



Running Stateful Applications in Containers with K8S Persistent Volumes and StatefulSets

[OpenStack Summit, Vancouver](#)

Kapil Arora
Solution Architect, NetApp EMEA
May 23, 2018





Agenda

- 1) Stateful Apps and Persistent Data
- 2) K8S Persistent Data Concepts
- 3) Static vs Dynamic Provisioning
- 4) NetApp Trident
- 5) PostgreSQL Demo



Stateful Apps and Persistent Data

Stateful Apps

For Example

- SQL Databases
 - MySQL, PostgreSQL, SQL Server
- NoSQL Databases
 - MongoDB, ElasticSearch, Redis, Cassandra
- Pub Sub Systems
 - Kafka
- Time series Databases
 - InfluxDB, Graphite
- Big Data
 - Splunk



Persistent Data

For Example

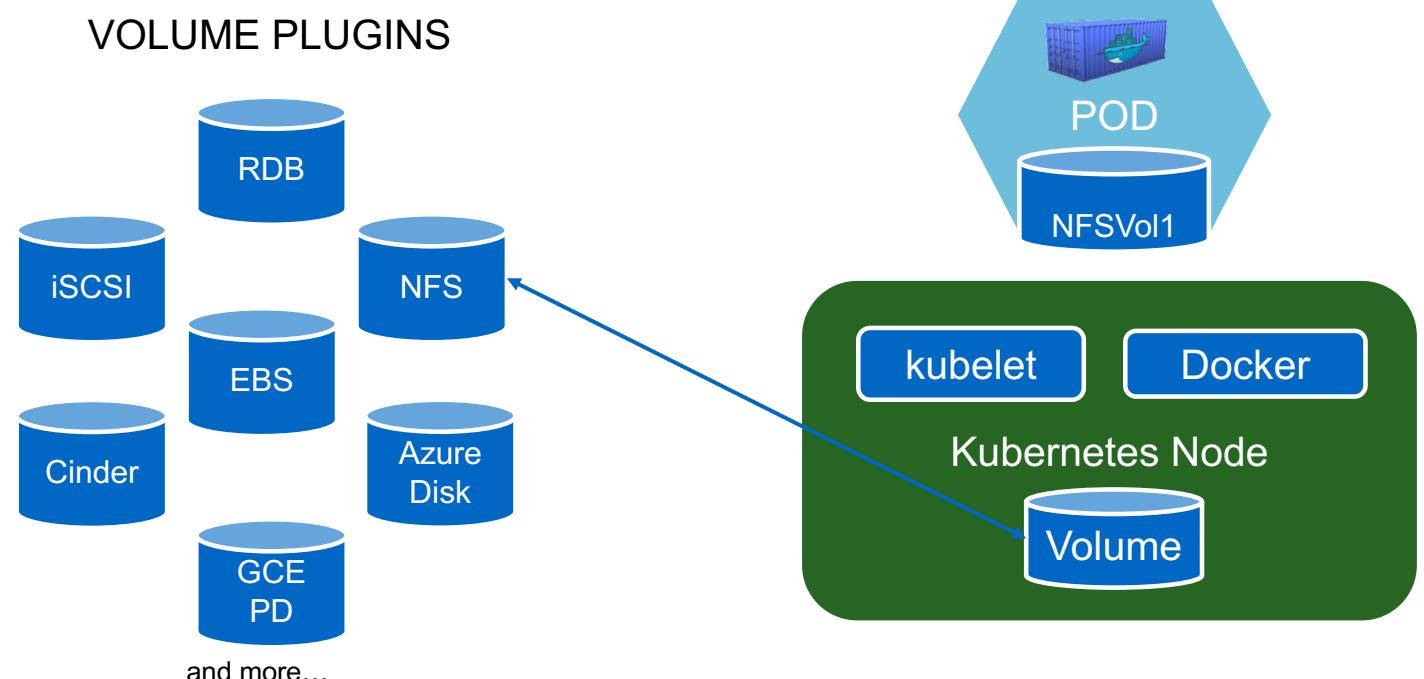
- Object
 - S3, Swift
- File
 - NFS, SMB, Manila, GlusterFS
- Block
 - iSCSI, FC, EBS, Cinder



Kubernetes Volume Plugins

Kubernetes Volume Plugins

- Reference block device and mounted file system
- Accessible by all containers in a pod
- Lifetime of a volume is same as the pod (or longer)



Kubernetes Volume Plugins

- Remote
 - GCEPersistentDisk
 - AWSElasticBlockStore
 - AzureFile, AzureDisk
 - iSCSI , FC
 - Flocker
 - NFS
 - RBD, CephFS, GlusterFS
 - Cinder
 - VsphereVolume, VMware Photon
 - Quobyte Volumes
 - Portworx Volumes
 - ScaleIO Volumes
 - StorageOS
- Ephemeral
 - Empty Dir
 - Secret
 - ConfigMap
- Local Persistent Volume (beta)
- Out-of-tree
 - CSI (beta)
- Other
 - Host path

Example

- Cinder Volume in a Pod

```
1  apiVersion: v1
2  kind: Pod
3  metadata:
4    name: cinder-web
5  spec:
6    containers:
7      - name: web
8        image: nginx
9        ports:
10       - name: web
11         containerPort: 80
12         protocol: tcp
13       volumeMounts:
14         - name: html-volume
15           mountPath: "/usr/share/nginx/html"
16     volumes:
17       - name: html-volume
18         cinder:
19           # Enter the volume ID below
20           volumeID: volume_ID
21           fsType: ext4
```

Example

- EBS Volume in a Pod

```
1  apiVersion: v1
2  kind: Pod
3  metadata:
4    name: aws-web
5  spec:
6    containers:
7      - name: web
8        image: nginx
9        ports:
10       - name: web
11         containerPort: 80
12         protocol: tcp
13       volumeMounts:
14         - name: html-volume
15           mountPath: "/usr/share/nginx/html"
16     volumes:
17       - name: html-volume
18         awsElasticBlockStore:
19           # Enter the volume ID below
20           volumeID: volume_ID
21           fsType: ext4
```

Example

- iSCSI volume in a Pod

```
1  ---
2  apiVersion: v1
3  kind: Pod
4  metadata:
5    name: iscsipd
6  spec:
7    containers:
8      - name: iscsipd-ro
9        image: kubernetes/pause
10       volumeMounts:
11         - mountPath: "/mnt/iscsipd"
12           name: iscsivol
13       volumes:
14         - name: iscsivol
15           iscsi:
16             targetPortal: 127.0.0.1
17             iqn: iqn.2015-02.example.com:test
18             lun: 0
19             fsType: ext4
20             readOnly: true
21             chapAuthDiscovery: true
22             chapAuthSession: true
23             secretRef:
24               name: chap-secret
```



PV and PVC Abstraction

Decouple storage implementation from consumption

Administrator creates a PV

- PV: PersistentVolume API Object
- Configured for backing storage device
 - NFS, iSCSI, Cinder, AWS EBS, GCE, Azure
- Includes connection information for the storage volume
- Example: NFS PV

```
1 apiVersion: v1
2 kind: PersistentVolume
3 metadata:
4   name: pv-5gig
5 spec:
6   capacity:
7     storage: 5Gi
8   accessModes:
9     - ReadWriteOnce
10  persistentVolumeReclaimPolicy: Recycle
11  nfs:
12    server: 10.0.101.107
13    path: /path_to/5gig
```

User Creates PVC

- PVC
 - PersistentVolumeClaim API Object
- Created by a user to request storage
- User provides:
 - Size
 - Access mode
 - Storage class (optional)
- Kubernetes assigns a PV to meet the requirements requested in the PVC
 - Does not require an exact match for capacity

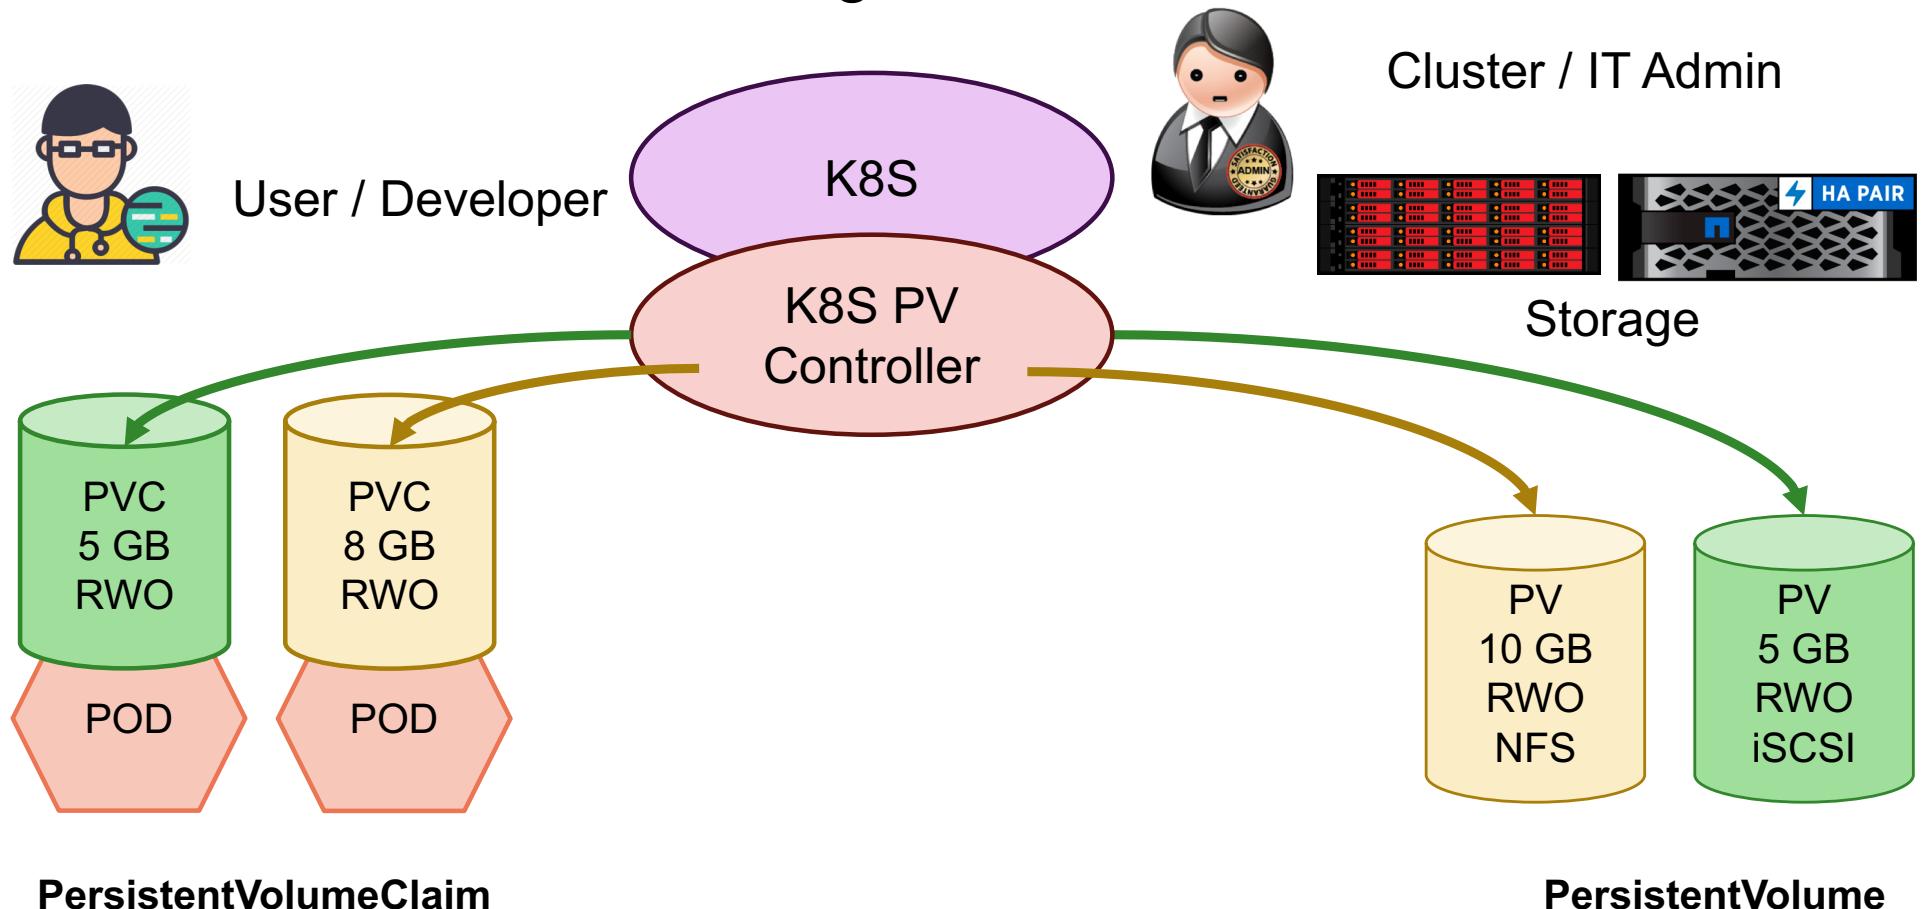
```
1 apiVersion: v1
2 kind: PersistentVolumeClaim
3 metadata:
4   name: pvc-5gig
5 spec:
6   accessModes:
7     - ReadWriteOnce
8   resources:
9     requests:
10      storage: 5Gi
```

User provides PVC name in the Pod/Deployment

- PVC claim name
- mount path

```
1  apiVersion: v1
2  kind: Pod
3  metadata:
4    name: pod-nginx
5  spec:
6    containers:
7      - image: nginx
8        name: nginx
9        ports:
10          - containerPort: 80
11            name: http
12        volumeMounts:
13          - name: my-vol
14            mountPath: /usr/share/nginx/html/
15        volumes:
16          - name: my-vol
17            persistentVolumeClaim:
18              claimName: pvc-5gig
```

Kubernetes Static Provisioning



PersistentVolumeClaim

PersistentVolume



Storage Classes and Dynamic Provisioning

StorageClass

- Administrators can classify different storage types in the K8s environment using Storage classes
- Example
 - SSD/HDD
 - QoS
 - Dedupe/compression
 - Protocol: iSCSI/NFS
 - Etc.
- Abstract underlying storage
- E.g. Gold::Silver::Bronze

SC Example

- Cinder

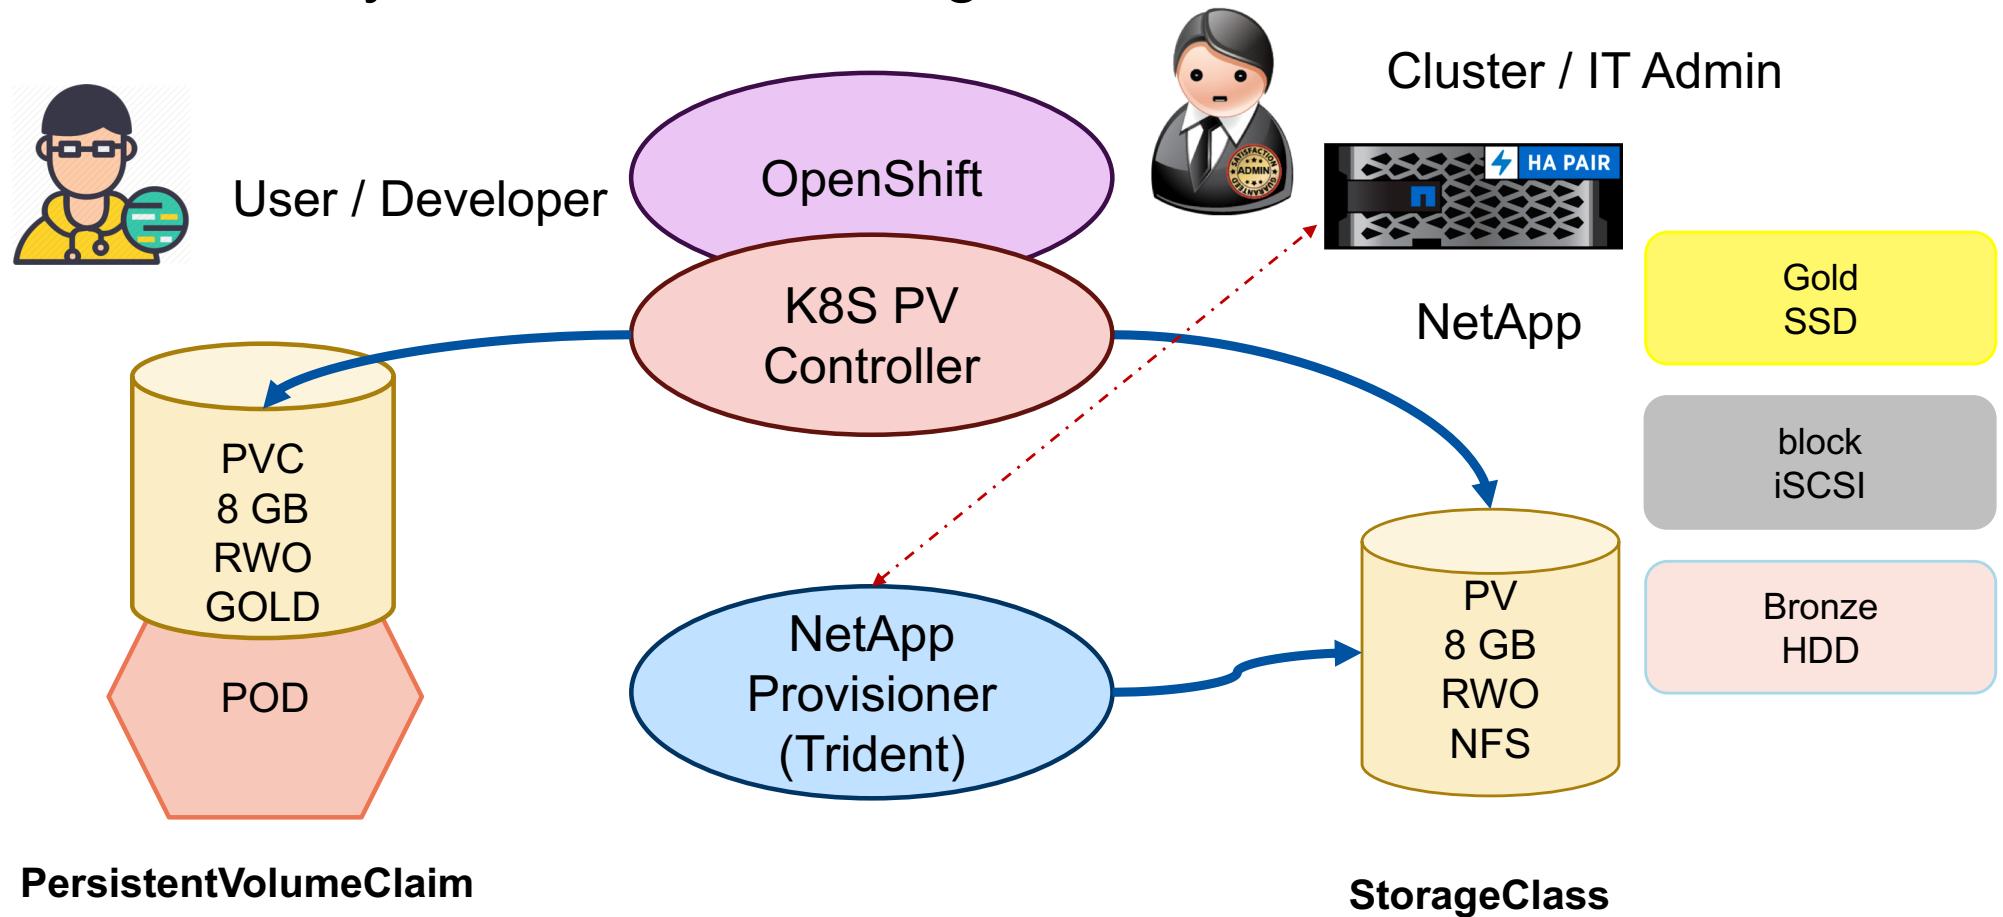
```
1 kind: StorageClass
2 apiVersion: storage.k8s.io/v1
3 metadata:
4   name: gold
5   provisioner: kubernetes.io/cinder
6 parameters:
7   type: fast
8   availability: nova
```

SC Example

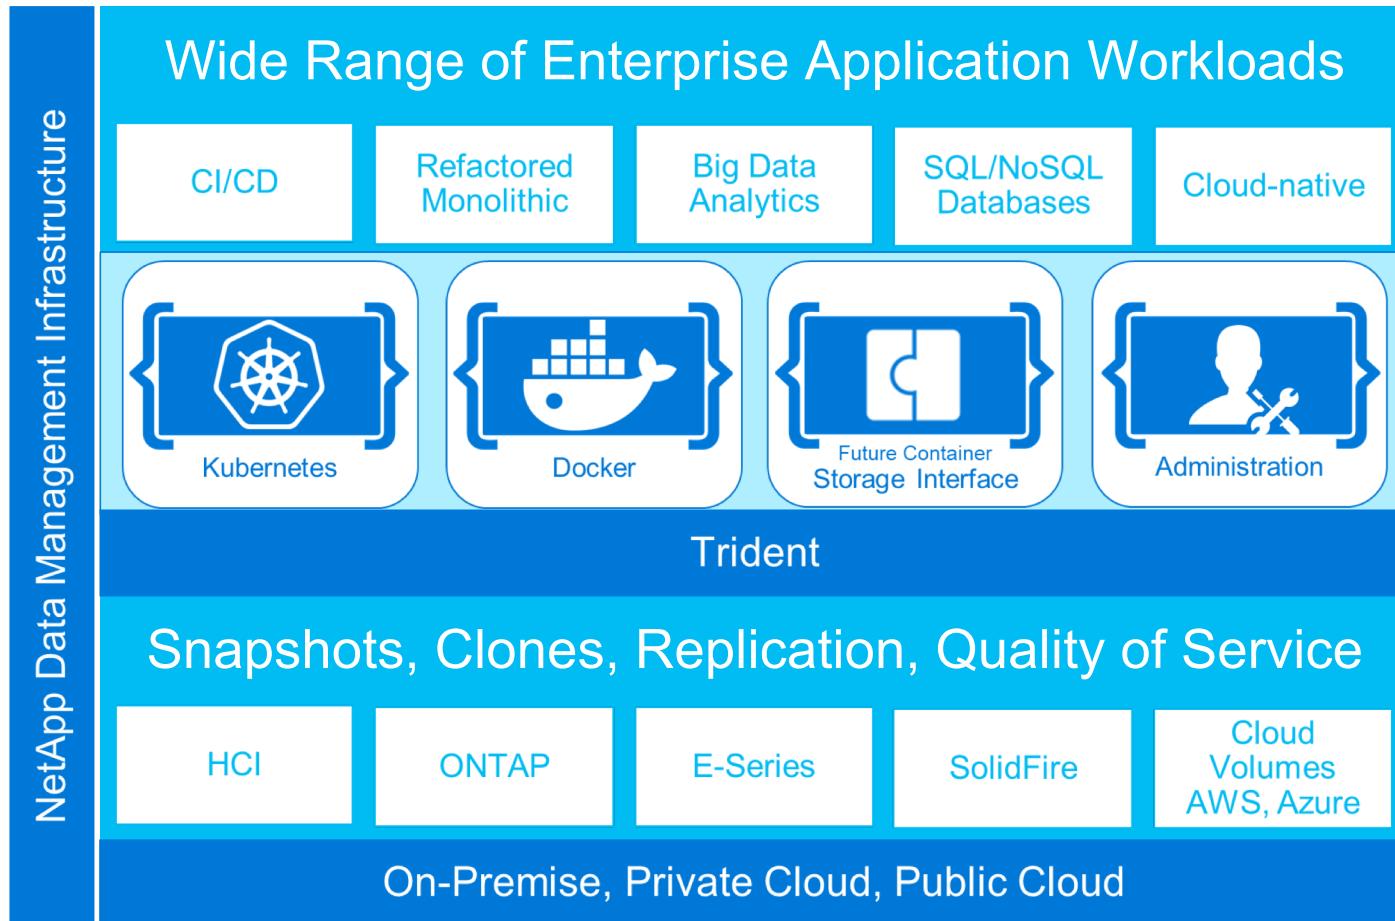
- NetApp ONTAP
 - NFS
 - SSD

```
1  apiVersion: storage.k8s.io/v1
2  kind: StorageClass
3  metadata:
4    name: ontap-gold
5  provisioner: netapp.io/trident
6  parameters:
7    backendType: "ontap-nas"
8    media: "ssd"
9    provisioningType: "thin"
10   snapshots: "true"
11
```

Kubernetes Dynamic Provisioning

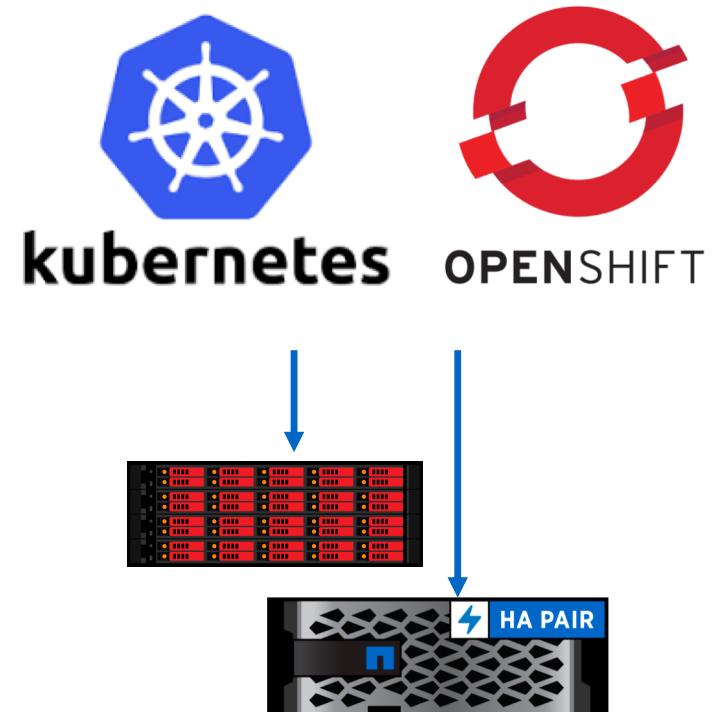


NetApp Trident



Trident for Kubernetes

- NetApp's Open source **dynamic storage provisioner** supports:
 - ONTAP
 - SolidFire
 - E-Series
- Automates Volume Creation and Mapping
- Compatible with:
 - Kubernetes
 - OpenShift Origin & RedHat OCP
- Snapshot and Cloning enabled
- Available on GitHub: <https://github.com/NetApp/trident>



Resource Quotas

- Kubernetes Resource Quotas
 - <https://kubernetes.io/docs/concepts/policy/resource-quotas/>
- Blog
 - <https://netapp.io/2017/06/09/self-provisioning-storage-kubernetes-without-worry/>
- Set Limits on
 - requests.storage
 - value.storageclass.storage.k8s.io/requests.storage
 - persistentvolumeclaims
 - value.storageclass.storage.k8s.io/persistentvolumeclaims

Future

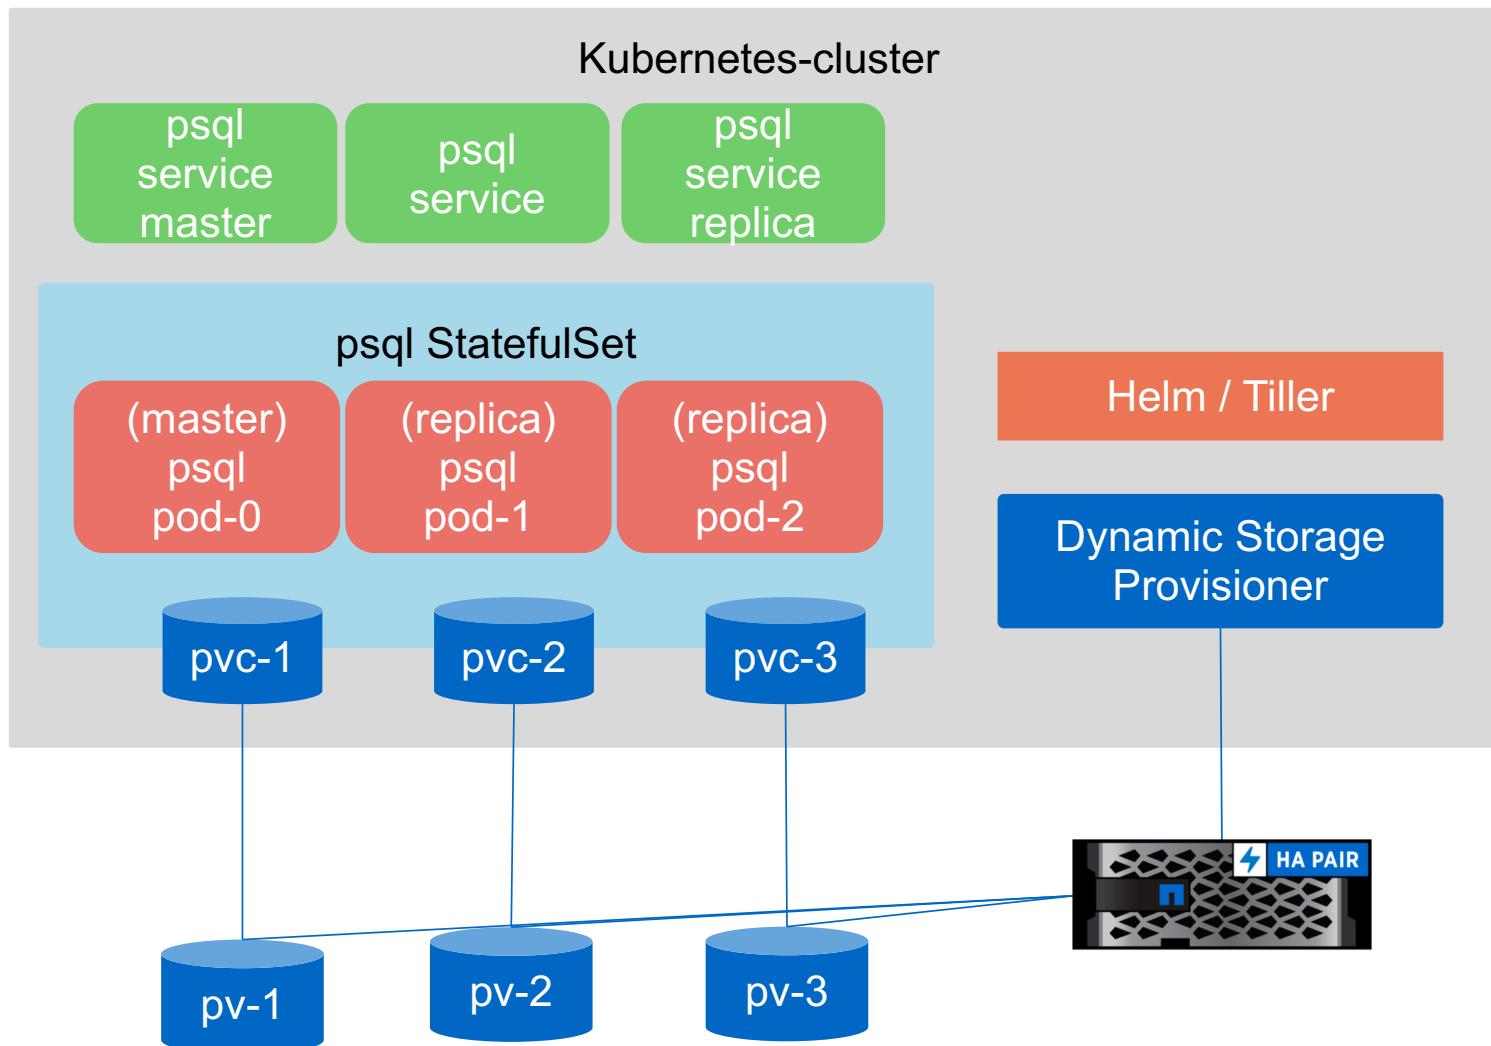
- Container Storage Interface (CSI)
 - Moving towards an industry standard, orchestrator-agnostic basic storage abstraction
 - Standard API
- Now Beta!
- Spec
 - <https://github.com/container-storage-interface/spec>
- Blog
 - <https://kubernetes.io/blog/2018/04/10/container-storage-interface-beta/>



CONTAINER
STORAGE
INTERFACE

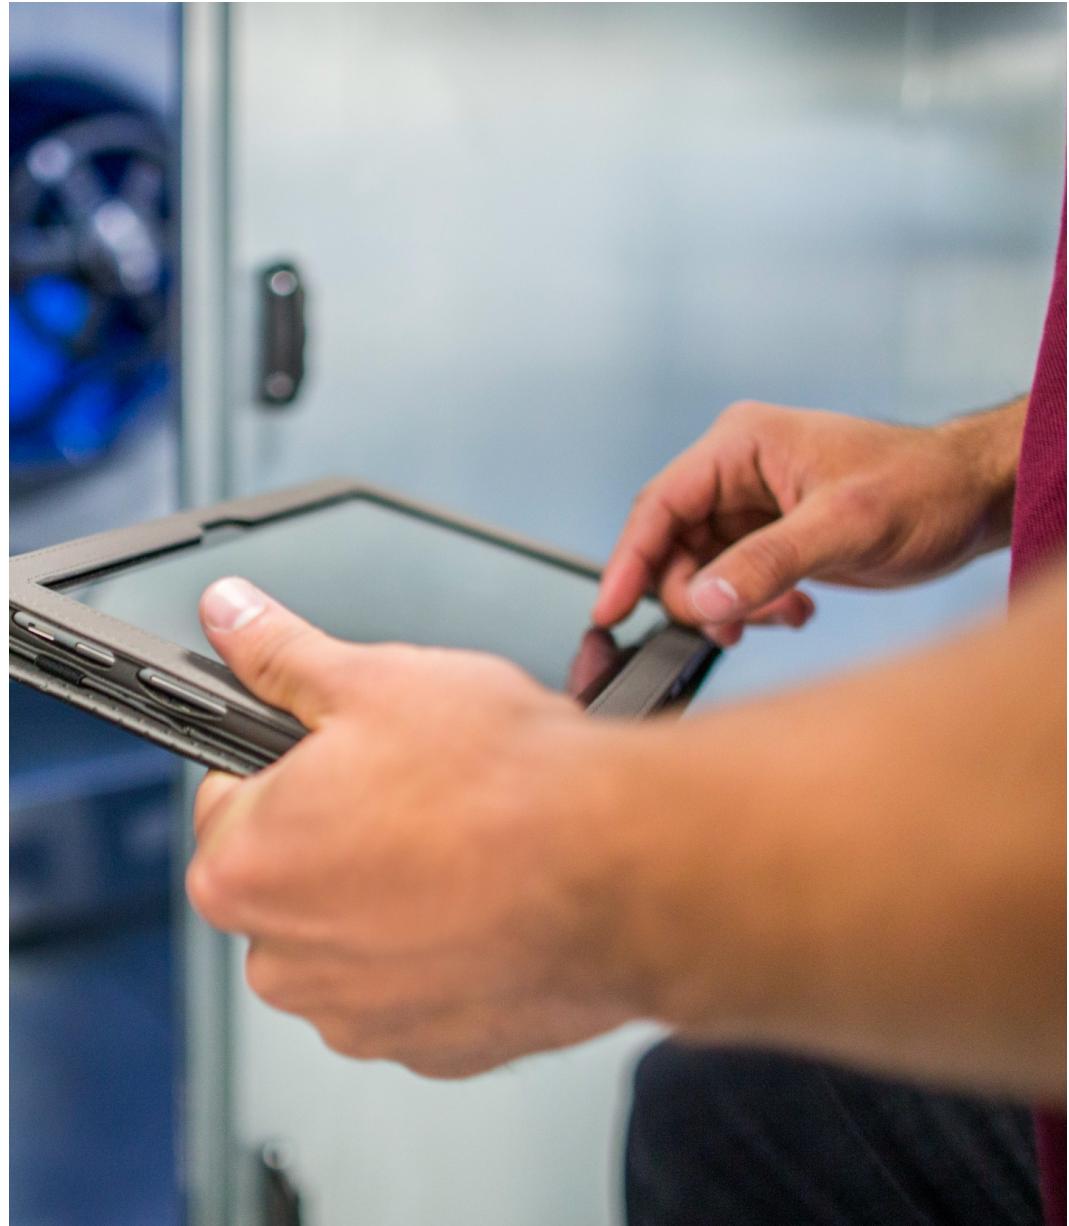


Demo : PostgreSQL



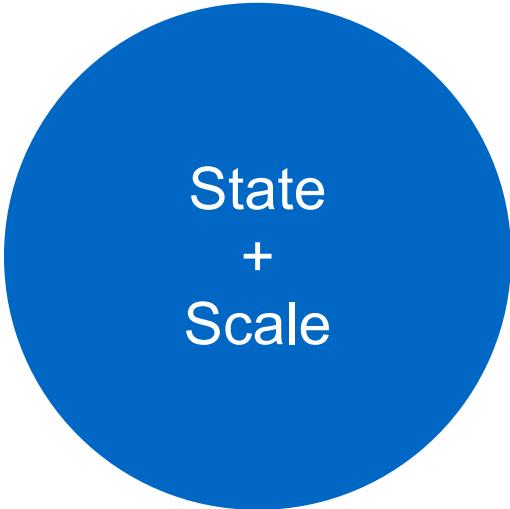
Why use Helm?

- Standardized deployments
- Build a catalog for users
- Repeatable and predictable deployments
- Follow best practices

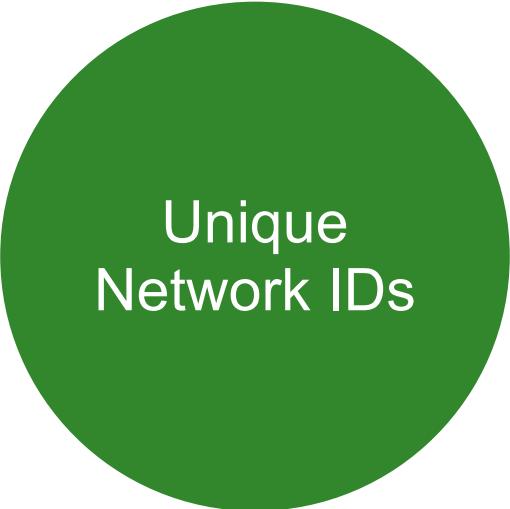


StatefulSets

[When and Why?](#)



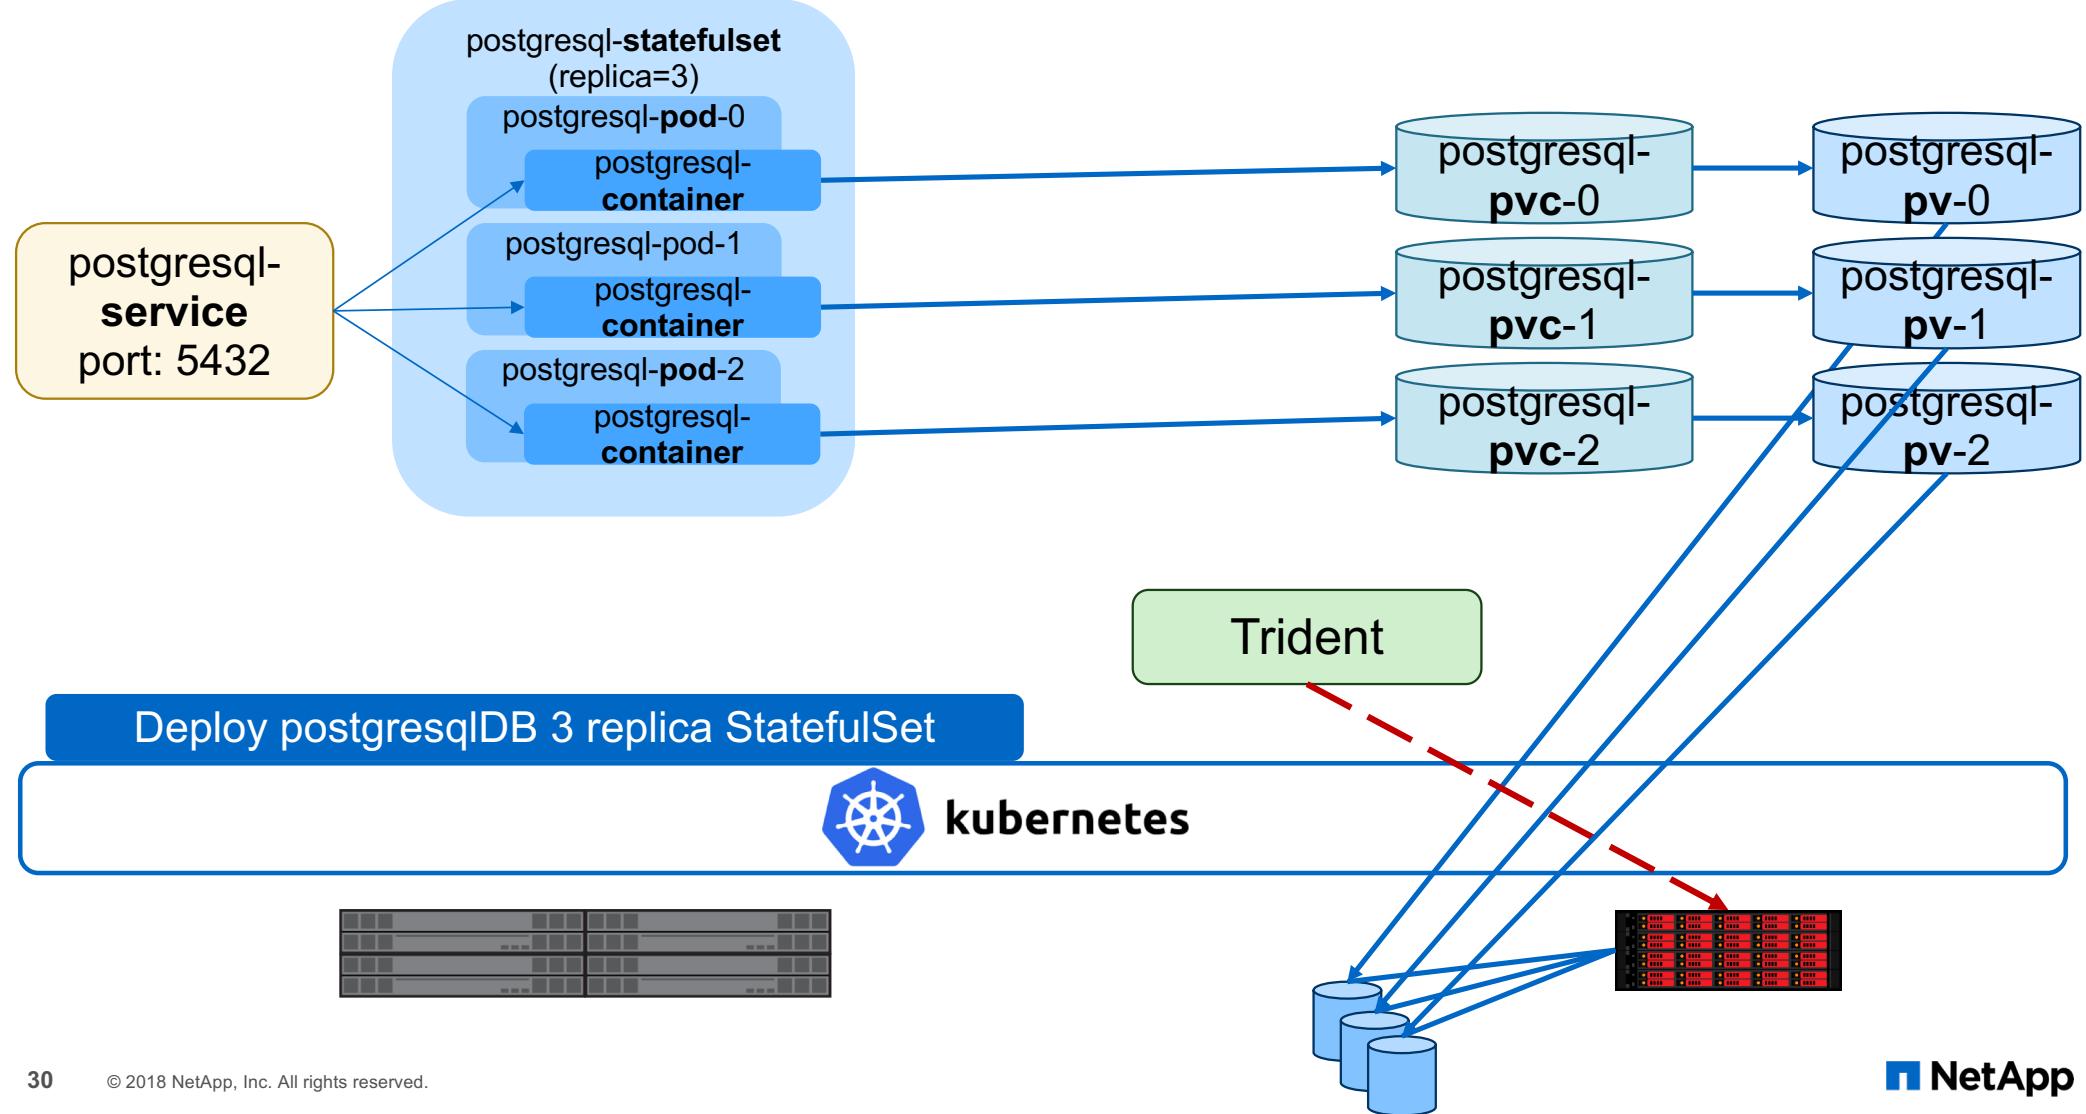
State
+
Scale



Unique
Network IDs



Ordered
Deployment
&
Termination



```
× root@akapila-kube-1... %1 × kapil@kapil-linux-ju... %2  
kapil@kapil-linux-jumphost:~$ _
```

Y

Up and running Kubernetes, Helm and NetApp trident



Summary

- Stateful Apps run on Kubernetes
- PV, PVC and StorageClass
- Static vs Dynamic provisioning
- NetApp Trident
- CSI



Resources

- Blog and Video
 - <https://netapp.io/2018/03/21/deploy-postgresql-using-kubernetes-helm-trident/>
- PostgreSQL sample Helm Chart
 - <https://github.com/NetAppEMEA/kubernetes-netapp>
- NetApp Trident
 - <https://github.com/NetApp/trident>

thePub: NetApp's Developer and Open Source Community

Engage with NetApp Engineers, Partners, and Customers



thePub
netapp.io



Slack
netapp.io/slack



Twitter
@NetAppPub



NetApp website
nt-ap.com/2CneVeR



GitHub
github.com/NetApp/trident



Thank You