

Kubernetes Backup

Adfinis**sy**Group

Be smart. Think open source.

Kubernetes Backup

Approaches to backup and restore

ected Recovery

etcd Recovery

Recover from single node loss

```
$ set -a; source /etc/sysconfig/etcdctl; set +a; export ETCDCCTL_API=3
$ etcdctl --endpoints $ETCDCTL_ENDPOINT member list
7795cbbd4193ea0, started, 18fc757fba5f45f39f91c348dc694c4c, https://vm-adsy-caasp-master-02:2380, https://vm-adsy-caasp-master-02:2379
add5572adc80e7cb, started, eb9b944a76d643109c587cf722162763, https://vm-adsy-caasp-master-03:2380, https://vm-adsy-caasp-master-03:2379
2cf8e5fd3cd49497, started, 13f79f8428fd49888efd33037d51c1a5, https://vm-adsy-caasp-master-01:2380, https://vm-adsy-caasp-master-01:2379
```

etcd Recovery

Remove member

```
$ etcdctl --endpoints $ETCDCTL_ENDPOINT member remove add5572adc80e7cb
Member 9ecc84c4bf3492 removed from cluster add5572adc80e7cb
```

etcd Recovery

Add new member with old name

```
CTL_ENDPOINT member add eb9b944a76d643109c587cf722162763 --peer-urls=https://vm-adsy-caasp-master-03:2380
added to cluster b0d903008a9ba35a
```

```
643109c587cf722162763"
```

```
8fc757fa5f45f39f91c348dc694c4c=https://vm-adsy-caasp-master-02:2380,eb9b944a76d643109c587cf722162763=https://vm-adsy-caasp-master-03:2380,13f79f8428fd49888efd33037d51c1a5=https://vm-adsy-caasp-master-01:2380"
```

```
PEER_URLS="https://vm-adsy-caasp-master-03:2380"
```

```
ATE="existing"
```

Compare the output with `/etc/sysconfig/etcd`

etcd Recovery

Check that member was added to the cluster

```
$ etcdctl --endpoints $ETCDCTL_ENDPOINT member list
7795cdbd4193ea0, started, 18fc757fba5f45f39f91c348dc694c4c, https://vm-adsy-caasp-master-02:2380, https://vm-adsy-caasp-master-02:2379
2cf8e5fd3cd49497, started, 13f79f8428fd49888efd33037d51c1a5, https://vm-adsy-caasp-master-01:2380, https://vm-adsy-caasp-master-01:2379
9eccc84c4bf3492, unstarted, , https://vm-adsy-caasp-master-03:2380,
```

etcd Recovery

Start etcd daemon on failed node

```
$ systemctl start etcd.service
```

etcd Recovery

Verify that node was added to the cluster

```
$ etcdctl --endpoints $ETCDCTL_ENDPOINT member list
7795cdbd4193ea0, started, 18fc757fba5f45f39f91c348dc694c4c, https://vm-adsy-caasp-master-02:2380, https://vm-adsy-caasp-master-02:2379
2cf8e5fd3cd49497, started, 13f79f8428fd49888efd33037d51c1a5, https://vm-adsy-caasp-master-01:2380, https://vm-adsy-caasp-master-01:2379
9ecc84c4bf3492, started, eb9b944a76d643109c587cf722162763, https://vm-adsy-caasp-master-03:2380, https://vm-adsy-caasp-master-03:2379
```


etcd Backup

Low-level backup of your Kubernetes cluster

etcd Backup

Snapshot etcd v3 keyspace

```
$ set -a; source /etc/sysconfig/etcdctl; set +a; export ETCDCTL_API=3  
$ etcdctl --endpoints $ETCDCTL_ENDPOINT snapshot save snapshot.db  
Snapshot saved at snapshot.db
```

etcd Restore

Perform restore command on every master

```
$ source /etc/sysconfig/etcd
$ ETCDCTL_API=3 etcdctl snapshot restore snapshot.db \
  --name $ETCD_NAME \
  --initial-cluster $ETCD_INITIAL_CLUSTER \
  --initial-cluster-token $ETCD_INITIAL_CLUSTER_TOKEN \
  --initial-advertise-peer-urls $ETCD_INITIAL_ADVERTISE_PEER_URLS
2018-11-28 23:56:37.918068 I | pkg/netutil: resolving vm-adsy-caasp-master-03:2380 to [::1]:2380
2018-11-28 23:56:37.918388 I | pkg/netutil: resolving vm-adsy-caasp-master-03:2380 to [::1]:2380
2018-11-28 23:56:37.955241 I | etcdserver/membership: added member 7795cdbd4193ea0 [https://vm-adsy-caasp-master-02:2380] to cluster b0d903008a9ba35a
2018-11-28 23:56:37.955346 I | etcdserver/membership: added member 2cf8e5fd3cd49497 [https://vm-adsy-caasp-master-01:2380] to cluster b0d903008a9ba35a
2018-11-28 23:56:37.955386 I | etcdserver/membership: added member add5572adc80e7cb [https://vm-adsy-caasp-master-03:2380] to cluster b0d903008a9ba35a
$ mv $ETCD_NAME.etcd/member /var/lib/etcd/
```

etcd Restore

Start etcd on every master

Start etcd on every master

```
$ systemctl start etcd.service
```

etcd Restore

Verify cluster health

```
$ set -a; source /etc/sysconfig/etcdctl; set +a; export ETCDCTL_API=3  
$ etcdctl --endpoints https://vm-adsy-caasp-master-01:2379,https://vm-adsy-caasp-master-02:2379,https://vm-adsy-caasp-master-03:2379 endpoint health  
https://vm-adsy-caasp-master-02:2379 is healthy: successfully committed proposal: took = 3.085025ms  
https://vm-adsy-caasp-master-01:2379 is healthy: successfully committed proposal: took = 3.494222ms  
https://vm-adsy-caasp-master-03:2379 is healthy: successfully committed proposal: took = 8.480983ms
```

Heptio Ark

Backup tool for Kubernetes objects

- Server component in the cluster
- CLI for controlling backups

Heptio Ark

Multiple storage backends supported

- AWS S3 (and compatible)
- Azure Blob Storage
- Google Cloud Storage

Cluster Backup

By default ark performs full cluster backups

```
$ ark backup create full-cluster  
Backup request "full-cluster" submitted successfully.  
Run `ark backup describe full-cluster` or `ark backup logs full-cluster` for more details.
```

Namespace Backup

Backups can be limited to a namespace

```
$ ark backup create demo --include-namespaces demo  
Backup request "demo" submitted successfully.  
  
Run `ark backup describe demo` or `ark backup logs demo` for more details.
```


Label based Backup

Using a label selector to select the objects

```
$ ark backup create lampp --selector app=lampp  
Backup request "lampp" submitted successfully.  
Run 'ark backup describe lampp' or 'ark backup logs lampp' for more details.
```

Scheduled backups

Scheduled backups are possible

```
$ ark schedule create hourly --schedule="0 * * * *"  
Schedule "hourly" created successfully.
```

Restore

Also a restore of the data is possible!

Also a restore of the data is possible:

```
$ ark restore create --from-backup demo
Restore request "demo-20181122133747" submitted successfully.
Run `ark restore describe demo-20181122133747` or `ark restore logs demo-20181122133747` for more details.
```

Restore

- Objects already present are skipped
- PV/PVC restores require special consideration
- Data in storage backend has priority
- Data from S3 is synced to Kubernetes Objects

Download

Backups can be downloaded for manual intervention

```
$ ark backup download demo  
Backup demo has been successfully downloaded to /home/example/demo-data.tar.gz
```

Restore PVC to use existing PV

To make a PVC use an existing PV the value for `claimRef.uid` needs to be removed:

```
$ kubectl patch pv example --type json \
-p [{"op": "remove", "path": "/spec/claimRef/uid"}]
pv/example patched
```

The PV will change to status `Available` and can be bound to the PVC defined in `claimRef`.

Attribution / License

- Slide Skeleton [https://de.wikipedia.org/wiki/Skeleton_\(Programmierung\)](https://de.wikipedia.org/wiki/Skeleton_(Programmierung))
- Slides Adfinis SyGroup AG, 2017, Attribution-NonCommercial 2.0 (CC BY-NC 2.0)

Feel Free to Contact Us

www.adfinis-sygroup.ch

[Tech Blog](#)

[GitHub](#)

info@adfinis-sygroup.ch

[Twitter](#)