

Postgresql in Kubernetes

- ▶ Armin Nesiren
 - ▶ Developer
 - ▶ 10 years as DBA
 - ▶ Devops
 - ▶ Lately dealing with Kubernetes in production

- ▶ Service oriented aritecture - SOA
- ▶ Future of kubernetes
- ▶ Test it home, it's free
- ▶ Postgresql available for everyone, quickly

- ▶ Kubernetes basics
- ▶ Kubernetes storage
- ▶ Deploying PostgreSQL inside kubernetes simple way
- ▶ Deploying PostgreSQL automated

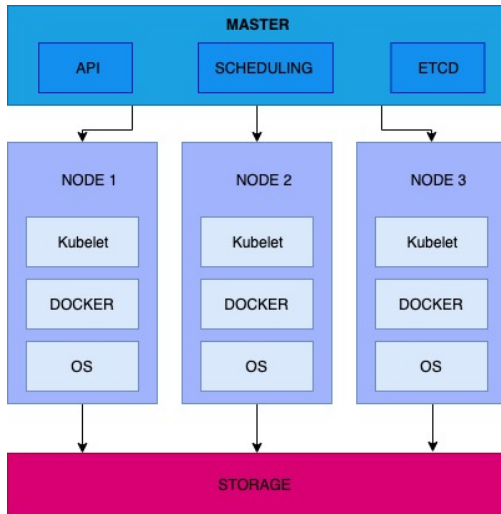
- ▶ Containers generally
 - ▶ We like containers because: lightweight, standalone, on any env
 - ▶ Resource efficient
- ▶ K8s is dealing with containers
 - ▶ Written in Golang
 - ▶ scaling, loadbalancing



Figure 1: k8s providers

- ▶ **Bare metal:**
 - ▶ as low as 3-5 nodes
 - ▶ needs to take care of persistent storage
 - ▶ exposing of services

K8s simple architecture



Rook Architecture

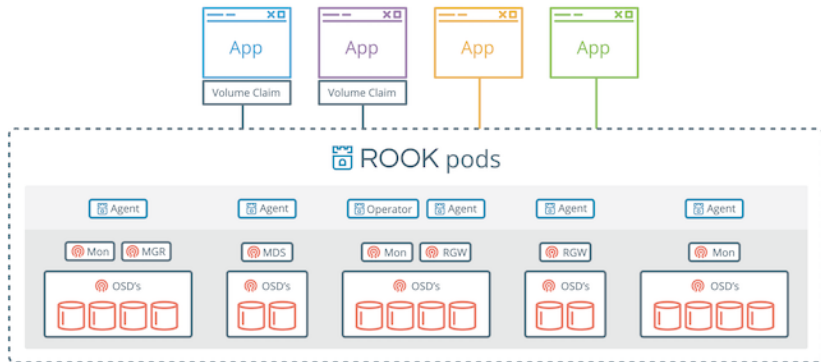


Figure 3: rook and ceph architecture

- ▶ **Concept of pods**
 - ▶ Group of one or more containers
 - ▶ shared storage/network
 - ▶ always co-located and co-scheduled
- ▶ **namespaces**
 - ▶ virtual cluster inside of a cluster
- ▶ **secrets**
 - ▶ storing sensitive information
- ▶ **Deployment of pods**
 - ▶ describing desired state of a pod
- ▶ **Services**
 - ▶ ClusterIP, NodePort, LoadBalancer, ExternalName

- ▶ kubectl
 - ▶ command line
- ▶ Yaml
- ▶ Client library
 - ▶ Go, Python, Java, dotnet, JavaScript and more

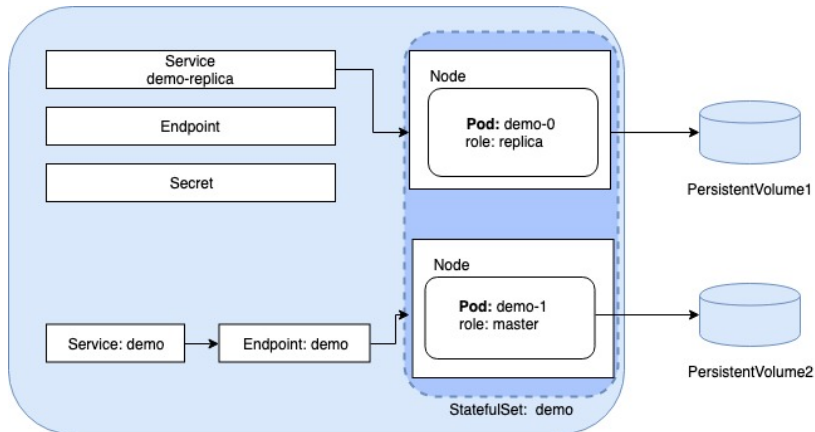
```
1  apiVersion: apps/v1beta1
2  kind: Deployment
3  metadata:
4  |   name: nginx-deployment
5  specs:
6  |   replicas: 3
7  |   template:
8  |     metadata:
9  |       labels:
10 |         app: nginx
11 |     specs:
12 |       containers:
13 |         - name: hello-world
14 |           image: hello-world:latest
15 |           ports:
16 |             - containerPort: 80
```

Figure 4: yaml example

- ▶ All supported PostgreSQL versions inside one image
- ▶ Plenty of extensions (pg_partman, pg_cron, postgis...)
- ▶ Additional tools (pgq, pgbouncer, wal-e/wal-g)
- ▶ PGDATA on external volume
- ▶ Environment-variables based configuration
- ▶ Light! ~120MB

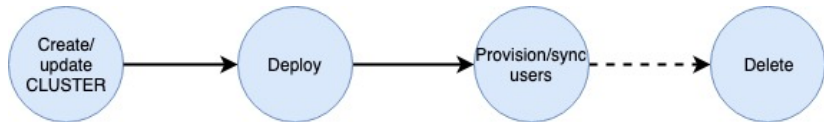
And than we also need... Patroni

- ▶ Automatic failover solution for streaming replication
- ▶ A python daemon that manages one PostgreSQL instance
- ▶ Keeps the cluster state in DCS (etcd, Zookeeper, Consul)
 - ▶ Uses DCS for leader elections
- ▶ Helps automate things like:
 - ▶ New cluster deployment
 - ▶ Scaling out and in
 - ▶ Postgresql configuration management



- ▶ Long YAML
- ▶ Hard to configure
- ▶ Not easy to manipulate with clusters (update, delete...)
- ▶ Manual generation of DB objects after setup

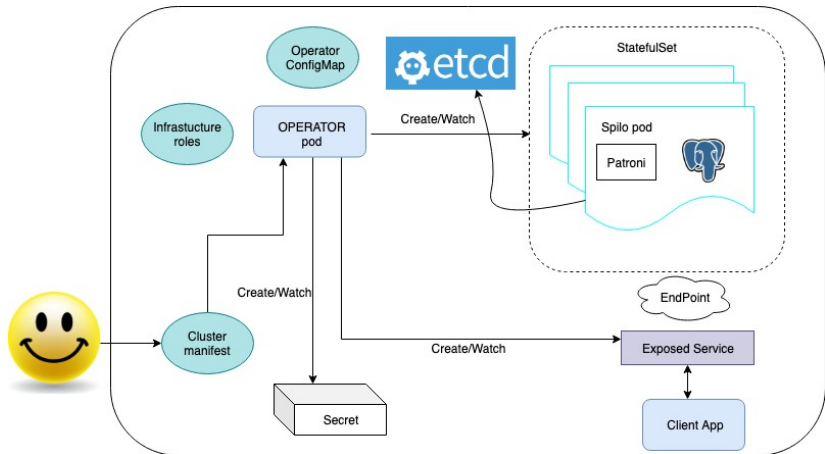
- ▶ Templated manifests
- ▶ Requires special pod (tiller) in your cluster
- ▶ Only one place for configuration
- ▶ Still problem with cluster update



- ▶ Encapsulates knowledge of a human operating the service
- ▶ Fully automated: deployments, cluster upgrades, user management

- ▶ Defines CDR (Postgresql)
- ▶ Watches instances, creates/updates/deletes
- ▶ Create DB objects, roles, passwords. . .
- ▶ Allows updating of resources (memory, cpu, volumes), postgresql configuration
- ▶ Auto-repairs, smart rolling updates

```
1  apiVersion: "acid.zalan.do/v1"
2  kind: postgresql
3  metadata:
4    name: ct-minimal-cluster
5  specs:
6    teamID: "CT"
7    volume:
8      size: 1G
9    numberOfInstances: 2
10   users:
11     testuser:
12       - creatorole
13       - createdb
14   database:
15     foo: testdb
16   postgresql:
17     version: "10"
```



End

CYBERTEC
DATA SCIENCE & POSTGRES SQL

THANK YOU