Hitachi Replication Plug-in for Containers

Version 1.1.0

Configuration Guide

This document provides an implementation overview and describes the usage requirements, installation, and configuration of Hitachi Replication Plug-in for Containers.

MK-92ADPTR155-01
March 2022

Contents

Preface

This document provides information and instructions to help you to setup and use Hitachi Replication Plug-in for Containers with Hitachi Storage Plug-in for Containers.

Intended audience

This document is intended for system administrators, Hitachi Vantara representatives, and authorized service providers who install, configure, and run Hitachi Replication Plug-in for Containers with Hitachi Storage Plug-in for Containers.

Readers of this document should be familiar with the following:

- Hitachi Virtual Storage Platform family
- Kubernetes and Red Hat OpenShift Platform

Document conventions

This document uses the following typographic conventions:

Convention	Description
Bold	▪ Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: Click **OK**. ▪ Indicates emphasized words in list items.
Italic	▪ Indicates a document title or emphasized words in text. ▪ Indicates a variable, which is a placeholder for actual text provided by the user or for output by the system. Example: `pairdisplay -g` *`group`* (For exceptions to this convention for variables, see the entry for angle brackets.)
`Monospace`	Indicates text that is displayed on screen or entered by the user. Example: `pairdisplay -g oradb`

Convention	Description
< > angle brackets	Indicates variables in the following scenarios: ▪ Variables are not clearly separated from the surrounding text or from other variables. Example: Status-<*report-name*><*file-version*>.csv ▪ Variables in headings.
[ ] square brackets	Indicates optional values. Example: [a | b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a | b } indicates that you must choose either a or b.
| vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a | b] indicates that you can choose a, b, or nothing. { a | b } indicates that you must choose either a or b.

This document uses the following icons to draw attention to information:

Icon	Label	Description
	Note	Calls attention to important or additional information.
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Caution	Warns the user of adverse conditions and/or consequences (for example, disruptive operations, data loss, or a system crash).
	WARNING	Warns the user of a hazardous situation which, if not avoided, could result in death or serious injury.

Conventions for storage capacity values

Physical storage capacity values (for example, disk drive capacity) are calculated based on the following values:

Physical capacity unit	Value
1 kilobyte (KB)	1,000 (10^3) bytes
1 megabyte (MB)	1,000 KB or $1,000^2$ bytes
1 gigabyte (GB)	1,000 MB or $1,000^3$ bytes
1 terabyte (TB)	1,000 GB or $1,000^4$ bytes
1 petabyte (PB)	1,000 TB or $1,000^5$ bytes
1 exabyte (EB)	1,000 PB or $1,000^6$ bytes

Logical capacity values (for example, logical device capacity, cache memory capacity) are calculated based on the following values:

Logical capacity unit	Value
1 block	512 bytes
1 cylinder	Mainframe: 870 KB Open-systems: ▪ OPEN-V: 960 KB ▪ Others: 720 KB
1 KB	1,024 (2^{10}) bytes
1 MB	1,024 KB or $1,024^2$ bytes
1 GB	1,024 MB or $1,024^3$ bytes
1 TB	1,024 GB or $1,024^4$ bytes
1 PB	1,024 TB or $1,024^5$ bytes
1 EB	1,024 PB or $1,024^6$ bytes

Accessing product documentation

Product user documentation is available on the Hitachi Vantara Support Website: https://knowledge.hitachivantara.com/Documents. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

The Hitachi Vantara Support Website is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to the Hitachi Vantara Support Website for contact information: https://support.hitachivantara.com/en_us/contact-us.html.

Hitachi Vantara Community is a global online community for Hitachi Vantara customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to community.hitachivantara.com, register, and complete your profile.

Comments

Please send us your comments on this document to doc.comments@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara LLC.

Thank you!

Chapter 1: Overview

Hitachi Replication Plug-in for Containers automates storage replication between two different Kubernetes clusters and storage systems located at different sites.

Traditional storage system replication deployments require tight communication among storage administrators, container users, and Kubernetes administrators. With Hitachi Replication Plug-in for Containers, container users and Kubernetes administrators can take a self-service approach to create replications using the Kubernetes command-line tool, **`kubectl`**.

The following diagram describes an overview of creating the replications with Hitachi Replication Plug-in for Containers and Hitachi Storage Plug-in for Containers.

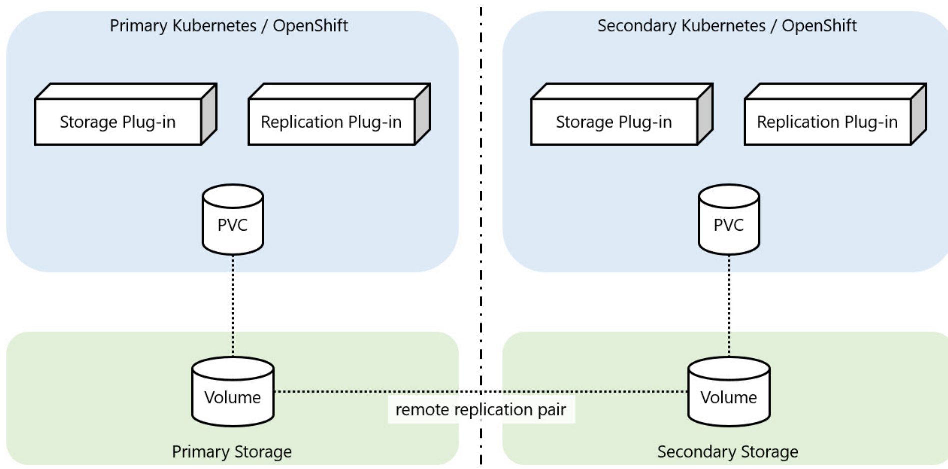

Pre-installation tasks

Before you install Hitachi Replication Plug-in for Containers, complete the following pre-installation requirements.

Components	Requirements
Kubernetes cluster	Verify that there are different Kubernetes clusters in the primary and secondary sites. A single cluster across the primary and secondary sites is not supported.
Replication target	Create the PersistentVolumeClaim file (PVC) in Storage Plug-in for Containers.
etcd	Verify that the etcd is not replicated by Hitachi Replication Plug-in for Containers.
Kubernetes cluster management	Verify that the declarative cluster is supported and imperative cluster is not supported.
The number of Kubernetes clusters	Primary: One Kubernetes cluster Secondary: One Kubernetes cluster
The number of storage systems	Primary: One storage system Secondary: One storage system
The number of replicas created from a PVC	One (Do not create any remote copy pair manually with the PVC that will be used for Hitachi Replication Plug-in for Containers.)
The number of Pods created from a PVC	One If the PVC is controlled by Hitachi Replication Plug-in for Containers, you can create only one Pod per PVC.
Used program product	Hitachi Universal Replicator (HUR).
Application settings (Applications using PVs created by Hitachi Storage Plug-in for Containers)	The application must be crash-consistent and shared-nothing because the data is replicated asynchronously.
Application redundancy (Applications using PVs created by Hitachi Storage Plug-in for Containers)	The application redundancy is Active-Passive. The secondary site is activated only when a failure occurs in the primary site.

Components	Requirements
Inter-site connection	Requires all connections listed below to setup replications: ▪ IP connection between Hitachi Replication Plug-in for Containers and Kubernetes. Hitachi Replication Plug-in for Containers in the primary site must communicate with the Kubernetes in the secondary site and vice versa. ▪ IP connection between Hitachi Replication Plug-in for Containers and the storage system. Hitachi Replication Plug-in for Containers in the primary site must communicate with the storage system in the secondary site and vice versa. ▪ IP connection between primary and secondary storage system REST API servers for management. ▪ FC or iSCSI connection between primary and secondary storage systems for data copy.
Reverse-resync	Reverse-resync is not supported. When performing failback from the secondary to primary site, delete the existing Hitachi Universal Replicator (HUR) pair first, and then create new HUR pair.
Snapshot or Clone	Do not take a snapshot or clone from a PVC that is controlled by Hitachi Replication Plug-in for Containers.

Supported software versions

<table>
<tr><th>Software</th><th>Supported Version</th></tr>
<tr><td>Hitachi Storage Plug-in for Containers</td><td>v3.9.0</td></tr>
<tr><td>Kubernetes and Red Hat OpenShift Platform</td><td>See Hitachi Storage Plug-in for Containers Release Notes.</td></tr>
<tr><td>Storage systems</td><td>For the list for the supported storage systems and versions, refer to the following documents:
Hitachi Storage Plug in for Containers Release Notes.
Hitachi Universal Replicator User Guide
Note: Hitachi Configuration Manager is not supported.</td></tr>
</table>

Chapter 2: Installing and configuring Hitachi Replication Plug-in for Containers

The following is the workflow for configuring and using Hitachi Replication Plug-in for Containers

1. Configuring the storage system (on page 13)
2. Configurations in Storage Plug-in for Containers (on page 13)
3. Creating the PersistentVolumeClaim file (on page 14)
4. Configuring Hitachi Replication Plug-in for Containers (on page 14)

Note: If you want to use the resource partitioning feature of Hitachi Storage Plug-in for Containers, see Using resource partitioning feature of Hitachi Storage Plug-in for Containers (on page 33).

Configure the storage system

Configure the storage system for replications.

1. Configure the storage system as described in the *Hitachi Storage Plug-in for Containers Quick Reference Guide*.
2. Configure the remote path between primary site and secondary site storage systems. For details, see *Hitachi Universal Replicator (HUR) User Guide*.
3. Configure journal volumes. For details, see *Hitachi Universal Replicator (HUR) User Guide.*

Configure the Storage Plug-in for Containers

Before you begin

Verify that you have the Quick Reference Guide of the supported version of Storage Plug-in for Containers

Procedure

1. Install Storage Plug-in for Containers in both the primary and secondary site. For details, see the *Quick Reference Guide of Storage Plug-in for Containers*.
2. Create StorageClass in both the primary and secondary sites.

Note:

- Name and fstype of the StorageClass must be the same for both sites.
- StorageClass in the primary site must point to the storage in the primary site.
- StorageClass in the secondary site must point to the storage in the secondary site.

Create the PersistentVolumeClaim file

Hitachi Replication Plug-in for Containers replicates a PersistentVolumeClaim file created in this section.

Procedure

1. Create a PVC with the StorageClass created in Configuration in Storage Plugin for Containers (on page 13). For details, see *Storage Plugin for Containers Quick Reference Guide*.
2. (Optional) Create a Pod that consumes the PVC. The PVC can be replicated even if it is attached to a Pod.

   Hitachi Replication Plug-in for Containers uses the following parameters of the PVC to create a PVC in the secondary site.

   - name
   - namespace
   - spec.accessModes
   - spec.resources
   - spec.storageClassName
   - spec.volumeMode

Configure the Hitachi Replication Plug-in for Containers

Before you begin

You must have a dedicated management machine that can access both the primary and secondary Kubernetes clusters.

Procedure

1. Get the kubeconfig file from both the primary and secondary sites.

```
# kubeconfig file paths for the primary and secondary site.
% KUBECONFIG_P=/path/to/primary-kubeconfig
% KUBECONFIG_S=/path/to/secondary-kubeconfig
```

```
# Secret file paths for the primary and secondary site
# These files are supposed to be created in the later steps.
% SECRET_KUBECONFIG_P=/path/to/primary-kubeconfig-secret.yaml
% SECRET_KUBECONFIG_S=/path/to/secondary-kubeconfig-secret.yaml
```

2. Configure an environment variable for the Secret file of the storage system.

```
# A Secret file for both storages.
# This file is supposed to be created in the later steps.
% SECRET_STORAGE=/path/to/storage-secret.yaml
```

3. Copy the Namespace manifest file to the management machine. This file is provided in the media kit. Do not edit it.

```
% vi hspc-replication-operator-namespace.yaml
apiVersion: v1
kind: Namespace
metadata:
  labels:
    control-plane: controller-manager
  name: hspc-replication-operator-system
```

4. Create a Secret manifest file with the secondary `kubeconfig` information to access the secondary Kubernetes cluster from Hitachi Replication Plug-in for Containers running in the primary Kubernetes cluster. For reference, see the `remote-kubeconfig-sample.yaml` file.

```
# base64 encoding
% cat ${KUBECONFIG_S} | base64 -w 0

% vi ${SECRET_KUBECONFIG_S}
apiVersion: v1
kind: Secret
metadata:
  name: hspc-replication-operator-remote-kubeconfig
  namespace: hspc-replication-operator-system
type: Opaque
data:
  remote-kubeconfig: <base64 encoded secondary kubeconfig>
```

5. Create a Secret manifest file the with the primary `kubeconfig` information to access the primary Kubernetes cluster from Hitachi Replication Plug-in for Containers running in the secondary Kubernetes cluster. For reference, see the `remote-kubeconfig-sample.yaml` file.

```
# base64 encoding
% cat ${KUBECONFIG_P} | base64 -w 0

% vi ${SECRET_KUBECONFIG_P}
apiVersion: v1
kind: Secret
```

```
metadata:
  name: hspc-replication-operator-remote-kubeconfig
  namespace: hspc-replication-operator-system
type: Opaque
data:
  remote-kubeconfig: <base64 encoded primary kubeconfig>
```

6. Create a Secret manifest file containing storage system information that enables access by Hitachi Replication Plug-in for Containers. For reference, see the `storage-secrets-sample.yaml` file.

```
% vi ${SECRET_STORAGE}
apiVersion: v1
kind: Secret
metadata:
  name: hspc-replication-operator-storage-secrets
  namespace: hspc-replication-operator-system
type: Opaque
stringData:
  storage-secrets.yaml: |-
    storages:
    - serial: 54321                 # (1)
      url: http://172.16.1.1        # (2)
      user: UserPrimary             # (3)
      password: PasswordPrimary     # (4)
      journal: 1                    # (5)
    - serial: 400100                # (6)
      url: http://172.16.1.2        # (7)
      user: UserSecondary           # (8)
      password: PasswordSecondary   # (9)
      journal: 2                    # (10)
```

This file includes information on the storage system that is used by the primary and secondary sites.

The following describes the parameters in the above example:

- (1) and (6) Storage system serial number for primary and secondary site.
- (2) and (7) URL for the REST API server.
- (3) and (8) Username for the primary and secondary storage system. The user must belong to the built-in storage administrator (View & Modify) user group (the same as Hitachi Storage Plug-in for Containers).
- (4)and (9) Password for the username.
- (5) and (10) Journal ID for HUR.

7. Modify the Hitachi Replication Plug-in for Containers manifest file provided in the media kit based on your requirement to use your private repository.

```
% vi hspc-replication-operator.yaml
…
```

```
  image: __REGISTRY__/hitachi/hspc-replication-operator:v<HRPC version>
…
```

8. Create Namespaces in the primary and secondary sites. Use the same manifest file in primary and secondary sites.

```
% KUBECONFIG=${KUBECONFIG_P} kubectl create -f hspc-replication-operator-namespace.yaml
% KUBECONFIG=${KUBECONFIG_S} kubectl create -f hspc-replication-operator-namespace.yaml
```

9. Create Secrets containing `kubeconfig` information in primary and secondary sites. Use the different manifest files in primary and secondary sites.

```
% KUBECONFIG=${KUBECONFIG_P} kubectl create -f ${SECRET_KUBECONFIG_S}
% KUBECONFIG=${KUBECONFIG_S} kubectl create -f ${SECRET_KUBECONFIG_P}
```

10. Create Secrets containing storage system information in primary and secondary sites. Use the same manifest file in primary and secondary sites.

```
% KUBECONFIG=${KUBECONFIG_P} kubectl create -f ${SECRET_STORAGE}
% KUBECONFIG=${KUBECONFIG_S} kubectl create -f ${SECRET_STORAGE}
```

11. Load the container (For example, docker load) `hrpc_1_1_0.tar` and push the loaded container to your private repository.
12. Create Hitachi Replication Plug-in for Containers in primary and secondary sites. Use the same manifest file for both the primary and secondary sites.

```
% KUBECONFIG=${KUBECONFIG_P} kubectl create -f hspc-replication-operator.yaml
% KUBECONFIG=${KUBECONFIG_S} kubectl create -f hspc-replication-operator.yaml
```

> **Note:** If you find any misconfigurations in Secrets, delete the replication controller first, and then modify the Secrets. The replication controller reads the Secrets when it starts.

13. Confirm that Hitachi Replication Plug-in for Containers containers are running in primary and secondary sites.

```
% KUBECONFIG=${KUBECONFIG_P} kubectl get pods -n hspc-replication-operator-system
% KUBECONFIG=${KUBECONFIG_S} kubectl get pods -n hspc-replication-operator-system
```

Chapter 3: Replicating storage volumes

This chapter describes how to replicate storage volumes by creating a Replication CustomResource (CR). Hitachi Replication Plug-in for Containers starts replicating the specified PVC and triggers the creation of a PVC in the secondary site.The data in the target PVC in the primary site is copied and protected by HUR.

Creating a mainfest file for Replication CR

Create the manifest file for a Replication CR to initiate the replication of the specified PVC.

Procedure

1. Specify a protection target `persistentVolumeClaimName` created in Create the PersistentVolumeClaim file (on page 14). This PVC is protected by HUR and must be created before creating a Replication CR.
2. Specify a `StorageClassName` from which the target PVC is created.

```
% vi hspc_v1_replication.yaml
apiVersion: hspc.hitachi.com/v1
kind: Replication
metadata:
  name: replication-sample
spec:
  persistentVolumeClaimName: pvc-xxxx    # (1)
  storageClassName: sc-hspc-dr           # (2)
```

> **Note:** After creating this manifest file, you cannot edit the parameters except for the `desiredPairState`. If you want to change the parameters, delete the existing file, and then re-create it with the new parameters.

Creating a Replication CR object

Create a Replication CR object with the Replication CR manifest file in the primary site. This triggers creation of an HUR pair and initial data copy.

For more details, see Details on HUR pair creation (on page 31).

```
% KUBECONFIG=${KUBECONFIG_P} kubectl create -f hspc_v1_replication.yaml
```

Note: When Hitachi Replication Plug-in for Containers detects a newly created Replication CR, it automatically creates a corresponding Replication CR and PVC in the secondary site. Hitachi Replication Plug-in for Containers considers the site in which the Replication CR is created as the primary site.

Verifying the status of Replication CR

The Replication CR status is changed to `Ready` when the initial replication is created, and the data protection is started. Transition to the `Ready` status depends on data size in target PVC and might take some time to change.

```
% KUBECONFIG=${KUBECONFIG_P} kubectl get replication
NAME                STATUS    DESIREDSTATE   OPERATION      AGE
replication-sample  Ready      none              none           12m

% KUBECONFIG=${KUBECONFIG_S} kubectl get replication
NAME                STATUS    DESIREDSTATE   OPERATION      AGE
replication-sample  Ready      none              none           12m
```

If you miscreated the Replication CR, delete the Replication CR and verify if a PVC with the specified name exists in the secondary site. If it exists, delete the PVC.

```
% KUBECONFIG=${KUBECONFIG_P} kubectl delete -f hspc_v1_replication.yaml
% KUBECONFIG=${KUBECONFIG_S} kubectl delete pvc <pvc name>
```

Chapter 4: Checking the data from the secondary site

After the Replication CR is created and is in Ready status, you can perform the split and resync operation to check the data from the secondary site. The data copy process from the primary site to the secondary site is stopped during the split and resync operation.

Splitting the Hitachi Universal Replicator pair

From the primary site, change the status of Replication CR to perform the split operation. This triggers Hitachi Replication Plug-in for Containers to split the HUR pair.

Procedure

1. Confirm that Replication CR status is **Ready**, and the Operation value is `none` .
2. Edit the Replication CR and change `spec.desiredPairState` to `split`.

```
% KUBECONFIG=${KUBECONFIG_P} kubectl edit replication
<replication_name>
# Before
…
spec:
  desiredPairState: <any state>
…

#After
…
spec:
  desiredPairState: split
…
```

3. Confirm that the Replication CR status is `Split` and the Operation value is `none`.

```
% KUBECONFIG=${KUBECONFIG_P} kubectl get replication
NAME                  STATUS   DESIREDSTATE   OPERATION   AGE
replication-sample    Split    split          none        6d20h
```

Confirming the replicated data

From the secondary site, deploy a Pod and mount the secondary volume to confirm the replicated data.

Procedure

1. Create a Pod that use the PVC of the secondary site to check the data.
2. Delete the Pod after the data check is complete.

Resynchronizing HUR pair

From the primary site, change the status of Replication CR to perform the `resync` opeartion. This triggers Hitachi Replication Plug-in for Containers to resync the HUR pair.

Note: You cannot perform the resync operation if a Pod mounts the secondary volume in the secondary site.

Procedure

1. Confirm that the Replication CR status is `Split` and the Operation value is `none`.
2. Edit the Replication CR and change `spec.desiredPairState` to `pair`.

```
% KUBECONFIG=${KUBECONFIG_P} kubectl edit replication
<replication_name>
# Before
…
spec:
  desiredPairState: split
…

#After
…
spec:
  desiredPairState: pair
…
```

3. Confirm that the Replication CR status is `Ready` and Operation value is `none`.

```
% KUBECONFIG=${KUBECONFIG_P} kubectl get replication
NAME                  STATUS   DESIREDSTATE   OPERATION   AGE
replication-sample    Ready    none           none        6d20h
```

Chapter 5: Failover and failback

When failures occur in the primary site, you can move the applications running in the primary site to the secondary site.

Performing failover to the secondary site

From the secondary site, change the status of Replication CR to perform failover operation. This triggers Hitachi Replication Plug-in for Containers to failover HUR pair.

Note: Do not modify Replication in the primary site during this operation.

Procedure

1. Confirm that the Replication CR status is `Ready` and operation value is `none`.
2. Edit the Replication CR and change `spec.desiredPairState` to `failover`.

```
% KUBECONFIG=${KUBECONFIG_S} kubectl edit replication
<replication_name>
# Before
…
spec:
  desiredPairState: <any state>
…

#After
…
spec:
  desiredPairState: failover
…
```

3. Confirm that Replication CR status is `Failover` and the operation value is `none`.

```
% KUBECONFIG=${KUBECONFIG_S} kubectl get replication
NAME                     STATUS    DESIREDSTATE    OPERATION    AGE
replication-sample       Failover     failover          none         6d20h
```

Running your applications in the secondary site

From the secondary site, deploy your applications.

Procedure

1. Prepare the manifest for the application using PVCs created by Hitachi Replication Plug-in for Containers. The Replication CR does not need to be created manually here. The Replication CR in the secondary site is automatically created when the Replication CR in the primary site is created.
2. Deploy the applications by using the manifests and confirm that the applications are running properly.

Next steps

After failover, proceed to either of the following chapters. If you plan to recover your primary site, proceed to Recovery and clean up of the primary site (on page 23) and Failback to the primary site (on page 24). If you plan not to recover your primary site, proceed to Running applications without primary site recovery (on page 22).

Recovery and cleanup of the primary site

Before you begin

Reset your Kubernetes cluster in the primary site before connecting to the network. If the Replication CRs remain in the cluster, an unexpected HUR pairs might be created automatically, and it might overwrite existing data being consumed in the secondary site.

Procedure

1. Initialize the Kubernetes cluster in the primary site before connecting it to the network.
2. Deploy the Hitachi Replication Plug-in for Containers in the primary site. For details, see Hitachi Replicator controller (on page 14).
3. Configure the Hitachi Replication Plug-in for Containers in the secondary site with the new `kubeconfig` of the primary site. For details, see Hitachi Replicator controller (on page 14).
4. Restart the Hitachi Replication Plug-in for Containers so that the new configuration is ready. For details, see Restarting Hitachi Replication Plug-in for Containers (on page 29).
5. Verify the storage system assets in the primary site. To identify the volume ID of the storage system in the primary site. For details, see Storage information of Replication CR (on page 30).
6. Delete the Replication CR in the secondary site.
7. Using Storage Navigator, delete the storage system assets you found in step 5 from the primary site.

Failback to the primary site

Create the Replication CR in the primary site after the recovery and cleanup of the primary site is completed.

Procedure

1. Create the Replication CR in the secondary site.
2. From the primary site, perform the following steps to copy data from the secondary site to the primary site.
   - Confirm that the Replication CR status is **Ready** and operation value is none.
   - Edit the Replication CR and change `spec.desiredPairState` to failover.
   - Confirm that Replication CR status is **Failover** and the operation value is none.
3. Delete the Replication CR and the PVC from the secondary site.
4. Create the Replication CR in the primary site.

Running applications without primary site recovery

After the failover, if you do not recover the primary site, refer to Deleting replication CRs (on page 29) to delete Replication CR and the HUR pair.

Deleting replication CR and the related objects

This chapter describes how to cleanup the Replication CR and the related objects.

> **Note:**
>
> - Deleting Replication CR and its related objects and storage system assets require accesses to Kubernetes clusters and storage systems in both the primary and secondary sites.
> - If your Kubernetes clusters and storage system are in a remote site (primary or secondary depending on where you are operating from) are not accessible, you can delete Replication CRs forcibly. For details, see Deleting replication CRs (on page 29).

Procedure

1. Delete the Replication CR in the primary site. This also deletes the Replication CR in the secondary site.

```
% KUBECONFIG=${KUBECONFIG_P} kubectl delete -f hspc_v1_replication.yaml
```

2. Confirm that the Replication CR has been deleted in the primary and secondary sites.

```
% KUBECONFIG=${KUBECONFIG_P} kubectl get replication
% KUBECONFIG=${KUBECONFIG_S} kubectl get replication
```

3. Delete any remaining PVCs from both the primary and secondary sites.

```
% KUBECONFIG=${KUBECONFIG_P} kubectl delete pod <pod name>
% KUBECONFIG=${KUBECONFIG_P} kubectl delete pvc <pvc name>
% KUBECONFIG=${KUBECONFIG_S} kubectl delete pod <pod name>
% KUBECONFIG=${KUBECONFIG_S} kubectl delete pvc <pvc name>
```

Chapter 6: Uninstalling Hitachi Replication Plug-in for Containers.

This chapter describes how to uninstall Hitachi Replication Plug-in for Containers

Procedure

1. Delete all the Replication CRs from the primary site.

```
% kubectl delete -n <namespace> replication <Replication name>
```

2. Confirm that all the Replication CRs are deleted from both the primary and secondary sites.
3. Delete Hitachi Replication Plug-in for Containers by deleting all the resources created in Configure the Hitachi Replication Plug-in for Containers (on page 14).
4. Uninstall Hitachi Storage Plug-in for Containers. For details, see *Hitachi Storage Plug-in for Containers Quick Reference Guide*.

Chapter 7: Troubleshooting

This chapter describes the error messages of Hitachi Replication Plug-in for Containers.

> **Note:** If you cannot create or delete PVCs or Pods using Storage Plug-in for Containers, see Troubleshooting chapter in the *Hitachi Storage Plug-in for Containers Quick Reference Guide*.

Troubleshooting Replication CR issues

Make sure you have read through this document and understood how Replication CRs are created. For more details, see Details on HUR pair creation (on page 31).

Creating a Replication CR triggers Hitachi Replication Plug-in for Containers to create a LocalVolume and a RemoteVolume CR. These CRs are used to gather information about the primary and secondary sites storage volumes.

When the status of Replication CR does not change to an expected status, then run the `kubectl describe` command for the event log.

If an error message is shown in the newest event log, follow the instructions of the error message to fix the problem. Also, check the event logs for the LocalVolume and RemoteVolume CRs in the primary and secondary sites.

Collecting information for troubleshooting

Use the logs and suggested procedures to troubleshoot issues with Hitachi Replication Plug-in for Containers.

Collecting logs for Hitachi Replication Plug-in for Containers

You can retrieve the logs for Hitachi Replication Plug-in for Containers by using `kubectl logs` command.

```
% kubectl logs -n hspc-replication-operator-system deploy/hspc-replication-operator-controller-manager
```

> **Note:** If necessary, set up cluster-level logging to save logs: https://kubernetes.io/docs/concepts/cluster-administration/logging/

To view recent events for Replication CR, run the following command:

```
% kubectl describe -n <namespace> <crd type> <cr name>
```

Note:

- There are three types of CRD: Replication, LocalVolume, and RemoteVolume.
- The name of these resources is the same as the Replication in the primary site.

Collecting storage system information for VSP family

If you are using an SVP, collect the regular dump files.

If you are not using an SVP, collect system dumps by using the maintenance utility. For details about how to collect the dump files of storage systems, see the Hitachi Virtual Storage Platform System Administrator Guide.

Information needed in case of contacting support

If you must contact customer support, collect the following information before contacting customer support:

Information	Procedure
Cluster information	Run the following command: `% kubectl cluster-info dump -A > dump.txt`
Hitachi Replication Plug-in for Containers logs	See Collecting logs for Hitachi Replication Plug-in for Containers
Hitachi Replication Plug-in for Containers manifests	Manifest files for Hitachi Replication Plug-in for Containers
Replication manifests	Get yaml files for Replication
PVC related manifests	Get YAML file for StorageClass, Secret, and PersistentVolumeClaim
Hitachi Storage Plug-in for Containers logs	See Hitachi Storage Plug-in for Containers Quick Reference Guide
Application manifests	Get the YAML files for applications that uses Storage Plug-in for Containers PVC
Storage logs	See Collect storage system information for VSP family (on page 28).

Restarting Hitachi Replication Plug-in for Containers

When you make configuration changes, you must restart the Hitachi Replication Plug-in for Containers. Use the following Kubernetes commands to restart the Hitachi Replication Plug-in for Containers:

```
# Restarting HRPC in the primary site
% KUBECONFIG=${KUBECONFIG_P} kubectl rollout restart deployment \
 -n hspc-replication-operator-system hspc-replication-operator-
controller-manager
# Restarting HRPC in the secondary site
% KUBECONFIG=${KUBECONFIG_S} kubectl rollout restart deployment \
 -n hspc-replication-operator-system hspc-replication-operator-
controller-manager
```

Deleting replication CRs without connection to a remote site

Deleting replication CRs requires connection to a remote site to gracefully delete all Kubernetes objects and storage system assets. However, there may be a situation where your remote site is not accessible, but you must delete the Replication CRs. In such cases, you can delete the Replication CRs forcibly.

Procedure

1. To enable Hitachi Replication Plug-in for Containers and to delete Replication CRs without connecting to a remote site, set SPC_REPLICATION_ALLOW_FORCE_DELETION environment variable to `true`.

   ```
   % kubectl set env -n hspc-replication-operator-system \
    deployments/hspc-replication-operator-controller-manager \
   SPC_REPLICATION_ALLOW_FORCE_DELETION=true
   ```

2. Run the following command to delete the Replication CRs forcibly.

   ```
   % kubectl delete -f hspc_v1_replication.yaml
   ```

3. Using Storage Navigator delete the HUR pair corresponding with the Replication CR since it remains after force deletion of Replication CRs.
4. Reset the SPC_REPLICATION_ALLOW_FORCE_DELETION environment variable to `false`.

Chapter 8: Appendix

Replication CR Status

The following table describes the status of the Replication CR.

Status	Description	HUR pair state
Pending	Accepted CR creation command.	NA
WaitingPrimarySite	Waiting for primary site to finish HUR operations. This only apply to secondary site.	NA
Copying	Initial copy is in progress.	COPY
Ready	Successfully created HUR pair and the state is “PAIR”	PAIR
Split	Successfully splitted HUR pair and the state is “PSUS” or “SSUS”	PSUS or SSUS
Failover	Successfully failover HUR pair.	SSWS
Error	Having errors in HUR pair.	PSUE, PFUL, or PFUS
Unknown	A pair does not exist or unexpected pair state.	Other than the above

Storage information of Replication CR

Information of the HUR pair is stored in the Replication CR under `.status.replicationPair` key. This information can be retrieved by `kubectl describe` command.

The following table lists the keys and descriptions.

Note: The local or remote depends on which site you executed `kubectl describe` from.

Key	Description
localSerialNumber	Serial number of the local site storage.
localStorageDeviceId	Storage device ID of the local site storage.
remoteSerialNumber	Serial number of the remote site storage.
remoteStorageDeviceId	Storage device ID of the remote site storage.
primaryVolumeId	Volume ID of the primary volume.
primaryVolumeNickname	Volume nickname of the primary volume.
secondaryVolumeId	Volume ID of the secondary volume.
secondaryVolumeNickname	Volume nickname of the secondary volume.

Details on HUR pair creation

When creating Replication CRs, other CRs are also created automatically. The following table shows a list of related CRs and an order of creation. This may help you to understand it and troubleshooting.

Site	CRD name	What the CR represents	The order in which the resource are ready
Primary	Replication	An HUR pair created by Hitachi Replication Plug-in for Containers. Physically the same as #4.	5
	LocalVolume	An HUR P-Vol created by users. Physically the same as #6.	1
	RemoteVolume	An HUR S-Vol created by Hitachi Replication Plug-in for Containers. Physically the same as #5.	4
Secondary	Replication	An HUR pair created by Hitachi Replication Plug-in for Containers. Physically the same as #1.	5
	LocalVolume	An HUR S-Vol created by Hitachi Replication Plug-in for Containers. Physically the same as #3.	2
	RemoteVolume	An HUR P-Vol created by users. Physically the same as #2.	3

When you run `kubectl create <replication>` command, the above CRs are created in the following order.

1. A LocalVolume CR in the primary site is created and becomes `true` when a PVC specified in `hspc_v1_replication.yaml` file exists.
2. A LocalVolume CR in the secondary site is created and becomes `true` when a PVC is successfully created.
3. A RemoteVolume CR in the secondary site is created and becomes `true` when Hitachi Replication Plug-in for Containers can confirm existence of the PVC in the primary site.
4. A RemoteVolume CR in the primary site is created and becomes `true` when Hitachi Replication Plug-in for Containers can confirm existence of the PVC in the secondary site.
5. Replication CRs in both primary and secondary site becomes `Ready` when both Hitachi Replication Plug-in for Containers can confirm HUR pair creation and initial data copy.

You will see the following output when all CRs are successfully created.

```
% kubectl get localvolume
NAME                 READY   AGE
replication-sample   true    6d20h

% kubectl get remotevolume
NAME                 READY   AGE
replication-sample   true    6d20h

% kubectl get replication
NAME                 STATUS   DESIREDSTATE   OPERATION  AGE
replication-sample   Ready      none              none         6d20h
```

Timezone settings

Replication Plug-in for Containers supports the following environment variables.

Variable	Description
TZ	Timezone used for logging. Default is UTC.
SPC_VERIFY_CERTIFICATE	If true, Replication Plug-in for Containers will verify TLS certificates. Default is false.

Using resource partitioning feature of Hitachi Storage Plug-in for Containers

If you want to use the resource partitioning feature of Hitachi Storage Plug-in for Containers, configure the following items in both primary and secondary sites. Configuring resource partitioning in only one site is not supported. For more information on the resource partitioning feature, see *Hitachi Storage Plug-in for Containers Quick Reference Guide* .

Item	Description
Hitachi Storage Plug-in for Containers	Enable resource partitioning.
Journal	Create a journal from the resource group that is configured in Hitachi Storage Plug-in for Containers

Item	Description
Host Group	Allocate one undefined host group ID for each port of the resource group that is configured in Hitachi Storage Plug-in for Containers. **Note:** A host group named "spc-replication" will be created when you create a Replication. If a host group with the same name but with a different resource group already exists in the same port, then delete the existing host group.
Storage system user	Use the same storage system user that is configured in Hitachi Storage Plug-in for Containers.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
HitachiVantara.com | community.HitachiVantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
HitachiVantara.com/contact