Kubernetes in production

Tomáš Kukrál

LinuxDays 2017

2017-10-07

About me

Tomáš Kukrál

tom@6shore.net

@tomkukral

Cloud Architect at Mirantis
MCP Kubernetes

Python developer

Prev: Infra engineer at FIT CTU

From Docker to Kubernetes

- Containers to pods
- Cluster network
- Modular architecture
- Desired state

Kubernetes resources

- node - machine
- pod - group of containers
- rc - replication-controller
- svc - service
- pv - persistent volume
- pvc - pv claim

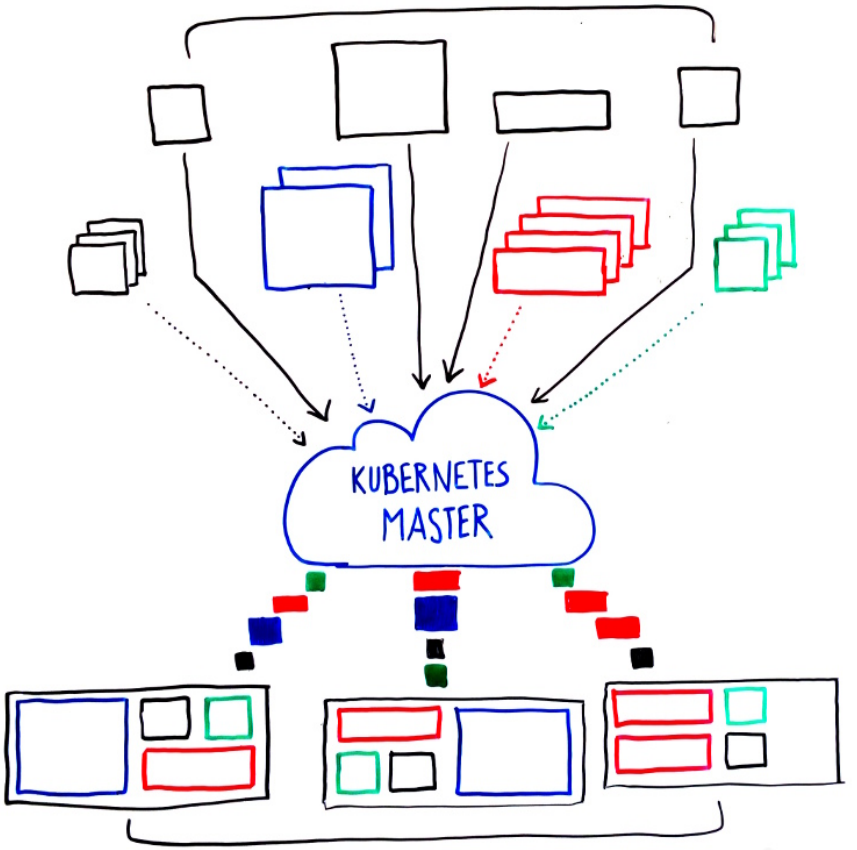
KUBERNETES
MASTER

```
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: flask
spec:
  replicas: 3
  template:
    metadata:
      labels:
        app: flask
    spec:
      containers:
      - name: nginx
        image: tomkukral/flask-app-demo
        imagePullPolicy: Always
        ports:
        - containerPort: 5000
        env:
          - name: REDIS_MASTER_SERVICE_HOST
            value: redis
```

Let's install Kubernetes

- Minikube
- Picokube
- Kubeadm
- Kubespray
- Salt formula Kubernetes
- Hosted solutions

Give me a HA!

Know you workloads

- Dynamic vs static
- Stateful vs stateless
- Multi vs single worker application

Kubernetes control plane

- Etcd
- Apiserver
- Scheduler, controller-manager
- Kubelet
- Proxy

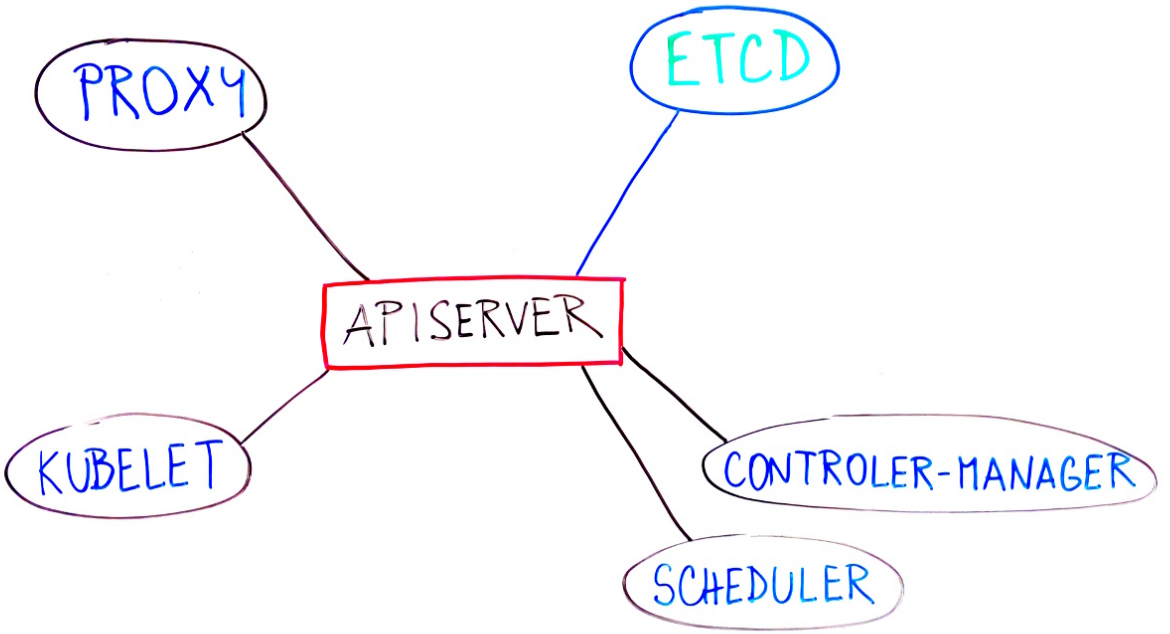
PROXY
ETCD
APISERVER
KUBELET
CONTROLER-MANAGER
SCHEDULER

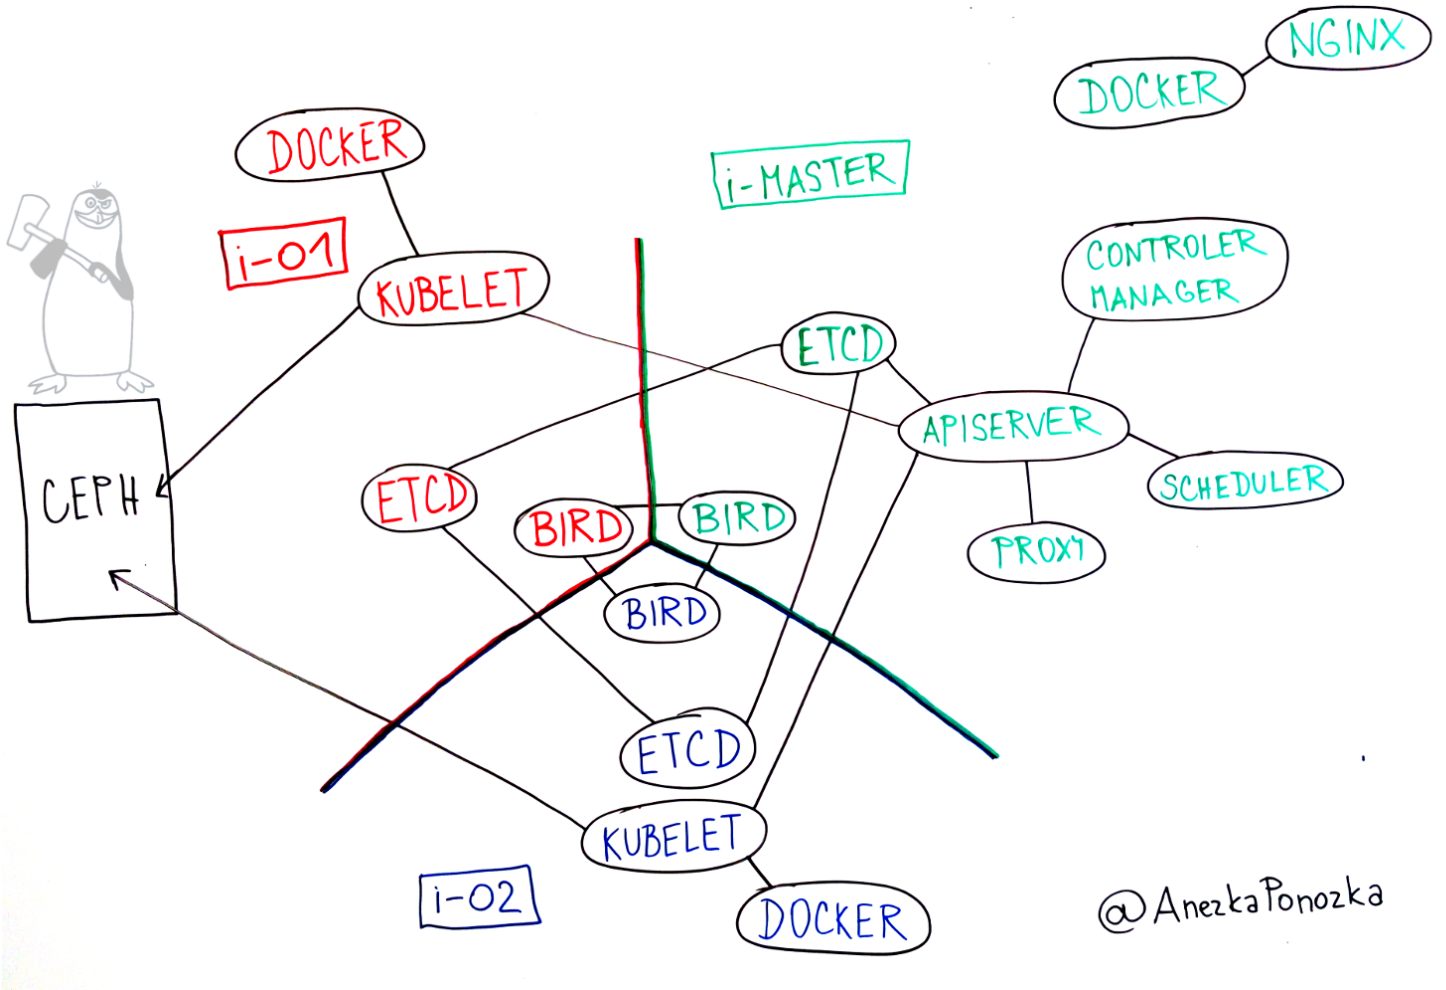
DOCKER
NGINX
i-MASTER
DOCKER
i-01
KUBELET
CONTROLER MANAGER
ETCD
APISERVER
CEPH
ETCD
SCHEDULER
BIRD
BIRD
PROXY
BIRD
ETCD
KUBELET
i-02
DOCKER
@AnezkaPonozka

Storage

- Kubernetes is dynamic, storage should be dynamic
- Ceph RBD or CephFS
- AWS EBS, GCP disks
- NFS, hostPath

Network

- Customer traffic
- Overlay vs underlay
- Don't use ClusterIP for customer traffic
- LoadBalancer or Ingress controller
- Advanced tools for microservice: istio, linkerd

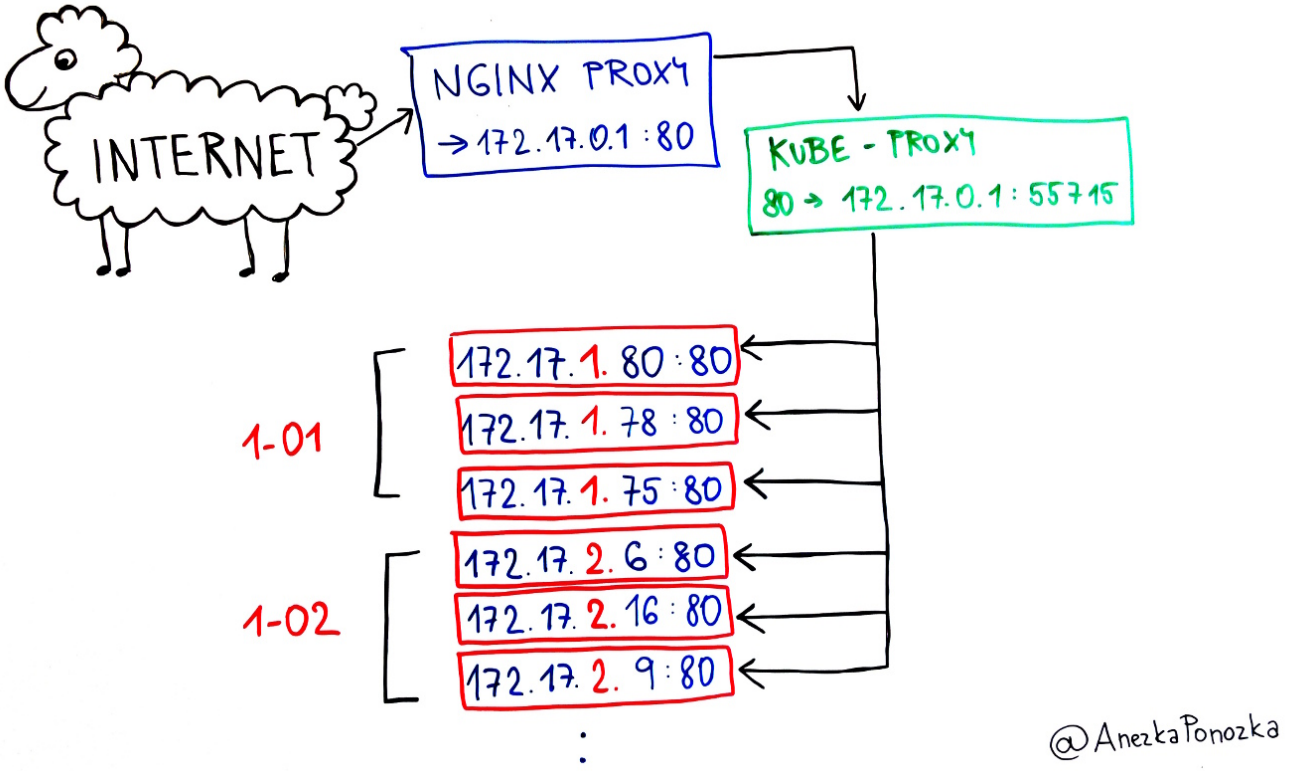
INTERNET
NGINX PROXY
→ 172.17.0.1 : 80
KUBE - PROXY
80 → 172.17.0.1 : 55715
1-01
172.17.1.80 : 80
172.17.1.78 : 80
172.17.1.75 : 80
1-02
172.17.2.6 : 80
172.17.2.16 : 80
172.17.2.9 : 80
@AnezkaPonozka

Know your images.

- Diversity vs unification
- Kernel features
- Different libraries and versions

- Image building pipeline

- Registries: registry:2, Atomic registry

- Audit images running in cluster

- What is `gcr.io/google_containers/pause-amd64:3.0`?

Tune your Kubernetes

- Etcd - `snapshot-count`
- Apiserver - `target-ram-mb`, `max-*requests-inflight`
- Controller-manager - `concurrent-*`
- Kubelet - `max-pods`, `*-reserved`
- All daemons - `kube-api-burst`, `kube-api-qps`
- Applatix: Making Kubernetes Production Ready – Part 2

LCM plan

- update packages?
- minor/major updates
- etcd updates
- networking updates

Component dependencies

- Everything is using apiserver
- apiserver is using etcd
- calico?
- Registry
- DNS service

Inteligent monitoring

- Pods starts and die
- Workload changes
- Adding (and removing) new minions

Backup

- Scope of the backups
- Recovery scenario
- Recover or reborn?
- Configuration

You don't need istio and linkerd for you web app with 3 containers.

Let's run everything in Kubernetes!

Great, let's parse `kubectl` output!

Demo time!

Thanks for you attention.

See you at Containers Meetup!