

Storage

Storage in Kubernetes (KCNA Relevant)

In Kubernetes, storage plays a crucial role in providing persistent storage solutions for applications running in containers. Unlike containers, which are ephemeral and can be destroyed and recreated, storage needs to persist across Pod restarts. Kubernetes provides a powerful system for managing and abstracting storage resources, allowing you to manage stateful applications effectively.

Here are the key **Storage** topics relevant to **Kubernetes** and the **KCNA exam**:

1. Persistent Storage in Kubernetes:

- **Persistent Volumes (PVs):**
 - A **Persistent Volume (PV)** is a piece of storage in the Kubernetes cluster that has been provisioned by an administrator or dynamically by a storage class.
 - **PV** is an abstraction of storage resources (e.g., network-attached storage, cloud storage).
 - PVs are independent of the lifecycle of Pods and are not tied to any specific Pod.
- **Persistent Volume Claims (PVCs):**
 - A **Persistent Volume Claim (PVC)** is a request for storage by a user or a Pod. It specifies the amount of storage and the access modes.
 - A PVC is used to claim a PV, and Kubernetes binds the claim to an available PV.
- **Storage Classes:**
 - A **StorageClass** defines a type of storage and how it should be provisioned (e.g., speed, replication, type of disk).
 - It allows dynamic provisioning of storage resources when a PVC is created.
 - **Examples:** `standard`, `fast`, `ssd`, or cloud-specific classes (e.g., `aws-efs`).
- **Dynamic Provisioning:**
 - When a PVC is created, Kubernetes can automatically provision the corresponding PV based on the storage class.

- If no PV exists that satisfies the PVC's request, Kubernetes dynamically provisions a PV.
 - **Access Modes:**
 - **ReadWriteOnce (RWO):** The volume can be mounted as read-write by a single node.
 - **ReadOnlyMany (ROX):** The volume can be mounted as read-only by many nodes.
 - **ReadWriteMany (RWX):** The volume can be mounted as read-write by many nodes.
-

2. Types of Storage in Kubernetes:

- **Ephemeral Storage:**
 - Temporary storage that is deleted when the Pod is terminated or deleted (e.g., emptyDir volumes).
 - Used for temporary data during the lifecycle of a Pod.
 - **Persistent Storage:**
 - Storage that persists beyond the Pod lifecycle (e.g., databases, logs).
 - This includes PVs and PVCs, which provide durable storage that survives Pod restarts.
 - **Volume Plugins (CSI):**
 - Kubernetes uses the **Container Storage Interface (CSI)** to support multiple storage providers (e.g., AWS EBS, GCE PD, NFS, GlusterFS, Ceph, etc.).
 - **CSI drivers** enable Kubernetes to work with external storage systems that support the CSI specification.
 - **Common Volume Types:**
 - **emptyDir:** A temporary directory that is created when a Pod starts and is deleted when the Pod terminates.
 - **hostPath:** Mounts a file or directory from the host node's filesystem into a Pod.
 - **nfs:** Mounts an NFS share into a Pod.
 - **awsElasticBlockStore (EBS):** Persistent block storage for AWS instances.
 - **gcePersistentDisk (GCE PD):** Persistent block storage for Google Cloud Engine.
 - **CephFS:** A shared file system that can be mounted on multiple Pods at once.
-

3. StatefulSets and Storage:

- **StatefulSets:**
 - **StatefulSets** are used for deploying stateful applications that require persistent storage.
 - StatefulSets automatically create and manage **Persistent Volume Claims (PVCs)** for each Pod in the set, ensuring each Pod has its own unique persistent storage.
 - The **PVC** created by a StatefulSet is bound to a **Persistent Volume (PV)** for each Pod, which persists across Pod restarts.
 - **Stable Network Identity:**
 - Unlike regular Deployments, StatefulSets provide a stable network identity for each Pod (e.g., `my-app-0`, `my-app-1`), which can be useful when accessing persistent data.
-

4. Accessing Storage:

- **Volume Mounts:**
 - Volumes can be mounted inside Pods to be accessible by containers.
 - The storage provided by the PVs and PVCs can be mounted to Pods as filesystems or raw block devices, depending on the configuration.
- **Example of Mounting a PVC to a Pod:**

```
apiVersion: v1
kind: Pod
metadata:
  name: mypod
spec:
  containers:
  - name: myapp
    image: myapp:latest
    volumeMounts:
    - mountPath: "/data"
      name: mydata-volume
  volumes:
  - name: mydata-volume
    persistentVolumeClaim:
      claimName: myclaim
```

5. Storage and High Availability:

- **Replication:**
 - Some storage solutions, like **NFS** or **Ceph**, provide replication features for high availability and fault tolerance.
 - Storage solutions can be configured to replicate data across multiple zones or nodes to ensure data redundancy.
- **Backup and Restore:**
 - Backup strategies should be in place for stateful applications, and solutions like **Velero** can be used for backup and restore of Kubernetes resources and persistent volumes.

6. Cloud-Native Storage:

- **Cloud Provider Storage Integration:**
 - Cloud platforms like **AWS**, **GCP**, **Azure**, and others offer cloud-native storage solutions that are integrated with Kubernetes.
 - These solutions include **block storage** (e.g., **AWS EBS**, **GCE PD**), **file storage** (e.g., **Azure Files**, **Google Cloud Filestore**), and **object storage** (e.g., **AWS S3**, **Google Cloud Storage**).
- **Cloud Provider-Specific Storage Classes:**
 - Kubernetes can use cloud-specific storage classes for dynamic provisioning of cloud-native volumes (e.g., **AWS EBS**, **Google Persistent Disk**).

7. Persistent Storage Lifecycle:

- **Binding:**
 - The process of associating a PVC with an available PV. This happens automatically if dynamic provisioning is configured.
- **Reclaim Policy:**
 - The behavior of the PV when the PVC is deleted. Common policies:
 - **Retain:** The PV is not deleted and must be manually cleaned up.

- **Delete:** The PV is deleted automatically when the PVC is deleted.
 - **Recycle:** The PV is cleaned and made available for reuse (deprecated in newer versions).
-

8. Troubleshooting Storage Issues:

- **kubectl describe pvc <claim-name>:**
 - Use this command to get detailed information about a PVC and its status.
 - **kubectl describe pv <volume-name>:**
 - Use this to check the status and details of a Persistent Volume.
 - **Pod logs:**
 - Check the logs of Pods using storage to diagnose mounting or access issues.
-

In the KCNA Exam:

In the **KCNA exam**, understanding **Kubernetes storage** fundamentals is crucial, especially how to:

- Create and use **Persistent Volumes (PVs)** and **Persistent Volume Claims (PVCs)**.
- Work with **Storage Classes** for dynamic provisioning.
- Handle **StatefulSets** and ensure persistent data in stateful applications.
- Recognize different **volume types** and their use cases.
- Manage **cloud-native storage solutions** and ensure the appropriate use of Kubernetes storage resources.

These topics are important to ensure that applications have the right kind of storage for their needs and are prepared for any potential data recovery, scaling, and performance needs in a Kubernetes

Revision #2

Created 18 November 2024 21:52:02 by Admin

Updated 21 November 2024 20:05:12 by Admin