

Backup on OpenShift, Tanzu and Kubernetes



Introduction

In a changing IT world where applications move to be containerized, it is important to consider what needs to be restored from backup in case of disaster or accidental delete in a Kubernetes, OpenShift or Tanzu environment.

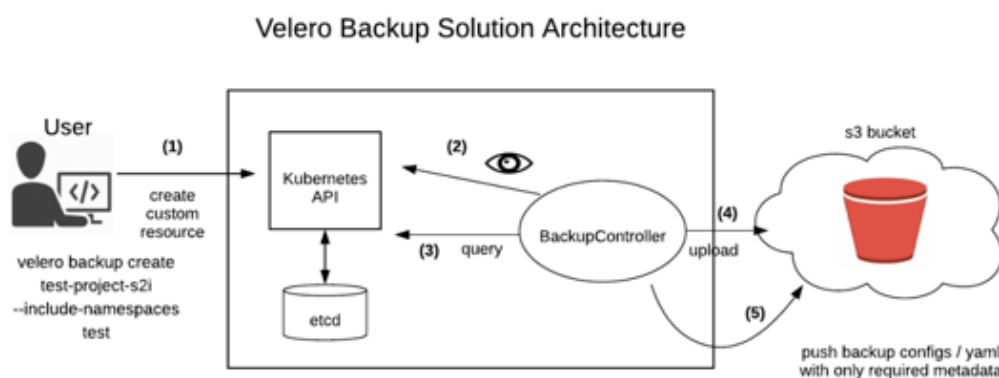
For this purpose, GBS has investigated how to best design and implement a backup for critical components in the Kubernetes, OpenShift or VMware Tanzu environment and which open-source tools (license free) are already available for it.

Velero

We have found Velero to suit our needs. Velero is a well-known Open-Source tool for (Disaster) recovery, Data Protection and Data Migration in a Kubernetes (K8s) environment. It can be used in a cloud environment, (AWS cluster) or on-premises (OpenShift). It also covers AWS and GCP as a storage provider for the backup and has scheduling functionality.

For a backup architecture in a Kubernetes environment, we depart from our regular centralized backup server. Instead, the Velero backup will be installed on K8s system itself. However, we can still use our (existing) GBS backup server as "offsite" Object Store for the backup. Our backup strategy is to store backups off-site at cloud storage or a remote backup server, not in the same environment as the original K8s cluster, to still have the backups available in case of a site disaster.

Velero uses an AWS- S3 compatible object store to store its backups. Therefore, we can use AWS cloud storage or a local S3 Object Store. For local object storage, we have integrated a MinIO Object Store that can be deployed on a centralized (existing) GBS backup server. In this case, the data will be stored in a subdirectory on the regular /backup partition.



For more information about how Velero works: <https://velero.io/docs/v1.9/how-velero-works/>

Velero can be installed and configured by GBS in order to cover the customers backup requirements in a K8s environment. GBS will inventory the customers backup requirements, and configure, deploy, test and document the Velero backup accordingly. We can also configure AWS backup facilities to be used as Velero object store if required.

Backup and Restore on a OpenShift / Kubernetes environment

With Velero installed on OpenShift/Kubernetes, the backup configuration is the key issue, as we need to make sure that all important resources are included in the backup. We suggest configuring dedicated backups for the important namespaces (projects). The backups can be scheduled on a daily/weekly basis in OADP. All resources in the namespace (Pods, Persistent Volumes etc.) will be included in the backup and can be restored to the same or a new namespace.

When restoring to a new namespace, a project copy is created, and the functionality can be tested without deleting the original project. This is fantastic restore functionality offered by Velero (OADP)! Demo is available on request.

Backup the etcd

In addition to the Velero/OADP backup, there is also the ETCD that needs to be backed up. With these two backups, we cover the full OpenShift/Kubernetes backup.

etcd is the key-value store for OpenShift Container Platform, which persists the state of all resource objects.

OpenShift includes a ETCD backup script that can be executed daily. The result is stored on a local folder.

The GBS backup server will transfer this backup to an off-site backup server, so that it is available in case of a disaster.

The OpenShift manual advises: "Back up your cluster's etcd data regularly and store in a secure location ideally outside the OpenShift Container Platform environment".

Virtual Machines in OpenShift

The latest release of OpenShift 4.1x includes OpenShift Virtualization, which allows running KVM based Virtual Machines in parallel with K8s on the same OpenShift cluster. This feature allows us to deploy a regular GBS virtualized backup server on the OpenShift cluster, thus allowing centralized backups of Virtual Machines, databases and MinIO Object Store.

For more information, please refer to this article: <https://www.redhat.com/en/technologies/cloud-computing/openshift/virtualization>

Further reading on Velero backup:

<https://patroware.medium.com/backing-up-openshift-5e8dc6b28ed9>

<https://velero.io/docs/v1.1.0/aws-config/>

<https://cloud.redhat.com/blog/velero-backup-and-restore-of-an-application-using-gp2-storageclass-on-rosa>

<https://cloud.redhat.com/blog/disaster-recovery-of-virtual-machines-within-openshift>

More GBS info

For more info on the GBS backup products and solution, please visit our website:

www.gbsbackup.com

Or feel free to contact our Technical Director:

Roel Otten (+31 20 3697972)

info@gbsbackup.com