

# Hitachi Block Storage Driver for Red Hat OpenStack Platform

---

## Install Guide

This guide provides information about installing Hitachi Block Storage Driver for Red Hat OpenStack Platform.

**MK-92ADPTR157-03**

**September 2024**

© 2022,2024 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara Corporation (collectively, "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara Corporation at [https://support.HitachiVantara.com/en\\_us/contact-us.html](https://support.HitachiVantara.com/en_us/contact-us.html).

**Notice:** Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara Corporation.

By using this software, you agree that you are responsible for:

- 1) Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals to access relevant data; and
- 2) Verifying that data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

**Notice on Export Controls.** The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

**EXPORT CONTROLS** - Licensee will comply fully with all applicable export laws and regulations of the United States and other countries, and Licensee shall not export, or allow the export or re-export of, the Software, API, or Materials in violation of any such laws or regulations. By downloading or using the Software, API, or Materials, Licensee agrees to the foregoing and represents and warrants that Licensee is not located in, under the control of, or a national or resident of any embargoed or restricted country.

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, Flash Copy, IBM, Lotus, MVS, OS/390, PowerPC, RS6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z/VM, BCPI™ and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft Corporate Logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or web site are properties of their respective owners.

# Table of Contents

<b>Preface .....</b>	<b>4</b>
About this document.....	4
Document conventions.....	4
Intended audience.....	4
<b>Installing Block Storage Driver for RHOSP .....</b>	<b>5</b>
<b>Verifying Block Storage Driver for RHOSP installation .....</b>	<b>6</b>
<b>Hitachi extended functions .....</b>	<b>8</b>
<b>Configuration options.....</b>	<b>13</b>
<b>Restrictions .....</b>	<b>19</b>
<b>Known Issues .....</b>	<b>19</b>

# Preface

## About this document

This document provides technical information and procedures for installing Hitachi Block Storage Driver, which is shipped as part of the Red Hat OpenStack Platform (RHOSP).

Target RHOSP versions of this document are as following:

- 16.2
- 17.1

## Document conventions

This document uses the following typographic convention:

Convention	Description
<b>Bold</b>	<ul style="list-style-type: none"><li>• Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: <b>Click OK</b>.</li><li>• Indicates emphasized words in list items.</li></ul>
<i>Italic</i>	Indicates a document title or emphasized words in text.
Monospace	Indicates text that is displayed on screen or entered by the user. Example: <code>pairdisplay -g oradb</code>

## Intended audience

This document is intended for operators and administrators who configure and operate cloud systems using Red Hat OpenStack Platform.

This document assumes basic knowledge of Linux operating systems.

# Installing Block Storage Driver for RHOSP

## Procedure

1. Create an environment file for Block Storage Driver for RHOSP.

To apply the settings of Block Storage Driver for RHOSP to a Cinder component during installation, create a new environment file (for example, `cinder-hitachi-config.yaml`) and then edit the file. For details about how to specify settings for an environment file, see the Red Hat document *Custom Block Storage Back End Deployment Guide* and *Hitachi block storage driver documentation* on the OpenStack Documentation portal.

**File:** `/home/stack/templates/cinder-hitachi-config.yaml`

```
# cinder-hitachi-config.yaml sample for Fibre Channel #
parameter_defaults:
  CinderEnableIscsiBackend: false
  ControllerExtraConfig:
    cinder::config::cinder_config:
      VSPG1000/volume_driver:
        value: cinder.volume.drivers.hitachi.hbsd_fc.HBSDFCDriver
      VSPG1000/hitachi_storage_id:
        value: 800000010199
      VSPG1000/san_ip:
        value: 172.16.18.108
      VSPG1000/san_login:
        value: HBSD_USER
      VSPG1000/san_password:
        value: password
      VSPG1000/hitachi_pool:
        value: 47
      VSPG1000/hitachi_target_ports:
        value: CL1-A,CL2-A
      VSPG1000/hitachi_group_create:
        value: true
      VSPG1000/volume_backend_name:
        value: VSPG1000
    cinder_user_enabled_backends: ['VSPG1000']
```

```
# cinder-hitachi-config.yaml sample for iSCSI #
parameter_defaults:
  CinderEnableIscsiBackend: false
  ControllerExtraConfig:
    cinder::config::cinder_config:
      VSPG1000/volume_driver:
        value: cinder.volume.drivers.hitachi.hbsd_iscsi.HBSDISCSIDriver
      VSPG1000/hitachi_storage_id:
        value: 800000010199
      VSPG1000/san_ip:
        value: 172.16.18.108
      VSPG1000/san_login:
        value: HBSD_USER
```

```
VSPG1000/san_password:
  value: password
VSPG1000/hitachi_pool:
  value: 47
VSPG1000/hitachi_target_ports:
  value: CL1-A,CL2-A
VSPG1000/hitachi_group_create:
  value: true
VSPG1000/volume_backend_name:
  value: VSPG1000
cinder_user_enabled_backends: ['VSPG1000']
```

## 2. Deploy the Hitachi Block Storage driver.

Run the `openstack overcloud deploy` command to deploy the components. When you run the command, specify the following environment setting file: `/home/stack/containers-prepare-parameter.yaml` and the environment file that you created in step 1.

```
openstack overcloud deploy --template \
-e /home/stack/containers-prepare-parameter.yaml \
-e <other templates> \
...
-e /home/stack/templates/cinder-hitachi-config.yaml
```

For information about the specifications of the `openstack overcloud deploy` command and deployment methods, see the Red Hat document *Red Hat OpenStack Platform Director Installation and Usage*.

# Verifying Block Storage Driver for RHOSP installation

Use the following procedure to verify installation and configuration of Block Storage Driver for RHOSP.

## Procedure

1. After the Cinder service restarts, make sure that Block Storage Driver for RHOSP is available for use by verifying the following:
  - Verify that the state of the `cinder-volume` service in the backend is `up`. To verify the `cinder-volume` service state, run the `openstack volume service list` command as shown in the following example. If the `cinder-volume` service state is `down`, then it cannot be used for volume operations. Check the log files located in `/var/log/containers/cinder` on the controller node.

```
$ openstack volume service list
+-----+-----+-----+-----+-----+-----+
| Binary          | Host          | Zone | Status | State | Updated At |
+-----+-----+-----+-----+-----+-----+
|                  |                |      |         |       |              |
```

cinder-scheduler	overcloud-controller-0	nova	enabled	up	2024-04-03T07:06:39.000000
cinder-backup	overcloud-controller-0	nova	enabled	up	2024-04-03T07:06:39.000000
cinder-volume	overcloud-controller-0@VSPG1000_ISCSI	nova	enabled	up	2024-04-03T07:06:45.000000
cinder-volume	overcloud-controller-0@VSPG1000_FC	nova	enabled	down	2024-01-23T08:17:19.000000
cinder-volume	overcloud-controller-0@VSPG1000_ISCSI2	nova	enabled	up	2024-04-03T07:06:40.000000

2. Verify that the following volume operations can be performed from the Horizon portal:

- Create Volume
- Delete Volume
- Attach Volume
- Detach Volume
- Create Snapshot
- Delete Snapshot
- Create Volume from Snapshot
- Create Volume from Volume (Clone)

If any of these volume operations fail, check the log files located in `/var/log/containers/cinder` on the controller node.

# Hitachi extended functions

## Set up global-active device and volume operation

Note:

- This function is supported on RHOSP 17.1.3 and later.

If you are using Global-Active Device(GAD), you can make the data of individual volumes redundant between two storage systems, thereby improving the availability of the storage systems. For details, see the *Global-Active Device User Guide*.

Note:

- You cannot apply global-active device configuration and remote replication configuration to the same backend.
- You cannot use Asymmetric Logical Unit Access (ALUA).

## Storage firmware versions

If you are using a VSP F350, F370, F700, F900 storage system or a VSP G350, G370, G700,G900 storage system in a global-active device configuration, make sure the firmware version is 88-03-21 or later .

## Creating a global-active device environment

Before using global-active device, create the prerequisite environment, such as connecting remote paths, configuring a quorum disk, and creating a virtual storage machine (VSM), by other storage system management tools. Hitachi blockstorage driver supports the following configurations.

- Configuration where the P-VOL is not registered to a VSM
- Configuration where the P-VOL is registered to a VSM

For details, see the *Workflow for creating a GAD environment in the Global-Active Device User Guide*.

Hitachi block storage driver automatically runs procedures that are described in the *Global-Active Device User Guide*. You can skip configuring the following procedures, which are described in the section *Workflow for creating a GAD environment*.

- The following steps of Setting up the secondary storage system:
  - Setting the GAD reserve attribute on the S-VOL
  - Creating a host group (Only if the `hitachi_group_create` parameter is True )
  - Creating the S-VOL

- Adding an LU path to the S-VOL
- Updating the CCI configuration definition files
- Creating the GAD pair
- Adding an alternate path to the S-VOL

You must register the information about the secondary storage system to the REST API server in the primary site and register the information about the primary storage system to the REST API server in the secondary site. For details about how to register the information, see the *Hitachi Command Suite Configuration Manager REST API Reference Guide* or the *Hitachi Ops Center API Configuration Manager REST API Reference Guide*.

Note:

- The users specified for the `san_login` parameter and the `hitachi_mirror_rest_user` parameter must have following roles:
  - Storage Administrator (View & Modify)
  - Storage Administrator (Remote Copy)
- Reserve unused host group IDs (iSCSI target IDs) for the resource groups related on the VSM. Reserve the IDs in ascending order. The number of IDs you need to reserve is 1 plus the sum of the number of controller odes and the number of compute nodes. For details on how to reserve a host group ID (iSCSI target ID), see the *Global-Active Device User Guide*.
- The LUNs of the host groups (iSCSI targets) of the specified ports on the primary storage system must match the LUNs of the host groups (iSCSI targets) of the specified ports on the secondary storage system. If they do not match, match the LUNs for the primary storage system with those for the secondary storage system.
- When you use a same storage system as secondary storage system for global-active device configuration and backend storage system for general use at the same time, you cannot use the same ports between different backend storage systems. Please specify different ports in `hitachi_target_ports` parameter, `hitachi_compute_target_ports` parameter, or `hitachi_rest_pair_target_ports` parameter between different backend storage systems.

## Create volume in a global-active device configuration

If you create a Cinder volume in a global-active device configuration, each global-active device pair is mapped to a Cinder volume.

To create a volume by using the Create Volume with the global-active device attribute specified, specify `hbsd:topology=active_active_mirror_volume` as an extra spec for the volume type as follows:

```
$ openstack volume type create <volume type name>
$ openstack volume type set --property \
hbsd:topology=active_active_mirror_volume <volume type name>
$ openstack volume create -type <volume type name> --size <size>
```

Note:

- In this case, the following restrictions apply:
  - You cannot create a volume for which the deduplication and compression function is enabled, or creating a volume will be failed with the error “MSGID0753-E: Failed to create a volume in a GAD environment because deduplication is enabled for the volume type.”.
- Note the following if the configuration is “P-VOL registered to a VSM”:
  - Be sure to create volumes with the global-active device attribute specified.
  - Before performing a volume operation, you must set a virtual LDEV number for every P-VOL.

## Unavailable Cinder functions in a global-active configuration

Following cinder functions are unavailable in a global-active configuration:

- Volume Migration (storage assisted)
- Manage Volume
- Unmanage Volume

Note:

- In addition, if the configuration is “P-VOL registered to a VSM”, the backup creation command of the Backup Volume functions cannot be run with the `-snapshot` option or the `-force` option specified.

## Data deduplication and compression

Note:

- This function is supported on RHOSP 17.1.3 and later.

You can use deduplication and compression to improve storage utilization using data reduction.

For details, see the *Capacity saving function: data deduplication and compression in the Provisioning Guide*.

## Enabling deduplication and compression

To use the deduplication and compression on the storage models, your storage administrator must first enable the deduplication and compression for the DP pool.

For details about how to enable this setting, see the description of pool management in the *Hitachi Command Suite Configuration Manager REST API Reference Guide* or the *Hitachi Ops Center API Configuration Manager REST API Reference Guide*.

Note:

- Do not set a subscription limit (`virtualVolumeCapacityRate`) for the DP pool.

## Creating a volume with deduplication and compression enabled

To create a volume with the deduplication and compression setting enabled, enable deduplication and compression for the relevant volume type.

Procedure

1. To enable the deduplication and compression setting, specify `deduplication_compression` for `hbsd:capacity_saving` in the extra specs for the volume type.
2. When creating a volume of the volume type created in the previous step, you can create a volume with the deduplication and compression function enabled.

## Deleting a volume with deduplication and compression enabled

The cinder delete command finishes when the storage system starts the LDEV deletion process. The LDEV cannot be reused until the LDEV deletion process is completed on the storage system.

## Maximum number of copy pairs and consistency groups

The maximum number of Thin Image pairs that can be created for each LDEV assigned to a volume (or snapshot) is restricted on a per-storage-system basis. If the number of pairs exceeds the maximum, copying cannot proceed normally.

For information about the maximum number of copy pairs and consistency groups that can be created, see the *Hitachi Thin Image User Guide*.

## Port scheduler

Note:

- This function is supported on RHOSP 17.1.3 and later.

You can use the port scheduler function to reduce the number of WWNs, which are storage system resource.

In Hitachi block storage driver, if host groups are created automatically, host groups are created for each compute node or VM (in an environment that has a WWN for each VM). If you do not use the port scheduler function, host groups are created and the same WWNs are registered in all of the ports that are specified for the parameter `hitachi_compute_target_ports` or for the parameter `hitachi_target_ports`. For

Hitachi storage devices, a maximum of 255 host groups and 255 WWNs can be registered for one port. When volumes are attached, the upper limit on the number of WWNs that can be registered might be unexpectedly exceeded.

For the port scheduler function, when the cinder-volume service starts, the Fibre Channel Zone Manager obtains the WWNs of active compute nodes and of active VMs. When volumes are attached, the WWNs are registered in a round-robin procedure, in the same order as the order of ports specified for the parameter `hitachi_compute_target_ports` or for the parameter `hitachi_target_ports`.

If you want to use the port scheduler function, set the `hitachi_port_scheduler` parameter. For details about the parameter, see Configuration options.

Note:

- If a host group already exists in any of the ports specified for the parameter `hitachi_compute_target_ports` or for the parameter `hitachi_target_ports`, no new host group will be created on those ports.
- Restarting the cinder-volume service re-initializes the round robin scheduling determined by the `hitachi_compute_target_ports` parameter or the `hitachi_target_ports` parameter.
- The port scheduler function divides up the active WWNs from each fabric controller and registers them to each port. For this reason, the number of WWNs registered may vary from port to port.

## Port assignment using extra specs

Note:

- This function is supported on RHOSP 17.1.3 and later.

Defining particular ports in the Hitachi-supported extra spec `hbsd:target_ports` determines which of the ports specified by the `hitachi_target_ports` or the `hitachi_compute_target_ports` parameters are used to create LUN paths during volume attach operations for each volume type.

Note:

- Use a comma to separate multiple ports.
- In a global-active device configuration, use the extra spec `hbsd:target_ports` for the primary storage system and the extra spec `hbsd:remote_target_ports` for the secondary storage system.
- In a global-active device configuration, the ports specified for the extra spec `hbsd:remote_target_ports` must be specified for the `hitachi_mirror_target_ports` parameter or the `hitachi_mirror_compute_target_ports` parameter.

# Configuration options

This table shows configuration options for Hitachi block storage driver.

Configuration option = Default value	Description	Note
hitachi_async_copy_check_interval = 10	(Integer(min=1, max=600)) Interval in seconds to check asynchronous copying status during a copy pair deletion or data restoration.	Supported RHOSP 17.1.3 and later
hitachi_compute_target_ports = []	(List of String) IDs of the storage ports used to attach volumes to compute nodes. To specify multiple ports, connect them by commas (e.g.CL1-A,CL2-A).	
hitachi_copy_check_interval = 3	(Integer(min=1, max=600)) Interval in seconds to check copying status during a volume copy.	Supported RHOSP 17.1.3 and later
hitachi_copy_speed = 3	(Integer(min=1, max=15)) Copy speed of storage system. 1 or 2 indicates low speed, 3 indicates middle speed, and a value between 4 and 15 indicates high speed.	Supported RHOSP 17.1.3 and later
hitachi_discard_zero_page = True	(Boolean) Enable or disable zero page reclamation in a DP-VOL.	
hitachi_exec_retry_interval = 5	(Integer) Retry interval in seconds for REST API execution.	Supported RHOSP 17.1.3 and later
hitachi_extend_timeout = 600	(Integer) Maximum wait time in seconds for a volume extension to complete.	Supported RHOSP 17.1.3 and later

hitachi_group_create = False	(Boolean) If True, the driver will create host groups or iSCSI targets on storage ports as needed.	
hitachi_group_delete = False	(Boolean) If True, the driver will delete host groups or iSCSI targets on storage ports as needed.	
hitachi_host_mode_options = []	(List of Integer) host mode option for host group or iSCSI target	
hitachi_ldev_range = None	(String) Range of the LDEV numbers in the format of 'xxxx-yyyy' that can be used by the driver. Values can be in decimal format (e.g.1000) or in colon-separated hexadecimal format(e.g. 00:03:E8).	
hitachi_lock_timeout = 7200	(Integer) Maximum wait time in seconds for storage to be logged in or unlocked.	Supported RHOSP 17.1.3 and later
hitachi_lun_retry_interval = 1	(Integer) Retry interval in seconds for REST API adding a LUN mapping to the server.	Supported RHOSP 17.1.3 and later
hitachi_lun_timeout = 50	(Integer) Maximum wait time in seconds for adding a LUN mapping to the server.	
hitachi_mirror_auth_password = None	(String) iSCSI authentication password	Supported RHOSP 17.1.3 and later
hitachi_mirror_auth_user = None	(String) iSCSI authentication username	Supported RHOSP 17.1.3 and later

hitachi_mirror_compute_target_ports = []	(List of String) Target port names of compute node for host group or iSCSI target	Supported RHOSP 17.1.3 and later
hitachi_mirror_ldev_range = None	(String) Logical device range of secondary storage system	Supported RHOSP 17.1.3 and later
hitachi_mirror_pair_target_number = 0	(Integer(min=0, max=99)) Pair target name of the host group or iSCSI target	Supported RHOSP 17.1.3 and later
hitachi_mirror_pool = None	(String) Pool of secondary storage system	Supported RHOSP 17.1.3 and later
hitachi_mirror_rest_api_ip = None	(String) IP address of REST API server	Supported RHOSP 17.1.3 and later
hitachi_mirror_rest_api_port = 443	(Port(min=0, max=65535)) Port number of REST API server	Supported RHOSP 17.1.3 and later
hitachi_mirror_rest_pair_target_ports = []	(List of String) Target port names for pair of the host group or iSCSI target	Supported RHOSP 17.1.3 and later
hitachi_mirror_rest_password = None	(String) Password of secondary storage system for REST API	Supported RHOSP 17.1.3 and later
hitachi_mirror_rest_user = None	(String) Username of secondary storage system for REST API	Supported RHOSP 17.1.3 and later
hitachi_mirror_snap_pool = None	(String) Thin pool of secondary storage system	Supported RHOSP 17.1.3 and later
hitachi_mirror_ssl_cert_path = None	(String) Can be used to specify a non default path to a CA_BUNDLE file or directory with certificates of trusted CAs, which will be used to validate the backend	Supported RHOSP 17.1.3 and later

hitachi_mirror_ssl_cert_verify = False	(Boolean) If set to True the http client will validate the SSL certificate of the backend endpoint.	Supported RHOSP 17.1.3 and later
hitachi_mirror_storage_id = None	(String) ID of secondary storage system	Supported RHOSP 17.1.3 and later
hitachi_mirror_target_ports = []	(List of String) Target port names for host group or iSCSI target	Supported RHOSP 17.1.3 and later
hitachi_mirror_use_chap_auth = False	(Boolean) Whether or not to use iSCSI authentication	Supported RHOSP 17.1.3 and later
hitachi_pair_target_number = 0	(Integer(min=0, max=99)) Pair target name of the host group or iSCSI target	Supported RHOSP 17.1.3 and later
hitachi_pool = None	(String) Pool number or pool name of the DP pool.	
hitachi_port_scheduler = False	(Boolean) Enable port scheduling of WWNs to the configured ports so that WWNs are registered to ports in a round-robin fashion.	Supported RHOSP 17.1.3 and later
hitachi_quorum_disk_id = None	(Integer(min=0, max=31)) ID of the Quorum disk used for global-active device	Supported RHOSP 17.1.3 and later
hitachi_rest_another_ldev_mapped_retry_timeout = 600	(Integer) Retry time in seconds when new LUN allocation request fails.	Supported RHOSP 17.1.3 and later
hitachi_rest_connect_timeout = 30	(Integer) Maximum wait time in seconds for connecting to REST API session.	

hitachi_rest_disable_io_wait = True	(Boolean) This option will allow detaching volume immediately. If set False, storage may take few minutes to detach volume after I/O.	Supported RHOSP 17.1.3 and later
hitachi_rest_get_api_response_timeout = 1800	(Integer) Maximum wait time in seconds for a response against sync methods, for example GET	
hitachi_rest_job_api_response_timeout = 1800	(Integer) Maximum wait time in seconds for a response against async methods from REST API, for example PUT and DELETE.	
hitachi_rest_keep_session_loop_interval = 180	(Integer) Loop interval in seconds for keeping REST API session.	
hitachi_rest_pair_target_ports = []	(List of String) Target port names for pair of the host group or iSCSI target	Supported RHOSP 17.1.3 and later
hitachi_rest_server_busy_timeout = 7200	(Integer) Maximum wait time in seconds when REST API returns busy.	
hitachi_rest_tcp_keepalive = True	(Boolean) Enables or disables use of REST API tcp keepalive	
hitachi_rest_tcp_keepcnt = 4	(Integer) Maximum number of transmissions for TCP keepalive packet.	
hitachi_rest_tcp_keepidle = 60	(Integer) Wait time in seconds for sending a first TCP keepalive packet.	

hitachi_rest_tcp_keepintvl = 15	(Integer) Interval of transmissions in seconds for TCP keepalive packet.	
hitachi_rest_timeout = 30	(Integer) Maximum wait time in seconds for each REST API request.	
hitachi_restore_timeout = 86400	(Integer) Maximum wait time in seconds for the restore operation to complete.	
hitachi_set_mirror_reserve_attribute = True	(Boolean) Whether or not to set the mirror reserve attribute	Supported RHOSP 17.1.3 and later
hitachi_snap_pool = None	(String) Pool number or pool name of the snapshot pool.	
hitachi_state_transition_timeout = 900	(Integer) Maximum wait time in seconds for a volume transition to complete.	
hitachi_storage_id = None	(String) Product number of the storage system.	
hitachi_target_ports = []	(List of String) IDs of the storage ports used to attach volumes to the controller node. To specify multiple ports, connect them by commas (e.g.CL1-A,CL2-A).	
hitachi_zoning_request = False	(Boolean) If True, the driver will configure FC zoning between the server and the storage system provided that FC zoning manager is enabled.	

## Restrictions

- The volume having snapshot(s) cannot be extended.

## Known Issues

- No compatibility copy pair name between RHOSP16.1 and RHOSP17.1.3 and later  
The identifier for GAD volume pair is different between

## Hitachi Vantara

Corporate Headquarters 2535 Augustine Drive  
Santa Clara, CA 95054 USA [www.HitachiVantara.com](http://www.HitachiVantara.com) [community.HitachiVantara.com](http://community.HitachiVantara.com)

### Regional Contact Information

Americas: +1 866 374 5822 or [info@hitachivantara.com](mailto:info@hitachivantara.com)

Europe, Middle East and Africa: +44 (0) 1753 618000 or [info.emea@hitachivantara.com](mailto:info.emea@hitachivantara.com)

Asia Pacific: +852 3189 7900 or [info.marketing.apac@hitachivantara.com](mailto:info.marketing.apac@hitachivantara.com)

