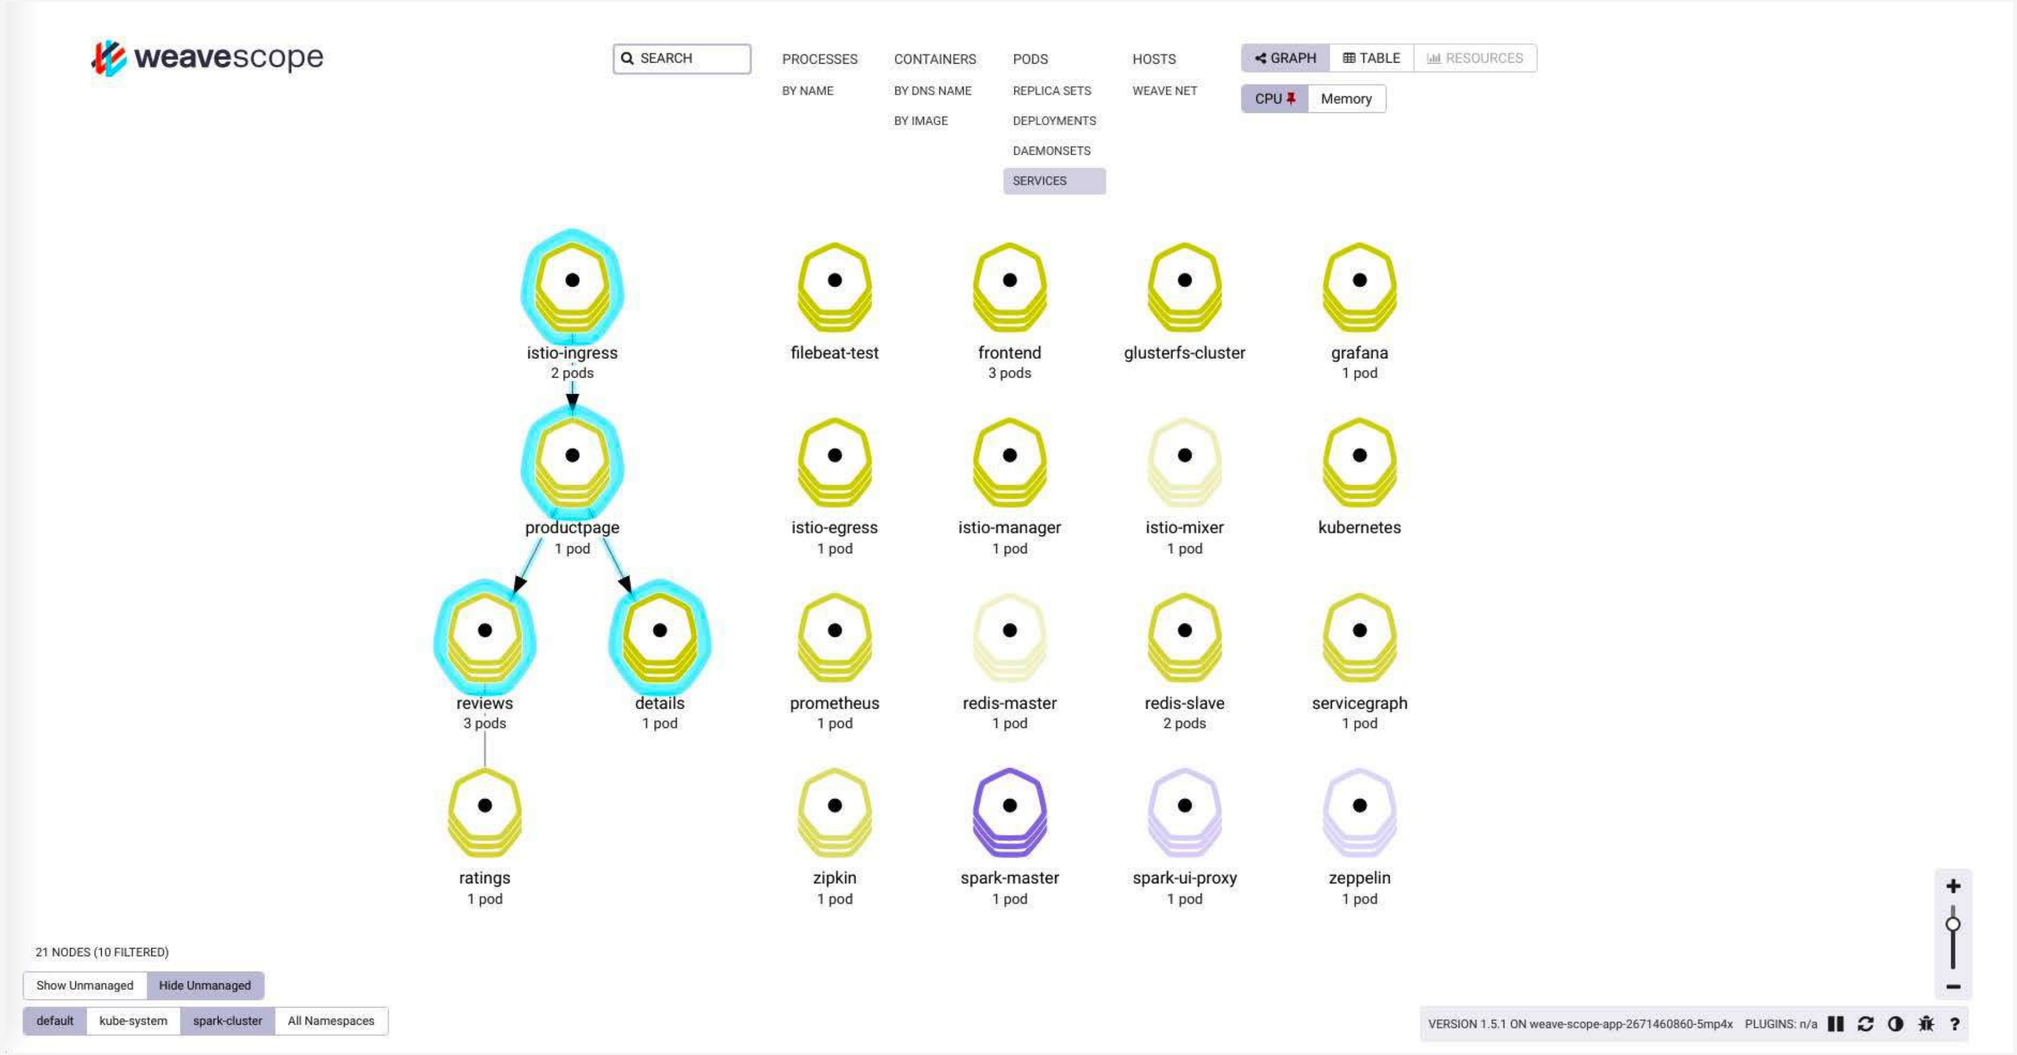
weavescope
SEARCH
PROCESSES
BY NAME
CONTAINERS
BY DNS NAME
BY IMAGE
PODS
REPLICA SETS
DEPLOYMENTS
DAEMONSETS
SERVICES
HOSTS
WEAVE NET
GRAPH
TABLE
RESOURCES
CPU
Memory
istio-ingress
2 pods
filebeat-test
frontend
3 pods
glusterfs-cluster
grafana
1 pod
productpage
1 pod
istio-egress
1 pod
istio-manager
1 pod
istio-mixer
1 pod
kubernetes
reviews
3 pods
details
1 pod
prometheus
1 pod
redis-master
1 pod
redis-slave
2 pods
servicegraph
1 pod
ratings
1 pod
zipkin
1 pod
spark-master
1 pod
spark-ui-proxy
1 pod
zeppelin
1 pod
21 NODES (10 FILTERED)
Show Unmanaged
Hide Unmanaged
default
kube-system
spark-cluster
All Namespaces
VERSION 1.5.1 ON weave-scope-app-2671460860-5mp4x
PLUGINS: n/a

Use Cases

- **Docker in TalkingData**
- **Containers**
- **Why Kubernetes?**
- **Microservices**
- **Cloud Native**
- **Service Mesh**
- **Use Cases**
- **Open Source**

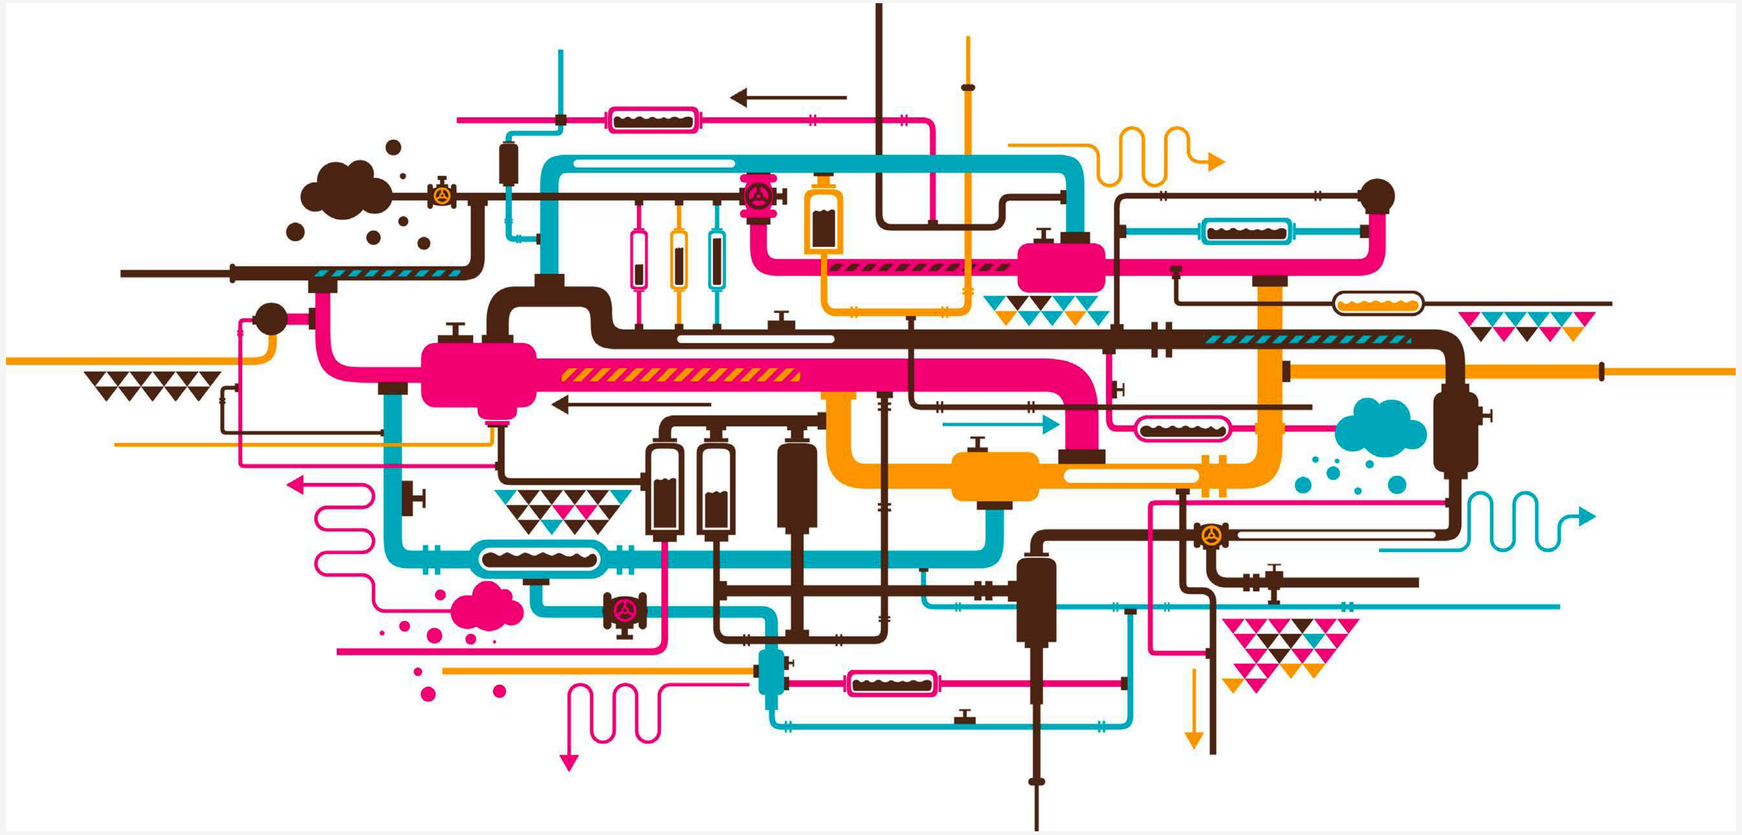

GitOps

- Infrastracture as code
- YAML! YAML! YAML!
- Git version control

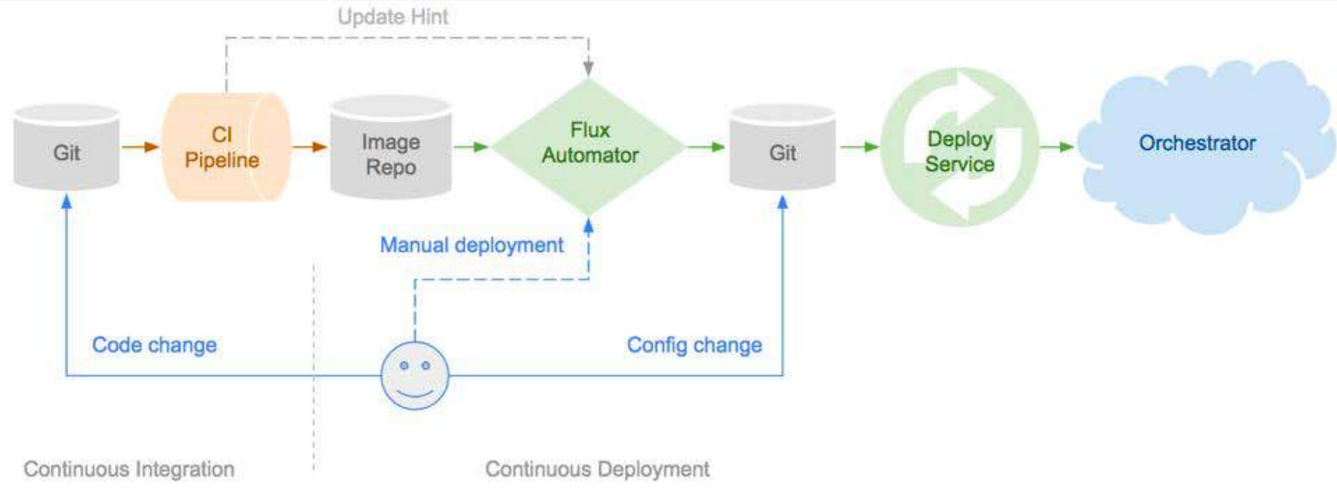

source https://www.weave.works

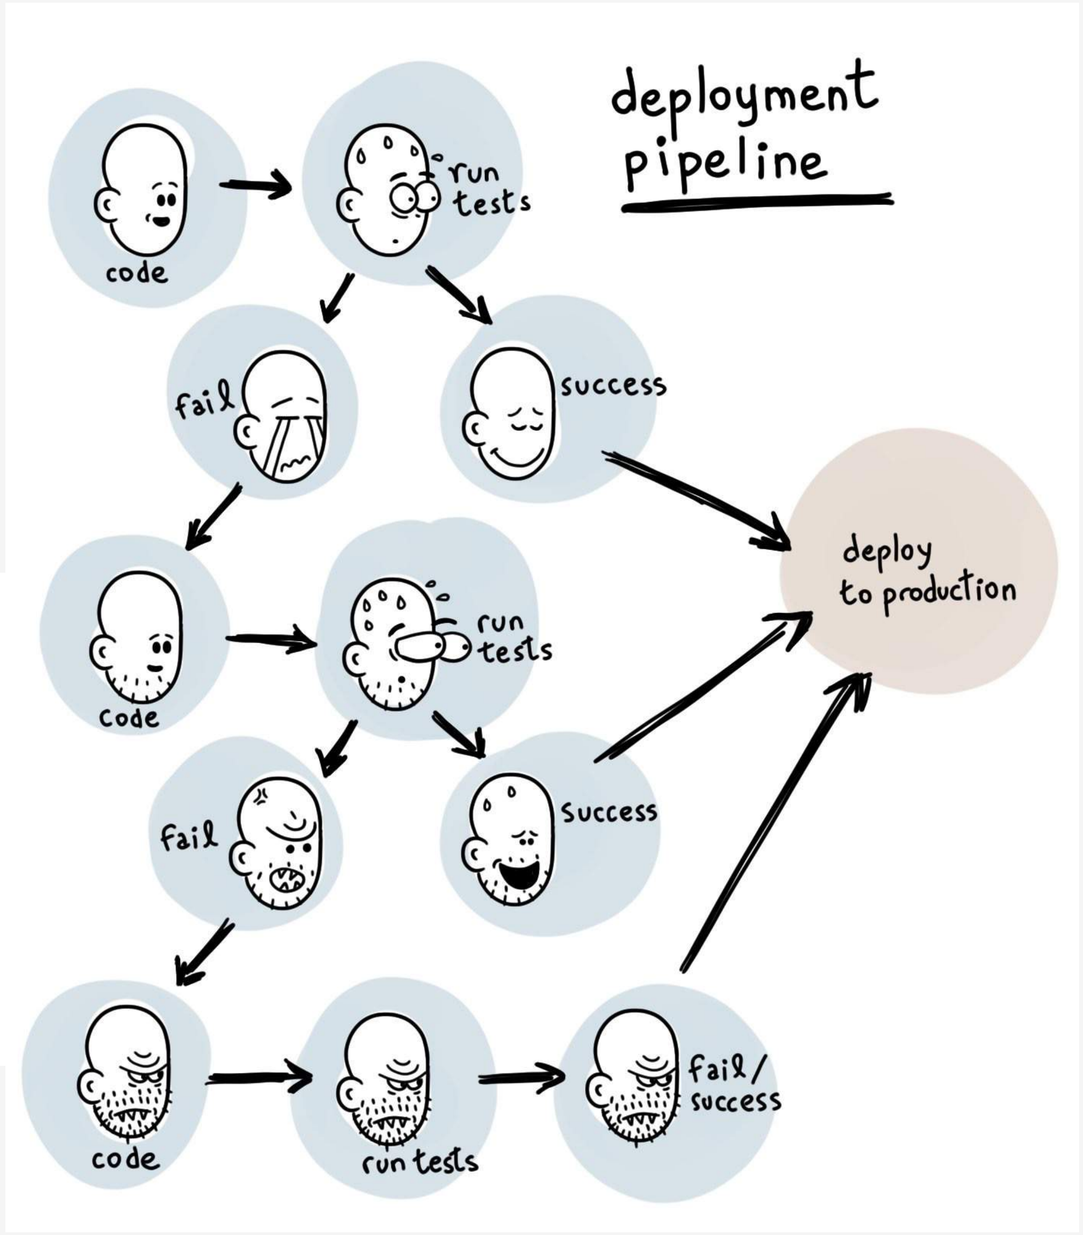

GitOps

- 上线/更新/下线
- YAML! YAML! YAML!
- Helm chart

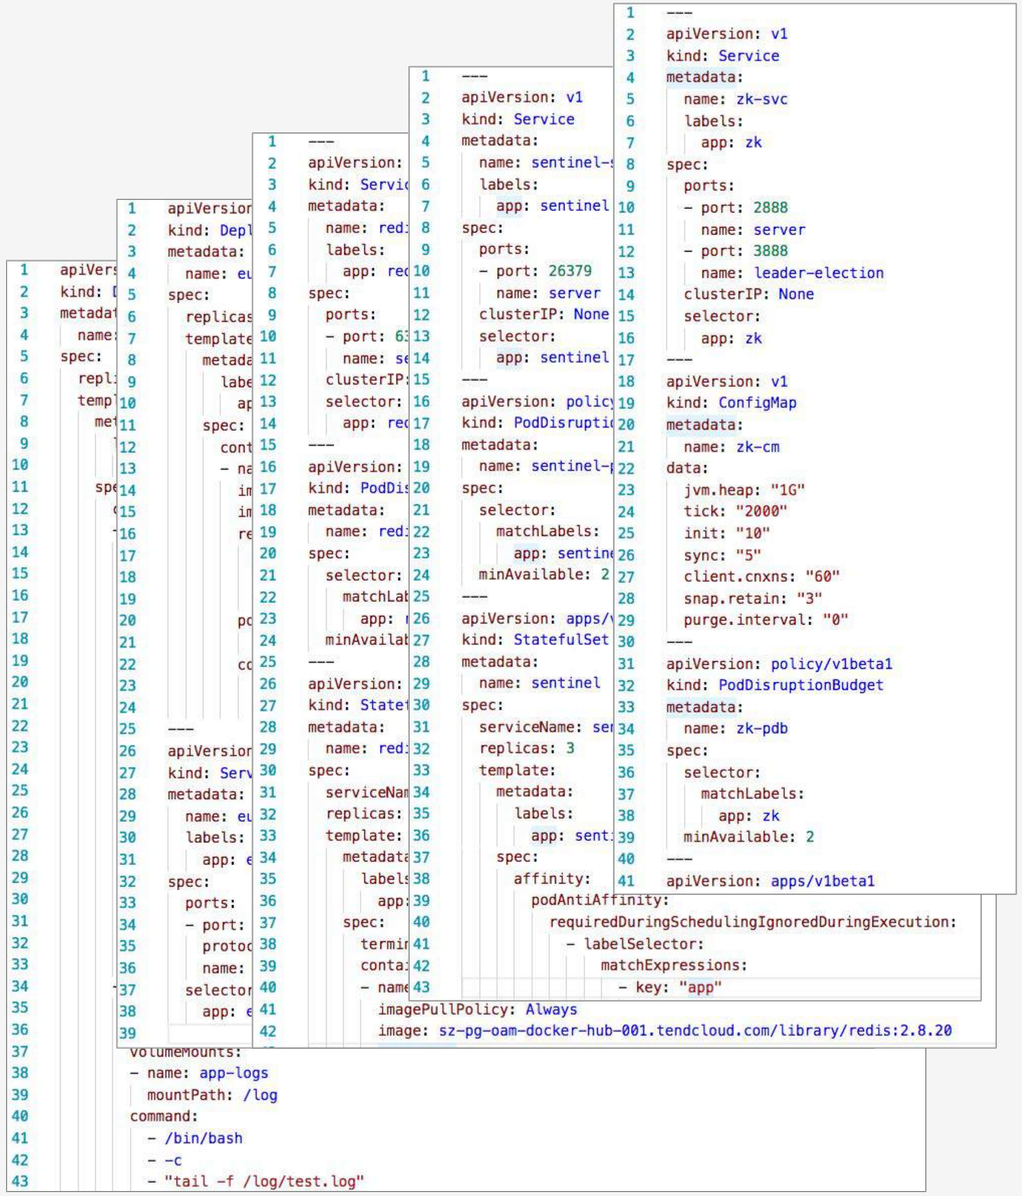

Spark on Kubernetes

Spark on Kubernetes with different schedulers

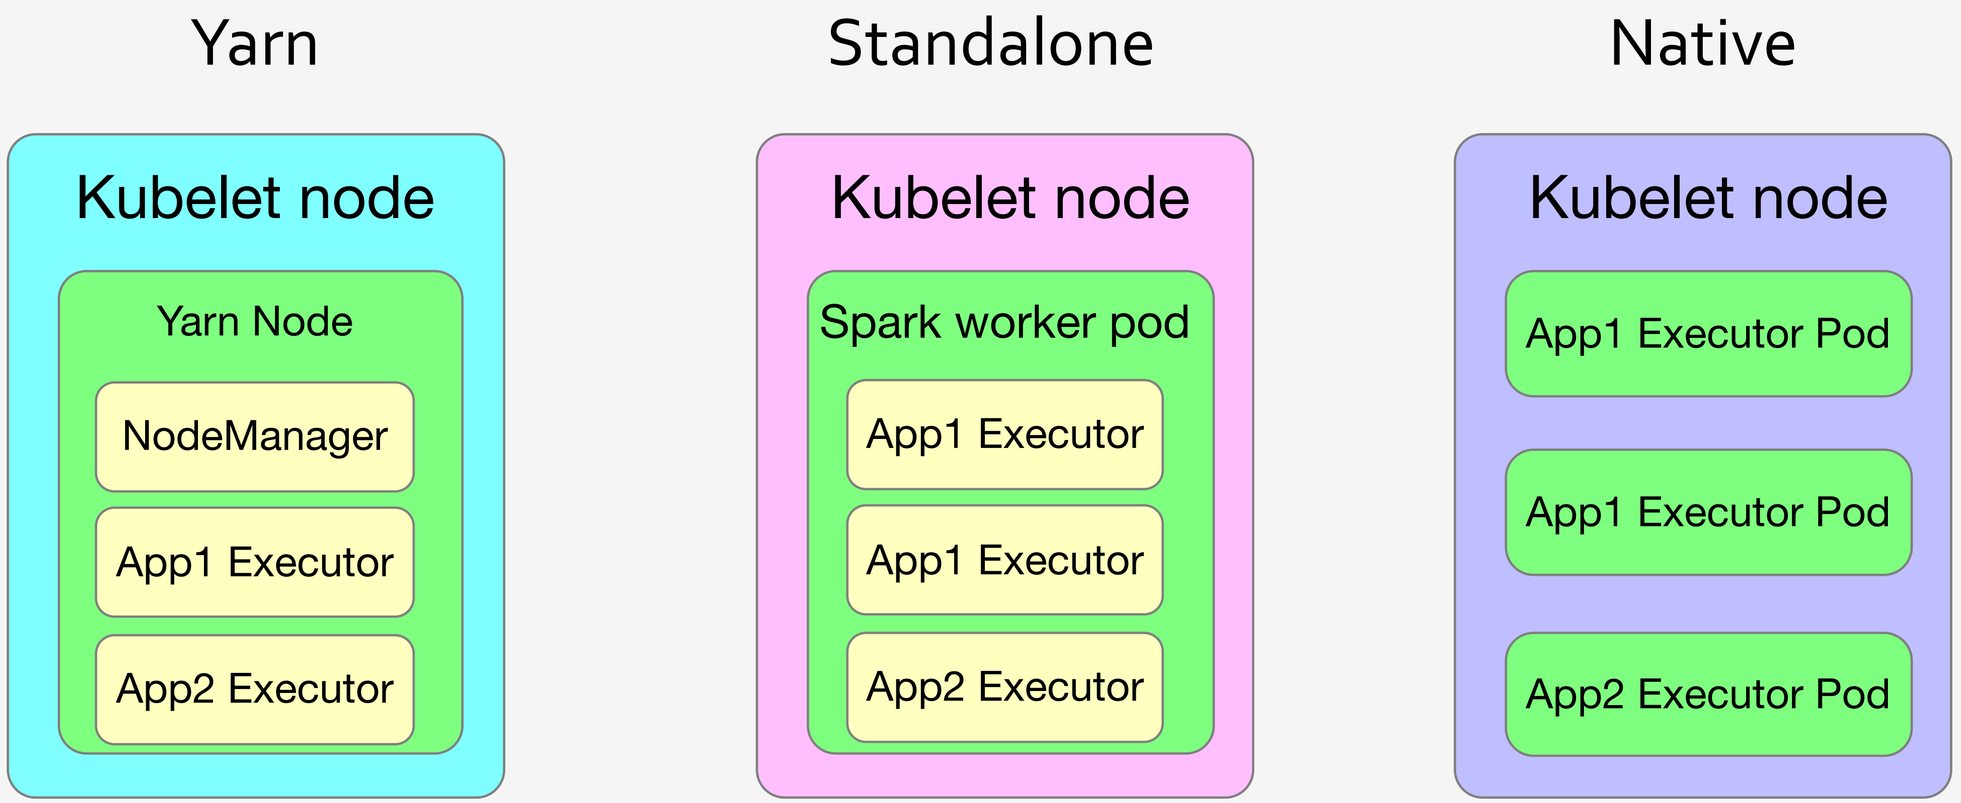

https://jimmysong.io

Spark on Kubernetes with yarn scheduler

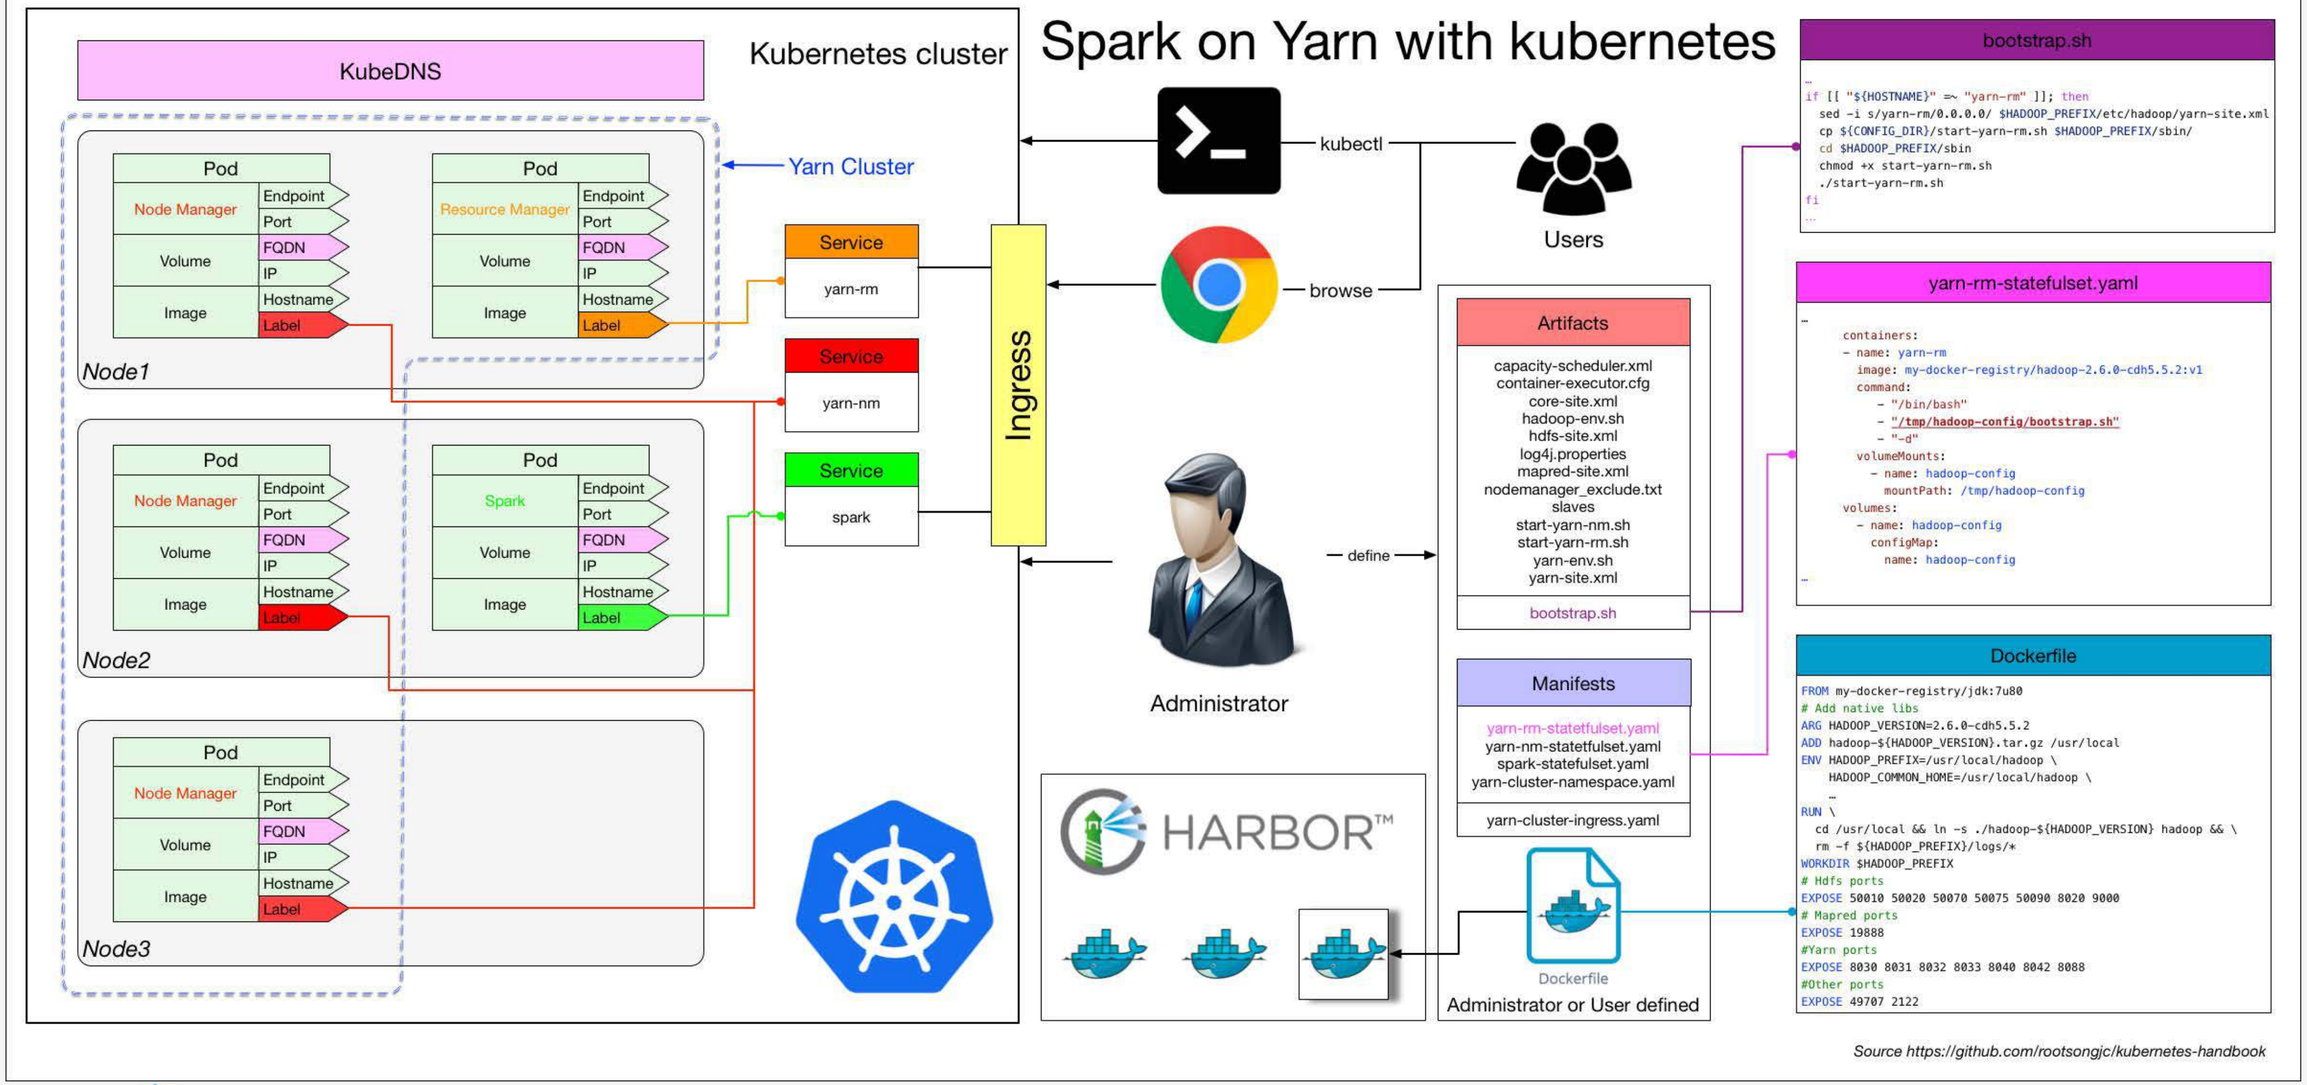

Source https://github.com/rootsongjc/kubernetes-handbook

Spark on Kuberentes with native scheduler

- Kubernetes原生调度：与yarn、mesos同级
- 资源隔离，粒度更细：以namespace来划分用户
- 监控的变革：单次任务资源计量
- 日志的变革：pod的日志收集

```
./spark-submit \
  --deploy-mode cluster \
  --class com.talkingdata.alluxio.hadooptest \
  --master k8s://https://172.20.0.113:6443 \
  --kubernetes-namespace spark-cluster \
  --conf spark.kubernetes.driverEnv.SPARK_USER=hadoop \
  --conf spark.kubernetes.driverEnv.HADOOP_USER_NAME=hadoop \
  --conf spark.executorEnv.HADOOP_USER_NAME=hadoop \
  --conf spark.executorEnv.SPARK_USER=hadoop \
  --conf spark.kubernetes.authenticate.driver.serviceAccountName=spark \
  --conf spark.driver.memory=100G \
  --conf spark.executor.memory=10G \
  --conf spark.driver.cores=30 \
  --conf spark.executor.cores=2 \
  --conf spark.driver.maxResultSize=10240m \
  --conf spark.kubernetes.driver.limit.cores=32 \
  --conf spark.kubernetes.executor.limit.cores=3 \
  --conf spark.kubernetes.executor.memoryOverhead=2g \
  --conf spark.executor.instances=5 \
  --conf spark.app.name=spark-pi \
  --conf spark.kubernetes.driver.docker.image=spark-driver:v2.1.0-kubernetes-0.3.1-1 \
  --conf spark.kubernetes.executor.docker.image=spark-executor:v2.1.0-kubernetes-0.3.1-1
\
  --conf spark.kubernetes.initcontainer.docker.image=spark-init:v2.1.0-kubernetes-0.3.1-1
\
  --conf spark.kubernetes.resourceStagingServer.uri=http://172.20.0.114:31000 \
~/Downloads/tendcloud_2.10-1.0.jar
```

Spark on Yarn V.S Spark on Kuberentes

Feature	Yarn	Kubernetes
queue	queue	namespace
instance	ExcutorContainer	Executor Pod
network	host	plugin
heterogeneous	no	yes
security	RBAC	ACL

https://github.com/apache-spark-on-k8s/spark

Outline

- **Docker in TalkingData**
- **Containers**
- **Why Kubernetes?**
- **Microservices**
- **Cloud Native**
- **Service Mesh**
- **Use Cases**
- **Open Source**