Hitachi Block Storage Driver for OpenStack Rocky

User Guide

This guide provides information about installing, configuring, and managing Hitachi Block Storage Driver for OpenStack.

MK-92ADPTR136-05
May 2019

Contents

Preface

This guide provides overview of and technical information for installing, configuring, and administering Hitachi Block Storage Driver for OpenStack.

Intended audience

This document is intended for operators and administrators who configure and operate cloud systems using the following software tools:

- Red Hat OpenStack Platform
- SUSE OpenStack Cloud
- Mirantis OpenStack
- Ubuntu OpenStack
- Configuration Manager REST API

This document assumes they have basic knowledge of Linux operating systems.

Software version

This document applies to Block Storage Driver for OpenStack version 7.0.0 or later for Rocky.

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.

Convention	Description
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`
< > 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).

Icon	Label	Description
	WARNING	Warns the user of a hazardous situation which, if not avoided, could result in death or serious injury.

Release notes

Read the release notes before installing and using this product. They may contain requirements or restrictions that are not fully described in this document or updates or corrections to this document. Release notes are available on Hitachi Vantara Support Connect: https://knowledge.hitachivantara.com/Documents.

Changes made in this revision

You can perform Cinder operations for volumes in a global-active device configuration.

Referenced documents

The documents below are referenced in this document or contain more information about the features described in this document.

Hitachi Vantara documents

Hitachi Command Suite Configuration Manager REST API:

REST API Reference Guide, MK-92HC229

Hitachi Virtual Storage Platform F series and G series:

- *Global-Active Device User Guide*, MK-97HM85037
- *Thin Image User Guide*, MK-97HM85032

Hitachi Virtual Storage Platform F1500 and Hitachi Virtual Storage Platform G1000, G1500:

Provisioning Guide for Open Systems, MK-92RD8014

Hitachi Virtual Storage Platform F400, F600, F800 and Hitachi Virtual Storage Platform G200, G400, G600, G800:

Provisioning Guide, MK-94HM8014

Hitachi Virtual Storage Platform F350, F370, F700, F900 and Hitachi Virtual Storage Platform G350, G370, G700, G900:

- *Hitachi Storage Advisor Embedded Guide*, MK-97HM85022
- *Provisioning Guide*, MK-97HM85026

Hitachi Vantara Support Connect, https://knowledge.hitachivantara.com/Documents

Red Hat documents

- *Red Hat OpenStack Platform Custom Block Storage Back End Deployment Guide*
- *Red Hat OpenStack Platform Director Installation and Usage*
- *Red Hat OpenStack Platform Keeping Red Hat OpenStack Platform Updated*
- *Red Hat OpenStack Platform Understanding Red Hat OpenStack Platform High Availability*

Related documents

The documents below are related to this document or contain more information about features described in this document.

Hitachi Vantara documents

Hitachi Virtual Storage Platform F1500 and Hitachi Virtual Storage Platform G1000, G1500:

- *Hardware Guide*, MK-92RD8007
- *System Administrator Guide*, MK-92RD8016

Hitachi Virtual Storage Platform F400, F600, F800 and Hitachi Virtual Storage Platform G200, G400, G600, G800:

- *Hardware Reference Guide: Hitachi Virtual Storage Platform F400, F600*, MK-94HM8045
- *Hardware Reference Guide: Hitachi Virtual Storage Platform F800*, MK-94HM8046
- *Hardware Reference Guide: Hitachi Virtual Storage Platform G200*, MK-94HM8020
- *Hardware Reference Guide: Hitachi Virtual Storage Platform G400, G600*, MK-94HM8022
- *Hardware Reference Guide: Hitachi Virtual Storage Platform G800*, MK-94HM8026
- *System Administrator Guide*, MK-94HM8016

Hitachi Virtual Storage Platform F350, F370, F700, F900 and Hitachi Virtual Storage Platform G350, G370, G700, G900:

- *Hardware Reference Guide: Hitachi Virtual Storage Platform F350 and F370*, MK-97HM85016
- *Hardware Reference Guide: Hitachi Virtual Storage Platform F700*, MK-97HM85017
- *Hardware Reference Guide: Hitachi Virtual Storage Platform F900*, MK-97HM85018
- *Hardware Reference Guide: Hitachi Virtual Storage Platform G350 and G370*, MK-97HM85012

- *Hardware Reference Guide: Hitachi Virtual Storage Platform G700*, MK-97HM85013
- *Hardware Reference Guide: Hitachi Virtual Storage Platform G900*, MK-97HM85014
- *System Administrator Guide*, MK-97HM85028

OpenStack

- *OpenStack Cloud Administrator Guide*
- *OpenStack Command-Line Interface Reference*
- *OpenStack Configuration Reference*
- *Red Hat OpenStack Platform Product Documentation*
- *SUSE OpenStack Cloud Product Manual*
- *Mirantis OpenStack Product Documentation*
- *Ubuntu OpenStack Documentation*

Conventions for capacity values

In this document, 1 KB, 1 MB, 1 GB, and 1 TB are used to denote 1 KiB, 1 MiB, 1 GiB, and 1 TiB, respectively.

1 KiB, 1 MiB, 1 GiB, and 1 TiB are equivalent to 1,024 bytes, 1,024 KiB, 1,024 MiB, and 1,024 GiB, respectively.

Getting help

Hitachi Vantara Support Connect is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to Hitachi Vantara Support Connect 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 Corporation.

Thank you!

Chapter 1: Overview

Block Storage Driver for OpenStack is based on the OpenStack Block Storage architecture (Cinder) and it enables you to provide storage services that use high-performance and high-reliability Hitachi storage systems to build and deliver Infrastructure as a Service (IaaS).

System Configuration

A typical system configuration for Block Storage Driver for OpenStack contains the main components shown in the figure below. This figure is an example of a multi-node configuration where the controller node and the compute nodes are configured on different nodes.

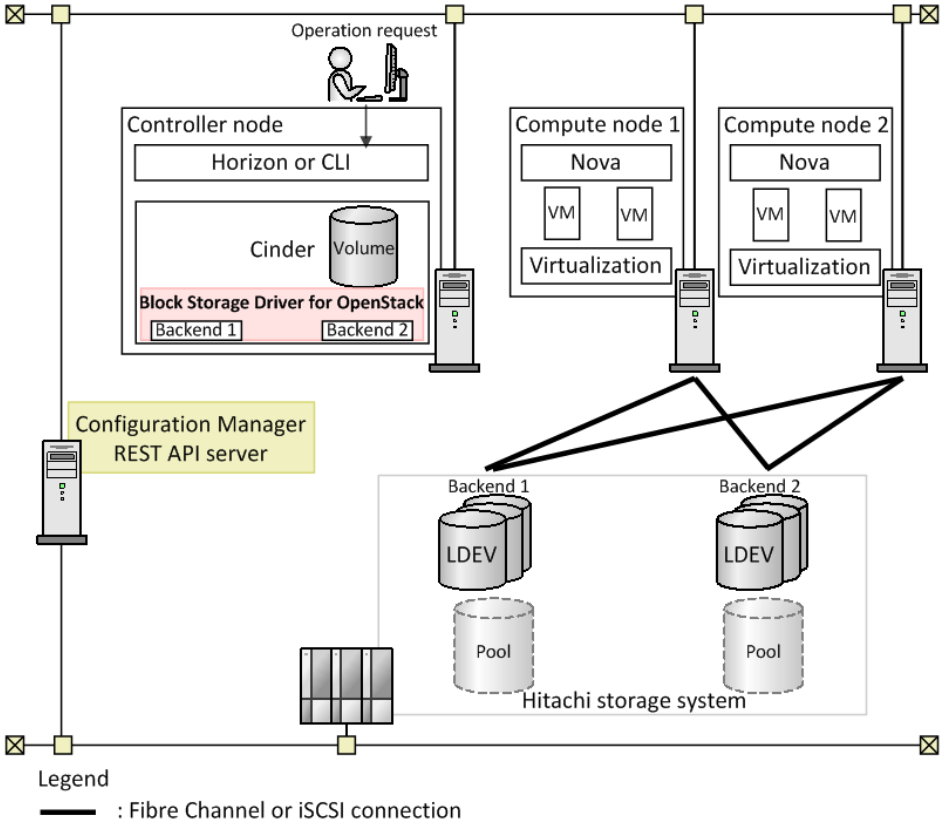

Controller node

A server that manages cloud systems. The cinder-volume services run on this node. Install and use Block Storage Driver for OpenStack on this node.

Compute node

A server managed by the controller node.

Configuration Manager REST API server

A server that receives REST API requests from Block Storage Driver for OpenStack, issues commands to the storage system, and returns the execution results. The REST API server can be installed independently or installed together as a part of the Hitachi Command Suite.

Backend storage system

A storage system managed by Cinder. Block Storage Driver for OpenStack uses thin-provisioned volumes (DP-VOLs) from Hitachi Dynamic Provisioning (DP) pools as backend storage. To use multiple DP pools, register each DP pool as a separate backend storage system.

LDEV

An individual logical data volume (on multiple drives in a RAID configuration) in a storage system.

Volume

Data unit managed by Block Storage Driver for OpenStack.

Supported Cinder functions

The following table lists the Cinder operations supported by Block Storage Driver for OpenStack.

Function	Description
Create Volume	Creates a volume.
Delete Volume[1]	Deletes a volume.
Attach Volume	Connects a volume to a VM instance.
Detach Volume	Disconnects a target volume from a VM instance.
Extend Volume	Resizes a volume.
Create Snapshot[2]	Creates a snapshot that stores point-in-time data.
Delete Snapshot	Deletes a snapshot.
List Snapshot	Lists snapshots.
Create Volume from Snapshot	Creates a volume with the same content as that of a snapshot.

Function	Description
Create Volume from Image	Retrieves image stored from Glance and loads it into the specified volume. A VM can be started based on the stored image.
Create Volume from Volume (Clone)	Creates a volume and then stores the data from the source volume.
Create Image from Volume	Creates an image from a volume and uploads the image (Glance).
Volume Migration (host assisted)[1]	Moves a volume to another DP pool managed by another storage system.
Consistency Group	Consistency groups are used to group volumes together for the purpose of data protection (snapshots, backups) and disaster recovery.
Over subscription	Uses thin provisioning to manage the capacity of the storage system.
Volume Migration (storage assisted)[1]	Moves a volume to another DP pool.
Backup using temp Snapshot	Uses a temporary snapshot or volume to perform non-disruptive backup. As a separate feature, this function can be used to back up snapshots. The Cinder backup service must be running.
Multi-attach Support	Simultaneously connects multiple hosts or servers to a volume. **Caution:** It is the responsibility of the user to ensure that a multiattach or clustered file system is used on the volumes. Otherwise there may be a high probability of data corruption.
Backup Volume	Backs up a volume to Swift. The Cinder backup service must be running.
Restore Backup	Restores data to a specified volume from Swift. The Cinder backup service must be running.
Revert to snapshot	Reverts a volume to a previously taken snapshot.
Manage Volume	Adds an LDEV created by a different OpenStack system to the volumes managed by Cinder.
Unmanage Volume	Removes a volume created by Cinder from Cinder management.

Function	Description
Retype Volume[1]	Changes the volume type.

Notes:

1. This function is not supported for volumes in a group.
2. Snapshots are created using Thin Image (Snapshot).

Support for additional OpenStack operations

The following table lists the OpenStack operations supported by Block Storage Driver for OpenStack.

OpenStack operation	Description
Live Migration	Moves a volume managed by Block Storage Driver for OpenStack and attached to a running VM from one compute node to another.
Multipath for I/O path	Multipath enabled for volumes managed by Block Storage Driver for OpenStack.
High availability for Cinder volume services	You can use the high availability cluster software to set the cinder-volume service in a redundant active-standby configuration. High availability is not supported for SUSE OpenStack Cloud or Mirantis OpenStack.
Fibre Channel Zone Manager	When a volume managed by Block Storage Driver for OpenStack is attached to a VM via Fibre Channel, Fibre Channel zoning can automatically be created or deleted between the compute node and the storage system during a volume attach or detach operation, respectively. See https://docs.openstack.org/cinder/latest/configuration/block-storage/fc-zoning.html for how to specify settings.

Backend storage management configuration

Block Storage Driver for OpenStack uses thin-provisioned volumes (DP-VOLs) from DP pools as backend storage (a Hitachi storage feature that provides virtual volume capability). To use multiple DP pools, register each DP pool as a separate backend storage system. Each DP pool will be identified in the `cinder.conf` file.

Volume operations

Block Storage Driver for OpenStack provides volume operation functions for managing LDEVs (logical devices or logical units) in the backend storage (DP pools). The figure below shows an overview of volume operations.

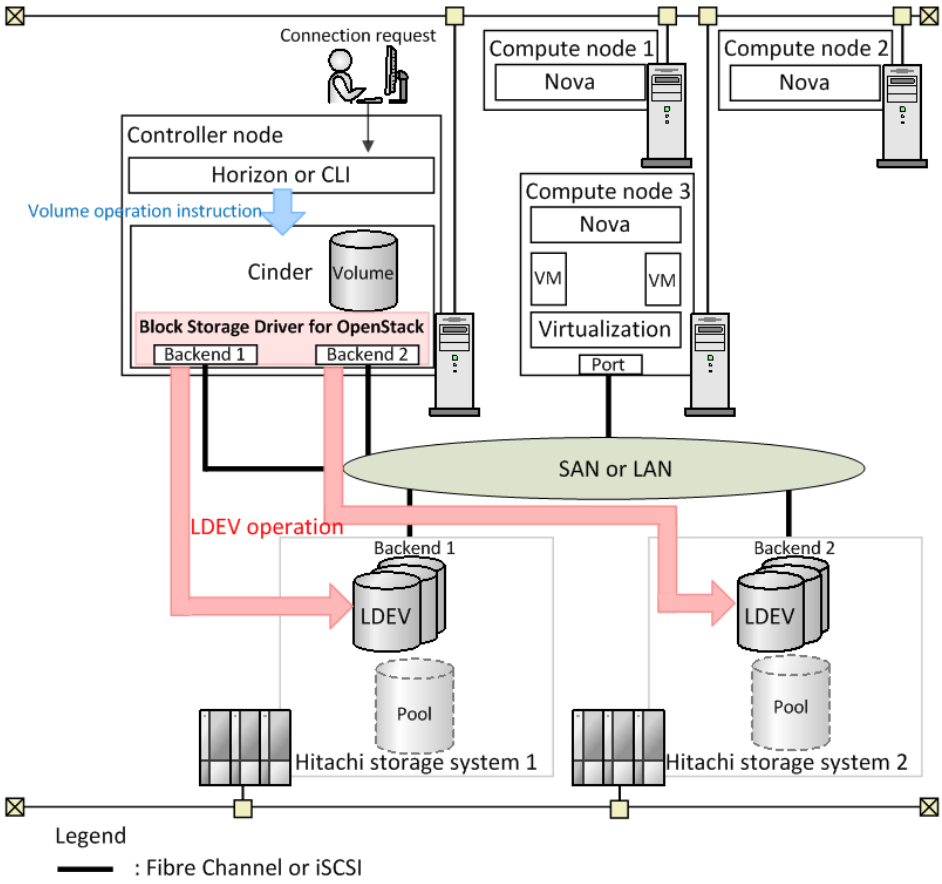

Figure 1 Overview of volume operations

Create and delete a volume

When Block Storage Driver for OpenStack receives the Cinder create volume request, it creates a virtual volume (DP-VOL) as a logical device (LDEV) from the DP pool. DP-VOLs allow the storage to be provisioned on an as-needed basis. Volumes can be deleted by using the Cinder delete volume command.

Extend a volume

Block Storage Driver for OpenStack allows you to resize the DP-VOL, ensuring that the volume can be expanded to meet changing business requirements.

Note: You cannot expand DP-VOLs created by using Thin Image (Clone) during asynchronous copying or DP-VOLs that have pair volumes that were created by using Thin Image (Snapshot).

Copy operations

When Block Storage Driver for OpenStack receives a Cinder copy volume request, Block Storage Driver for OpenStack copies the volume in full by using Thin Image (Clone), and takes a point-in-time snapshot of the volume by using Thin Image (Snapshot). In addition, snapshots for multiple volumes can be created by using a single command. This group of snapshots either guarantees or does not guarantee consistency, depending on the property of the group type to which the group belongs.

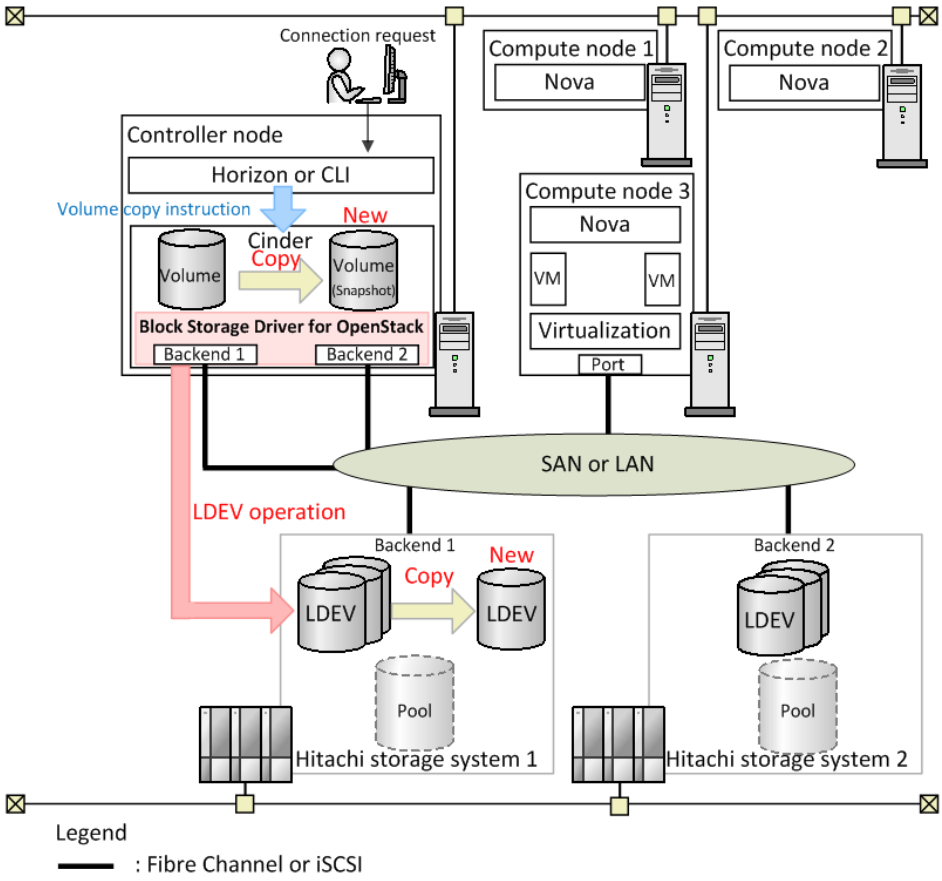

Figure 2 Snapshot operation example

Before you attempt to copy a volume attached to a VM, stop I/O on the VM.

The following table describes the copy methods available for Block Storage Driver for OpenStack. To perform copy operations, you will need a license for the Thin Image software product. A maximum of 1,024 clones and snapshots can be created for each volume.

Copy method	Description
Thin Image (Clone)	All data in the copy source volume (DP-VOL) is asynchronously copied to the destination volume. Both volumes can be used during the asynchronous copy.

Copy method	Description
Thin Image (Snapshot)	Differential data with the copy destination volume is stored as needed in a virtual volume (V-VOLv2) as an LDEV. Both volumes can be used. Consistency groups can also be copied. The virtual volume (V-VOLv2) can be used in a cascade configuration.

Create a snapshot

Upon receiving a Cinder create snapshot request, Block Storage Driver for OpenStack creates an LDEV linked to the snapshot, and then copies the snapshot data from the volume to the LDEV by Thin Image (Snapshot). The snapshot cannot be attached to a VM, but it can be kept as a data backup.

Create a volume from a snapshot

Upon receiving a Cinder create volume from snapshot request, Block Storage Driver for OpenStack creates an LDEV linked to the volume, and then stores the specified snapshot data in the LDEV by using one of the available copy methods (see the previous table). You can restore data by using the created snapshot.

Create a volume from a volume (clone)

Upon receiving a Cinder create volume from volume (clone) request, Block Storage Driver for OpenStack creates an LDEV linked to the volume, and then stores the specified volume data in the LDEV by using one of the available copy methods (see the previous table). Volume cloning is useful when multiple users use the same data.

Restore a volume from a snapshot

If data in a volume is corrupted, the Cinder snapshot function restores the data by creating a new volume from the snapshot. During this process, free space is temporarily consumed by the new volume and you must change the VM from the corrupted volume to the newly created volume. By using the snapshot or volume restore operation, you can restore the corrupted volume without having to create a new one.

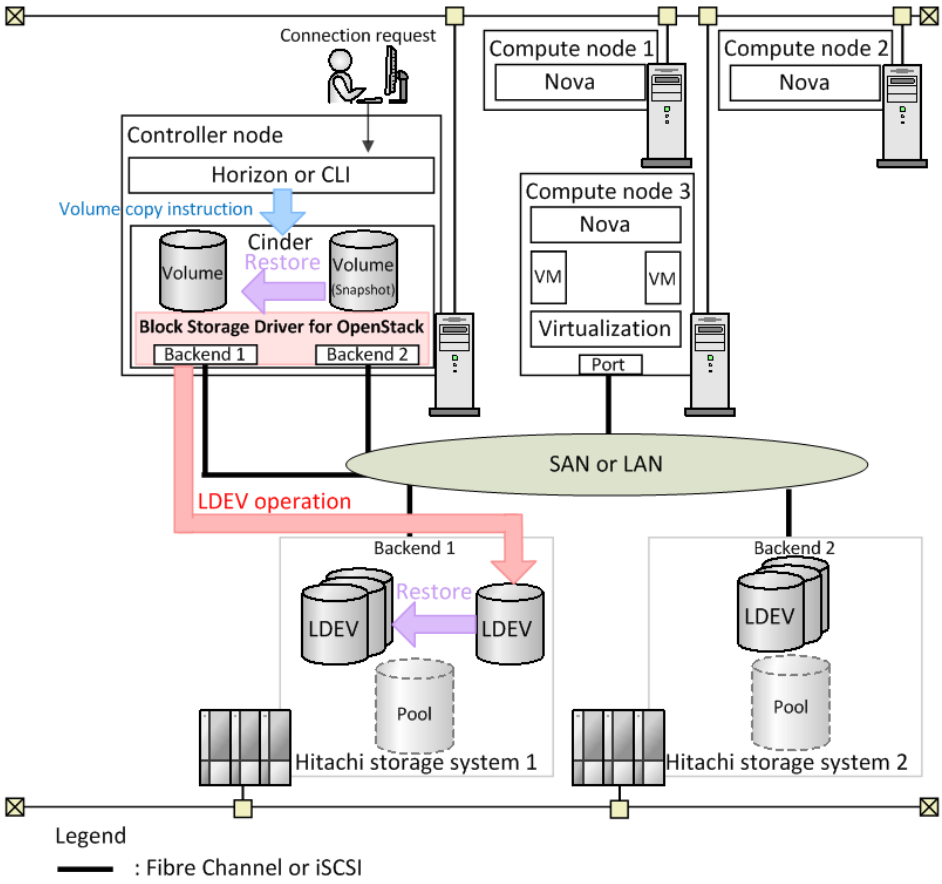

Figure 3 Snapshot and volume restore operation

Create or delete a group

Block Storage Driver for OpenStack enables you to create or delete a group of volumes. Deleting a group does not make changes to the backend storage if there are no volumes in the group. If you try to delete a group that has volumes in it by using **`cinder`** `consisgroup-delete --force` or **`cinder`** `group-delete --delete-volumes`, the driver deletes the LDEVs allocated to the volumes in the group.

Create a group from another group

You can create a group from another group. Block Storage Driver for OpenStack creates DP-VOLs in the backend storage and copies data from the volumes in the source group to the volumes in the destination group. When data copy is complete, the driver updates the status of the destination group.

Create a group from a snapshot of another group

You can create a group from a snapshot of another group. Block Storage Driver for OpenStack creates DP-VOLs in the backend storage and copies data from the snapshots of the source group to the volumes in the destination group. When data copy is complete, the driver updates the status of the destination group.

Note: Data consistency between source and destination groups is not maintained even for groups that are consistency groups.

Update a group

You can modify name and description of a group, as well as add or remove volumes to or from a group.

Create a snapshot of a group

You can create snapshots of multiple volumes by using a single command (**cinder** `cgsnapshot-create` or **cinder** `group-snapshot-create`). These snapshots are created at the same point in time if the group is a consistency group. If the group is not a consistency group, snapshots are created separately and the points in time at which the data for each snapshot is created might differ.

Note:

- There are two methods for creating a consistency group. A group created by using **cinder** `consisgroup-create` command is always a consistency group. A group created by using **cinder** `group-create` is a consistency group only if the `consistent_group_snapshot_enabled` property of the group type to which the group belongs is set to *<is>* `True`.
- There are two types of group-related Cinder commands: CG commands (commands that begin with **cinder** `consisgroup-` or **cinder** `cgsnapshot-`) and GVG commands (commands that begin with **cinder** `group-`). It is not recommended to use CG commands. In the future, CG commands will become obsolete and will no longer be supported. In addition, CG commands and GVG commands are not compatible with each other. Snapshots created by using CG commands cannot be operated by using GVG commands.
- GVG-related commands require volume API version 3.20 or later (**cinder** `--os-volume-api-version 3.20` or **export** `OS_VOLUME_API_VERSION=3.20`).
- **openstack** commands support CG-related operations, but do not support GVG-related operations yet.

Volume attachment

Block Storage Driver for OpenStack supports storage attachment capability for storage systems that use the following protocols: Fibre Channel, iSCSI, and Fibre Channel over Ethernet. Block Storage Driver for OpenStack supports multipath configurations. The port scheduler function is also supported. If this function is used, then, when volumes are attached, the active WWNs on the obtained compute nodes are distributed and registered to the host groups of each port of the storage system. Volumes will be assigned based on this attachment information. The following figure shows an overview of the volume attachment operations.

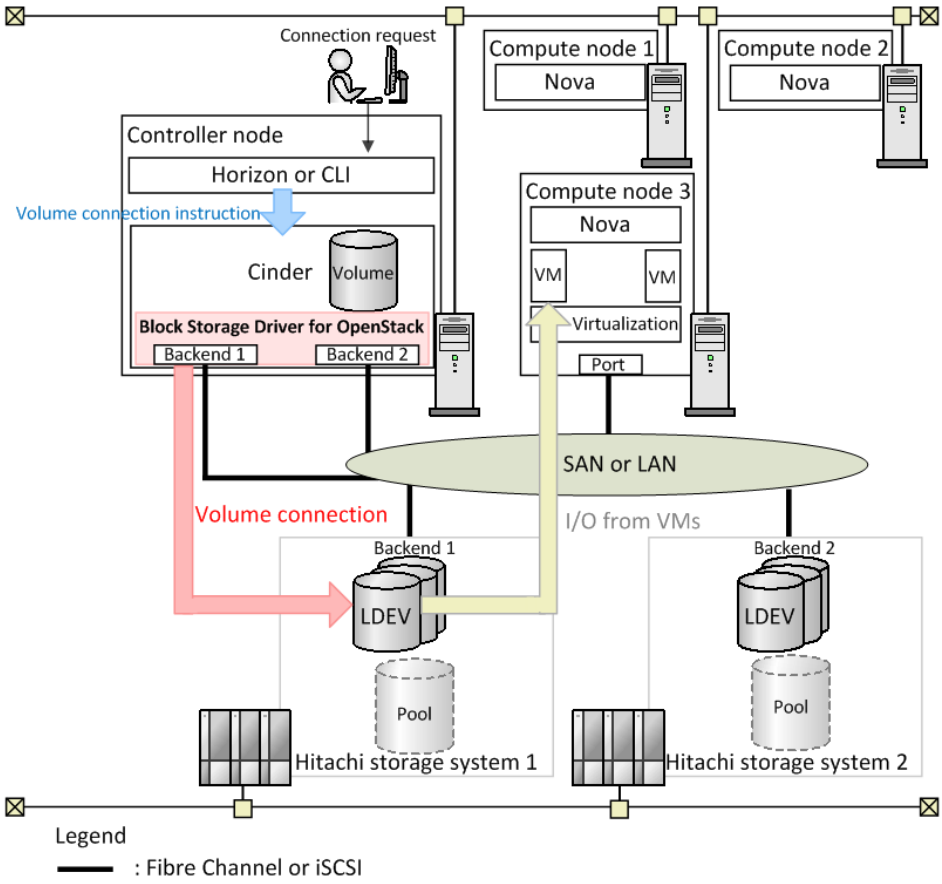

Figure 4 Volume attachment

Upon receiving a Cinder attach volume request, Block Storage Driver for OpenStack attaches a volume to a VM in the specified compute node. Block Storage Driver for OpenStack then directly connects the Hitachi storage system and the compute node. This allows the VM to access the attached volume to store data. The volume can also be detached from the VM. The volume can be attached again to another VM. Block Storage Driver for OpenStack allows volume attach and detach operations in a multipath configuration.

Backup operations

You can use standard Cinder functions to back up or restore volumes created by Block Storage Driver for OpenStack. If you restore a volume from external storage by using standard OpenStack functions, the metadata for the volume will be overwritten. To prevent this, perform the following procedure to restore a volume from external storage.

Procedure

1. Create a volume to which the data is to be restored. Note the LDEV metadata value, which can be acquired by using the **`cinder`** `metadata-show` command.
2. Restore the volume from external storage. If the volume ID of the restore target is not specified, or if the specified volume ID is different from the backed-up volume, the type and LDEV metadata for the restore target volume will not be consistent with the actual LDEV number.

3. Remove the metadata by using the **cinder** `metadata` command. Use the same command to restore the type and LDEV metadata values. For details about this command, see the help messages or the Command-Line Interface Reference.

Manage and unmanage a volume

If you plan to move a volume from one OpenStack Cinder instance to another, you can use the **cinder** `manage` and **cinder** `unmanage` commands. Specify `source-id <LDEV Number>` in the `--id-type` option.

Volume migration

When Block Storage Driver for OpenStack receives a Cinder migrate volume request, a virtual volume (DP-VOL) is created as a logical device (LDEV) in the pool specified as the migration destination, and then the data on the source volume is copied.

When volumes are migrated within the same storage system, storage-assisted copy processing is performed. When volumes are migrated between different storage systems, host-assisted copy processing is performed.

Chapter 2: System requirements

Before setting up your environment for use with Block Storage Driver for OpenStack, review the supported hardware, software, and license requirements.

Operating systems

For details about the supported operating systems, see the Hitachi Block Storage OpenStack driver support matrix at: https://support.HitachiVantara.com/en_us/interoperability.html.

OpenStack

For details about the supported OpenStack, see the Hitachi Block Storage OpenStack driver support matrix at: https://support.HitachiVantara.com/en_us/interoperability.html.

Storage devices

For details about the supported storage systems, see the Hitachi Block Storage OpenStack driver support matrix at: https://support.HitachiVantara.com/en_us/interoperability.html.

Storage firmware versions

The following table provides the Hitachi storage device firmware versions for each Block Storage Driver for OpenStack version.

Storage model	Firmware version
VSP F1500 VSP G1000, VSP G1500	80-05-43 or later
VSP F400, F600, F800 VSP G200, G400, G600, G800	83-04-43 or later
VSP F350, F370, F700, F900	88-01-04 or later

Storage model	Firmware version
VSP G350, G370, G700, G900	

Software

This section describes the software required for each version of the Block Storage Driver for OpenStack.

Prerequisite packages

The following table lists the prerequisite packages for Block Storage Driver for OpenStack.

Distribution	Node	Item	Contents
SUSE	Controller node	Linux environment	sysfsutils [1], sg3_utils [1], open-iscsi [2], multipath-tools [3]
	Compute node	Linux environment	sysfsutils[1], sg3_utils[1], open-iscsi[2], multipath-tools[3]
Ubuntu and Mirantis	Controller node	Linux environment	sysfsutils[1], sg3-utils[1], open-iscsi[2], multipath-tools[3]
	Compute node	Linux environment	sysfsutils[1], sg3-utils[1], open-iscsi[2], multipath-tools[3]

Notes:

1. For configuring FC.
2. For configuring iSCSI.
3. For configuring multipath.

For the supported storage platforms and the distribution releases, see: https://support.HitachiVantara.com/en_us/interoperability.html.

Storage management software

The following table shows the versions of Configuration Manager REST API that can be used to manage Hitachi storage devices.

Storage model	Configuration Manager REST API Version
VSP F1500	8.5.4-00 or later

Storage model	Configuration Manager REST API Version
VSP G1000, VSP G1500 VSP F400, F600, F800 VSP G200, G400, G600, G800	
VSP F350, F370, F700, F900 VSP G350, G370, G700, G900	8.5.6-00 or later

Storage software licenses

The following software licenses are required:

- Hitachi Storage Virtualization Operating System (SVOS)
  - Hitachi LUN Manager
  - Hitachi Dynamic Provisioning
- Hitachi Local Replication (Hitachi Thin Image)

The following software licenses are optional:

- Deduplication and compression
- Global-active device

Software restrictions

Do not use Hitachi Dynamic Link Manager (HDLM) in controller nodes or compute nodes. Use DM-Multipath instead.

Chapter 3: Installation and configuration

This chapter describes how to install and configure Block Storage Driver for OpenStack.

Installation and setup workflow

This topic describes the workflow for preparing, installing and configuring Block Storage Driver for OpenStack.

Prepare for installation

1. Connect nodes and the storage devices. See Connect nodes (on page 29).
2. Configure the node settings. See Configure node setting (on page 31).
3. Configure the node settings for storage resources used by Block Storage Driver for OpenStack. See Configure storage resources (on page 32).
4. Install and configure the Configuration Manager REST API server. See Configure the Configuration Manager REST API server (on page 37).

Deploy Block Storage Driver for OpenStack (Red Hat)

1. Deploy Block Storage Driver for OpenStack. See Deploy Block Storage Driver for OpenStack (Red Hat) (on page 39).
2. Configure Block Storage Driver for OpenStack. See Initial setup (on page 43).
3. Verify the installation. See Verify Block Storage Driver for OpenStack installation (on page 46).

Install Block Storage Driver for OpenStack (SUSE/Mirantis/Ubuntu)

1. Check the required memory size and disk space, and verify the OpenStack package installation. See Installation prerequisites (on page 41).
2. Install Block Storage Driver for OpenStack. See Install Block Storage Driver for OpenStack (SUSE/Mirantis/Ubuntu) (on page 41).
3. Configure Block Storage Driver for OpenStack. See Initial setup (on page 43).
4. Restart Block Storage Driver for OpenStack. See Restart Cinder (SUSE/Mirantis/Ubuntu) (on page 45).
5. Verify the installation. See Verify Block Storage Driver for OpenStack installation (on page 46).

Supported configurations for Block Storage Driver for OpenStack

Hitachi recommends an active-standby cluster configuration of OpenStack controller nodes and the Configuration Manager REST API servers. The following figure illustrates such a high-availability configuration.

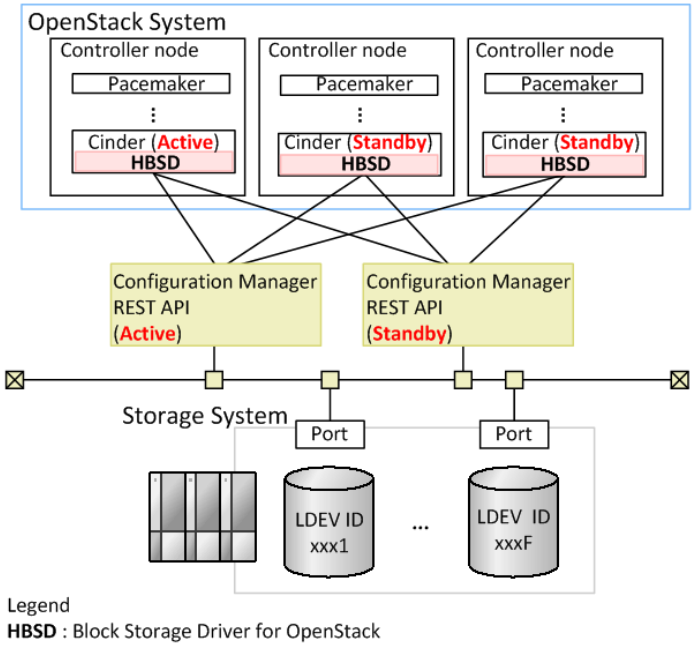

The following table describes valid OpenStack system configurations.

Note: Services other than the cinder-volume service can be used in an active-active configuration in accordance with OpenStack recommendations.

Configuration	Description
Single-controller configuration	This configuration has a single controller node.
Active-standby configuration	This configuration has multiple controller nodes. In this configuration, only one node on which the cinder-volume service is running is an active node. The other nodes are standby nodes. The storage system is shared amongst all nodes.

The following table describes the valid system configurations of the Configuration Manager REST API when used with Block Storage Driver for OpenStack.

Configuration	Description
Single-server configuration	This configuration has one Configuration Manager REST API server. You can install the Configuration Manager REST API on a controller node. For the prerequisites for installing the Configuration Manager REST API on a controller node, see *System Requirements* for Hitachi Command Suite.
Active-standby configuration	This configuration has multiple Configuration Manager REST API servers. Only one server is an active server, and the other servers are standby servers.

For details about high-availability cluster configurations of the Configuration Manager REST API, see the Hitachi Command Suite Configuration Manager REST API Reference Guide.

Management software restrictions

Block Storage Driver for OpenStack does not lock storage systems. For this reason, if storage resources are also registered with another management software, an error might occur due to management conflicts. To avoid such errors, you can use the following parameters to specify resources that can be used by Block Storage Driver for OpenStack: `hitachi_ldev_range` to specify a range of LDEV IDs.

Do not use another management software to perform operations on the resources specified by this parameter.

You can use Configuration Manager REST API to set up resource groups to configure access control on resources managed by various management software. For details, see the description of resource groups in the Hitachi Command Suite Configuration Manager REST API Reference Guide.

Environment setup

It is necessary to configure the storage device to allow Block Storage Driver for OpenStack to use it. Refer to the documentation provided with the storage device for details.

Connect nodes

Connect the controller nodes and compute nodes to the ports of the storage device. The compute nodes use the ports of the storage device as data paths. You can use Block Storage Driver for OpenStack to manage both configurations in which controller nodes and compute nodes use the same ports of a storage device, and configurations in which controller nodes and compute nodes use different ports. Design the connection environment based on your operation policy.

The following figure shows a Fibre Channel (FC) connection example:

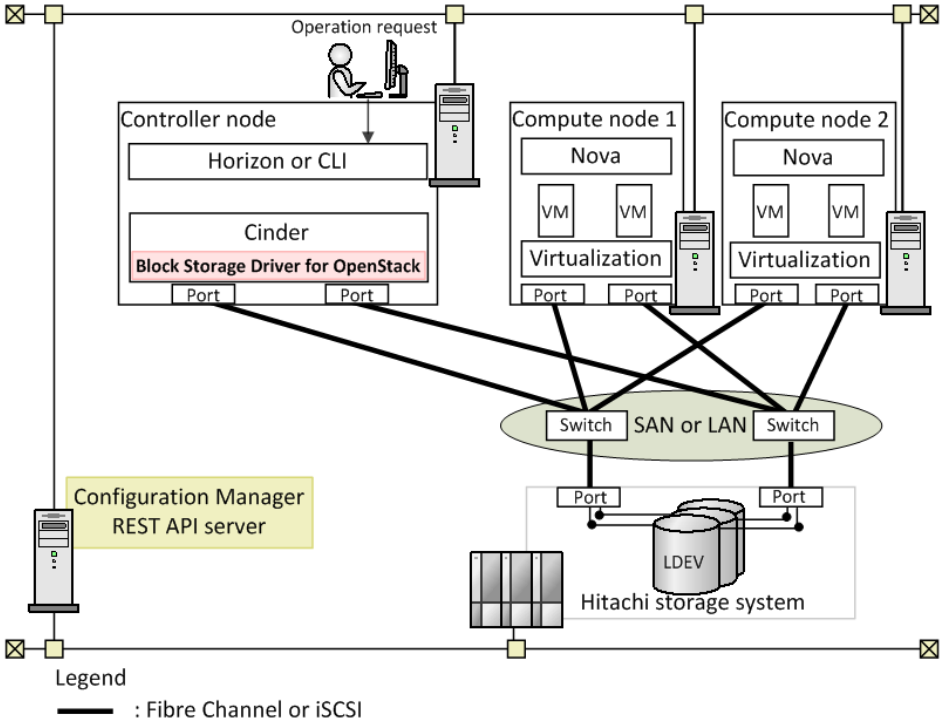

Figure 5 Connection example of node configuration

The ports of a storage device can be shared with other management software.

Single node configuration

Block Storage Driver for OpenStack allows single node configurations where a single server is configured to serve the role of the controller node and the compute node. There are a few operation restrictions to note for a single node environment. These operations cause I/O disruptions and require you to detach a volume to perform the operation, and reattach the volume.

If a DP-VOL is attached, you cannot copy the volume to images.

Controller node without I/O data paths

You can use the configuration on the Controller node without I/O data paths, as shown in the following image. This configuration is a viable solution if you do not want to use Fibre Channel ports for management.

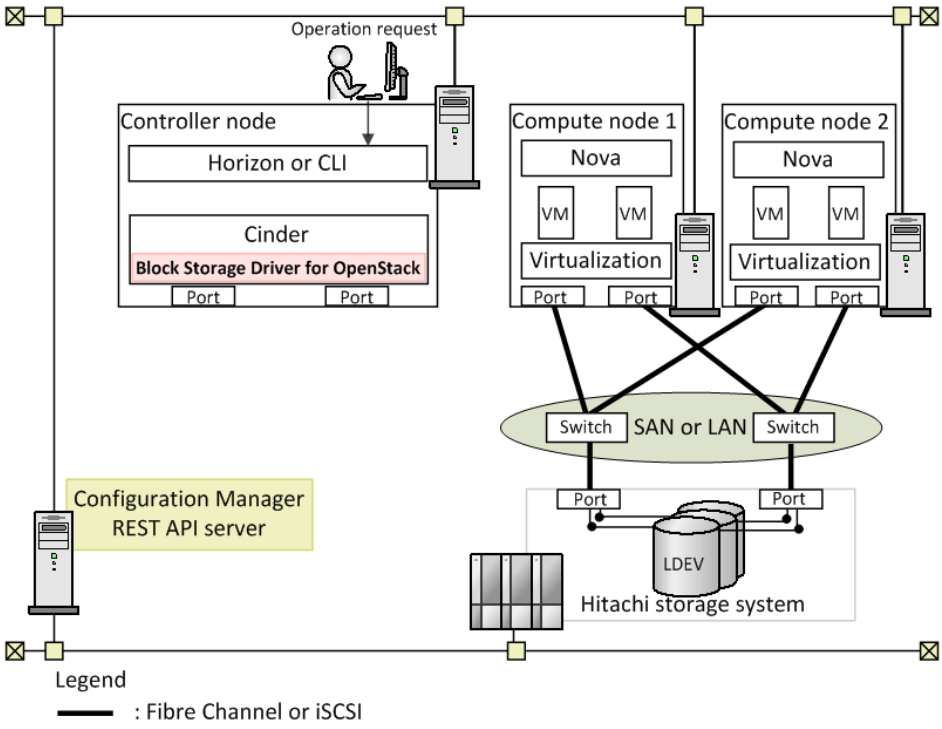

Figure 6 Connection example without I/O data paths for controller node

Use the following settings to configure the controller node without I/O data path.

- Specify the `hitachi_compute_target_ports` parameter.
- Specify the `hitachi_rest_pair_target_ports` parameter.
- Do not specify the `hitachi_target_ports` parameter.

The following operations are not supported by the controller node without I/O data paths:

- Perform a volume migration (host assisted)
- Create a volume from an image
- Create an image from a volume
- Back up a volume
- Restore a backup
- Retype a volume

Configure node setting

You can use the OpenStack configuration installer provided by each distribution to build the controller node and compute nodes beforehand. It is recommended to use a multi-node configuration in which the controller node and the compute nodes use separate data paths. Some OpenStack configuration installers might require you to specify a Cinder driver when deploying a Cinder component. When prompted, specify another driver such as NFS (as you cannot select Block Storage Driver for OpenStack at this time) and install the Cinder driver after the deployment. Refer to online resources for the usage of the OpenStack configuration installer provided by each distribution. After building the controller node and the compute nodes, configure each node as described in the following table.

Node	Configuration item	Instruction
Controller node	my_ip for cinder service (`/etc/cinder/cinder.conf`)	Specify a unique value for the IPv4 address for the management LAN of the node.
	Initiator IQN (`/etc/iscsi/initiatorname.iscsi`)	Specify a unique value for the Initiator IQN. This item is for an iSCSI connection.
Compute node	my_ip for nova compute service (`/etc/nova/nova.conf`)	Specify a unique value for the IPv4 address for the management LAN of the node.
	Initiator IQN (`etc/iscsi/initiatorname.iscsi`)	Specify a unique value for the Initiator IQN. This item is for an iSCSI connection.

About the communication mode of the Configuration Manager REST API server

The Block Storage Driver for OpenStack uses Configuration Manager REST API server for storage operations. The Block Storage Driver for OpenStack supports both in-band and out-of-band methods.

In the in-band management method, the Configuration Manager REST API server connects to the storage system through a SAN.

In the out-of-band management method, the Configuration Manager REST API server connects to the storage system through a LAN. This is the default communication mode for the Configuration Manager REST API server.

If the out-of-band method is used to connect the Configuration Manager REST API server and the storage system, set the `isSecure` property to `True` to encrypt communications between the Configuration Manager REST API server and the storage system. For the setting, see Changing the communication mode of the REST API server in Hitachi Command Suite Configuration Manager REST API Reference Guide.

Configure storage resources

Common resources

The following table describes resource requirements for storage systems:

Resource type	Contents
All resources	All storage resources, such as DP pools and host groups, must have a name in order for the driver to use them (name fields cannot be left blank).
LDEVs	If a virtual LDEV number is assigned to an LDEV used by Block Storage Driver for OpenStack, volume operations might fail. If you specify the `hitachi_ldev_range` parameter, do not assign a virtual LDEV number to any of the LDEVs in the range specified by the `hitachi_ldev_range` parameter. If you do not specify `hitachi_ldev_range` parameter, do not assign a virtual LDEV number for any of the LDEVs used by Block Storage Driver for OpenStack. However, if you want to use a global-active device (GAD) in a "P-VOL registered to a VSM" configuration, you need to assign a virtual LDEV number for the primary volume (P-VOL). For details, see Cinder operations in a global-active device (GAD) configuration (on page 61) .
Resource groups	If using a new resource group for exclusive use by an OpenStack system, create a new resource group, and assign the necessary resources, such as LDEVs (including LDEVs for a command device or pool), port, and host group (iSCSI target) to the created resource group.

Resource type	Contents
User accounts	▪ If you are using a backend storage system for which `REST` is specified for the parameter `hitachi_storage_cli`, the user account must be registered to the Administrator user group (a built-in user group). Even if a user account was created by using another management tool, the account can be used if it is registered to a user group to which all of the following roles have been assigned: - Storage Administrator (View Only) - Storage Administrator (Provisioning) - Storage Administrator (Local Copy) - Storage Administrator (Performance Management) Permissions for all resource groups must also be granted to the user account. You can also grant permissions for all resource groups to the user group that you want to use by assigning any of the following roles to the user group. - Audit Log Administrator (View & Modify) - Audit Log Administrator (View Only) - Security Administrator (View & Modify) - Security Administrator (View Only) - Support Personnel - User Maintenance ▪ If you are using a backend storage system for which `SIMPLE_REST` is specified for the parameter `hitachi_storage_cli`, the user account must be registered to the Administrator user group (a built-in user group) and the Maintenance user group (a built-in user group). Even if a user account was created by using another management tool, the account can be used if it is registered to a user group to which all of the following roles have been assigned: - Storage Administrator (View Only) - Storage Administrator (Provisioning) - Storage Administrator (Local Copy) - Storage Administrator (Performance Management) - User Maintenance

Resource type	Contents
DP Pool	▪ Create a DP pool that is used by Block Storage Driver for OpenStack backend. The driver manages virtual capacity (capacity reserved for over-provisioning of the actual pool capacity). However you can set the percentage of DP-VOL capacity that can be created to the total capacity of the pool as the subscription limit. This can help prevent DP-VOL blocking caused by a full pool. For example, when the subscription limit is set to 100%, the total DP-VOL capacity is equal to the DP pool capacity. To set the subscription limit, set the `hitachi_over_subscription` parameter to `True` and specify the ratio by the `max_over_subscription_ratio` parameter if the storage system model is VSP F350, F370, F700, F900, VSP G350, G370, G700, or G900. Otherwise, set the subscription limit on the storage system. ▪ Use a pool that is not assigned to a virtual storage machine. You cannot use DP-VOLs that are assigned to a virtual storage machine or volumes that have been migrated from another system using nondisruptive migration.
TI pool	Normally, you do not need to create a TI pool. In this case, differential data is stored in the DP pool. If you have a particular reason to store differential data for volumes and snapshots in a dedicated pool, create a TI pool.

Fibre Channel connectivity

The following table describes Fibre Channel (FC) storage configuration requirements.

Resource type	Contents
Ports	▪ Enable Port Security for the port used by the driver. ▪ To use the port scheduler function, enable fabric, and then set Point to Point (P to P) for Connection Type. ▪ If you change the port configuration used by the driver for connection, restart the `openstack-cinder-volume` service. If the volume is attached or detached without restarting the service, the operation might fail.

Resource type	Contents
Host group for I/O data path	Manual configuration: ▪ Create a host group for the port that connects the controller node and compute nodes to the storage system. In a multipath configuration, create host groups for all connecting ports. ▪ For a multipath configuration, the LDEV numbers and the connection order (LUN) must match within each host group of the target node. If they do not match, change each LDEV connection accordingly to match each LUN. ▪ The host group must be named `HBSD-<target host WWPN>` (for example, `HBSD-0123456789ABCDEF`). ▪ Specify `HBSD-<target host WWPN>` using the WWN of the smallest WWPN used for the connecting node. ▪ Register the WWN of the connecting node in the newly created host group. ▪ If the host groups for a controller node and compute node are created, reboot the nodes to recognize these paths. If a new compute node is added and the host group for the node is created, reboot is necessary for only this node. Automatic configuration: ▪ To enable automatic creation of the host groups, set the `hitachi_group_request` parameter to be `True`. In addition, to add LDEVs to a host group being created at the same time, Fibre Channel switches must be used to connect the storage and the node, and the Fibre Channel Zone Manager must be enabled. ▪ After receiving a cinder volume attach request, the driver automatically creates host groups on ports specified by the `hitachi_target_ports` parameter or the `hitachi_compute_target_ports` parameter if there is no host group that is connected to the VM instance. The host group name is `HBSD-<host WWPN>`. The driver registers all host WWPNs with the created host group. ▪ Automatic deletion of host groups requires the `hitachi_group_delete` parameter to be set to `True`.

Resource type	Contents
Fibre Channel zoning	Manual configuration: ▪ For Fibre Channel switches between a node and the storage device, manually configure zoning. Automatic configuration: ▪ When the driver attaches a volume to the target VM, Fibre Channel Zone Manager automatically creates zoning between the target node and the port used by the driver. ▪ Enable Fibre Channel Zone Manager and set the `hitachi_zoning_request` to `True`.

iSCSI connectivity

The following table describes iSCSI storage configuration requirements.

Resource type	Contents
Ports	▪ Enable port security for the ports used by the driver. ▪ Assign an IP address and a TCP port number to the port. ▪ If you change the port configuration used by the driver for connection, restart the `openstack-cinder-volume` service. If the volume is attached or detached without restarting the service, the operation might fail.
iSCSI target for I/O data path	Manual configuration: ▪ Create an iSCSI target for the port that connects the controller node and compute nodes to the storage system. In a multipath configuration, create iSCSI targets for all connecting ports. ▪ The iSCSI target must be named `HBSD-<my_ip>` (for example, `HBSD-10.20.30.40`). Use the same IP address for *my_ip* that you use for a cinder or nova compute service in each node. ▪ Register the Initiator IQN of the connecting node in the newly created iSCSI target. ▪ Register the Target IQN in the newly created iSCSI target. In a multipath configuration, register the same target IQN in all iSCSI targets connected to a node.

Resource type	Contents
	▪ In a multipath configuration, use the storage system management software to set the host mode option to `83`. ▪ If CHAP authentication is used, register the CHAP user name and password with the iSCSI target. Use the same CHAP user and password for all iSCSI targets. Automatic configuration: ▪ Automatic creation of iSCSI targets requires the `hitachi_group_request` parameter to be `True`. ▪ Upon a cinder volume attach request, the driver automatically creates iSCSI targets on ports specified by the `hitachi_target_ports` parameter or the `hitachi_compute_target_ports` parameter if there is no iSCSI target that is connected to the VM instance. The target IQN of the created iSCSI target is *<Initiator IQN>*`.hbsd-target`. The driver registers the initiator IQN with the created iSCSI target. ▪ Automatic deletion of iSCSI targets requires the `hitachi_group_delete` parameter to be `True`. Considerations for multipath configurations: ▪ For a multipath configuration, the LDEV numbers and the connection order (LUN) must match within each iSCSI target of the target node. If they do not match, change each LDEV connection accordingly to match each LUN. ▪ If the settings are automatic, the driver will use the smallest available iSCSI target number. Align the LUNs that are already being used for each port to force the LUNs created by the driver to match.

Configure the Configuration Manager REST API server

Use the procedure below to configure the Configuration Manager REST API server.

Before you begin

Before using Block Storage Driver for OpenStack, start the Configuration Manager REST API services. Additionally, if your site has installed the Configuration Manager REST API server on the Controller node, then create a service startup sequence in the OS. This ensures that the Configuration Manager REST API services have started before the cinder-volume services start.

Procedure

1. Install the Configuration Manager REST API server.

   For details, see *installing the REST API* in the Hitachi Command Suite Configuration Manager REST API Reference Guide.

2. Register information about the storage system to the Configuration Manager REST API server.

   For details, see *Registering a storage system* in the Hitachi Command Suite Configuration Manager REST API Reference Guide.

```
# cat storage.json
{
   "model" : "VSP F400",
   "serialNumber" : 411111,
   "svpIp" : "10.197.74.6"
}
# curl -H "Accept:application/json" -H "Content-Type:application/json" -u USER_ID:PASSWORD -X POST --data-binary
@./storage.json https://192.0.2.100:23451/ConfigurationManager/v1/objects/storages
{
 "storageDeviceId" : "834000411111",
 "model" : "VSP F600",
 "serialNumber" : 411111
}
```

Using high-speed REST API requests

If you are using a VSP F350, F370, F700, F900, or VSP G350, G370, G700, G900 storage system, you can perform some storage system operations in Block Storage Driver for OpenStack by using REST API requests that enable faster processing than regular REST API requests. These API requests, which get information in a simple format and return responses at a high speed, can be used when volumes are attached or when volumes are created (as clones or snapshots). For details about these REST API requests, see the *Hitachi Storage Advisor Embedded Guide*.

If you want to use high-speed REST API requests for Block Storage Driver for OpenStack, specify `SIMPLE_REST` for the `hitachi_storage_cli` parameter.

Note: You can specify `REST` or `SIMPLE_REST` for the `hitachi_storage_cli` parameter. However, after you specify a value and start operation, you will not be able to change the value, because different storage systems have different management models.

Secure your storage system password

Your password file should be kept out of the file directories that are subject to data collection for logs. Follow the procedure below to keep your password from being shared with unauthorized users.

Procedure

1. As a root user, create a password file using a text editor. For example:

   File: `/opt/hitachi/hbsd/rest_password`

   Password: *!gh8M59Wz#*

2. Confirm the password.

```
# cat /opt/hitachi/hbsd/rest_password
!gh8M59Wz#
```

3. Hide the password file from all users except the Cinder group.

```
# sudo chgrp cinder /opt/hitachi/hbsd/rest_password
# sudo chmod 640 /opt/hitachi/hbsd/rest_password
```

4. Specify the `hitachi_rest_password_path` parameter in the `cinder.conf` file (`/etc/cinder/cinder.conf`). For example, `hitachi_rest_password_path=`*`/opt/hitachi/hbsd/rest_password`*.

   Note: Do not use the `hitachi_rest_password` parameter. The `hitachi_rest_password` and `hitachi_rest_password_path` parameters are mutually exclusive and using `hitachi_rest_password` will conflict with the password file path you have added to the `cinder.conf` file.

Deploy Block Storage Driver for OpenStack (Red Hat)

In Red Hat OpenStack Platform 13 and later versions, RPM packages are no longer installed on controller nodes. Instead, Cinder components (including various drivers) are deployed as a Docker image to create a virtual machine on which Cinder runs. As a result, you can now use Block Storage Driver for OpenStack by deploying the container image of Block Storage Driver for OpenStack.

Use the `openstack overcloud deploy` command to deploy Block Storage Driver for OpenStack. For details about using this command, see the online documentation on the Red Hat website.

The following is a sample deployment procedure.

Procedure

1. Log in to `registry.connect.redhat.com` by using the command `docker login`.
2. Use the command `docker pull` to download the container image of Block Storage Driver for OpenStack from the following registry.

   `registry.connect.redhat.com/hitachi/rhosp14-openstack-cinder-volume-hbsd`

3. Use the commands `docker tag` and `docker push` to add a tag for the container image of Block Storage Driver for OpenStack to the Local Registry and to push the tagged Block Storage Driver for OpenStack container image to the Local Registry.
4. Specify the container image of Block Storage Driver for OpenStack on the Red Hat website.

   Edit `containers-prepare-parameter.yaml` in the following location to exclude the container image of the Cinder volume and to include the container image of Block Storage Driver for OpenStack.

   File: `/home/stack/templates/containers-prepare-parameter.yaml`

   For details about how to specify settings for `containers-prepare-parameter.yaml`, see the Red Hat document *Red Hat OpenStack Platform Director Installation and Usage*.
5. Create an environment file for Block Storage Driver for OpenStack.

   To apply the settings of Block Storage Driver for OpenStack to a Cinder component during deployment, create a new environment file (for example, `custom_env.yaml`) and then edit the file.

   File: `/home/stack/templates/custom_env.yaml`

   For information about how to set the environment file, see the Red Hat document *Red Hat OpenStack Platform Custom Block Storage Back End Deployment Guide*.
6. Deploy the container image.

   Run the `openstack overcloud deploy` command to deploy the components. When you run the command, specify the following environment setting files: `containers-prepare-parameter.yaml` and the environment file that you created in step 5.

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

Install Block Storage Driver for OpenStack (SUSE/Mirantis/Ubuntu)

Installation prerequisites

The installation prerequisites for Block Storage Driver for OpenStack are listed below.

- Verify that you have sufficient memory on the server. The physical memory size required by the software is: 1 MB x the number of backends x 2.
- Verify that you have sufficient disk space required to install the driver. The following table describes the required disk space.

Item	Size	Location
Block Storage Driver for OpenStack file	1 MB	`/usr/share/pyshared/hitachi-hbsd`
hbsdgetinfo	1 MB	`/opt/hitachi/hbsd/`
Log files	100 MB	`/var/log/cinder/`

- Check the OpenStack package installation for any missing or invalid packages (see Software (on page 24)). Use the **rpm** command with the `-V` option specified (debsums for Ubuntu) to confirm that these packages are installed and are in a normal state. If a package is not installed or an invalid package is found, install the package again.

Block Storage Driver for OpenStack installation for SUSE

Use the **rpm** command to install Block Storage Driver for OpenStack. You must log in as a superuser (root) on the controller node where you want to perform the installation. For details about how to use the **rpm** command, see the online documentation on the Linux website.

Procedure

1. Use the **mount** command to mount the CD file system when installing from a CD:

   ```
   # /bin/mount -r -o mode=0544 /dev/cdrom /media
   ```

2. Before installing the driver, stop the Cinder volume service.

   ```
   # /sbin/service openstack-cinder-volume stop
   ```

3. After the service is stopped, run the **rpm** command to install the driver.

   Use the following RPM package:

   RPM package name: `hbsd-7.0.0-0_13.0.noarch.rpm`

RPM storage directory: `RPMS/noarch`

An example of installing Block Storage Driver for OpenStack:

```
# /bin/rpm -ivh /media/RPMS/noarch/<RPM package name>
```

4. After the installation is completed, use the **umount** command to unmount the CD.

```
# /bin/umount/media
```

Block Storage Driver for OpenStack installation for Mirantis and Ubuntu

Use the **dpkg** command to install Block Storage Driver for OpenStack. You must log in as a superuser (root) on the controller node where you want to perform the installation. For details about how to use the **dpkg** command, see the Mirantis or Ubuntu online documentation.

Procedure

1. Use the **mount** command to mount the CD file system when installing from a CD:

```
# /bin/mount -r -o mode=0544 /dev/cdrom /media
```

2. Before installing the driver, stop the cinder-volume service.

```
# /usr/sbin/service cinder-volume stop
```

3. After the service is stopped, run the **dpkg** command to install the driver.

   Use the following DEB package:

   DEB package name: `hbsd_7.0.0-0-13.0_all.deb`

   DEB storage directory: `DEBS/all`

   An example of installing Block Storage Driver for OpenStack:

```
# /usr/bin/dpkg -i /media/DEB/UBUNTU/all/<DEB package name>
```

4. After the installation is completed, use the **umount** command to unmount the CD.

```
# /bin/umount/media
```

Initial setup

Add a Block Storage Driver for OpenStack configuration (SUSE/ Mirantis/Ubuntu)

Normally, a Block Storage Driver for OpenStack configuration is added to the configuration file `/etc/cinder/cinder.conf` provided by an openstack-cinder package. However, some OpenStack configuration installers might automatically change the configurations of `cinder.conf` by running configuration management software such as Puppet. The following describes the procedure required in such situations.

Using Crowbar (SUSE OpenStack Cloud)

A Chef agent runs on the controller node when an OpenStack environment is built with Crowbar. Chef resets `cinder.conf` for the controller node, and any configuration you have built in the driver will be lost. To resolve this issue, follow the procedure below to add your driver configuration to the **Add New Cinder Backend** page in the Crowbar WebGUI.

Procedure

1. Launch the WebGUI of Crowbar, and select **OpenStack** in the **Barclamps** menu.
2. Select **Edit** in Cinder.
3. Select **Other Driver** in the **Type of Volume** menu in the **Add new Cinder Backend**. Click **Add Backend** to enter the name of the backend.
4. Enter the driver path in the **Driver** and the parameters in the **Options** for the backend you have added.
5. In the **Deployment**, specify the controller node (only one selectable) in the **cinder-controller**. In addition, specify the backend in the **cinder-volume**.
6. Click **Apply** to save the configuration.

Using FUEL (Mirantis OpenStack)

To configure Block Storage Driver for OpenStack using Mirantis OpenStack, edit the configuration file (`/etc/cinder/cinder.conf`) on the controller node.

Using MAAS (Ubuntu OpenStack)

To configure the driver using Ubuntu, edit the configuration file (`/etc/cinder/cinder.conf`) on the controller node.

Linking a volume type and a backend

For creating an environment with multiple backends, you can create a volume type and link it to a particular backend.

You can configure the backend specification (`volume_backend_name`) during backend setup by using the **`cinder`** command in the controller node where a Cinder client is installed.

```
 # /usr/bin/cinder type-create <volume type name>
 # /usr/bin/cinder type-key <volume type name> set
volume_backend_name=<volume backend name>
```

The following commands show an example of creating a volume type called `HBSD_SAMPLE1` and linking it to the backend `hbsd_backend1`.

```
 # /usr/bin/cinder type-create HBSD_SAMPLE1
 # /usr/bin/cinder type-key HBSD_SAMPLE1 set
volume_backend_name=HBSD_backend1
```

Use the **`cinder`** `extra-specs-list` command to check the volume type you created and the value registered for the `volume_backend_name` key.

```
# /usr/bin/cinder extra-specs-list
+-----------------+--------------+-----------------------------------------+
|     ID          |     Name     |                extra_specs              |
+-----------------+--------------+-----------------------------------------+
|<Volume Type ID> | HBSD_SAMPLE1 | {u'volume_backend_name': u'hbsd_backend1'}|
+-----------------+--------------+-----------------------------------------+
```

Edit cinder.conf (SUSE/Mirantis/Ubuntu)

A driver configuration is added using the OpenStack configuration installer provided by the distribution in use (see Add a Block Storage Driver for OpenStack configuration (SUSE/Mirantis/Ubuntu) (on page 43)). This topic describes how to edit the configuration file `/etc/cinder/cinder.conf` provided by the openstack-cinder package.

Specify the driver settings in the configuration file `/etc/cinder/cinder.conf`. You need to manually update the configuration file with values that reflect the latest storage configuration. Specify "*`parameter=value`*" pair per line in the following format:

```
parameter=value
```

Specify the driver-specific parameters and Cinder-provided parameters for each backend. An example Block Storage Driver for OpenStack configuration is shown below, followed by the description of each section of the file. For details about the parameters, see Parameters (on page 65) .

```
# cinder.conf sample #
[DEFAULT]
: (Omitted)
enabled_backends=VSPG1000
logging_context_format_string=%(asctime)s.%(msecs)03d % (process)d %
(thread)s
%(levelname)s %(name)s [%(request_id)s %(user_identity)s] % (instance)s%
(message)s
: (Omitted)
[VSPG1000]
volume_driver=cinder.volume.drivers.hitachi.hbsd.hbsd_fc.HBSDFCDriver
volume_backend_name=hbsd_backend1
suppress_requests_ssl_warnings=True
hitachi_storage_id=800000054321
hitachi_pool=0
hitachi_rest_api_ip=192.0.2.100
hitachi_rest_user=user
hitachi_rest_password_path=/opt/hitachi/hbsd/rest_password
hitachi_target_ports=CL1-A,CL2-A
hitachi_compute_target_ports=CL1-B,CL2-B
```

Note:

- If you want to share a port between multiple backends in a multipath configuration, specify the same port name for all the ports used by the backends. If the same port name is specified for only some of the backends, volume-connection operations might fail.
- If you use the Configuration Manager REST API server's default server authentication, set `suppress_requests_ssl_warnings=True`.
- If you use attach operations for multiple volumes, it is recommended that you set `rpc_response_timeout` to a larger value.

Restart Cinder (SUSE/Mirantis/Ubuntu)

To complete the configuration, use the **`cinder`** command to restart the cinder-volume service.

Run the following commands to restart the cinder-volume service.

- For SUSE, use the **`service`** command:

```
 # /sbin/service openstack-cinder-volume restart
Currently stopped openstack-cinder-volume: [ OK ]
Currently starting openstack-cinder-volume: [ OK ]
```

- For Ubuntu and Mirantis, use the **`service`** command:

```
# /usr/sbin/service cinder-volume restart
Currently stopped cinder-volume: [ OK ]
Currently starting cinder-volume: [ OK ]
```

Verify Block Storage Driver for OpenStack installation

This topic describes the methods that you can use to verify installation and configuration of Block Storage Driver for OpenStack.

Procedure

1. After the Cinder service restarts, make sure that Block Storage Driver for OpenStack is available for use by verifying the following:
   - The message `MSGID0003-I` is output to the following log file.

     Red Hat OpenStack Platform: `/var/log/cinder/volume.log`

     SUSE OpenStack Cloud, Mirantis OpenStack, and Ubuntu OpenStack: `/var/log/cinder/cinder-volume.log`

     ```
      # /bin/grep MSGID0003-I /var/log/cinder/volume.log
     2018-09-21 13:53:06.067 137596 WARNING
     cinder.volume.drivers.hitachi.hbsd.hbsd_utils [req-6d5ebee0-8940-
     43b0-b3d4-1d320c2ad447 - - - - -] MSGID0003-I: The storage backend
     can be used. (config_group: VSP_G1000)
     ```

   - The state of the cinder-volume service of the backend is `up`. This can be displayed by the **`cinder`** `service-list` command as shown in the following example. The backend whose state is `up` can be used for volume operations. If the state is `down`, check the log file for relevant messages.

     ```
     # /usr/bin/cinder service-list
     +------------------+-------------------------------+------+---------+--------+----------------------------+
     |      Binary      |        Host                   | Zone | Status  | State  |         Updated_at         |
     +------------------+-------------------------------+------+---------+--------+----------------------------+
     |  cinder-backup   | hostname.localdomain          | nova | enabled |   up   | 2014-07-07T12:31:10.000000 |
     | cinder-scheduler | hostname.localdomain          | nova | enabled |   up   | 2014-07-07T12:31:12.000000 |
     |  cinder-volume   | hostname.localdomain@VSPG1000 | nova | enabled |   up   | 2014-07-02T02:55:47.000000 |
     +------------------+-------------------------------+------+---------+--------+----------------------------+
     ```

2. Verify that the following volume operations can be performed from the Horizon portal:
   - Create a volume
   - Delete a volume
   - Attach a volume
   - Detach a volume
   - Create a snapshot
   - Delete a snapshot
   - Create a volume from snapshot
   - Create a volume from volume (clone)

   If any of these volume operations fail, check the error messages output to the OpenStack Cinder log file.

Update the Block Storage Driver for OpenStack

Update the Block Storage Driver for OpenStack software for Red Hat

You can update the container image by running the commands `openstack overcloud update prepare` and `run`. In the same way as in the procedure to deploy Block Storage Driver for OpenStack, run the `openstack overcloud update prepare` command with the environment file for Block Storage Driver for OpenStack specified.

For a detailed update procedure, see the Red Hat document *Red Hat OpenStack Platform Keeping Red Hat OpenStack Platform Updated*.

Update the Block Storage Driver for OpenStack software for SUSE

Use the procedure described below to update the driver software.

Procedure

1. Stop the cinder-volume service.

```
# /sbin/service openstack-cinder-volume stop
```

2. After the service has stopped, install a new version of the driver as shown below:

```
# /bin/rpm -Uvh /media/RPM-storage-directory/RPM-package-name
```

   For the RPM package name and RPM storage name, see Block Storage Driver for OpenStack installation for SUSE (on page 41).

3. After updating the driver software, check the parameters in the `cinder.conf` file and update them as needed.

   For details about each parameter, see Parameters (on page 65).

Update the Block Storage Driver for OpenStack software for Mirantis and Ubuntu

Use the procedure described below to update the driver software for Mirantis and Ubuntu.

Procedure

1. Stop the cinder-volume service.

```
# /usr/sbin/service cinder-volume stop
```

2. After the service stops, run the **dpkg** command to install the driver.

   For the DEB package name and DEB storage name, see Block Storage Driver for OpenStack installation for Mirantis and Ubuntu (on page 42).

3. After updating the driver software, check the parameters in the `cinder.conf` file and update them as needed.

   For details about each parameter, see Parameters (on page 65).

Uninstall Block Storage Driver for OpenStack (SUSE/Mirantis/Ubuntu)

Follow the procedure below to uninstall the driver.

Procedure

1. Stop the cinder-volume service.
   - For SUSE:

```
# /sbin/service openstack-cinder-volume stop
```

   - For Ubuntu and Mirantis:

```
# /usr/sbin/service cinder-volume stop
```

2. Uninstall the driver.
   - For SUSE:

```
# /bin/rpm -e hbsd
```

   - For Ubuntu and Mirantis:

```
# /usr/bin/dpkg -purge hbsd
```

Live migration

The Cinder live migration capability allows you to migrate a VM to another host for hardware maintenance and other purposes without disrupting workloads. See online documentation such as *OpenStack Cloud Administrator Guide* for details about the live migration function on Nova components and the usage of the OpenStack configuration installer provided by the distribution in use.

When you perform live migration using Block Storage Driver for OpenStack, follow these guidelines:

- Configure DM-Multipath in compute node by editing the configuration file (for example, `/etc/multipath.conf`) and then restarting the multipath daemon. See Configure multipathing (on page 49) .
  - `user_friendly_names` is set to "*no*"
  - `find_multipaths` is set to "*no*"

  DM-Multipath is required for performing live migration for both single and multipath configurations.

Note: Volume migration finishes even if it fails to delete the LDEV of the original volume. In this case, delete it manually because the LDEV of the original volume remains in storage.

Configure multipathing

Redundant I/O data paths to the storage device enable continuous operations when one of the paths fails. Device mapper multipathing (DM-Multipath) provides path redundancy in OpenStack as illustrated in the figure in Connect nodes (on page 29) .

DM-Multipath creates and attaches a multi-path device (dm-X) to the VM.

The details of multipath configuration for controller and compute nodes are described below. See Cinder-specific parameters (on page 75) for the parameter details. In addition, see the Linux website for documentation on setting up DM-Multipath.

Controller nodes:

- Set `True` for the `use_multipath_for_image_xfer` parameter.
- Configure DM-Multipath for the port specified by the Block Storage Driver for OpenStack parameter `hitachi_target_ports` and `hitachi_compute_target_ports`.
- Start multipath daemon after configuring the DM-Multipath setting.

Compute nodes:

- Set `True` for the `volume_use_multipath` parameter. List this parameter in the libvirt section in the Nova configuration file.
- Start the multipath daemon after configuring the DM-Multipath setting.

Note:

- The recommended value for the DM-Multipath varies depending on the storage device. Refer to multipath configuration guidelines of your storage device.
- Configure the live-migration parameters to use live migration (see Live migration (on page 49)).
- Attach or detach volume operation of the volume may fail when one of the iSCSI paths fails. However, VMs that the volume is attached to can be used with the other path.

Configure high availability

High availability configuration overview

You can configure multiple controller nodes in a cluster to ensure high availability. If one node fails, the system fails over to another node to continue operations without any disruption to the workloads. Each service can select active-standby or active-active configuration with Pacemaker.

However, active-standby is the only controller-node fail-over configuration supported by Block Storage Driver for OpenStack.

Pacemaker configurations for Cinder volumes

Block Storage Driver for OpenStack supports only active-standby configuration for Cinder volumes, and fencing actions reboot when the system fails over. To allow this, configure Pacemaker parameters for cinder volumes as described in the following table.

Parameter	Value
monitor	fencing
stop	fencing
start	fencing
start-delay	10s
clone-set	Delete this item

Configure the Pacemaker configuration using the **pcs** commands. You can reconfigure node behaviors after the node has been rebooted.

Build high availability environment of controller node

For information about how to build high availability environment, see the Red Hat document *Red Hat OpenStack Platform Understanding Red Hat OpenStack Platform High Availability*.

Recoverable failures

When the system detects a storage control path failure that causes the Cinder volume service to abend, it switches the failed component to the normal one to continue the operations. Other types of failures might not be detected and as a result the Cinder volume service becomes unavailable.

When the failed system switches over, the node associated with the failed system will reboot. During the reboot, all services on the node will become unavailable. You can reconfigure the node behaviors after it has been rebooted. If you reconfigure the node to be added back to the cluster, the Active-Active services can be used as Active again.

Note: If the failed node is recovered and you want to switch the service back to the recovered node, make sure that there are no running snapshot volumes on the failover node by using the **`cinder`** `snapshot-list volume` command before you switch the node back.

Handling high availability post-failover

When a controller node fails while storage operations are in progress, the operation target volumes might be left unprocessed in the storage. This may cause other operations that are involved with these resources to fail. Follow the procedure listed below to determine the correct actions.

Procedure

1. Check Pacemaker logs for the failover status and the time it has occurred.
2. If the failover caused Cinder volume processing to stop, use the **`cinder`** `list` command to find the volumes or snapshot being processed.
3. If the status of a volume is "creating" then go to step 4. Otherwise, perform the following steps:
   a. Get the metadata of the volume or the snapshot by using the **`cinder`** `metadata-show` command and identify the LDEV number.
   b. Determine the storage operation from the volume status.
   c. Based on the obtained LDEV number, in the storage system management software, ensure that the volume is restored to the status it was in before the operation.
   d. Change the status for the corresponding entry on CinderDB to "available" using the **`cinder`** `reset-state` command.
   e. Go to step 6.
4. Identify the incomplete volume processing from the file `volume.log` or `debug.log` (for Red Hat, the file `hbsd_debug.log`) around the time when the failover occurred.

5. Check whether any LDEV is associated with the volume. If no LDEV is associated, go to the following step d. If an LDEV is associated, perform the following steps:
   a. Obtain the pair information from the LDEV ID by using the management software. If a pair exists, delete it.
   b. If the LDEV is mapped to a host group or an iSCSI target, disable the mapping.
   c. Delete the LDEV by using the storage management software.
   d. Change the state of the volume to "available" using the **`cinder`** `reset-state` command and delete the corresponding entry on CinderDB using the **`cinder`** `delete` command.
6. If there are any other volumes with incomplete processing, repeat steps 3 through 5.

Chapter 4: Hitachi extended functions

This section provides details and examples for using the extended storage operation capabilities provided by Hitachi storage devices.

Copy methods

When cloning volumes or creating snapshots, you can choose from one of the two copy methods, Thin Image (Clone) and Thin Image (Snapshot), provided by storage systems. When creating snapshots, the copy method is set to Thin Image (Snapshot). The following table describes which copy method is used under each setting. The copy method is selected by the `hitachi_default_copy_method` parameter and the `copy_method` metadata for the destination volume. For example, setting pattern 1 indicates that Thin Image (Clone) is used when you clone a volume if the `hitachi_default_copy_method` parameter is set to `FULL` and the `copy_method` metadata is not set.

Item		Value	Setting Pattern									
			1	2	3[1]	4[1]	5	6[1]	7	8	9[1]	10
Settings	Operation	Create volume from volume (clone) or create volume from snapshot	T	T	T	T	T	T	F	T	T	F
		Create snapshot	F	F	F	F	F	F	T	F	F	T
	Copy source	Volume or snapshot	T	T	T	T	T	T	T	F	F	F
		Group or group snapshot	F	F	F	F	F	F	F	T	T	T
	`hitachi_default_copy_method` parameter	`FULL`	T	T	T	F	F	F	F	T	F	F
		`THIN`[3]	F	F	F	T	T	T	F	F	T	F
		`FULL` or `THIN`[3]	F	F	F	F	F	F	T	F	F	T
	`copy_method`	Not set	T	F	F	T	F	F	F	T[2]	T[2]	F

Item		Value	Setting Pattern									
			1	2	3[1]	4[1]	5	6[1]	7	8	9[1]	10
	metadata	FULL	F	T	F	F	T	F	F	F	F	F
		THIN[3]	F	F	T	F	F	T	F	F	F	F
		FULL or THIN[3]	F	F	F	F	F	F	T	F	F	T
Selected copy method		Thin Image (Clone)	T	T	F	F	T	F	F	T	F	F
		Thin Image (Snapshot)	F	F	T	T	F	T	T	F	T	T

Legend:

- T: True
- F: False

Note:

1. Creating volumes by using Thin Image (Snapshot) is not recommended because you cannot delete source volumes or snapshots without deleting the destination volumes, and you cannot extend or unmanage source volumes, source snapshots, or destination volumes.
2. You cannot set `copy_method` metadata because all processing (creation of the destination volumes and data copy) takes place in the program before destination volumes become available.
3. When you create a global-active device pair volume by the operation create volume from volume or create volume from snapshot, you cannot specify `THIN` for the `hitachi_default_copy_method` parameter or the `copy_method` metadata.

Note:

- Do not set or change metadata other than `copy_method`.
- You can use Thin Image (Snapshot) in cascade configurations. You can connect up to 64 levels.
- When using Thin Image (Snapshot) in a cascade configuration, you cannot delete any intermediary volumes. To delete an intermediary volume, you first must delete all volumes under the target volume.
- For the following storage system models, if you use Thin Image (Clone) to copy data to a volume of a different size, after asynchronous copy is performed, the volume is expanded. For this reason, processing might take longer than the time required to copy to a volume of the same size.

  - VSP F1500

  - VSP G1000

  - VSP F400, F600, F800

  - VSP G200, G400, G600, G800

  The following describes the differences between the two types of processing:

  - Cloning volumes of the same size: (volume creation process) + (initialization of asynchronous copy process)

  - Cloning volumes of different sizes: (volume creation process) + (initialization of asynchronous copy process) + (asynchronous copy completion wait time) + (volume capacity extension process)

- By using Thin Image (Clone), you can repeat the data copy operation for the target volume while it is being copied. You can repeat the operation up to three times. The fourth time will cause an error.

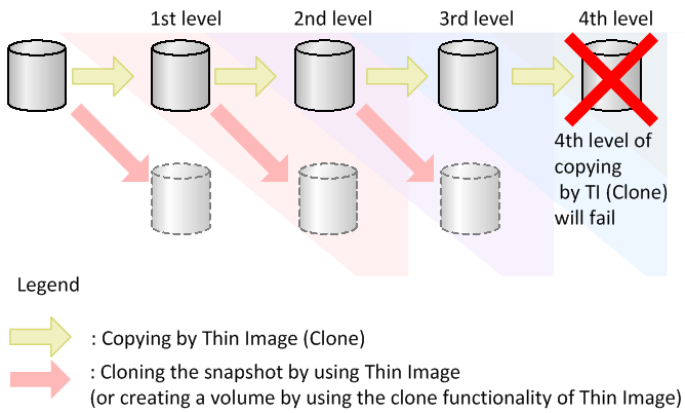

Maximum number of copy pairs and consistency groups

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, copy operations will not be executed successfully.

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

Port assignment using extra specs

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.

The following example shows how to use commands to specify ports in the extra specs. If you want to add a new port, specify a different volume type.

```
$ cinder type-create HUS-VM1
$ cinder type-key HUS-VM1 set volume_backend_name=HUS-VM1
$ cinder type-key HUS-VM1 set hbsd:target_ports="CL1-A, CL2-A"
$ cinder type-create HUS-VM2
$ Cinder type-key HUS-VM2 set volume_backend_name=HUS-VM2
$ cinder type-key HUS-VM2 set hbsd:target_ports="CL3-A, CL4-A"
$ cinder extra-specs-list
+--------------------------------------+--------+----------------------------------------------------------------------------------+
|                  ID                  |  Name  |                                   extra_specs                                    |
+--------------------------------------+--------+----------------------------------------------------------------------------------+
| b0d66cce-2766-43ef-868c-0f3b17f61111 | HUS-VM1| {u'hbsd:target_ports': u'CL1-A, CL2-A', u'volume_backend_name': u'HUS-VM1'} |
| b0d66cce-2766-43ef-868c-0f3b17f62222 | HUS-VM2| {u'hbsd:target_ports': u'CL3-A, CL4-A', u'volume_backend_name': u'HUS-VM2'} |
+--------------------------------------+--------+----------------------------------------------------------------------------------+
```

Note: Use a comma to separate multiple ports.

Port scheduler

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

In Block Storage Driver for OpenStack, 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.

The figure below shows an example of automatically created host groups.

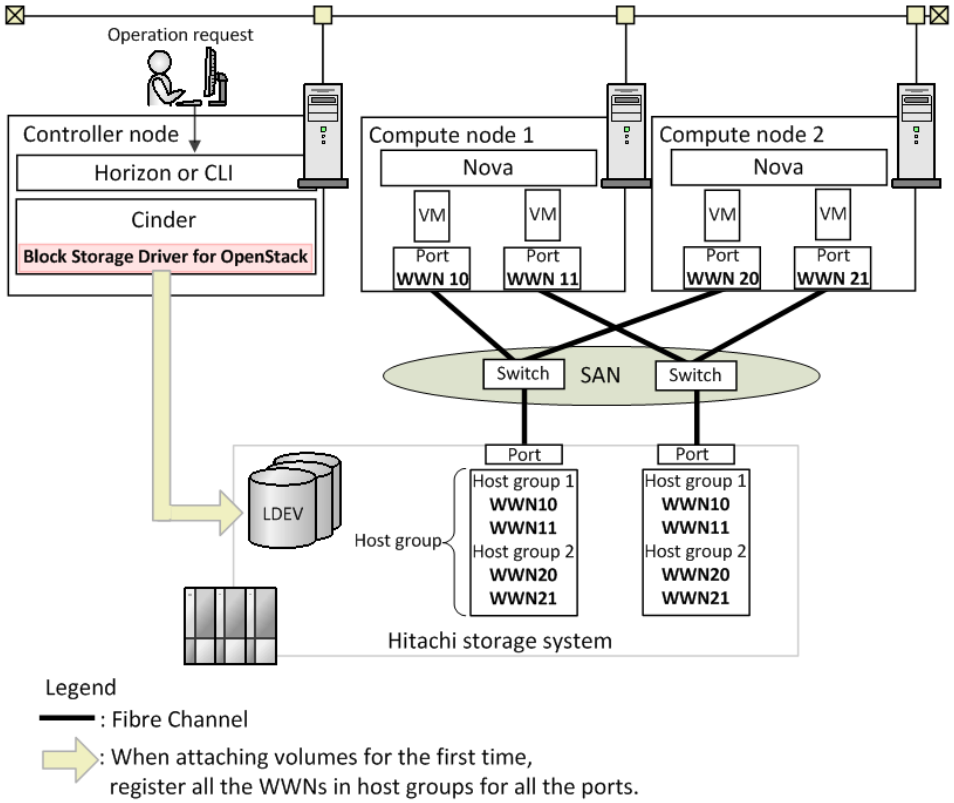

For the port scheduler function, when the cinder-volume service starts, the Fibre Channel Zone Manager obtains the WWNs of active compute nodes and of active VMs. When volumes are attached, the WWNs are registered in a round-robin procedure, in the same order as the order of ports specified for the parameter `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 Block Storage Driver for OpenStack specific parameters (on page 65) .

The figure below shows an example of host groups created from the port scheduler.

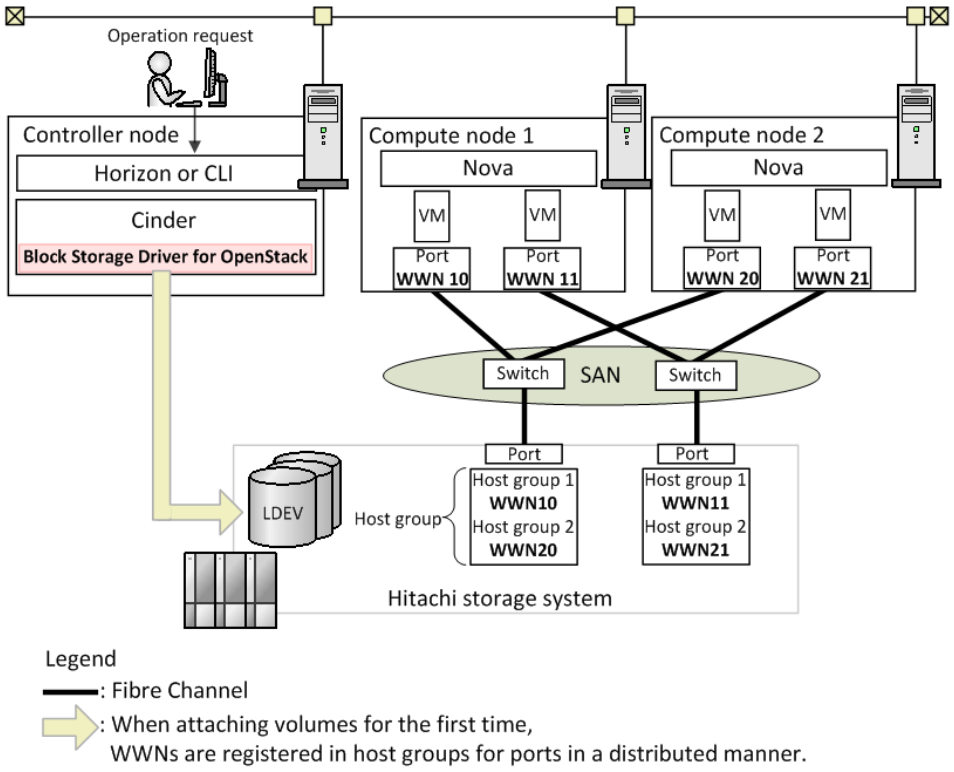

Note:

- Fibre Channel is supported. For details about ports, see Fibre Channel connectivity (on page 34) .
- 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.

Data deduplication and compression

Deduplication and compression is a method of improving storage utilization by data reduction. Once your storage administrator has enabled deduplication and compression for the DP pool, you can create or delete system data volumes for use with deduplication and compression.

For details, see *Capacity saving function: data deduplication and compression* in the *Provisioning Guide* for the applicable storage system, as listed in Referenced documents (on page 9) .

Enabling deduplication and compression

To use the deduplication and compression function, your storage administrator must first enable the deduplication and compression function 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.

Note: Configure the virtualization-excess ratio to be unlimited. If you want to set some other value for the virtualization-excess ratio, keep in mind that the capacity of the system data volume is 40 TB.

Creating a volume with deduplication and compression enabled

To create a volume with the deduplication and compression setting enabled, you need to 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.
   The example below shows how to enable the deduplication and compression function for volume type *VSP-G1500-1*.

```
# cinder type-create VSP-G1500-1
# cinder type-key VSP-G1500-1 set hbsd:capacity_saving=deduplication_compression
# cinder extra-specs-list
+--------------------------------------+----------------+-----------------------------------------------------------+
|                  ID                  |      Name      |                        extra_specs                        |
+--------------------------------------+----------------+-----------------------------------------------------------+
| fb994628-3983-4a1b-8b88-d1926a14a43d |   VSP-G1500-1  |   {'hbsd:capacity_saving': 'deduplication_compression'}  |
+--------------------------------------+----------------+-----------------------------------------------------------+
```

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.

   The example below shows how to create a volume (10 GB) with the deduplication and compression function enabled, where the volume type is *VSP-G1500-1*.

```
# cinder create --volume-type VSP-G1500-1 10
+--------------------------------+--------------------------------------+
|            Property            |                Value                 |
+--------------------------------+--------------------------------------+
|          attachments           |                  []                  |
|       availability_zone        |                 nova                 |
|            bootable            |                false                 |
|      consistencygroup_id       |                 None                 |
|           created_at           |      2016-10-25T02:24:05.000000      |
|          description           |                 None                 |
|           encrypted            |                False                 |
|               id               | edeb50d8-d340-4e3a-924b-0febd6f2b948 |
|            metadata            |                  {}                  |
|       migration_status         |                 None                 |
|          multiattach           |                False                 |
|              name              |                 None                 |
|     os-vol-host-attr:host      |                 None                 |
| os-vol-mig-status-attr:migstat |                 None                 |
| os-vol-mig-status-attr:name_id |                 None                 |
|  os-vol-tenant-attr:tenant_id  |   2afa8c77df6245b5939896a95b34f84c   |
|       replication_status       |                disabled              |
|              size              |                  10                  |
|          snapshot_id           |                 None                 |
|          source_volid          |                 None                 |
|             status             |                creating              |
|           updated_at           |                 None                 |
|            user_id             |   60361daadb8d4aee8034c541f5089323   |
|          volume_type           |               VSP-G1500-1            |
+--------------------------------+--------------------------------------+
```

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.

Chapter 5: Cinder operations in a global-active device (GAD) configuration

Even for volumes in a global-active device configuration, you can use Block Storage Driver for OpenStack to perform Cinder operations.

Overview and configuration of global-active device

If you use 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*.

In a global-active device configuration, if you want to use Block Storage Driver for OpenStack to perform operations on volumes, we recommend that you use a configuration in which the Configuration Manager REST API server is available at both the primary site and the secondary site. The following figure gives an example of such a configuration:

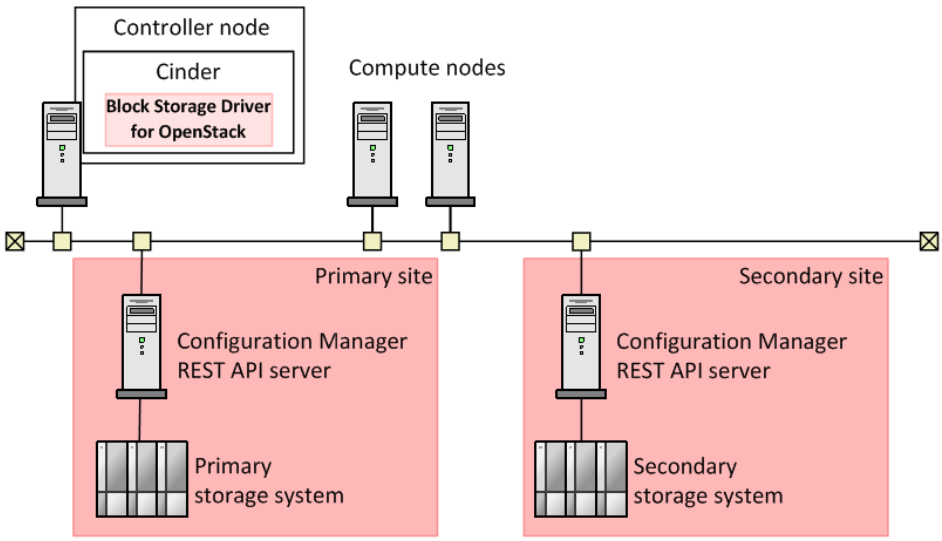

A global-active device pair is mapped as a Cinder volume, on which you can perform operations from a server in a single-server configuration or in a cross-path configuration. You cannot perform operations from a server in a server-cluster configuration.

Note:

- You cannot use Asymmetric Logical Unit Access (ALUA).
- 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.

Creating a global-active device environment

Before you can use global-active device, you need to use other storage system management tools to perform operations for creating the prerequisite environment, such as connecting remote paths, configuring a quorum disk, and creating a virtual storage machine (VSM). Block Storage Driver for OpenStack supports the following configurations:

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

  In this configuration, a VSM is configured on the secondary storage system by using the actual information of the primary storage system, and the P-VOL is not registered to a VSM on the primary storage system.

- Configuration where the P-VOL is registered to a VSM

  In this configuration, VSMs are configured on both the primary storage system and the secondary storage system, and the volumes registered to each VSM are assigned the same virtual LDEV number. The P-VOL is registered to the VSM on the primary storage system.

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

Block Storage Driver for OpenStack automatically executes procedures that are described in the *Global-Active Device User Guide*. You can skip configuring the following procedures that 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*

    If the `hitachi_group_request` parameter is `True`, you can skip this procedure.
  - *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.

Volume operations 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 as a Cinder volume. You can then perform operations on Cinder volume without thinking of it as a global-active device pair.

The following figure shows the behavior of the volumes in a global-active device pair in two storage systems and of a Cinder volume in a compute node when a Create Volume or Create Snapshot request is executed.

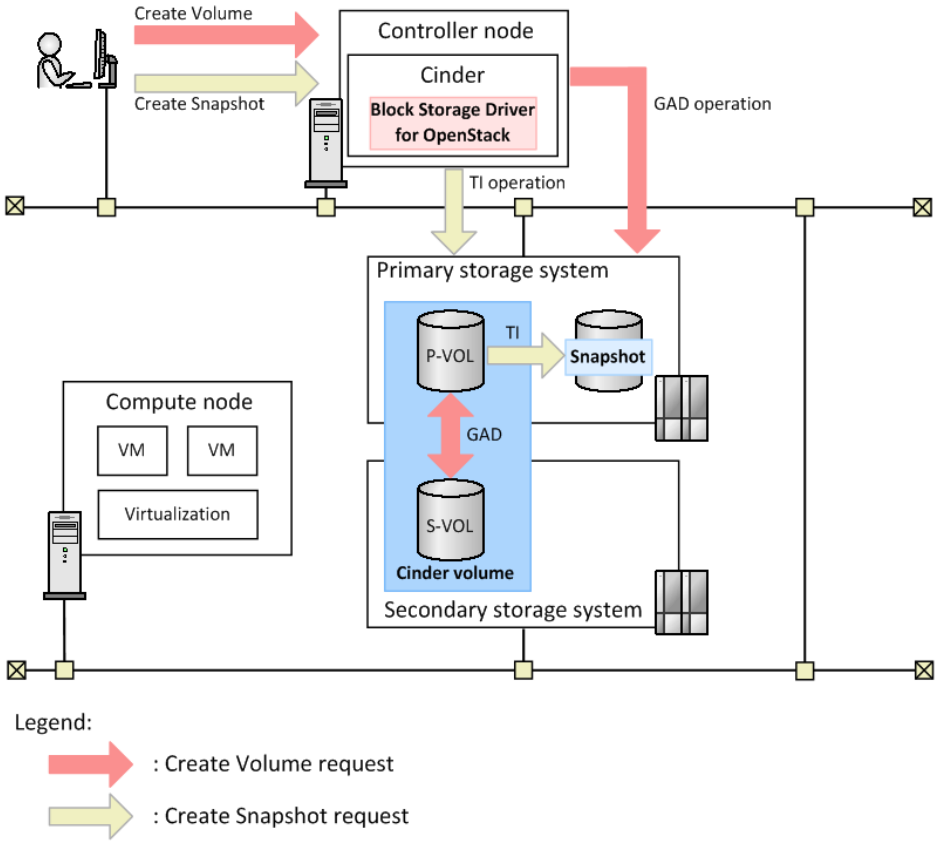

If you create a volume by using the Create Volume with the global-active device attribute specified, a global-active device pair is created as a Cinder volume. Using OpenStack, you cannot operate each volume of a global-active device pair as a separate Cinder volume.

If you use the Create Snapshot to obtain a snapshot of a Cinder volume, a snapshot of the P-VOL of the global-active device pair is created.

Note: Before performing operations on volumes, check the permissions of your user account.

Permission to access storage systems from Block Storage Driver for OpenStack, as well as the Security Administrator (View & Modify) and Storage Administrator (Remote Copy) permission.

Create 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:

```
# cinder type-create <volume type name>
# cinder type-key <volume type name> \
set hbsd:topology=active_active_mirror_volume
# cinder create --volume-type <volume-type name>
```

In this case, the following restrictions apply:

- You cannot create a volume for which the deduplication and compression function is enabled.
- When you create a global-active device pair volume by the operation create volume from volume or create volume from snapshot, you cannot specify `THIN` for the `hitachi_default_copy_method` parameter or the `copy_method` metadata.

Note: 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.

Unsupported Cinder functions

If a global-active device configuration is used, the following Cinder functions are not supported.

- Volume Migration (storage assisted)
- Consistency Group
- Generic volume group
- Revert to snapshot
- Manage Volume
- Unmanage Volume

In addition, if the configuration is "P-VOL registered to a VSM", of the Backup Volume functions, execution of the backup creation command with the `--snapshot` option or the `--force` option specified is not supported.

Chapter 6: Parameters

This section provides syntax for both Block Storage Driver for OpenStack-specific parameters and the Cinder-specific parameters. See the example configuration file in Edit cinder.conf (SUSE/Mirantis/Ubuntu) (on page 44).

Block Storage Driver for OpenStack specific parameters

The table below describes parameters specific to Block Storage Driver for OpenStack, which are specified in the backend section of the applicable configuration file.

- For Red Hat: The environment file for configuring Block Storage Driver for OpenStack (example: `custom_env.yaml`)

  For information about how to set the environment file, see the Red Hat document *Red Hat OpenStack Platform Custom Block Storage Back End Deployment Guide*.

- For SUSE/Mirantis/Ubuntu: `cinder.conf`

Parameter	Usage	Default value	Description
`hitachi_storage_id` (`hitachi_gad_storage_id`)*	Required	None	Specifies the storage device ID (12 digits) registered on the REST API server. Storage device ID: fixed value per device type (6 digits) + serial number (6 digits) If the serial number has fewer than 6 digits, add zeroes (0) to the front of the serial number.
`hitachi_storage_cli`*	Required	`REST`	If you specify `REST`, the Configuration Manager REST API is used for all storage-system operations. If you are using a VSP F350, F370, F700, F900 or VSP G350, G370, G700, G900 storage system, specify `SIMPLE_REST`. For global-active device configurations, you cannot enable `SIMPLE_REST`.

Parameter	Usage	Default value	Description
`hitachi_pool` (`hitachi_gad_pool`)*	Required	None	Specifies the ID (integer) or the name of the DP pool that stores LDEVs for volumes (or snapshots). Do not use the DP pool specified for `hitachi_pool` for other purposes.
`hitachi_thin_pool` (`hitachi_gad_thin_pool`)*	Optional	None	If you use the TI pool for creating snapshots, specify the ID (integer) or the name of the TI pool that stores LDEVs for volumes or snapshots. If you omit this parameter, the snapshots created by using Thin Image (Snapshot) are stored in the pool specified for `hitachi_pool`.
`hitachi_ldev_range` (`hitachi_gad_ldev_range`)*	Optional	None	Specifies a range of usable LDEV numbers in the *<integer-value-1>* – *<integer-value-2>* format. The value of *<integer-value-1>* must be equal to or smaller than the value of *<integer-value-2>*. If no value is specified, one of the following ranges applies: (1) If no resource group is used, the entire range permitted by the storage system can be used, or (2) The entire range permitted in a resource group for the user specified by the `hitachi_rest_user` can be used. You can specify integer values by using a decimal format or colon-separated hexadecimal format (*xx:yy:zz*).
`hitachi_rest_number` (`hitachi_gad_rest_number`)*	Optional	`0`	Specifies the number for identifying the Cinder driver instance. The value is used in names of storage resources (copy groups, device groups, and snapshots). Integer value (`0-255`).
`hitachi_rest_user` (`hitachi_gad_rest_user`)*	Required	None	Specifies the user name used for authentication to access the storage device. If the driver manages multiple storage devices, use the same user name, password, resource group and permission for all storage systems.

Parameter	Usage	Default value	Description

`hitachi_rest_password`
or
`hitachi_rest_password_path`
(`hitachi_gad_rest_password`
or
`hitachi_gad_rest_password_path`)*
	Required	None	Specifies the password used for authentication to access the storage device. Only one of the parameters should be specified, not both. If neither parameter is specified, startup of the backend will fail. For security reasons, it is strongly recommended that you use the `hitachi_rest_password_path` parameter. For more information about how to secure the password, see Secure your storage system password (on page 38).

`hitachi_rest_name_only_discovery`*	Optional	`False`	Sets a search restriction on host groups or iSCSI targets when attaching and detaching volumes. If `True` is specified, the driver searches only the host groups or the iSCSI targets that are named `HBSD-<target host WWPN>` or `HBSD-<IP Address of the connecting host (my_ip)>`, and ignores the rest. If host groups or iSCSI targets are created under different names, attaching and detaching volumes will result in an error. If `False` is specified, the driver searches the rest if it does not find specific host groups or iSCSI targets that match these names. In this case, a search might take more time.

Parameter	Usage	Default value	Description
`hitachi_target_ports` (`hitachi_gad_target_ports`)*	Optional	None	Specifies the controller port name to search host groups (iSCSI targets) when attaching volumes. Specify the port name in a format conforming to the storage system that uses it (for example, `VSP G1000: CL1-A`). To use multipath connection, specify controller port names, separated by a comma (for example, `CL1-A,CL2-A`). When attaching volumes to the Controller node, the controller port specified for this parameter is used. When attaching to the Compute node, the controller port specified for this parameter is used if `hitachi_compute_target_ports` is None (default). When this parameter is not specified and `hitachi_compute_target_ports` as well as `hitachi_rest_pair_target_ports` are specified, the driver operates as the environment without the I/O data path of the Controller node. If you have a backend storage system for which `SIMPLE_REST` is specified for the `hitachi_storage_cli` parameter and a backend storage system for which a different value is specified for the parameter, do not set the same port number for those backend storage systems.
`hitachi_compute_target_ports` (`hitachi_gad_compute_target_ports`)*	Optional	None	Specifies the controller port name to search host groups (iSCSI targets) when attaching volumes to a compute node. Specify the port name in a format conforming to the storage system that uses it (for example, `VSP G1000: CL1-A`). To use the multipath connection, specify the controller port names separated by a comma. If None (default) is specified, the controller port specified by `hitachi_target_ports` is used.

Parameter	Usage	Default value	Description
			If you have a backend storage system for which `SIMPLE_REST` is specified for the `hitachi_storage_cli` parameter and a backend storage system for which a different value is specified for the parameter, do not set the same port number for those backend storage systems.
`hitachi_rest_pair_target_ports` (`hitachi_gad_rest_pair_target_ports`)*	Optional	None	Specifies the names of one or more storage ports where host groups (iSCSI targets) are created for conventional copy functions. The name of the created host group (iSCSI target) is `HBSD-pair00`. Specify the port names in a format conforming to the corresponding storage system (for example, VSP G1000: CL1-A). To specify more than one port name, separate the names with commas. If this parameter is omitted, the ports specified by the `hitachi_target_ports` parameter are used.
`hitachi_group_request`*	Optional	`False`	Specifies `True` or `False` to determine whether a host group (iSCSI target) is created automatically if a host group (iSCSI target) corresponding to the connection target node does not exist for the port specified by the `hitachi_target_ports` or `hitachi_compute_target_ports` parameter. If `False` (default) is specified, the driver does not create the host group automatically.

Parameter	Usage	Default value	Description
`hitachi_group_delete`*	Optional	`False`	Defines specifications for automatically deleting empty host groups or iSCSI targets assigned to the target ports specified by the `hitachi_target_ports` or `hitachi_compute_target_ports` parameter. Set `False` to not delete automatically. Set `True` to delete automatically. If you specify `True`, you must also specify `True` for the `hitachi_group_request` parameter. If the Fibre Channel Zone Manager is used (the `hitachi_zoning_request` parameter must be set to `True`), Fibre Channel zones configured for the host groups or iSCSI targets will be deleted.
`hitachi_zoning_request`*	Optional	`False`	Specify this parameter if you are using a Fibre Channel connection. Specifies `True` or `False` to determine whether Fibre Channel zoning between the storage system and connection target node is configured automatically by collaborating with Fibre Channel Zone Manager. To configure automatically, you need to enable Fibre Channel Zone Manager. See the online manual for the configuration of Fibre Channel Zone Manager. If `False` (default) is specified, Fibre Channel zoning is not configured regardless of the availability of Fibre Channel Zone Manager.
`hitachi_use_chap_auth` (`hitachi_gad_use_chap_auth`)*	Optional	None	Specify this parameter if you are using an iSCSI connection. When using CHAP authentication of the iSCSI targets, set this parameter to `True`. For details, see *CHAP authentication* in the *Provisioning Guide* for the applicable storage system, as listed in Referenced documents (on page 9).

Parameter	Usage	Default value	Description
`hitachi_auth_user` (`hitachi_gad_auth_user`)*	Optional	None	Specify this parameter if you are using an iSCSI connection. Specifies the CHAP user name used for authentication of the iSCSI target. See the relevant user documentation. If `True` is specified for the `hitachi_use_chap_auth` parameter, this parameter is required.
`hitachi_auth_password` (`hitachi_gad_auth_password`)*	Optional	None	Specify this parameter if you are using an iSCSI connection. Specifies the password for the `hitachi_auth_user` setting. See the relevant user documentation. If `True` is specified for the `hitachi_use_chap_auth` parameter, this parameter is required.
`hitachi_auth_password_path` (`hitachi_gad_auth_password_path`)*	Optional	None	Specify this parameter if you are using an iSCSI connection. Specifies the file path that contains the password hint for `hitachi_auth_user`. Either the `hitachi_auth_password` or the `hitachi_auth_password_path` parameter should be specified if using CHAP authentication. Block Storage Driver for OpenStack returns an error if both parameters are specified. The file must be placed in a directory where the Cinder user or group can access but outside the directory in which the files are collected for troubleshooting purposes (for example, `/opt/hitachi/hbsd/auth_password`).
`hitachi_default_copy_method`*	Optional	`FULL`	Specifies the default `copy_method` for copying volumes. You can specify `FULL` for Thin Image (Clone) or `THIN` for Thin Image (Snapshot). For details about this parameter, see Copy methods (on page 53).

Parameter	Usage	Default value	Description
`hitachi_consisgroup_copy_method`*	Optional	The copy method specified in the `hitachi_default_copy_method` parameter.	Specifies the volume copy method to create a consistency group from a consistency group snapshot. You can specify `FULL` for Thin Image (Clone) or `THIN` for Thin Image (Snapshot). If nothing is specified, the copy method specified in the `hitachi_default_copy_method` is used.
`hitachi_copy_speed`*	Optional	`3`	Specifies the copy speed for copying volumes by using Thin Image (Clone) functions. You can specify a value in the range from `1` to `15`. If nothing is specified, `3` is set. Specify `1` or `2` to select slow copy speed. Specify `3` to select normal speed, and specify `4` or larger to specify high speed (prior).
`hitachi_copy_check_interval`*	Optional	`3`	Specifies the interval (seconds) at which pair creation is confirmed during a volume copy. You can specify a value in the range from `1` to `600`. If nothing is specified, `3` (seconds) is set.
`hitachi_async_copy_check_interval`*	Optional	`10`	Specifies the interval (seconds) at which copy pair deletion or data restoration is confirmed. You can specify a value in the range from `1` to `600`. If nothing is specified, `10` (seconds) is set.
`hitachi_rest_disable_io_wait`*	Optional	`False`	Specifies whether to enable the I/O check suppression option for a host group when the host group is created. If this parameter is set to `True`, before performing volume detach operations, make sure that no storage I/O operations are being performed. By setting this parameter to `True`, you can avoid the problem described in Volume detach operation fails (on page 87) .
`hitachi_rest_api_ip` (`hitachi_gad_rest_api_ip`)*	Required	None	Specifies the REST API server hostname or IP address.

Parameter	Usage	Default value	Description
`hitachi_rest_api_port` (`hitachi_gad_rest_api_port`)*	Optional	`23451`	Specifies the REST API server port number. For details about the REST API server port number, see the Hitachi Command Suite Configuration Manager REST API Reference Guide.
`hitachi_rest_tcp_keepalive`*	Optional	`True`	`True`: Use TCP keepalive functionality in the communication with the REST API server. With this setting, REST API server failures are detected in a short time. `False`: Do not use TCP keepalive functionality in the communication with the REST API server.
`hitachi_over_subscription`*	Optional	`False`	Enables the over-subscription function for VSP F350, F370, F700, F900 and VSP G350, G370, G700, G900 storage systems.
`hitachi_server_auto_create`*	Optional	`True`	Enables automatic connection between the attached volume and the server. If True is specified, the object is automatically created or updated every time a Compute node requests to connect to a volume. This parameter is enabled only if `SIMPLE_REST` is specified for the `hitachi_storage_cli` parameter. If `False` is specified, also specify `False` for the `hitachi_server_auto_delete` parameter.
`hitachi_server_auto_delete`*	Optional	`False`	Enables automatic disconnection of the server from the detached volume. If this parameter is set to `True`, the object is automatically deleted if no volumes are attached to the relevant host. This parameter is enabled only if `SIMPLE_REST` is specified for the `hitachi_storage_cli` parameter.
`hitachi_port_scheduler`	Optional	`False`	Specify this parameter to enable or disable the port scheduler function. If you specify `True`, when an attach request is received, the active WWNs that are obtained by Fibre Channel Zone Manager will be distributed and

Parameter	Usage	Default value	Description
			registered to the host groups of each port of the storage system. To use this function, specify `True` for the `hitachi_group_request` and `hitachi_rest_name_only_discovery` parameter. If you specify `False` for the `hitachi_rest_name_only_discovery` parameter, it will take a long time to attach volumes.
`hitachi_discard_zero_page`*	Optional	`True`	Enables discarding of zero-data pages after a volume is copied or a snapshot is created by using the **dd** command. If this parameter is set to `True`, zero-data pages are discarded.
`hitachi_host_mode_options`*	Optional	None	Specifies the host mode options for the host group or iSCSI target to be created. To set multiple host mode options, separate them by using commas. This parameter is enabled only if `hitachi_group_request` is `True` or `hitachi_server_auto_create` is `True`. For iSCSI targets, the host mode option `83` (Report iSCSI Full Portal List Mode) is automatically set regardless of this parameter. If `True` is specified for the `hitachi_rest_disable_io_wait` parameter, the host mode option `91` (Disable I/O wait for OpenStack Option) is automatically set regardless of this parameter.
`hitachi_gad_quorum_disk_id`	Optional (Required for global-active device configurations)	None	Specify the ID of the quorum disk used by global-active device. You can specify a value in the range from `0` to `31`. This parameter is enabled only if the `hitachi_gad_storage_id` parameter is specified (if a global-active device configuration is used).

Parameter	Usage	Default value	Description
`hitachi_gad_path_group_id`	Optional	`0`	Specify the ID of the path group used by global-active device. You can specify a value in the range from `0` to `255`. This parameter is enabled only if the `hitachi_gad_storage_id` parameter is specified (if a global-active device configuration is used).
`hitachi_gad_copy_speed`	Optional	`3`	Specify the copy speed used by global-active device. You can specify a value in the range from `1` to `15`. This parameter is enabled only if the `hitachi_gad_storage_id` parameter is specified (if a global-active device configuration is used).
* You can also use this parameter when performing operations on volumes in a global-active device configuration. In this case, the `hitachi_gad_storage_id` parameter needs to be specified. The parameters enclosed in parentheses are used to perform operations on a storage system at a secondary site.			

Cinder-specific parameters

Default section (SUSE/Mirantis/Ubuntu)

The table below describes details for the Cinder parameters set in the default section of the `cinder.conf` file. Only the parameters relevant to Block Storage Driver for OpenStack are described here.

Parameter	Usage	Default value	Description
`enabled_backends`	Optional	None	Specifies the backend definition section to be used. To specify multiple labels, separate them by a comma (for example, `VSP_G1000_1`, `VSP_G1000_2`). Do not specify more than 32 backends for the `enabled_backends` parameter.*
`my_ip`	Optional	`127.0.0.1`	Specifies the IP address of the local host (of the management path). You must use a unique address that is not used in other nodes.
`logging_context_format_string`	Optional	See the default value given in Block Storage Driver for OpenStack log output format.	Specifies the output format of logs.
`rpc_response_timeout`	Optional	`60`	Specifies the timeout period for attach or detach processing on a VM. If attach or detach processing times out, the processing is restarted. This parameter is specified in the DEFAULT section on the controller node that the cinder-api is running. To reflect the configuration, restart the openstack-cinder-api service.

Parameter	Usage	Default value	Description
`host`	Optional	None	Specifies a common host name for all backends. In the high availability configuration for controller nodes, specify the same value (for example, `ha_cluster`) for each controller node. This parameter is not required if the `backend_host` parameter is specified in the backend definition section.
* A maximum of 32 backends supported by Block Storage Driver for OpenStack can be used concurrently.			

Block Storage Driver for OpenStack log output format

For Block Storage Driver for OpenStack, add thread information after the process information in the default format for log analysis.

Default:

```
%(asctime)s.%(msecs)03d %(process)d %(levelname)s
%(name)s [%(request_id)s %(user_identity)s]
%(instance)s%(message)s
```

Block Storage Driver for OpenStack:

```
%(asctime)s.%(msecs)03d %(process)d %(thread)s
%(levelname)s %(name)s[%(request_id)s %(user_identity)s]
%(instance)s%(message)s
```

Backend definition section

The table below describes details for the Cinder parameters set in the backend section of the following configuration files.

- For Red Hat: The environment file for configuring Block Storage Driver for OpenStack (example: `custom_env.yaml`)

  For information about how to set the environment file, see the Red Hat document *Red Hat OpenStack Platform Custom Block Storage Back End Deployment Guide*.

- For SUSE/Mirantis/Ubuntu: `cinder.conf`

Name	Usage	Default value	Description

`volume_driver`	Required	None	Specifies the path of the driver to be used as follows:
			To use the FC driver:
			`cinder.volume.drivers.hitachi.hbsd.hbsd_fc.HBSDFCDriver`
			To use the iSCSI driver:
			`cinder.volume.drivers.hitachi.hbsd.hbsd_iscsi.HBSDISCSIDriver`
			If no path is specified, startup of the backend will fail.

`volume_backend_name`*	Optional	None	Specifies the name of the backend associated with the volume type.
`reserved_percentage`*	Optional	`0`	Specifies the reserved area size for the back-end storage, using a value in the range from `0` to `100` (%). To provide no reserved area, specify `0` (default).
`use_multipath_for_image_xfer`*	Optional	`False`	When the controller node and storage are connected via multipath, specify `True` or `False` to determine whether to set up a redundant configuration using a DM-Multipath in the controller node. To set up a redundant configuration, you must start the multipath daemon in advance. If you do not want to set up a redundant configuration, specify `False` (default).

Name	Usage	Default value	Description
`backend_host`	Optional	None	Specifies a host name for a target backend. In the high availability configuration for controller nodes, specify the same value (for example, `backend_ha_cluster`) for each backend of the controller node. If this parameter and the host parameter are specified, this parameter is used.
`image_volume_cache_enabled`	Optional	`False`	This parameter is for Image Volume Cache, which cannot be used with Block Storage Driver for OpenStack. Set the value of this parameter to `False`. If you want to quickly create a volume from an image, you can use the Volume-backed Image function.
`image_upload_use_cinder_backend`	Optional	`False`	This parameter is for the Volume-backed Image, which can be used if the parameter below is set to the value shown. `hitachi_default_copy_method=FULL` To use this function, set the value of this parameter to `True` and set the parameter listed above to the value shown. To not use this function, set the value of the parameter to `False`.
`suppress_requests_ssl_warnings`	Optional	`False`	Suppress SSL certificate warnings.
`driver_ssl_cert_path` (`hitachi_gad_ssl_cert_path`)*	Optional	None	Can be used to specify a non-default path to a `CA_BUNDLE` file or directory with

Name	Usage	Default value	Description
			certificates of trusted Certificate Authorities.
`driver_ssl_cert_verify` (`hitachi_gad_ssl_cert_verify`)*	Optional	`False`	If set to `True`, the SSL certificate of the REST API server will be enabled.
`max_over_subscription_ratio`*	Optional	`20.0`	Specify the maximum over-subscription ratio of the pool of the storage system. The specified value must be `1.0` or more. This parameter is enabled only if `True` is specified for the `hitachi_over_subscription` parameter and the storage system model is VSP F350, F370, F700, F900, or VSP G350, G370, G700, G900.

* You can also use this parameter when performing operations on volumes in a global-active device configuration. In this case, the `hitachi_gad_storage_id` parameter needs to be specified.

The parameters enclosed in parentheses are used to perform operations on a storage system at a secondary site.

Chapter 7: Troubleshooting

When errors are returned or if you encounter problems, you can troubleshoot them. For errors that require additional investigation, see Getting help (on page 11).

Note: For details about troubleshooting and disaster recovery for global-active device, see the *Global-Active Device User Guide*.

Check the version of Block Storage Driver for OpenStack

The first step in troubleshooting is to make sure you are running the desired version of Block Storage Driver for OpenStack that provides the Cinder functions you want to use.

If you encounter a problem, perform the following procedure to check the version of the driver and make sure the intended version is installed. If you are using an older version of the driver, update the software and perform the operation to confirm that the update has resolved the problem.

- For Red Hat:
  1. Run the `docker ps` command to obtain the container image ID for `rhosp14-openstack-cinder-volume-hbsd`.
  2. Run the following command:

     ```
     # docker inspect <container-image-ID>
     ```

     Check `Config->Labels->version` in the command execution result.
- For SUSE:

  Run the following command:

  ```
  # /bin/rpm -qi hbsd
  ```
- For Ubuntu and Mirantis:

  Run the following command:

  ```
  # /usr/bin/dpkg -s hbsd
  ```

hbsdgetinfo command

Function

The command **`hbsdgetinfo`** collects data needed to investigate driver errors and outputs the data to `hbsdgetinfo-`*`aaa.bbb.ccc.ddd-yyyymmddhhmmss`*`.tar.gz` (where *`aaa.bbb.ccc.ddd`* is the control node IP address and *`yyyymmddhhmmss`* is the date and time of collection).

Format

```
/opt/hitachi/hbsd/sbin/hbsdgetinfo [-d directory] [-e cinder]
```

Parameter	Description
-d *directory*	Specify the directory to save the data to be collected. The directory name can be a maximum of 512 bytes long. The directory can be specified as an absolute or relative path. If you omit this parameter, the `/tmp` directory will be used.
-e *cinder*	Specify this parameter to remove data belonging to the Cinder group (data to be collected by the data collection command cinder) from the collection target.

Exit code

- 0: Success
- 1: Fail

Collected data

Data in the common group is always collected. You can remove data in the Cinder group from the collection target by specifying the `-e` option.

Column	Collected data	Groups
Cinder (including Block Storage Driver for OpenStack) setting	All files below `/etc/cinder/`	Common
	`/etc/sudoers.d/cinder*`	Common
Cinder (including Block Storage Driver for OpenStack) log	All files below `/var/log/cinder/`	Common
Setting related to FC	`/etc/multipath.conf`	Common
	All files below `/etc/multipath.d/`	Common

Column	Collected data	Groups
System log	`/var/log/messages*`	Common
Setup of FC-related (Ubuntu only)	`/etc/multipath` directory	Common
Build date (Ubuntu only)	`/opt/hitachi/hbsd/build_date/*`	Common

Column	Command	Output file name	Groups
Working status of Cinder	**cinder** `absolute-limits`	`cinder_absolute-limits`	Cinder
	cinder `availability-zone-list`	`cinder_availability-zone-list`	Cinder
	cinder `backup-list`	`cinder_backup-list`	Cinder
	cinder `cgsnapshot-list`	`cinder_cgsnapshot-list`	Cinder
	cinder `consisgroup-list`	`cinder_consisgroup-list`	Cinder
	cinder `encryption-type-list`	`cinder_encryption-type-list`	Cinder
	cinder `extra-specs-list`	`cinder_extra-specs-list`	Cinder
	cinder `group-list`	`cinder_group-list`	Cinder
	cinder `group-snapshot-list`	`cinder_group-snapshot-list`	Cinder
	cinder `group-specs-list`	`cinder_group-specs-list`	Cinder
	cinder `group-type-list`	`cinder_group-type-list`	Cinder
	cinder `list`	`cinder_list`	Cinder

Column	Command	Output file name	Groups
	cinder qos-list	cinder_qos-list	Cinder
	cinder rate-limits	cinder_rate-limits	Cinder
	cinder service-list	cinder_service-list	Cinder
	cinder snapshot-list	cinder_snapshot-list	Cinder
	cinder transfer-list	cinder_transfer-list	Cinder
	cinder type-list	cinder_type-list	Cinder
	cinder list-extensions	cinder_list-extensions	Cinder
	pip show cinder	pip_show_cinder	Cinder
Main status of All component of openstack	**openstack-status**	openstack-status	Common
Working status of hfcldd	**cat** /proc/scsi/hfcldd/*	cat_proc_scsi_hfcldd	Common
	cat /sys/class/scsi_host/host*/hfcldd_proc	cat_class_scsi_host_hfcldd_proc	Common
Working status of iSCSI	**iscsiadm** -m discovery -P 1	iscsiadm_-m_discovery_-P_1	Common
	iscsiadm -m session -P 3	iscsiadm_-m_session_-P_3	Common
Working status related to FC	**ls** -l /dev/disk/by-path	ls -l_dev_disk_by-path	Common
	multipath -ll	multipath_-ll	Common

Column	Command	Output file name	Groups
	cat /sys/class/fc_host /host*/node_name	cat_sys_class_fc_host_node_name	Common
	cat /sys/class/fc_host host*/port_state	cat_sys_class_fc_host_port_state	Common
System information	**uname** -a	uname_-a	Common
	dmidecode	dmidecode	Common
	rpm -qa [1]	rpm_-qa	Common
	rpm -qi openstack-cinder [1]	rpm_-qi_openstack-cinder	Common
	rpm -qi hbsd [1]	rpm_-qi_hbsd	Common
	rpm -V python-cinder [1,4]	rpm -V python-cinder	Common
	rpm -V hbsd [1,4]	rpm -V hbsd [4]	Common
	dpkg-l [2]	dpkg_l	Common
	dpkg-s cinder-volume [2]	dpkg_s_cinder-volume	Common
	dpkg-s hbsd [2]	dpkg_s_hbsd	Common
	ifconfig -a	ifconfig_-a	Common
	ethtool_all [3]	ethtool_all	Common
System log	**journalctl**	journalctl	Common

Notes:

1. SUSE only.
2. Ubuntu, Mirantis only.
3. Script that collects **ethtool** information about all network interfaces in a system.
4. The output format is SM5DLUGT c *<file>*. See the **rpm** online documentation for more details about the output format.

Example

The following example shows how to use the command **hbsdgetinfo** to collect data.

```
# /opt/hitachi/hbsd/sbin/hbsdgetinfo
```

Output example

The command **hbsdgetinfo** outputs as follows:

```
# /opt/hitachi/hbsd/sbin/hbsdgetinfo
hbsdgetinfo command started.
creating directory ... succeeded.
copying "/etc/cinder" ... succeeded.
copying "/etc/sudoers.d/cinder" ... succeeded.
copying "/var/log/cinder" ... succeeded.
: (Omitted)
executing "/usr/bin/cinder transfer-list" ... succeeded.
executing "/usr/bin/cinder type-list" ... succeeded.
executing "/usr/bin/cinder list-extensions" ... succeeded.
archiving files ... succeeded.
output to "/tmp/hbsdgetinfo-aaa.bbb.ccc.ddd-yyyymmddhhmmss.tar.gz"
hbsdgetinfo command completed successfully.
```

Note:

- The data that can be collected varies depending on the environment. If data collection fails, the command outputs error messages.
- The command outputs the message "`hbsdgetinfo command completed successfully`" to confirm successful completion of the command.
- To collect the data for the Cinder group, Keystone certification for the administrator user must be kept in the environment variables. To change settings for the environment variables, use the **env** command. If that is not possible, specify the `-e cinder` option and delete the Cinder group data from the output.

Configuration Manager logs

If error messages starting with `KART-` are output to `/var/log/cinder/volume.log` and processing ends due to an error, Configuration Manager REST API log files might be required to investigate the problem.

For details about getting the logs, see *Collecting maintenance information about the REST API* in the Hitachi Command Suite Configuration Manager REST API Reference Guide.

Openstack-cinder-volume service does not start

If an error message for Block Storage Driver for OpenStack is output to `/var/log/cinder/volume.log`, see *Messages* to troubleshoot the problem.

If no error message is output for Block Storage Driver for OpenStack, a Cinder error message might be output to syslog.

Volume (snapshot) cannot be created

Check for an error message in `volume.log`, and see Messages (on page 89) to troubleshoot problem.

Volume (snapshot) cannot be attached

For error messages of Block Storage Driver for OpenStack in `volume.log`, see Messages (on page 89) to troubleshoot the problem.

Volume operations fail in a multi-controller node environment

Verify that your OpenStack configuration is valid. Refer to Supported configurations for Block Storage Driver for OpenStack (on page 27) .

Volume attach and detach operations fail

Volume attach and detach operations might fail when multiple volume operations are performed concurrently. This can be caused by time outs that occur during the processing of the Cinder requests in the backends.

You can increase the timeout threshold for all the nodes by modifying the `rpc_response_timeout` parameter to `3600` in the `nova.conf` (`/etc/nova/nova.conf`) and the `cinder.conf` file (`/etc/cinder/cinder.conf`).

Note: Changing the value of the `rpc_response_timeout` parameter to a larger value may increase the time needed to complete processing if an error occurs during attach or detach operations.

Volume detach operation fails

If you are using a software that detaches volumes in which no data is written, when you perform a volume detach operation, the target volume might have already been detached. In such cases, the volume detach operation fails because of a timeout error. To prevent this from occurring, set the I/O check suppression option (host mode option *91*) for the host group. If the `hitachi_group_request` parameter is set to `True`, and if new host groups are created automatically, specify this option by setting the `hitachi_rest_disable_io_wait` parameter to `True`.

If setting host groups or iSCSI targets for manual configuration, use the storage system management software to set `91` for the host mode option. For details about setting the host mode option, see the Hitachi Command Suite Configuration Manager REST API Reference Guide.

Volume attach performance issues after HBA replacement

Replacing a host bus adapter (HBA) in a multipath configuration can impact processing of volume attach requests in backend storage devices.

Make sure that the host group for the HBA has been updated after the replacement. If you have not already updated the host group name, with Storage Navigator, change the name to "HBSD-*xxx*" where *xxx* is the smallest WWPN number available for the server.

Volume deletion causes error status

When deleting many volumes in one operation, some of the volumes may result in an `error_deleting` status. If this happens, enter the **`cinder`** `reset-state` command. Verify that the volume status has returned to `available`, and then try to delete the volumes again.

Collecting Block Storage Driver for OpenStack reports and logs

If you need to contact Hitachi Vantara customer support, provide as much information about the problem as possible.

Collect the files below and provide them along with the actions you have taken to resolve the problem.

- Use the **`sosreport`** (Red Hat, Ubuntu, Mirantis) or **`supportconfig`** (SUSE) command to collect data on all nodes.
- Use the **`hbsdgetinfo`** command to collect data on all controller nodes.
- Copies of the operation logs of high availability cluster software that collect data on the controller node in the high availability cluster. For example, `/var/log/pacemaker`.

Chapter 8: Messages

Message output format

This topic describes the format of messages that are output to the `volume.log` file.

The following is Block Storage Driver for OpenStack message format:

```
storage_id yyyy-mm-dd hh:mm:ss.fff PID loglevel filepath [request_id user tenant] MSGIDn1n2n3n4-t: message
```

Where:

- *storage_id*: Storage system ID (the last six digits of the parameter `hitachi_storage_id`)

  The storage system ID is output only for a global-active device configuration.
- *yyyy-mm-dd hh:mm:ss.fff* : Date and time of output
- *PID*: Process ID
- *loglevel*: Log level
- *filepath*: File path of the output source
- *request_id*: Request ID
- *user*: User ID
- *tenant*: Tenant ID
- *MSGID*: Indicates that the message is a driver message
- *n1n2n3n4*: The serial number of the message
- *t*: The message type
- *message*: Message

The message type indicates the operation corresponding to the message. The message types are as follows:

Type	Description
E	Error message indicating that processing cannot continue due to a driver error or an incorrectly specified option.
W	Warning message indicating that a driver error occurred or that an option was incorrectly specified, but that processing will continue.

Type	Description
I	Message providing additional information about processing. Messages that are neither error messages nor warning messages.

List of messages

Message ID	Message	Driver status	Required actions
0000-I	The method <*method name*> was called. (config_group: <*backend definition name*>)	Continues processing the request.	None.
0001-I	The parameter of the storage backend. (config_group: <*backend definition name*>)	Continues processing the request.	None.
0002-I	The method <*method name*> completed successfully. (config_group: <*backend definition name*>)	Continues processing the request.	None.
0003-I	The storage backend can be used. (config_group: <*backend definition name*>)	Continues processing the request.	None.
0300-W	Failed to configure the internal logging. (ret: <*return value*>, stderr: <*standard error output*>).	Continues processing the request.	Verify that the user has sufficient permission for the directory.
0301-W	A LUN (HLUN) was not found. (LDEV: <*LDEV number*>)	Continues processing the request.	None.
0302-W	Failed to specify a logical device for the volume <*volume ID*> to be unmapped.	Continues processing the request.	None.
0303-W	An iSCSI CHAP user could not be deleted. (username: <*CHAP user name*>)	Continues processing the request.	If the CHAP user is no longer needed, delete it using the storage management software.

Message ID	Message	Driver status	Required actions
0304-W	Failed to specify a logical device to be deleted. (method: <*method name*>, id: <*volume ID* | *snapshot ID*>)	Continues processing the request.	Make sure that the logical devices do not exist for unnecessary volumes.
0305-W	The logical device for specified <*volume*|*snapshot*> <*volume ID* | *snapshot ID*> was already deleted.	Continues processing the request.	Make sure that no invalid volumes or snapshots exist.
0306-W	A host group could not be deleted. (port: <*port name*>, gid: <*group ID*>)	Continues processing the request.	Take the actions specified in the previous message.
0308-W	A host group could not be added. (port: <*port name*>)	Continues processing the request for other ports.	Take the actions specified in the previous message.
0309-W	An iSCSI target could not be added. (port: <*port name*>)	Continues processing the request for other ports.	Take the actions specified in the previous message.
0310-W	Failed to unmap a logical device. (LDEV: <*LDEV number*>)	Continues processing the request for other ports.	Make sure that the logical devices do not exist for unnecessary volumes.
0311-W	A free LUN (HLUN) was not found. Add a different host group. (LDEV: <*LDEV number*>)	Continues processing the request.	None.
0312-W	Failed to get a storage resource. The system will attempt to get the storage resource again. (resource: <*resource*>)	Continues processing the request.	None.
0313-W	Failed to delete a logical device. (LDEV: <*LDEV number*>)	Continues processing the request.	Make sure that the logical devices do not exist for unnecessary volumes or snapshots.

Message ID	Message	Driver status	Required actions
0314-W	Failed to map a logical device. (LDEV: <*LDEV*>, LUN: <*LUN number*>, port: <*port name*>, id: <*group ID*>)	Continues processing the request for the other host groups (iSCSI targets).	Take the actions specified in the previous message.
0315-W	Failed to perform a zero-page reclamation. (LDEV: <*LDEV number*>)	Continues processing the request.	None.
0317-W	Failed to assign the WWN. (port: <*port name*>, gid: <*group ID*>, wwn: <*WWN*>)	Continues processing the request for the other ports.	Take the actions specified in the previous message.
0318-W	Failed to copy meta data of destination volume <*volume ID*> to source volume <*volume*>. (reason: <*detailed message*>)	Continues processing the request.	None.
0319-W	The logical device does not exist in the storage system. (LDEV: <*LDEV*>)	Continues processing the request.	None.
0321-W	Failed to perform user authentication of the REST API server. (user: <*user*>)	Continues processing the request.	Verify user authentication information. Enter the correct values and retry the operation.
0324-W	Failed to delete full copy pair. (P-VOL: <*LDEV number*>, S-VOL: <*LDEV number*>)	Continues processing the request.	Take the actions specified in the previous message.
0325-W	Failed to delete thin copy pair. (P-VOL: <*LDEV number*>, S-VOL:<*LDEV number*>)	Continues processing the request.	Take the actions specified in the previous message.
0326-W	Failed to change the status of full copy pair. (P-VOL: <*LDEV number*>, S-VOL:<*LDEV number*>)	Continues processing the request.	Take the actions specified in the previous message.

Message ID	Message	Driver status	Required actions
0329-W	Failed to detach the logical device. (LDEV: <*LDEV number*>, reason: <*detailed message*>)	Continues processing the request for the other ports.	Take the actions specified in the previous message.
0330-W	The port name specified for the extra spec key "hbsd:target_ports" of the volume type is not specified for the hitachi_target_ports or hitachi_compute_target_ports parameter in cinder.conf. (port: <*port*>, volume type: <*volume type*>)	Continues processing the request.	Correct the port name for the extra spec key hbsd:target_ports.
0331-W	Volume copy to a larger volume started. This process waits until all data is copied to the destination volume, and therefore it may take a long time before the destination volume becomes available. Consider copying to a volume with the same size to save time because it performs data copy in the background and therefore ends much faster. (source <*volume* | *snapshot*>: <*volume-ID* | *snapshot-ID*>, destination volume: <*volume-ID*>, source size: <*source-size*>, destination size: <*destination-size*>)	Continues processing the request.	Include the time required for data copying to the estimated time until the copy destination volume can be used. For faster subsequent copying of the volume, set the source volume and the destination volume to the same size.
0332-W	Failed to change a volume type. The storage system does not support deduplication, or deduplication is not enabled for the pool. (pool: <*pool*>)	Continues processing the request using standard Cinder functions.	Check if the storage system supports deduplication. If it does, enable it, and then retry the operation.

Message ID	Message	Driver status	Required actions
0333-W	Retyping the volume will be performed using migration because the specified volume is being rehydrated. This process may take a long time depending on the data size. (volume: <*volume-ID*>, volume type: <*volume type*>)	Continues processing the request using standard Cinder functions.	To change volume type quickly, do not change volume types that enable the deduplication and compression function on volumes that just had the function disabled.
0334-W	Retyping the volume will be performed using migration because inconsistency was found in the deduplication system data volume. This process may take a long time depending on the data size. (volume: <*volume-ID*>, volume type: <*volume type*>)	Continues processing the request using standard Cinder functions.	To resolve the inconsistency in the system data volume, see your system administrator.
0335-W	Failed to create the host group because the host group maximum of the port is exceeded. (port: <*port*>)	Continues processing by using another port.	Check whether the port has any unnecessary host groups, and then delete those host groups.
0336-W	Failed to add the wwns to the host group port because the the WWN maximum of the port is exceeded. (port: <*port*>, WWN: <*WWN*>)	Continues processing by using another port.	Check whether the port has any unnecessary WWNs, and if it does, cancel the assignment of those WWNs.
0337-W	Failed to {attach | detach} the {P-VOL | S-VOL} in a GAD pair. (volume: <*volume-ID*>, reason: <*reason*>)	Continues processing the request.	Take the actions specified in the previous message.
0338-W	Failed to initialize the driver for the GAD {primary | secondary} storage system.	Continues processing the request.	Take the actions specified in the previous message.

Message ID	Message	Driver status	Required actions
0339-W	Port <*port*> will not be used because its settings are invalid. ({portType: <*port type*>|portAttributes: <*port attributes*>|(portType, portAttributes): (<*port type*>, <*port attributes*>)}[, lunSecuritySetting: <*LUN security setting*>[, {fabricMode: <*fabric mode*>, portConnection: <*port connection*>|ipv4Address: <*IPv4 address*>, tcpPort: <*TCP port*>}]])	Continues processing the request.	Change the settings of the port configuration specified in the message.
0340-W	Port <*port*> will not be used because it is not considered to be active by the Fibre Channel Zone Manager.	Continues processing the request.	Check the status of the port specified in the message.
0600-E	The Command <command> failed. (ret: <*return value*>, stdout:<*standard output*>, stderr:<*standard error output*>)	Stopped processing the request.	Remove the cause of the error by following the message output by the **`external`** command.
0601-E	A parameter is invalid. (<*parameter name* | *detailed message*>)	Stopped processing the request.	Check and, if necessary, revise the specified parameter values, and then try the operation again with the correct values specified.
0602-E	A parameter value is invalid. (<*meta data name*>)	Stopped processing the request.	Check and, if necessary, revise the specified parameter values, and then try the operation again with the correct values specified.

Message ID	Message	Driver status	Required actions
0603-E	Failed to acquire a resource lock. (serial: <*serial number*>, inst: <*HORCM instance number*>, ret: <*return value*>, stderr: <*standard error output*>)	Stopped processing the request.	Remove the cause of the error by following the message output by the **`external`** command.
0606-E	The snapshot <*snapshot ID*> cannot be deleted, because a read-only volume for the snapshot exists.	Stopped processing the request.	Make sure the snapshot operation target is correct. Alternatively, delete the target snapshot after deleting read-only volume related to the target snapshot.
0610-E	The status change of full copy pair could not be completed. (S-VOL: <*LDEV number*>)	Stopped processing the request.	Take the actions specified in the previous message.
0611-E	The status change of thin copy pair could not be completed. (S-VOL: <*LDEV number*>)	Stopped processing the request.	Take the actions specified in the previous message.
0612-E	The source logical device to be replicated does not exist in the storage system. (LDEV: <*LDEV number*>)	Stopped processing the request.	Specify the correct volume and retry the operation.
0613-E	The volume <*volume ID*> to be extended was not found.	Stopped processing the request.	Specify the correct volume and retry the operation.
0614-E	No WWN is assigned. (port: <*port name*>, gid: <*group ID*>)	Stopped processing the request.	Take the actions specified in the previous message.
0615-E	A pair could not be created. The maximum number of pair is exceeded. (copy method: <*copy-method name*>, P-VOL: <*LDEV number*>)	Stopped processing the request.	If `FULL` is specified, wait for copying to finish, and then try the operation again. If `THIN` is specified, delete the unnecessary pairs, and then try the operation again.

Message ID	Message	Driver status	Required actions
0616-E	A pair cannot be deleted. (P-VOL: <*LDEV number*>, S-VOL: <*LDEV number*>(<*pair status*>) [, <*LDEV number*>(<*pair status*>)...])	Stopped processing the request.	Wait until the pair can be deleted, and then try the operation again. For Thin Image (Snapshot), you cannot delete only the P-VOL. You must first delete the S-VOL. If the pair status is PSUE, resolve the issue on the storage system, and then retry the operation.
0617-E	Failed to create a volume from <*volume* | *snapshot*>. The size of the new volume must be equal to or greater than the size of the original <*volume* | *snapshot*>. (new volume: <*volume-ID*>)	Stopped processing the request.	Set the size of the destination volume equal to or greater than the size of the original volume or snapshot.
0618-E	The volume <*volume-ID*> could not be extended. The volume type must be Normal and the volume must not be an S-VOL in a cascade pair.	Stopped processing the request.	Verify the target volume is correct. If the type metadata is `Normal`, after the copy operation completes and the pair is deleted, retry the operation. If otherwise, since the volume cannot be extended, set the `copy_method` metadata to `FULL`, and then clone the volume. Extend and use this newly created volume.
0619-E	The volume <*volume ID*> to be mapped was not found.	Stopped processing the request.	Specify the correct volume, and then try the operation again. If the volume is no longer needed, delete the volume.

Message ID	Message	Driver status	Required actions
0620-E	Failed to provide information about a pool. (pool: <*pool*>)	Stopped processing the request.	Take the actions specified in the previous message.
0621-E	Failed to create a volume from <*volume* | *snapshot*>. The size of the new volume must be equal to the size of the original <*volume* | *snapshot*> when the new volume is created by Thin copy. (new volume: <*volume-ID*>)	Stopped processing the request.	Set the size of the destination volume equal to the size of the original volume or snapshot.
0624-E	HBSD has stopped processing the requests.	Stopped processing the request.	Specify the correct copy-source volume, and then try the operation again.
0631-E	Failed to create a file. (file: <*file name*>, error: <*error number*>, strerror: <*error message*>)	Stopped processing the request.	Check that the user has sufficient permission to perform the action on the directory.
0633-E	Permission denied.	Stopped processing the request.	Check that the user has sufficient permission to perform the action on the file.
0634-E	Failed to attach the logical device. (LDEV: <*LDEV number*>, reason: <*detailed message*>)	Stopped processing the request.	Take the actions specified in the previous message.
0636-E	Failed to add the logical device.	Stopped processing the request.	Take the actions specified in the previous message.
0638-E	Failed to add the pair target.	Stopped processing the request.	Take the actions specified in the previous message.
0639-E	Failed to map a logical device to any pair targets. (LDEV: <*LDEV number*>)	Stopped processing the request.	Take the actions specified in the previous message.

Message ID	Message	Driver status	Required actions
0640-E	A pool could not be found. (pool: <*pool*>)	Stopped processing the request.	Check and, if necessary, revise the parameter, and then restart the service.
0643-E	The iSCSI CHAP user <*CHAP user name*> does not exist.	Stopped processing the request.	Create the specified CHAP user. Alternatively, enable the parameter `hitachi_add_chap_user` to enable automatic creation of a CHAP user when creating an iSCSI target.
0648-E	There are no resources available for use. (resource: <*resource*>)	Stopped processing the request.	Check and, if necessary, revise the storage device settings and parameter settings to secure the target resources.
0649-E	The host group or iSCSI target was not found.	Stopped processing the request.	Create the host group or iSCSI target. Alternatively, if you are using iSCSI, enable the parameter `hitachi_group_request` to enable automatic creation of an iSCSI target.
0650-E	The resource <*resource*> was not found.	Stopped processing the request.	Check the setting for the target resources.
0651-E	The IP Address was not found.	Stopped processing requests.	Make sure an IP address has been set for the storage device iSCSI port.
0652-E	Failed to delete a logical device. (LDEV: <*LDEV number*>)	Stopped processing the request.	Take the actions specified in the previous message.

Message ID	Message	Driver status	Required actions
0653-E	The creation of a logical device could not be completed. (LDEV: <*LDEV number*>)	Stopped processing the request.	Take the actions specified in the previous message.
0654-E	A volume status is invalid. (status: <*status*>)	Stopped processing the request.	Make sure the operation target volume is correct. Alternatively, wait until the status of the target volume becomes Available, and then try the operation again.
0655-E	A snapshot status is invalid. (status: <*status*>)	Stopped processing the request.	Make sure the snapshot operation target is correct. Alternatively, wait until the status of the target snapshot becomes Available, and then try the operation again.
0656-E	The volume <*volume ID*> could not be restored. (reason: <*detailed message*>) **1.** Invalid input:<*volume ID | snapshot ID*> **2.** Status of target volume <*volume ID*> is not "available":<*status of volume*> **3.** Status of source volume <*volume ID | snapshot ID*> is not "available":<*status of volume*>	Stopped processing the request.	Make sure that the status and metadata for the specified volume are correct.
0657-E	A read-only volume cannot be created from the snapshot <snapshot ID>. A read-only volume already exists.	Stopped processing the request.	Make sure that the specified snapshot is correct.

Message ID	Message	Driver status	Required actions
0658-E	A new volume cannot be created from the snapshot <*snapshot ID*>. A read-only volume of the snapshot is attached.	Stopped processing the request.	Make sure that the specified snapshot is correct. Alternatively, detach the read-only volume from the instance before creating a new volume from the specified snapshot. After the new volume is created, reattach the read-only volume.
0659-E	A host group is invalid. (host group: <*host group name*>)	Stopped processing the request.	Create a host group.
0660-E	The specified <*volume* | *snapshot*> <*volume ID* | *snapshot ID*> is busy.	Stopped processing the request.	Make sure that the specified volume or snapshot is correct. Alternatively, wait for the other processing to finish, and then try the operation again.
0700-E	No valid value is specified for "storage_id". A valid value must be specified for "storage_id" to manage the volume.	Stopped processing the request.	Make sure that the storage identifier information or the command is correct, and then try the operation again.
0701-E	No valid value is specified for "ldev". A valid value must be specified for "ldev" to manage	Stopped processing the request.	Make sure that the LDEV information or the command is correct and retry the operation.

Message ID	Message	Driver status	Required actions
0702-E	Failed to manage the specified LDEV (<*LDEV number*>). The LDEV must be an unpaired DP-VOL <*volume type*>.	Stopped processing the request.	Make sure that the target LDEV information is correct. Alternatively, if the target LDEV has a pair, delete the pair and retry the operation. The driver does not manage LDEVs that are not DP-VOL. Create a DP-VOL, copy the data from the LDEV, and then retry the operation.
0703-E	Failed to manage the specified LDEV (<*LDEV number*>). The LDEV size must be expressed in gigabytes.	Stopped processing the request.	Make sure that the target LDEV information is correct. Extend the LDEV so that it is expressed in gigabytes, and then retry the operation.
0704-E	Failed to manage the specified LDEV (<*LDEV number*>). The LDEV must not be mapped.	Stopped processing the request.	Make sure that the target LDEV information is correct. Unmap the LDEV and try the operation.
0705-E	Failed to manage the specified LDEV (<*LDEV number*>). The LDEV must not be mapped.	Stopped processing the request.	Make sure the target LDEV information is correct. Clone the volume and retry the operation.

Message ID	Message	Driver status	Required actions
0706-E	Failed to unmanage volume <*volume ID*>. The volume type must be *Normal* and the volume must not be an S-VOL in a cascade pair.	Stopped processing the request.	Verify the target volume is correct. If the type metadata is `Normal`, retry the operation after copy completes and the pair is deleted. If otherwise, since the volume cannot be removed from Cinder management, set the `copy_method` metadata to `FULL`, and then clone the volume. Remove the newly created volume from Cinder management.
0707-E	No valid value is specified for "source-id". A valid LDEV number must be specified in "source-id" to manage the volume.	Stopped processing the request.	Make sure that the LDEV information or the command is correct, and then try the operation again.
0710-E	Failed to create a cloned volume for the volume <*volume ID*>. The volume type must be Normal.	Stopped processing the request.	Check the `type` metadata of the volume. If the `type` metadata of the volume is not `Normal`, replication of the volume might be in progress. Retry the operation after the replication finishes.
0711-E	A source volume for clone was not found. (volume_uuid: <*volume ID*>)	Stopped processing the request.	Specify the volume ID of an existing bootable volume in `volume_uuid`.

Message ID	Message	Driver status	Required actions
0712-E	Failed to create a consistency group snapshot. (copy method: <*copy-method*>) The number of pairs in the consistency group or the number of consistency group snapshots has reached the limit.	Stopped processing the request.	Verify that the number of pairs in the consistency group is not exceeding the max permitted (8,192). If the copy method is `FULL`, make sure there are no other consistency snapshot operations in progress, and then retry the operation. If the copy method is `THIN`, delete any unnecessary snapshots for the consistency group and retry the operation.
0713-E	V-VOLs cannot be added to a consistency group.	Stopped processing the request.	Add only DP-VOLs to the consistency group. Change the copy method so that DP-VOLs are added to the consistency group.
0714-E	Failed to delete unused full copy pair. (Copy Group: <*copy group* >)	Stopped processing the request.	Take the actions specified in the previous message.
0716-E	No logical device exists in the storage system for the volume <*volume-ID*> to be added to the consistency group <*consistencygroup-ID*>.	Stopped processing the request.	Wait for the LDEVs to be created, and then retry the operation.
0717-E	Multiple mappings for volume <*volume_id*> found; no host specified.	Stopped processing the request.	Verify the volume using the storage management software. Manually delete any unnecessary mappings.
0718-E	Failed to create a thin or full copy pair because specified volume (<*volume-ID*>) has cascade pair of clone or thin copy pair of snapshot.	Stopped processing the request.	Delete the clone or snapshot cascade pair for the specified LDEV, and then retry the operation.

Message ID	Message	Driver status	Required actions
0719-E	Failed to create a cascade pair because the micro of storage system does not support the copy function specified by the hitachi_copy_version parameter. (hitachi_storage_id: <*hitachi_storage_id*>,hitachi_copy_version: <*hitachi_copy_version*>, hitachi_thin_pool: <*hitachi_thin_pool*>, P-VOL: <*pvol*>, S-VOL:<*svol*>, Micro version:<*micro_version*>)	Stopped processing the request.	Update the micro version of the storage system software, and then retry the operation.
0722-E	Failed to unmanage the snapshot <*snapshot-ID*>. This driver does not support unmanaging snapshots.	Stopped processing the request.	Use the **`cinder`** `snapshot-reset-state` command as necessary to modify the status of the snapshot.
0723-E	Failed to create a volume. An invalid value is specified for extra spec key "<*key*>" of the volume type. (value: <*value*>)	Stopped processing the request.	Set the correct value for the extra spec key for the volume type, and retry the operation.
0724-E	Failed to create a deduplicated volume. The storage system does not support deduplication, or deduplication is not enabled for the pool. (pool: <*pool*>)	Stopped processing the request.	Verify the storage system supports deduplication. If supported, enable deduplication of the pool, and retry the operation.
0725-E	Failed to copy the volume. (copy method: <*copy-method*>, P-VOL: <*pvol*>, S-VOL: <*svol*>)	Stopped processing the request.	Take the actions specified in the previous message.

Message ID	Message	Driver status	Required actions
0726-E	A volume or snapshot cannot be deleted. The logical device consistency for the volume or snapshot cannot be guaranteed. (LDEV: <*LDEV*>)	Stopped processing the request.	After backing up any required data, refer to the storage user's guide, and then initialize data reduction.
0727-E	Failed to change the volume type. An invalid value is specified for extra spec key "<*key*>" of the volume type after change. (value: <*value*>)	Stopped processing the request.	Specify a valid value for extra spec key `<key>` for the volume type after the change, and then retry the operation.
0728-E	The volume type change could not be completed. (LDEV: <*LDEV*>)	Stopped processing the request.	Take the actions specified in the previous message.
0730-E	A newer version of the REST API is required. (current:<*current_version*>, required: <*required_version*> or newer)	Stopped processing the request.	Update the REST API server to the required version.
0731-E	Failed to communicate with the REST API server. (exception: <*exception*>, message: <*message*>, method: <*method*>, url: <url>, params: <*params*>, body: <*body*>)	Stopped processing the request.	Remove the cause of the error by following the message output by the HTTP library.
0732-E	The REST API failed. (source: <*errorSource*>, ID: <*messageId*>, message: <*message*>, cause: <*cause*>, solution: <*solution*>, code: <*errorCode*>, method: <*method*, url: <*url*>, params: <*params*>, body: <*body*>)	Stopped processing the request.	Remove the cause of the error by following the message output by the REST API server. If the "code" value in the message is {u'SSB1': u'2E13', u'SSB2': u'6003'}, run the **`cinder`** `migrate` command with the `--force-host-copy True` option.

Message ID	Message	Driver status	Required actions
0733-E	The REST API timed out. (job ID: <*job_id*>, job status: <*status*>>, job state: <*state*>, method: <*method*>, url: <*url*>, params: <*params*>, body: <*body*>)	Stopped processing the request.	Verify the REST API operation log and storage log, and then respond to the cause of the timeout. Verify volume status by using the storage management software. If any unnecessary volumes exist, remove them manually.
0734-E	The REST API failed. (HTTP status code: <*status_code*>, response body: <*response_body*>, method: <*method*>, url: <*url*>, params: <*params*>, body: <*body*>)	Stopped processing the request.	Verify the REST API server operation log.
0735-E	Failed to create a consistency group snapshot. V-VOLs exist in the consistency group. (group: <*group-ID*>, group snapshot: <*group-snapshot-ID*>, group type: <*group-type-ID*>, volumes: <*volume-ID*>[, <*volume-ID*>...], snapshots: <*snapshot-ID*>[, <*snapshot-ID*>...])	Stopped processing the request.	Snapshots for volumes in the consistency group cannot be created because V-VOLs exist in the consistency group. Delete the volumes indicated in the message or delete the consistent_group_snapshot_enabled attribute of the group type, and then retry the operation.
0736-E	Failed to delete a {volume|snapshot} in a {group|group snapshot}. ({group|group snapshot}: <*{group|group-snapshot}-ID*>, {volume|snapshot}: <*{volume|snapshot}-ID*>, LDEV: <*LDEV*>, reason: <*reason*>)	Stopped processing the request.	Take the actions specified in the previous message.

Message ID	Message	Driver status	Required actions
0737-E	Failed to create a volume snapshot in a group snapshot that does not guarantee consistency. (group: <*group-ID*>, group snapshot: <*group-snapshot-ID*>, group type: <*group-type-ID*>, volume: <*volume-ID*>, snapshot: <*snapshot-ID*>)	Stopped processing the request.	Refer to previous messages and eliminate the cause of the failure. Then create a snapshot of the volume specified in this message.
0738-E	Failed to migrate a volume. The volume status must be 'available'. (volume: <*volume-ID*>, status: <*status*>)	Stopped processing the request.	Run the **`cinder`** `migrate` command with the `--force-host-copy True` option.
0739-E	The Simple REST API failed. (source: <*error-source*>, ID: <*message-ID*>, message: <*message*>, cause: <*cause*>, solution: <*solution*>, code: <*error-code*>)	Stopped processing the request.	Take the actions specified in the previous message. If `SIMPLE_REST` is specified for the `hitachi_storage_cli` parameter, make sure that the model of the storage system whose ID is specified by `hitachi_storage_id` is VSP F350, F370, F700, F900 or VSP G350, G370, G700, G900. If the model is neither VSP F350, F370, F700, F900 nor VSP G350, G370, G700, G900, specify `REST` for the `hitachi_storage_cli` parameter or omit the parameter value.
0740-E	The Simple REST API timed out. (request ID: <*request-ID*>, command progress: <*progress*>, command result: <*result*>)	Stopped processing the request.	Check the REST API server operation log or the storage log to address the cause of the timeout.

Message ID	Message	Driver status	Required actions
0741-E	The Simple REST API failed. (HTTP status code: <*status-code*>, response body: <*response-body*>)	Stopped processing the request.	Check the REST API server operation log.
0742-E	Failed to attach the volume to the host because the server connection management object corresponding to the host was not found in the storage system. (server nickname: <*server-nickname*>, volume: <*volume-ID*>)	Stopped processing the request.	Create a server connection management object in the storage system. Alternatively, specify `True` for the `hitachi_server_auto_create` parameter.
0743-E	Failed to initialize volume connection because the extra spec key "hbsd:target_ports" is set for the volume type when hitachi_storage_cli parameter is set to "SIMPLE_REST". (volume: <*volume-ID*>, volume type: <*volume type*>, hbsd:target_ports: <*hbsd:target_ports*>)	Stopped processing the request.	Delete the extra spec key `hbsd:target_ports`, and then retry the operation.
0744-E	Failed to create a server connection management object in the storage system because the number of used objects has reached the limit. (server nickname: <*server-nickname*>, protocol: <*protocol*>, OS type: <*OS-type*>, OS type option: <*OS-type-option*>, iSCSI authentication method: <*iSCSI-authentication-method*>, iSCSI mutual setting: <*iSCSI-mutual-setting*>)	Stopped processing the request.	Delete unnecessary server connection management objects in the storage system.

Message ID	Message	Driver status	Required actions
0745-E	Failed to associate a server connection management object with a port in the storage system because the number of objects that can be associated with a port has reached the limit. (server nickname: <*server-nickname*>, connector ID: <*connector ID*>, ports: <*ports*>)	Stopped processing the request.	Delete unnecessary server connection management objects that are associated with the port.
0746-E	Failed to attach the volume to the host because the path setting for the server connection management object corresponding to the host is invalid. (server nickname: <*server-nickname*>, volume: <*volume-ID*>, ports: <*port*>*[, <port>…]*)	Stopped processing the request.	From the path setting for the server connection management object, delete any ports that are not specified for the `hitachi_target_ports` parameter or for the `hitachi_compute_target_ports` parameter. Alternatively, specify `True` for the `hitachi_server_auto_create` parameter.
0747-E	Failed to initialize volume connection because no active WWN was found for the connector. (WWN: <*WWN*>[, <*WWN*> …], volume: <*volume-ID*>)	Stopped processing the request.	Check the fabric network configuration, and configure the network so that active WWNs can be found.
0748-E	Failed to initialize volume connection because no port with an active WWN was found. (port, WWN: <*port*>, <*WWN*>[, port, WWN: <*port*>, <*WWN*>…], volume: <*volume-ID*>)	Stopped processing the request.	Check the fabric network configuration, and configure the network so that the ports with an active WWN can be found.

Message ID	Message	Driver status	Required actions
0749-E	The Fibre Channel Zone Manager is not available. The Fibre Channel Zone Manager must be up and running when hitachi_port_scheduler parameter is set to True.	Stopped processing the request.	Revise the settings of Fibre Channel Zone Manager so that it can run.
0750-E	Failed to initialize volume connection because no available resource of host group or wwn was found. (ports: <*port*>)	Stopped processing the request.	Revise the host groups or WWNs assigned to the port to secure available resources.
0751-E	The driver is not initialized for the GAD {primary | secondary} storage system.	Stopped processing the request.	Perform the actions indicated by the message that was output when the service started.
0752-E	Failed to create the {P-VOL | S-VOL} for a GAD pair. (volume: <*volume-ID*>, volume type: <*volume type*>, size: <*size*>)	Stopped processing the request.	Follow the instructions in the previously output message to eliminate the cause of the failure to create the LDEV, and then retry the operation.
0753-E	Failed to create a volume in a GAD environment because deduplication is enabled for the volume type. (volume: <*volume-ID*>, volume type: <*volume type*>, size: <*size*>)	Stopped processing the request.	Delete the specification `hbsd:capacity_saving` for the extra spec key of the volume type, or change the value specified for the extra spec key to `disable`, and then retry the operation.
0754-E	Failed to create a GAD pair or to mirror data in a GAD pair. (P-VOL: <*pvol*>, S-VOL: <*svol*>, copy group: <*copy group*>, pair status: <*pair status*>)	Stopped processing the request.	Check the REST API server operation log.

Message ID	Message	Driver status	Required actions
0755-E	Failed to split a GAD pair. (P-VOL: <*pvol*>, S-VOL: <*svol*>, copy group: <*copy group*>, pair status: <*pair status*>)	Stopped processing the request.	Check the REST API server operation log.
0756-E	A timeout occurred before the status of the GAD pair changes. (P-VOL: <*pvol*>, S-VOL: <*svol*>, copy group: <*copy group*>, current status: <*current status*>, expected status: <*expected status*>, timeout: <*timeout*> seconds)	Stopped processing the request.	Check the REST API server operation log.
0757-E	Failed to create a volume from a <*volume* | *snapshot*>. In a GAD environment, volumes cannot be cloned or created from snapshots if the copy method is THIN. (new volume: <*volume-ID*>, original <*volume* | *snapshot*>: <*volume-ID* | *snapshot-ID*>, copy method: <*copy-method*>)	Stopped processing the request.	Specify `FULL` for the `copy_method` metadata, and then retry the operation.
0758-E	Failed to extend a volume. The LDEVs for the volume are in a GAD pair and the volume is attached. (volume: <*volume-ID*>, LDEV: <*ldev*>, source size: <*source-size*>, destination size: <*destination-size*>, P-VOL: <*pvol*>, S-VOL: <*svol*>, P-VOL[numOfPorts]: <*pvol[numOfPorts]*>, S-VOL[numOfPorts]: <*svol[numOfPorts]*>)	Stopped processing the request.	Detach the target volume, and then retry the operation.
0759-E	Failed to add a volume to a group. The LDEVs for the volume are in a GAD pair. (volume: <*volume-ID*>, LDEV: <*ldev*>, group: <*group-ID*>)	Stopped processing the request.	Add a volume whose `type` metadata is anything other than `GAD` to the group.

Messages output by the command hbsdgetinfo

Message	Driver status	Required actions
hbsdgetinfo command started.	Continues processing the request.	None.
hbsdgetinfo command completed successfully.	Continues processing the request.	None.
output to *pathDesignated*/hbsdgetinfo-yyyymmddhhmmss.tar.gz.	Continues processing the request.	None.
creating directory ... {succeeded | failed}	Continues processing the request.	None.
"<*list of collection files*> | <*list of collection commands*>" is not found	Has stopped processing the requests.	Reinstall the driver.
"<*collection relevant files*> | <*collection commands*>" is not found	Continues processing the request.	None.
copying "<*fileName*>" ... {succeeded | failed}.	Continues processing the request.	None.
executing "<*commandName*>" ... {succeeded | failed | skipped}.	Continues processing the request.	None.
archiving files ... {succeeded | failed}.	Continues processing the request.	None.
hbsdgetinfo is already running	Has stopped processing the requests.	None.
hbsdgetinfo [-d directory] [-e cinder]	Has stopped processing the requests.	Check and, if necessary, revise the specified arguments, and then try the operation again.
directory name is too long.	Has stopped processing the requests.	Check and, if necessary, revise the specified directory.
"<*pass which you designated in -d*>" is not a directory	Has stopped processing the requests.	Specify a directory.

Message	Driver status	Required actions
"*<group name which you designated in -e>*" is not a group name	Has stopped processing the requests.	Check and if necessary, revise the specified group name.

Messages output by the installer

Message	Driver installer status	Required actions
---- <*service name*> service is running	Installer has stopped.	Stop the Cinder service.

Appendix A: OpenStack overview

This appendix provides overviews of OpenStack and Cinder.

About OpenStack

OpenStack is open source software for setting up a cloud environment of the IaaS layer. By linking components, OpenStack automatically combines necessary resources from the hardware resource pool that contains servers, storages, and network devices to provide users or groups with a private-server environment.

The list of the main OpenStack components is shown in the following table.

Component	Feature	Description
OpenStack Compute (Nova)	Hypervisor control	Manages virtual machines in a multi-hypervisor environment.
OpenStack Image (Glance)	Image, Snapshot control	Manages VM images.
OpenStack Network (Neutron)	Virtual network control	Manages L2/L3 networks.
OpenStack Object Storage (Swift)	Object storage control	Provides highly-available distributed object storage.
OpenStack Block Storage (Cinder)	Block storage control	Provides permanent block storage to VM.
OpenStack Identity (Keystone)	Integrated authentication infrastructure	Manages authentication and operation authorization common to each component.
OpenStack Dashboard (Horizon)	Self-service WebUI	Provides the self-service Web portal.

Cinder overview

Cinder is an OpenStack service component that provides persistent block device storage to VM OpenStack Compute instances as a volume. Data in a Cinder volume persists (remains undeleted) when the volume is detached from a VM instance. Since the data remains on the block device corresponding to the volume, it can always be reattached by changing the connection to other VM instances. In addition, high I/O performance is expected because it has direct access to the block device corresponding to the volume from the VM.

Storage providers include a Cinder driver for their storage devices so they can be used to provide storage in an OpenStack environment.

Cinder commands

The OpenStack Cinder commands can operate backends through the OpenStack Cinder API. Before using the commands, you must configure the available Keystone user authentication for the environment variables.

The OpenStack configuration installer included in each distribution creates the file containing the authentication information about the controller node. Load this file into the environment variables using the **`source`** command. The following is an example of an OpenStack environment configuration with the **`packstack`** command for the Red Hat OpenStack Platform.

```
# source /root/keystonerc_admin
```

Authentication information for admin user is described in `keystonerc_admin` and placed under `/root`.

For details about **`cinder`** command options, see the help message output by entering the following:

```
# /usr/bin/cinder help
```

Cinder functions available via CLI vary according to the version of the API. Refer to online help, such as OpenStack Command-Line Interface Reference, for details about CLI support for Cinder functions.

OpenStack commands

Commands of all OpenStack projects are being integrated into **`openstack`** commands. For example, **`cinder`** `create` is being replaced by **`openstack`** `volume create`. Using **`cinder`** commands is recommended because not all volume use cases are available in the **`openstack`** commands, such as GVG functions.

Appendix B: Block Storage Driver for OpenStack License

Apache License 2.0 is applied to Block Storage Driver for OpenStack.

```
# Copyright (C) 2014, 2019, Hitachi, Ltd.
#
# Licensed under the Apache License, Version 2.0 (the "License"); you may
# not use this file except in compliance with the License. You may obtain
# a copy of the License at
#
# http://www.apache.org/licenses/LICENSE-2.0
#
# Unless required by applicable law or agreed to in writing, software
# distributed under the License is distributed on an "AS IS" BASIS, WITHOUT
# WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the
# License for the specific language governing permissions and limitations
# under the License.
```

Glossary

Consistency group (CG)

A consistency group allows a storage device to take a snapshot of multiple volumes simultaneously so that all of the snapshots have the same point-in-time data consistency.

DP pool

Area in which Dynamic Provisioning manages virtual volumes (DP-VOL).

DP-VOL

A virtual volume where physical storage space is only consumed when needed (thin provisioning).

Dynamic Provisioning

Dynamic Provisioning is Hitachi's implementation of thin provisioning. Instead of reserving a fixed amount of storage for a volume, capacity from an available pool (DP pool) is used at the time data is written to the storage.

Generic volume group (GVG)

A group of volumes. This concept is an extension of consistency group. A GVG either guarantees or does not guarantee consistency. A GVG that guarantees consistency is similar to a consistency group, but is operated through different commands. A GVG is also simply referred to as a group.

Global-active device (GAD)

This software allows you to make data (in units of volumes) redundant and retain the same data on two storage systems, thereby improving the availability of storage systems.

In-band

In the in-band management method, the Configuration Manager REST API server connects to the storage system through a SAN.

Out-of-band

In the out-of-band management method, the Configuration Manager REST API server connects to the storage system through a LAN. This is the default communication mode for the Configuration Manager REST API server.

Snapshot

Data unit in which data of a volume at any point in time is stored.

Thin Image (Clone)

A copy method where the data on the source volume (DP-VOL) is copied asynchronously to the destination volume (DP-VOL).

Thin Image (Snapshot)

A copy method where a snapshot of the data on the source volume (DP-VOL) is created on the copy destination volume (DP-VOL).

TI pool

Thin Image pool used by Thin Image.

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