时速云
tenxcloud.com

Kubernetes在企业中的场景运用及管理实践

—— 时速云 王磊

CNUTCon2016 全球容器技术大会
主办方 Geekbang 极客邦科技 InfoQ

Agenda

- Why Enterprise need Private PaaS?
- What Enterprise Cares about
  - Network
  - Storage
  - CI/CD
  - Application Package and Deployment
  - Configuration Management
  - Secret Data
- How to Manage Your PaaS

Why Enterprise Need Private PaaS?

Hybrid Cloud

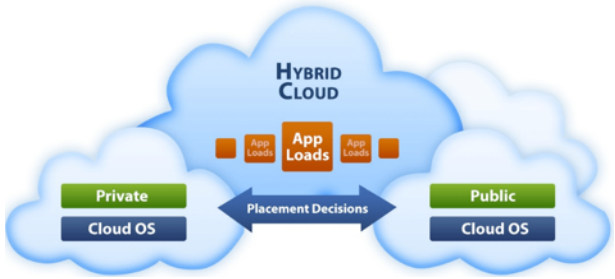

Cloud Enablement of Existing and New Applications

Private Cloud Strategy

Microservices and APIs

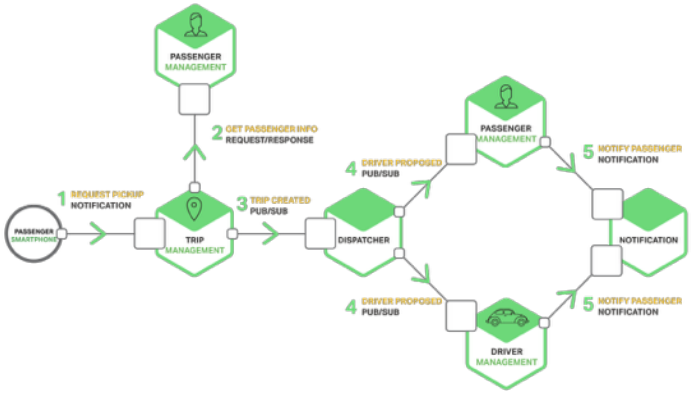

Requirements for Private PaaS?

Ease of deployment and management

Data layer separation

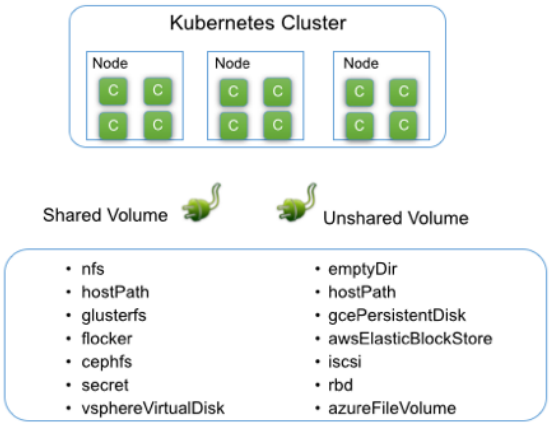

Integration with Legacy Systems

Extensibility

Portability through Standards

What Enterprise Cares about?

- Network
- Storage
- CI/CD
- Application Package and Deployment
- Configuration Management
- Secret Data
- Monitor and Alert

Network - Flannel

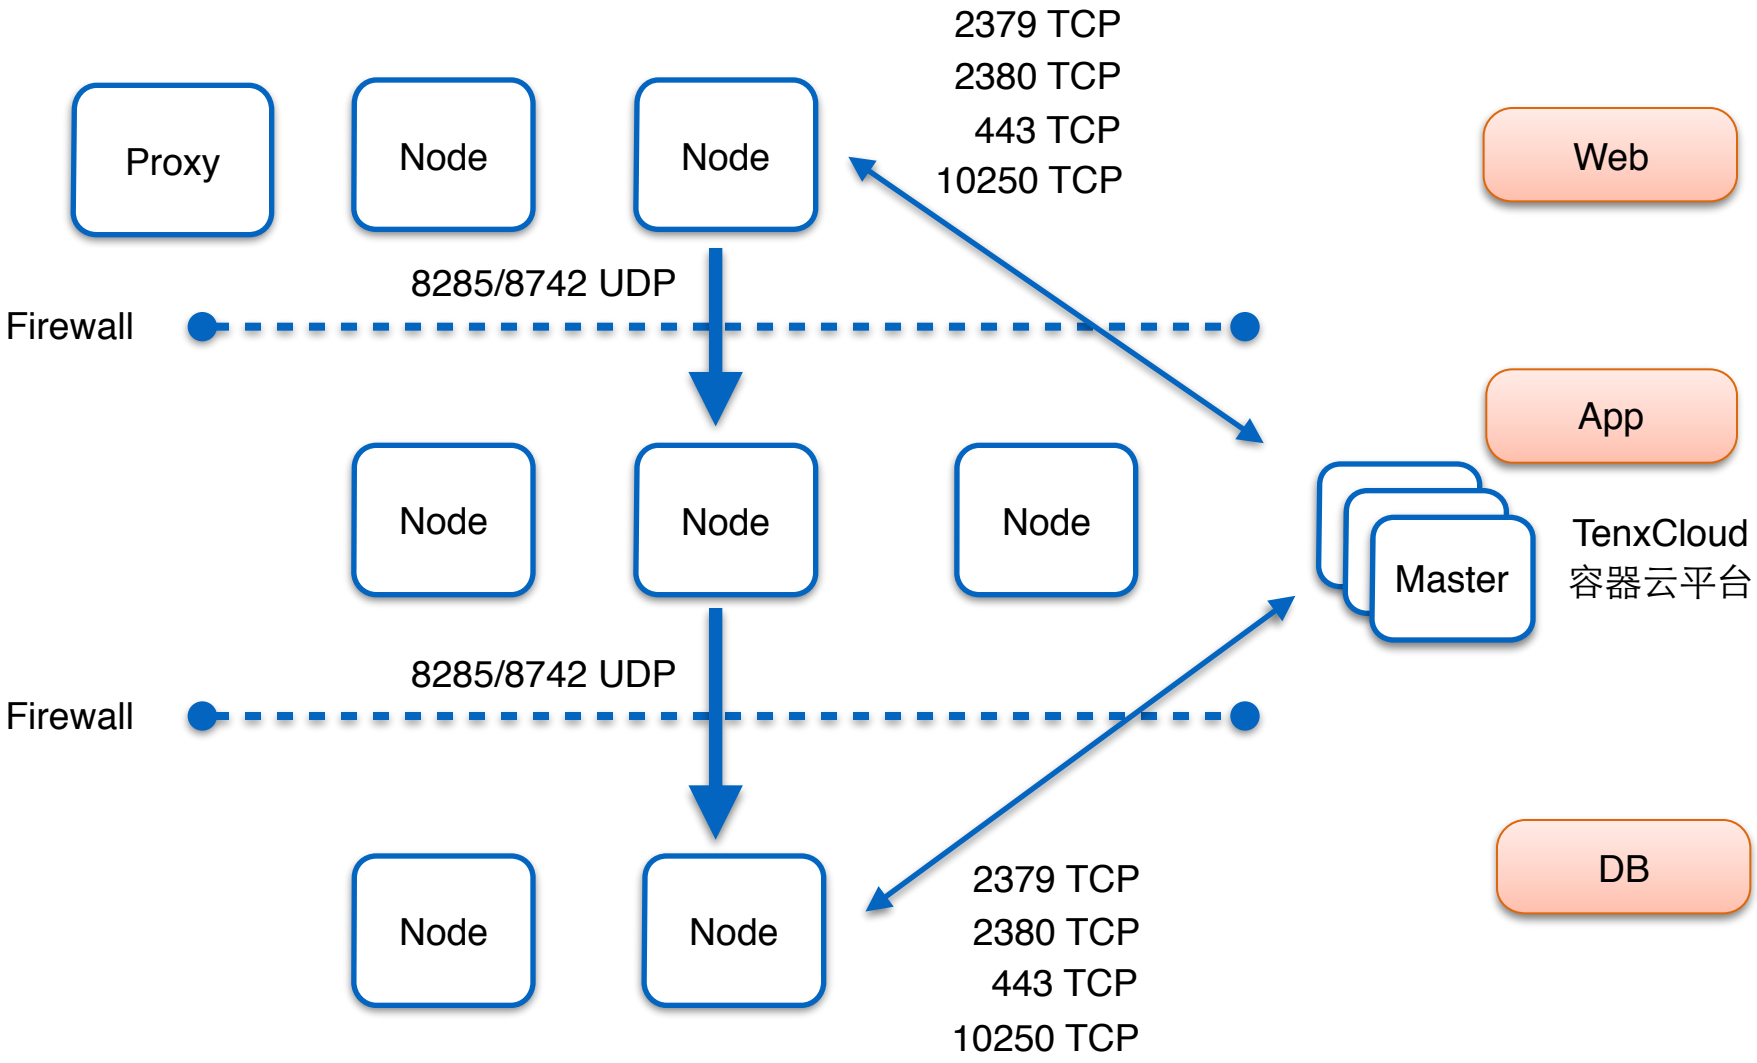

Network - Docker Bridge

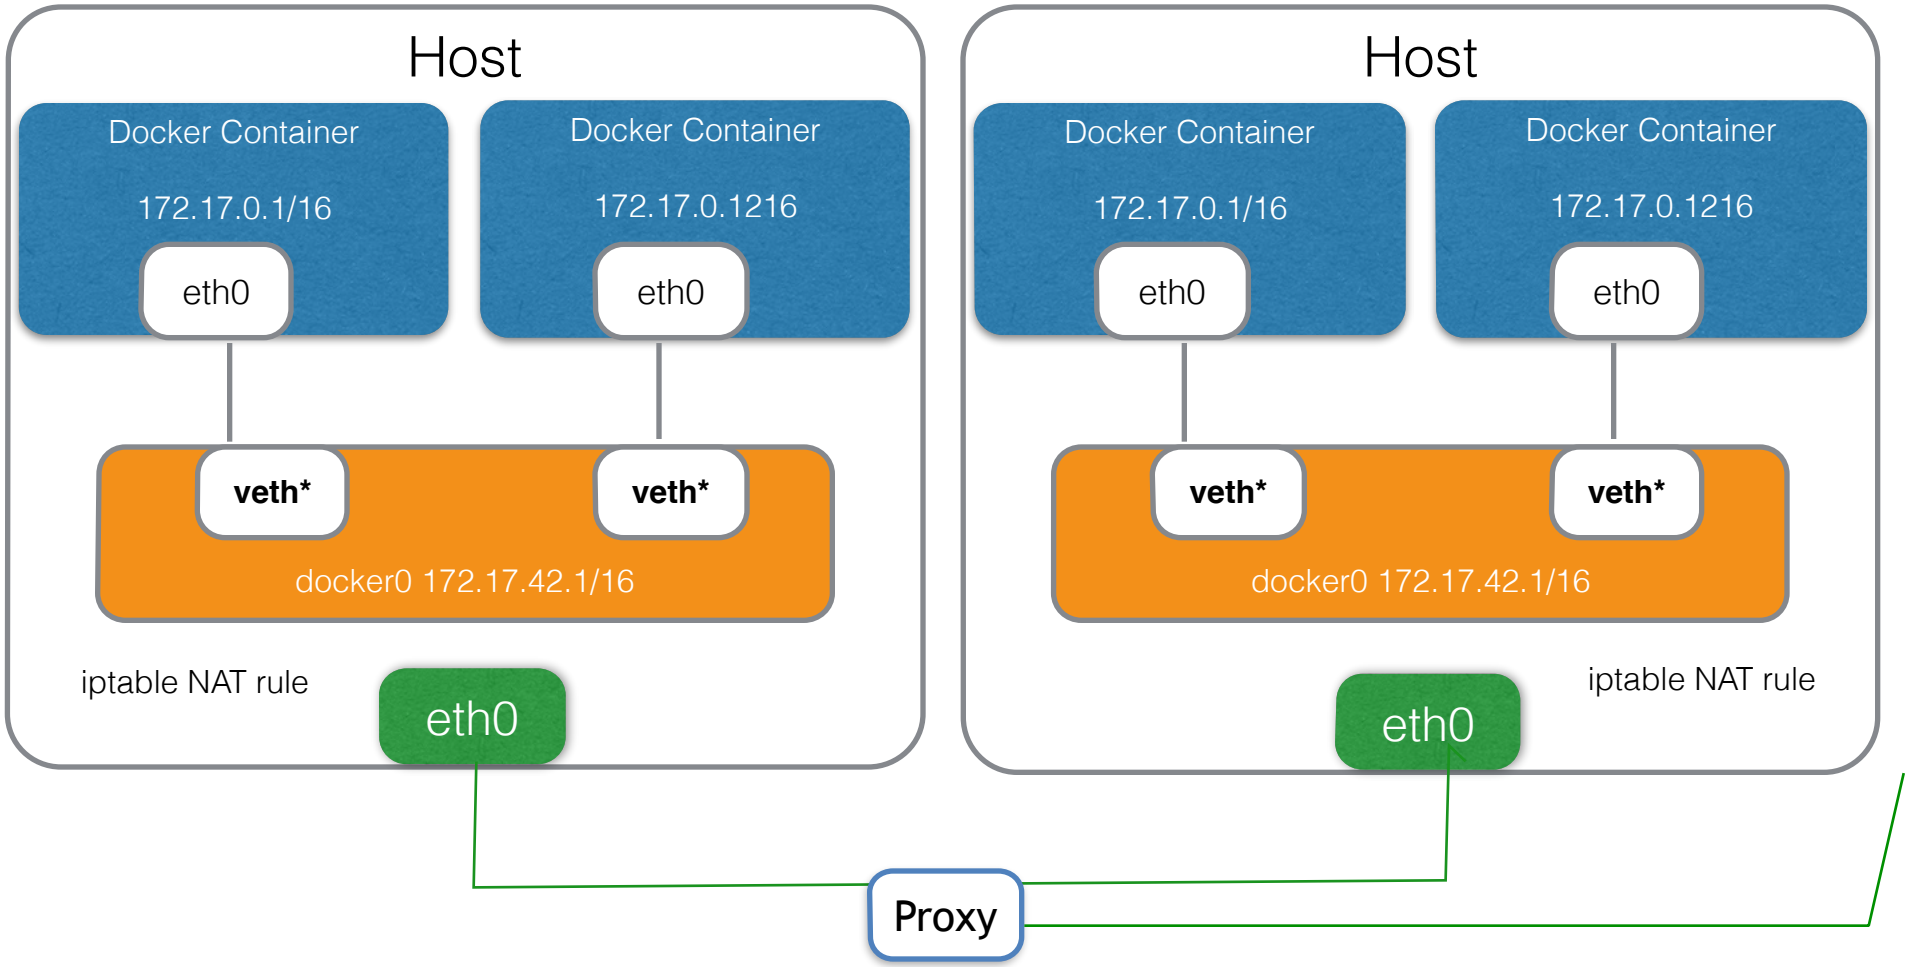

Network - Calico

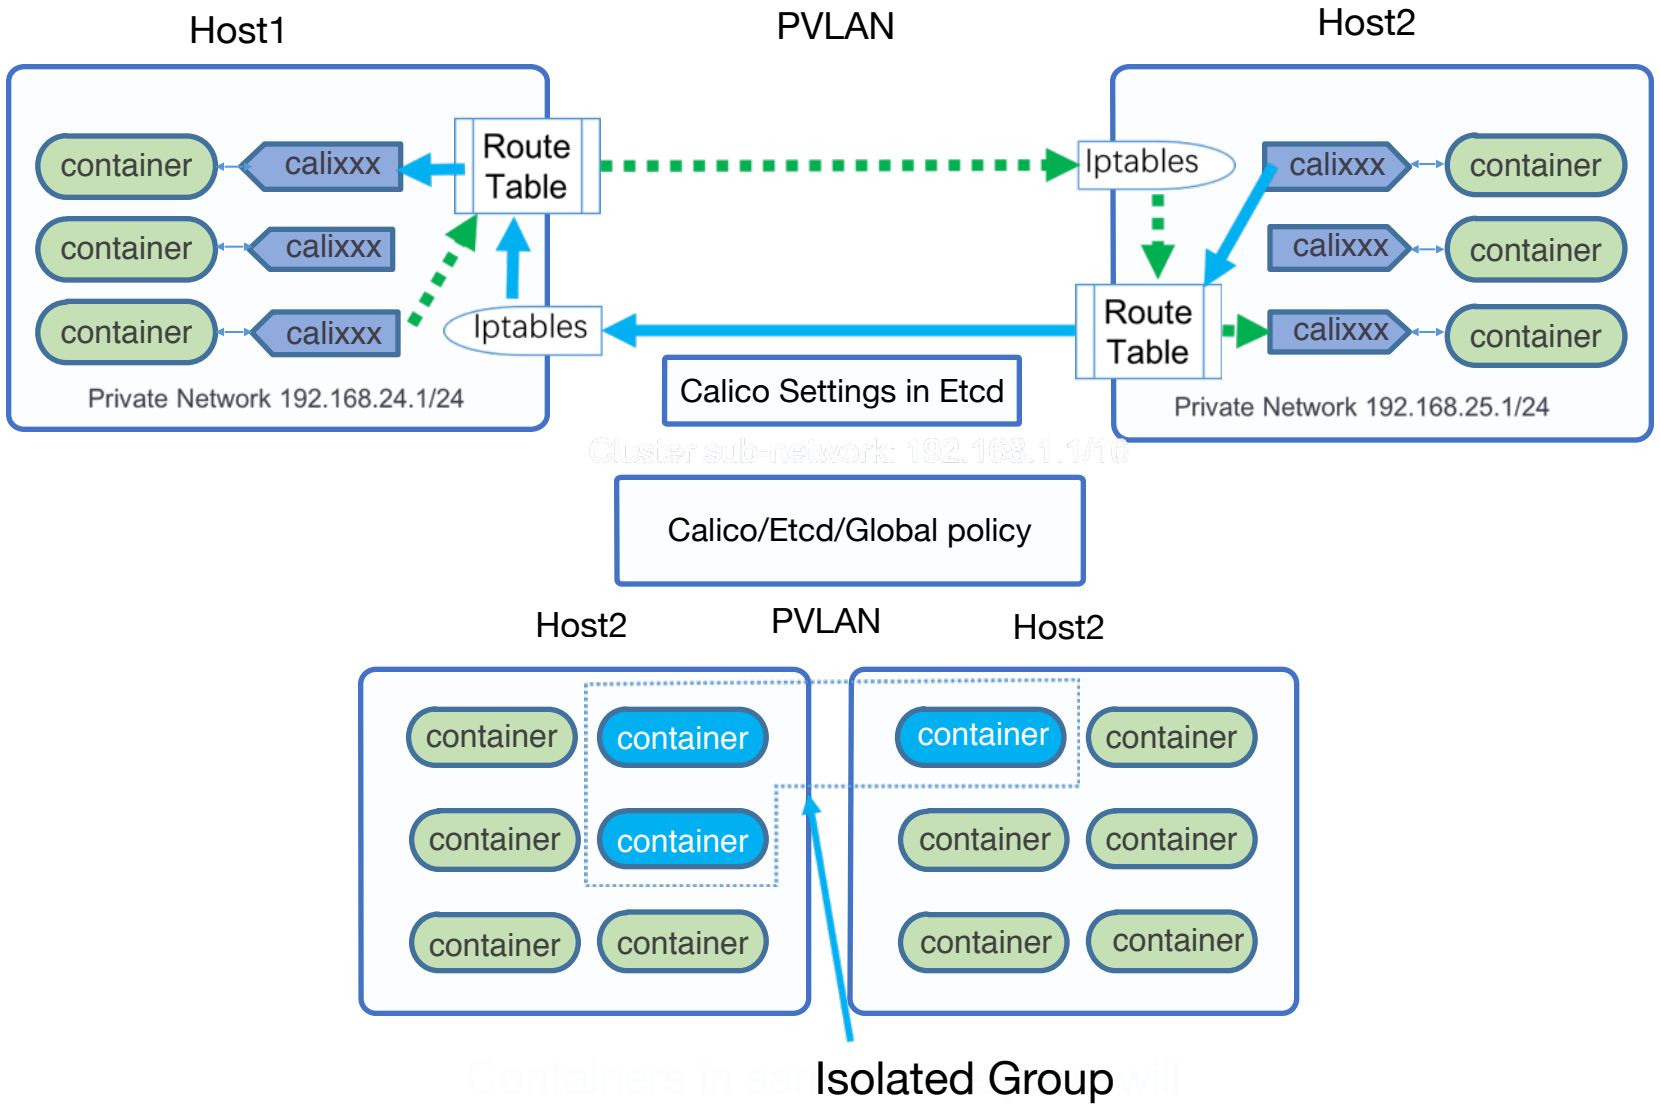

Storage - Extensible Volume Support

- Easy to add customized plugin to use other storage backend following volume interface

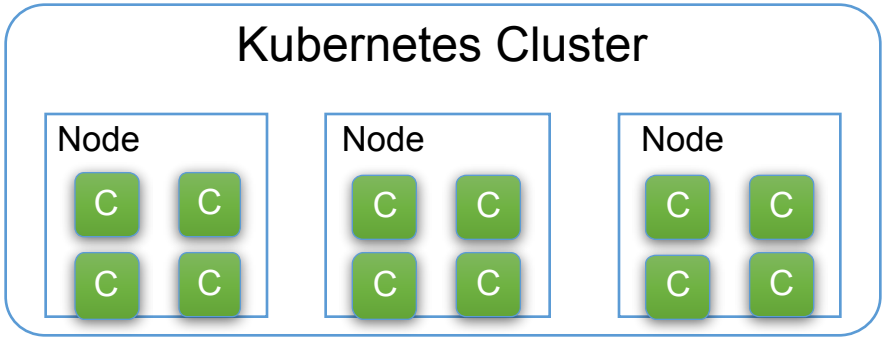

Various volume plugins

Shared Volume

Unshared Volume

- nfs
- hostPath
- glusterfs
- flocker
- cephfs
- secret
- vsphereVirtualDisk

- emptyDir
- hostPath
- gcePersistentDisk
- awsElasticBlockStore
- iscsi
- rbd
- azureFileVolume

Storage - Add your volume plugin

1. Register in kubelet entry

   kubelet/app/plugins.go

   allPlugins = append(allPlugins, customizedPlugin.ProbeVolumePlugins()...)

2. Implement it in the package below:

   pkg/volume/<your_plugin>, interface can refer to volume.go

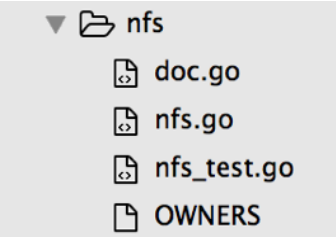

- ProbeVolumePlugins
- Init
- CanSupport
- SetUp
- TearDown
- …

3. Update API Spec

CI/CD - Typical

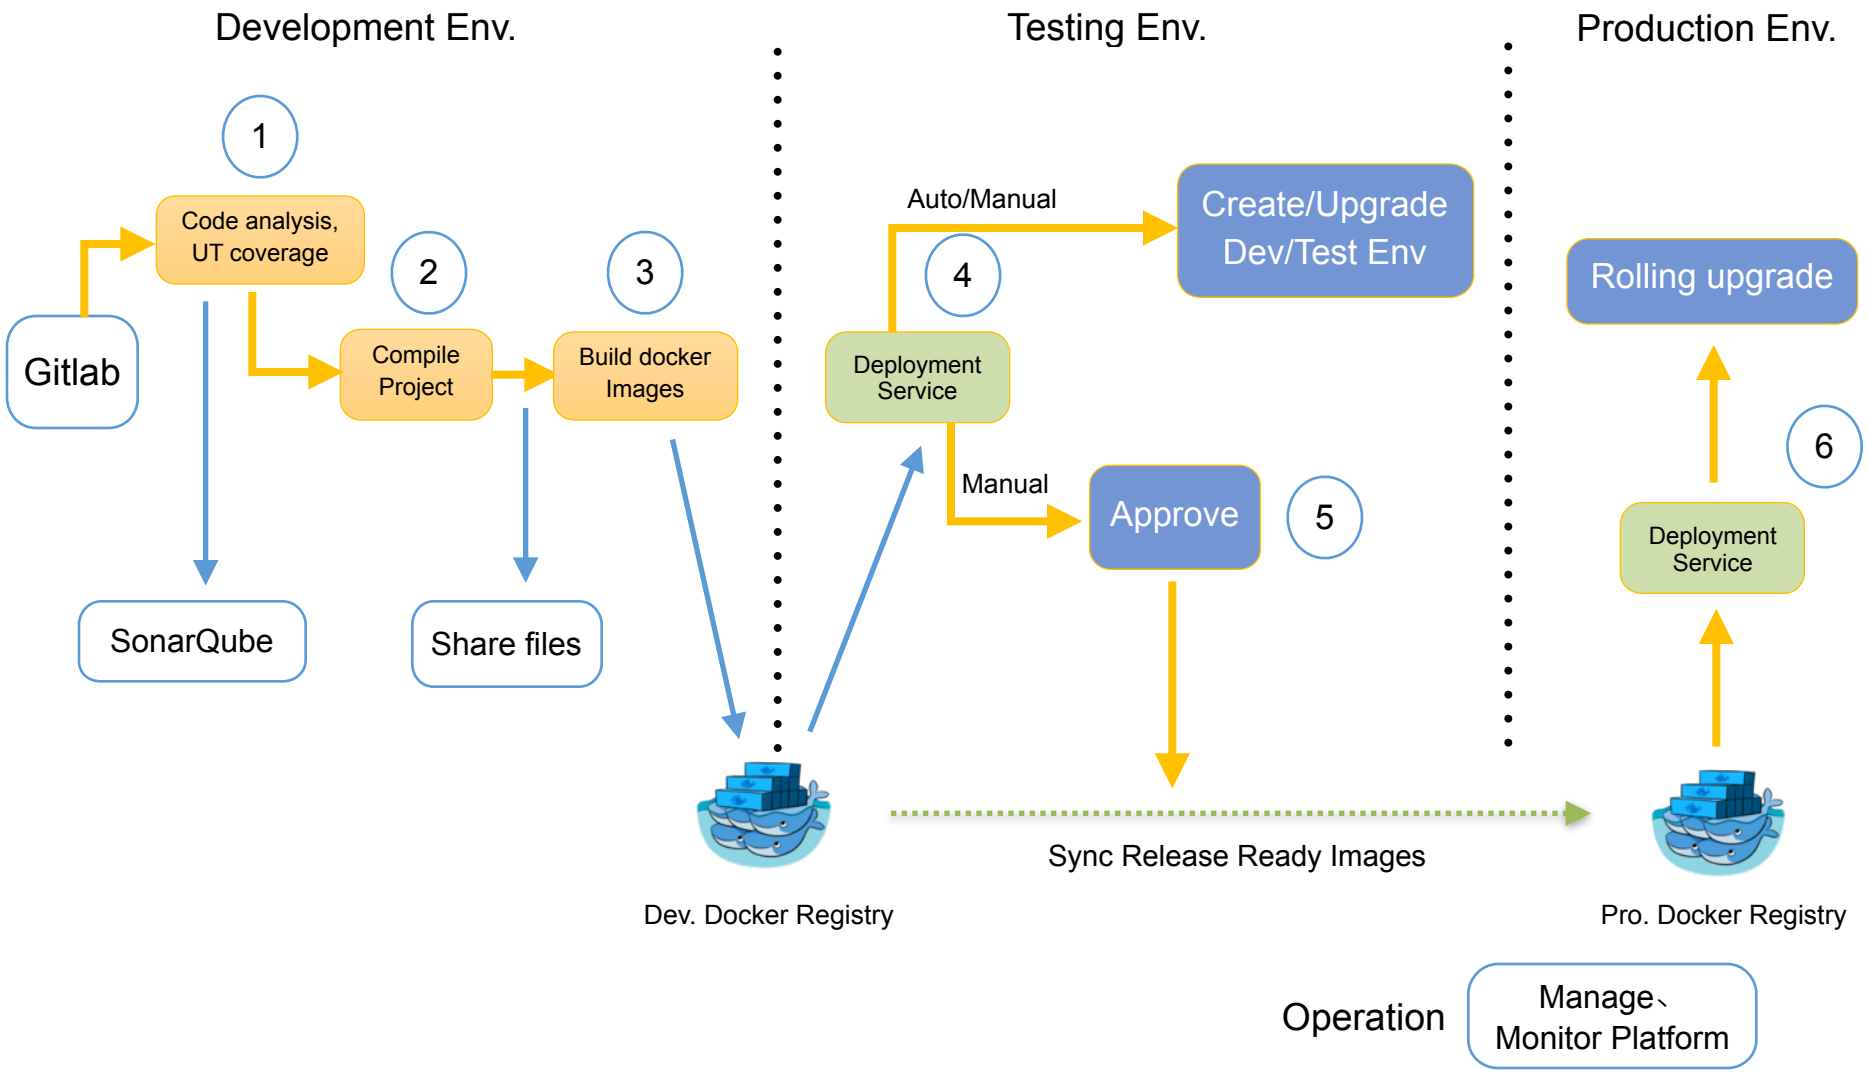

CI/CD - Example

1

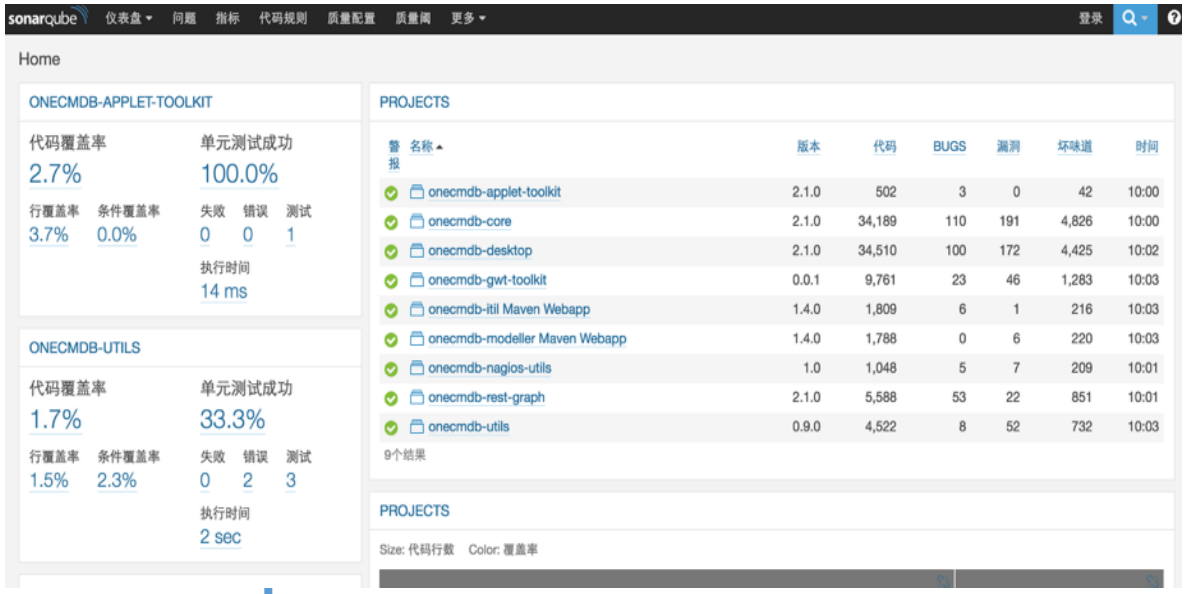

3

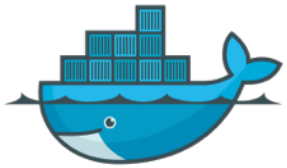

tenxcloud/tomcat-base:8.0
ADD *.war /tomcat/webapps/

Share war files to final image

Compile to generate war files

2

itil.war
onecmdb-itil-1.4.0.war
onecmdb-modeller.war
ROOT.war

No source code in final images
Only compiled code

Application Package and Deployment

Common Deployment Approach

FROM tomcat:8.0
Add v1.war /tomcat/webapps

FROM tomcat:7.0
Add v1.war /tomcat/webapps

FROM tomcat:8.0
Add v2.war /tomcat/webapps

FROM tomcat:7.0
Add v2.war /tomcat/webapps

FROM tomcat:8.0
Add v3.war /tomcat/webapps

FROM tomcat:7.0
Add v3.war /tomcat/webapps

M * N number of images

Deploy using Pod Orchestration Approach

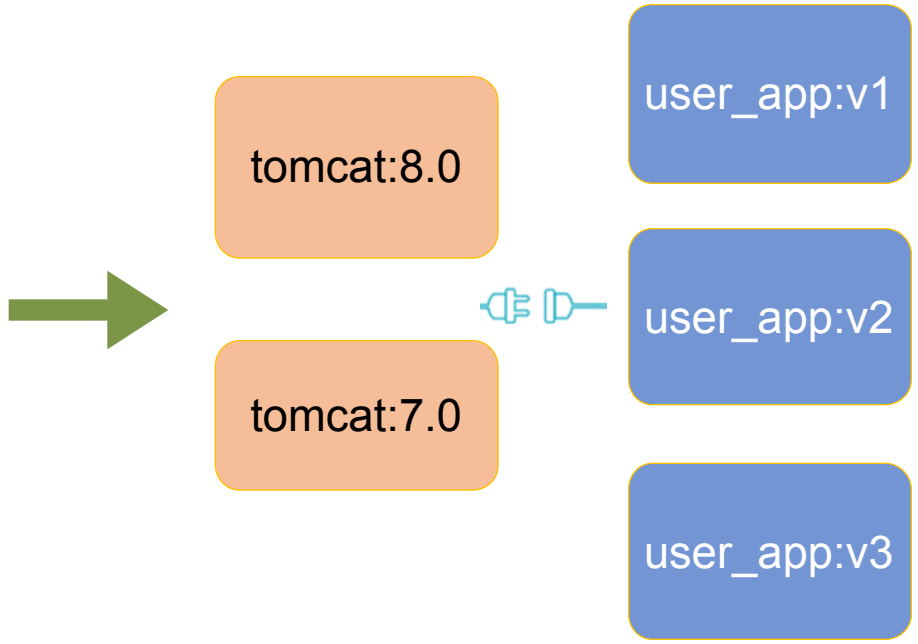

M + N number of images

CI/CD - Example

1

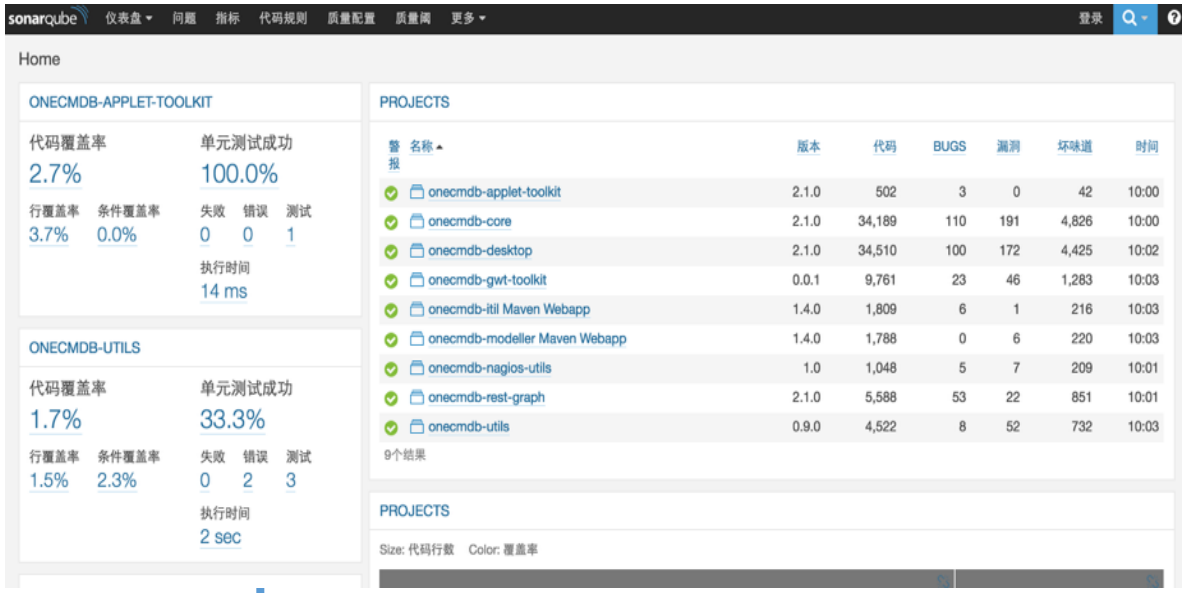

3

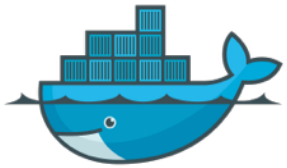

tenxcloud/busybox
ADD *.war /app_dir

Share war files to final image

Compile to generate war files

2

itil.war
onecmdb-itil-1.4.0.war
onecmdb-modeller.war
ROOT.war

No source code in final images
Only compiled code

Configuration Management

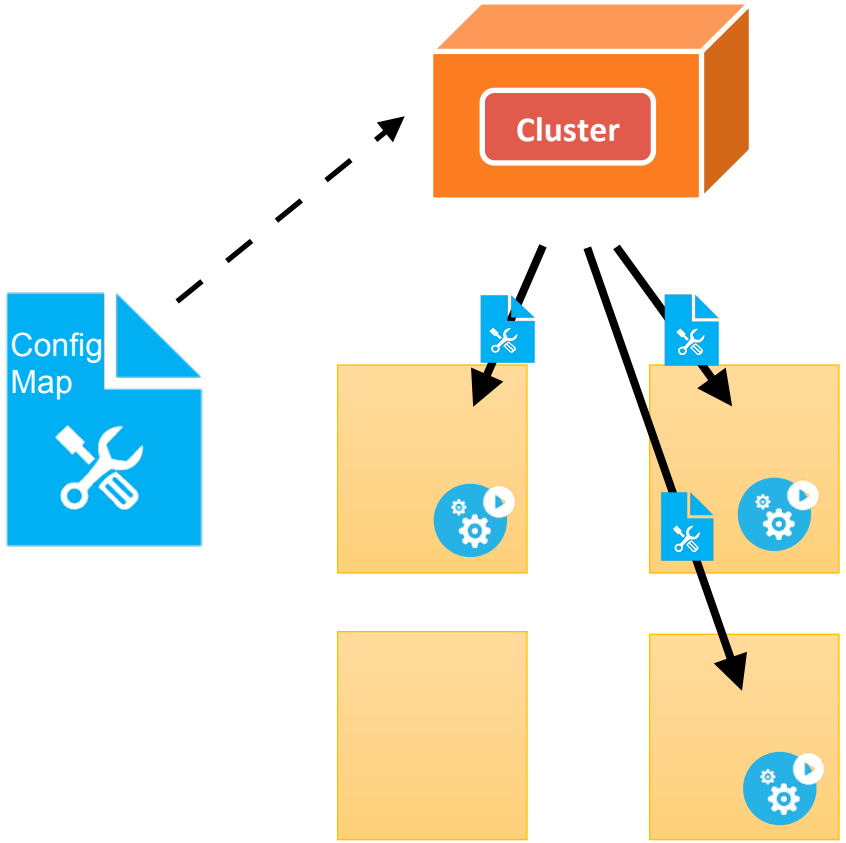

Centralized and cluster level configuration management

Features:

- Decouple config from image content
- Support environment variables, command-line arguments or as a volume
- Update associated application node automatically with new config once updated

Secret Data

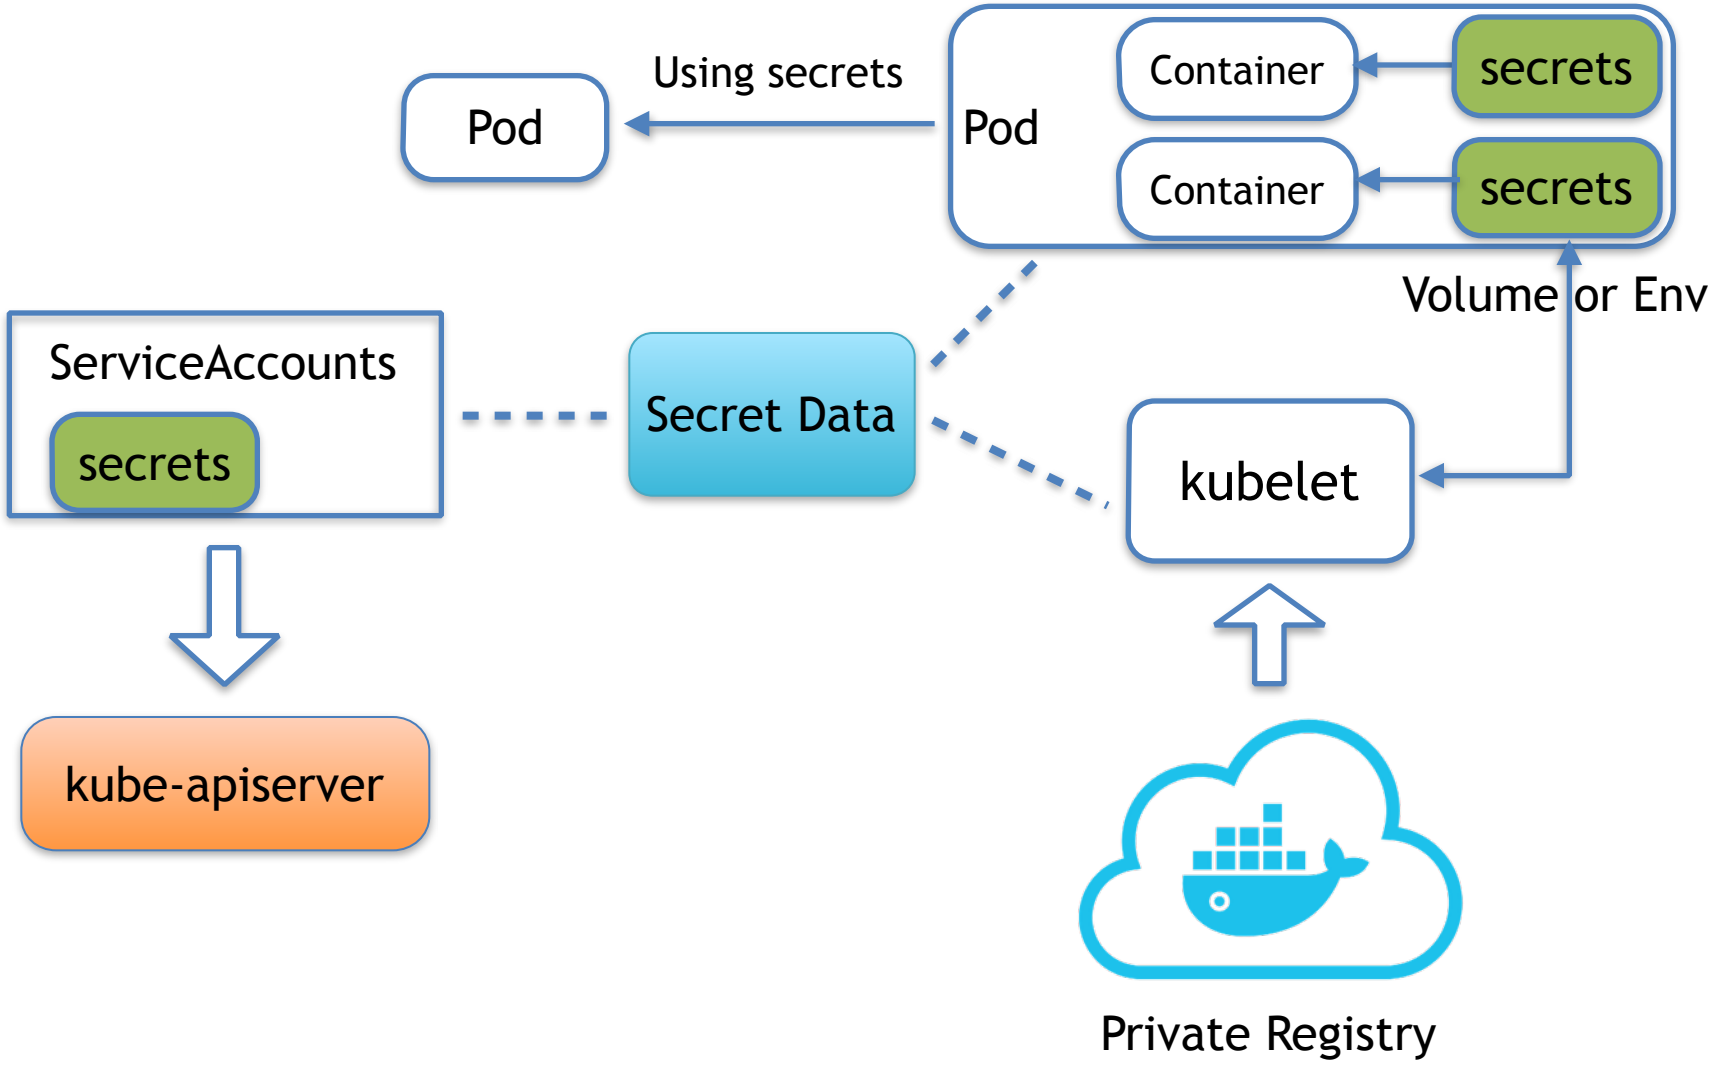
Pod
Using secrets
Pod
Container
secrets
Container
secrets
Volume or Env
ServiceAccounts
secrets
Secret Data
kubelet
kube-apiserver
Private Registry

Agenda

- Why Enterprise need Private PaaS?
- What Enterprise Cares about
  - Network
  - Storage
  - CI/CD
  - Application Package and Deployment
  - Configuration Management
  - Secret Data
  - Monitor and Alert
- **How to Manage Your PaaS**

Monitor and Alert

- Platform
- Kubernete Cluster- Components
- Node
- Service

- CPU
- Memory
- Disk
- QPS

Node Problem Detector

- A DaemonSet detects node problems and reports them to APIServer.

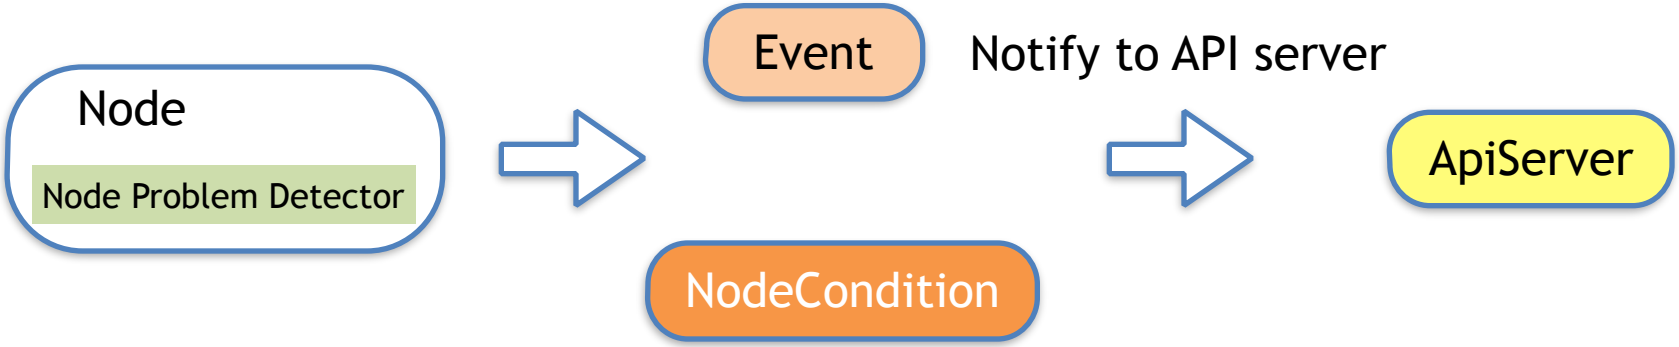


```
Events:
  FirstSeen     LastSeen        Count   From                            SubobjectPath   Type            Reason                  Message
  ---------     --------        -----   ----                            -------------   --------        ------                  -------
  1d            46m             5       {controllermanager }                            Normal          NodeNotReady            Node 192.168.1.189 status is now: NodeNotReady
  1d            46m             4       {kubelet 192.168.1.189}                         Normal          NodeHasSufficientDisk   Node 192.168.1.189 status is now: NodeHasSufficient
Disk
  1d            46m             3       {kubelet 192.168.1.189}                         Normal          NodeReady               Node 192.168.1.189 status is now: NodeReady
  4m            1m              2       {kernel-monitor 192.168.1.189}                  Warning         TaskHung                task docker:20744 blocked for more than 120 seconds
```

Operation

KUBECTL

- Debug
  - kubectl create/get/describe/delete (type name) —namespace=<ns_name>
- Node maintenance
  - kubectl drain/uncordon <node>
- kubectl patch
  - kubectl patch (type name) -p PATCH
- Resource annotation
  - kubectl annotate (type name) key=value
- kubectl label
  - kubectl annotate (type name) key=value

KOPS

Clusters operation tool

- kops create cluster
- kops update cluster
- kops get cluster
- kops delete cluster

Upgrade

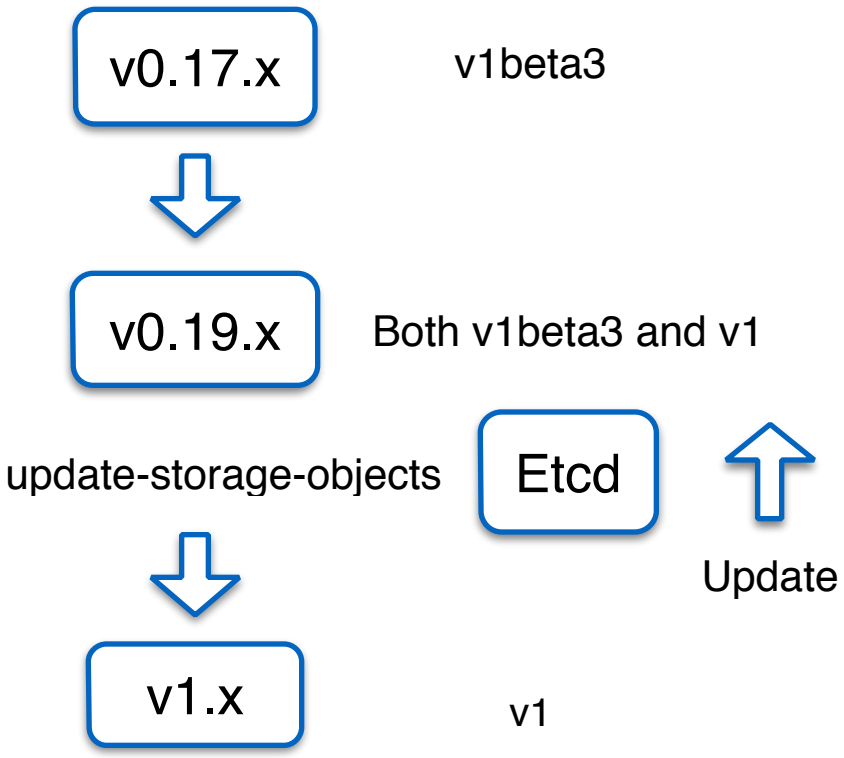

- No hard breaking changes over version boundaries

- Kubernetes Version Definition
  - X.Y.Z: X=>major, Y=>minor, Z=>patch
  - API vX[betaY]
- Upgrade
  - ✓ Upgrade from 1.x to any other 1.x release as rolling upgrade
    - Master first and then each node
  - ✓ Run latest patch releases of a given minor release
  - ✗ Upgrade more than two minor releases at a time

THANKS!