

What Are Kubernetes Persistent Volumes

Description

In the early days of computers, people used the term “volume” a lot. This was the term referred to the computer hard drive. It is where all your data is stored. Let me set the context by explaining a simple example. While working on Spring boot with H2 database (in-memory database) at your local IDE level, whenever we perform CRUD operations against the database, the data will be wiped out when you restart your spring boot application. Obviously in production environment/scenario we will use others databases like Oracle/Mysql/Postgresql and we will ensure that the system is a highly available.

In Kubernetes terms, lets say you have a MySQL pod where the data is stored, the data will be wiped out when you restart the pod. In the kubernetes architecture, you have to take care of your applications data store requirements. In short, the data storage should not be dependent on the POD lifecycle nor the cluster state in the event of crash.

Data storage requirements can be anything from database to storing files in a directory.

To sum up in a single line, Kubernetes Persistent Volume is a component in Kubernetes architecture, which is just like a cluster resource which helps with the storage requirements of the containers in the pod. Just imagine if Volumes wouldn't exist, the data which containers created would be lost if the containers died. This is the reason why there was a need for a separate component in the kubernetes cluster i.e. Persistent Volumes. Now that you have some idea about Persistent volumes, lets explore each aspect.

Characteristics of Volumes

- Since the data is an essential component that your applications relies on, volumes should be highly available in nature.
- Even if the cluster dies or crashes completely the Persistent Volumes needs to stay intact.
- Volumes should be available to all the nodes.

Persistent Volumes (PV)

Persistent volumes are just an abstract component or an interface, the actual storage will be taken care by various cloud providers such as AWS, Azure etc. Actual storage can be actual physical storage from local Hard disk drive of the cluster node or NFS servers outside the cluster. Persistent volumes gets created by using kubernetes yml file. In this file we can specify various properties like the storage provider, storage volume etc

You can take a look at the [list of Persistent Volumes](#) that kubernetes supports. Lets look at one example of Persistent volume yaml file configuration.

I hope you liked the post. Happy Learning 🙌

Category

1. Cloud
2. Devops
3. Kubernetes

Date Created

February 14, 2022

Author

kk-ravi144gmail-com

default watermark