

What is Persistent Storage?

Persistent storage is any data storage device that retains data after power to that device is shut off. It is also sometimes referred to as nonvolatile storage.

In containerization, persistent storage refers to storage volumes—usually associated with stateful applications such as databases—that remain available beyond the life of individual containers. Persistent storage volumes can be contrasted with ephemeral storage volumes that live and die with containers and are associated with stateless apps.

Types of persistent storage

Because persistent storage is designed to survive independently of any running instance, it can be used for any data that needs to be reused. Reuse can be either by different instances or beyond the life of a specific instance.

Containers

The OpenStack object storage service (openstack-swift) is a fully distributed storage solution that can be used to store any kind of static data or binary object, such as media files, large datasets, and disk images. The solution organizes these objects by using containers. Although a volume's contents can only be accessed through instances, the objects inside a container can be accessed through the object storage REST API. Therefore, the object storage can be used as a repository by almost every service in a public or private cloud.

Kubernetes

OpenShift is a platform-as-a-service (PaaS) offering that is built around Docker container packaging and uses Kubernetes to orchestrate containerized application deployments across a cluster. Previously, storage volumes had to be pre-provisioned and manually registered as PersistentVolume (PV) objects before they could be consumed by an application. StorageClass objects enable specification of a provisioner that creates storage volumes for applications on demand.

How does persistent storage work with containers?

PersistentVolume (PV)

PVs are resources in the cluster that have a lifecycle independent of any individual pod that uses a PV. This is the “physical” volume on the host machine that stores persistent data. PersistentVolumes provide storage resources in a cluster, allowing the storage resource to persist even when the pods that use them are cycled. PersistentVolumes can be statically or dynamically provisioned, and they can be customized for use by defining properties such as performance, size, and access mode.

PersistentVolumeClaim (PVC)

PVCs are requests for resources that act as claim checks to the resource. A PVC is a request for the platform to create a PersistentVolume and to attach PVs to pods via a PVC.



Trident

Trident integrates with Kubernetes PersistentVolume (PV) framework and provides a unified interface for storage provisioning from NetApp® ONTAP® and SolidFire® scale-out storage systems. It is deployed as a highly available Kubernetes “deployment” and communicates with the API server to process pod storage requests in the form of PersistentVolumeClaims (PVCs).

Persistent storage use cases

Certain functions generate data that must be kept. Other functions, such as load balancing, may not generate data that needs to be maintained. For example, a sales tax calculating function creates data that needs to be maintained and shared within an application. Critical data that must persist (data that you want to keep) can be managed with the following methods.

Data infrastructure management

Data infrastructure automation enables the agility to keep pace as IT ecosystems rapidly expand and change. This agility can be accomplished through multicloud, DevOps, and virtualization environments.

Continuous integration and continuous delivery

Continuous integration and continuous delivery are two approaches to software development that are designed to improve code quality and enable rapid delivery and deployment of code. They are usually deployed together to ensure rapid delivery of new software features and fixes.

Migrating traditional applications into the cloud

The “lift-and-shift” approach opens paths to IT modernization by moving to an open and more extensible architecture in the cloud. Companies employ lift and shift for solid business reasons, including reduced costs and improved performance and resiliency.

For more information on NetApp storage solutions, visit masterclass.eplus.com today



Not many partnerships are as tried-and-true as ours.

No need to start from scratch when you choose ePlus and NetApp for your data management and integrity needs. From ransomware resiliency and data protection to maximizing the value of hybrid cloud and integrating new technologies, we'll bring broad experience mixed with strategic vision that will help you build upon and extend your success.