NetBackup™ Web UI Kubernetes Administrator's Guide

Release 9.1

VERITAS™

Last updated: 2021-06-04

Legal Notice

This product may contain third-party software for which Veritas is required to provide attribution to the third party (“Third-party Programs”). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Veritas product or available at:

https://www.veritas.com/about/legal/license-agreements

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. Veritas Technologies LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

http://www.veritas.com

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

https://www.veritas.com/support

You can manage your Veritas account information at the following URL:

https://my.veritas.com

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)	CustomerCare@veritas.com
Japan	CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

https://sort.veritas.com/documents

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

NB.docs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

http://www.veritas.com/community/

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1 Introducing the NetBackup web user interface 6

About the NetBackup web UI 6
Terminology 8
Sign in to the NetBackup web UI 9
Sign out of the NetBackup web UI 11

Chapter 2 Monitoring NetBackup 12

The NetBackup dashboard 12
Monitoring jobs 12
Filter jobs in the job list 13

Chapter 3 Overview of NetBackup for Kubernetes 14

Overview 14
Features of NetBackup support for Kubernetes 15

Chapter 4 Deploying and configuring the NetBackup Kubernetes operator 16

Configuring clusters for Kubernetes 16
Prerequisites for deployment 21
Deploying the NetBackup Kubernetes operator 21
Upgrading the NetBackup Kubernetes operator deployment 24
Deleting a NetBackup Kubernetes operator deployment 24
Operator configuration on NetBackup side 24
Operator configuration on Kubernetes side 25
Getting token for adding clusters 25
About expired images 27

Chapter 5 Managing Kubernetes assets 28

Adding Kubernetes clusters 28
Configuring settings 29
Managing Kubernetes assets 30

Chapter 6 Protecting Kubernetes assets 32

Kubernetes protection plans 32
Configuring backup options for Kubernetes protection plans 32

Chapter 7 Recovering Kubernetes assets 34

Recovering Kubernetes assets 34

Chapter 8 Troubleshooting Kubernetes issues 37

Connecting to the primary server with a short host name 37
Cluster discovery fails 38
Error during backup: Namespace has been marked for deletion 38
Error during restore: Final job status is partially failed 38
Backup stuck in InProgress state 39
Restore stuck in InProgress state 39

Introducing the NetBackup web user interface

This chapter includes the following topics:

- About the NetBackup web UI
- Terminology
- Sign in to the NetBackup web UI
- Sign out of the NetBackup web UI

About the NetBackup web UI

The NetBackup web user interface provides the following features:

- Ability to access the primary server from a web browser, including Chrome and Firefox. For details on supported browsers for the web UI, see the NetBackup Software Compatibility List.
  Note that the NetBackup web UI may behave differently for different browsers. Some functionality, for example a date picker, may not be available on all browsers. These inconsistencies are due to the capabilities of the browser and not because of a limitation with NetBackup.
- A dashboard that displays a quick overview of the information that is important to you.
- Role-based access control (RBAC) that lets the administrator configure user access to NetBackup and to delegate the tasks for workload protection.
- Protection of assets is achieved through protection plans, job management, and visibility of the protection status of assets.

Alternatively, policy management is also available for a limited number of policy types. More information about these policy types is available:

- Workload administrators can create protection plans, subscribe assets to the protection plans that meet the SLO, monitor protection status, and perform self-service recovery of assets.

Note: The NetBackup web UI is best viewed at a 1280x1024 or higher screen resolution.

Access control in the NetBackup web UI

NetBackup uses role-based access control to grant access to the web UI. Access control is accomplished through roles.

- A role defines the operations that a user can perform and the features that the user can access in the web UI. For example, access to any workload assets, protection plans, or credentials.
- RBAC is only available for the web UI and the APIs.

  Other access control methods for NetBackup are not supported for the web UI and APIs, with the exception of Enhanced Auditing (EA).

Monitor NetBackup jobs

The NetBackup web UI lets administrators more easily monitor NetBackup job operations and identify any issues that need attention.

Protection plans: One place to configure schedules, storage, and storage options

Protection plans offer the following benefits:

- A default workload administrator can select the protection plans to use to protect assets.
- With the necessary RBAC permissions, a workload administrator can create and manage protection plans, including the backup schedules and storage that is used.
- In addition to schedules for backups, a protection plan can also include a schedule for replication and long-term retention.
- When you select from your available storage, you can see any additional features available for that storage.

Self-service recovery

The NetBackup web UI makes it easy for a workload administrator to recover VMs, databases, or other asset types applicable to that workload.

Terminology

The following table describes the concepts and terms in web user interface.

Table 1-1 Web user interface terminology and concepts

Term	Definition
Asset group	See *intelligent group*.
Asset	The data to be protected, such as physical clients, virtual machines, and database applications.
Backup now	An immediate backup of an asset. NetBackup performs a one-time, full backup of an asset using the selected protection plan. This backup does not affect any scheduled backups.
Intelligent group	Allows NetBackup to automatically select assets for protection based on the criteria (queries) that you specify. An intelligent group automatically stays up-to-date with changes in the production environment. These groups are also referred to as asset groups. These groups appear under the tab **Intelligent VM groups** or **Intelligent groups**.
Protection plan	A protection plan defines when backups are performed, how long the backups are retained, and the type of storage to use. Once a protection plan is set up, assets can be subscribed to the protection plan.
RBAC	Role-based access control. The role administrator can delegate or limit access to the NetBackup web UI through the roles that are configured in RBAC. **Note:** The roles that you configure in RBAC do not control access to the NetBackup Administration Console or the CLIs.
Role	For RBAC, defines the operations that a user can perform and the assets or objects that they can access. For example, you can configure a role to manage recovery of specific databases and the credentials that are needed for backups and restores.

Table 1-1 Web user interface terminology and concepts *(continued)*

Term	Definition
Storage	The storage to which the data is backed up, replicated, or duplicated (for long-term retention).
Subscribe, to a protection plan	The action of selecting an asset or an asset group to subscribe to a protection plan. The asset is then protected according to the schedule in the plan. The web UI also refers to *Subscribe* as *Add protection*.
Unsubscribe, from a protection plan	*Unsubscribe* refers to the action of removing protection or removing an asset or asset group from a plan.
Workload	The type of asset. For example, VMware, RHV, AHV, or Cloud.

Sign in to the NetBackup web UI

Authorized users can sign in to a NetBackup primary server from a web browser, using the NetBackup web UI.

The following sign-in options are available:

- Sign in with a username and password
- Sign in with a certificate or smart card
- Sign in with single sign-on (SSO)

Sign in with a username and password

Only authorized users can sign in to NetBackup web UI. Contact your NetBackup security administrator for more information.

To sign in to a NetBackup primary server using a username and password

1. Open a web browser and go to the following URL.

   https://*primaryserver*/webui/login

   The *primaryserver* is the host name or IP address of the NetBackup primary server that you want to sign in to.

2. Enter your credentials and click **Sign in**.

   For example:

For this type of user	Use this format	Example
Local user	*username*	**jane_doe**
Windows user	*DOMAIN\username*	**WINDOWS\jane_doe**
UNIX user	*username@domain*	**john_doe@unix**

Sign in with a certificate or smart card

You can sign in to NetBackup web UI with a smart card or digital certificate if you are an authorized user. Contact your NetBackup security administrator for more information.

To use a digital certificate that is not on a smart card, you must first upload the certificate to the browser's certificate manager. See the browser documentation for instructions or contact your certificate administrator for more information.

To sign in with a certificate or smart card

1. Open a web browser and go to the following URL.

   https://*primaryserver*/webui/login

   The *primaryserver* is the host name or IP address of the NetBackup primary server that you want to sign in to.

2. Click **Sign in with certificate or smart card**.

3. When your browser prompts you, select the certificate.

Sign in with single sign-on (SSO)

You can sign in to NetBackup web UI with the single sign-on (SSO) option if SAML is configured as an identity provider in your NetBackup environment. Contact your NetBackup security administrator for more information.

To sign in to a NetBackup primary server using SSO

1. Open a web browser and go to the following URL.

   https://*primaryserver*/webui/login

   The *primaryserver* is the host name or IP address of the NetBackup primary server that you want to sign in to.
2. Click **Sign in with single sign-on**.
3. Follow the steps as provided by your administrator.

   On subsequent logons, NetBackup signs you in automatically to the primary server.

Sign out of the NetBackup web UI

Note that NetBackup automatically signs you out of the web UI after 24 hours, which is the maximum time that is allowed for a user session. After that time NetBackup requires that you sign in again. You can also sign out if you want to change the sign-in option that you want to use (username and password, smart card, or single sign-on (SSO)).

To sign out of the NetBackup web UI

◆ On the top right, click the profile icon and click **Sign out**.

Monitoring NetBackup

This chapter includes the following topics:

- The NetBackup dashboard
- Monitoring jobs
- Filter jobs in the job list

The NetBackup dashboard

The NetBackup dashboard provides a quick view of the details that are related to your role in your organization.

Table 2-1 The NetBackup dashboard

Dashboard widget	Description
Jobs	Lists job information, including the number of active and queued jobs and the status of attempted and completed jobs.

Monitoring jobs

Use the **Jobs** node to monitor the jobs in your NetBackup environment and view the details for a specific job.

To monitor a job

1 Click on a job name that you want to view.

On the **Overview** tab you can view information about a job.

- The **File List** contains the files that are included in the backup image.

- The **Status** section shows the status and the status codes that are related to the job. Click the status code number to view information about this status code in the Veritas Knowledge Base.
  See the NetBackup Status Codes Reference Guide.

2 Click the **Details** tab to view the logged details about a job. You can filter the logs by error type using the drop-down menu.

See "Filter jobs in the job list" on page 13.

Filter jobs in the job list

You can filter the jobs to display the jobs in a specific state. For example, you can display all of the active jobs or all of the suspended jobs.

To filter the job list

1 Click **Jobs**.

2 Above the job list, click the **Filter** option.

3 In the **Filter** window, select a filter option to dynamically change the jobs that are displayed. The filter options are as follows:

- **All**
- **Active**
- **Done**
- **Failed**
- **Incomplete**
- **Partially Successful**
- **Queued**
- **Successful**
- **Suspended**
- **Waiting for Retry**

4 Click **Apply Filters**.

5 To remove the selected filters, click **Clear All**.

Overview of NetBackup for Kubernetes

This chapter includes the following topics:

- Overview
- Features of NetBackup support for Kubernetes

Overview

The NetBackup web UI provides the capability for backups and restores of Kubernetes applications in the form of namespaces. The protectable assets in the Kubernetes clusters are automatically discovered in the NetBackup environment and administrators can select one or more protection plans that contain the wanted schedule, backup, and retention settings.

The NetBackup web UI lets you perform the following operations:

- Add Kubernetes cluster for protection
- View discovered namespaces.
- Manage permissions for roles
- Set resource limits to optimize load on your network
- Select protection plans to protect Kubernetes assets.
- Restore namespaces and persistent volumes.
- Monitor backup and restore operations.

Features of NetBackup support for Kubernetes

Table 3-1 NetBackup for Kubernetes

Feature	Description
Integration with NetBackup role-based access control (RBAC)	The NetBackup web UI provides RBAC roles to control which NetBackup users can manage Kubernetes operations in NetBackup. The user does not need to be a NetBackup administrator to manage Kubernetes operations.
Licensing	Capacity-based licensing.
Protection plans	The following benefits are included: ■ Use a single protection plan to protect multiple Kubernetes namespaces. The assets can be spread over multiple clusters. ■ Ability to retain or discard partially successful backups. ■ You are not required to know the Kubernetes commands to protect the Kubernetes assets.
Intelligent management of Kubernetes assets	NetBackup automatically discovers the namespaces, persistent volumes, persistent volume claims, and so on, in the Kubernetes clusters. You can also perform manual discovery. After the assets are discovered, the Kubernetes workload administrator can select one or more protection plans to protect them.
Kubernetes specific credentials	Kubernetes service accounts used to authenticate and manage the clusters.
Backup and restore features	The following features are available for backup and restore: ■ Backups and restore are managed entirely by the NetBackup server from a central location. Administrators can schedule automatic, unattended backups for namespaces on different Kubernetes clusters. ■ The NetBackup web UI supports the backup and restore of namespaces from one interface. ■ Backup schedules for full backups. ■ Manual backups and snapshot only backups. ■ Restore Kubernetes namespaces and persistent volumes to different locations. ■ Resource throttling for each cluster to improve the performance of backups.
Snapshot backups	NetBackup can perform backups of Kubernetes namespaces with snapshot methodology, achieving faster recovery time objectives.

Chapter 4

Deploying and configuring the NetBackup Kubernetes operator

This chapter includes the following topics:

- Configuring clusters for Kubernetes
- Prerequisites for deployment
- Deploying the NetBackup Kubernetes operator
- Upgrading the NetBackup Kubernetes operator deployment
- Deleting a NetBackup Kubernetes operator deployment
- Operator configuration on NetBackup side
- Operator configuration on Kubernetes side
- Getting token for adding clusters
- About expired images

Configuring clusters for Kubernetes

You need to configure the clusters before you can deploy the NetBackup™ Kubernetes operator. You can deploy the NetBackup Kubernetes operator using Helm charts in three different platforms:

- Red Hat OpenShift
- Google Kubernetes Engine (GKE)

- VMware Tanzu

Configuring OpenShift for NetBackup

Before you begin make sure that you have the required privileges in your OpenShift account to perform these operations.

To configure OpenShift:

- Log on to the OpenShift OC using the CLI, using the following command:

  ```
  oc login --token=<TOKEN> --server=<URL>
  ```

  Where:

  - <TOKEN>: Is your logon token
  - <URL>: Is your OpenShift server URL

Note: You can get the Token and the URL by logging on to your OpenShift account. Click the name of your OpenShift admin account by which you logged into the console, on top right of the home page, and then click the **Copy Login Command** option. On the new page that opens, click **Display Token**, to see the command.

This command adds a new kubectl context in the `~/.kube/config` file, and sets this new context as kubectl current context.

Configuring GKE for NetBackup

Before you begin make sure that you have the required privileges in your GKE account to perform these operations

Prerequisites:

- The port number for GKE clusters can be 443, 6443, or 8443. Default port is 443. Verify the correct secure port number before adding.
- When creating persistent volumes or persistent volume claims on GKE, specify a storage class whose `Provisioner` is `kubernetes.io/gce-pd`.

To log on using an existing account:

1. Use the following command to logon to the GKE account by using an existing user account:

   ```
   gcloud auth login <account>
   ```

2. Enter the logon credentials interactively or non-interactively.

3 To list all the clusters and find the cluster name, run the command:

```
gcloud container clusters list
```

The output looks like this:

```
NAME               LOCATION       MASTER_VERSION    MASTER_IP
csi-cluster        us-central1-c  1.17.14-gke.400   35.238.135.170
sailor             us-central1-c  1.16.15-gke.6000  35.224.28.128
surens-cluster     us-east1-b     1.17.14-gke.1600  35.231.17.183
bw-kube-cluster-1  us-east1-c     1.16.15-gke.6000  35.196.24.132
```

4 To get credentials for the cluster and add it to `.kube/config`, run this command:

```
gcloud container clusters get-credentials <cluster name>
```

For example: gcloud container clusters get-credentials bw-kube-cluster-1

Alternatively, you can create and use a dedicated service account for your cluster to logon.

To create a dedicated service account:

1 To create an account, run this command:

```
gcloud iam service-accounts create <account name> --display-name
"<account description>"
```

For example: gcloud iam service-accounts create veritas-netbackup-k8s-sa --display-name "Veritas NetBackup K8s Service Account"

2 To list the users, run this command:

```
gcloud iam service-accounts list --filter <email ID>@<project
ID>.gserviceaccount.com
```

For example: gcloud iam service-accounts list --filter veritas-netbackup-k8s-sa@projectID.gserviceaccount.com

3 To download the service account key, run this command:

```
gcloud iam service-accounts keys create <key json file name>
--iam-account <e-mail address of the service account>
```

For example: gcloud iam service-accounts keys create veritas-netbackup-k8s-sa-key.json --iam-account <e-mail ID of the service account>

4. To associate a role, run this command:

   ```
   gcloud iam roles create <role name> --project <project ID> --file
   ./<role name>.yaml
   ```

   For Example: gcloud iam roles create rolename --project projectID --file ./rolename.yaml

5. To activate the service account, run this command:

   ```
   gcloud auth activate-service-account --project=<project ID>
   --key-file=<key file name>
   ```

   For example: gcloud auth activate-service-account --project=<YOUR PROJECT ID> --key-file=veritas-netbackup-k8s-sa-key.json

6. To list all the clusters and find the cluster name, run the command:

   ```
   gcloud container clusters list
   ```

   The output looks like this:

   ```
   NAME               LOCATION       MASTER_VERSION    MASTER_IP
   csi-cluster        us-central1-c  1.17.14-gke.400   35.238.135.170
   sailor             us-central1-c  1.16.15-gke.6000  35.224.28.128
   surens-cluster     us-east1-b     1.17.14-gke.1600  35.231.17.183
   bw-kube-cluster-1  us-east1-c     1.16.15-gke.6000  35.196.24.132
   ```

7. To get credentials for the cluster and add it to `.kube/config`, run this command:

   ```
   gcloud container clusters get-credentials <cluster name>
   ```

   For example: gcloud container clusters get-credentials bw-kube-cluster-1

Configuring VMware Tanzu for NetBackup

Before you begin make sure that you have the required privileges in your Tanzu account to perform these operations. Make sure that you have the TKG client installed.

Add an existing Tanzu management cluster to your local TKG instance:

1. Copy the `kube-tkg/config` file from the management cluster to local user home directory: `~/`

2. Run the command: `chmod 775 ~/.kube-tkg/config`

3. Run the command: `export KUBECONFIG=.kube-tkg/config`

4. To get a list of the contexts, run the command: `tkg get mc`. The output looks like this:

```
  MANAGEMENT-CLUSTER-NAME  CONTEXT-NAME               STATUS
  tkg-mgmt *               tkg-mgmt-admin@tkg-mgmt    Success
  tkg1-mgmt                tkg1-mgmt-admin@tkg1-mgmt  Success
  tkg2-mgmt                tkg2-mgmt-admin@tkg2-mgmt  Success
```

5. To switch to the TKG context, run the command: `tkg set mc tkg1-mgmt`

   The current management cluster context is switched to **tkg1-mgmt**.

6. To check the kubectl context, run the command: `kubectl config get-contexts`. The output looks like this:

```
 CURRENT NAME            CLUSTER AUTHINFO   NAMESPACE
tkg1-mgmt-admin@tkg1-mgmt      tkg1-mgmt    tkg1-mgmt-admin
tkg2-mgmt-admin@tkg2-mgmt      tkg2-mgmt    tkg2-mgmt-admin
```

7. To check the management clusters in the local TKG instance, run the command: `tkg get mc`. The output looks like this:

```
 [dxxxx@xxxxxxxxx01vm1392 ~]$ tkg get mc
  MANAGEMENT-CLUSTER-NAME  CONTEXT-NAME               STATUS
  tkg-mgmt                tkg-mgmt-admin@tkg-mgmt    Success
  tkg1-mgmt *                tkg1-mgmt-admin@tkg1-mgmt  Success
  tkg2-mgmt                tkg2-mgmt-admin@tkg2-mgmt  Success
```

8. To get all the clusters in your current context, run the command: `tkg get clusters`. The output looks like this:

```
 NAME         NAMESPACE  STATUS  CONTROLPLANE  WORKERS  KUBERNETES
  tkg1-cluster1  default  running  3/3        3/3   v1.19.3+vmware.1
  tkg1-cluster2  default  running  3/3        3/3   v1.19.3+vmware.1
  tkg1-cluster3  default  running  3/3        3/3   v1.19.3+vmware.1
```

9. To add credentials to `kubectl` configuration file, run the command: `tkg get credentials tkg1-cluster1`

   This saves the credentials of the workload cluster **tkg1-cluster1** in the configuration file. To access the cluster, run the command: `kubectl config use-context tkg1-cluster1-admin@tkg1-cluster1`

10 To switch to the **kubectl** context, run the command: `kubectl config use-context tkg1-cluster1-admin@tkg1-cluster1`

11 To check the kubectl context, run the command: `kubectl config get-contexts`

The output looks like this:

```
 CURRENT NAME              CLUSTER AUTHINFO         NAMESPACE
tkg1-cluster1-admin@tkg1-cluster1   tkg1-cluster1   tkg1-cluster1-admin
tkg1-mgmt-admin@tkg1-mgmt           tkg1-mgmt       tkg1-mgmt-admin
tkg2-mgmt-admin@tkg2-mgmt           tkg2-mgmt       tkg2-mgmt-admin
```

Now you can use any kubectl commands in the **tkg1-cluster1**.

Prerequisites for deployment

Download and install Velero in the clusters where you want to deploy the NetBackup Kubernetes operator.

Note: For the supported Velero versions refer to the NetBackup Software Compatibility List. For Velero installation and configuration refer to the Velero documentation.

Deploying the NetBackup Kubernetes operator

After configuring your cluster(s), you can deploy the NetBackup Kubernetes operator in them. You must deploy the operator in each cluster, where you want to use NetBackup.

Configuring the Helm Chart

You can use Helm Chart to deploy NetBackup Kubernetes operator. You can create a chart for the NetBackup Kubernetes operator. Here is the Helm chart and tree structure layout.

```
netbackupkops-helm-chart

   ├── charts

   ├── Chart.yaml

   ├── templates
```

```
│    └── deployment.yaml

└── values.yaml
```

To deploy the NetBackup Kubernetes operator:

1. Download the operator service package.
2. Extract the package to the home directory. The `netbackupkops-helm-chart` folder should be in the home directory.
3. To list all cluster contexts, run the command: `kubectl config get-contexts`
4. To switch to the cluster where you want to deploy the operator service, run the command: `kubectl config use-context <cluster-context-name>`
5. To change the current directory to your home directory, run `cd ~`
6. If you use a private docker registry, follow the instructions in this step to create a secret `nb-docker-cred` in Velero namespace. Otherwise, skip to the next step.
   - To log on to the private docker registry, run the command: `docker login -d <user name> -p <password>`
     After logon, the `config.json` file containing the authorization token is created or updated. To view the `config.json` file, run the command: `cat ~/.docker/config.json`
     The output looks like:

     ```
     {

        "auths": {

            "https://index.docker.io/v1/": {

                "auth": "c3R...zE2"

            }

        }

     }
     ```

   - To create a secret named as `netbackupkops-docker-cred` in the Velero namespace, run the command:

     ```
     kubectl create secret generic netbackupkops-docker-cred \
     --from-file=.dockerconfigjson=.docker/config.json \
     ```

```
--type=kubernetes.io/dockerconfigjson -n velero
```

- To check if the secret `netbackupkops-docker-cred` is created in the Velero namespace, run the command: `kubectl get secrets -n velero`
- If you use an image tar file, to load the image to the docker cache and push the image to the docker image repository, run the following commands:

  ```
  docker load -i <name of the tar file>
  docker tag <image name:tag of the loaded image>
  <repo-name/image-name:tag-name>
  docker push <repo-name/image-name:tag-name>
  ```

- Open the `netbackupkops-helm-chart/values.yaml` file in a text editor, and then replace the value *image* in the *manager* section, with your image name with tag (*repo-name/image-name:tag-name*) and save the file.

7 To deploy the NetBackup Kubernetes operator service, run the following command in a single line:

```
helm install <release name of the deployment>
./netbackupkops-helm-chart -n <namespace in which NetBackup
operator service will run>
```

For example: `helm install veritas-netbackupkops ./netbackupkops-helm-chart -n netbackup`

- You can change the release name of the deployment as required.
- The `-n` option is required to specify the namespace in which NetBackup operator service is intended to run. It must be the same namespace where Velero is intended to run.

8 To check the status of the deployment, run the command:

```
helm list -n <namespace in which NetBackup operator service will
run>
```

For example:

```
helm list -n netbackup
```

9 To check the release history, run the command: `helm history veritas-netbackupkops -n <namespace in which NetBackup operator service will run>`

For example:

```
helm history veritas-netbackupkops -n netbackup
```

Upgrading the NetBackup Kubernetes operator deployment

You can upgrade the NetBackup Kubernetes operator deployment using Helm commands.

```
helm upgrade <release name> ./<directory of the chart> -n <namespace>
```

For example: `helm upgrade veritas-netbackupkops ./nbukops-helm-chart -n netbackup`

Deleting a NetBackup Kubernetes operator deployment

You can delete a NetBackup Kubernetes operator deployment from a cluster as required.

To delete a NetBackup Kubernetes operator deployment, run the following command:

```
helm uninstall <release name> -n <namespace>
```

For example: `helm uninstall veritas-netbackupkops -n netbackup`

Operator configuration on NetBackup side

NetBackup 9.1 introduces two new default RBAC roles:

- **Default NetBackup Kubernetes operator** : This role provides the required permissions to the operator running in the Kubernetes cluster to communicate with the NetBackup Web Services. The security administrator or the NetBackup administrator can assign the required users to this role. This API key has limited permissions, defined as part of the role. The API key and CA certificate are required in the configuration at the Kubernetes cluster end.

  Note: Contact the Security Administrator or the NetBackup Administrator to fetch the NetBackup CA certificate.

- **Default Kubernetes administrator**: This role has all the required permissions for NetBackup web UI and APIs.

Operator configuration on Kubernetes side

You must create a Kubernetes secret resource, as a part of the required configuration at the Kubernetes cluster for running the NetBackup Kubernetes operator. The file name of the secret must be the same as of the configured primary server and the namespace must be the same where NetBackup Kubernetes operator is running. The API key and CA certificate of the NetBackup primary server are required in the configuration at the Kubernetes cluster end.

Note: Contact the Security Administrator or the NetBackup Administrator to fetch the NetBackup CA certificate.

Here is the format of the secret. Provide the values as specified in the pointed brackets:

```
apiVersion: v1
kind: Secret
metadata:
  name: <NetBackp primary server host name or IP address.>>
  namespace: <Namespace name where the NetBackup Kubernetes operator
is deployed>>


type: Opaque
stringData:
  apiKey: <API key of the primary server>>
  caCert: "<CA certificate of the primary server>>"
```

Contact your security administrator for the API key and the CA certificate.

Getting token for adding clusters

To add Kubernetes clusters in NetBackup you need CA Certificate and a token. To get the CA Certificate and the token, run the following command in the Kubernetes cluster:

```
kubectl get secret <[namespace-name]-backup-server-token-<id>> -n
<namespace name> -o yaml
```

Select the token without an annotation field

Here is a sample CA certificate:

```
LS0tLS1CRUdJTiBDRVJUSUZJQ0FURS0tLS0tCk1JSURCakNDQWU2
Z0F3SUJBZ0lCQVRBTkJna3Foa2lHOXcwQkFRc0ZBREFWTVJNd0VR
```

```
WURWUVFERXdwdGFXNXAKYTNWaVpVTkJNQjRYRFRJd01URXhPVEV3
TURZeU1sb1hEVE13TVRFeE9ERXdNRFl5TWxvd0ZURVRNQkVHQTF
VRQpBeE1LYldsdWFXdDFZbVZEUVRDQ0FTSXdEUVlKS29aSWh2Y05
BUUVCQlFBRGdnRVBBRENDQVFvQ2dnRUJBTklaClduc0MvTEpjaUV
NOGx0ZnU0dzFPcmNaeTVZemhOTXoxQWV0V09xRmUrQ0VxblFVY3h
mVEpwOElWMFRTei9yYmYKSHVBdWlmWTd2ZGxNdC9zREJUbDlIMGF
xUkxLdG9KMDZaUHVBRzN0WjA5Nm1VUzV5bXYzRktWV2kvaVMyYz
I0ZQpFc2NENTBRaTRyYUM5YTlHK1NuSWVRNXYrQzZGUU9vYnBuS
ERXOTNIMlRpK3gyaEMrTHVoSndVVlRldG1EbzkyCktHendENU5OU
kV0L1FPYnVtaTN0QnFPMTdpSThua2xwb0tBd0RYQWlYd2ZjeFpQ
RXNrKytkajRBNVo2bWFGRHMKMUxxVkQ3ZFpkYk1vM08rTDJ6bzB
KdFIzWXYzenY3L0tYM0JDVmdzQWduQWdNeWJUYWMyenRRYzhsWH
hwLzZxcAowbTRPT0h0ME1KYnRSMmo4bWJrQ0F3RUFBYU5oTUY4d
0RnWURWUjBQQVFIL0JBUURBZ0trTUIwR0ExVWRKUVFXCk1CUUdD
Q3NHQVFVRkJ3TUNCZ2dyQmdFRkJRY0RBVEFQQmdOVkhSTUJBZjh
FQlRBREFRSC9NQjBHQTFVZERnUVcKQkJSUkVNM3JxQjhFTjRJW
FFiQ3RJd2hhb3hzeW1EQU5CZ2txaGtpRzl3MEJBUXNGQUFPQ0F
RRUFsY2ZURGNObwpoZi9EM3BQYmx6V3BXUm5xbUc2aTF5eG0wT
2V3OUJWMjhVeDc0S1ppcGEvUm5valpaVDlRdmwvcmg3YW5rRHd
NCllDSlVsZUNHVWJkc1dwRHpycFlqa0lJVlMybHkxeHpUWkNLY0
FWOENEWWkzdjdHdWswR2R2SXc0VENndk5XajIKanBDbC9QWkFp
ZUFXdTlYL2R4THU1S01FN05uTnlGNWx4Uy85cTVvMkRUSS8reD
RncEQwQ09rQVl3SDZ4SzViUgp2WGNabFJ3NmNlMWlTZG43dVE1
V1dxcU50ZEQ1MHRNRHlzWERqUzI4WVh6WjlRYThkMEVnR1E1dW
JMZnZZdzJkCm9xcmZIZjN6bitYajNFVUg0eXRORkdOc1hMN0t4
NVdtNjNjTGlrSzBLV1dOQjMwdVpsVEljUXIyenQ2MGFjK28KbC9
0dFhsUWdoaGUwaFE9PQotLS0tLUVORCBDRVJUSUZJQ0FURS0tLS0tCg==
```

Here is a sample token:

```
ZXlKaGJHY2lPaUpTVXpJMU5pSXNJbXRwWkNJNkluRjZNMUl5UlR
reWJVSmFkbTFqV1Zaa1FtNVNkbkZwWjFSSWIzWmpTa1ZhWm5KSV
NVSnBRVEJUWVhjaWZRLmV5SnBjM01pT2lKcmRXSmxjbTVsZEdW
ekwzTmxjblpwWTJWaFkyTnZkVzUwSWl3aWEzVmlaWEp1WlhSbGN
5NXBieTl6WlhKMmFXTmxZV05qYjNWdWRDOXVZVzFsYzNCaFkyVW
lPaUoyWld4bGNtOGlMQ0pyZFdKbGNtNWxkR1Z6TG1sdkwzTmxjbl
pwWTJWaFkyTnZkVzUwTDNObFkzSmxkQzV1WVcxbElqb2lkbVZz
WlhKdkxXSmhZMnQxY0MxelpYSjJaWEl0ZEc5clpXNHRhSEpqT0
cwaUxDSnJkV0psY201bGRHVnpMbWx2TDNObGNuWnBZMlZoWTJOd
mRXNTBMM05sY25acFkyVXRZV05qYjNWdWRDNXVZVzFsSWpvaWRt
VnNaWEp2TFdKaFkydDFjQzF6WlhKMlpYSWlMQ0pyZFdKbGNtNWx
kR1Z6TG1sdkwzTmxjblpwWTJWaFkyTnZkVzUwTDNObGNuWnBZMl
```

```
V0WVdOamIzVnVkQzUxYVdRaU9pSTFNVEptWldNd09DMWlaRFV5T
FRRd01HRXRZV1V3TWkwMlpUbGpaVGhpWmpObE1Ea2lMQ0p6ZFdJ
aU9pSnplWE4wWlcwNmMyVnlkbWxqWldGalkyOTFiblE2ZG1Wc1p
YSnZPblpsYkdWeWJ5MWlZV05yZFhBdGMyVnlkbVZ5SW4wLnFEWm
t2bDNmSHlabTQzNUQyakZGX2Q5MlA2RkdFb1R0Mmx2V1J6RGR5V
GItYngxSnZ3S25Wa1M0MGswRF9jeGlMcEx5X3liNGVqelZJM2dz
UG0xM0hJUlV2bWhiSEZaUzhlX0FvOTdnbGhOd3VpQlhncjRjNW0
3dUd3eGVKOWs4eERWazVhUVhUalM4cWJlMHB4QXhpVG9EOUt4aF
BtMTVoNy1EaUtjbHBlZEJkZ2N1V2JHenRuci1uXzAzYUFFbkY3Y
zU2c1Z5N1VrV2ZUQXZMZXBZUG9jZkJoRjY3cUR4eEMza2d0S2U4
SnJUNlItclgxYWRnQVhnRnJ5WDJYNGM0RUI3WE14NFd6SFMzQXR
RdEFqNno3eEVMNXQ4eHlQZ1EtMnlpeGJudzVUTXVac1JLcnZyak
1OX2FxUTRDTEJlM29BWEVXaEdJaW1uaXgydkdpNVVVdw==
```

About expired images

To reclaim the storage space occupied by the expired Kubernetes images, NetBackup submits a delete request to Velero by creating the `deletebackuprequest` custom resource. However, NetBackup does not wait for the deletion to complete or track the delete request status. Velero accepts the delete request and performs delete.

You can track the progress of the delete request by listing the `deletebackuprequest.velero.io` CRs within the Kubernetes cluster. Once the Velero backup is deleted, the `deletebackuprequest` CR is also deleted.

Note: Manual expiration of the images is supported only through CLI and the API, not through web UI and Java UI.

Chapter 5

Managing Kubernetes assets

This chapter includes the following topics:

- Adding Kubernetes clusters
- Configuring settings
- Managing Kubernetes assets

Adding Kubernetes clusters

You can add Kubernetes clusters in NetBackup and discover all the assets inside the cluster automatically. For asset discovery to take place after you add the cluster, you need to add operator configuration to the cluster.

See “Operator configuration on Kubernetes side” on page 25.

To add a cluster

1. On the left click **Kubernetes**, under **Workloads**.
2. Click the **Kubernetes clusters** tab, click **Add**.
3. In the Add Kubernetes cluster page, enter the following:
   - **Cluster name**: Enter a name for the cluster. The name should be a DNS resolvable value or an IP address.
   - **Port**: Enter the Kubernetes API server port number.
   - **Controller namespace**: Enter the namespace where the NetBackup Kubernetes operator is deployed in the Kubernetes cluster.
4. Click **Next**. In the **Manage credentials** page, you can add credentials to the cluster.

- To use an existing credential, choose **Select from an existing credential**, and click **Next**. In the next page, select the required credentials, and click **Next**.
- To create a new credential, click **Add credential**, and click **Next**. In the **Manage credentials** page, enter the following:
  - **Credential name**: Enter a name of the credential.
  - **Tag**: Enter a tag to associate with the credential.
  - **Description**: Enter a description of the credential.
  - **Token**: Enter the authentication token value in Base64 encoded form.
    See "Getting token for adding clusters" on page 25.
  - **CA certificate**: Enter the CA certificate file contents.
    See "Getting token for adding clusters" on page 25.

5 Click **Next**.

The credentials are validated and on successful validation, the cluster is added. After the cluster is added, autodiscovery runs to discover available assets in the cluster.

Configuring settings

The Kubernetes settings let you configure the various aspects of the Kubernetes deployment.

Setting the Kubernetes resource limit

With this setting you can control the number of backups that can be performed simultaneously on Kubernetes clusters. For example, if you protect 20 assets, and you have set the limit to 5, only five assets can perform backup simultaneously, rest of the 15 assets stand in a queue. After one of the first 5 assets completes the backup, an asset from the queue takes its place.

The default value for this resource limit is 1. Indicating that only one backup job per cluster can be in progress, while the rest of the assets are the queued state.

Configuring this setting is recommended for optimized use of your system and network resources. The settings apply to all Kubernetes backups for the selected primary server.

To set the resource limit

1 On the left, **Workloads > Kubernetes**.

2 On top right, click **Kubernetes settings > Resource limits**.

3 Click **Edit**, next to **Backup jobs per Kubernetes cluster**.

4 In the **Edit Kubernetes cluster** dialog:

- Enter a value in the **Global** field, to set a global limit for all the clusters. This limit denotes the number of backup jobs that are performed simultaneously on a cluster.
- You can add individual limits to the clusters that override the global limit for that cluster. To set individual limits to the clusters, click **Add**.
- Select a cluster from the list and enter a value for the limit. You can add limits to each available cluster in your deployment.
- Click **Save** to save the changes.

Configuring autodiscovery frequency

Autodiscovery keeps a count of the NetBackup protected assets in your clusters. This setting lets you set the frequency by which NetBackup runs autodiscovery to locate new assets in your clusters and gather count of the assets that are removed or deleted from the clusters.

Possible values are between 5 minutes to one year. The default value is 30 minutes.

To set the autodiscovery frequency

1 On the left, click **Workloads > Kubernetes**.

2 On top-right, click **Kubernetes settings > Autodiscovery**.

3 Click **Edit**, near **Frequency**.

4 Enter the number of hours after which NetBackup runs autodiscovery. Click **Save**.

Configuring permissions

Using manage permissions, you can assign different access privileges to the user roles. For more information see the *Managing role-based access control* chapter in the *NetBackup Web UI Administrator's Guide*.

Managing Kubernetes assets

The **Namespaces** tab (**Workloads** > **Kubernetes**), lets you monitor the assets in your Kubernetes clusters, see their protection status, and easily add protection to any unprotected assets. You can also take a quick backup of asset using the backup now feature. This feature creates a one-time backup of the selected asset without affecting any scheduled backups.

The Namespaces tab displays with all the discovered Kubernetes assets that NetBackupcan protect. This tab displays the following information:

- **Namespaces**: Display name of the asset.
- **Cluster**: The cluster to which the asset belongs.
- **Protected by**: Name of the protection plan applied to the asset.
- **Last successful backup**: Date and time of the last successful backup of the asset.

You can perform the following action in the **Namespaces** tab.

To add protection to an unprotected asset

1. On the left, click **Workloads > Kubernetes**.
2. Select the option in the rows of the assets. Click **Add protection** on top right. Alternatively, click the Actions menu in the row of the asset and click **Add protection**.
3. Select a protection plan from the list and click **Next**. In the next page, click **Protect**.

To quickly back up an asset

1. Select the option in the rows of the assets, click **Backup now** on top right. Alternatively, click the Actions menu in the row of the asset and click **Backup now**.
2. In the next page,
   - If you backup an already protected asset, select a protection plan from the list of plans to which the asset is already subscribed, and click **Start backup**.
   - If you are backing up an unprotected asset, select a protection plan from the available plans for the asset, click **Start backup**.

Chapter 6

Protecting Kubernetes assets

This chapter includes the following topics:

- Kubernetes protection plans
- Configuring backup options for Kubernetes protection plans

Kubernetes protection plans

Like other NetBackup workloads, you need to create protection plans to protect Kubernetes workloads. Kubernetes protection plans:

- Does not require any storage to be specified in the protection plan.
- Supports only full-backup schedule.

Configuring backup options for Kubernetes protection plans

Kubernetes protection plans allow you to distinguish partially successful backups, and retain or discard them as required. A partially successful backup may not have successfully backed up all the resources that you intend to backup. You can decide whether to keep such backups or discard them, and make this specification for each protection plan separately.

See the *Managing protection plans* section of the *NetBackup Web UI Administrator's Guide*, for details of how to create a protection plan.

To configure the backup options while configuring a protection plan for Kubernetes, in the **Backup options** page, select the option **Fail a backup job, if any of the**

resources fail to get protected. This setting discards any partially-successful backup jobs.

Recovering Kubernetes assets

This chapter includes the following topics:

- Recovering Kubernetes assets

Recovering Kubernetes assets

Using NetBackup you can recover Kubernetes namespaces and persistent volumes.

Note: In NetBackup 9.1, exclusive recovery of persistent volumes is only supported for the Velero plug-in for Google Cloud Platform (GCP).

Note: After recovery, the newly created namespaces, Persistent volumes, and other resources get new system-generated UIDs.

To recover namespaces

1. On the left, click **Kubernetes**, under **Workloads**.
2. In the **Namespaces** tab, click the namespace of the asset that you want to recover. Click the **Recovery points** tab.
3. The **Recovery points** tab shows you all the recovery points with the date and time of the backup. You can set filters to filter the displayed recovery points. Click the date in the **Date** column, to view the details of the recovery point. The **Recovery points details** dialog shows the resources that were backed up, like ConfigMaps, namespaces, secrets, persistent volumes, pod, and so on. For details about these resources, see https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/.

4. Click the ellipsis menu (three dots), in the row of the recovery point that you want to recover. To recover a namespace, click **Restore namespace**.

5. In the **Recovery target** page, to recover the asset to the same source cluster, click **Next**. To recover to an alternate cluster, click **Select cluster**. In the **Select cluster** dialog, select the target cluster, and click **Select**. Click **Next**.

   Note: If you select the target cluster to be different than the original cluster, the object storage used by the Velero plug-ins on both the clusters must be the same.

6. In the **Recovery options** page, do the following:

   - To recover to the original namespace, select **Use original namespace**. To restore the asset to an alternate namespace, select **Use alternate namespace**, and enter the new namespace name. The name should follow Kubernetes specification for namespace names.
   - To allow restore even if a same namespace already exists, select **Proceed with restore if namespace already exists**. This option helps you to restore any missing resources in an already existing namespace. If any resource is missing in an asset inside the cluster, and the same resource exists in the backup copy, you can use this option to restore the missing resources in the asset. This option does not overwrite any of the existing resources in the asset, only restores missing ones.
   - To restore all resources within the asset, select **Recover all resources**. To restore selected resource types of the asset, choose **Select resource types**, and select the resource types that you want to recover. Note that you cannot restore individual resources or instances, you need to select the types of the resources, that may contain any number of individual resources or instances.

     Note: The **Select resource types** option is for advanced users. If you are not careful in selecting the resources that you want to restore, you may not get a fully functional namespace after the restoration.

7. Click **Next**.

8. In the **Recovery overview** page, review all the recovery option that you have selected. To go back and change any of the settings, click **Previous**. Once all the parameters are changed, click **Start recovery**.

To recover persistent volumes

1. Perform the steps 1-3 of the previous procedure.
2. Click the ellipsis menu (three dots), in the row of the recovery point that you want to recover. To recover persistent volumes, click **Restore persistent volumes**.
3. In the **Recovery target** page, to recover the persistent volume to the same source cluster, click **Next**. To recover to an alternate cluster, click **Select cluster**. In the **Select cluster** dialog, select the target cluster, and click **Select**. Click **Next**.

   Note: If you select the target cluster to be different than the original cluster, the object storage used by the Velero plug-ins on both the clusters must be the same.

4. In the **Recovery options** page, do any of the following:
   - To recover to the original namespace, select **Use original namespace**. To allow restore even if a same namespace already exists, select **Proceed with restore if namespace already exists**. This option helps you to restore any missing persistent volumes in an already existing namespace. If any persistent volume is missing in an asset inside the cluster, and the same persistent volume exists in the backup copy, you can use this option to restore the missing persistent volume in the asset. This option does not overwrite any of the existing persistent volumes in the asset, only restores the missing ones.
   - Select **Use temporary system generated namespace**, to use a unique system-generated namespace. The namespace is deleted after the restore operation is complete. Use this option when your priority is to restore persistent volumes data, instead of namespaces.
5. Click **Next**.
6. In the **Recovery overview** page, review all the recovery option that you have selected. To correct any of the settings, click **Edit** against that option, or click **Previous**. Once all the parameters are correct, click **Start recovery**.

Chapter 8

Troubleshooting Kubernetes issues

This chapter includes the following topics:

- Connecting to the primary server with a short host name
- Cluster discovery fails
- Error during backup: Namespace has been marked for deletion
- Error during restore: Final job status is partially failed
- Backup stuck in InProgress state
- Restore stuck in InProgress state

Connecting to the primary server with a short host name

The NetBackup primary server must be reachable from the NetBackup Kubernetes operator. NetBackup primary server name can be FQDN or short host name, it should be resolvable at NetBackup Kubernetes operator.

You can use hostAliases in the Velero controller manager deployment to make short host name-based NetBackup primary server reachable from NetBackup Kubernetes operator.

```
hostAliases:
- hostnames:
  - falcon
    ip: 10.x.x.x
```

Cluster discovery fails

Discovery reconciler (NetBackup Kubernetes operator) posts asset data to NetBackup. To perform this task, the operator requires an API key and caCert in a secret file.

Recommended action:

- Ensure that the secret with same name as NetBackup primary server name is present in the namespace where the NetBackup Kubernetes operator is deployed.
- Make sure that the secret file exists in the same namespace as the NetBackup Kubernetes operator pod.
- Make sure that the API key and CA certificate are valid.

Error during backup: Namespace has been marked for deletion

This error appears when the namespace that you tried to backup was deleted from the Kubernetes cluster. NetBackup asset service also deletes that asset from the asset database. But if there was a backup available for that namespace then the asset service does not delete the asset but rather marks it as DELETED. In this case, we don't want to take the backup of such assets or namespaces.

Recommended action: Verify if the namespace exists on the cluster.

Error during restore: Final job status is partially failed

Final job status is partially failed with few warnings specific to the resource RoleBinding.

These warnings are specific to the resource RoleBinding for API group authorization. `openshift.io` and `rbac.authorization.Kubernetes.io` are thrown because these RoleBinding are auto managed by the controller and gets created when we create a new Namespace.

Recommended action: You can avoid these warnings by excluding the relevant RoleBinding resources from the restore.

Backup stuck in InProgress state

Scenario 1: Happens if a Velero pod running in a Kubernetes cluster, restarts when a backup job is running.

Recommended Action:

Follow these steps in Velero documentation, to identify if the job is stuck or running slowly. See Velero documentation. Delete the NetBackup backup ("backups.netbackup.veritas.com") CRD job and the Velero backup ("backups.velero.io") CRD job.

Scenario 2: Snapshots are in *UploadFailed* state or uploads fail, which may result in retry of the upload job.

Recommended Action:

Verify the reason for the issue by looking at the `datamanager` logs (on the processing node) and the backup driver logs to identify the root cause of the upload errors. To enable NetBackup jobs, delete the NetBackup backup CR ("backups.netbackup.veritas.com"), this marks the backup jobs as failed. Also clean the corresponding velero backup job ("backups.velero.io"), snapshot job ("snapshots.backupdriver.cnsdp.vmware.com") and upload jobs ("uploads.datamover.cnsdp.vmware.com").

Restore stuck in InProgress state

The issue occurs if the Velero pod running in Kubernetes cluster, restarts while the restore job is running.

Recommended Action:

For Velero plugin for vSphere,

If the restore contains persistent volumes, verify the state of CloneFromSnapshot ("clonefromsnapshots.backupdriver.cnsdp.vmware.com") CRD corresponding to the restore job. If the request for downloads ("downloads.datamover.cnsdp.vmware.com") fails, there may be an issue with the object store access or the namespace to which the restore is targeted. The namespace may already contain an existing persistent volume claim with the same name. Delete the NetBackup restores ("restores.netbackup.veritas.com") CRD job and Velero restores ("restores.velero.io") CRD job along with the corresponding clonefromsnapshots CRD if applicable. Deleting corresponding data mover download ("downloads.datamover.cnsdp.vmware.com") CRD is optional.

For Velero plugin for GCP:

Delete the NetBackup restores ("restores.netbackup.veritas.com") CRD job and Velero restores ("restores.velero.io") CRD job. Optionally, delete the corresponding download ("downloadrequests.velero.io") CRD.