

Hitachi Block Storage Driver for OpenStack Pike

User Guide

© 2016, 2018 Hitachi, Ltd. All rights reserved.

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

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

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

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

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

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

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

Contents

Preface.....	7
Intended audience.....	7
Software version.....	7
Document conventions.....	7
Release notes.....	9
Changes made in this revision.....	9
Related documents.....	9
Conventions for storage capacity values.....	10
Getting help.....	11
Comments.....	11
 Chapter 1: Overview.....	 12
System Configuration.....	12
Supported Cinder functions.....	13
Support for additional OpenStack operations.....	16
Backend storage management configuration.....	16
Volume operation.....	16
Copy operations	17
Volume attachment.....	21
Backup operations.....	21
Manage and unmanage a volume.....	22
Volume migration	22
 Chapter 2: Operating environment.....	 23
Storage devices.....	23
Software.....	23
Prerequisite packages.....	23
Storage management software.....	24

Storage software licenses.....	24
Software restrictions.....	25
Storage firmware versions.....	25
Chapter 3: Installation and configuration.....	26
Installation and setup workflow.....	26
Supported configurations for Hitachi Block Storage Driver for OpenStack	26
Management software restrictions	28
Environment setup.....	28
Pre-check storage setting.....	28
Connect nodes.....	29
Single node configuration.....	30
Controller node without I/O data paths	31
Configure node setting.....	32
Configure Configuration Manager REST API server setting	33
Configure storage resources.....	33
Fibre Channel connectivity	33
iSCSI connectivity	35
Configure Configuration Manager REST API server	38
Secure your storage system password.....	38
Install Hitachi Block Storage Driver for OpenStack.....	39
Installation prerequisites.....	39
Block Storage Driver for OpenStack installation for Red Hat and SUSE....	40
Block Storage Driver for OpenStack installation for Mirantis and Ubuntu....	41
Initial setup.....	42
Add an Block Storage Driver for OpenStack configuration.....	42
Using Director (Red Hat OpenStack Platform).....	42
Using Crowbar (SUSE OpenStack Cloud).....	42
Using FUEL (Mirantis OpenStack).....	43
Using MAAS (Ubuntu OpenStack).....	43
Linking a volume type and a backend.....	43
Edit cinder.conf.....	43
Edit policy.json.....	45

Restart Cinder.....	45
Verify Block Storage Driver for OpenStack installation.....	46
Update the Hitachi Block Storage Driver for OpenStack.....	48
Update the Hitachi Block Storage Driver for OpenStack software for Red Hat and SUSE.....	48
Update the Block Storage Driver for OpenStack software for Mirantis and Ubuntu.....	48
Updating the management parameters in the cinder.conf file.....	48
Uninstall Hitachi Block Storage Driver for OpenStack.....	49
Live migration.....	50
Configure multipathing.....	50
Configure High Availability.....	51
High Availability overview.....	51
Pacemaker configurations for Cinder volumes.....	52
Build High Availability environment of controller node.....	52
Recoverable failures.....	53
Handling HA post-failover	54
Chapter 4: Hitachi extended functions.....	55
Copy methods	55
Maximum number of copy pairs and consistency groups.....	57
Create a volume from snapshot or create volume from volume.....	57
Port assignment for each volume type that can be attached.....	59
Data deduplication and compression.....	59
Enabling deduplication and compression.....	60
Creating a volume with dedupe/comp enabled.....	60
Deleting a volume with dedupe/comp enabled.....	62
Chapter 5: Parameters.....	63
Hitachi Block Storage Driver for OpenStack specific parameters.....	63
Cinder-specific parameters.....	70
Chapter 6: Troubleshooting.....	76
Check the version of Block Storage Driver for OpenStack	76
hbsdgetinfo command.....	76

Configuration Manager logs.....	80
Openstack-cinder-volume service does not start.....	80
Volume (snapshot) cannot be created.....	81
Volume (snapshot) cannot be attached.....	81
Volume operations fail in a multi-controller node environment.....	81
Volume attach and detach operations fail.....	81
Volume detach operation fails.....	81
Volume attach performance issues after HBA replacement.....	83
Volume deletion causes error status.....	83
Contacting Hitachi Vantara Corporation customer support.....	84
Chapter 7: Messages.....	85
Message output format.....	85
List of messages.....	86
Messages output by the command <code>hbsdgetinfo</code>	103
Messages output by the installer.....	104
Appendix A: OpenStack overview.....	105
About OpenStack.....	105
Cinder overview.....	105
Cinder commands.....	106
OpenStack commands.....	106
Glossary.....	107

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 software tools listed below.

- Red Hat OpenStack Platform
- 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 Hitachi Block Storage Driver for OpenStack version 5.1.0 or later for Pike.




Document conventions


This document uses the following typographic conventions:

Convention	Description
Bold	<ul style="list-style-type: none">▪ 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
<i>Italic</i>	<ul style="list-style-type: none"> 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: <pre>pairdisplay -g group</pre> <p>(For exceptions to this convention for variables, see the entry for angle brackets.)</p>
Monospace	Indicates text that is displayed on screen or entered by the user. Example: <code>pairdisplay -g oradb</code>
< > angle brackets	<p>Indicates variables in the following scenarios:</p> <ul style="list-style-type: none"> Variables are not clearly separated from the surrounding text or from other variables. Example: <pre>Status-<report-name><file-version>.csv</pre> 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	<p>Indicates that you have a choice between two or more options or arguments. Examples:</p> <p>[a b] indicates that you can choose a, b, or nothing.</p> <p>{ a b } indicates that you must choose either a or b.</p>

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

The following changes have been made to this document:

- The appearance and wording of this document were modified.
- The structure of this document was improved.

Related documents

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

Hitachi Virtual Storage Platform G1000, G1500, and F1500:

- *Provisioning Guide for Open Systems*, MK-92RD8014
- *Hitachi Thin Image User Guide*, MK-92RD8011
- *System Administrator Guide*, MK-92RD8016
- *Hardware Guide*, MK-92RD8007

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

- *Provisioning Guide*, MK-94HM8014
- *Hitachi Thin Image User Guide*, MK-92RD8011
- *System Administrator Guide*, MK-94HM8016
- *Hitachi Virtual Storage Platform G200 Hardware Reference Guide*, MK-94HM8020
- *Hitachi Virtual Storage Platform G400, G600 Hardware Reference Guide*, MK-94HM8022
- *Hitachi Virtual Storage Platform G800 Hardware Reference Guide*, MK-94HM8026
- *Hitachi Virtual Storage Platform F400, F600 Hardware Reference Guide*, MK-94HM8045
- *Hitachi Virtual Storage Platform F800 Hardware Reference Guide*, MK-94HM8046

Configuration Manager REST API

- *Hitachi Command Suite Configuration Manager REST API Reference Guide, MK-92HC229*

OpenStack

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

Conventions for storage capacity values

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

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

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

Logical capacity unit	Value
1 block	512 bytes
1 cylinder	Mainframe: 870 KB Open-systems: <ul style="list-style-type: none"> ▪ OPEN-V: 960 KB ▪ Others: 720 KB
1 KB	$1,024 (2^{10})$ bytes

Logical capacity unit	Value
1 MB	1,024 KB or 1,024 ² bytes
1 GB	1,024 MB or 1,024 ³ bytes
1 TB	1,024 GB or 1,024 ⁴ bytes
1 PB	1,024 TB or 1,024 ⁵ bytes
1 EB	1,024 PB or 1,024 ⁶ bytes

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.

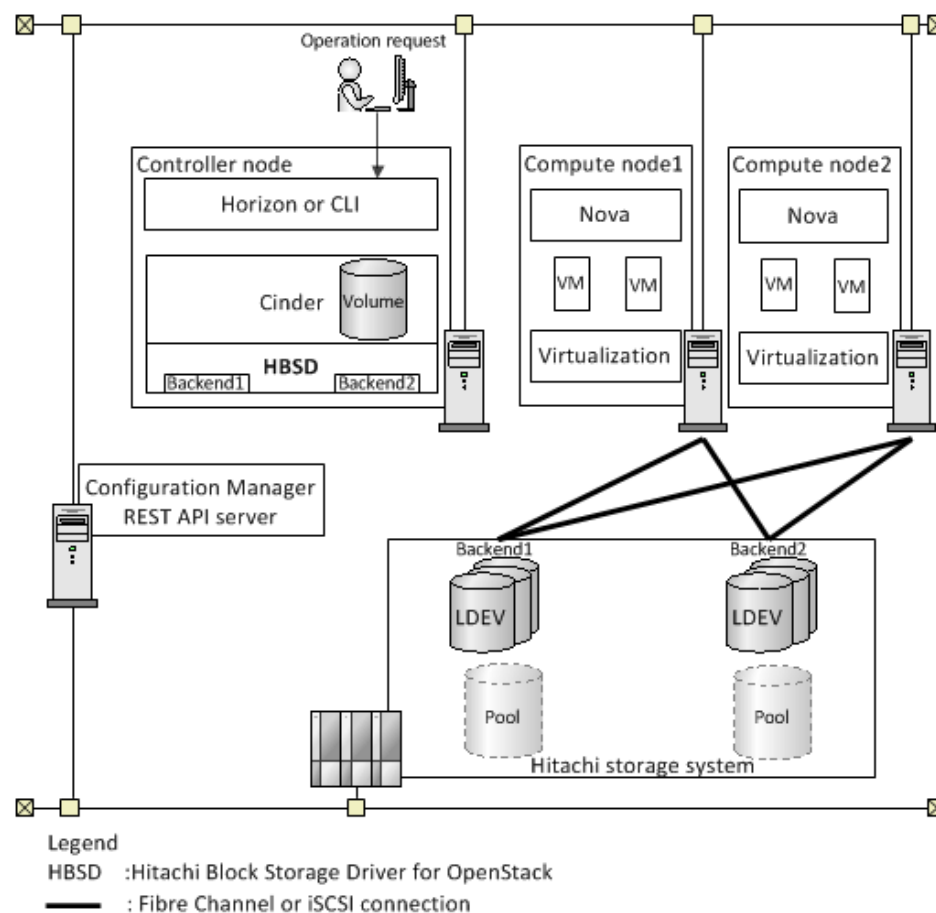
Thank you!

Chapter 1: Overview

Hitachi 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 Hitachi 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 node are configured in different nodes.



Controller node

The server which manages cloud systems. The Cinder-volume service is started in this node. Install and use Hitachi Block Storage Driver for OpenStack in this node.

Compute node

The server managed by the controller node.

Configuration Manager REST API server

A server that receives REST API requests from Hitachi 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. Hitachi Block Storage Driver for OpenStack uses thin-provisioned volumes (DP-VOLs) from Hitachi Dynamic Provisioning (Hitachi) pools as backend storage. To use multiple DP pools, register each DP pool as a separate backend storage.

LDEV

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

Volume

Data unit managed by Hitachi Block Storage Driver for OpenStack.

Supported Cinder functions

The following table lists the Cinder functions supported by Hitachi 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, 3}	Creates a snapshot that stores point-in-time data.
Delete Snapshot	Deletes a snapshot.
List Snapshot	Lists snapshots

Function	Description
Create Volume from Snapshot ^{3, 4}	Creates a volume with the same content of a snapshot.
Create Volume from Image	Retrieves image stored from Glance and loads it onto the specified volume. A VM can be started based on the stored image.
Create Volume from Volume (Clone) ^{3, 4}	Creates a volume and then stores the data from the source volume.
Create Consistency Group	Creates a consistency group.
Delete Consistency Group	Deletes a consistency group.
Create Consistency Group from Source Consistency	Creates a consistency group from a source consistency group.
Create Consistency Group from Snapshot	Creates a consistency group from a snapshot of another consistency group.
Create Snapshot for Consistency Group	Creates a snapshot for a consistency group.
Delete Snapshot of Consistency Group	Deletes a snapshot of a consistency group.
Update Consistency Group	Modifies a consistency group.
Create Image from Volume	Creates an image from a volume and uploads it (Glance).
Volume Migration ¹ (host assisted)	Moves a volume to another HDP pool managed by another storage system.
Volume Migration ¹ (storage assisted)	Moves a volume to another HDP pool.
Get Volume Status	<p>Returns storage information, such as the total virtual capacity and free virtual space. Virtual capacity is the maximum capacity that can be subscribed, and it is calculated using the maximum subscription rate against the actual Dynamic Provisioning pool capacity.</p> <p>For example, if the actual capacity is 100 GB and the maximum subscription rate is 150%, the virtual capacity is calculated as follows:</p> <p>100 GB x 1.50 = 150 GB</p>
Backup Volume	Backs up a volume to Swift. Starting Cinder backup service is required.

Function	Description
Restore Backup	Restores data to a specified volume from Swift. Starting Cinder backup service is required.
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.
Retype Volume ¹	Changes the volume type.
Create Group	Creates a generic volume group.
Delete Group	Deletes a generic volume group.
Create Group from Source Group	Creates a generic volume group from another generic volume group.
Create Group from Group Snapshot	Creates a generic volume group from a generic volume group snapshot.
Update Group	Updates a generic volume group.
Create Group Snapshot	Creates a generic volume group snapshot.
Delete Group Snapshot	Deletes a generic volume group snapshot.
Revert to snapshot	Reverts a volume to a previously taken snapshot.
Notes: <ol style="list-style-type: none"> 1. Not supported if a volume is in a group. 2. Snapshots are created using Thin Image (Snapshot) 3. Block Storage Driver for OpenStack can utilize block copy functions for faster performance (license required) or host-based <code>dd</code> copy command for slower performance. 4. When creating a volume with Thin Image, its size can be greater than its source volume. Cloning a volume to a different size takes longer than cloning one to the same size because the asynchronous copy must complete before the volume can be extended. The time needed to extend the volume capacity is in addition to the time needed to complete the asynchronous copy. <ul style="list-style-type: none"> ▪ 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) 	

Support for additional OpenStack operations

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

Table 1 OpenStack operations supported by Hitachi Block Storage Driver for OpenStack

OpenStack Operations	Description
Live Migration	Moves Block Storage Driver for OpenStack-managed volume attached to a running VM from one compute node to another.
Multipath for I/O path	Multipath enabled for Block Storage Driver for OpenStack-managed volumes.
Fibre Channel Zoning 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.
High Availability for Cinder volume service	Cinder service can be set to Active-Standby redundant using High Availability cluster software. High Availability is not supported for SUSE OpenStack Cloud or Mirantis OpenStack.

Backend storage management configuration

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

Volume operation

Hitachi 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.

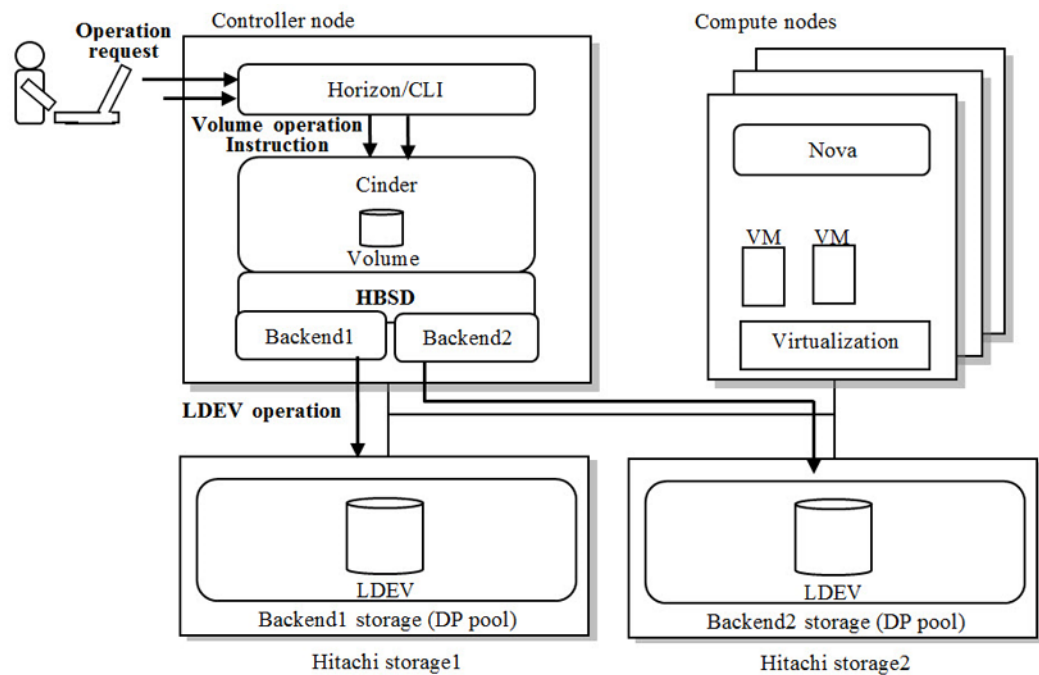


Figure 1 Overview of volume operations

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 Thin Image (Clone) during asynchronous copying or DP-VOLs which have pair volumes that were created by Thin Image (Snapshot).

Copy operations

When the Block Storage Driver for OpenStack receives a Cinder copy volume request, the Block Storage Driver for OpenStack copies the volume in full using Thin Image (Clone), and takes a point-in-time snapshot of the volume using Thin Image (Snapshot). In addition, snapshots for multiple volumes can be created with 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.

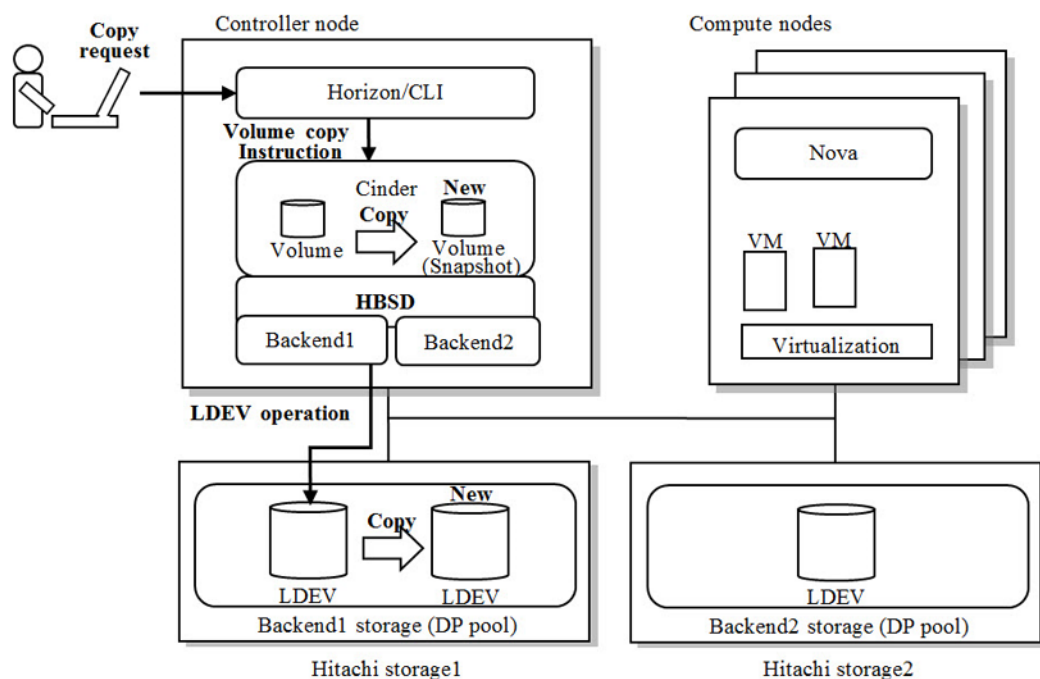


Figure 2 Snapshot operation example

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.
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 in a cascade configuration.

Create a snapshot

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

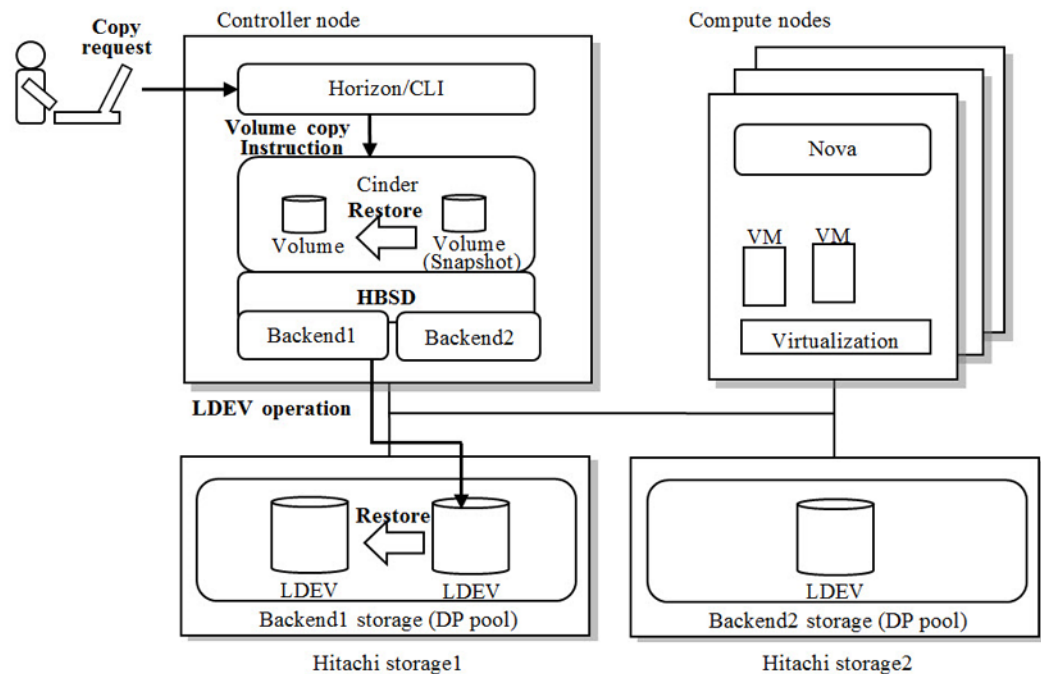


Figure 3 Snapshot and volume restore operation

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 with **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. The 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 group.

Create a group from a snapshot of another group

You can create a group from a snapshot of another group. The 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 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 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. Otherwise, snapshots are created separately and the points in time at which the data for each snapshot is created may differ.



Note:

- There are two methods for creating a consistency group. A group created by **cinder** `consisgroup-create` command is always a consistency group. A group created by **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. CG commands will be obsolete and not be supported in the future. In addition, CG commands and GVG commands are not compatible with each other. Snapshots created with CG commands cannot be operated with 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 does not support GVG-related operations yet.

Volume attachment

Block Storage Driver for OpenStack supports storage attachment capability for Fibre Channel, iSCSI, and Fibre Channel over Ethernet storage. Block Storage Driver for OpenStack supports multipath configurations. The following figure shows an overview of the volume attachment operations.

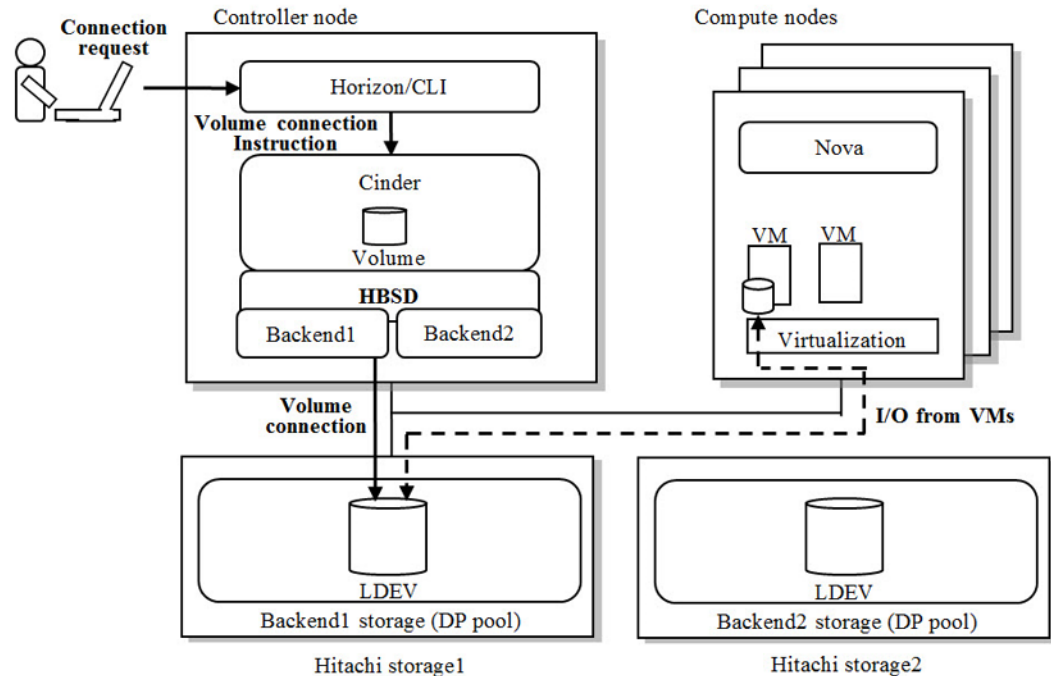


Figure 4 Volume attachment

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 and the compute. 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 backup or restore volumes created by HBSD; however, if you restore a volume from external storage using standard OpenStack functions, the metadata for the volume will be overwritten. To prevent this, use the procedure listed below to restore a volume from external storage.

Procedure

1. Create the volume for the restore target beforehand. Note the LDEV metadata value that can be acquired 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 using the **cinder metadata** command. Use the same command to restore the type and LDEV metadata values. For details of these commands, see help messages or the Command-Line Interface Reference.

Manage and unmanage a volume

If you plan to move a volume from one OpenStack Cinder 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: Operating environment

This chapter describes the hardware and software necessary for using Hitachi Block Storage Driver for OpenStack.

Storage devices

The following storage systems are supported as backend storage devices.

Storage devices
Hitachi Virtual Storage Platform G1000, G1500
Hitachi Virtual Storage Platform F1500
Hitachi Virtual Storage Platform G200, G400, G600, G800
Hitachi Virtual Storage Platform F400, F600, F800

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

Software

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

Prerequisite packages

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

Distribution	Node	Item	Contents
Red Hat	Controller Node	Linux Environment	sysfsutils ¹ , sg3_utils ¹ , iscsi-initiator-utils ² , device-mapper-multipath ³

Distribution	Node	Item	Contents
	Compute Node	Linux Environment	sysfsutils ¹ , sg3_utils ¹ , iscsi-initiator-utils ² , device-mapper-multipath ³
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: <ol style="list-style-type: none"> 1. For configuring FC. 2. For configuring iSCSI. 3. For configuring multipath. 			

For other Hitachi driver specifications, see: <https://www.HitachiVantara.com/en-us/products-solutions/application-solutions/hitachi-and-openstack.html>.

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

Storage management software

Configuration Manager REST API version 8.5.4-00 or later is required for managing Hitachi storage devices.

Storage software licenses

The software licenses listed below are required on each storage device.

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

For VSP F400, F600, and F800, the Hitachi Local Replication software is included in the Hitachi Storage Virtualization Operating System.

- Dedupe and compression (optional)

Software restrictions

Do not use Hitachi Dynamic Link Manager (HDLM) in controller nodes or compute nodes. Use DM-Multipath instead. Simultaneous use of the Hitachi Block Storage Driver and the Hitachi Device Manager (HDvM) is not supported.

Storage firmware versions

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

Table 2 Supported storage firmware versions

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

Chapter 3: Installation and configuration

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

Installation and setup workflow

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

Prepare for installation

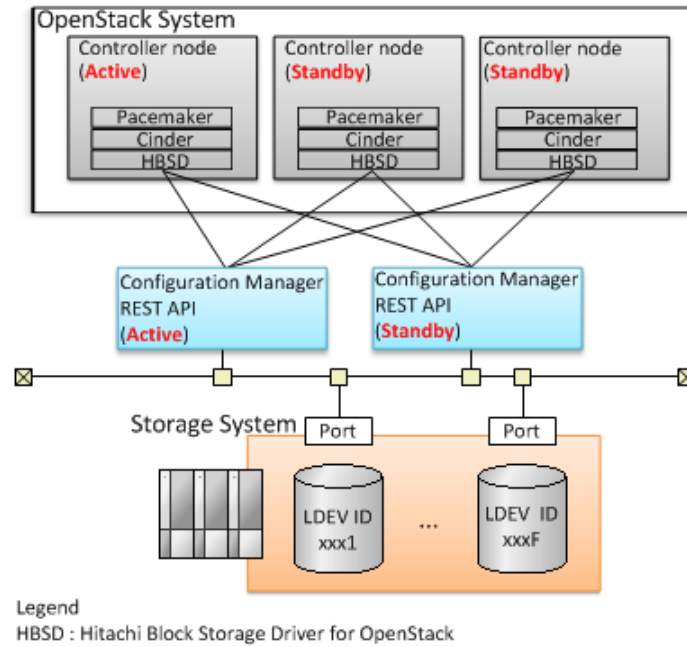
1. Verify that basic volume operations can be performed on the storage. See [Pre-check storage setting \(on page 28\)](#).
2. Connect each nodes and the storage. See [Connect nodes \(on page 29\)](#).
3. Configure node setting. See [Configure node setting \(on page 32\)](#).
4. Configure setting for storage resources used by Block Storage Driver for OpenStack. See [Configure storage resources \(on page 33\)](#).
5. Install and configure the REST API server.

Install Hitachi Block Storage Driver for OpenStack

1. Check the required memory size and disk space, and verify the OpenStack package installation. See [Installation prerequisites \(on page 39\)](#).
2. Install Block Storage Driver for OpenStack. See [Install Hitachi Block Storage Driver for OpenStack \(on page 39\)](#).
3. Configure Block Storage Driver for OpenStack. See [Initial setup \(on page 42\)](#).
4. Restart Block Storage Driver for OpenStack. See [Restart Cinder \(on page 45\)](#).
5. Verify the installation. See [Verify Block Storage Driver for OpenStack installation \(on page 46\)](#).

Supported configurations for Hitachi 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 (HA) configuration.



The following table describes valid OpenStack system configurations.

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 by amongst all nodes.
Note: If multiple OpenStack components are used in the controller node, services other than the Cinder-volume service and components other than Cinder can be used in an active-active configuration, according to the recommended system configuration.	

The following table describes the valid system configurations of the Configuration Manager REST API when used with Hitachi 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 <i>System Requirements</i> for Hitachi Command Suite.

Configuration	Description
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

Hitachi 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 conflict. To avoid such errors, you can use the following parameters to specify resources that can be used by Hitachi Block Storage Driver for OpenStack:

- `hitachi_ldev_range` to specify a range of LDEV IDs.
- `hitachi_target_ports` and `hitachi_compute_target_ports` to specify ports.

Do not use other management software to perform operations on the resources specified for these parameters.

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.

Pre-check storage setting

Make sure that the storage device is configured and ready for volume operations.

Use the Storage Navigator software to confirm that LDEVs can be created, connected to the controller node or compute nodes, asynchronously copied, and that host groups can be created. After confirming, delete any host groups that you created.

For details, see the documentation listed below.

- *Provisioning Guide for Open Systems: Hitachi Virtual Storage Platform G1000, G1500, and F1500*, MK-92RD8014
- *Provisioning Guide: Hitachi Virtual Storage Platform Gx00 and Fx00 Models*, MK-94HM8014

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 an Fibre Channel (FC) connection example:

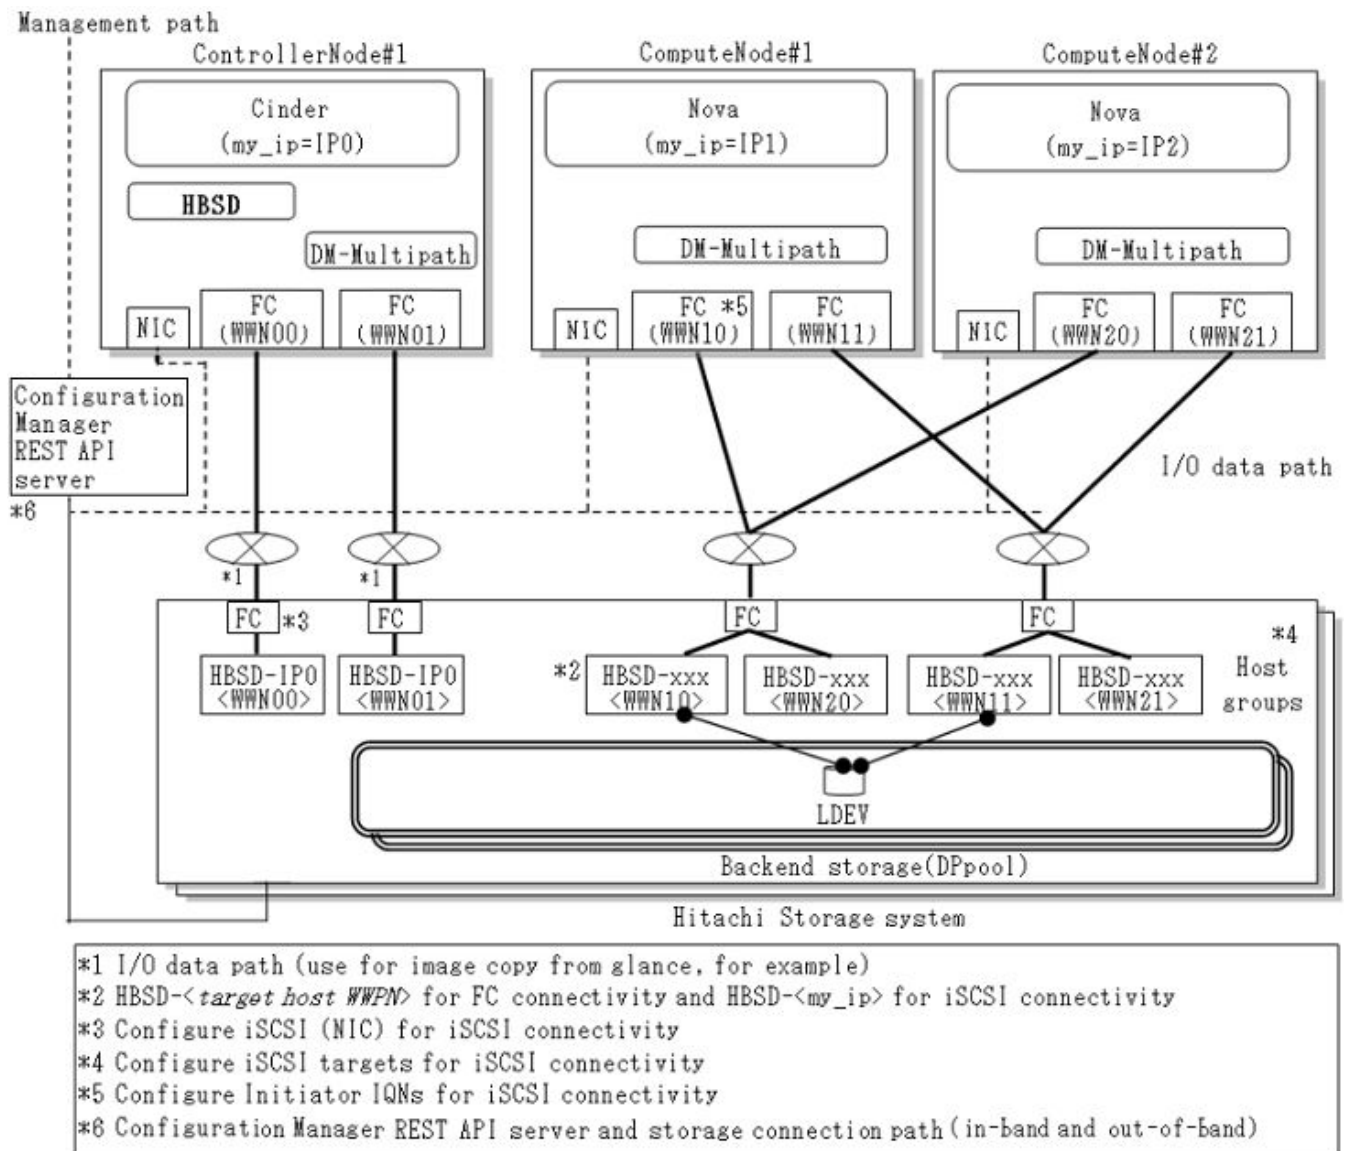


Figure 5 Connection example of node configuration

Single node configuration

Hitachi 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.
- If a read-only volume is attached, you cannot create volumes for the original snapshot for the volume.

Controller node without I/O data paths

In Red Hat OpenStack Platform and Ubuntu OpenStack, you can use the configuration on the Controller node without I/O data paths such as shown in the following image. This configuration is a viable solution if you do not want to give up Fibre Channel ports for management.



Note: You cannot use this configuration in SUSE OpenStack Cloud or Mirantis OpenStack.

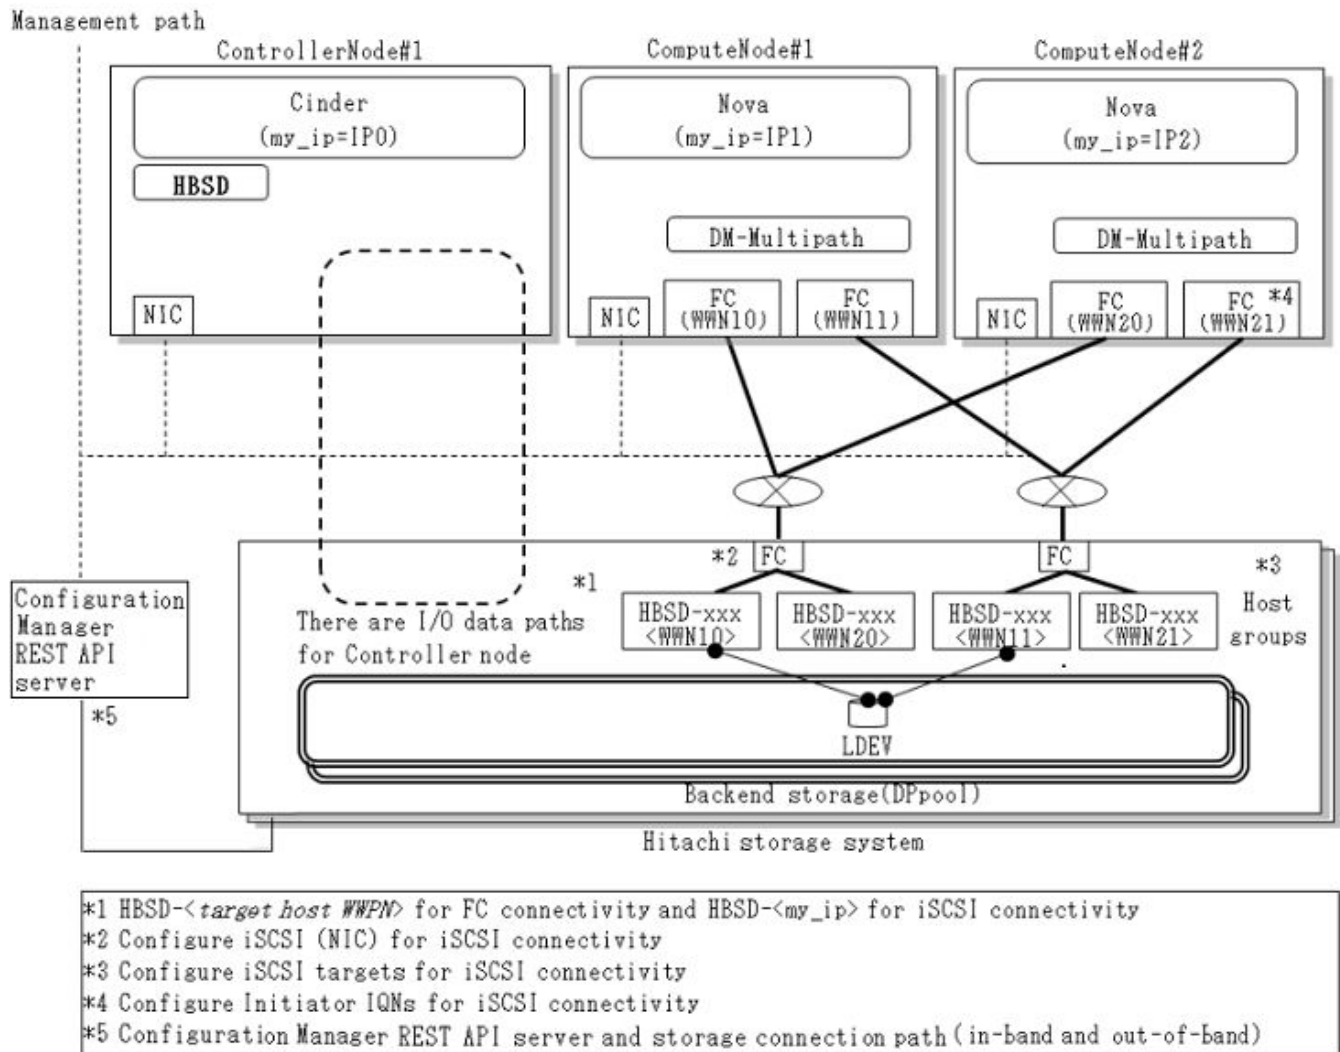


Figure 6 Connection example without I/O data paths for controller node

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

- Specify `hitachi_compute_target_ports` parameter
- Specify `hitachi_rest_pair_target_ports` parameter
- Do not specify `hitachi_target_ports` parameter for HBSD.

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

- Migrate a volume with a host assisted
- Create a volume from an image
- Create an image from a volume
- Backup 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 node use separate data path. Some OpenStack configuration installers might require you to specify a cinder driver when deploying 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 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	Items	Instruction
Controller node	my_ip for cinder service (/etc/cinder/cinder.conf)	Specify a unique value for the IPv4 address for 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 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.

Configure Configuration Manager REST API server setting

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 an in-band management the Configuration Manager REST API server connects to the storage system through SAN.

In an out-of-band management the Configuration Manager REST API server connects to the storage system through 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 Common operations in [Hitachi Command Suite Configuration Manager REST API Reference Guide](#).



Note: The software used for internal processing has changed from Command Control Interface (CCI) to Configuration Manager REST API in the version 5. See [Updating the management parameters in the cinder.conf file \(on page 48\)](#) for updating cinder.conf.

Configure storage resources

Fibre Channel connectivity

The following table describes several Fibre Channel storage configuration requirements.

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).
Resource groups	<p>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 (for command device and pool included), port, and host group.</p> <p>Note: The driver waits for up to two hours while another software has resource lock on the resource group used by its users.</p>
User accounts	<p>To access the storage system from the driver, you must create an account and be assigned to a user group with the following roles:</p> <ul style="list-style-type: none"> Storage Administrator (View Only) Storage Administrator (Provisioning)

Resource type	Contents
	<ul style="list-style-type: none"> Storage Administrator (Local Copy) Storage Administrator (Performance Management) <p>The user groups have management privileges for the created Resource-Group. You can also use the built-in Storage Administrator (View & Modify) user group.</p>
DP pool	<ul style="list-style-type: none"> 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 Thin Image pool capacity. 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	Create TI pool for copying by Thin Image.
Ports	<ul style="list-style-type: none"> Enable Port Security for port used by the driver. If you change port configuration used by the driver for connection, restart the <code>openstack-cinder-volume</code> service. If the volume is attached or detached without restarting the service, the operation might fail.
Host group for I/O data path	<p>Manual configuration:</p> <ul style="list-style-type: none"> Create a host group for the port that connects the controller node and compute nodes. In multipath configuration, create host group 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 to 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.

Resource type	Contents
	<ul style="list-style-type: none"> 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. <p>Automatic configuration:</p> <ul style="list-style-type: none"> Automatically creating the host group requires that the <code>hitachi_group_request</code> parameter to be <code>True</code>. 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 Zoning Manager must be enabled. Upon a cinder volume attach request, the driver automatically creates a host group for the VM port specified in <code>hitachi_target_ports</code> or <code>hitachi_compute_target_ports</code> if there is no host group defined for the port. In absence of existing host groups, the driver automatically creates one when copying a volume using copy functions on the storage system. Do not manually create or delete host names HBSD-pairXX (XX: number). The driver creates host groups and registers all WWNs for all ports specified in the <code>hitachi_target_ports</code> or <code>hitachi_compute_target_ports</code> parameter. Deleting host groups automatically requires the <code>hitachi_group_delete</code> parameter to be set to <code>True</code>.
Fibre Channel zoning	<p>Manual configuration:</p> <ul style="list-style-type: none"> For Fibre Channel switches between a node and the storage device, manually configure zoning. <p>Automatic configuration:</p> <ul style="list-style-type: none"> When the driver attaches a volume to the target VM, Fibre Channel zoning manager automatically creates zoning between the target node and the port used by the driver. Fibre Channel zoning manager must be enabled, and <code>hitachi_zoning_request</code> is <code>True</code>.

iSCSI connectivity

The following table describes the several iSCSI storage configuration requirements.

Resource type	Contents
All resources	All storage resources, such as DP pools and iSCSI targets, must have a name in order for HBSD to use them (name fields cannot be left blank).
Resource groups	<p>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 (for command device and pool included), port, and host group.</p> <p>Note: HBSD waits for up to two hours while software other than HBSD has resource lock on the resource group used by its users.</p>
User accounts	<p>To access the storage system from HBSD, you must create an account and be assigned to a user group with the following roles:</p> <ul style="list-style-type: none"> ▪ Storage Administrator (View Only) ▪ Storage Administrator (Provisioning) ▪ Storage Administrator (Local Copy) ▪ Storage Administrator (Performance Management) <p>The user groups have management privileges for the created Resource-Group. You can also use the built-in Storage Administrator (View & Modify) user group.</p>
DP pool	Create a DP pool that is used by HBSD backend. HBSD manages virtual capacity (capacity reserved for over-provisioning of the actual pool capacity). Set a proper over-subscription rate for your environment. For example, a 100% over-subscription rate means that space for only actual capacity is guaranteed.
TI pool	Create TI pool for copying by Thin Image.
Ports	<p>Enable port security for the ports used by HBSD. Configure the IP address and the TCP port number.</p> <p>If you change the port configuration used by HBSD for connection, restart the <code>openstack-cinder-volume</code> service. If the volume is attached or detached without restarting the service, the operation might fail.</p>

Resource type	Contents
iSCSI target for I/O data path	<p>Manual configuration:</p> <ul style="list-style-type: none"> ▪ Create an iSCSI target for the port that connects the controller node and compute nodes. In multipath configuration, create host group for all connecting ports. ▪ The iSCSI target must to be named "HBSD-<my_ip>" (for example, HBSD-10.20.30.40). Use the same IP address for <i>my_ip</i> that you use for 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 the multipath configuration, register the same target IQN in all iSCSI targets connected to a node. ▪ In the multipath configuration, set HostModeOption=83 on the created iSCSI targets. ▪ 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. <p>Automatic configuration:</p> <ul style="list-style-type: none"> ▪ When HBSD attaches a volume to a VM instance, HBSD automatically creates an iSCSI target for the VM ports that are specified in <code>hitachi_target_ports</code> or <code>hitachi_compute_target_ports</code> if no SCSI target exists for the ports. ▪ In absence of existing iSCSI targets, HBSD automatically creates one when copying a volume using copy functions on the storage. Do not manually create or delete the iSCSI targets HBSD-pairXX (XX: number). ▪ Automatic creation of iSCSI targets requires that the <code>hitachi_group_request</code> parameter is <code>True</code>. HBSD creates iSCSI targets and registers all Initiator IQNs for all ports used by a node that are specified in the <code>hitachi_target_ports</code> or <code>hitachi_compute_target_ports</code> parameter. HBSD registers the target IQN <code><Initiator IQN for the target node>.target</code> with the iSCSI targets. ▪ Deleting iSCSI targets automatically requires the parameter <code>hitachi_group_delete</code> to be <code>True</code>.

Resource type	Contents
	<p>Considerations for multipath configurations:</p> <ul style="list-style-type: none"> 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, HBSD 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 HBSD to match.

Configure Configuration Manager REST API server

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

For details, see Registering a storage system in [Hitachi Command Suite Configuration Manager REST API Reference Guide](#).

Procedure

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

```
# cat storage.json
{
  "model": "VSP F400",
  "serialNumber": 411111,
  "svplp": "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
}
```

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.

Install Hitachi Block Storage Driver for OpenStack

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
HBSD file	1 MB	<PATH> /usr/share/pyshared/ hitachi-hbsd
rootwrap file	1 MB	/etc/cinder/rootwrap.d/
sudo file	1 MB	/etc/sudoers.d
hbsdgetinfo	1 MB	/opt/hitachi/hbsd/
Log files	1 MB	<PATH> /var/log/hbsd/

- Check the OpenStack package installation for any missing or invalid packages (see Software). 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 Red Hat and SUSE

Use the **rpm** command to install Block Storage Driver for OpenStack. You must log in as a super user (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
```

For Red Hat, use **systemctl** command.

3. After the service is stopped, run the **rpm** command to install the driver. Install the appropriate RPM package for your OpenStack environment. The following table shows the Block Storage Driver for OpenStack RPM package for each supported OpenStack distribution.

RPM package name	RPM storage directory
hbsd.5.1.0-0_11.0.noarch.rpm	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 super user (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. The following table shows the DEB package for each supported OpenStack distribution.

DEB package name	DEB storage directory
hbsd-5.1.0-0-11.0_all.deb	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 an Block Storage Driver for OpenStack configuration

Normally, an 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. This topic describes additional necessary preparation steps and how to add the Block Storage Driver for OpenStack configuration, which vary depending on which OpenStack configuration installer you use.

Using Director (Red Hat OpenStack Platform)

A Puppet agent runs on the controller node when an OpenStack environment is built with the Director. The Puppet agent resets `cinder.conf` for the controller node and any configuration you have built in the driver will be lost. To avoid this, disable Puppet agent on all the controller nodes.

```
# puppet agent --disable
```

To configure the driver, edit the configuration file `/etc/cinder/cinder.conf` on all controller nodes.



Note: While the Puppet agent is being stopped, nodes cannot be added by the Director. For this reason, if you need to add nodes, first back up the `cinder.conf` and run the Puppet agent, and then add nodes. When you have finished adding nodes, disable the Puppet agent and restore the `cinder.conf` from the backup.

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** field. Click **Add Backend** to enter the name of the backend.
4. Enter the driver path in the **Driver** field and the parameters in the **Options** field for the backend you have added.

5. In the **Deployment** field, specify the controller node (only one selectable) in the **cinder-controller** field. And specify the backend(s) in the **cinder-volume** field.
6. Click **Apply** to save the configuration.

Using FUEL (Mirantis OpenStack)

To configure HBSD 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`.

`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

A driver configuration is added using the OpenStack configuration installer provided by the distribution in use (see [Add an Block Storage Driver for OpenStack configuration \(on page 42\)](#)). 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 HBSD configuration is shown below, followed by the description of each section of the file. For details about the parameters, see [Parameters \(on page 63\)](#).

```
# 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_thin_pool=1
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.

Edit policy.json

The consistency group features are disabled by default. To perform consistency group operations, you must edit the policy file provided by the openstack-cinder package (/etc/cinder/policy.json) as shown below and restart the Cinder API service:

```
/sbin/service openstack-cinder-api restart
```

Default settings (consistency group functions disabled):

```
"consistencygroup:create": "group:nobody",
"consistencygroup:delete": "group:nobody",
"consistencygroup:update": "group:nobody",
"consistencygroup:get": "group:nobody",
"consistencygroup:get_all": "group:nobody",
"consistencygroup:create_cgsnapshot" : "group:nobody",
"consistencygroup:delete_cgsnapshot": "group:nobody",
"consistencygroup:get_cgsnapshot": "group:nobody",
"consistencygroup:get_all_cgsnapshots": "group:nobody",
```

Modified settings (consistency group functions enabled):

```
"consistencygroup:create" : "",
"consistencygroup:delete": "",
"consistencygroup:update": "",
"consistencygroup:get": "",
"consistencygroup:get_all": "",
"consistencygroup:create_cgsnapshot" : "",
"consistencygroup:delete_cgsnapshot": "",
"consistencygroup:get_cgsnapshot": "",
"consistencygroup:get_all_cgsnapshots": "",
```

Restart Cinder

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 Red Hat, use the **systemctl** command:

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

- For SUSE, use **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 **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 HBSD.

Procedure

- After the Cinder service is restarted, confirm that HBSD is in use by verifying the following:
 - The message `MSGID0003-I` is output to the log file `/var/log/cinder/volume.log` (Red Hat OpenStack Platform) or `/var/log/cinder/cinder-volume.log` (SUSE OpenStack Cloud, Mirantis OpenStack, and Ubuntu OpenStack) for each backend that has been set.

```
# /bin/grep MSGID0003-I /var/log/cinder/volume.log
2014-03-10 11:34:05.564 28060 WARNING
cinder.volume.drivers.hitachi.hitachi_common [req-c9c94628
-988a-4bcb-8c6d-ac48910e5052 None None] MSGID0003-I: The
storage backend can be used. (config group: VSP G1000)
```

- The host status for the specified backend is up by using the `cinder service-list` command as shown in the example. The verified backend(s) can be used for volume operations. If the message is not output, review and update the configuration file.

```
# /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

- 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

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

Update the Hitachi Block Storage Driver for OpenStack

Update the Hitachi Block Storage Driver for OpenStack software for Red Hat and SUSE

Use the procedure described below to update the HBSD software for Red Hat and SUSE.

Procedure

1. Stop the Cinder volume service.

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

For Red Hat, use the **systemctl** command.

2. After the service has stopped, install a new version of the driver as shown below. The following table provides the information about the RPM packages and supported OpenStack versions.

```
# /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 Red Hat and SUSE \(on page 40\)](#).

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

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

Procedure

1. Stop the Cinder volume.

```
# /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 \(on page 41\)](#).

Updating the management parameters in the cinder.conf file

The software used for internal processing was updated from CCI to Configuration Manager REST API in Block Storage Driver for OpenStack version 5.0. If you are updating from a version earlier than Block Storage Driver for OpenStack version 5.0, you must update the parameters in the *cinder.conf* file.

Procedure

1. Uninstall the Command Control Interface.
2. Modify the parameters in the *cinder.conf* file according to the table below. For details on editing the *cinder.conf* file, see [Edit cinder.conf \(on page 43\)](#). For details on each parameter, see [Parameters \(on page 63\)](#).

Old Parameter	New Parameter
hitachi_storage_cli	Delete this parameter.
hitachi_storage_id	The format of this parameter value has changed.
hitachi_horcm_numbers	hitachi_rest_number The format of this parameter value has changed.
hitachi_horcm_user	hitachi_rest_user
hitachi_horcm_password or hitachi_horcm_password_path	hitachi_rest_password or hitachi_rest_password_path
hitachi_horcm_add_conf	Delete this parameter.
hitachi_horcm_name_only_discovery	hitachi_rest_name_only_discovery
hitachi_horcm_disable_io_wait	hitachi_rest_disable_io_wait
hitachi_horcm_pair_target_ports	hitachi_rest_pair_target_ports
None	Add this parameter: hitachi_rest_api_ip

3. Prepare the Configuration Manager REST API. See [Configure Configuration Manager REST API server \(on page 38\)](#).
4. Install the Block Storage Driver. See [Update the Hitachi Block Storage Driver for OpenStack \(on page 48\)](#).

Uninstall Hitachi Block Storage Driver for OpenStack

Follow the procedure below to uninstall the driver.

Procedure

1. Stop the Cinder volume service.

- For SUSE:

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

- For Red Hat, use **systemctl** command.
- For Ubuntu and Mirantis:

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

2. Uninstall the driver.

- For Red Hat and 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 \(on page 50\)](#).
 - `user_friendly_names` is set by "no"
 - `find_multipaths` is set by "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 parameter \(on page 70\)](#) 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. List this parameter in each backend definition section in the Cinder configuration file.
- Configure DM-Multipath for the port in the Block Storage Driver for OpenStack parameter `hitachi_target_ports` and `hitachi_compute_target_ports`. List this parameter in each backend definition section in the Cinder configuration file.
- 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 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 \(on page 50\)](#)).
- 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 overview

You can configure multiple controller nodes in a cluster to ensure high availability. If one node fails, the system falls 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 HBSD.

**Note:**

HBSD cannot be used in a High Availability (HA) configuration in SUSE OpenStack Cloud or Mirantis OpenStack. In addition, HA can detect single storage-control path failures, which can cause the cinder service to abend, but it may not be able to detect other types of failures.

Pacemaker configurations for Cinder volumes

HBSD 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** command. The following example shows configuring 10s for the start-delay. You can reconfigure node behaviors after the node has been rebooted.

Build High Availability environment of controller node

Build HA in controller node by referring to the documentation of the HA cluster software in use. You can use the procedure below to configure HBSD after building HA configuration using Foreman (Red Hat OpenStack Platform 10).

Procedure

1. Stop Puppet agent by disabling Puppet agent on all controller nodes.

```
# puppet agent -disable
```

2. To run the HBSD backend, configure the pacemaker of the cinder volume by referring to Pacemaker configurations for Cinder volumes.
3. Remove the cinder volume from the monitoring list of Pacemaker temporarily.

```
# pcs resource disable cinder-volume
```

4. Configure the cinder volume service and HBSD on all controller nodes. See the configuration values for the cinder volume services in the following table.

Item	Content
my_ip parameter	Configure unique IP address for each node.
host parameter	If you specify a common host name for all backend, write this parameter to the DEFAULT section for the configuration file of Cinder. Configure the same value (for example, <code>ha_cluster</code>) for each Controller node. This parameter is unnecessary when <code>backend_host</code> parameter is specified.
backend_host parameter	If you specify a host name for each backend, write this parameter to each backend section for the configuration file of Cinder. Configure the same value (for example, <code>backend_ha_cluster</code>) for each backend of the Controller node. If this parameter and the <code>host</code> parameter are specified, this parameter is used.
backend of HBSD	Configure same value for each Controller node.

5. Configure the control paths to the storage device.
6. Restart all the controller nodes.

```
# pcs cluster stop
# reboot
```

7. After starting the controller node, confirm that the backend of HBSD is running properly. See [Verify Block Storage Driver for OpenStack installation \(on page 46\)](#). In addition, confirm that the `cinder-volume` service is running with Active-Standby configuration by using the `pcs status` command.



Note: Once the Puppet agent is stopped, nodes cannot be added by Foreman. Back up the `cinder.conf` file and run Puppet agent and then add the node. After adding nodes, disable Puppet agent and restore `cinder.conf` from the backup.

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 HA 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 has caused Cinder volume processing to stop, find the volumes being processed using the `cinder list` command.
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 using the `cinder metadata-show` command and identify the LDEV number
 - b. Find out the storage operation from the volume status.
 - c. From the LDEV number, get the storage operation details using the management software in use to restore to a normal state.
 - d. Change the status for the corresponding entry on CinderDB to "available" using the `cinder reset-state` command.
 - e. Go to the step 6.
4. Identify the incomplete volume processing from the `volume.log/debug.log` around the time the failover has occurred.
5. Check if there are any LDEVs associated with the volume. If none, then go to step 5d; if otherwise, perform the following steps:
 - a. Obtain the pair information from the LDEV ID 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 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 3-5.

Chapter 4: Hitachi extended functions

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

Copy methods

When cloning volumes or creating volumes from snapshots, you can choose from one of the two copy methods provided by backend storage. Thin Image (Clone) and Thin Image (Snapshot). When creating snapshots, the copy method is set to Thin Image (Snapshot). The copy method is determined by the `hitachi_default_copy_method` parameter and the `copy_method` metadata for the destination volume. The following table describes which copy method is used under each condition. 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.

Table 3 Copy methods used in volume data copy

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	F	F	F	T	T	T	F	F	T	F
		FULL or THIN	F	F	F	F	F	F	T	F	F	T

Item		Value	Setting Pattern									
			1	2	3 ¹	4 ¹	5	6 ¹	7	8	9 ¹	10
	copy_method metadata	Not set	T	F	F	T	F	F	F	T ²	T ²	F
		FULL	F	T	F	F	T	F	F	F	F	F
		THIN	F	F	T	F	F	T	F	F	F	F
		FULL or THIN	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 with Thin Image (Snapshot) is not recommended because source volumes or snapshots cannot be deleted until the destination volumes are deleted and neither source volumes, source snapshots, nor destination volumes can be extended or unmanaged.
2. It is impossible for the user to 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 to the user.

**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 a cascade configuration with Thin Image (Snapshot), you cannot delete any intermediary volumes. If you want to delete an intermediary volume, you must delete the volumes starting from the last volume until the volume you want to delete is in the last position, then delete it, too.
- With 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 as shown in the following table. If the number of pairs exceeds the maximum, copy operations will not be executed successfully.

Max. copy pairs per volume	Max. copy pairs per storage system	Max. copy pairs per consistency group	Max. consistency group per storage system
Thin Image (Clone): 1,024 Thin Image (Snapshot): 1,024	The number of ports used by the driver x 1024	Thin Image (Clone): 8,192 ¹ Thin Image (Snapshot): 8,192 ¹	Thin Image (Clone): 2,048 ^{1, 2} Thin Image (Snapshot): 2,048 ^{1, 2}
Note: <ol style="list-style-type: none"> 1. Including the number of pairs or consistency groups that are not created by the driver. 2. The maximum number of consistency groups for Thin Image combined for each storage system is 2,048. 			

Create a volume from snapshot or create volume from volume

To create a volume from a snapshot or create volume from volume, set the `copy_method` metadata for the volume (`dst_volid`) to be created. The following is the syntax for specifying the copy method.

```
# /usr/bin/cinder create [--name name-of-the-volume-to-be-created] [--
snapshot-id src_snapid | --source-uuid src_uuid] size --metadata
copy_method=FULL
# /usr/bin/cinder metadata {name-of-the-volume-to-be-created | dst_uuid} unset
copy_method
```

The following section provides an example of creating a volume from a snapshot as illustrated in the following diagram.

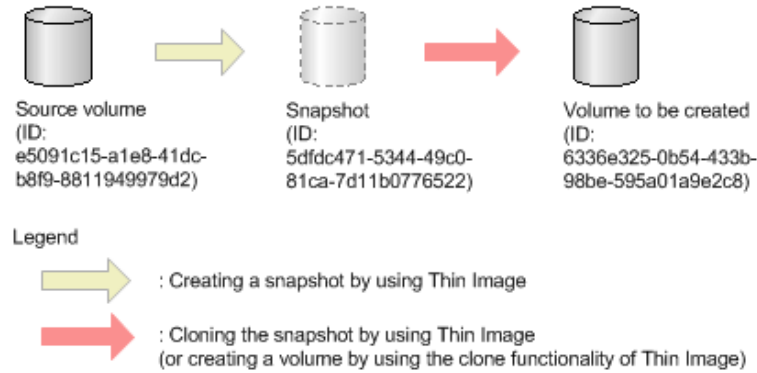


Figure 7 Creation of volume from snapshot example

Procedure

1. Create a volume called "create_volume_from_snapshot_test" from the snapshot using Thin Image (Clone).

```
# cinder create --name create_volume_from_snapshot_test --snapshot_id
5dfdc471-5344-49c0-81ca-7d11b0776522 1 --metadata copy_method=FULL
```

Property	Value
attachments	[]
availability_zone	nova
bootable	false
created_at	2014-06-12T10:39:05.000000
description	None
encrypted	False
id	6336e325-0b54-443b-98be-595a01a9e2c8
metadata	{}
name	create_volume_from_snapshot_test
os-vol-host-attr:host	OpenStack-LPAR08.localdomain@VSP
os-vol-mig-status-attr:migstat	None
os-vol-mig-status-attr:name_id	None
os-vol-tenant-attr:tenant_id	8223735aee1a418798a84362249f9c8f
size	1
snapshot_id	5dfdc471-5344-49c0-81ca-7d11b0776522
source_uuid	None
status	creating
user_id	e645e97e84554ca2bec2d1205b53b080
volume_type	VSP

- Using the `cinder list` command, verify that the volume has been successfully created.

**Note:**

- You can clone volumes using either Thin Image (Clone) or Thin Image (Snapshot). However, if you use Thin Image (Snapshot), you will not be able to delete the intermediary volume once it is in the cascade configuration. Therefore, it is generally recommended that you use Thin Image (Clone).
- When you have cloned a volume or created a volume from a snapshot, volume data copy may be in process internally even after the command has returned. You can delete the destination volume or the snapshot that are in such status, in which situation internal volume data copy is stopped and the volume or the snapshot is removed.

Port assignment for each volume type that can be attached

Defining particular ports in the Hitachi supported extra spec `hbsd:target_ports` determines which of the ports specified in 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 commands show an example of specifying ports using the extra spec. 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.

Data deduplication and compression

Deduplication and compression (dedupe/comp) is a method of improving storage utilization by data reduction. Once your storage administrator has enabled dedupe/

comp for the HDP 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 following documentation from [Related documents \(on page 9\)](#):

- *Provisioning Guide for Open Systems: Hitachi Virtual Storage Platform G1000, G1500, and F1500, MK-92RD8014*
- *Provisioning Guide: Hitachi Virtual Storage Platform Gx00 and Fx00 Models, MK-94HM8014*

Enabling deduplication and compression

Storage Systems

For a list of the storage devices that support the dedupe/comp functionality, see [Storage devices \(on page 23\)](#)

Enabling dedupe/comp functionality

To use the dedupe/comp function, your storage administrator must first enable the dedupe/comp function for the HDP pool.

For example, to enable dedupe/comp for a Hitachi Dynamic Provisioning (HDP) pool with a pool ID of 6 and specifying a yet unused LDEV with an LDEV ID of 1000 for creating a system data volume used for dedupe/comp, perform both steps listed below. The dedupe/comp function uses the specified LDEV as the volume for managing the deduplication and compression.

```
# raidcom modify pool -pool 6 -deduplication yes -ldev_id 1000
# raidcom modify pool -pool 6 -subscription 65535
```



Note:

In the example above, `-subscription 65535` sets the virtualization-excess ratio for the pool ID, where 65535 indicates that the virtualization-excess ratio is unlimited. If you set the virtualization-excess ratio to some other value, keep in mind that capacity for the system data volume is 40TB.

After enabling dedupe/comp for a Hitachi Dynamic Provisioning (HDP) pool, reference the pool information and verify that dedupe/comp is set. Running the command listed below displays all the pool information. In the information for all the pools, if the Num column of the target pool line is 1 or more, dedupe/comp has been enabled for the target pool.

```
# raidcom get pool -key saving
```

Creating a volume with dedupe/comp enabled

To create a volume with the dedupe/comp settings enabled, you need to enable dedupe/comp for the relevant volume type.

Procedure

1. To enable the dedupe/comp settings, specify `deduplication_compression` for `hbsd:capacity_saving` in the extra specs for the volume type. The example below shows how to enable the dedupe/comp 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 using the volume type created in the previous step, you can create a volume with the dedupe/comp function enabled.

The example below shows how to create a volume (10 GB) with the dedupe/comp 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_vol_id	None
status	creating
updated_at	None
user_id	60361daadb8d4aee8034c541f5089323
volume_type	VSP-G1500-1

Deleting a volume with dedupe/comp enabled

The **cinder delete** command completes 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: Parameters

This section provides syntax for both the Hitachi Block Storage Driver for OpenStack-specific parameters and the Cinder-specific parameters.

Hitachi Block Storage Driver for OpenStack specific parameters

The table below describes the parameters specific to Block Storage Driver for OpenStack software.

Parameter	Required or optional	Default value	Description
<code>hitachi_storage_id</code>	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.
<code>hitachi_pool</code>	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 <code>hitachi_pool</code> for other purposes.
<code>hitachi_thin_pool</code>	Required	None	Specifies the ID (integer) or the name of the TI pool that stores LDEVs for volumes or snapshots. If this parameter is not specified, Thin Image cannot be used to copy volumes. Do not use the TI pool specified for <code>hitachi_thin_pool</code> for other purposes.

Parameter	Required or optional	Default value	Description
<code>hitachi_ldev_range</code>	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, the entire range permitted by the storage is used, or, this would be the permitted range in a resource group for the user specified in the <code>hitachi_rest_user</code> . You can specify integer values by using a decimal format or colon-separated hexadecimal format (<code>xx:yy:zz</code>).
<code>hitachi_rest_number</code>	Optional	0	Specifies the number for identifying the Cinder driver instance. Created by adding together the storage device copy group name, device group name, and snapshot name. Integer value (0-255).
<code>hitachi_rest_user</code>	Required	None	Specifies the username 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.
<code>hitachi_rest_password</code> or <code>hitachi_rest_password_path</code>	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 <code>hitachi_rest_password_path</code> parameter. For more information about how to secure the password, see Secure your storage system password (on page 38) .

Parameter	Required or optional	Default value	Description
<code>hitachi_rest_name_only_discovery</code>	Optional	False	<p>Sets 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-<i><target host WWPN></i>" or "HBSD-<i><IP Address of the connecting host (my_ip)></i>", and ignores the rest. If host groups or iSCSI targets are created under different names, attaching and detaching volumes will result in an error.</p> <p>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, search might take more time.</p>
<code>hitachi_target_ports</code>	Optional	None	<p>Specifies the controller port name to search host groups (iSCSI targets) when attaching volumes. Specify the format of port name conformed to the using storage (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 <code>hitachi_compute_target_ports</code> is None (default).</p> <p>When this parameter is not specified and <code>hitachi_compute_target_ports</code> as well as <code>hitachi_rest_pair_target_ports</code> are specified, the driver operates as the environment without the I/O data path of the Controller node.</p>

Parameter	Required or optional	Default value	Description
<code>hitachi_compute_target_ports</code>	Optional	None	Specifies the controller port name to search host groups (iSCSI targets) when attaching volumes to compute node. Specify the format of port name conformed to the using storage (for example, VSP G1000: CL1-A). To use the multipath connection, specify the controller port names separated by a comma. If <code>None</code> (default) is specified, the controller port specified in <code>hitachi_target_ports</code> is used.
<code>hitachi_rest_pair_target_ports</code>	Optional	None	Specifies a port name for creating a host group (iSCSI target) of the HBSD-pair XX (XX is number) to register an LDEV when the volume is copied with storage function. Specify the format of port name conformed to the using storage (for example, VSP G1000: CL1-A). To specify more than one port name, separate the names with commas. If <code>None</code> (default) is specified, the controller port specified in <code>hitachi_target_ports</code> is used.
<code>hitachi_group_request</code>	Optional	False	Specifies <code>True</code> (true) or <code>False</code> (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 in the <code>hitachi_target_ports</code> or <code>hitachi_compute_target_ports</code> parameter. If <code>False</code> (default) is specified, it does not created the host group automatically.

Parameter	Required or optional	Default value	Description
<code>hitachi_group_delete</code>	Optional	False	Defines specifications for automatically deleting empty host groups or iSCSI targets assigned to the target ports specified in the <code>hitachi_target_ports</code> or <code>hitachi_compute_target_ports</code> parameter. Set false to not delete automatically. Set true to delete automatically. If you specify true, you must also specify true in the <code>hitachi_group_request</code> parameter. If the Fibre Channel Zone Manager is used (the <code>hitachi_zoning_request</code> parameter must be enabled), Fibre Channel zones configured for the host groups or iSCSI targets will be deleted.
<code>hitachi_zoning_request</code>	Optional	False	Specify this parameter if you are using a Fibre Channel connection. Specifies <code>True</code> (true) or <code>False</code> (false) to determine whether Fibre Channel zoning between storage and connection target node is configured automatically by collaborating with Fibre Channel zoning manager. To configure automatically, you need to enable Fibre Channel zoning manager. See the online manual for the configuration of Fibre Channel zoning manager. If <code>False</code> (default) is specified, Fibre Channel zoning is not configured regardless of availability of Fibre Channel zoning manager.

Parameter	Required or optional	Default value	Description
<code>hitachi_use_chap_auth</code>	Optional	None	<p>Specify this parameter if you are using an iSCSI connection. When using CHAP authentication of the iSCSI targets, set this parameter to <i>True</i>. For details about CHAP authentication, see the following documents:</p> <ul style="list-style-type: none"> Provisioning Guide for Open systems: Hitachi Virtual Storage Platform G1000, G1500, and F1500, MK-92RD8014 Provisioning Guide: Hitachi Virtual Storage Platform Gx00 and Fx00 Models, MK-94HM8014
<code>hitachi_auth_user</code>	Optional	None	<p>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 nothing is specified, HBSD-CHAP-user is set.</p> <p>If <code>True (true)</code> is specified for the <code>hitachi_use_chap_auth</code> parameter, this parameter is required</p>
<code>hitachi_auth_password</code>	Optional	None	<p>Specify this parameter if you are using an iSCSI connection. Specifies the password for the <code>hitachi_auth_user</code> setting. See the relevant user documentation. If nothing is specified, the HBSD-CHAP-password is set.</p> <p>If <code>True (true)</code> is specified for the <code>hitachi_use_chap_auth</code> parameter, this parameter is required.</p>

Parameter	Required or optional	Default value	Description
hitachi_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 <code>hitachi_auth_user</code> . Either the <code>hitachi_auth_password</code> parameter or the <code>hitachi_auth_password_path</code> should be specified if using CHAP authentication. HBSD 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, <code>/opt/hitachi/hbsd/rest_password</code>).
hitachi_default_copy_method	Optional	FULL	Specifies the volume copy method. You can specify <code>FULL</code> for Thin Image (Clone) or <code>THIN</code> for Thin Image (Snapshot). For snapshot creation, <code>THIN</code> is always prioritized over the specified value.
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 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 synchronization is confirmed. You can specify a value in

Parameter	Required or optional	Default value	Description
			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 disabled, before performing volume detach operations, make sure that no storage I/O operations are being performed. By enabling this parameter, you can avoid the problem described in the Volume detach operation fails (on page 81) .
hitachi_rest_api_ip	Required	None	Specifies the REST API server hostname or IP address.
hitachi_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. REST API server failure is detected in a short time. False: Do not use TCP keepalive functionality in the communication with the REST API server.

Cinder-specific parameters

Default section

The table below describes details for the Cinder parameters set in the default section of the `cinder.conf` file. Only the parameters relevant to the Block Storage Driver for OpenStack are described here. See the example configuration file in [Edit cinder.conf \(on page 43\)](#).

Parameter	Required or optional	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 <code>enabled_backends</code> parameter.*
my_ip	Optional	127.0.0.1	Specifies the IP address of own host (of management path). You must use unique address that is not used in other nodes.
logging_context_format_string	Optional	See 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 controller node that the cinder-api is running. To reflect the configuration, restart the openstack-cinder-api service.

Parameter	Required or optional	Default value	Description
host	Optional	None	Specifies a common host name for all backends. In the HA configuration for controller nodes, configure the same value (for example, <code>ha_cluster</code>) for each controller node. This parameter is not required if the <code>backend_host</code> parameter is specified in the backend definition section.
Notes: * 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 default format for analysis of log.

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 `cinder.conf` file. See the example configuration file in [Edit cinder.conf \(on page 43\)](#).

Name	Required or optional	Default value	Description
volume_driver	Required	None	<p>Specifies the path of the driver to be used as follows:</p> <p>To use the FC driver:</p> <pre>cinder.volume.drivers.hitachi.hbsd.hbsd_fc.HBSEFCDriver</pre> <p>To use the iSCSI driver:</p> <pre>cinder.volume.drivers.hitachi.hbsd.hbsd_iscsi.HBSEISCSIDriver</pre> <p>If it is not specified, startup of the backend will fail.</p>
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 <code>True</code> or <code>False</code> 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 <code>False</code> (default).

Name	Required or optional	Default value	Description
backend_host	Optional	None	Specifies a host name for a target backend. In the HA configuration for controller nodes, configure the same value (for example, <code>backend_ha_cluster</code>) 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	None	<p>This parameter is for Image Volume Cache, which cannot be used with the Hitachi Block Storage driver . Set the value of this parameter to <code>False</code>.</p> <p>If you want to quickly create a volume from an image, you can use the Volume-backed Image function.</p>
image_upload_use_cinder_backend	Optional	None	<p>This parameter is for the Volume-backed Image, which can be used if the parameter below is set to the value shown.</p> <p><code>hitachi_default_copy_method=FULL</code></p> <p>To use this function, set the value of this parameter to <code>True</code> and set the parameter listed above to the value shown. To not use this function, set the value of the parameter to <code>False</code>.</p>
suppress_requests_ssl_warnings	Optional	False	Suppress SSL certificate warnings.

Name	Required or optional	Default value	Description
<code>driver_ssl_cert_path</code>	Optional	None	Can be used to specify a non-default path to a CA_BUNDLE file or directory with certificates of trusted Certificate Authorities.
<code>driver_ssl_cert_verify</code>	Optional	False	If set to <i>True</i> , the SSL certificate of the REST API server will be validated.

Chapter 6: Troubleshooting

This section describes troubleshooting procedures for Hitachi Block Storage Driver for OpenStack.

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, run the **rpm** or **dpkg** command to check the version of the driver to 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 if the update has resolved the problem.

- For Red Hat and SUSE:

```
# /bin/rpm -qi hbsd
```

- For Ubuntu and Mirantis:

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

hbsdgetinfo command

Function

The command **hbsdgetinfo** collects the data needed to investigate the 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 <i>directory</i>	Specify the directory to save the data is 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 <code>/tmp</code> directory will be used.
-e <i>cinder</i>	Specify this parameter to remove data belonging to the Cinder group (data to be collected by the data collection command <code>cinder</code>) 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 HBSD) setting	All files below <code>/etc/cinder/</code>	Common
	<code>/etc/sudoers.d/cinder *</code>	Common
Cinder log	All files below <code>/var/log/cinder/</code>	Common
HBSD log	All files below <code>/var/log/hbsd/</code>	Common
Setting related to FC	<code>/etc/multipath.conf</code>	Common
	All files below <code>/etc/multipath.d/</code>	Common
System log	<code>/var/log/messages*</code>	Common
Setup of FC-related ¹	<code>/etc/multipath</code> directory	Common
Build date ¹	<code>/opt/hitachi/hbsd/build_date/*</code>	Common
Notes: * 1. Ubuntu only.		

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
	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

Column	Command	Output file name	Groups
Working status related to FC	ls -l /dev/disk/by-path	ls -l_dev_disk_by-path	Common
	multipath -ll	multipath_-ll	Common
	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 ⁴	rpm -V python-cinder	Common
	rpm -V hbsd 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: <ol style="list-style-type: none"> 1. Red Hat, SUSE only. 2. Ubuntu, Mirantis only. 3. Script that collects ethtool information for all network interfaces in a system. 4. The output format is SM5DLUGT c <file>. See the rpm online documentation for more details for 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.
copying "/var/log/hbsd" ... 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 getting the logs, see Collecting maintenance information about the REST API about 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 of the problem.

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

Volume (snapshot) cannot be created

Check for an error message in the `volume.log`, and see [Messages \(on page 85\)](#) to troubleshoot the problem.

Volume (snapshot) cannot be attached

For error messages of Block Storage Driver for OpenStack in `volume.log`, see [Messages \(on page 85\)](#) 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 Hitachi Block Storage Driver for OpenStack \(on page 26\)](#).

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

Volume detach operations may fail because LUNs are in a state where they cannot be unmapped during the 3-minute period following I/O. 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`.

To set this option for host groups that already exist, follow the usage listed below.

Usage

To set the option:

```
# raidcom modify host_grp -port <port number> <host group name> -host_mode
LINUX -host_mode_opt 91
```

To verify the option:

```
# raidcom get host_grp -port <port number>
```

91 is displayed for HMO_BITS

To delete the option:

```
# raidcom modify host_grp -port <port number> <host group name> -host_mode
LINUX -host_mode_opt 100
```

Procedure

1. Obtain the host group name for each port used by the Hitachi Block Storage Driver. If `hitachi_group_request` is set to `True`, the auto-created hostgroup (iSCSI target) name will be *HBSD-<target host WWPN>* or *HBSD-<IP Address of the connecting host (my_ip)>*.

Example

```
# raidcom get host_grp -port CL1-A
PORT  GID GROUP_NAME          Serial#  HMD      HMO_BITS
CL1-A 120 HBSD-0123456789ABCDEF 200075  LINUX/IRIX
```

```
# raidcom get host_grp -port CL2-A
PORT  GID GROUP_NAME          Serial#  HMD      HMO_BITS
CL2-A 120 HBSD-0123456789ABCDEF 200075  LINUX/IRIX
```

2. Set the I/O suppression option for all host groups used by the Hitachi Block Storage Driver.

Example

```
# raidcom modify host_grp -port CL1-A HBSD-0123456789ABCDEF -host_mode
LINUX -host_mode_opt 91
# raidcom get host_grp -port CL1-A
PORT  GID GROUP_NAME          Serial#  HMD          HMO_BITS
CL1-A 120 HBSD-0123456789ABCDEF 200075   LINUX/IRIX  91
```

```
# raidcom modify host_grp -port CL2-A HBSD-0123456789ABCDEF -host_mode
LINUX -host_mode_opt 91
# raidcom get host_grp -port CL2-A
PORT  GID GROUP_NAME          Serial#  HMD          HMO_BITS
CL2-A 120 HBSD-0123456789ABCDEF 200075   LINUX/IRIX  91
```



Note: Setting the I/O check suppression option (host mode option 91) enables volumes to be detached even during disk input/output, which may damage the file system. Set it so that volumes are detached after input/out has completed.

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 volume status has returned to *available*, then retry deleting the volume(s).

Contacting Hitachi Vantara Corporation customer support

If you need to contact Hitachi Vantara customer support, provide as much information about the problem as possible. Collect the files below and provide 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.
- Copy of the operation logs of High Availability cluster software to collect data on the controller node in the High Availability cluster. For example, `/var/log/pacemaker`.

Chapter 7: Messages

Message output format

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

The following is the Block Storage Driver for OpenStack message format:

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

Where:

- *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
- *message*: Message content

The actual message content (indicated by 8 in the message format above) is output as follow:

```
MSGIDn1n2n3n4-t: Message text
```

Where:

- *MSGID*: Indicates that the message is an driver message
- *n1n2n3n4*: The serial number of the message
- *t*: The message type

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.

Type	Description
W	Warning message indicating that a driver error occurred or that an option was incorrectly specified, but that processing will continue.
I	Message providing additional information about processing. Messages that are neither error messages nor warning messages.

List of messages

Message ID	Message text	Driver status	Required actions
0000-I	The method <i><method name></i> was called. (config_group: <i><backend definition name></i>)	Continues processing the request.	None.
0001-I	The parameter of the storage backend. (config_group: <i><backend definition name></i>)	Continues processing the request.	None.
0002-I	The method <i><method name></i> completed successfully. (config_group: <i><backend definition name></i>)	Continues processing the request.	None.
0003-I	The storage backend can be used. (config_group: <i><backend definition name></i>)	Continues processing the request.	None.
0300-W	Failed to configure the internal logging. (ret: <i><return value></i> , stderr: <i><standard error output></i>).	Continues processing the request.	Verify that the user has sufficient permission for the directory.
0301-W	A LUN (HLUN) was not found. (LDEV: <i><LDEV number></i>)	Continues processing the request.	None.
0302-W	Failed to specify a logical device for the volume <i><volume ID></i> to be unmapped.	Continues processing the request.	None.

Message ID	Message text	Driver status	Required actions
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.
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.

Message ID	Message text	Driver status	Required actions
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.
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.

Message ID	Message text	Driver status	Required actions
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.
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 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 copy to the time estimated before 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, then retry the operation.

Message ID	Message text	Driver status	Required actions
0333-W	Retrying 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 dedupe/ compression function on volumes that just had the function disabled.
0334-W	Retrying 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.
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 text	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 text	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, 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 <code>copy_method</code> metadata to FULL, then clone the volume. Extend and use this newly created volume.

Message ID	Message text	Driver status	Required actions
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.
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.

Message ID	Message text	Driver status	Required actions
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.
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 <code>hitachi_add_chap_user</code> 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 <code>hitachi_group_request</code> 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.

Message ID	Message text	Driver status	Required actions
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.
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.

Message ID	Message text	Driver status	Required actions
0656-E	<p>The volume <volume ID> could not be restored. (reason: <detailed message>)</p> <ol style="list-style-type: none"> 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.
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.

Message ID	Message text	Driver status	Required actions
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.
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, delete the pair if the target ldev has a pair and retry the operation. The driver does not manage ldevs that are not DP-VOL. Create a DP-VOL and copy the data from the ldev and 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, 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.

Message ID	Message text	Driver status	Required actions
0705-E	Failed to manage the specified LDEV (< <i>LDEV number</i> >). 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.
0706-E	Failed to unmanage volume < <i>volume ID</i> >. The volume type must be <code>Normal</code> 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 <code>Normal</code> , retry the operation after copy completes and the pair is deleted. If otherwise, since the volume cannot be removed from Cinder management, set the <code>copy_method</code> metadata to <code>FULL</code> , 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 < <i>volume ID</i> >. The volume type must be <code>Normal</code> .	Stopped processing the request.	Make sure the target ldev information is correct. If cloned image the volume except <code>Normal</code> type, try the operation after having cloned a volume.
0711-E	A source volume for clone was not found. (volume_uuid: < <i>volume ID</i> >)	Stopped processing the request.	Specify the volume ID of existing bootable volume in <code>volume_uuid</code> .

Message ID	Message text	Driver status	Required actions
0712-E	Failed to create a consistency group snapshot. (copy method: <i><copy-method></i>) 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), then retry the operation if the copy method is FULL after making sure there are no other consistency snapshot operations. If the copy method is THIN, delete any unnecessary snapshots for the consistency group and retry the operation.
0714-E	Failed to delete unused full copy pair. (Copy Group: <i><copy group ></i>)	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 <i><volume-ID></i> to be added to the consistency group <i><consistencygroup-ID></i> .	Stopped processing the request.	Wait for the LDEVs to be created, then retry the operation.
0717-E	Multiple mappings for volume <i><volume_id></i> 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 (<i><volume-ID></i>) 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, then retry the operation.

Message ID	Message text	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, 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.
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 then initialize data reduction.

Message ID	Message text	Driver status	Required actions
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, 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>)	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: <source>, ID: <message ID>, message: <message>, cause: <cause>, solution: <solution>, code: <error code>)	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 <code>--force-host-copy True</code> option.

Message ID	Message text	Driver status	Required actions
0733-E	The REST API timed out. (job ID: <job_id>, job status: <status>, job state: <state>)	Stopped processing the request.	Verify the REST API operation log and storage log, then respond to the cause of the timeout. Verify volume status 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>)	Stopped processing the request.	Verify the REST API server operation log.
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.
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.

Messages output by the command **hbsdgetinfo**

Message text	Driver status	Required actions
hbsdgetinfo command started.	Continues processing the request.	None
hbsdgetinfo command completed successfully.	Continues processing the request.	None
output to <i>pathDesignated/</i> hbsdgetinfo- <i>yyyymmddhhmmss.tar.gz.</i>	Continues processing the request.	None
creating directory ... {succeeded failed}	Continues processing the request.	None
"< <i>list of collection files</i> > < <i>list of collection commands</i> >" is not found	Has stopped processing the requests.	Reinstall the driver.
"< <i>collection relevant files</i> > < <i>collection commands</i> >" is not found	Continues processing the request.	None
copying "< <i>fileName</i> >" ... {succeeded failed}.	Continues processing the request.	None
executing "< <i>commandName</i> >" ... {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.
"< <i>pass which you designated in -d</i> >" is not a directory	Has stopped processing the requests.	Specify a directory.

Message text	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 text	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.

Table 4 OpenStack main components

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 on 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 an 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.

Glossary

Cascade pair

A pair created by the Thin Image function.

Consistency group

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 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 (HDP) is the 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.

In-band

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

LDEV

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

Out-of-band

In an out-of-band management the Configuration Manager REST API server connects to the storage system through 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

2845 Lafayette Street

Santa Clara, CA 95050-2639 USA

www.HitachiVantara.com | community.HitachiVantara.com

Regional Contact Information

Americas: +1 866 374 5822 or info@hitachivantara.com

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

Asia Pacific: + 852 3189 7900 or info.marketing.apac@hitachivantara.com