

Dell EMC PowerProtect Data Manager: Protect and Restore a Kubernetes Master Node etcd database

Abstract

This white paper describes how to protect a Kubernetes etcd database, stage it to Dell EMC PowerProtect Appliance using Dell EMC PowerProtect Data Manager, and restore the database.

September 2021

Revisions

Date	Description
October 2020	Initial release
September 2021	Version 2

Acknowledgments

Author: Debyeet Bagchi

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

This document may contain certain words that are not consistent with Dell's current language guidelines. Dell plans to update the document over subsequent future releases to revise these words accordingly.

This document may contain language from third party content that is not under Dell's control and is not consistent with Dell's current guidelines for Dell's own content. When such third-party content is updated by the relevant third parties, this document will be revised accordingly.

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [9/23/2021] [Technical White Paper] H18571.1

Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents	3
Executive summary.....	4
1 Overview.....	5
2 Configuring etcdctl utility	6
3 Back up the etcd database.....	7
4 Restore etcd database after a disaster	8
A Technical support and resources	9
A.1 Related resources.....	9

Executive summary

Kubernetes is a portable, extensible, open-source platform for managing containerized workloads and services, and it facilitates both declarative configuration and automation. It has a large, rapidly growing ecosystem, which continues to increase the importance of protecting Kubernetes clusters and resources.

Dell EMC PowerProtect Data Manager solves the requirement to protect Kubernetes clusters, pods, persistent volume claims, namespaces, and other resources. This document discusses Kubernetes resources and the etcd database that resides in the primary node or nodes, and describes how to protect and restore the etcd database.

1 Overview

The etcd database is an important resource because all cluster data is stored there, including all Kubernetes objects such as cluster-scoped resources, and deployment and pod information. If there is a disaster that causes the primary node to go down, the Kubernetes cluster suffers unavailability, and no operations can be scheduled or managed. The DevOps team must have a procedure to follow regularly to protect the etcd database. The team also requires a procedure to restore the database from the protection system to bring the system back online and work to the last-known backup timestamp.

You can deploy the etcd database as a pod in the primary node and also deploy it externally to enable resiliency and security. Also, Kubernetes deployments can be of different types such as bare-metal Linux nodes, or they can be deployed as a virtual machine.

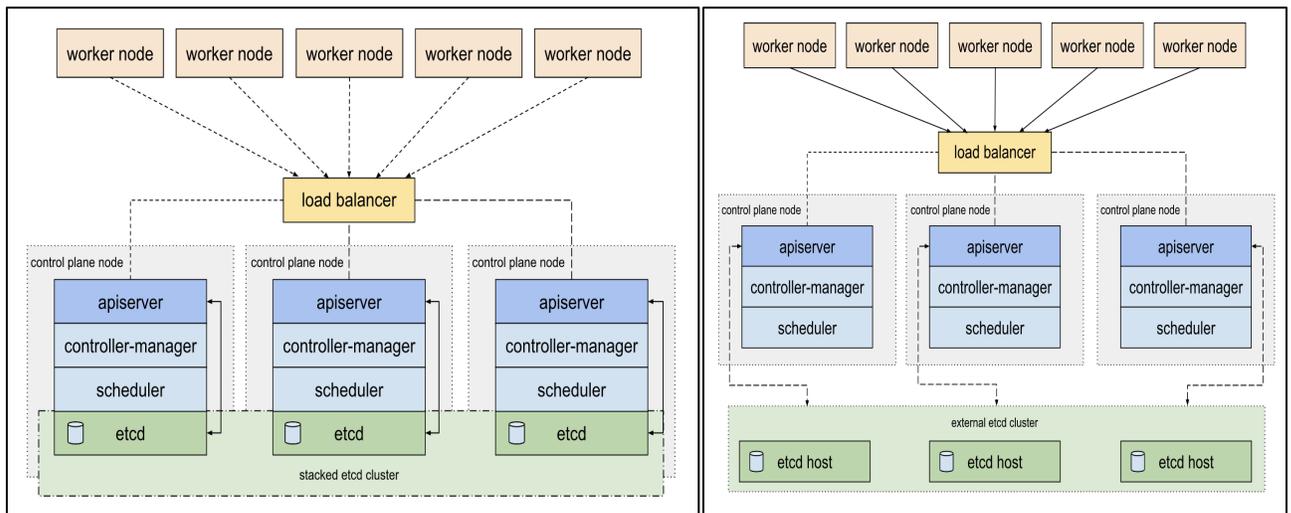


Figure 1 Deployment topologies for etcd

2 Configuring the etcdctl utility

etcdctl is a command line client for etcd. It is required to take an etcd snapshot dump at a specific time. The v3 API (ETCDCTL_API=3) of etcdctl will be used while running the snapshot command. The easiest way to get the etcdctl utility is to take it from the etcd binary. The following steps ensure that the etcdctl v3 API is installed properly.

Note: Run the following commands as superuser.

1

Download the etcd binary and place it in the /tmp directory –
`curl -L https://storage.googleapis.com/etcd/v3.4.13/etcd-v3.4.13-linux-amd64.tar.gz -o /tmp/etcd-v3.4.13-linux-amd64.tar.gz`

2

Untar the binary and change the directory:

- a. `cd /tmp`
- b. `tar -xvf etcd-v3.4.13-linux-amd64.tar.gz`
- c. `cd etcd-v3.4.13-linux-amd64`

3

Grant the utility global permissions –

- a. `chmod +x ./etcdctl`
- b. `mv ./etcdctl /usr/local/bin/etcdctl`

3 Back up the etcd database

Use the following steps to take a snapshot dump of the etcd database:

1. Create a directory (mkdir) on the primary node, for example:

```
mkdir "k8backupforDELLDR"
```

2. List all pods in the system:

```
kubectl get pods -A -o wide
```

3. Describe the etcd pod to get the etcd server IP address endpoint:

```
kubectl describe <etcd-pod> -n kube-system
```

4. Copy all certificates and manifests to the **k8backupforDELLDR** directory:

```
cp -r /etc/kubernetes /k8backupforDELLDR
```

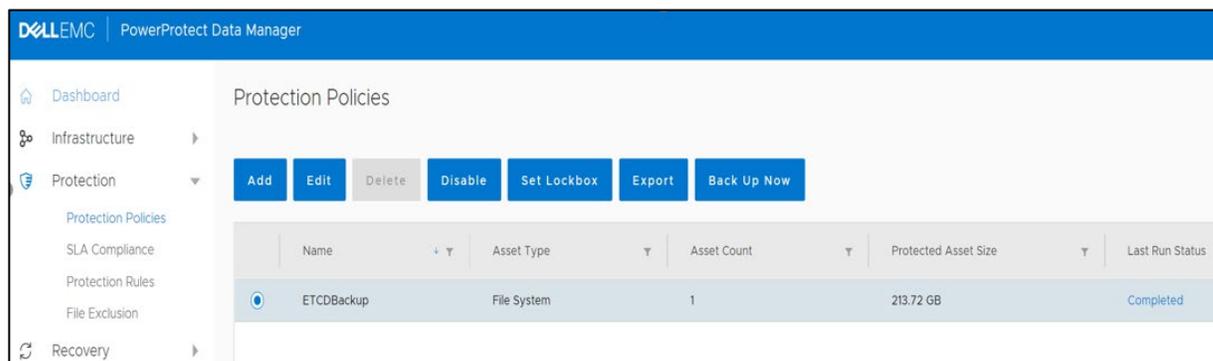
5. Run the following command to take an etcd snapshot:

```
ETCDCTL_API=3 etcdctl --endpoints=https://<etcd-server-IP>:2379 --cacert=/etc/kubernetes/pki/etcd/ca.crt --cert=/etc/kubernetes/pki/etcd/server.crt --key=/etc/kubernetes/pki/etcd/server.key snapshot save /k8backupforDELLDR/etcd-snapshot.db
```

6. View the status and details of the snapshot:

```
ETCDCTL_API=3 etcdctl --write-out=table snapshot status /k8backupforDELLDR/etcd-snapshot.db
```

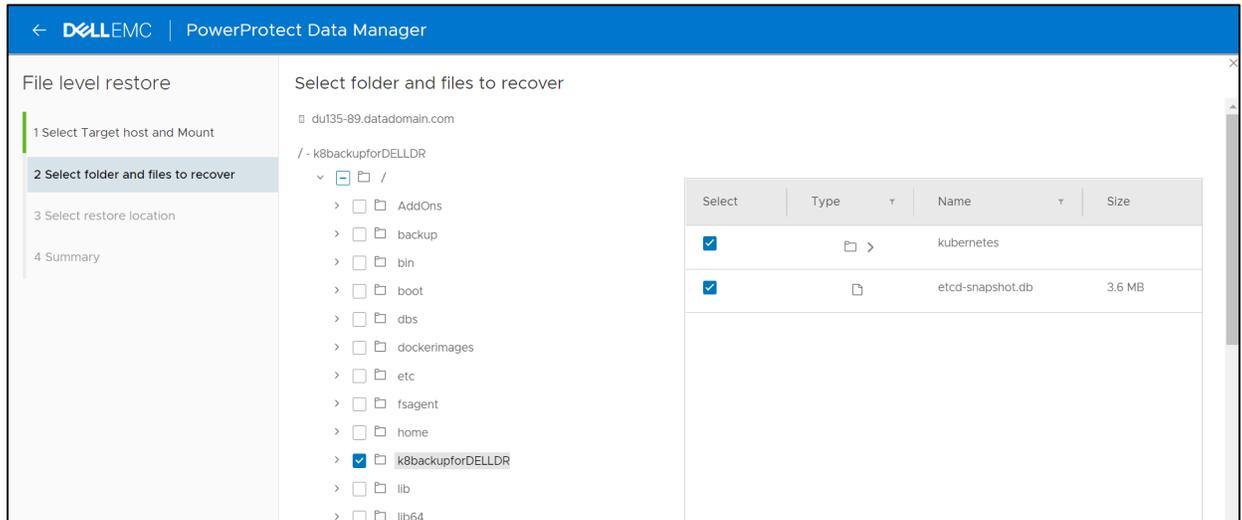
Note: The certificates and the etcd snapshot are available in the **k8backupforDELLDR** directory. Follow the previous steps to back up the volume containing the directory using the PowerProtect Data Manager File System Agent for Linux or VM Image protection. View the volume on which the k8backupforDELLDR directory is mounted, and run a file system backup on a specific volume asset or a VM image backup on a VM asset.



4 Restore the etcd database after a disaster

Use the following steps to restore the backed-up directory to the new primary node after a disaster:

1. Use PowerProtect Data Manager to perform a file-level recovery on the VM Image backup or the file system backup, and restore the contents of the specific directory to the new primary node.



2. After the k8backupforDELLDR directory is restored, run the following etcdctl command to restore the etcd database to the point when the backup was taken:

```
ETCDCTL_API=3 etcdctl --endpoints=https://<etcd-server-IP>:2379 --
cacert=/etc/kubernetes/pki/etcd/ca.crt \
--name=master \
--cert=/etc/kubernetes/pki/etcd/server.crt --
key=/etc/kubernetes/pki/etcd/server.key \
--data-dir /var/lib/etcd-from-backup \
--initial-cluster=master=https://<etcd-server-IP>:2380 \
--initial-cluster-token etcd-cluster-1 \
--initial-advertise-peer-urls=https://<etcd-server-IP>:2380 \
snapshot restore /k8backupforDELLDR/etcd-snapshot.db
```

3. Modify the file **/etc/kubernetes/manifests/etcd.yaml**:

```
--data-dir=/var/lib/etcd-from-backup
Add entry : --initial-cluster-token=etcd-cluster-1
Change volumeMounts: -mountPath to /var/lib/etcd-from-backup
Change volumes: -hostPath to /var/lib/etcd-from-backup
```

4. Refresh the **serviceaccount** tokens for the application, default, and kube-system namespaces by removing the existing tokens in the serviceaccount config:

```
kubectl edit serviceaccount <serviceaccount-name> -n <namespace-name>
Verify new tokens are generated : kubectl get secrets -n <namespace-name>
```

A Technical support and resources

[Dell.com/support](https://www.dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage and data protection technical white papers and videos](#) provide expertise that helps to ensure customer success with Dell EMC storage and data protection products.

A.1 Related resources

- <https://superuser.openstack.org/articles/a-guide-to-kubernetes-etcd-all-you-need-to-know-to-set-up-etcd-clusters/>
- <https://github.com/etcd-io/etcd/tree/master/etcdctl>
- <https://kubernetes.io/docs/tasks/administer-cluster/configure-upgrade-etcd/#backing-up-an-etcd-cluster>