



KubeCon



CloudNativeCon

North America 2019

# How to Backup and Recover Your Kubernetes Cluster

*Dylan Murray & Annette Clewett*



# Overview



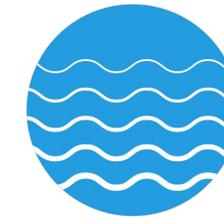
KubeCon



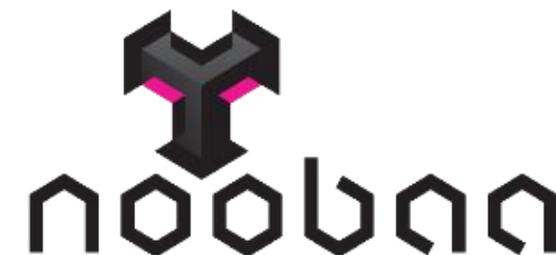
CloudNativeCon

North America 2019

- What is Velero?
  - How does it work?
    - Kubernetes Resources
    - Persistent Volumes
  - How is it integrated with Kubernetes?
  - What alternatives exist?
- What is Rook?
  - Architecture
  - Ceph on Kubernetes with Rook
- What is NooBaa?
- Example of Backup/Restore workflow
- Demo



VELERO



# What is Velero?



KubeCon



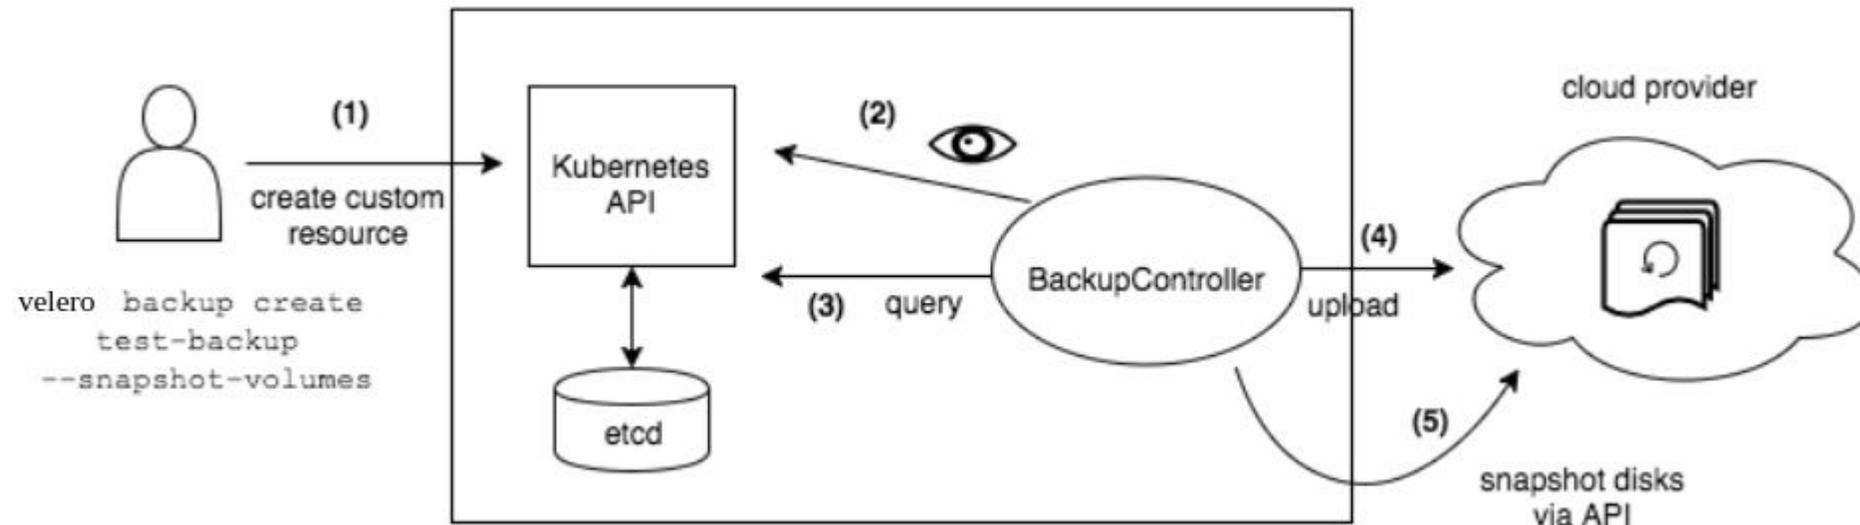
CloudNativeCon

North America 2019

- **Kubernetes Tool for Backup/Restore**
  - “Velero is an open source tool to safely backup and restore, perform disaster recovery, and migrate Kubernetes cluster resources and persistent volumes.”
    - Aiming to help with:
      - Disaster Recovery: *Recover from an issue*
      - Data Migration: *Migrate apps between clusters*
      - Data Protection: *Scheduled Actions*
    - GitHub: <https://github.com/vmware-tanzu/velero>
      - ~3 Years Old, ~100 contributors
      - Recent 1.2 release
    - Project Page: <https://velero.io/>
- **Originally ‘Ark’ from Heptio**, then Heptio acquired by VMWare.
  - Changed name to ‘Velero’ early 2019.

# How does it work?

- **Handles**
  - Kubernetes Resources
  - Persistent Volumes
- **Extends k8s API** and allows users to create notions of a **Backup or Restore object**
  - Velero will watch for these Backup or Restore objects
    - When it detects one it will process the information in the object which tells it what exactly to Backup or to Restore



# How does it work? - Resources



KubeCon



CloudNativeCon

North America 2019

- **Kubernetes Resources**
  - **Backup:**
    - Velero is able to export the objects in a namespace to their YAML
      - Similar to
        - `kubectl get pod mypod -o yaml` for all objects in a namespace, then storing that data together in a 'Backup'
          - Discovers all API groups present in a cluster, then queries for resources for each in the namespace. With this it's able to support various Custom Resources
  - **Restore:**
    - Parses a 'Backup' and is able to recreate the object definitions
      - Similar to
        - `kubectl create -f mypod.yml` for each resource yaml in the backup
- **Supports Plugins for customizing backup/restore actions per resource type**
  - Example on restoring an OpenShift route - change the domain used to match the cluster restored to
    - <https://github.com/fusor/openshift-velero-plugin/blob/master/velero-plugins/route/restor>  
[e.go](https://github.com/fusor/openshift-velero-plugin/blob/master/velero-plugins/route/restor)

# How does it work? - PVs



KubeCon



CloudNativeCon

North America 2019

- **Persistent Volumes**

- Able to backup/restore data inside of a Persistent Volume
- **Plugin approach**, supports multiple kinds of PVs

- Cloud Provider Snapshots

- AWS EBS
- Azure Managed Disk
- Google Compute Engine Disks
- Portworx
- DigitalOcean, etc...

- Filesystem Based

- **Restic**: NFS, Gluster, Ceph, etc

- Restic is a like an rsync implementation that will copy data at filesystem level and store in a de-duplicated efficient format in object storage, supporting later restores



# How does it work? - PVs



KubeCon



CloudNativeCon

North America 2019

- **Restic**

- Runs within Velero server
- During backup, Velero looks for an annotation on all backed up pods which tells it to create a `PodVolumeBackup` resource
- Restic operates on all `PodVolumeBackup` resources and stores the filesystem copy in object storage
- On restore, Restic restores the filesystem data to a new Persistent Volume
  - Managed by creation of a `PodVolumeRestore` which is initiated by Velero during the restore process



# How does it work? - Object Storage



KubeCon



CloudNativeCon

North America 2019

- Velero stores data to an 'Object Storage' provider, typically:
  - S3 Compatible API (NooBaa, Minio, AWS S3, etc)
  - Azure Blob Storage
  - Google Cloud Storage
- Persistent Volume data may be stored in the native mechanism for a cloud snapshot or if using restic will exist in the object storage along with the backup of k8s objects

# Kubernetes Integration



KubeCon



CloudNativeCon

North America 2019

## CLI/API Focus

Native k8s experience using Custom Resource Definitions.

```
kubectl create -f backup.yml
```

```
kubectl get backup mybackup
```

OR

```
velero create backup <flags>
```

```
$ oc get backup d8a692e0-a8ab-11e9-ac74-15162c64f165-d7bfx -n mig -o yaml
apiVersion: velero.io/v1
kind: Backup
metadata:
  annotations:
    openshift.io/migrate-copy-phase: stage
    openshift.io/migrate-quietse-pods: "true"
    openshift.io/migration-registry: 172.30.154.7:5000
    openshift.io/migration-registry-dir: /test-pv-4-registry-1bb598f5-a7e4-11e9-b6b2-0ac4ab43e1ee
  creationTimestamp: "2019-07-17T16:08:50Z"
  generateName: d8a692e0-a8ab-11e9-ac74-15162c64f165-
  generation: 1
  labels:
    migration: d9002c80-a8ab-11e9-b6b2-0ac4ab43e1ee
    migration-stage-backup: d9002c80-a8ab-11e9-b6b2-0ac4ab43e1ee
    velero.io/storage-location: jwms3-8dpmm
  name: d8a692e0-a8ab-11e9-ac74-15162c64f165-d7bfx
  namespace: mig
  resourceVersion: "2669296"
  selfLink: /apis/velero.io/v1/namespaces/mig/backups/d8a692e0-a8ab-11e9-ac74-15162c64f165-d7bfx
  uid: 2a7f74cb-a8ad-11e9-ac0a-063ca044fe88
spec:
  excludedNamespaces: []
  excludedResources: []
  hooks:
    resources: []
  includeClusterResources: null
  includedNamespaces:
  - mysql-persistent-ns
  includedResources:
  - namespaces
  - persistentvolumes
  - persistentvolumeclaims
  - imagestreams
  - imagestreamtags
  - secrets
  - configmaps
  - pods
  - namespaces
  labelSelector:
    matchLabels:
      migration-included-stage-backup: d9002c80-a8ab-11e9-b6b2-0ac4ab43e1ee
  storageLocation: jwms3-8dpmm
  ttl: 720h0m0s
  volumeSnapshotLocations:
  - jwms3-5452b
```

# What is Rook?



KubeCon



CloudNativeCon

North America 2019

- Open Source (Apache 2.0)
- Cloud-Native Computing Foundation (CNCF)
  - Incubation Project
- Extends Kubernetes with Operators and custom types
- Framework for many storage providers and solutions



# Rook Operators



KubeCon



CloudNativeCon

North America 2019

- Implements the **Operator Pattern** for storage solutions
- Defines *desired state* for the storage resource
  - Storage Cluster, Pool, Object Store, etc.
- The Operator runs reconciliation loops
  - Watches for changes in desired state
  - Watches for changes in the cluster
  - Applies changes to the cluster to make it match desired

# Rook Architecture

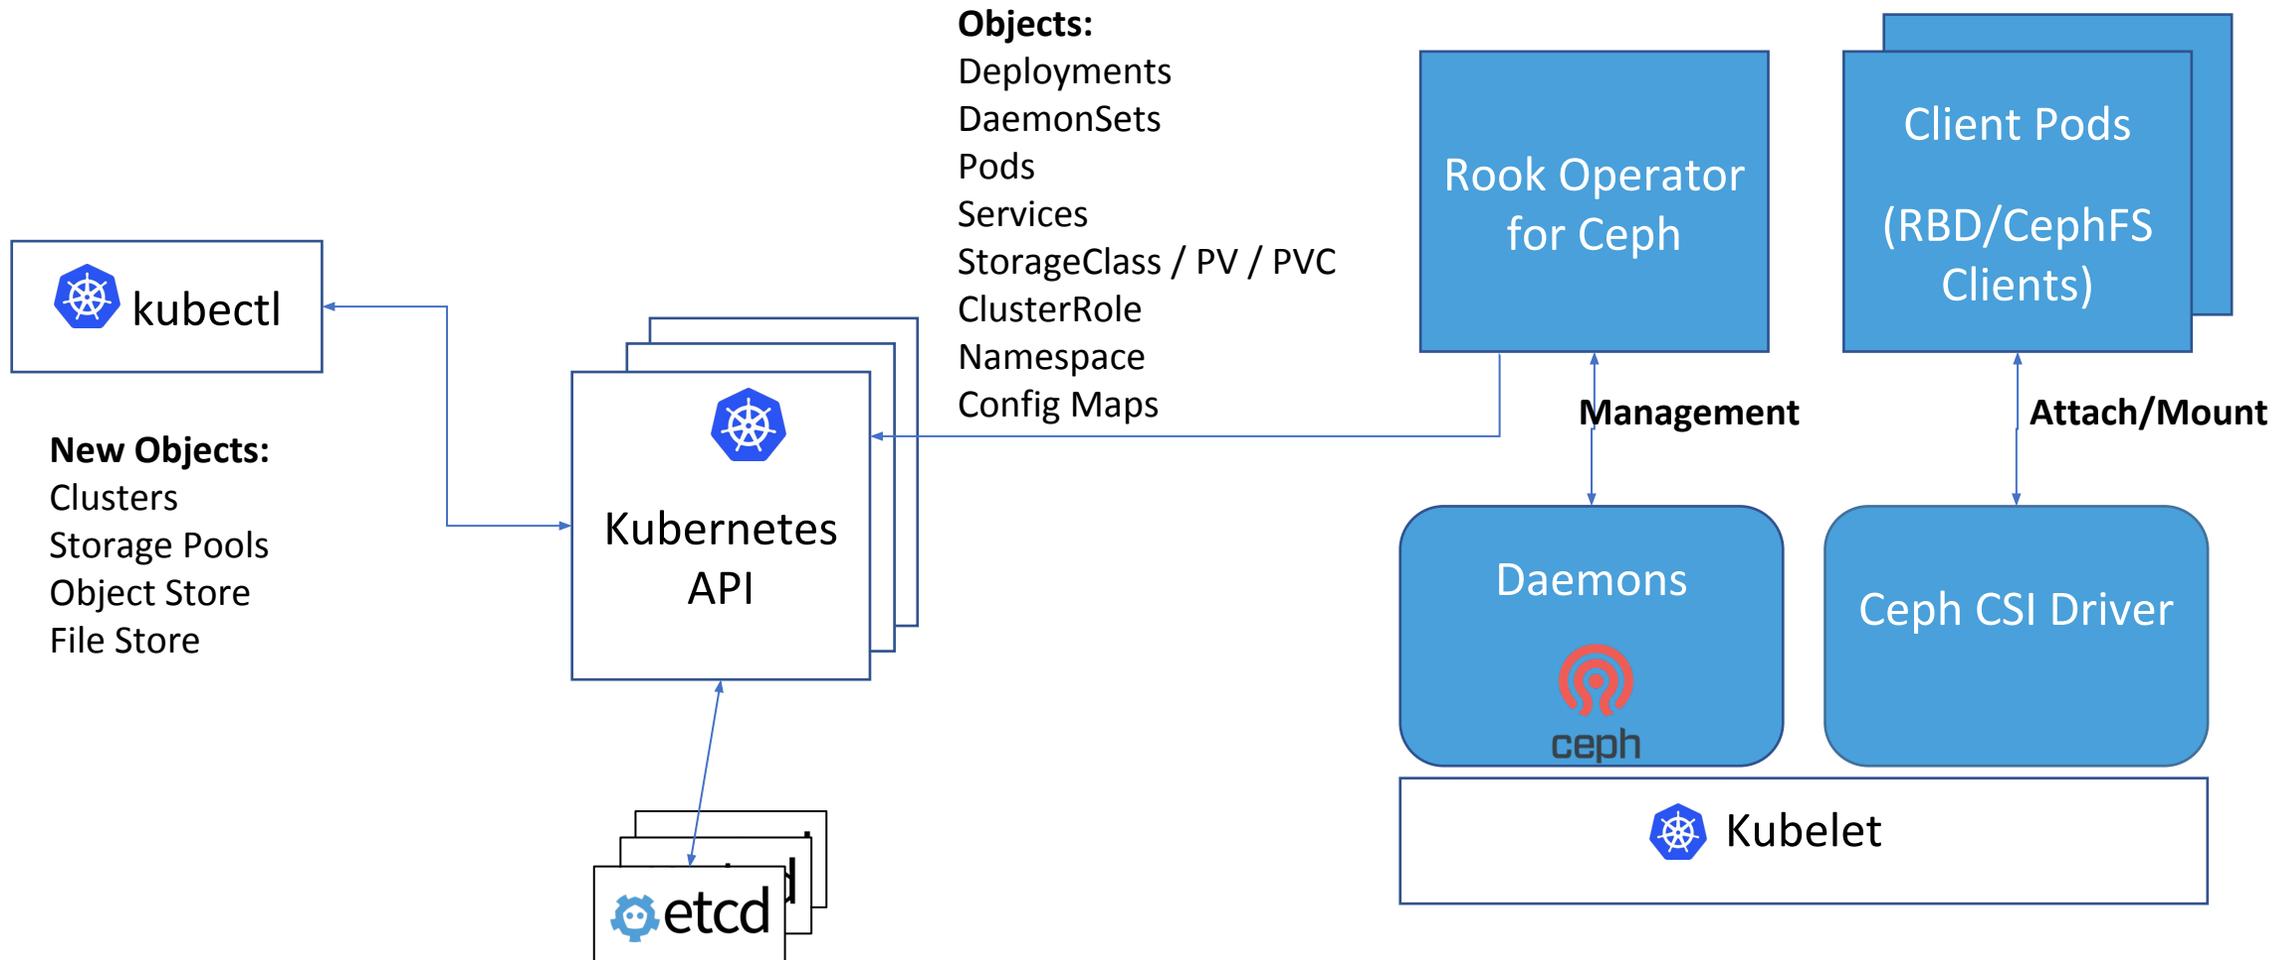


KubeCon



CloudNativeCon

North America 2019



# Ceph on Kubernetes with Rook



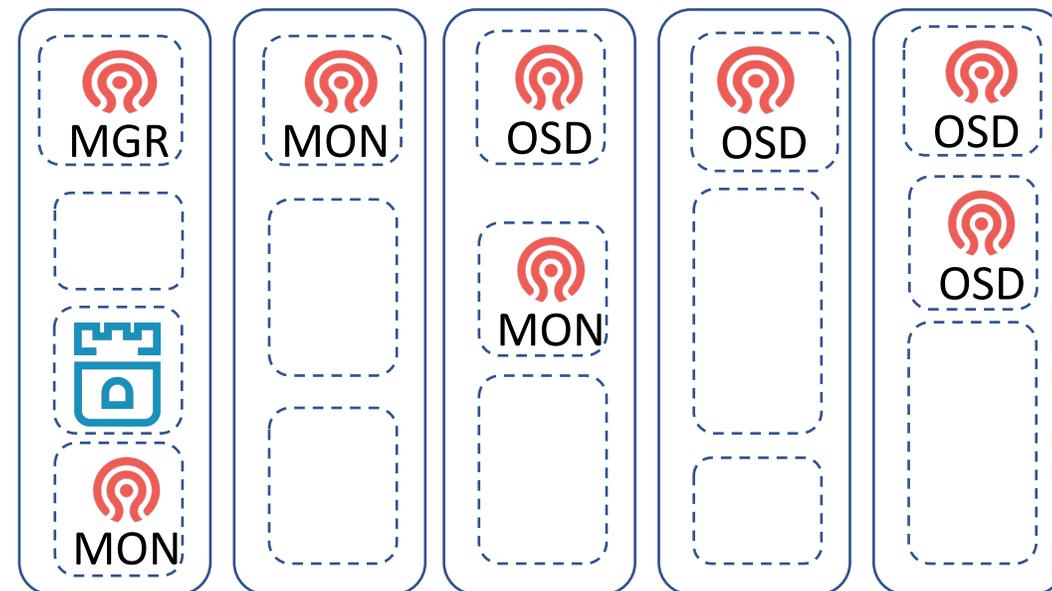
KubeCon



CloudNativeCon

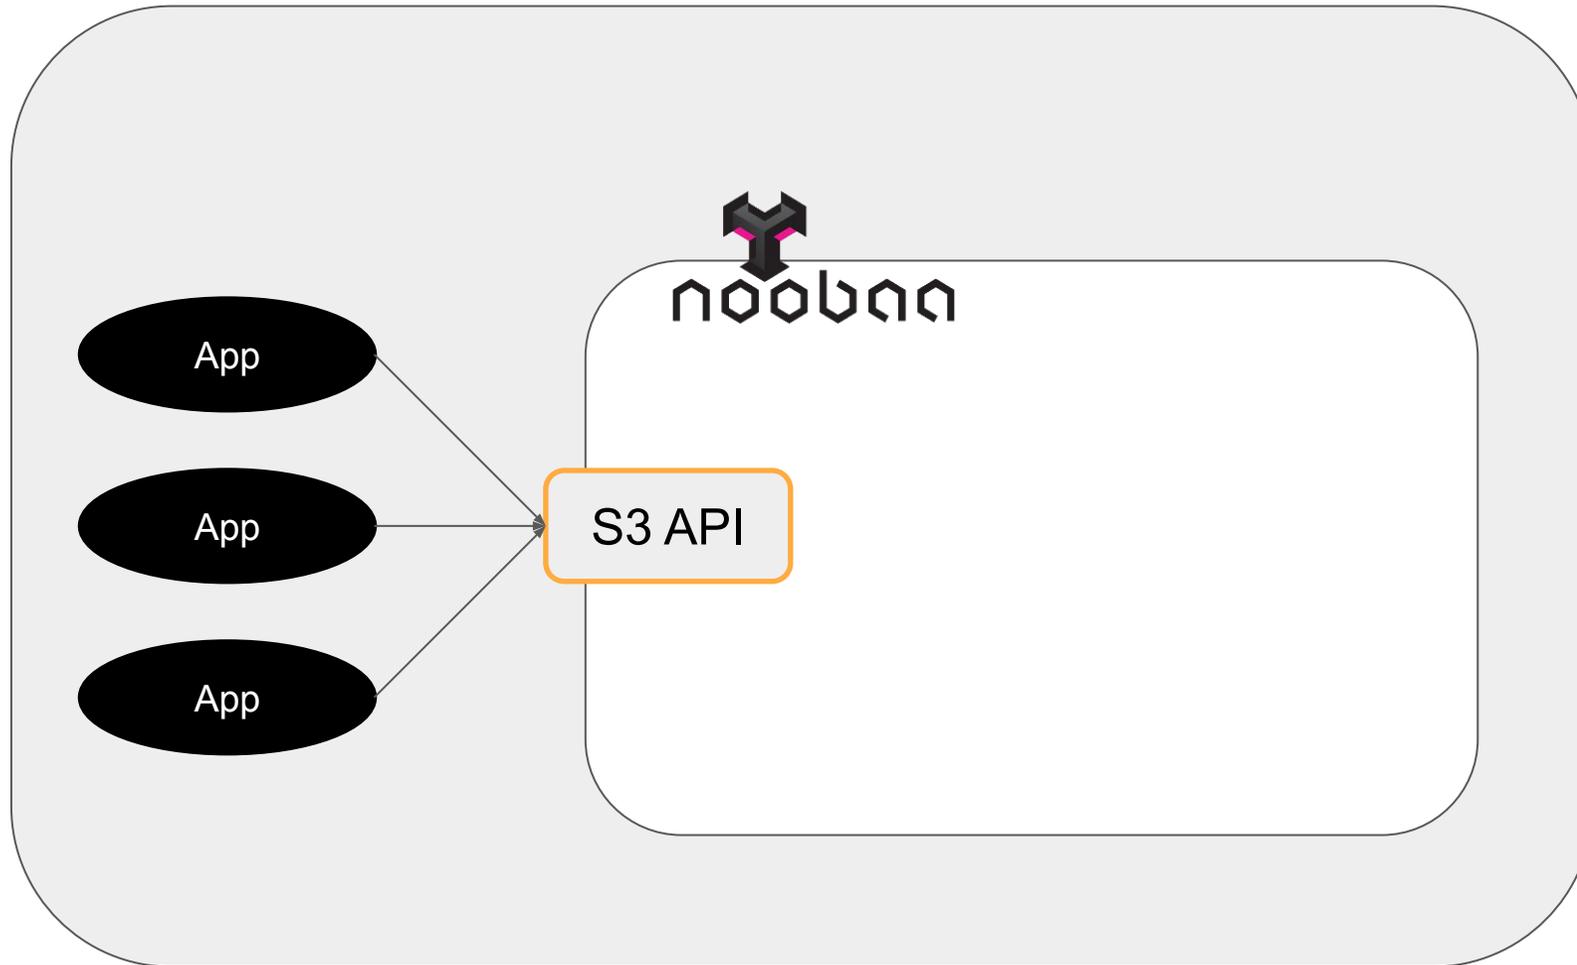
North America 2019

```
apiVersion: ceph.rook.io/v1
kind: Cluster
metadata:
  name: rook-ceph
spec:
  cephVersion:
    image: ceph/ceph:v14
  mon:
    count: 3
  network:
    hostNetwork: false
  storage:
    useAllNodes: true
```



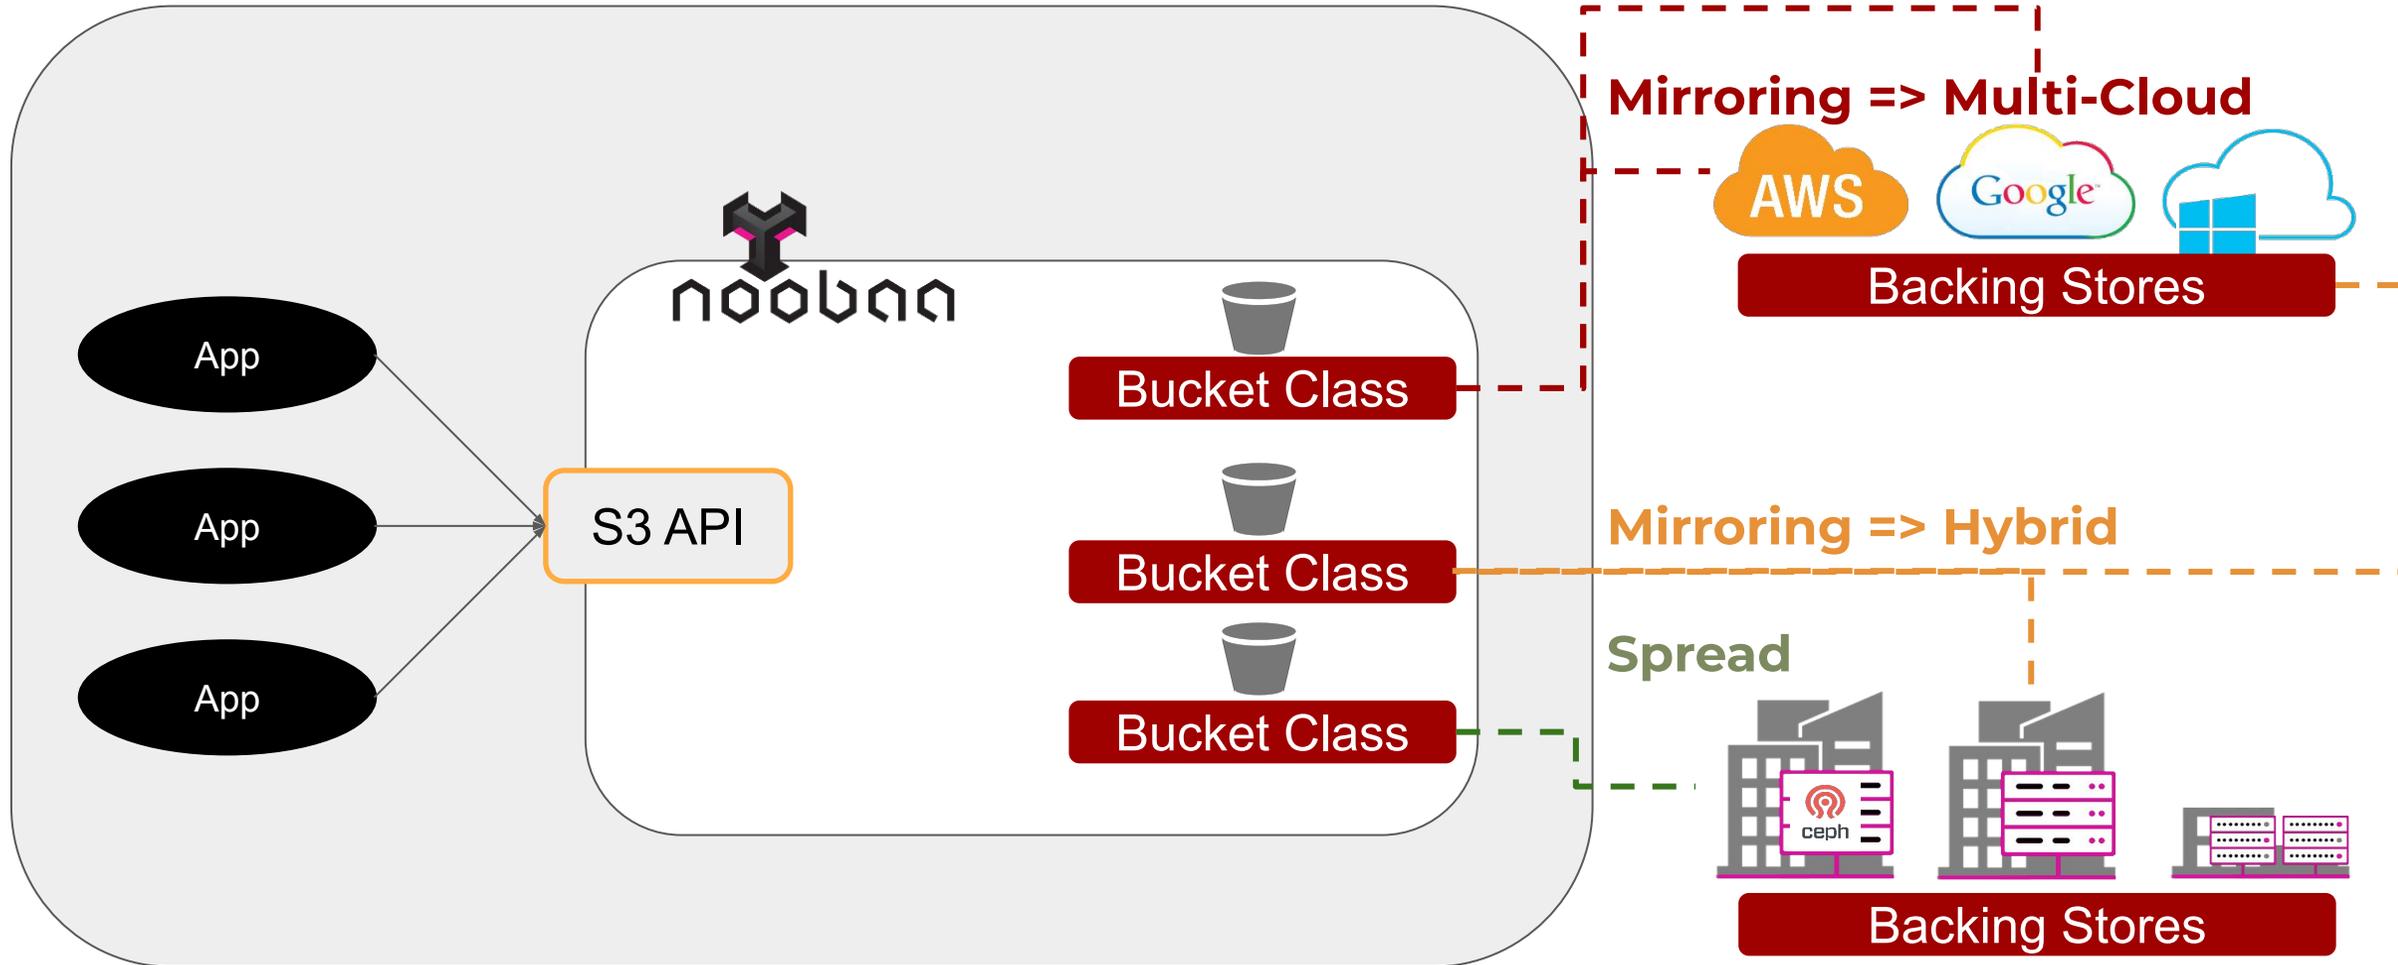
# What is Noobaa?

## DEPLOY AND MANAGE DATA SERVICES



# Bucket Class & Backing Stores

## DEPLOY AND MANAGE DATA SERVICES





**KubeCon**



**CloudNativeCon**

North America 2019

# K8s Backup and Recovery Example



# Backup Workflow

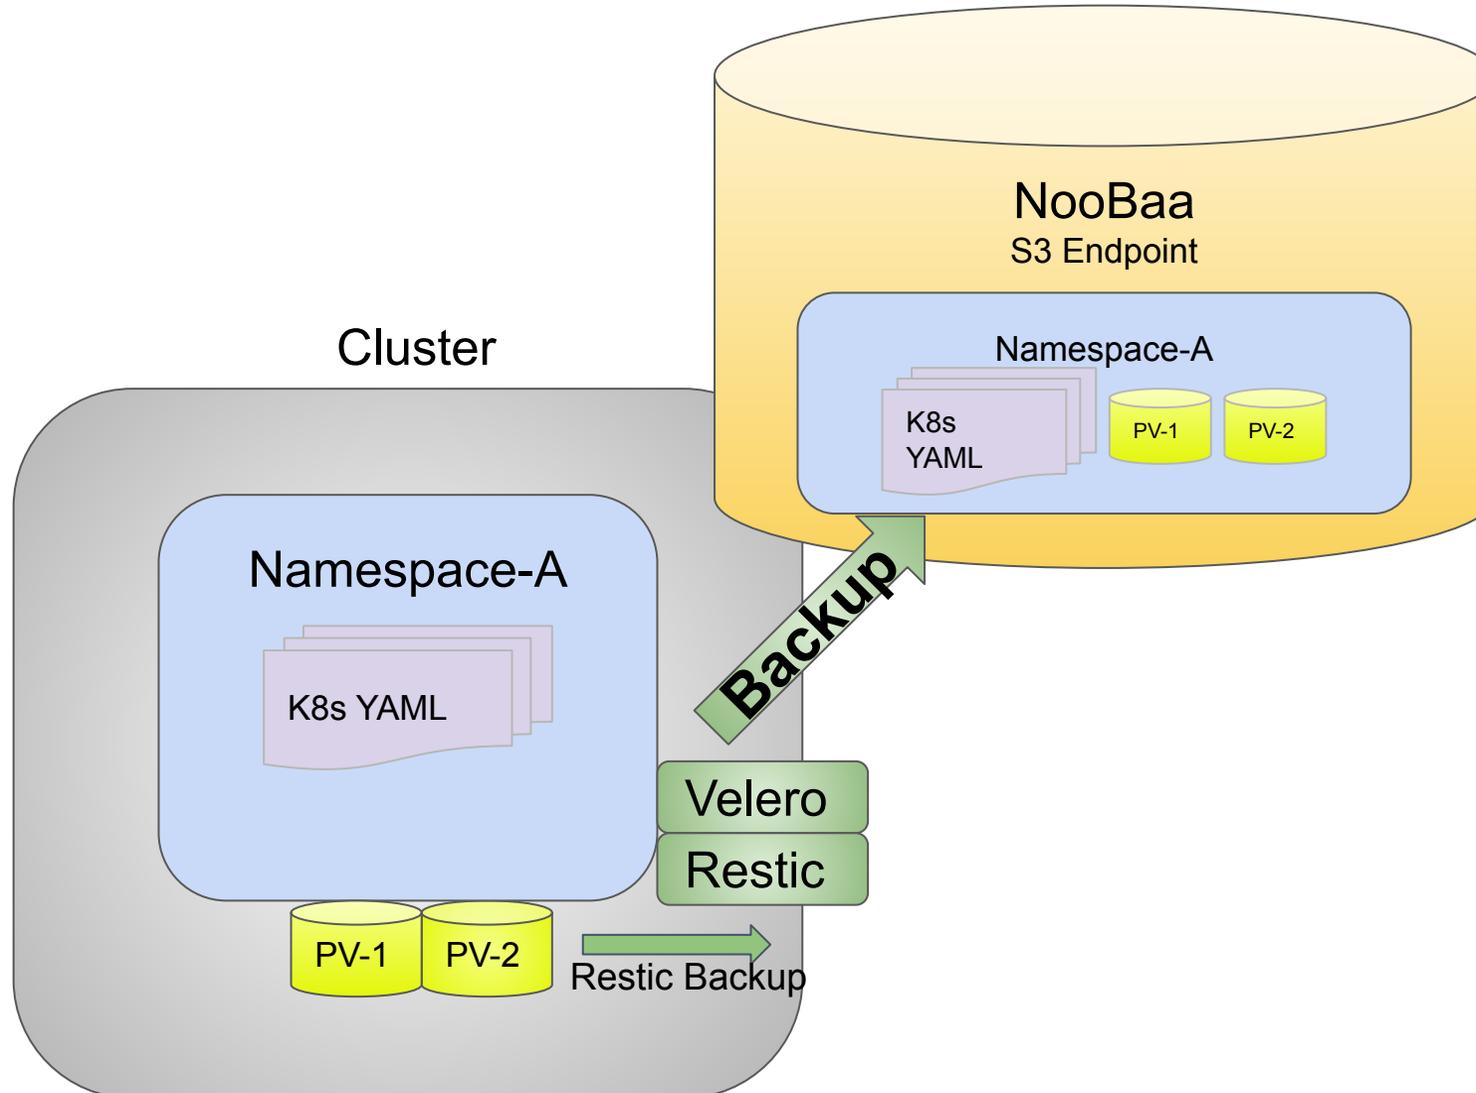


KubeCon



CloudNativeCon

North America 2019



# Disaster Scenario

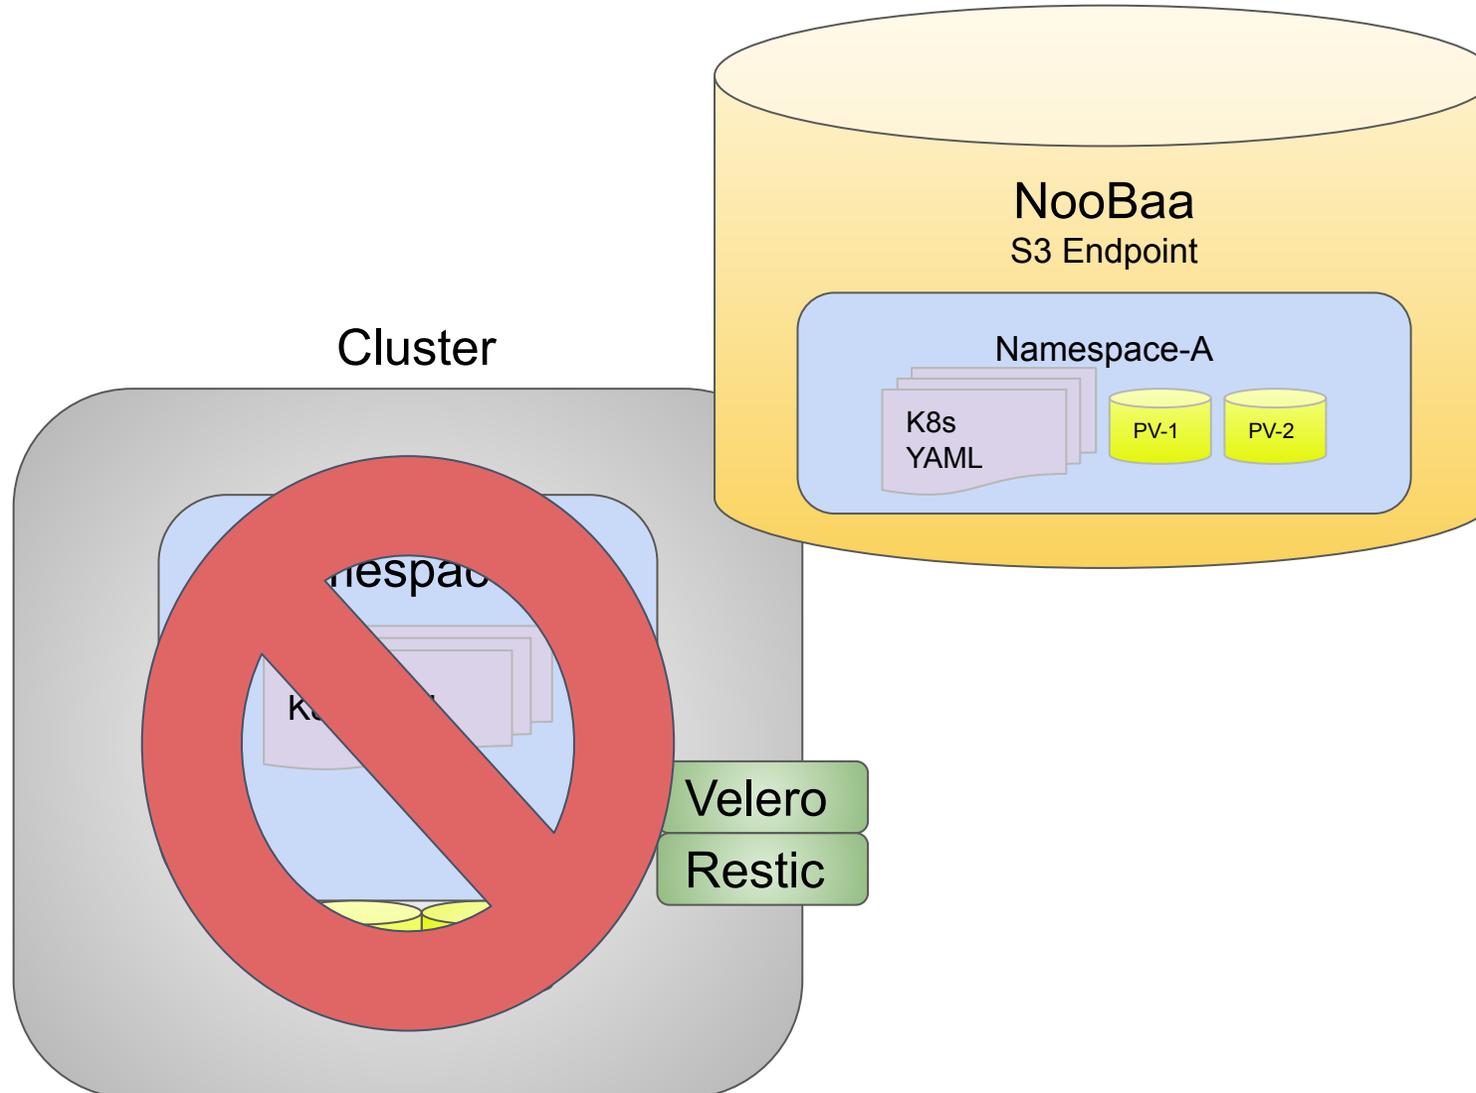


KubeCon



CloudNativeCon

North America 2019



# Disaster Scenario

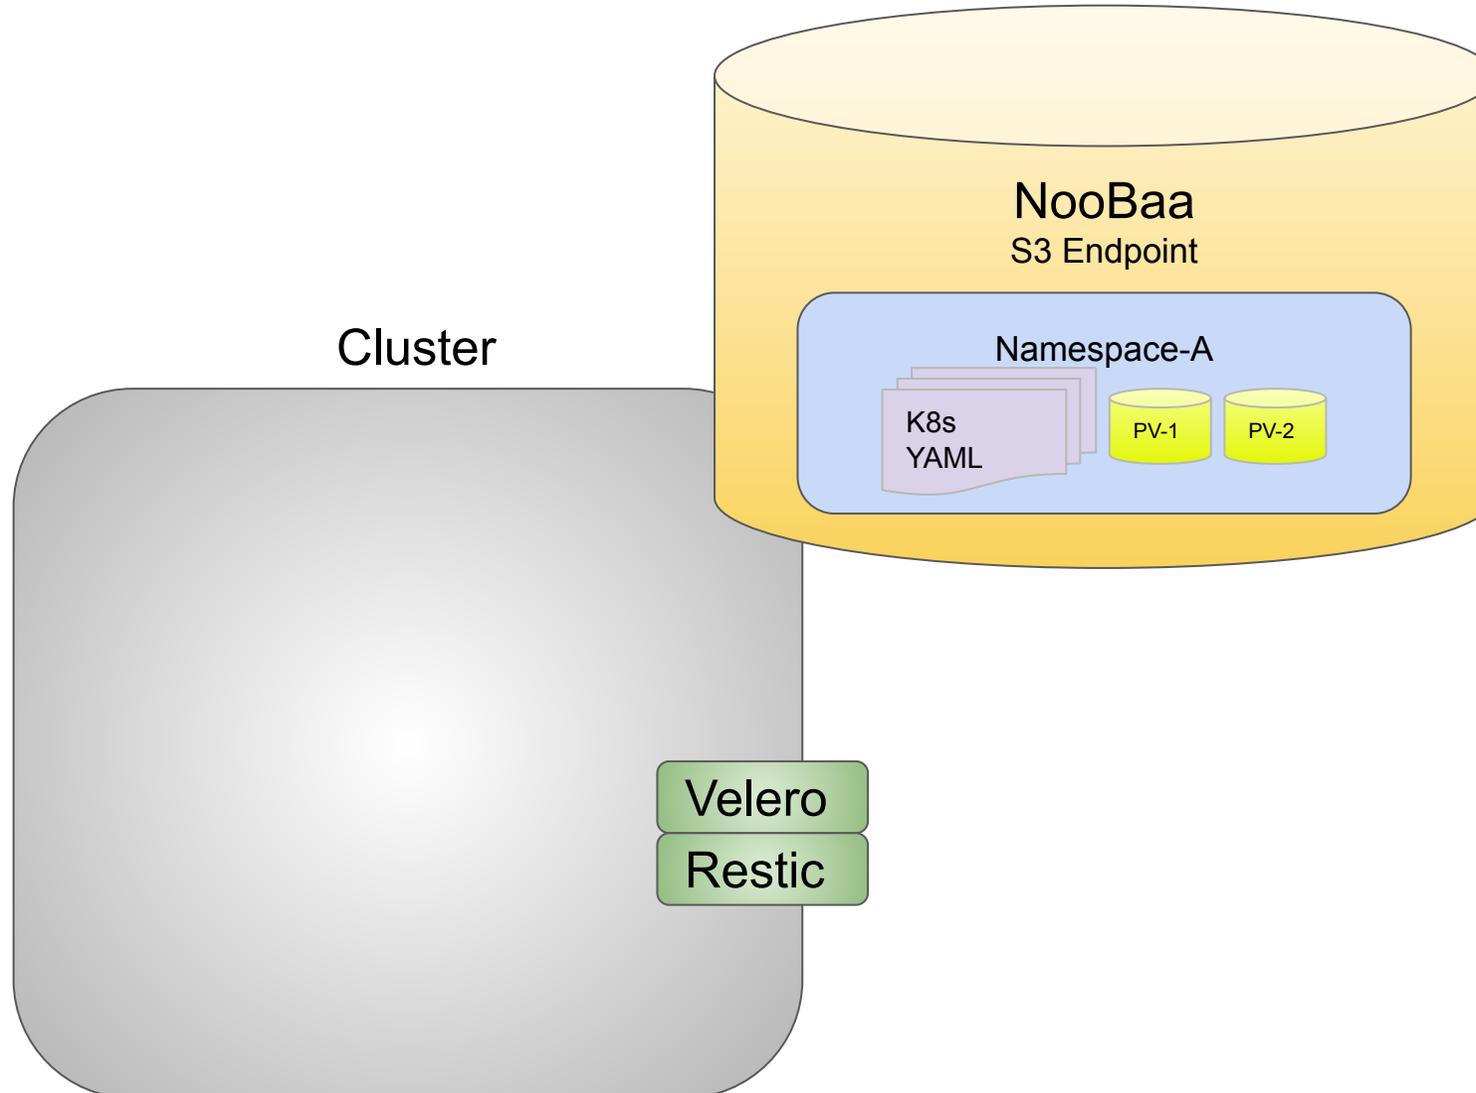


KubeCon



CloudNativeCon

North America 2019



# Restore Workflow

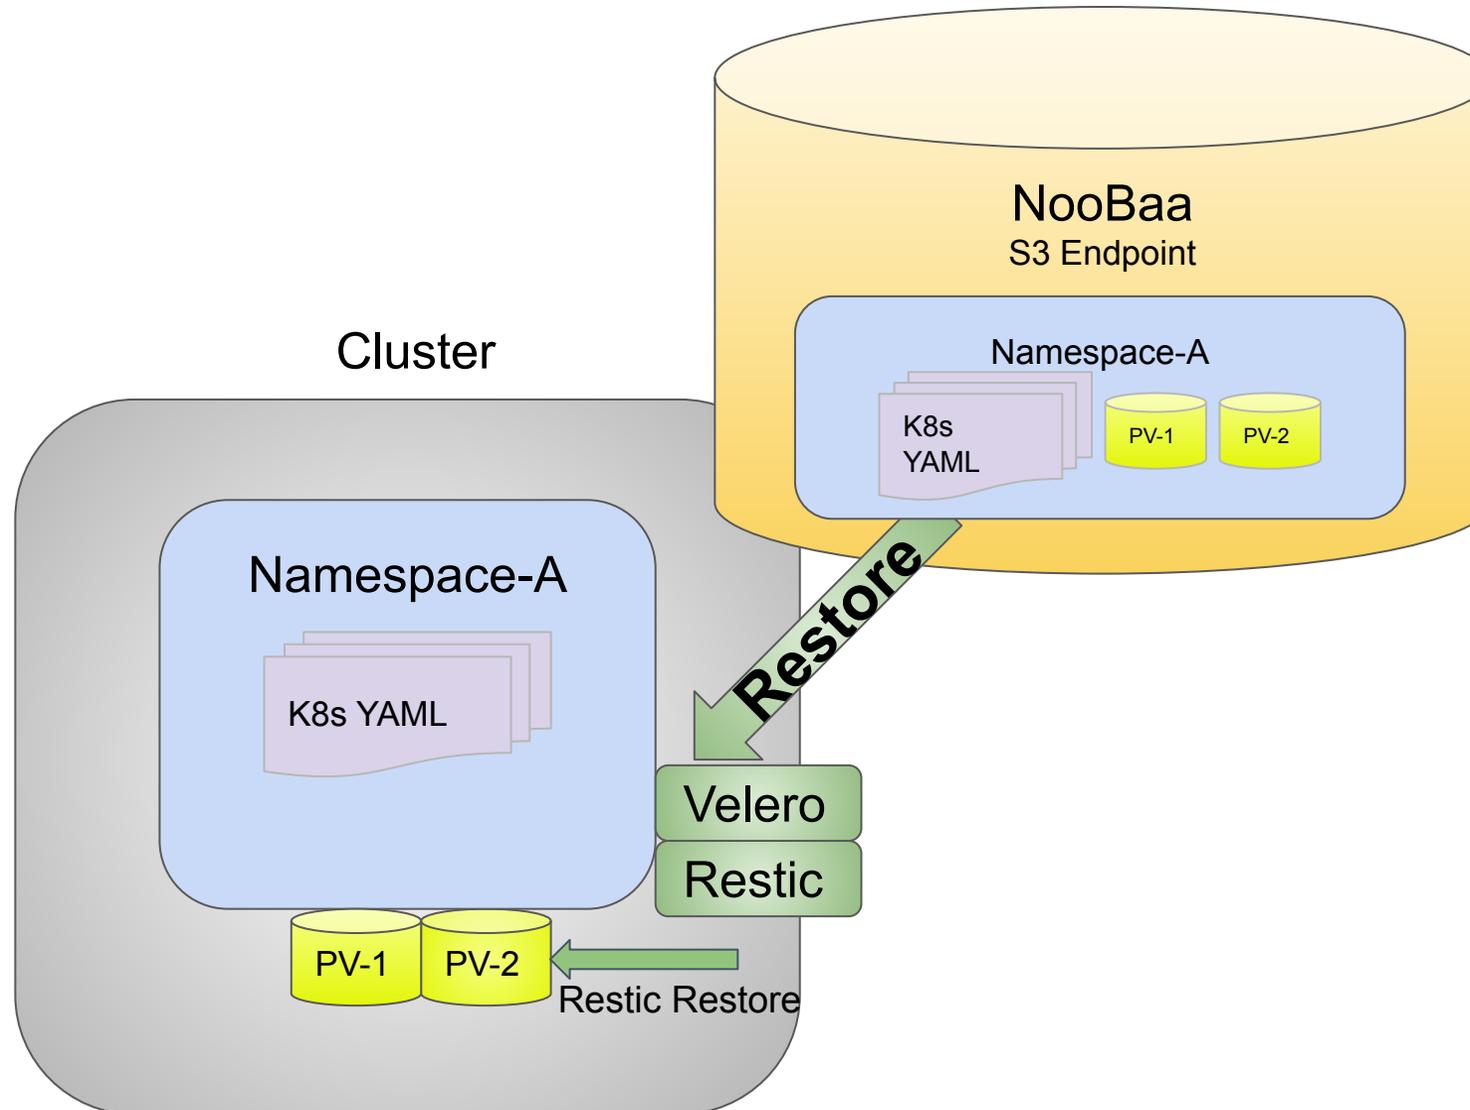


KubeCon



CloudNativeCon

North America 2019





**KubeCon**



**CloudNativeCon**

North America 2019

# Demo



# Resources/Links



KubeCon



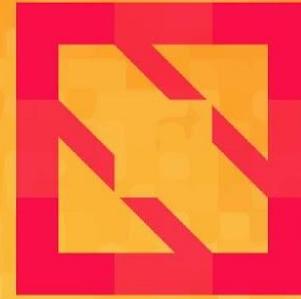
CloudNativeCon

North America 2019

- Velero
  - <https://velero.io/>
  - <https://github.com/vmware-tanzu/velero>
  - Community information
    - <https://velero.io/community/>
- Restic
  - <https://restic.net/>
  - <https://github.com/restic/restic/>
- Rook
  - <https://github.com/rook/rook>
  - Slack - <https://rook-io.slack.com>
  - Twitter - @rook\_io
  - Community Meetings - Every Tue at 9am Pacific Time
    - <https://zoom.us/j/392602367>
  - Forums: <https://groups.google.com/forum/#!forum/rook-dev>
- Noobaa
  - <https://github.com/noobaa/noobaa-core>
  - Slack - <http://noobaa.slack.com>



**KubeCon**



**CloudNativeCon**

**North America 2019**

**Thank You!**

