

Hitachi Virtual Storage Software

1.12

Block Storage Administrator Guide

This guide describes and provides instructions for performing storage administration tasks for Hitachi Virtual Storage Software block (VSS block), including managing drives, volumes, and storage pools, allocating volumes, and using snapshots to back up and restore volumes.

© 2022, 2023 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 LLC (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 LLC 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 LLC.

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; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

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

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

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

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

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found in the product documentation, at <https://www.hitachivantara.com/en-us/company/legal.html> or https://knowledge.hitachivantara.com/Documents/Open_Source_Software.

Contents

Preface.....	9
Intended audience.....	9
Product version.....	9
Release notes.....	9
Changes made in this revision.....	10
Document conventions.....	10
Accessing product documentation.....	12
Getting help.....	12
Comments.....	12
Chapter 1: Getting started.....	13
Registering a license (CLI or REST API).....	13
Allocating volumes to compute nodes.....	14
Expanding storage pools (CLI or REST API).....	15
About Virtual Storage Software block GUI.....	18
Logging in to and logging out of the GUI.....	20
Dashboard window.....	22
Navigation bar.....	24
Displaying performance information about storage systems.....	26
Monitor link.....	28
Performing maintenance.....	30
Viewing the List window.....	32
Viewing the detailed information window.....	34
Refresh display.....	35
Help display.....	35
Pop-up message display.....	36
Chapter 2: Managing storage pools.....	37
Capacity information about storage pools.....	37
Obtaining a list of information about storage pools (CLI or REST API).....	39
Obtaining information about individual storage pools (CLI or REST API).....	42
Obtaining a list of capacity information about volumes (CLI or REST API).....	45
Obtaining capacity information about individual volumes (CLI or REST API)...	45
Obtaining a list of performance information about volumes (CLI or REST API).....	46

Obtaining performance information about individual volumes (CLI or REST API).....	47
Chapter 3: Managing your license.....	48
License management overview.....	48
Deleting a license (CLI or REST API).....	51
Obtaining a list of license information (CLI or REST API).....	52
Obtaining information about individual licenses (CLI or REST API).....	53
Editing the license warning threshold (CLI or REST API).....	54
Obtaining license settings (CLI or REST API).....	54
Chapter 4: Managing compute nodes and allocating volumes to compute nodes.....	56
Installing a compute node.....	56
Registering information about compute nodes (CLI or REST API).....	58
Obtaining a list of information about compute nodes (CLI or REST API).....	59
Obtaining information about individual compute nodes (CLI or REST API).....	60
Editing information about compute nodes (CLI or REST API).....	61
Deleting information about individual compute nodes (CLI or REST API).....	62
Registering information about the initiators for compute nodes (CLI or REST API).....	64
In the case of iSCSI connection (CLI or REST API).....	64
In the case of FC connection (CLI or REST API) (Virtual machine).....	65
Obtaining a list of information about initiators for compute nodes (CLI or REST API).....	67
Obtaining information about individual initiators for compute nodes (CLI or REST API).....	67
Deleting information about the initiators for compute nodes (CLI or REST API).....	68
Registering path information about compute nodes (CLI or REST API).....	70
Obtaining a list of path information about compute nodes (CLI or REST API)..	72
Obtaining specific path information about compute nodes (CLI or REST API).....	73
Deleting path information about compute nodes (CLI or REST API).....	74
Allocating volumes to compute nodes (CLI or REST API).....	76
Obtaining a list of information about allocation of volumes to compute nodes (CLI or REST API).....	79
Obtaining information about allocation of individual volumes to individual compute nodes (CLI or REST API).....	80
Canceling allocation of volumes to compute nodes (CLI or REST API).....	81
Releasing multiple connections between the volumes and compute nodes (CLI or REST API).....	82
Managing compute nodes.....	84
Overview.....	84

Registering compute nodes (GUI).....	85
Editing compute nodes (GUI).....	87
Deleting compute nodes (GUI).....	88
Reconnecting compute nodes and all the compute ports in full mesh (GUI).....	89
Allocating and canceling allocation of volumes to and from compute nodes.....	90
Creating volumes and allocating them to compute nodes (GUI).....	90
Allocating volumes to compute nodes (Volume) (GUI).....	91
Allocating volumes to compute nodes (Compute Node) (GUI).....	92
Canceling allocation of volumes to compute nodes (GUI).....	94
Chapter 5: Managing storage nodes.....	96
Obtaining a list of information about storage nodes (CLI or REST API).....	96
Obtaining information about individual storage nodes (CLI or REST API).....	97
Overview of capacity management of a storage node.....	98
Editing capacity management settings of a storage node (CLI or REST API).....	101
Obtaining a list of the capacity management settings of storage nodes (CLI or REST API).....	102
Obtaining capacity management settings of individual storage nodes (CLI or REST API).....	103
Performing maintenance operation for storage nodes.....	103
Overview.....	103
Performing storage node maintenance recovery (GUI).....	104
Verifying the conditions for performing storage node maintenance blocking (GUI).....	107
Performing storage node maintenance blocking (GUI).....	115
Chapter 6: Verifying the status of drive data relocation.....	118
Overview of drive data relocation.....	118
Approximate time until drive data relocation is completed.....	118
Verifying the drive data relocation status (CLI or REST API).....	120
Suspending drive data relocation (CLI or REST API).....	121
Resuming drive data relocation (CLI or REST API).....	122
Chapter 7: Managing compute ports.....	124
Obtaining a list of information about compute ports (CLI or REST API).....	124
Obtaining information about individual compute ports (CLI or REST API).....	125
Modifying settings of compute ports (CLI or REST API).....	126
Chapter 8: Managing drives.....	131
Overview of managing drives.....	131
Obtaining a list of information about drives (CLI or REST API).....	131
Obtaining information about individual drives (CLI or REST API).....	132
Removing drives (CLI or REST API).....	133

Replacing drives (CLI or REST API).....	136
Adding drives.....	138
Overview of adding drives.....	138
Inserting drives (CLI or REST API).....	140
Expanding a storage pool (CLI or REST API).....	140
Operating drives.....	145
Overview.....	145
Removing drives (GUI).....	145
Adding drives (GUI).....	147
Replacing drives (GUI).....	149
Turning on/off the locator LED (GUI).....	149
Importing a system requirements file.....	151
Overview.....	151
Importing a system requirements file (GUI).....	151
Chapter 9: Managing volumes.....	153
Creating volumes (CLI or REST API).....	153
Deleting volumes (CLI or REST API).....	155
Expanding volumes (CLI or REST API).....	157
Obtaining a list of volumes (CLI or REST API).....	158
Obtaining information about individual volumes (CLI or REST API).....	160
Modifying settings of individual volumes (CLI or REST API).....	162
Managing volumes.....	164
Creating volumes (GUI).....	164
Modifying settings of individual volumes (GUI).....	165
Expanding volumes (GUI).....	167
Deleting volumes (GUI).....	168
Chapter 10: Backing up and restoring volumes by using snapshots....	170
Overview of snapshots.....	170
Preparing to take a snapshot.....	182
Taking a snapshot (CLI or REST API).....	184
Deleting snapshots (CLI or REST API).....	186
Restoring P-VOLs from a snapshot (CLI or REST API).....	189
Obtaining P-VOL information for the volume (CLI or REST API).....	194
Obtaining a list of S-VOL information for the volume (CLI or REST API).....	195
Chapter 11: Obtaining information about storage clusters.....	198
Obtaining information about storage clusters (CLI or REST API).....	198
Obtaining a list of fault domain information (CLI or REST API).....	200
Obtaining information about individual fault domain (CLI or REST API).....	200
Obtaining health status (CLI or REST API).....	201

Verifying the cluster master node (primary) (CLI or REST API).....	201
Chapter 12: Managing protection domains.....	203
Obtaining a list of information about protection domains (CLI or REST API)..	203
Obtaining information about individual protection domains (CLI or REST API).....	204
Changing usage of internal processing I/O resources (CLI or REST API).....	205
Chapter 13: Managing storage controllers.....	207
Overview of the features.....	207
Capacity status of storage controllers.....	208
Obtaining a list of information about storage controllers (CLI or REST API)...	209
Obtaining information about individual storage controllers (CLI or REST API).....	210
Editing the settings for the storage controller (CLI or REST API).....	212
Chapter 14: Verifying the status of Rebuild (data reconstruction).....	214
Overview of Rebuild.....	214
Cases when the Rebuild is possible.....	216
Estimated time before completion of Rebuild.....	216
Verifying Rebuild status (CLI or REST API).....	222
Chapter 15: Troubleshooting.....	224
First points to confirm.....	224
Action to be taken for a volume failure.....	228
Resolving other error conditions.....	229
Action to be taken if KARS06132-C occurs when starting the storage cluster	231
Action to be taken when "Alerting" is shown in the license health status.....	235
Action to be taken when "Alerting" is shown in the Fault Domains health status.....	236
Determining the maintenance priority of storage nodes.....	237
Operations with dump log files (GUI).....	239
Overview.....	239
Creating dump log files (GUI).....	241
Downloading dump log files (GUI).....	243
Deleting dump log files (GUI).....	245
Appendix A: Capacity design (for HPEC 4D+1P).....	247
Capacity design principles (1) : Physical capacity (for HPEC 4D+1P).....	247
Capacity design principles (2) : Logical capacity (for HPEC 4D+1P).....	248
Capacity design method (1)-1 : Logical capacity from physical capacity (for HPEC 4D+1P).....	249
Capacity design method (2) : Physical capacity from logical capacity (for HPEC 4D+1P).....	250

Appendix B: Capacity design (for HPEC 4D+2P)	252
Capacity design principles (1) : Physical capacity (for HPEC 4D+2P).....	252
Capacity design principles (2) : Logical capacity (for HPEC 4D+2P).....	253
Capacity design method (1)-1 : Logical capacity from physical capacity (for HPEC 4D+2P).....	254
Capacity design method (2) : Physical capacity from logical capacity (for HPEC 4D+2P).....	255
Appendix C: Capacity design (for Mirroring)	257
Capacity design principles (1) : Physical capacity (for Mirroring).....	257
Capacity design principles (2) : Logical capacity (for Mirroring).....	258
Capacity design method (1)-1 : Logical capacity from physical capacity (for Mirroring).....	259
Capacity design method (2) : Physical capacity from logical capacity (for Mirroring).....	260
Appendix D: Capacity design when adding storage nodes	262
Adding capacity.....	262
Capacity design when adding storage nodes (for HPEC 4D+1P).....	262
Capacity design when adding drives (for HPEC 4D+1P).....	263
Capacity design when adding storage nodes (for HPEC 4D+2P).....	264
Capacity design when adding drives (for HPEC 4D+2P).....	265
Capacity design when adding storage nodes (for Mirroring).....	266
Capacity design when adding drives (for Mirroring).....	267
Appendix E: ALUA configuration guidelines	269
Required preparations.....	269
VMware.....	269
Windows.....	269
Linux.....	270
ALUA setting procedure.....	270
VMware (in the case of NMP).....	270
Windows (in the case of MPIO).....	271
Linux system (in the case of Device Mapper).....	271
Appendix F: Processing time for snapshot operations	273
Estimated processing time for snapshot operations.....	273
Glossary	276

Preface

This guide describes and provides instructions for performing storage administration tasks for Hitachi Virtual Storage Software block (VSSB), including managing drives, volumes, and storage pools, allocating volumes, and using snapshots to back up and restore volumes.

This manual applies to both the virtual machine and bare metal models of VSSB.

- Sections in this manual marked with (Virtual machine) apply to the virtual machine model.
- Sections in this manual marked with (Bare metal) apply to the bare metal model.

Please read this document carefully to understand how to use these products and maintain a copy for reference purposes.

Intended audience

This document is intended for system administrators, Hitachi Vantara representatives, and authorized service providers who are involved in installing, configuring, and operating Virtual Storage Software block.

Readers of this document should have at least the following knowledge and experience:

Skills using CLI and REST API

- Knowledge of hypervisor type virtualization environment
- Skills using CLI and REST API

Product version

This document revision applies to Virtual Storage Software block version 1.12 (01.12.xx.xx).

The version in this document is described only by [aa.bb], and [aa.bb.cc.dd] is used only when required.

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 the Hitachi Vantara Support Website: <https://knowledge.hitachivantara.com/Documents>.

Changes made in this revision

- Added information about multi-tenancy.
- Added an explanation of when drive blockage and node blockage rebuilding are fixed.
- Added information about the license when the storage cluster is rebuilt.
- Added information that can be obtained in volume information acquisition.
- Revised the description of portSpeedDuplex, which is the acquisition information of compute ports.
- Revised the procedure for taking snapshots.
- Updated information about the rebuild processing time.
- Added information about license registration when rebuilding a storage cluster.
- Updated the explanation for disabling a storage node.
- Revised the description of SSL communication.







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.
<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	<p>Indicates text that is displayed on screen or entered by the user. Example: <code>pairdisplay -g oradb</code></p>

Convention	Description
< > 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 additional information.
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Important	Highlights information that is essential to the completion of a task.
	Caution	Warns the user of adverse conditions and/or consequences (for example, disruptive operations, data loss, or a system crash).
	CAUTION	Warns the user of a hazardous situation that, if not avoided, could result in major or minor injury.
	WARNING	Warns the user of a hazardous situation which, if not avoided, could result in death or serious injury.

Accessing product documentation

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

Getting help

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

[Hitachi Vantara Community](https://community.hitachivantara.com) 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 comments to doc.comments@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara LLC.

Thank you!

Chapter 1: Getting started

Registering a license (CLI or REST API)

To request and register a license, contact your sales account representative to purchase a license key. To purchase a license key, you must provide the storage cluster ID and type of license you require. Depending on the type of the license, you also must notify the licensed capacity and period.



Tip:

You can obtain the ID of the storage cluster by running the command described in *Obtaining information about storage clusters* or clicking the Cluster information icon in the GUI.



Caution:

After reconfiguring a storage cluster, the license key issued before reconfiguring a storage cluster cannot be used for a new storage cluster. In this case, contact your sales account representative to re-issue a license key that corresponds to the new storage cluster ID, and then register the new license.

Before you begin

- Required role: Storage

Procedure

1. Contact your sales account representative to purchase a license key.

To purchase a license key, you must provide the storage cluster ID and type of license you require. Depending on the type of the license, you also must notify the licensed capacity and period.

You can obtain the ID of the storage cluster by running the command described in *Obtaining information about storage clusters* or clicking the Storage Cluster Information icon in the GUI.

2. Register a license.

Run one of the following commands with the license key specified.

REST API: POST /v1/objects/licenses

CLI: license_install

Verify the job ID which is displayed after the command is run.

3. Verify the state of the job by specifying the job ID.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

4. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.


Allocating volumes to compute nodes

Perform the following procedures to connect compute nodes to the Virtual Storage Software block storage system:

- Verify the requirements for compute nodes to be installed.
- Register information about all compute nodes that handle the storage system.
- Register initiator information of compute nodes.
- Register path information of compute nodes.

In addition, to use the storage system from a compute node, in the storage pool, create a volume for storing user data.

To enable the compute node to handle the created volume, set up a path between the volume and the compute node.



Category	See:	Is there a GUI operation ?
Allocate volumes to compute nodes	1. Installing a compute node (on page 56)	
	2. Registering information about compute nodes (CLI or REST API) (on page 58)	Yes
	3. Registering information about the initiators for compute nodes (CLI or REST API) (on page 64)	
	4. Registering path information about compute nodes (CLI or REST API) (on page 70)	
	5. Creating volumes (CLI or REST API) (on page 153)	Yes
	6. Allocating volumes to compute nodes (CLI or REST API) (on page 76)	Yes

Expanding storage pools (CLI or REST API)

When you finish setting up Virtual Storage Software block, the storage pool is created with the name specified in the VSSB configuration file (SystemConfigurationFile.csv). The storage pool is empty. Therefore, you must add drives to the storage pool.

To perform this operation, you need the Storage role. If setup is completed, a user is created in ServiceAdministrators. Because the Service and Storage roles are set for ServiceAdministrators, you can expand storage pools by using the user created in ServiceAdministrators. By using this user and user ID as credentials when running commands, perform the following procedures.



Caution:

If there is an operation that was performed before this procedure, confirm that the process is completed or terminated due to an error before performing this procedure. If the operation that was performed before this procedure is terminated due to an error, and the event log indicates the action that needs to be taken before this procedure, take that action, and then perform this procedure.

Before you begin

Required role: Storage

Procedure

1. Before you expand a storage pool, verify the logical capacity, free capacity, ID, rebuild capacity policy, and tolerable number of drive failures of the intended storage pool.

If you use the CLI to specify a storage pool by name, check the name of the storage pool.

REST API: GET /v1/objects/pools

CLI: pool_list

2. Change the rebuild capacity policy as necessary to suit your operation.

For details, see *Editing storage pool settings*.

**Caution:**

If you change the rebuild capacity policy after expanding storage pools, sufficient rebuild capacity might not be secured. Therefore, if you want to change the rebuild capacity policy, do so before expanding storage pools.

3. Verify the status of the storage node containing the drive you want to add.

REST API: GET /v1/objects/storage-nodes

CLI: storage_node_list

If the status of the storage node is "Ready" or "RemovalFailed", go to step 3.

4. Verify drive IDs.

REST API: GET /v1/objects/drives

CLI: drive_list

5. Add drives to the storage pool.

Run either of the following commands with the ID of the intended storage pool and the IDs of the drives to be added specified.

If you use the CLI, you can specify a name instead of the ID of the storage pool.

REST API: POST /v1/objects/pools/<id>/actions/expand/invoke

CLI: pool_expand

Verify the job ID which is displayed after the command is run.

**Caution:**

To perform other operations, make sure that the storage pool has been expanded as described in step 7 before doing so.

6. Verify the state of the job by specifying the job ID.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

After running the command, if you receive a response indicating "Succeeded" as the state, the job is completed. Subsequently, capacity allocation to the storage controller, which is internal processing, is performed. Upon completion of these processes, the added capacity will be available for use.

The processing time (approximate) for capacity allocation to storage controller can be calculated by the following formula.

The processing is executed in each storage controller in parallel. Therefore, the total processing time is almost the same as the processing time for a single storage controller.

- For HPEC 4D+1P:
30 [sec] × (total physical capacity of drives per storage node / 121 [GiB])
- For HPEC 4D+2P:
30 [sec] × (total physical capacity of drives per storage node / 145 [GiB])
- For Mirroring Duplication:
30 [sec] × (total physical capacity of drives per storage node / 194 [GiB])



Caution:

If physical capacity of a storage node exceeds the capacity that can be allocated to a storage controller, the calculated approximate processing time might be longer. For details about the maximum capacity that can be allocated to a storage controller, contact customer support

7. Obtain event logs to make sure that the drives specified in step 5 can now be used as storage pool capacity.

Wait for the event log KARS16020-I to output as many storage nodes as possible. Until then, you may not be able to perform other operations.

If you get as many KARS16020-I event logs as the number of storage nodes, the storage pool has been expanded as specified.

When you expand a storage pool, if the KARS16046-I message is output to the event log, the logical capacity has not been configured. To configure the logical capacity, add necessary drives to the storage pool. For details, see *Capacity design (for HPEC 4D +1P)*, *Capacity design (for HPEC 4D+2P)*, or *Capacity design (for Mirroring)*.

8. Obtain information about the storage pool that has the newly-added drives.

REST API: GET /v1/objects/pools

CLI: pool_list

Verify that the logical capacity and free capacity of the storage pool has increased as a result of adding the drives.

9. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.



Tip:

A drive whose status is "Offline" is not incorporated into the storage pool. You can verify the drive status from drive information.

About Virtual Storage Software block GUI

Virtual Storage Software block GUI is a software product that allows for easy verification of the status and total configuration of a storage system and information about individual resources managed by Virtual Storage Software block with easy navigation and fast response. You can also perform various operations related to volumes, drives, storage nodes, and compute nodes, and operations for dump log files.

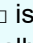
For user permissions, see the required role in *Before you begin* for each operation.

Requirements for web browsers to use the GUI

The following table lists the requirements for web browsers for using the GUI.

Web browser	OS
Microsoft Edge (Latest version of stable channel)	Windows platform supported by the web browser.
Google Chrome (Latest version of stable channel)	Windows or Linux platform supported by the web browser.

**Note:**

- Do not use "Refresh" of the web browser. Using it might cause an unexpected window to be displayed. If an unexpected window is displayed, close the web browser, and then retry login.
- After you have updated the storage software, re-open the GUI
- Some screen magnification settings might cause screen display abnormality. For example, buttons might become unclickable. In such a case, use the zoom function of the web browser to adjust the magnification.
- If GUI screens are displayed in multiple windows or if tabs that contain GUI screens are inactive, the refresh interval for charts and other screens pertaining to performance information might become longer. You might be able to mitigate the issue by performing the following operations:
 - Reduce the number of tabs and windows as much as possible.
 - Make sure that charts and other screens pertaining to performance information are displayed in the active tab of each window as much as possible.
- If "Waiting for available socket..." is displayed in the web browser, and the operation slows down, you may be able to resolve the issue by restarting the web browser.
- If "Loading..." is displayed and stays in the top left of the window, the network might not be operating normally. Verify whether the network is operating normally, close the tab once, and then open the tab again to use the GUI.
- If an icon is not displayed or displayed abnormally (for example, when  is displayed instead of the icon), the network might not be operating normally. Verify whether the network is operating normally, close the tab once, and then open the tab again to use the GUI.
- If the operation icon button is not displayed, verify that the user is assigned appropriate permissions. When the user has appropriate permissions, the network might not be operating normally. Verify whether the network is operating normally, close the tab once, and then open the tab again to use the GUI.
- When the number of resources to be displayed in the volumes window or other windows becomes large, it might take longer for the windows to be displayed.

How to set the browser to use the GUI

To use the GUI, set the browser as follows.

Browser to be used	Required setting	Setting Description
Microsoft Edge	Enabling JavaScript	For details about how to set the web browser, see Help of your web browser.
	Enabling Cookies	
	Enabling automatic download of multiple files	
	Disabling sleep status of inactive tabs	
Google Chrome	Enabling JavaScript	
	Enabling Cookies	
	Enabling automatic download of multiple files	

Logging in to and logging out of the GUI

This section describes how to log in to and log out from the GUI.

When you log in to the GUI, a new session is generated. When you log out of the GUI, the session ends.

Before you begin

- Required role: Security, Storage, Monitor, or Service

Procedure

1. In the web browser, enter the following URL:

https://<<IP-address-or-corresponding-FQDN>>[:443]/hsds/

- <<IP-address-or-corresponding-FQDN>>: Specify one of the following IP addresses set in the VSSB configuration file (SystemConfigurationFile.csv) or the corresponding FQDN :
 - Representative IP address (ClusterIpv4Address) of the storage cluster or the corresponding FQDN if representative IP address (ClusterIpv4Address) of the storage cluster is set
 - Any IP address (ControlNWIPv4) for the storage node control network or the corresponding FQDN if no representative IP address (ClusterIpv4Address) of the storage cluster is set



Tip:

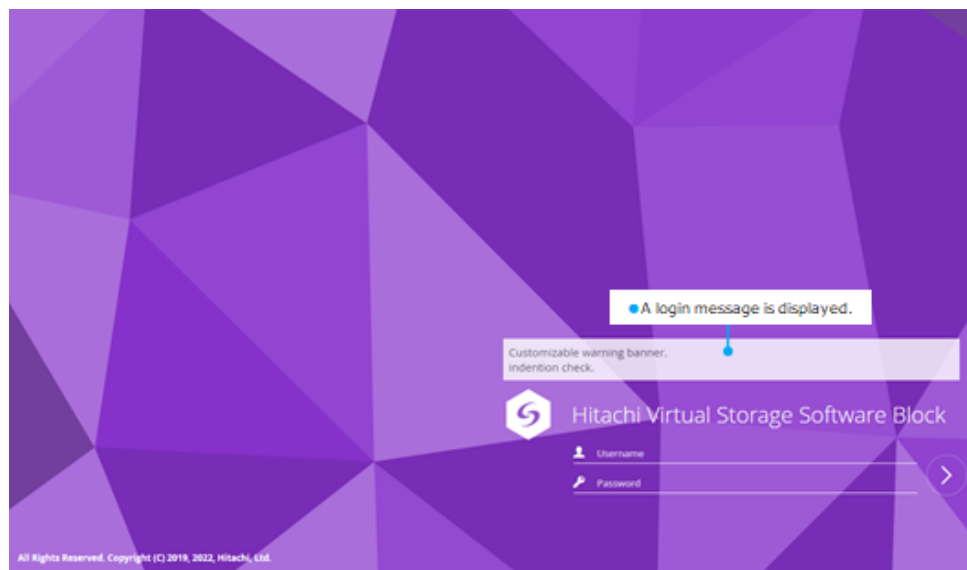
For details about server certificate verification, see *Client requirements for SSL/TLS communication*. When you specify the FQDN, the DNS that enables the IP address for the control network to be resolved by forward lookup from the FQDN of the storage node must be registered in the system from which the web browser is started.

- "443" is a port number. ":443" can be omitted.


2. When the Login window opens, enter your user ID and password to log in.

Logging in shows the dashboard window. The dashboard window displays total information about the Virtual Storage Software block storage system.

- Your account will be locked if your login is unsuccessful multiple consecutive times. Wait until your account is unlocked, and then retry the login.



3. To log out, click the user icon in the navigation bar and then select **Logout**.



Dashboard window

The dashboard window displays the following information about the Virtual Storage Software block storage system.

- Health status for each resource type
- Logical capacity and used capacity of the storage pool
- Used capacity reduction effect
- Data reduction effect
- System performance information
- Number of resources

Health status for each resource type:
 The health of all resource types is "Normal".
 There is at least one resource type whose status is "Warning" or "Error".

The dashboard displays the following information:

- Health Status:** Licenses, Fault Domains, Storage Nodes, Internode Ports, Drives, Compute Ports, Storage Cluster, Storage Pools, Volumes, Control Ports. All are shown as "Normal".
- Capacity:** A donut chart showing 64.5 TB used out of 151.5 TB total.
- Total Efficiency:** 10.00 : 1
- Data Reduction:** 1.00 : 1. A bar chart shows 100% reduction with 900.0 MB written and 900.0 MB consumed.
- Performance:** IOPS of 16.46 x. A line graph shows IOPS over time.
- Resource Counts:**
 - Fault Domains: 3
 - Storage Nodes: 9
 - Drives: 256
 - Compute Nodes: 50
 - Volumes: 100
 - Storage Pools: 1
 - Compute Ports: 18

Logical capacity and used capacity of the storage pool

Ratio of the data capacity before and after data reduction (in the case of the storage software version 01.12.xx.xx, the data reduction function is not supported and "1.00:1" is always displayed)

The sum of the number of Read I/Os per second and the number of write I/Os per second for all volumes in the storage cluster

Capacity efficiency by the volume creation function and snapshot function

Number of each resource. Clicking each resource displays the list window of the resource.

**Note:**

- Performance information is displayed only for users who have the Storage, Monitor, or Resource role.
- If volume migration is running, a value including the number of I/Os issued by the volume migration is displayed in the performance information.
- When a storage node has a high load, the chart of performance information might not be fully displayed, but it will be fully displayed after a while.
- Total Efficiency indicates the used capacity reduction effect achieved by the volume creation function and snapshot function.

This is an average ratio, where the ratio of usedCapacity (used capacity of the storage pool) to totalVolumeCapacity (total capacity of created volumes)* is calculated for each storage controller, and then weighted according to the total capacity of the created volumes for each storage controller. The larger the total capacity of created volumes of a storage controller is, the more the capacity reduction effect of that storage controller are reflected to the Total Efficiency value.

* You can get the value of used capacity of the storage pool (usedCapacity) and the value of total capacity of created volumes (totalVolumeCapacity) by obtaining information about the storage pool. You can verify the used capacity of the storage pool (usedCapacity) in USED of USED/TOTAL in Summary on the Storage Pool window, and the total capacity of created volumes (totalVolumeCapacity) in TOTAL CAPACITY of VOLUME CAPACITY INFORMATION on the Storage Pool window.

When the ratio is larger than 99999.99:1, it is displayed as ">99999.99:1".

- When snapshot volumes are used, the used storage pool capacity can be reduced by storing only the differential data between P-VOL and S-VOL. Therefore, the value of totalEfficiency is larger than that when snapshot volumes are not included, even if the total capacity of created volumes is the same.
- When storage pool expansion processing is being performed and KARS16017-I, KARS16020-I, KARS16022-I, or KARS16081-I is output, the value of totalEfficiency might become larger even when snapshot operations are not being performed. Note that snapshot operations refer to preparation (for obtaining snapshots) and obtaining, deleting, and restoring snapshots.
- Volumes whose type is "ExternalMigrationOrigin" are not included in the number of volumes.

Health status

The health status indicates the status ("Normal" or "Alerting") of each resource. For example, "Normal" displayed for resource type "Storage Nodes" in the health status display area indicates that the STATUS SUMMARY of all storage nodes is "Normal". If "Alerting" is displayed for resource type "Nodes", there is at least one storage node whose STATUS SUMMARY is "Warning" or "Error".


If the health status could not be obtained, the following is displayed in the health status display area.



Unknown The health status could not be obtained.

Also, a health status summary is displayed in the navigation bar. Each status description is as follows:

- "Normal": The health status of all resource types is "Normal".
- "Alerting": There is at least one resource type whose status is "Alerting".
- "Unknown": The health status could not be obtained.



Clicking the health status summary opens the dashboard window.



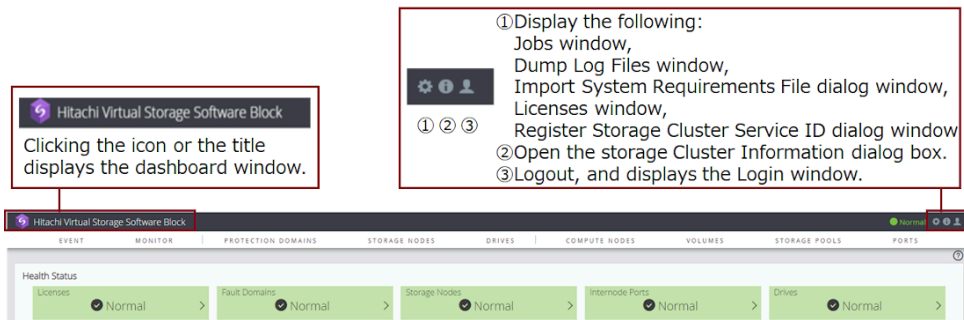
Caution:

The health status of the storage cluster and the health status summary displayed in the navigation bar do not indicate the boot status of the storage cluster.

For details about how to start a storage cluster, contact customer support

Navigation bar


You can perform the following operations from the navigation bar at the top of the GUI window.

**Caution:**

The ID displayed when clicking "Register Storage Cluster Service ID" from the gear icon is the one set by service personnel or maintenance personnel. This function is unavailable for users.

**Note:**

- Clicking the gear icon displays the following menu items:
 - Jobs: Displayed only to users with the Storage, Service, or Security role.
 - Dump Log Files: Displayed only to users with the Service role.
 - Import System Requirements File: Displayed only to users with the Service role.
 - Licenses: Displayed only to users with the Storage, Monitor, or Resource role.
 - Register Storage Cluster Service ID: Displayed only to users with the Storage or Service role.
- You can verify the model name of Virtual Storage Software block by clicking the Storage Cluster Information icon.
 - MODEL NAME = "VSSB": Virtual machine model
 - MODEL NAME = "VSSBB1": Bare metal model



Clicking each item displays the following information.

- **EVENT:** Displays the event logs.
- **MONITOR:** Displays performance information in a chart. "MONITOR" is displayed only for users who have the Storage, Monitor, or Resource role.

- PROTECTION DOMAIN: Displays information about the protection domain and fault domain.
- STORAGE NODES: Displays information about the storage nodes.
- DRIVES: Displays information about the drives.
- COMPUTE NODES: Displays information about the compute nodes.
- VOLUMES: Displays information about the volumes.
- STORAGE POOLS: Displays information about the pools.
- PORTS: Displays information about the compute ports, internode ports, and control ports.



Note:

Do not use "Back", "Forward", or "Refresh" of the web browser. Using any of these functions might cause an unexpected window to be displayed. If an unexpected window is displayed, close the web browser, and then retry login.

Displaying performance information about storage systems

The performance information for each resource collected by Virtual Storage Software block can be displayed on the System Monitor window. Performance information can be displayed by specifying a range for the last 2 hours.

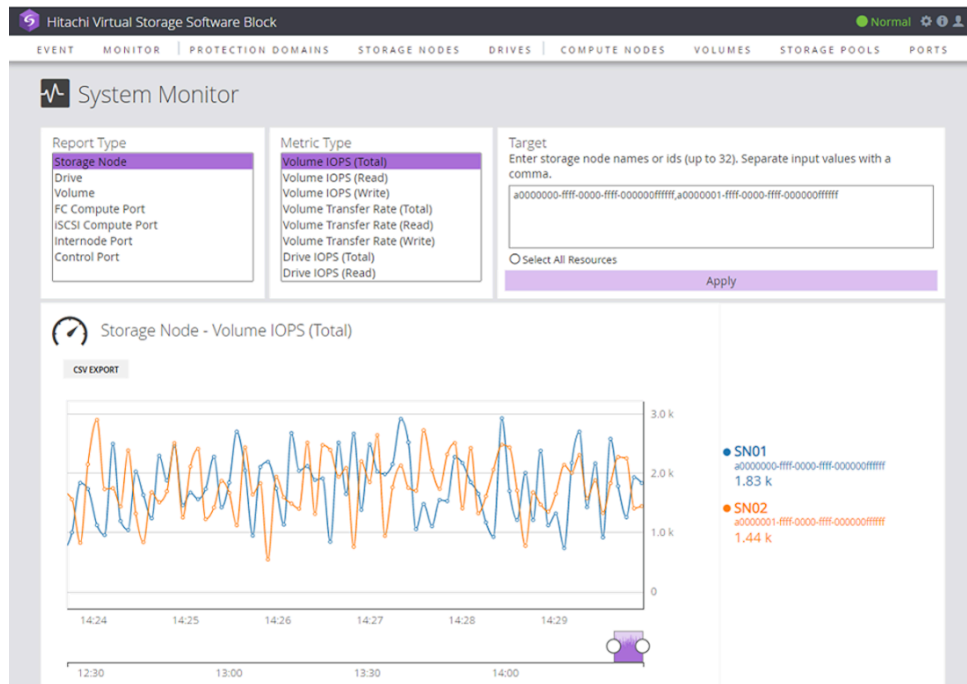
Before you begin

- Required role: Storage, Monitor, or Resource role

Procedure

1. Click **MONITOR** on the navigation bar.

The System Monitor window opens.



2. Select Report Type.
3. Select Metric Type.
4. In Target, specify resource IDs or names (up to 32, delimited by commas). You can specify a combination of IDs and names.

If you selected Storage Node for Report Type, clicking the check box to the left of [Select All resources] selects all the resources.

5. Click **Apply**.

The performance information is displayed.

If multiple resource IDs are specified in step 4, a chart for the specified multiple resources will be displayed with a legend.

**Note:**


- If a processing that involves I/O operations, such as volume migration and drive data relocation, is running, values including performance information for the I/Os are displayed in the performance information.
- Click **CSV EXPORT** to save the displayed performance information as a CSV file.

- When a storage node has a high load, the chart of performance information might not be fully displayed, but it will be fully displayed after a while.


Monitor link

The monitor link icon is displayed in the list window or the detailed information window of nodes, drives, volumes, or ports. It is also displayed in the list of compute ports, internode ports, or control ports, which is related information about the node detailed information window.

However, the Monitor link icon is displayed only for users who have the Storage, Monitor, or Resource role.

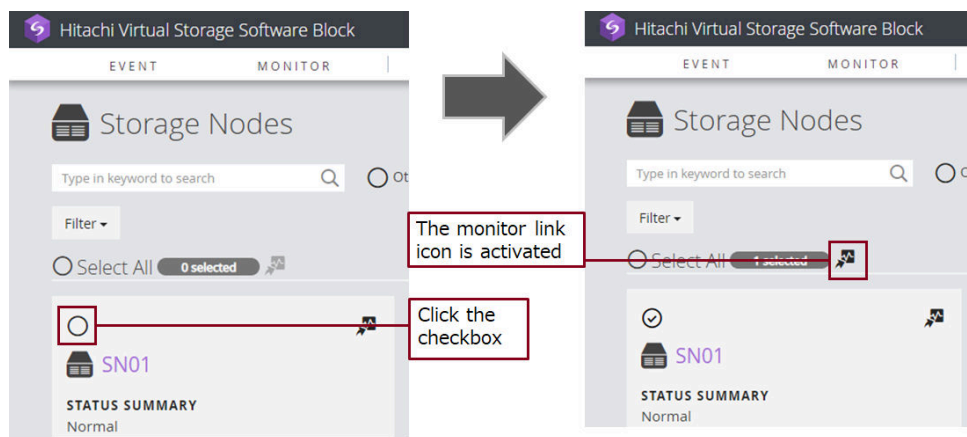


Clicking the activated monitor link icon selects the resource ID having this icon and displays the System Monitor window with the report type and resource ID set.

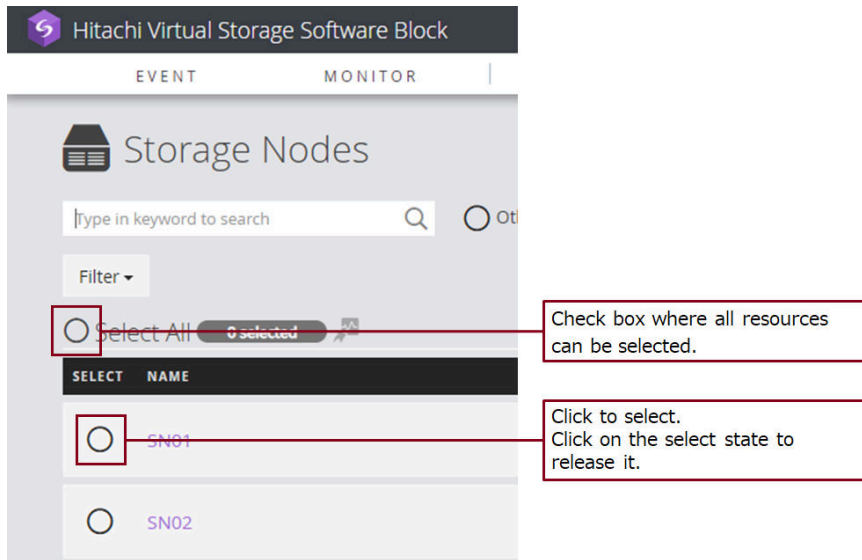


You can select a maximum of 32 resources. To select multiple resources, click the checkbox in each resource. The Monitor link icon to the right of the [Select All] will be activated.

Clicking on the Monitor link icon opens the System Monitor window with a chart of the specified multiple resources displayed with a legend.



To select all resources, click the checkbox to the left of [Select All]. If 33 or more resources are selected, the Monitor link to the right of [Select All] is not activated because up to 32 resources can be selected. Clear the selection by clicking the checkbox in each resource and set the number of selections to 32 or less.



- Be careful when clicking the checkbox to the left of [Select All]. Resources will be selected over multiple pages.

Performing maintenance

The Storage Node list window, Storage Node detailed information window, Drive list window, and Drive detailed information window have a Maintenance link icon.

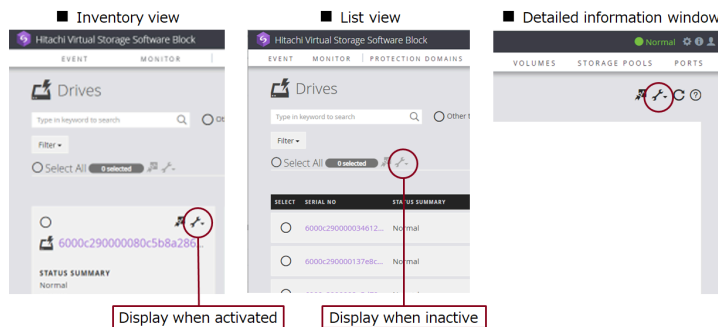
The Maintenance link icon is displayed only to users that can perform maintenance. For details about user's privilege, see "Required role" as a prerequisite for each procedure.



The following table shows the menus that are displayed when you click an active Maintenance icon and their respective transition-destination windows (displayed when you click a menu).

Window	Menu to be displayed	Transition-destination window
Drives list window Drive detailed information window	Remove Drives	Drive removal dialog. This dialog is displayed with SERIAL NO of the drive to be removed set.
	Add Drives	Drive addition dialog. This dialog is displayed with SERIAL NO of the drive to be added set.


Window	Menu to be displayed	Transition-destination window
	Turn On/Off Locator LEDs	Locator LED on/off dialog. This dialog is displayed with SERIAL NO of the drive (whose Locator LED is to be turned on/off) set.
Storage Node list window Storage Node detailed information window	Maintenance Recovery	Storage node maintenance recovery dialog. This dialog is displayed with the name of the storage node subject to maintenance recovery set.
	Maintenance Blockade	Storage node maintenance blocking dialog. This dialog is displayed with the name of the storage node subject to maintenance blocking set.



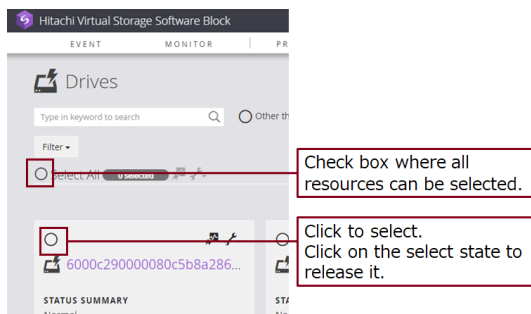
How many resources can be selected varies depending on the displayed menu. To select multiple resources, click the check box in each resource. Doing this activates the Maintenance icon to the right of the Monitor link icon.

Clicking an activated Maintenance icon shows menus from which you can open transition-destination windows that display the specified number of resources.

Check the maximum number of selectable items for each menu to be displayed by operating each menu.



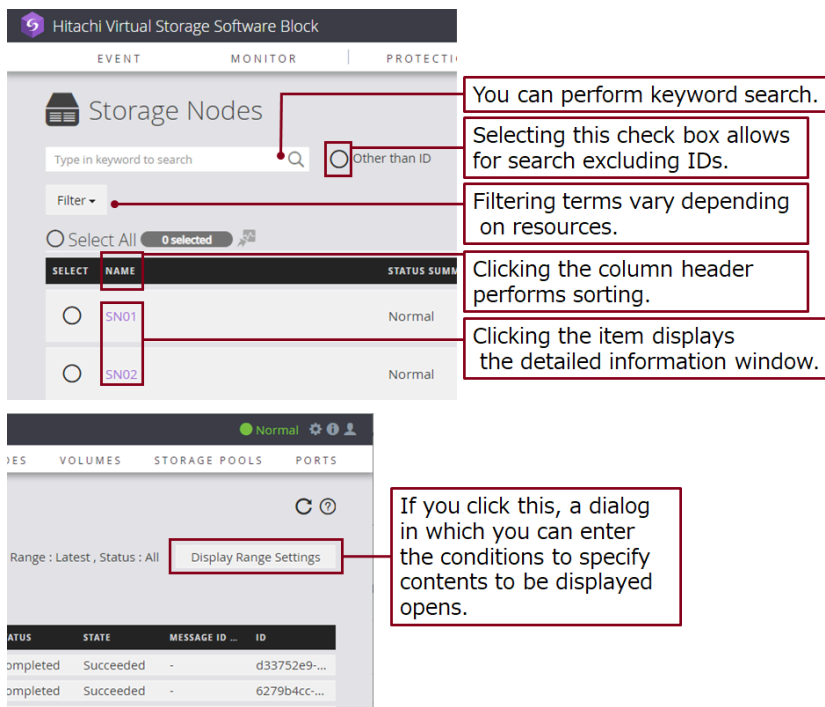
To select all resources, click the check box to the left of [Select All]. If you select more resources than are allowed, the Maintenance icon to the right of the Monitor link icon is deactivated. Clear the selection by clicking the check box in each resource, and then set the number of resource selections to the maximum or less.



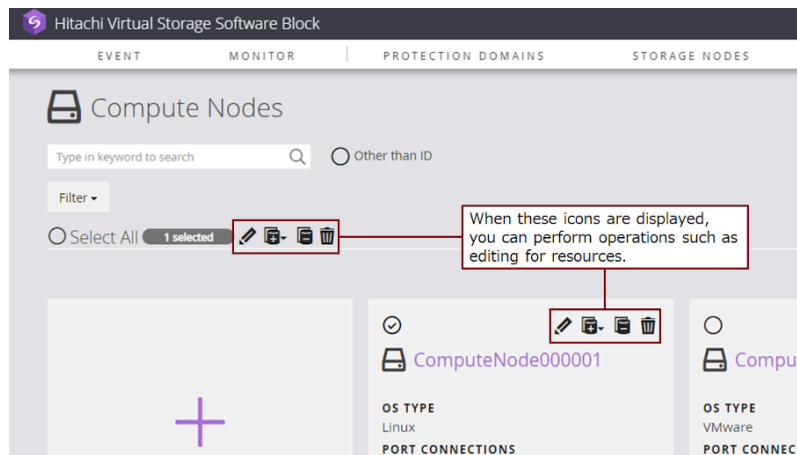
Viewing the List window

Selecting an item in the list view opens a window showing detailed information.

You can perform keyword search, filtering, and sort in the list window.

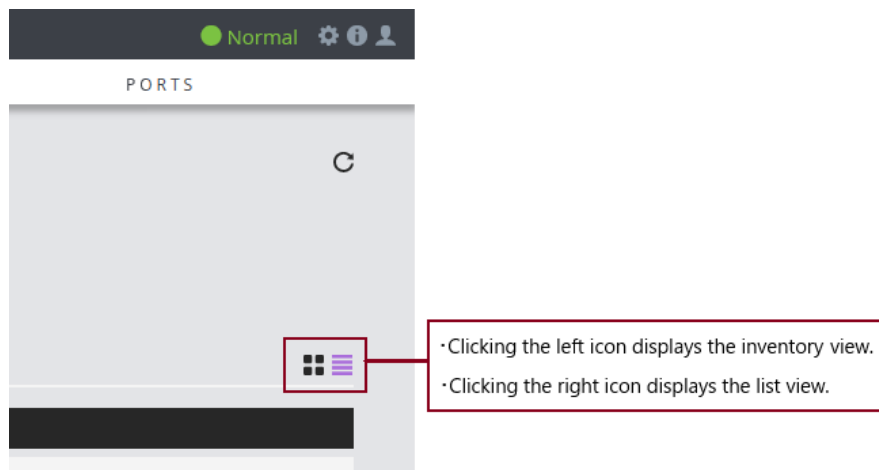


Depending on the resource, the following icons are shown so that you can perform editing and other operations.



A set of icons shown to the right of [Select All] are used for operations on the selected resources. A set of icons shown for each resource are for operations on that resource.

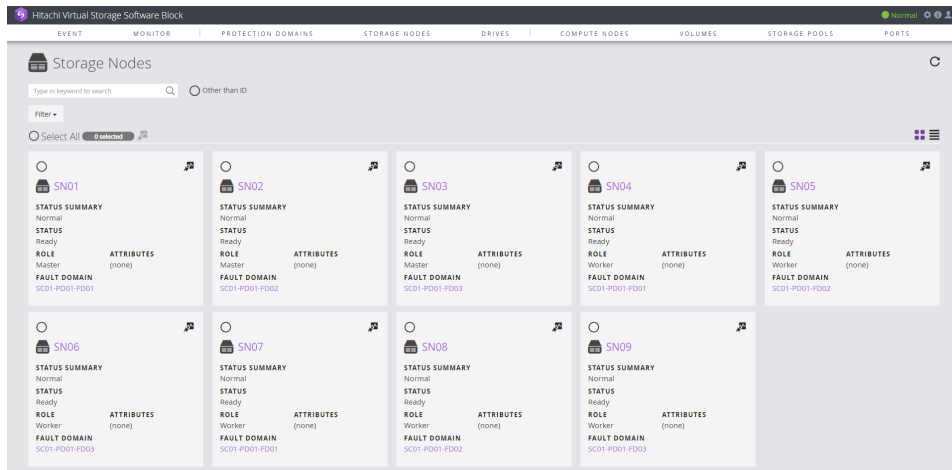
Clicking the switching icon toggles between Inventory view and List view.



■ List view

SELECT	NAME	STATUS SUMMARY	STATUS	FAULT DOMAIN	ROLE	ATTRIBUTES	ID
<input type="radio"/>	SN01	Normal	Ready	SC01-PD01-FD01	Master	(none)	a0000001-fff-0000-fff-00000000ffff
<input type="radio"/>	SN02	Normal	Ready	SC01-PD01-FD02	Master	(none)	a0000001-fff-0000-fff-00000000ffff
<input type="radio"/>	SN03	Normal	Ready	SC01-PD01-FD03	Master	(none)	a0000002-fff-0000-fff-00000000ffff
<input type="radio"/>	SN04	Normal	Ready	SC01-PD01-FD01	Worker	(none)	a0000003-fff-0000-fff-00000000ffff
<input type="radio"/>	SN05	Normal	Ready	SC01-PD01-FD02	Worker	(none)	a0000004-fff-0000-fff-00000000ffff
<input type="radio"/>	SN06	Normal	Ready	SC01-PD01-FD03	Worker	(none)	a0000005-fff-0000-fff-00000000ffff
<input type="radio"/>	SN07	Normal	Ready	SC01-PD01-FD01	Worker	(none)	a0000006-fff-0000-fff-00000000ffff
<input type="radio"/>	SN08	Normal	Ready	SC01-PD01-FD02	Worker	(none)	a0000007-fff-0000-fff-00000000ffff
<input type="radio"/>	SN09	Normal	Ready	SC01-PD01-FD03	Worker	(none)	a0000008-fff-0000-fff-00000000ffff

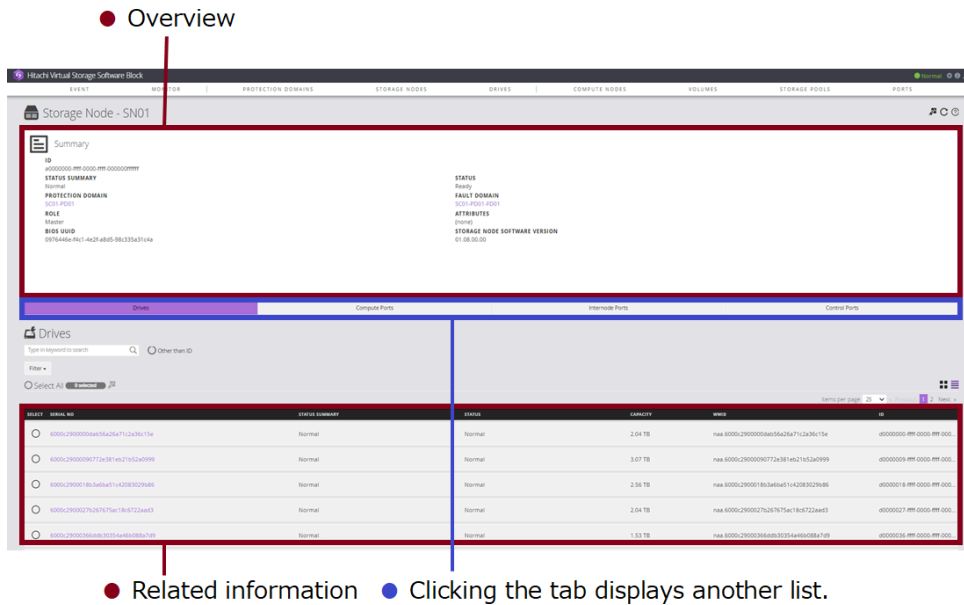
■ Inventory view



Viewing the detailed information window

Clicking an item in the list window opens the detailed information window related to that item. (An overview is displayed at the top, and related information is displayed at the bottom.)

The related information area allows you to perform keyword search and filtering. Selecting an item in the displayed list opens its corresponding detailed information window. In some windows, clicking the tab toggles between the lists of related information.




Depending on the resource, the following icons are shown so that you can perform editing and other operations.




Refresh display

In the window with the refresh icon, clicking the refresh icon refreshes the displayed information.

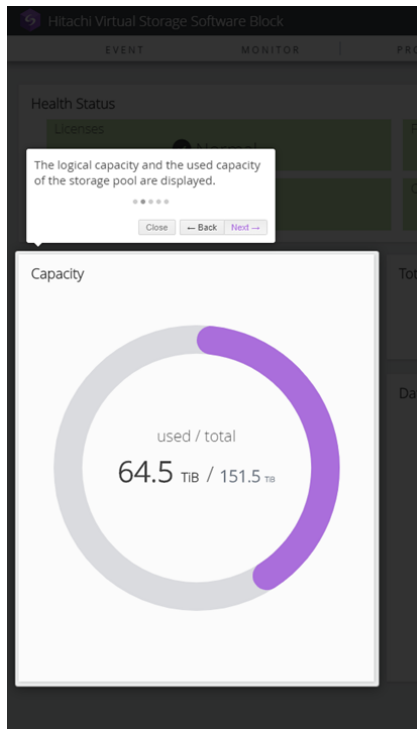


Help display

In the window with the help icon, clicking the help icon shows the help information.



■ Help display example



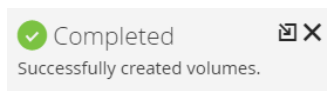

Pop-up message display

You can check the operation result on the pop-up message displayed at the right bottom of the window.

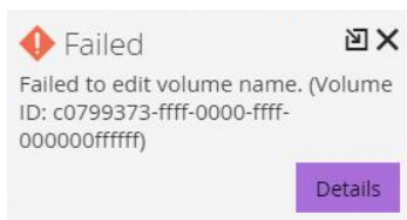
Clicking the **x** icon closes the pop-up message.

If the operation fails, you can check the details by clicking **Details** on the pop-up message.

■ Example display of pop-up messages when operations succeed



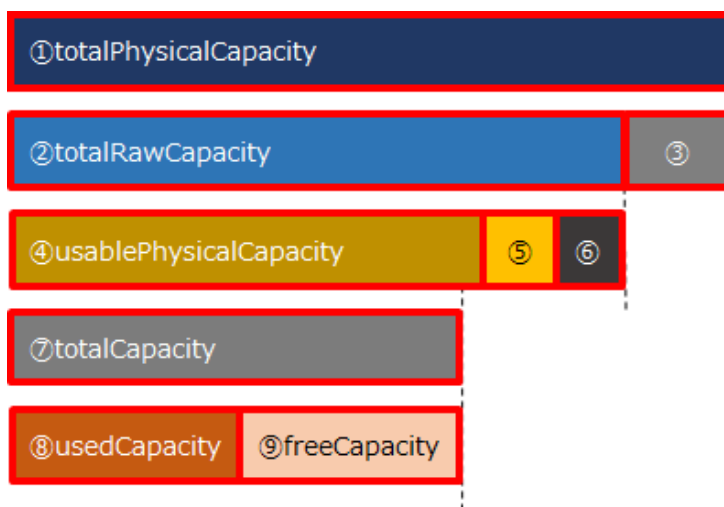
■ Example display of a pop-up message when an operation fails



Chapter 2: Managing storage pools

Capacity information about storage pools

Among the information obtained in *Obtaining a list of information about storage pools* or *Obtaining information about individual storage pools*, you can obtain the following capacity information ((1) to (13)). You can also verify these information through GUI.



- (1) totalPhysicalCapacity: Total capacity of the storage pool.
- (2) totalRawCapacity: The amount of space that is now available by adding drives to the storage pool.
- (3) metaDataPhysicalCapacity: Capacity for control information in totalPhysicalCapacity.
- (4) usablePhysicalCapacity: Capacity that can be used as logical capacity in totalRawCapacity.
- (5) reservedPhysicalCapacity: Capacity not used for logical capacity in totalRawCapacity.
- (6) blockedPhysicalCapacity: Blocked capacity in totalRawCapacity.
- (7) totalCapacity: Capacity that can be used in the storage pool.
- (8) usedCapacity: Capacity of storage that actually stores user data and metadata related to user data.
- (9) freeCapacity: Capacity that is not used in the storage pool.

You can verify the total capacity of created volumes from the capacity information about storage pools.

⑩totalVolumeCapacity

⑪provisionedVolumeCapacity

⑫otherVolumeCapacity

⑬

Name	Description	Type of the included volumes and snapshot attributes
(10) totalVolumeCapacity	Total capacity of created volumes (Total of created provisioned volume capacity, other volume capacity, and temporary volume capacity)	<ul style="list-style-type: none"> ▪ volumeType: Normal ▪ volumeType: Snapshot ▪ volumeType: MigrationDestination
(11) provisionedVolumeCapacity	Capacity of provisioned volumes that have been created (Total capacity of normal volumes and snapshot volumes (P-VOLs))	<ul style="list-style-type: none"> ▪ volumeType: Normal ▪ volumeType: Snapshot and snapshotAttribute: P-VOL
(12) otherVolumeCapacity	Capacity of other volumes that have been created (Total capacity of snapshot volumes (S-VOLs and P/S-VOLs))	<ul style="list-style-type: none"> ▪ volumeType: Snapshot and snapshotAttribute: P/S-VOL ▪ volumeType: Snapshot and snapshotAttribute: S-VOL
(13) temporaryVolumeCapacity	Temporary capacity of volumes that have been created (Virtual machine) Total capacity of volumes created temporarily by Data migration and Capacity balance (Bare metal) Total capacity of volumes created temporarily by Capacity balance	<ul style="list-style-type: none"> ▪ volumeType: MigrationDestination



Note:

- The used capacity of snapshot volumes might be larger than the capacity of created volumes.
- Each capacity listed in the preceding table does not include control information of the volume. Capacity including the control information can be calculated by referring to *Maximum storage pool capacity consumed by a volume*.



Caution:

Even if the totalVolumeCapacity is smaller than the totalCapacity (logical capacity of the storage pool), the capacity of the storage controller might become full if the used capacity of each storage controller is uneven.

When the capacity of storage controller is full, the following failures might occur:

- Operations other than deleting the normal and snapshot volumes fail
- I/O to the normal and snapshot volumes fails
- Data in the snapshot volume is lost

Obtaining a list of information about storage pools (CLI or REST API)

The following information can be obtained.

- id: IDs (uuid) of storage pools
- name: Names of storage pools
- protectionDomainId: IDs (uuid) of protection domains containing storage pools
- statusSummary: Summary of statuses of storage pools
- status: Status of each storage node
- totalCapacity: Logical capacity of the storage pool
- totalRawCapacity: Valid physical capacity in the total physical capacity of the storage pool
- usedCapacity: Used capacity of the storage pool
- freeCapacity: Free capacity of the storage pool
- totalPhysicalCapacity: Total physical capacity of the storage pool
- metaDataPhysicalCapacity: Capacity for control information in the total physical capacity of the storage pool

- reservedPhysicalCapacity: Area that is reserved internally for other than control information and that cannot be used for redundancy in the valid physical capacity of the storage pool
- usablePhysicalCapacity: Capacity that can be used as logical capacity in the valid physical capacity of the storage pool
- blockedPhysicalCapacity: Blocked capacity in the valid physical capacity of the storage pool
- usedCapacityRate: Usage rate
- maximumReserveRate: Maximum reserve rate
- thresholdWarning: Warning threshold
- thresholdDepletion: Depletion threshold
- thresholdStorageControllerDepletion: Depletion threshold of the storage controller that manages the storage pool
- savingEffects: Effect of the data reduction function on the storage pool
- numberOfVolumes: Number of volumes in each storage pool
- redundantPolicy: User data protection method
- redundantType: User data protection type
- dataRedundancy: Redundancy of user data
- storageControllerCapacitiesGeneralStatus: Summary information about the capacity status (capacityStatus) of all the storage controllers managing the storage pool
- totalVolumeCapacity: Total capacity of volumes that have been created on the applicable storage controller
- provisionedVolumeCapacity: Capacity of provisioned volumes that have been created on the applicable storage controller
- otherVolumeCapacity: Capacity of other volumes that have been created on the applicable storage controller
- temporaryVolumeCapacity: Temporary capacity of volumes that have been created on the applicable storage controller
- rebuildCapacityPolicy: Rebuild capacity policy
- rebuildCapacityResourceSetting: Resource type and the number of resources for the rebuild capacity
 - numberOfTolerableDriveFailures: Tolerable number of drive failures

- rebuildCapacityStatus: Rebuild capacity allocation status
- rebuildableResources: Rebuildable resources
 - numberOfDrives: Rebuildable resources



Note:

totalEfficiency indicates the used capacity reduction effect achieved by the volume creation function and snapshot function.

This is an average ratio (unit: %), where the ratio of usedCapacity (used capacity of the storage pool) to totalVolumeCapacity (total capacity of created volumes) is calculated for each storage controller, and then weighted according to the total capacity of the created volumes for each storage controller. The larger the total capacity of created volumes of a storage controller is, the more the capacity reduction effect of that storage controller are reflected to the totalEfficiency value.

For example, if the total capacity of created volumes is 1,000 for each of the storage controllers and the used capacity of the storage pool is 50, totalEfficiency of 2,000 (unit: %) is output by using a REST API or CLI.

In GUI, the value corresponding to totalEfficiency is displayed as "20.00:1" in the dashboard window (Total Efficiency) and the Storage Pool window (TOTAL EFFICIENCY). Also, in the dashboard window of the GUI, if the ratio is greater than 9,999,999:100, ">99999.99:1" is displayed.

When snapshot volumes are used, the used storage pool capacity can be reduced by storing only the differential data between P-VOL and S-VOL. Therefore, the value of totalEfficiency is larger than that when snapshot volumes are not included, even if the total capacity of created volumes is the same.

When storage pool expansion processing is being performed and KARS16017-I, KARS16020-I, KARS16022-I, or KARS16081-I is output, the value of totalEfficiency might become larger even when snapshot operations are not being performed. Note that snapshot operations refer to preparation (for obtaining snapshots) and obtaining, deleting, and restoring snapshots.

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Obtain a list of information about storage pools.

REST API: GET /v1/objects/pools

CLI: pool_list

Obtaining information about individual storage pools (CLI or REST API)

The following information can be obtained.

- id: ID (uuid) of the intended storage pool
- name: Name of the intended storage pool
- protectionDomainId: ID (uuid) of the protection domain containing the storage pool
- statusSummary: Status summary of the intended storage pool
- status: Status of the intended storage pool
- totalCapacity: Logical capacity of the storage pool
- totalRawCapacity: Valid physical capacity in the total physical capacity of the storage pool
- usedCapacity: Used capacity of the storage pool
- freeCapacity: Free capacity of the storage pool
- totalPhysicalCapacity: Total physical capacity of the storage pool
- metaDataPhysicalCapacity: Capacity for control information in the total physical capacity of the storage pool
- reservedPhysicalCapacity: Area that is reserved internally for other than control information and that cannot be used for redundancy in the valid physical capacity of the storage pool
- usablePhysicalCapacity: Capacity that can be used as logical capacity in the valid physical capacity of the storage pool
- blockedPhysicalCapacity: Blocked capacity in the valid physical capacity of the storage pool
- usedCapacityRate: Usage rate
- maximumReserveRate: Maximum reserve rate
- thresholdWarning: Warning threshold value
- thresholdDepletion: Depletion threshold value
- thresholdStorageControllerDepletion: Depletion threshold of the storage controller that manages the storage pool
- savingEffects: Effect of the data reduction function on the storage pool

- numberOfVolumes: Number of volumes in the storage pool
- redundantPolicy: User data protection method
- redundantType: User data protection type
- dataRedundancy: Redundancy of user data
- storageControllerCapacitiesGeneralStatus: Summary information about the capacity status (capacityStatus) of all the storage controllers managing the storage pool
- totalVolumeCapacity: Total capacity of created volumes on the applicable storage controller
- provisionedVolumeCapacity: Capacity of provisioned volumes that have been created on the applicable storage controller
- otherVolumeCapacity: Capacity of other volumes that have been created on the applicable storage controller
- temporaryVolumeCapacity: Temporary capacity of volumes that have been created on the applicable storage controller
- rebuildCapacityPolicy: Rebuild capacity policy
- rebuildCapacityResourceSetting: Resource type and the number of resources for the rebuild capacity
 - numberOfTolerableDriveFailures: Tolerable number of drive failures
- rebuildCapacityStatus: Rebuild capacity allocation status
- rebuildableResources: Rebuildable resources
 - numberOfDrives: Rebuildable resources

**Note:**

totalEfficiency indicates the used capacity reduction effect achieved by the volume creation function and snapshot function.

This is an average ratio (unit: %), where the ratio of usedCapacity (used capacity of the storage pool) to totalVolumeCapacity (total capacity of created volumes) is calculated for each storage controller, and then weighted according to the total capacity of the created volumes for each storage controller. The larger the total capacity of created volumes of a storage controller is, the more the capacity reduction effect of that storage controller are reflected to the totalEfficiency value.

For example, if the total capacity of created volumes is 1,000 for each of the storage controllers and the used capacity of the storage pool is 50, totalEfficiency of 2,000 (unit: %) is output by using a REST API or CLI.

In GUI, the value corresponding to totalEfficiency is displayed as "20.00:1" in the dashboard window (Total Efficiency) and the Storage Pool window (TOTAL EFFICIENCY). Also, in the dashboard window of the GUI, if the ratio is greater than 9,999,999:100, ">99999.99:1" is displayed.

When snapshot volumes are used, the used storage pool capacity can be reduced by storing only the differential data between P-VOL and S-VOL. Therefore, the value of totalEfficiency is larger than that when snapshot volumes are not included, even if the total capacity of created volumes is the same.

When storage pool expansion processing is being performed and KARS16017-I, KARS16020-I, KARS16022-I, or KARS16081-I is output, the value of totalEfficiency might become larger even when snapshot operations are not being performed. Note that snapshot operations refer to preparation (for obtaining snapshots) and obtaining, deleting, and restoring snapshots.

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Verify the ID of the intended storage pool.

If you use the CLI to specify a storage pool by name, check the name of the storage pool.

REST API: GET /v1/objects/pools

CLI: pool_list

2. Obtain information about the intended storage pool.

Run either of the following commands with the storage pool ID specified.

If you use the CLI, you can specify a name instead of the ID of the storage pool.

REST API: GET /v1/objects/pools/<id>

CLI: pool_show

Obtaining a list of capacity information about volumes (CLI or REST API)

The following information can be obtained.

- id: Volume IDs (uuid)
- capacityUsage: Consumed amount

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Obtain a list of capacity information about volumes.
REST API: GET /v1/objects/performances/volume-capacities
CLI: volume_capacity_performance_list

Obtaining capacity information about individual volumes (CLI or REST API)

The following capacity information is obtained in low-resolution for the volume with the specified ID.

- id: Volume IDs (uuid)
- capacityUsage: Consumed amount

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Verify the ID of the volume for which you want to obtain capacity information.
If you use the CLI to specify a volume by name, check the name of the volume.
REST API: GET /v1/objects/volumes
CLI: volume_list
2. Obtain capacity information.
Run either of the following commands with the volume ID specified.
If you use the CLI, you can specify a name instead of the ID of the volume.
REST API: GET /v1/objects/performances/volume-capacities/<id>
CLI: volume_capacity_performance_show

Obtaining a list of performance information about volumes (CLI or REST API)

The following information can be obtained.

- id: Volume IDs (uuid)
- readIOPS: Number of read operations per second
- writeIOPS: Number of write operations per second
- readTransferRate: Read transfer amount per second
- writeTransferRate: Write transfer amount per second
- readResponseTime: Average time required to respond to the read command of the volume
- writeResponseTime: Average time required to respond to the write command of the volume

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Verify the IDs of the volumes for which you want to obtain performance information.

If you use the CLI to specify a volume by name, check the name of the volume.

REST API: GET /v1/objects/volumes

CLI: volume_list

2. Obtain a list of performance information.

To obtain high-resolution performance information, run the command with the volume IDs specified in the query parameter. You can specify a maximum of 32 volume IDs.

If you use the CLI, you can specify a name instead of the ID of the volume.

REST API (low-resolution): GET /v1/objects/performances/volumes

REST API (high-resolution): GET /v1/objects/detail-performances/volumes

CLI (low-resolution): volume_performance_list

CLI (high-resolution): volume_detail_performance_list

Obtaining performance information about individual volumes (CLI or REST API)

The following information can be obtained for the volume with the ID specified.

- id: Volume ID (uuid)
- readIOPS: Number of read operations per second
- writeIOPS: Number of write operations per second
- readTransferRate: Read transfer amount per second
- writeTransferRate: Write transfer amount per second
- readResponseTime: Average time required to respond to the read command of the volume
- writeResponseTime: Average time required to respond to the write command of the volume

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Verify the ID of the volume for which you want to obtain performance information.

If you use the CLI to specify a volume by name, check the name of the volume.

REST API: GET /v1/objects/volumes

CLI: volume_list

2. Obtain performance information.

Run one of the following commands with the volume ID specified.

If you use the CLI, you can specify a name instead of the ID of the volume.

REST API (low-resolution): GET /v1/objects/performances/volumes/<id>

REST API (high-resolution): GET /v1/objects/detail-performances/volumes/<id>

CLI (low-resolution): volume_performance_show

CLI (high-resolution): volume_detail_performance_show

Chapter 3: Managing your license

License management overview

To use Virtual Storage Software block, you need a license. After Virtual Storage Software block is set up, it is installed with a trial license (valid for 60 days). You need to register a license for continued use after the trial license expires.

Requesting and registering a license

To request and register a license, contact your sales account representative to purchase a license key. To purchase a license key, you must provide the storage cluster ID and type of license you require. Depending on the type of the license, you also must notify the licensed capacity and period.

You can obtain the ID of the storage cluster by running the command described in *Obtaining information about storage clusters* or clicking the Storage Cluster Information in the GUI.



Caution:

After reconfiguring a storage cluster, the license key issued before configuring a storage cluster cannot be used for a new storage cluster. In this case, contact your sales account representative to re-issue a license key that corresponds to the new storage cluster ID.

License types

The following license types are available for different capacities and periods.

Type	Expiration time	Capacity limit	Description
Perpetual *	No	Yes	The specified capacity (registrable storage pool capacity) can be used indefinitely. If the type of the existing license is "Perpetual", the type of license you can register newly is "Perpetual" or "Subscription". If you register a "Subscription" license later, your license returns to a "Perpetual" license when the "Subscription" license expires or when you delete the "Subscription" license.

Type	Expiration time	Capacity limit	Description
Subscription	Yes	Yes	Determine the period during which the specified capacity (registrable storage pool capacity) can be used. If the type of the existing license is "Subscription", the type of license you can register newly is "Perpetual" or "Subscription".
* Not currently sold			

Program product

There are two types of licenses: the base license, which enables the basic functions of Virtual Storage Software block, and the program product licenses, which enable the corresponding value-added functions.

Program product name	Type	Description
Hitachi SDS Core Base	Base license	This license manages the validity period for using Virtual Storage Software block and the size of available space in the storage pool. When this product is installed, the trial license, which enables unlimited capacity use for 60 days, is applied.

Valid period of a license

The valid period of a license starts when Virtual Storage Software block is set up or a license is registered. The period ends at 23:59 of the last day. For example, if you register a Subscription license which is valid for one year at 13:00 on August 5, 2020, the license is valid until 23:59 on August 5, 2021.

Licensed capacity

The capacity enabled by a Perpetual or Subscription license is "totalCapacity" (logical capacity of the storage pool) obtained as described in *Obtaining a list of information about storage pools* or *Obtaining information about individual storage pools*.

status and statusSummary of a license

"status" of a license shows whether the license is valid. "statusSummary" shows whether an error occurred for the license. You can obtain status and statusSummary of a license through an operation as described in *Obtaining a list of license information* or *Obtaining information about individual licenses*.

Status	Description	statusSummary
Active	<p>The license is invalid.</p> <ul style="list-style-type: none"> ▪ The license becomes active upon registration. ▪ If another type of license with an overlapping valid period is registered when the existing license is active, the existing license is overwritten during the overlapping period. 	Normal
Warning	<p>Although the license is valid, a warning threshold is exceeded.</p> <p>In the case of a Perpetual license:</p> <ul style="list-style-type: none"> • Warning threshold \leq Consumed capacity < Licensed capacity <p>In the case of a Subscription license (one of the following cases):</p> <ul style="list-style-type: none"> • Warning threshold \leq Consumed capacity < Licensed capacity • Warning threshold \leq Use period < Licensed period <p>You can determine which one is the cause by verifying the "cause" obtained as described in <i>Obtaining a list of license information</i> or <i>Obtaining information about individual licenses</i>.</p> <p>In a warning status, a warning event log is output. The status becomes active when the user changes the warning threshold or removes the cause of the warning threshold being exceeded.</p>	Warning
GracePeriod	<p>The license has expired or the licensed capacity is exceeded.</p> <p>In the case of a Perpetual license:</p> <ul style="list-style-type: none"> • Licensed capacity \leq Consumed capacity <p>In the case of a Subscription license (one of the following cases):</p> <ul style="list-style-type: none"> • Licensed capacity \leq Consumed capacity • Licensed period \leq Use period <p>You can determine which one is the cause by verifying the "cause" obtained as described in <i>Obtaining a list of license information</i> or <i>Obtaining information about individual licenses</i>.</p> <p>In the GracePeriod, a warning event log is output.</p>	Error

Status	Description	statusSummary
	If 30 days pass in the GracePeriod state, the license becomes invalid. If a license with an additional period or capacity is registered during the GracePeriod, the license becomes active.	
Invalid	The license is invalid. An error event log is notified. I/O operations are possible even when the license is invalid. The allowed management operations are license management and user management only.	Error

**Note:**

In addition to the statuses mentioned so far, an "Overwritten" status exists. An "Overwritten" status occurs when the license is temporarily invalid because a conflicting license exists.

If a license having an overlapping period is registered when the existing license is active, the existing license is overwritten. In this case, the license that is registered later becomes active. If the license registered later expires or is deleted, the existing license becomes active from the "Overwritten" status. However, if the existing license has expired or exceeded the licensed capacity at this time, it enters the GracePeriod from the "Overwritten" status.

For example, when you register a Perpetual license of 500 TiB when you are using 600 TiB, the Perpetual license enters the GracePeriod. If you register a Subscription license for 600 TiB in this state, the Perpetual license is overwritten and the Subscription license becomes active.

Setting a warning threshold

The default valid period of a license is 30 days and the default capacity threshold is 80%. If you set a warning threshold, you can receive a warning notification in an event log when the threshold is reached. To change the threshold setting, see *Editing the license warning threshold*.

Deleting a license (CLI or REST API)

Delete a license as follows. If no license exists after you delete a license, license deletion is not allowed.

Before you begin

Required role: Storage

Procedure

1. Verify the ID of the license to be deleted.

REST API: GET /v1/objects/licenses

CLI: license_list

2. Delete a license.

Run one of the following commands with the license ID specified.

REST API: DELETE /v1/objects/licenses/<id>

CLI: license_uninstall

Verify the job ID which is displayed after the command is run.

3. Verify the state of the job by specifying the job ID.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

4. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Obtaining a list of license information (CLI or REST API)

Obtain a list of information about a license as follows.

- id: License ID (uuid)
- programProductName: Program product name
- status: License status
- statusSummary: License status summary
- cause: Reason for "Warning" or "Error" license status summary
- keyType: License key type
- permittedCapacityInTiB: Total logical capacity of all storage pools allowed by the license (null if no limitation is applied to the capacity)
- totalPoolCapacityInGiB: Total logical capacity of all storage pools (null if no limitation is applied to the capacity)
- remainingDays: Number of remaining days (null if the license is valid infinitely)

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Obtain a list of information about a license.

REST API: GET /v1/objects/licenses

CLI: license_list

Obtaining information about individual licenses (CLI or REST API)

The following information can be obtained.

- License ID (uuid)
- Program product name
- Reason for "Warning" or "Error" license status or statusSummary
- License key type
- Total logical capacity of all storage pools allowed by the license (null if no limit is applied to the capacity)
- Total logical capacity of all storage pools (null if no limitation is applied to the capacity)
- Number of remaining days (null if the license is valid infinitely)

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Verify the ID of the license to be obtained.

REST API: GET /v1/objects/licenses

CLI: license_list

2. Obtain license information.

Run one of the following commands with the license ID specified.

REST API: GET /v1/objects/licenses/<id>

CLI: license_show

Editing the license warning threshold (CLI or REST API)

You can set a warning threshold as follows.

The initial values of the Virtual Storage Software block system (defaults after setup) are as follows:

- Number of remaining days: 30
- Warned capacity rate: 80

Before you begin

Required role: Storage

Procedure

1. Edit the license settings.

Run either of the following commands with the number of days and capacity rate until a warning will be issued specified.

REST API: PATCH /v1/objects/license-setting

CLI: license_setting_set

Verify the job ID which is displayed after the command is run.

2. Verify the state of the job by specifying the job ID.

REST API: GET /v1/objects/jobs/<jobID>

CLI: job_show

If the job state is "Succeeded", the job is completed.

3. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Obtaining license settings (CLI or REST API)

The following information can be obtained.

- warningThresholdSetting: License warning threshold settings (number of remaining days and warned capacity rate)

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Obtain the license settings.

REST API: GET /v1/objects/license-setting

CLI: license_setting_show

Chapter 4: Managing compute nodes and allocating volumes to compute nodes

Installing a compute node

A compute node is a node that runs the application of the user and instructs input/output of user data to the storage node. The compute node is connected to the storage system through the compute network.

After the connection is complete, confirm the following requirements and schedule the tasks. Then, register the information about the compute node according to *Registering information about compute nodes* in this manual.

The following table describes the requirements for compute nodes.

For details about how to enable ALUA, see *Appendix E. ALUA configuration guidelines*. For details about other items, see the documentation for your OS.

Item	Requirement
Supporting and enabling ALUA	Asymmetric logical unit access (ALUA) must be supported. In addition, ALUA must be enabled.
SCSI timeout setting	The SCSI timeout setting of the compute node must be 120 seconds or longer to prevent I/O operations from stopping if a node failure occurs. ¹
Multipath setting (Linux DM-Multipath)	Meet all of the following: <ul style="list-style-type: none">▪ The retry count setting (no_path_retry) for the same path is 6 or more.▪ The path polling interval (polling_interval) setting value is 30 or more.▪ The setting value of fast_io_fail_tmo is other than off.²▪ The setting value of dev_loss_tmo is the OS maximum value.³▪ The failback setting is "immediate".⁴▪ The path_checker setting is readsector0.⁵

Item	Requirement
	<p>1. The SCSI Timeout setting value is the time until the OS determines that a SCSI command has not responded. When the OS detects no response to the SCSI command, it performs recovery processing for the failed path. If the failed path cannot be recovered by the recovery process, the path is blocked and switched to another path. During the recovery process, I/O operations for the relevant path are stopped. Therefore, in the application layer timeout design, it is necessary to consider the recovery processing time in addition to the SCSI Timeout setting value.</p> <p>2. The default value for <code>fast_io_fail_tmo</code> is 5. When set to off, path switching does not occur until <code>dev_loss_tmo</code> seconds have elapsed when a path failure occurs, and it does not operate as expected.</p> <p>3. The maximum value depends on your distribution. See your distribution manual for the maximum value.</p> <p>4. Set the failback policy ("failback") of the path group to "immediate" (immediate failback). If the failback policy is not set to "immediate", I/O will still be issued to the non-priority path after recovery from a path fault, thus requiring a manual path switch operation.</p> <p>5. If a failure occurs in a storage node of Virtual Storage Software block, when a value other than <code>readsector0</code> is set, paths might be wrongly blocked. For this reason, make sure that you set <code>readsector0</code>.</p>

Virtual Storage Software block supports the following VMware vStorage APIs for Array Integration (VAAI) primitives.

VAAI primitive	Description	Support status
Block Zeroing	When a virtual machine is created, blocks are formatted (zeroed out). This primitive enables the storage to perform this procedure, speeding up provisioning of virtual machines. When used with the thin provisioning feature, the primitive can free the block areas that are usually reserved when a virtual machine is created, enabling efficient use of disk capacity.	Supported
Hardware Assisted Locking	If multiple virtual machines share one VMFS volume in your environment, SCSI reservation conflicts might occur and performance might degrade when Storage vMotion or a virtual machine is turned on. With this primitive, you can avoid such a problem. The primitive offers efficient locking.	Supported

VAAI primitive	Description	Support status
Full Copy	Traditionally, VMware ESXi copy data between volumes. This primitive offloads the operation to the storage, greatly reducing the time required for virtual machine cloning and Storage vMotion processing.	Not supported

Registering information about compute nodes (CLI or REST API)

Register information about compute nodes.

The maximum number of registered compute nodes is 1024 per protection domain. This maximum does not change even if storage nodes are added or removed.

When using the GUI

You can use the GUI to perform the tasks described in this section unless the registration destination is a VPS.

For the procedure and details, see the following in the *Hitachi Virtual Storage Software Block GUI Guide*.

- *Registering compute nodes*
- When registering compute nodes to a VPS: Scope of the VPS

Before you begin

Required role: Storage

Procedure

1. When registering a compute node to a VPS, verify the registration-target VPS ID and conditions set for the VPS (upper limit for the number of compute nodes).

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list

2. Register information about compute nodes.

Run either of the following commands with the nicknames of the intended compute nodes and the type of the OS running on the compute nodes specified.

Conventions to be followed when setting a nickname

- Number of characters: 1 to 229
- Characters that can be used: Numbers (0 to 9), uppercase alphabet (A to Z), lowercase alphabet (a to z), symbols (\ . : @ _) for the first character. In addition to these, a hyphen (-) can be used for the second and subsequent characters.
- Each compute node must have a unique nickname.

REST API: POST /v1/objects/servers

CLI: server_create

Verify the job ID which is displayed after the command is run.

3. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobID>

CLI: job_show

If the job state is "Succeeded", the job is completed.

4. Obtain a list of compute nodes and verify that the information about the intended compute nodes is registered.

REST API: GET /v1/objects/servers

CLI: server_list

5. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Obtaining a list of information about compute nodes (CLI or REST API)

The following information can be obtained.

- id: IDs (uuid) of compute nodes
- nickname: Nicknames of compute nodes
- osType: OS types of compute nodes
- totalCapacity: Total capacity of the volumes on the storage pool allocated to the compute node
- usedCapacity: Consumed amount of the volumes on the storage pool allocated to the compute node

- numberOfPaths: Number of registered paths
- vpsId: ID of the VPS to which compute nodes belong
- vpsName: Name of the VPS to which compute nodes belong

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Obtain a list of information about compute nodes.

REST API: GET /v1/objects/servers

CLI: server_list

Obtaining information about individual compute nodes (CLI or REST API)

The following information can be obtained for the compute node with the ID specified.

- numberOfVolumes: Number of allocated volumes
- paths: List of information about registered paths (WWN or iSCSI name of the initiator for the intended compute node, list of IDs of compute ports of targets with which the applicable initiator negotiates)
- id: ID (uuid) of the intended compute node
- nickname: Nickname of the intended compute node
- osType: OS type of the intended compute node
- totalCapacity: Total capacity of the volumes on the storage pool allocated to the compute node
- usedCapacity: Consumed amount of the volumes on the storage pool allocated to the compute node
- numberOfPaths: Number of registered paths
- vpsId: ID of the VPS to which the intended compute node belongs
- vpsName: Name of the VPS to which the intended compute node belongs

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Verify the ID of the intended compute node.

If you use the CLI to specify a compute node by nickname, check the nickname of the compute node.

REST API: GET /v1/objects/servers

CLI: server_list

2. Obtain information about the intended compute node.

Run either of the following commands with the compute node ID specified.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: GET /v1/objects/servers/<id>

CLI: server_show

Editing information about compute nodes (CLI or REST API)

Edit information about the intended compute node. You can edit the nickname and OS type of the intended compute node.

Each compute node must have a unique nickname.

**Caution:**

While the information about a compute node is edited, I/O processing to a volume specified by a volume path in the compute node is temporarily stopped.

Before you begin

- Required role: Storage
- When editing compute nodes belonging to a VPS: Scope of the VPS

Procedure

1. When editing a compute node that belongs to a VPS, verify the VPS ID.

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list

2. Verify the ID of the intended compute node.

If you use the CLI to specify a compute node by nickname, check the nickname of the compute node.

REST API: GET /v1/objects/servers

CLI: server_list

3. Edit information about the intended compute node.

Run either of the following commands with the compute node ID specified.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: PATCH /v1/objects/servers/<id>

CLI: server_set

Verify the job ID which is displayed after the command is run.

4. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

After running the command, if you receive a response indicating "Succeeded" as the state, the job is completed.

5. Obtain a list of compute nodes and verify that the information about the intended compute node is edited.

REST API: GET /v1/objects/servers

CLI: server_list

6. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Deleting information about individual compute nodes (CLI or REST API)

Delete information about the intended compute node. Deleting the compute node information also deletes all compute node initiator information, all compute node path information, and all volume path information.

Before you begin

Required role: Storage

Procedure

1. When deleting a compute node that belongs to a VPS, verify the VPS ID.

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list

2. Verify the ID of the intended compute node.

If you use the CLI to specify a compute node by nickname, check the nickname of the compute node.

REST API: GET /v1/objects/servers

CLI: server_list

3. Delete information about the intended compute node.

Run either of the following commands with the compute node ID specified.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: DELETE /v1/objects/servers/<id>

CLI: server_delete

Verify the job ID which is displayed after the command is run.

4. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobID>

CLI: job_show

If the job state is "Succeeded", the job is completed.

5. Obtain a list of compute nodes and verify that the information about the intended compute node is deleted.

REST API: GET /v1/objects/servers

CLI: server_list

6. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

If tasks are scheduled according to *Scheduling the tasks* in this manual, go to the next step. If no tasks are scheduled, this completes the procedure.

7. From the Start menu, select [Administrative Tools], and then [Task Scheduler].

The Task Scheduler window appears.

8. Select [Task Scheduler Library] in the left side of the window, select the task created in *Scheduling the tasks*, and then click [Delete] from the right-click menu.

Delete all the tasks created in *Scheduling the tasks*.

9. Close the Task Scheduler window.

Registering information about the initiators for compute nodes (CLI or REST API)

After operating the compute node to investigate the initiator name (iSCSI name or WWN) of the compute node, register the compute node initiator information in Virtual Storage Software block from the controller node.

You can use the GUI to perform the tasks described in this section unless the registration destination is a VPS. For the procedure and details, see the following in the *Hitachi Virtual Storage Software Block GUI Guide*.

The maximum number of registered initiators is 4 per compute node.

The procedure differs depending on whether iSCSI connection or FC connection is made between the compute node and the storage node.

In the case of iSCSI connection (CLI or REST API)



Caution:

The initiator name (iSCSI name) of the compute node to be registered with Virtual Storage Software block must be unique in the system. Verify it before registration. It is impossible to see the volume from a compute node whose initiator name (iSCSI name) is the same as that of another compute node.

Also, "iqn" and "eui" at the beginning of the initiator name (iSCSI name) cannot be uppercase. It can only be specified in lower case.

Before you begin

- Required role: Storage
- When registering initiators to a VPS: Scope of the VPS
- The information about the applicable compute node must be registered beforehand.

Procedure

1. Verify the initiator name (iSCSI name) of the applicable compute node.
For details, see the documentation for the OS used on the compute node.
2. Verify that the initiator name (iSCSI name) verified in step 1 is not the same as the initiator name (iSCSI name) of another compute node.
If they are the same, change the initiator name (iSCSI name).
3. When registering initiators to a VPS, verify the VPS ID and conditions set for the VPS (upper limit for the number of initiators).
If you want to specify a VPS by its name in the CLI, verify the VPS name.
REST API: GET /v1/objects/virtual-private-storages
CLI: vps_list

4. Verify the ID of the applicable compute node.

If you use the CLI to specify a compute node by nickname, check the nickname of the compute node.

REST API: GET /v1/objects/servers

CLI: server_list

5. Register information about the intended initiator.

Run either of the following commands with the ID of the compute node, connection protocol for the initiator, and iSCSI name of the initiator.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: POST /v1/objects/servers/<id>/hbas

CLI: hba_create

Verify the job ID which is displayed after the command is run.

6. Verify the state of the job.

Run either of the following commands with the job ID specified.

GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

7. Obtain a list of information about initiators and verify that the information about the intended initiator is registered.

Run either of the following commands with the compute node ID specified.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: GET /v1/objects/servers/<id>/hbas

CLI: hba_list

8. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

In the case of FC connection (CLI or REST API) (Virtual machine)

Before you begin

- Required role: Storage
- When registering initiators to a VPS: Scope of the VPS
- The information about the applicable compute node must be registered beforehand.

Procedure

1. Verify WWN in documentation and so on, from each HBA vendor.
2. When registering initiators to a VPS, verify the VPS ID and conditions set for the VPS (upper limit for the number of initiators).

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list

3. Verify the ID of the applicable compute node.

If you use the CLI to specify a compute node by nickname, check the nickname of the compute node.

REST API: GET /v1/objects/servers

CLI: server_list

4. Register information about the intended initiator.

Run either of the following commands with the ID of the compute node, connection protocol for the initiator, and the initiator name (WWN) specified.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: POST /v1/objects/servers/<id>/hbas

CLI: hba_create

Verify the job ID which is displayed after the command is run.

5. Verify the state of the job.

Run either of the following commands with the job ID specified.

GET /v1/objects/jobs/<jobid>

CLI: job_show

If the job state is "Succeeded", the job is completed.

6. Obtain a list of information about initiators and verify that the information about the intended initiator is registered.

Run either of the following commands with the compute node ID specified.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: GET /v1/objects/servers/<id>/hbas

CLI: hba_list

7. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Obtaining a list of information about initiators for compute nodes (CLI or REST API)

The following information can be obtained.

- id: IDs (uuid) of initiators
- serverId: IDs (uuid) of compute nodes
- name: WWN or iSCSI names of initiators
- protocol: Connection protocols for initiators
- portIds: List of IDs (uuid) of compute ports of targets with which initiators negotiate
- vpsId: ID of the VPS to which initiators belong
- vpsName: Name of the VPS to which initiators belong

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Verify the IDs of compute nodes.

If you use the CLI to specify a compute node by nickname, check the nickname of the compute node.

REST API: GET /v1/objects/servers

CLI: server_list

2. Obtain a list of information about initiators.

Run either of the following commands with a compute node ID specified.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: GET /v1/objects/servers/<id>/hbas

CLI: hba_list

Obtaining information about individual initiators for compute nodes (CLI or REST API)

The following information can be obtained for the initiator with the specified ID.

- id: ID (uuid) of the intended initiator
- serverId: ID (uuid) of the applicable compute node

- name: WWN or iSCSI names of initiators
- protocol: Connection protocol for the intended initiator
- portIds: List of IDs (uuid) of compute ports of targets with which the intended initiator negotiates
- vpsId: ID of the VPS to which the intended initiator belongs
- vpsName: Name of the VPS to which the intended initiator belongs

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Verify the ID of the applicable compute node.

If you use the CLI to specify a compute node by nickname, check the nickname of the compute node.

REST API: GET /v1/objects/servers

CLI: server_list

2. Verify the ID of the intended initiator.

If you use the CLI to specify an initiator by WWN or iSCSI name, check the WWN or iSCSI name of the initiator.

Run either of the following commands with the compute node ID specified.

If you use the CLI, you can specify a nickname instead of the COMPUTE node ID.

REST API: GET /v1/objects/servers/<id>/hbas

CLI: hba_list

3. Obtain information about the intended initiator.

Specify the ID of the applicable compute node and ID of the intended initiator, and run either of the following commands.

If you use the CLI, you can specify a nickname instead of the compute node ID, or a WWN or iSCSI name instead of the initiator ID.

REST API: GET /v1/objects/servers/<id>/hbas/<hbald>

CLI: hba_show

Deleting information about the initiators for compute nodes (CLI or REST API)

Delete the initiator information of a compute node. Deleting information about an initiator also deletes the path information of all the related compute nodes.

Before you begin

Required role: Storage

Procedure

1. When deleting initiators that belong to a VPS, verify the VPS ID.

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list

2. Verify the ID of the applicable compute node.

If you use the CLI to specify a compute node by nickname, check the nickname of the compute node.

REST API: GET /v1/objects/servers

CLI: server_list

3. Verify the ID of the intended initiator.

If you use the CLI to specify an initiator by WWN or iSCSI name, check the WWN or iSCSI name of the initiator.

Specify the ID of the applicable compute node and run either of the following commands.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: GET /v1/objects/servers/<id>/hbas

CLI: hba_list

4. Delete information about the intended initiator.

Specify the ID of the applicable compute node and ID of the intended initiator, and run either of the following commands.

If you use the CLI, you can specify a nickname instead of the compute node ID, or a WWN or iSCSI name instead of the initiator ID.

REST API: DELETE /v1/objects/servers/<id>/hbas/<hbald>

CLI: hba_delete

Verify the job ID which is displayed after the command is run.

5. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

6. Obtain a list of information about initiators and verify that the information about the intended initiator is deleted.

Run either of the following commands with the compute node ID specified.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: GET /v1/objects/servers/<id>/hbas

CLI: hba_list

7. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Registering path information about compute nodes (CLI or REST API)

The maximum number of registered compute node paths is 4096 per compute node.

Depending on how parameters portId and hbaId are specified, the following are possible:

- Both portId and hbaId are omitted: All the initiators of the intended compute node are allocated to all the compute ports.
- Only hbaId is specified: The specified initiator is allocated to all the compute ports.
- Only portId is specified: The specified compute port is allocated to all the initiators of the intended compute nodes.
- Both portId and hbaId are specified: The initiator specified by hbaId is allocated to the compute port specified by portId.



Caution:

- If you have changed the path information of a compute node, perform a rescan of the storage in that compute node before you start I/O operations.
- If using VMware ESXi as a compute node, set up a path between the compute node and all computer ports. Unless all paths have been set, some volumes might be invisible from a compute node.

If not using VMware ESXi as a compute node, it is also recommended to set up a path between the compute node and all computer ports to prevent I/O performance from being deteriorated.

Before you begin

- Required role: Storage
- The information about the intended compute node and its initiator must be registered beforehand.
- When registering compute node path information to a VPS: Scope of the VPS

Procedure

1. When registering compute node path information to a VPS, verify the registration-target VPS ID.

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list

2. Verify the ID of the intended compute node.

If you use the CLI to specify a compute node by nickname, check the nickname of the compute node.

REST API: GET /v1/objects/servers

CLI: server_list

3. When you specify the hbaId parameter, verify the ID of the initiator for the intended compute node.

Run either of the following commands with the compute node ID specified.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: GET /v1/objects/servers/<id>/hbas

CLI: hba_list

4. When you specify the portId parameter, verify the ID of the compute port to be allocated to the intended compute node.

REST API: GET /v1/objects/ports

CLI: port_list

5. Register compute node path information.

Run either of the following commands with the compute node ID specified.

If you use the CLI, you can specify a nickname instead of the compute node ID. You can also specify the compute port of the target behavior to which it is assigned by WWN or iSCSI Name instead of the ID.

REST API: POST /v1/objects/servers/<id>/paths

CLI: path_create

Verify the job ID which is displayed after the command is run.

6. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobID>

CLI: job_show

If the job state is "Succeeded", the job is completed.

7. Obtain a list of path information and verify that the intended path information is added.

Run either of the following commands with the compute node ID specified.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: GET /v1/objects/servers/<id>/paths

CLI: path_list

8. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Obtaining a list of path information about compute nodes (CLI or REST API)

The following information can be obtained.

- id: Path IDs (character string of initiator ID of the computer node and ID of the compute port for the target operation connected by a comma)
- serverId: IDs (uuid) of compute nodes
- hbaName: WWNs or iSCSI names for compute nodes
- hbalId: IDs (uuid) of initiators for compute nodes
- portId: IDs (uuid) of compute ports of targets with which initiators negotiate
- portName: WWNs or iSCSI names of compute ports of targets with which initiators negotiate
- portNickname: Nicknames of compute ports of targets with which initiators negotiate
- vpsId: ID of the VPS to which compute node paths belong
- vpsName: Name of the VPS to which compute node paths belong

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Verify the IDs of compute nodes.

REST API: GET /v1/objects/servers

CLI: server_list

2. Obtain a list of path information about compute nodes.

Run either of the following commands with a compute node ID specified.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: GET /v1/objects/servers/<id>/paths

CLI: path_list

Obtaining specific path information about compute nodes (CLI or REST API)

The following information can be obtained for the path information about compute nodes with the specified ID.

- id: Path IDs (character string of initiator ID of the computer node and ID of the compute port for the target operation connected by a comma)
- serverId: IDs (uuid) of compute nodes
- hbaName: WWNs or iSCSI names of initiators for compute nodes
- hbaId: IDs (uuid) of initiators for compute nodes
- portId: IDs (uuid) of compute ports of targets with which initiators negotiate
- portName: WWNs or iSCSI names of compute ports of targets with which initiators negotiate
- portNickname: Nicknames of compute ports of targets with which initiators negotiate
- vpsId: ID of the VPS to which the compute node path belongs
- vpsName: Name of the VPS to which the compute node path belongs

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Verify the ID of the applicable compute node.

If you use the CLI to specify a compute node by nickname, check the nickname of the compute node.

REST API: GET /v1/objects/servers

CLI: server_list

2. Verify the ID of the initiator for the applicable compute node and the ID of the compute port.

If you use the CLI to specify initiators and compute ports by WWN or iSCSI name, check the WWN or iSCSI name of the initiator and compute port.

Run either of the following commands with the compute node ID specified.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: GET /v1/objects/servers/<id>/hbas

CLI: hba_list

3. Obtain path information about the applicable compute node.

Run either of the following commands with the ID of the compute node, ID of the initiator, and ID of the compute port for the target operation specified.

If you use the CLI, you can specify a nickname instead of the compute node ID, a WWN or iSCSI name instead of the initiator ID, and a WWN or iSCSI name instead of the compute port ID.

REST API: GET /v1/objects/servers/<id>/paths/<hbald>,<portId>

CLI: path_show

Deleting path information about compute nodes (CLI or REST API)



Caution:

- When compute nodes are clustered and the volumes recognized by the compute nodes are online, set the intended volume offline and then delete path information.
- Before you delete path information from a compute node, verify whether the volumes that can be accessed from the compute node are in SCSI-2 Reserve status or SCSI-3 Persistent Reserve status. If the volumes are in either status, release them from the status, and then delete path information.
- When you change the path information of a compute node, perform a rescan of the storage on that compute node. If the path information already deleted remains on the compute node, it might cause malfunction.

Before you begin

- Required role: Storage
- When deleting path information of the compute nodes belonging to a VPS: Scope of the VPS

Procedure

1. When deleting path information about compute nodes that belong to a VPS, verify the VPS ID.

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list

2. Verify the ID of the applicable compute node.

If you use the CLI to specify a compute node by nickname, check the nickname of the compute node.

REST API: GET /v1/objects/servers

CLI: server_list

3. Obtain a list of path information about compute nodes.

Run either of the following commands with the compute node ID specified.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: GET /v1/objects/servers/<id>/paths

CLI: path_list

4. Delete path information from the applicable compute node.

Run either of the following commands with the ID of the compute node, ID of the initiator, and ID of the compute port for the target operation specified.

If you use the CLI, you can specify a nickname instead of the compute node ID, a WWN or iSCSI name instead of the initiator ID, and a WWN or iSCSI name instead of the compute port ID.

REST API: DELETE /v1/objects/servers/<id>/paths/<hbaid>,<portId>

CLI: path_delete

Verify the job ID which is displayed after the command is run.

5. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

6. Obtain a list of path information and verify that the intended path information is deleted.

Run either of the following commands with the compute node ID specified.

If you use the CLI, you can specify a nickname instead of the compute node ID.

REST API: GET /v1/objects/servers/<id>/paths

CLI: path_list

7. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Allocating volumes to compute nodes (CLI or REST API)

Set paths (volume paths) between volumes and compute nodes.

The maximum number of volume paths that can be registered is as follows: 8,192 per compute node; 2,048 per storage controller; and 65,536 per storage cluster.

The maximum number of registered volume paths is 8,192 per compute node, and 65,536 per storage cluster.

The following two combinations are possible for specifying parameters.

- Combination of volumeld, serverId, and lun: Allocates a volume specified by volumeld to a compute node specified by serverId (lun is optional).
- Combination of volumelds, serverIds, and startLun: Allocates all the volumes specified with volumelds to all the compute nodes specified with serverIds (startLun can be omitted).

**Caution:**

- When using VMware ESXi as a compute node, set up a path between the compute node and all computer ports when registering path information of the compute node. Unless all paths have been set, some volumes might be invisible from a compute node.

Even when not using VMware ESXi as a compute node, we recommend that you set up a path between the compute node and all computer ports to prevent deterioration in I/O performance.

For how to set up a path, see *Registering path information about compute nodes*.

- If you have changed the connection information between a volume and compute node, perform a rescan of the storage in that compute node before you start I/O operations.
- If you specify a volume path by omitting lun or startLun, the smallest LUN number that is not used is automatically allocated. However, if multiple volume paths are set without specifying lun at the same time, or if volume paths are set for multiple volumes, LUNs might not be assigned in ascending order they are set. To assign a specific LUN, set a volume path by specifying lun.
- If some volumes cannot be recognized by the OS, the system might behave as follows. Resolve the state where volumes cannot be recognized, and then perform a rescan of the storage on the compute node.
 - During recognition of LUNs in the order from the smallest, if an unrecognizable LUN exists, LUN recognition stops.
 - When LUN=0 cannot be recognized, LUN recognition stops.
- Processing performance when registering volume paths

The processing time varies depending on the number of volume paths registered on the applicable compute node.

 - When registering the first volume path, the processing completes in one or two seconds.
 - For every 1,000 registered volume paths, the processing time for registering one volume path will increase by 1 to 2 seconds.
 - When registering the 8,192th volume path (upper limit), the processing takes approximately 15 seconds.
- If any of the following occurs, volumes might not be recognized from the host.
 - Capacity depletion of storage controller

If volumes cannot be recognized by the host, verify if KARS06003-E is issued. If either event log is issued, resolve the problem with the storage controller according to the indicated action.

Before you begin

- Required role: Storage
- The intended volume must be created and the information about the intended compute node must be registered beforehand.
- When registering volume paths to a VPS: Scope of the VPS

Procedure

1. When registering a volume path to a VPS, verify the registration-target VPS ID and conditions set for the VPS (upper limit for the number of volume paths).
If you want to specify a VPS by its name in the CLI, verify the VPS name.
REST API: GET /v1/objects/virtual-private-storages
CLI: vps_list
2. Verify the ID of the volume to be allocated to the compute node.
If you use the CLI to specify a volume by name, check the name of the volume.
REST API: GET /v1/objects/volumes
CLI: volume_list
3. Verify the ID of the compute node to which the volume is to be allocated.
If you use the CLI to specify a compute node by nickname, check the nickname of the compute node.
REST API: GET /v1/objects/servers
CLI: server_list
4. Allocate the volume to the compute node.
REST API: POST /v1/objects/volume-server-connections
CLI: volume_server_connection_create
Verify the job ID which is displayed after the command is run.
5. Verify the state of the job.
Run either of the following commands with the job ID specified.
REST API: GET /v1/objects/jobs/<jobId>
CLI: job_show
If the job state is "Succeeded", the job is completed.
6. Obtain a list of information about allocation of volumes to compute nodes, and verify that the intended volume is allocated to the intended compute node.
REST API: GET /v1/objects/volume-server-connections
CLI: volume_server_connection_list

7. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Obtaining a list of information about allocation of volumes to compute nodes (CLI or REST API)

The following information can be obtained.

- id: ID of volume paths (each ID is a string consisting of a volume ID and a compute node ID connected by a Comma (,))
- serverId: IDs (uuid) of compute nodes
- volumeId: IDs (uuid) of volumes
- lun: LUN
- vpsId: ID of the VPS to which volume paths belong
- vpsName: Name of the VPS to which volume paths belong

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Verify the IDs of volumes.

REST API: GET /v1/objects/volumes

CLI: volume_list

2. Verify the IDs of compute nodes.

REST API: GET /v1/objects/servers

CLI: server_list

3. Obtain a list of information about allocation of volumes to compute nodes.

Run one of the following commands with the volume ID and compute node ID specified in the query parameter.

REST API: GET /v1/objects/volume-server-connections

CLI: volume_server_connection_list

Obtaining information about allocation of individual volumes to individual compute nodes (CLI or REST API)

Obtain the information about allocation of volumes to compute nodes with the specified volume ID and compute node ID.

- `id`: ID of volume paths (a string consisting of a volume ID and a compute node ID connected by a Comma (,))
- `serverId`: ID of the VPS to which the volume path belongs
- `volumeId`: Name of the VPS to which the volume path belongs
- `lun`: LUN
- `vpsId`: ID of a virtual private storage (VPS) to which volume paths belong
- `vpsName`: Name of a virtual private storage (VPS) to which volume paths belong

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Verify the ID of the intended volume.

If you use the CLI to specify a volume by name, check the name of the volume.

REST API: GET /v1/objects/volumes

CLI: `volume_list`

2. Verify the ID of the intended compute node.

If you use the CLI to specify a compute node by nickname, check the nickname of the compute node.

REST API: GET /v1/objects/servers

CLI: `server_list`

3. Obtain information about allocation of the intended volume to the intended compute node.

Specify the ID of the intended volume and ID of the intended compute node, and run either of the following commands.

If you use the CLI, you can specify a name instead of the ID of the volume, but a nickname instead of the ID of the compute node.

REST API: GET /v1/objects/volume-server-connections/<volumeId>,<serverId>

CLI: `volume_server_connection_show`

Canceling allocation of volumes to compute nodes (CLI or REST API)

Disconnects by removing the path (volume path) between the volume and the compute node. Before you run the command, verify that no I/O operation is being performed between the intended compute node and the intended volume.



Caution:

- When compute nodes are clustered and the volumes recognized by the compute nodes are online, set the intended volume offline and then cancel allocation of the volume.
- Before you cancel allocation of a volume to a compute node, verify whether the volumes that can be accessed from the compute node are in SCSI-2 Reserve status or SCSI-3 Persistent Reserve status. If the volumes are in either status, release them from the status, and then cancel allocation of the intended volume.
- When you change the path information of a compute node, perform a rescan of the storage on that compute node. If the path information already deleted remains on the compute node, it might cause malfunction.

Before you begin

- Required role: Storage
- When removing volume paths belonging to a VPS: Scope of the VPS

Procedure

1. When removing volume paths that belong to a VPS, verify the VPS ID.

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list

2. Verify the ID of the intended volume.

If you use the CLI to specify a volume by name, check the name of the volume.

REST API: GET /v1/objects/volumes

CLI: volume_list

3. Verify the ID of the intended compute node.

If you use the CLI to specify a compute node by nickname, check the nickname of the compute node.

REST API: GET /v1/objects/servers

CLI: server_list

4. Obtain a list of information about allocation of volumes to compute nodes.

Run one of the following commands with the volume ID and compute node ID specified in the query parameter.

If you use the CLI, you can specify a name instead of the ID of the volume, but a nickname instead of the ID of the compute node.

REST API: GET /v1/objects/volume-server-connections

CLI: volume_server_connection_list

5. Cancel allocation of the volume to the compute node.

Specify the ID of the intended volume and ID of the intended compute node, and run either of the following commands.

If you use the CLI, you can specify a name instead of the ID of the volume, but a nickname instead of the ID of the compute node.

REST API: DELETE /v1/objects/volume-server-connections/<volumeId>,<serverId>

CLI: volume_server_connection_delete

Verify the job ID which is displayed after the command is run.

6. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

7. Obtain a list of information about allocation of volumes to compute nodes and verify that the intended allocation is canceled.

REST API: GET /v1/objects/volume-server-connections

CLI: volume_server_connection_list

8. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Releasing multiple connections between the volumes and compute nodes (CLI or REST API)

Releases the connections by removing all specified volumes and the path (volume path) of all specified compute nodes. Before you execute the operation, verify that no I/O operation is being performed between the intended compute node and the intended volume.



Caution:

- When compute nodes are clustered and the volumes recognized by the compute nodes are online, set the intended volume offline, and then cancel allocation of the volume.
- Before you cancel allocation of a volume to a compute node, verify whether the volumes that can be accessed from the compute node are in SCSI-2 Reserve status or SCSI-3 Persistent Reserve status. If the volumes are in either status, release them from the status, and then cancel allocation of the intended volume.
- When you change the path information of a compute node, perform a rescan of the storage on that compute node. If the path information already deleted remains on the compute node, it might cause malfunction.

Before you begin

- Required role: Storage
- When removing volume paths belonging to a VPS: Scope of the VPS

Procedure

1. When removing volume paths that belong to a VPS, verify the VPS ID.

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list

2. Verify the IDs of volumes.

If you use the CLI to specify a volume by name, check the name of the volume.

REST API: GET /v1/objects/volumes

CLI: volume_list

3. Verify the ID of the intended compute node.

If you use the CLI to specify a compute node by nickname, check the nickname of the compute node.

REST API: GET /v1/objects/servers

CLI: server_list

4. Cancel allocation of the volumes to the compute nodes.

Specify the IDs of the volumes and IDs of the compute node (for which you want to cancel allocation), and then run either of the following commands.

If you use the CLI, you can specify a name instead of the ID of the volume, but a nickname instead of the ID of the compute node.

REST API: POST /v1/objects/volume-server-connections/actions/release/invoke

CLI: volume_server_connection_release_connections

Verify the job ID which is displayed after the command is run.

5. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

After running the command, if you receive a response indicating "Succeeded" as the state, the job is completed.

6. Obtain a list of information about allocation of volumes to compute nodes and verify that the intended allocation is canceled.

REST API: GET /v1/objects/volume-server-connections

CLI: volume_server_connection_list

7. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.


Managing compute nodes

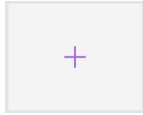



Overview

The following table describes the operations that you can perform for compute nodes.

For details about the procedure flow and prerequisites for operating a compute node, contact customer support.

This manual describes the procedure that can be performed by using the GUI.

Operation	Window to operate on	Operation icon	Dialog
Registering compute nodes	Compute Nodes window	For List view: 	Register Compute node

Operation	Window to operate on	Operation icon	Dialog
		For Inventory view: 	
Editing compute nodes	Compute Nodes window Compute Node detailed information window		Edit Compute node
Deleting compute nodes	Compute Nodes window Compute Node detailed information window		Delete Compute nodes
Reconnecting compute nodes and all the compute ports in full mesh	Compute Node detailed information window	 *	Configure Port Connections
<p>* This icon appears on the Compute Node detailed information window when the applicable compute node is not connected to the compute ports in full mesh.</p> <p>In this manual, full mesh means that path information has been configured for all combinations of initiators and compute ports of a compute node.</p>			

Registering compute nodes (GUI)

Register a compute node, initiator information, and path information about the compute node as follows.

The path of the compute node is configured in full mesh.



Caution:

Performing the following procedure enables setting a path of a compute node in full mesh without physical wiring. If you want to set a compute node path between only a specific initiator and specific compute node, do so from the REST API or CLI. For details about how to set a path of a compute node by using the REST API or CLI, see *Compute node connection management* in the *Hitachi Virtual Storage Software Block REST API Reference* and *Hitachi Virtual Storage Software Block CLI Reference*.

Before you begin

Required role: Storage


Procedure

1. Open the Compute Nodes window, and click either of the following icons.

For List view:



For Inventory view:



The following dialog appears.

2. Enter each of the following parameters:
 - **COMPUTE NODE NAME:** The nickname of the compute node
 - **OS TYPE:** The OS type of the intended compute node
 - **WWNS:**
 - (Virtual machine) WWN in the case of FC connection
 - (Bare metal) Not used. Do not enter this parameter.
 - **iSCSI INITIATOR NAMES:** iSCSI name in the case of iSCSI connection

Clicking **+ Add WWN** or **+ Add iSCSI Initiator Name** displays an additional input field. To the right of the input field, there is an **x** icon. To remove the entry field, click the **x** icon.

3. Click **Submit**.

If you want to continue registering other compute nodes, click **+Submit and add another compute node**.

4. When the following "Completed" message is displayed, processing is completed.
 - Successfully configured port connections.

Editing compute nodes (GUI)

Edit information about the intended compute node as follows.

Before you begin

- Required role: Storage

Procedure

1. In the Compute Nodes window or the Compute Node detailed information window, edit the information by using one of the following methods:



- In the Compute Nodes window, select the editing-target compute node (one node), and then click the preceding icon shown to the right of Select All.
- Change the view mode of the Compute Nodes window to Inventory view, and then click the preceding icon shown for the editing-target compute node.
- In the Compute Node detailed information window for the editing-target compute node, click the preceding icon.

The following dialog appears.

The screenshot shows the 'Edit Compute Node' dialog box. It contains the following fields and controls:

- COMPUTE NODE NAME:** A text input field containing 'ComputeNode000001'.
- OS TYPE:** A dropdown menu with 'Linux' selected.
- WWNS:** A list of WWN addresses: '0000000000000080' and '0000000000000081'. Each has a delete 'X' icon. There is an '+ Add WWN' button below.
- ISCSI INITIATOR NAMES:** A list of iSCSI initiator names: 'iqn.1996-04.de.suse:129c2c1eab83' and 'iqn.1996-04.de.suse:129c2c1eab82'. Each has a delete 'X' icon. There is an '+ Add iSCSI Initiator Name' button below.
- Buttons:** 'Cancel' and 'Submit' buttons are located at the bottom right of the dialog.

2. You can edit the following parameters:
 - **COMPUTE NODE NAME:** The nickname of the compute node
 - **OS TYPE:** The OS type of the intended compute node

- WWNS:
(Virtual machine) WWN in the case of FC connection
(Bare metal) Not used. Do not enter this parameter.
- ISCSI INITIATOR NAMES: The iSCSI name in the case of iSCSI connection

Clicking **+ Add WWN** or **+ Add iSCSI Initiator Name** displays an additional input field. To the right of the input field, there is an **x** icon. To remove the entry field, click the **x** icon.

3. Click **Submit**.
4. When one of the following "Completed" messages is displayed, processing is completed.
 - When you modified COMPUTE NODE NAME or OS TYPE: Successfully edited compute node info. (Compute node Name: XXX)
 - When you added or modified the HBA information: Successfully configured port connections.
 - When you only deleted the HBA information: Successfully deleted initiator of compute node.

Deleting compute nodes (GUI)

Delete information about the intended compute node as follows. Deleting the compute node information also deletes all of the compute node initiator information and all of the compute node path information.

Before you begin

- Required role: Storage


Procedure

1. In the Compute Nodes window or the Compute Node detailed information window, delete the compute node by using one of the following methods:



- In the Compute Nodes window, select the deletion-target compute nodes (1 to 25 nodes), and then click the preceding icon shown to the right of Select All.
- Change the view mode of the Compute Nodes window to Inventory view, and then click the preceding icon shown for each deletion-target compute node.
- In the Compute Node detailed information window for each deletion-target compute node, click the preceding icon.

The following dialog appears.



2. Click **Submit**.
3. When the following "Completed" message is displayed, processing is completed.
 - Successfully deleted compute nodes.

Reconnecting compute nodes and all the compute ports in full mesh (GUI)

Reconnect compute nodes and all the compute ports in full mesh as follows.

A compute node not in a full-mesh connection is marked as (not fullmeshed) in the PORT CONNECTIONS item on the Compute Node detailed information window. In this case, you can perform the following operations:



Caution:

Performing the following procedure enables setting a path of a compute node in full mesh without physical wiring. If you want to set a compute node path between only a specific initiator and specific compute node, do so from the REST API or CLI. For details about how to set a path of a compute node by using the REST API or CLI, see *Compute node connection management* in the *Hitachi Virtual Storage Software Block REST API Reference* and *Hitachi Virtual Storage Software Block CLI Reference*.

Before you begin


- Required role: Storage

Procedure

1. In the Compute Node detailed information window, click the following icon.



The following dialog appears.



2. Click **Submit**.
3. When the following "Completed" message is displayed, processing is completed.
 - Successfully configured port connections.

Allocating and canceling allocation of volumes to and from compute nodes

Creating volumes and allocating them to compute nodes (GUI)

Create volumes and allocate them to compute nodes.

Before you begin

- Required role: Storage

Procedure

- In the Compute Nodes window or the Compute Node detailed information window, allocate volumes to compute nodes by using one of the following methods:



- In the Compute Nodes window, select the allocation-target compute nodes (1 to 100 nodes), click the preceding icon shown to the right of Select All, and then, from the menu that appears, select Create and Attach Volumes.
- Change the view mode of the Compute Nodes window to Inventory view, click the preceding icon shown for each allocation-target compute node, and then, from the menu that appears, select Create and Attach Volumes.
- In the Compute Node detailed information window for each allocation-target node, click the preceding icon, and then, from the menu that appears, select Create and Attach Volumes.

The following dialog appears.

Hitachi Virtual Storage Software Block

EVENT MONITOR PROTECTION DOMAINS STORAGE NODES DRIVES COMPUTE NODES VOLUMES STORAGE POOLS PORTS

< Create and Attach Volumes

COMPUTE NODES
ComputeNode000001

Specify volume information to create and attach.

CAPACITY (0.05 GiB to 9978.40 GiB) GiB

NUMBER OF VOLUMES (1 to 1000)

VOLUME NAME

SUFFIX START NUMBER NUMBER OF DIGITS (Not specified)

VOLUME NICKNAME (Optional)

START LUN (Optional)

Cancel Submit

2. Enter each of the following parameters:
 - CAPACITY: Logical capacity of the volume and its unit
 - NUMBER OF VOLUMES: The number of volumes to be created
 - VOLUME NAME: The name of the volume
 - SUFFIX START NUMBER: The first sequential number suffixed to a volume name or nickname, in the case multiple volumes are created with the same volume name or nickname. If omitted, no number is added.
 - NUMBER OF DIGITS: The number of digits of a number suffixed to a name or nickname
 - VOLUME NICKNAME: The nickname of the volume. If omitted, VOLUME NAME is used.
 - START LUN: LUN start number

If you specify the LUN start number, unused LUN numbers are allocated in ascending order from the specified start number. If omitted, unused LUNs are allocated automatically in ascending order.
3. Click **Submit**.
4. When the following "Completed" message is displayed, processing is completed.
 - Successfully attached volumes to compute nodes.

Allocating volumes to compute nodes (Volume) (GUI)

Set paths (volume paths) between volumes and compute nodes.

The following describes the steps to be performed from the Volumes window or Volume detailed information window. For the steps to be taken from the Compute Nodes window or Compute Node detailed information window, see *Allocating volumes to compute nodes (Compute Node)*.

Before you begin

- Required role: Storage

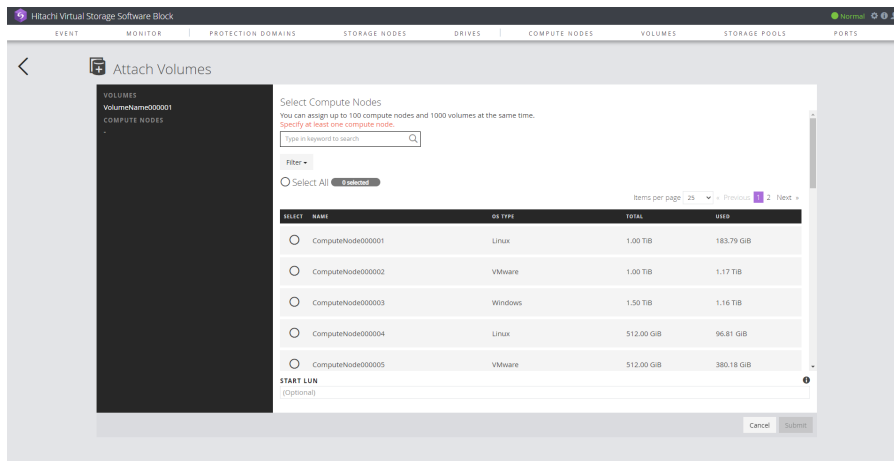
Procedure

1. In the Volumes window or the Volume detailed information window, allocate volumes by using one of the following methods:



- In the Volumes window, select the allocation-target volumes (1 to 1,000 volumes), and then click the preceding icon shown to the right of Select All.
- Change the view mode of the Volumes window to Inventory view, and then click the preceding icon shown for each allocation-target volume.
- In the Volume detailed information window for each allocation-target volume, click the preceding icon.

The following dialog appears.



2. Select one or more compute nodes (100 nodes at maximum) to be allocated, enter the following as required, and then click **Submit**.
 - **START LUN**: LUN start number
 If you specify the LUN start number, unused LUN numbers are allocated in ascending order from the specified start number. If omitted, unused LUNs are allocated automatically in ascending order.
3. When the following "Completed" message is displayed, processing is completed.
 - Successfully attached volumes to compute nodes.

Allocating volumes to compute nodes (Compute Node) (GUI)

Set paths (volume paths) between volumes and compute nodes.

The following describes the steps to be performed from the Compute Nodes window or Compute Node detailed information window. For the steps to be taken from the Volumes window or Volume detailed information window, see *Allocating volumes to compute nodes (Volume)*.

Before you begin

- Required role: Storage

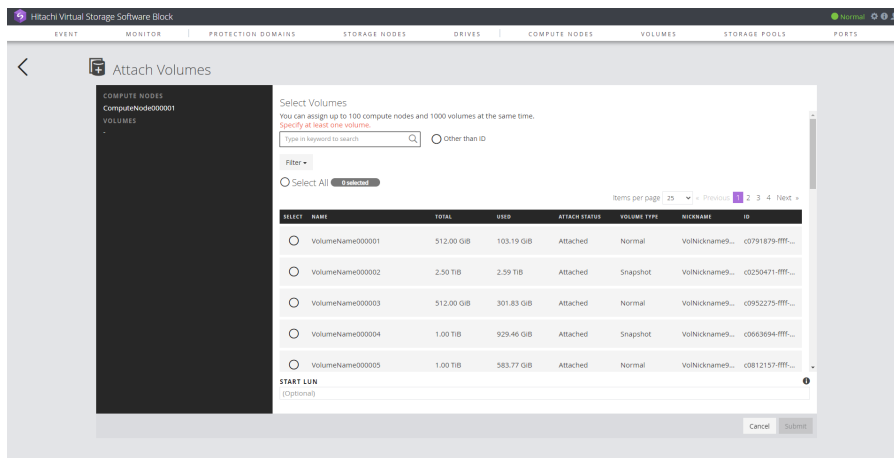
Procedure

1. In the Compute Nodes window or the Compute Node detailed information window, allocate volumes to compute nodes by using one of the following methods:



- In the Compute Nodes window, select the allocation-target compute nodes (1 to 100 nodes), click the preceding icon shown to the right of Select All, and then, from the menu that appears, select Attach Volumes.
- Change the view mode of the Compute Nodes window to Inventory view, click the preceding icon shown for each allocation-target compute node, and then, from the menu that appears, select Attach Volumes.
- In the Compute Node detailed information window for each allocation-target node, click the preceding icon, and then, from the menu that appears, select Attach Volumes.

The following dialog appears.



2. Select one or more compute nodes (1,000 nodes at maximum) to be allocated, enter the following as required, and then click **Submit**.
 - **START LUN:** LUN start number
 If you specify the LUN start number, unused LUN numbers are allocated in ascending order from the specified start number. If omitted, unused LUNs are allocated automatically in ascending order.
3. When the following "Completed" message is displayed, processing is completed.
 - Successfully attached volumes to compute nodes.

Canceling allocation of volumes to compute nodes (GUI)

Cancel allocation of the volumes to the compute nodes as follows. Before you start operation, verify that no I/O operation is being performed between the intended compute node and the intended volume.

Before you begin

- Required role: Storage

Procedure

1. In the Compute Nodes window or the Compute Node detailed information window, cancel allocation by using one of the following methods:

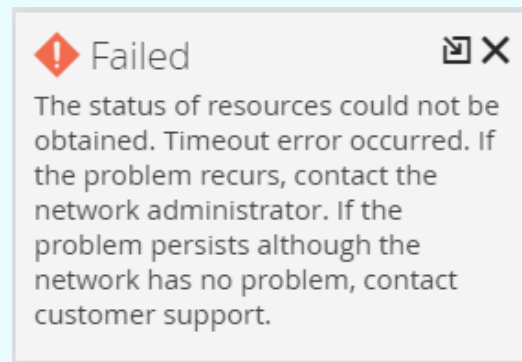


- In the Compute Nodes window, select cancellation-target compute nodes (1 to 100 nodes), and then click the preceding icon shown to the right of Select All.



Note:

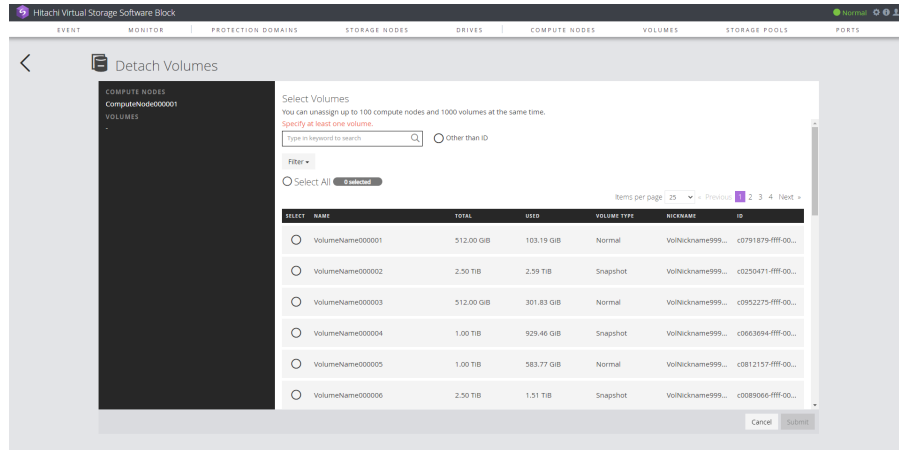
In the Compute Nodes window, if you select multiple cancellation-target compute nodes and click the icon, a timeout error message (shown below) might be displayed. However, because this is due to a transient problem that does not affect cancellation of allocation between volumes and compute nodes, you do not need to take any action.




- Change the view mode of the Compute Nodes window to Inventory view, and then click the preceding icon shown for each cancellation-target compute node.
- In the Compute Node detailed information window for each cancellation-target node, select the volumes (1 to 1,000 volumes), and then click the preceding icon shown to the right of Select All.
- In the Compute Node detailed information window for each cancellation-target node, in the list of volumes displayed in Inventory view, click the preceding icon shown for each cancellation-target volume.

The following dialog appears.

If you operate from the Compute Nodes window:



If you operate from the Compute Node detailed information window:



2. If you operate from the Compute Nodes window, select volumes (1,000 volumes at maximum), and then click **Submit**.

If you operate from the Compute Node detailed window, click **Submit**.

3. When the following "Completed" message is displayed, processing is completed.
 - Successfully detached volumes.

Chapter 5: Managing storage nodes

Obtaining a list of information about storage nodes (CLI or REST API)

The following information can be obtained.

- id: IDs (uuid) of storage nodes
- biosUuid: UUIDs of storage nodes
- protectionDomainId: IDs (uuid) of protection domains containing storage nodes
- faultDomainId: IDs (uuid) of fault domains containing storage nodes
- faultDomainName: Names of fault domains containing storage nodes
- name: Names of storage nodes
- clusterRole: Role of the storage nodes in the storage cluster
- storageNodeAttributes: Attributes of storage nodes
- statusSummary: Summary of statuses of storage nodes
- status: Status of each storage node
- driveDataRelocationStatus: Status of drive data relocation
- controlPortIpv4Address: IP addresses (IPv4) of control ports
- internodePortIpv4Address: IP addresses (IPv4) of internode ports
- softwareVersion: Version of the storage software
- modelName: Model name of the server on which the intended storage node runs
- serialNumber: Serial number of the server on which the intended storage node runs
- memory: Size of memory installed on the storage node
(Virtual machine) Memory capacity of VM
(Bare metal) Memory capacity available in the physical server

- insufficientResourcesForRebuildCapacity: Resources with insufficient rebuild capacity
 - capacityOfDrive: Insufficient drive capacity for rebuild capacity
 - numberOfDrives: Insufficient number of drives for rebuild capacity
- rebuildableResources: Rebuildable resources
 - numberOfDrives: Number of rebuildable drives

Before you begin

Required role: Security, Storage, Monitor, Service, Audit, or Resource

Procedure

1. Obtain a list of information about storage nodes.

REST API: GET /v1/objects/storage-nodes

CLI: storage_node_list

Obtaining information about individual storage nodes (CLI or REST API)

The following information can be obtained.

- id: ID (uuid) of the intended storage node
- biosUuid: UUID of the intended storage node
- protectionDomainId: ID (uuid) of the protection domain containing the intended storage node
- faultDomainId: ID (uuid) and name of the fault domain containing the intended storage node
- faultDomainName: ID and name of the fault domain containing the intended storage node
- name: Name of the intended storage node
- clusterRole: Role of the intended storage node in the storage cluster
- storageNodeAttributes: Attributes of the intended storage node
- statusSummary: Status summary of the intended storage node
- status: Status of the intended storage node
- driveDataRelocationStatus: Status of drive data relocation

- controlPortIpv4Address: IP addresses (IPv4) of the control port
- internodePortIpv4Address: IP addresses (IPv4) of the internode port
- softwareVersion: Version of the storage software
- modelName: Model name of the server on which the intended storage node runs
- serialNumber: Serial number of the server on which the intended storage node runs
- memory: Size of memory installed on the storage node
(Virtual machine) Memory capacity of VM
(Bare metal) Memory capacity available in the physical server
- insufficientResourcesForRebuildCapacity: Resources with insufficient rebuild capacity
 - capacityOfDrive: Insufficient drive capacity for rebuild capacity
 - numberOfDrives: Insufficient number of drives for rebuild capacity
- rebuildableResources: Rebuildable resources
 - numberOfDrives: Number of rebuildable drives

Before you begin

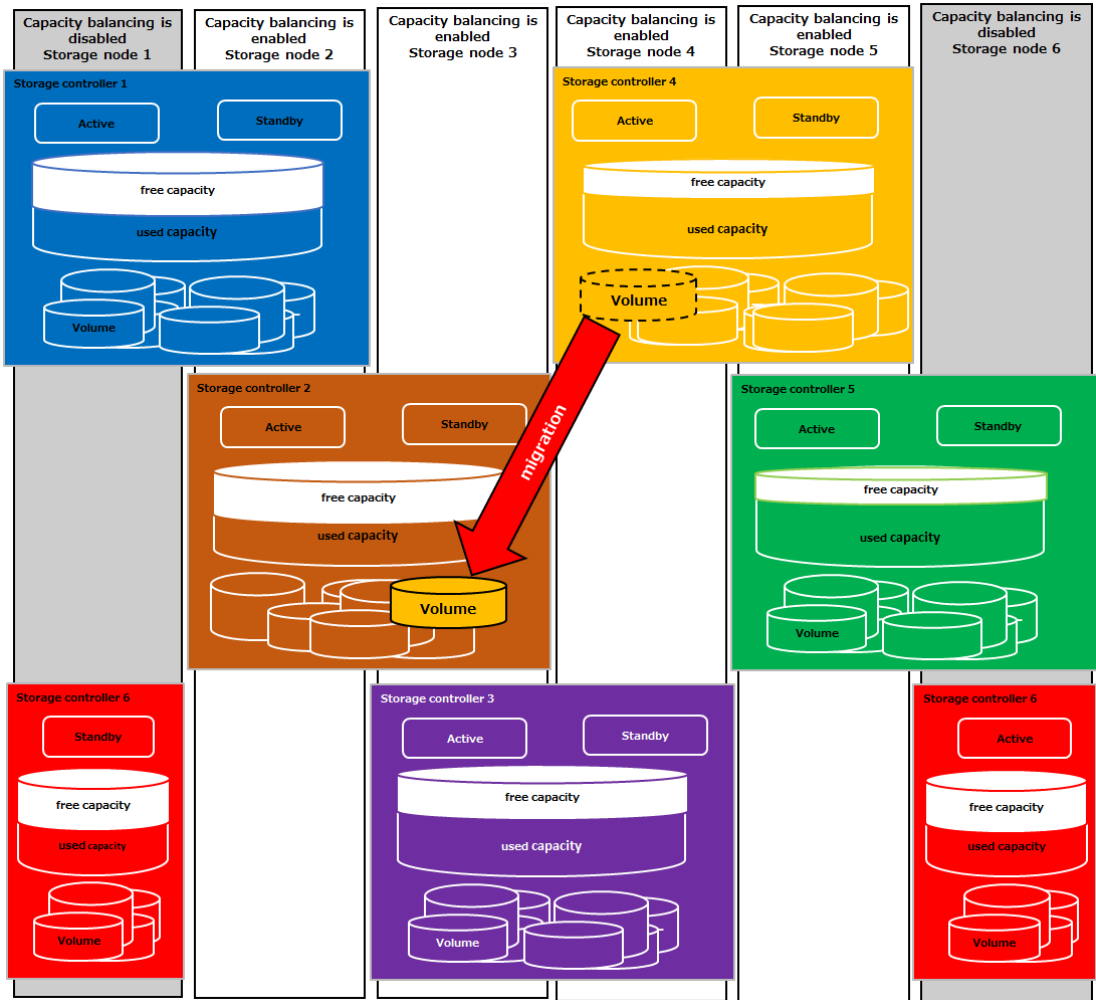
Required role: Security, Storage, Monitor, Service, Audit, or Resource

Procedure

1. Verify the ID of the intended storage node.
REST API: GET /v1/objects/storage-nodes
CLI: storage_node_list
2. Obtain information about the intended storage node.
Run either of the following commands with the storage node ID specified.
REST API: GET /v1/objects/storage-nodes/<id>
CLI: storage_node_show

Overview of capacity management of a storage node

If capacity usage is not balanced among storage controllers, Virtual Storage Software block moves volumes automatically from high capacity usage storage controllers to low capacity usage storage controllers. This is called Capacity balance. Capacity balance is enabled by default. If Capacity balance for a storage node is disabled, volumes cannot be moved to and from the storage controller that runs on the storage node in active-standby mode.



In the figure, note that storage nodes 1 and 6 have the Capacity balance setting disabled. Therefore, volumes managed by storage controllers 1, 5, and 6 are not subject to movement by Capacity balance, and these storage controllers will not be selected as a movement destination.

Also note that capacity usage is high in storage controllers 4 and 5 and low in storage controllers 1, 2, and 6. At this time, Capacity balance is performed on only the storage controllers associated with the storage nodes on which Capacity balance is enabled. As a result, Capacity balance moves volumes from storage controller 4 to storage controller 2.

Operating requirements of Capacity balance

Capacity balance operates when all the following conditions are satisfied:

1. Capacity usage of the storage pool is 80% or less.

2. A storage controller that satisfies all of the following three conditions exists.
 - Capacity usage of the storage controller is 70% or more.
 - The storage controller satisfies one of the following conditions:
 - The storage controller has capacity (including the management area of all the volumes) equal to allocatableCapacity or more.
 - The storage controller has a snapshot volume.
 - Capacity usage of the storage controller is equal to that of the storage pool or more.
3. Capacity balance setting is enabled.

Determining whether to enable or disable the Capacity balance setting

Enable Capacity balance in the following cases:

- You want to automatically manage capacity according to the capacity usage of the storage pool and storage controllers.
- Capacity usage is high so you want to expand capacity (by adding more storage nodes) and, if possible, automatically move the volumes in the existing storage controller.

Disable Capacity balance in the following cases.

When Capacity balance is disabled, monitor the capacity usage per storage controller to prevent capacity depletion.

- You place higher priority on I/O performance than capacity and do not want I/O performance to be degraded due to volume movement performed by Capacity balance.
- You do not want to move volumes because you created them with the intention of fully utilizing their capacity.
- You do not want to move volumes because you created them in the specified storage controller.
- You do not want to move volumes to a newly added storage node.



Note:

- Capacity balancing does not move snapshot volumes (P-VOLs, P/S-VOLs or S-VOLs) among storage nodes.
- The following volumes will not be moved between the storage nodes by Capacity balance:
 - Snapshot volumes (P-VOL, P/S-VOL or S-VOL)
- If the following volume exists, capacity might not be sufficiently balanced by Capacity balance.
 - A volume whose capacity exceeds 20% of allocatableCapacity of the storage controller

The capacity balancing performance and the number of volumes that can be moved in parallel are as follows:

- Capacity balancing performance (approximate performance per volume): 14.7 MiB/s
- Number of volumes that can be moved in parallel (per storage node): 42



Caution:

Capacity balancing among storage nodes might affect the I/O performance of existing volumes.

Editing capacity management settings of a storage node (CLI or REST API)

Enable or disable the Capacity balance setting. Capacity balance is enabled by default.

Before you begin

Required role: Storage

Procedure

1. Verify the ID of the intended storage node.
 - REST API: GET /v1/objects/storage-nodes
 - CLI: storage_node_list

2. Edit the capacity management settings of the storage node.

Run either of the following commands with the storage node ID specified and whether to enable capacity balancing specified.

REST API: PATCH /v1/objects/storage-node-capacity-settings/<id>

CLI: storage_node_capacity_setting_set

Verify the job ID which is displayed after the command is run.

3. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobID>

CLI: job_show

If the job state is "Succeeded", the job is completed.

4. Obtain the capacity management settings for the storage node and verify that the settings have been edited.

Run either of the following commands with the storage node ID specified.

REST API: GET /v1/objects/storage-node-capacity-settings/<id>

CLI: storage_node_capacity_setting_show

5. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Obtaining a list of the capacity management settings of storage nodes (CLI or REST API)

The following information can be obtained.

- id: Storage node ID (uuid)
- capacityBalancingSetting: Capacity balancing settings

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Run either of the following commands to obtain a list of the capacity management settings of storage nodes.

REST API: GET /v1/objects/storage-node-capacity-settings

CLI: storage_node_capacity_setting_list

Obtaining capacity management settings of individual storage nodes (CLI or REST API)

The following information can be obtained.

- id: Storage node ID (uuid)
- capacityBalancingSetting: Capacity balancing settings

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Verify the ID of the intended storage node.

REST API: GET /v1/objects/storage-nodes

CLI: storage_node_list

2. Obtain capacity management settings of the storage node.

Run either of the following commands with the storage node ID specified.

REST API: GET /v1/objects/storage-node-capacity-settings/<id>

CLI: storage_node_capacity_setting_show


Performing maintenance operation for storage nodes


Overview

The following table describes the maintenance operations you can perform on storage nodes.

For details about the procedure flow and prerequisites for operating storage nodes, contact customer support.

This manual describes the procedure that can be performed by using the GUI.

Operation	Window to operate on	Operation icon	Dialog
Performing storage node maintenance recovery	Storage Node list window Storage Node detailed information window	 Select the [Maintenance Recovery] menu.	Maintenance Recovery for Storage Node

Operation	Window to operate on	Operation icon	Dialog
Performing storage node maintenance blocking	Storage Node list window Storage Node detailed information window	 Select the [Maintenance Blockade] menu.	Maintenance Blockade for Storage Node

Performing storage node maintenance recovery (GUI)

When customer support determines that storage node maintenance recovery is required, you can use the GUI the perform maintenance recovery as follows:

Before you begin

- Required role: Service

Procedure

1. Verify the ID and STATUS of the storage node in the Storage Node list window or Storage Node detailed information window.

Verify that STATUS of the storage node to be recovered is "MaintenanceBlockage", "TemporaryBlockage", "PersistentBlockage", "RemovalFailedAndTemporaryBlockage", "RemovalFailedAndMaintenanceBlockage", or "RemovalFailedAndPersistentBlockage", and then go to the next step.

2. Power on the storage node.

- (Virtual machine) In VMware vCenter Server, power on the storage node VM to be recovered.

If the storage node VM you want to recover was active at the time of starting this procedure, shut down the guest OS from VMware vCenter Server or VMware ESXi Client, verify that the storage node VM is stopped, and then power it on again.

If the guest OS do not stop after approximately 5 minutes of the shutdown operation, stop the OS by powering off VMware vCenter Server or VMware ESXi Client, and then power on the storage node VM again.

In VMware vCenter Server, verify that the storage node VM you want to recover is active.

Go to the next step after finishing verifying that the VM (storage node) you want to recover is active.

- (Bare metal) Power on the storage node to be recovered.

If the storage node you want to recover was active at the time of starting this procedure, shut down the storage node by OS shutdown operation. After confirming that the storage node is stopped, power on the storage node again.

If the OS does not stop after approximately 5 minutes of the shutdown operation, stop the storage node by performing forced shutdown, and then power on the storage node again. For details about how to perform forced shutdown, see the hardware documentation.

After powering on the storage node and verifying that it has started, proceed to the next step.



Note:

- (Virtual machine) If the storage node VM cannot be started and you have replaced parts on the target storage node, the replacement parts might be the cause. For the action to be taken for each replacement part, contact customer support. See also the VMware documentation. If the storage node VM does not become active after taking action according to the preceding documentation, contact customer support.
- (Bare metal)
 - If the storage node cannot be started and you have replaced parts on the target storage node, the replacement parts might be the cause.
 - If you are using Hitachi Advanced Server HA820, to shut down the OS, click the **Momentary Press** button in the **Power & Thermal** section in iLO. To shut down the OS forcibly, click the **Press and Hold** button in the **Power & Thermal** section in iLO. For detailed operations, see the hardware documentation.

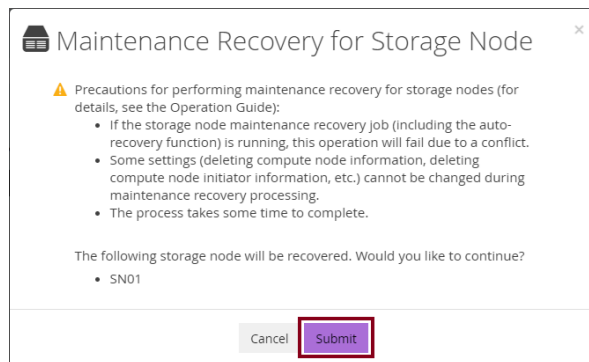
3. Recover the storage node in the Storage Node list window or Storage Node detailed information window by either of the following means.



- Select a storage node to be recovered in the Storage Node list window (one node), click the preceding icon to the right of the Monitor link icon, and then select **Maintenance Recovery**.
- Click the preceding icon displayed for the storage node to be recovered in the Storage Node list window (displayed as inventory view), and then select **Maintenance Recovery**.
- Click the preceding icon in the Storage node detailed information window for the storage node to be edited, and then select **Maintenance Recovery**.

The storage node maintenance recovery dialog is displayed.

4. Verify the displayed information, and then click **Submit**.



The storage node maintenance recovery dialog is closed and storage node maintenance recovery is performed.

5. When the following "Completed" message is displayed, processing is completed.
 - Successfully maintenance recovered for storage node.
6. If you are logged out during recovery processing, verify the Jobs window.

In some cases, storage node maintenance recovery takes a long time, which can cause logout.

In the Jobs window, verify a job whose PATH is `"/v1/objects/storage-nodes/<id>/actions/recover/invoke"`. If the STATUS of the job is "Completed", storage node maintenance recovery is successfully completed.

* `<id>` represents the ID of the target storage node.



Caution:

The Jobs window only shows jobs performed by the user. Therefore, once you are logged out, log in as the same user as when performing storage node maintenance recovery.

7. If you performed hardware replacement between storage node blocking and storage node recovery, verify the following:

Replaced hardware component	Verifying and taking action
(Virtual machine) When the FCHBA was replaced	<p>Verify the Compute Ports list window or Compute Port detailed information window.</p> <p>If PORT SPEED is "LinkDown", the set link speed might not be supported by the FCHBA. Verify the set link speed in CONFIGURED PORT SPEED and whether the replacement FCHBA supports the set link speed.</p> <p>If the FCHBA does not support the set link speed, contact customer support to replace it with an appropriate FCHBA. If PORT SPEED is still "LinkDown" after replacing with an appropriate FCHBA that supports the set link speed, contact customer support.</p>
When the user data drive has been replaced with a drive that is different from the one that was used when the storage node was blocked	Contact. customer support.

8. Back up the configuration information.

For details about the procedure, contact customer support.

Note that, if you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Verifying the conditions for performing storage node maintenance blocking (GUI)

Use the following procedure to make sure that storage node maintenance blocking can be performed.



Caution:

- Performing storage node maintenance blocking lowers the degree of redundancy for the elements that are made redundant, such as the user data, storage controller, or cluster master node. This means that fault tolerance is degraded until the storage node blocked for maintenance is recovered. Therefore, perform only the minimum necessary scope of maintenance blocking.
- Even after the following confirmation procedure has been performed, maintenance blocking might be unsuccessful depending on various conditions. In this case, see the output event logs, and then take appropriate action.

Before you begin

- Required role: Service

Procedure

1. Verify the ID and STATUS of the storage node in the Storage Node list window or Storage Node detailed information window.

If STATUS of the storage node to be blocked for maintenance is "Ready" or "RemovalFailed", go to the next step.

If there are storage nodes in any other statuses, take action as follows:

- If STATUS is "TemporaryBlockage", "MaintenanceBlockage", "PersistentBlockage", "InstallationFailed", "RemovalFailedAndTemporaryBlockage", "RemovalFailedAndMaintenanceBlockage", or "RemovalFailedAndPersistentBlockage", the storage node is blocked properly and separation from the storage cluster is completed. Therefore, you do not have to perform maintenance blocking.

Because of this, if you need to perform maintenance blocking again, you must recover the target storage node first.

- If the GUI displays "Alerting" for "Health status" of a storage node, contact customer support.

However, "Alerting" is also displayed for "Health status" when only a storage node with the status "RemovalFailed" exists. In this case, no action is required.

- If there is a storage node in any status other than those mentioned earlier, a process is being performed for such a storage node. Wait until STATUS changes, and then verify STATUS again.

2. Determine whether Rebuild operation is being performed by confirming the following statuses in the Protection Domain window.

REBUILD STATUS indicates the operation status of the Rebuild operation, and REBUILD PROGRESS RATE indicates the progress rate of the Rebuild operation.

- REBUILD STATUS

- Stopped: Status in which Rebuild operation is not being performed.
- Running: Status in which Rebuild operation is being performed. The Rebuild operation cannot be stopped. Wait until the Rebuild operation is completed, and then verify the Rebuild status in the Protection Domain window again.
- Error: Status in which Rebuild processing cannot be performed due to an error. Verify the event logs and perform troubleshooting.

- REBUILD PROGRESS RATE

Shows the progress rate (%) of Rebuild operation. The progress rate is updated when it fluctuates by 1 point or more. (When progress is made in a short period of time, such as when Fast Rebuild is performed, the progress rate might be updated by several points rather than by 1 point.)

If the Rebuild operation is not being performed, go to the next step.

If the Rebuild operation is being performed, wait until the operation completes. Then, go to the next step.

However, if the storage node to be blocked for maintenance was recovered immediately on a last occasion, you can go to the next step without waiting.

3. Verify that the Drive data relocation operation is not being performed by confirming the following statuses in the Protection Domain window.

DRIVE DATA RELOCATION STATUS shows the Drive data relocation status, and DRIVE DATA RELOCATION PROGRESS RATE shows the progress rate of Drive data relocation processing.

- DRIVE DATA RELOCATION STATUS

- Stopped: Status in which Drive data relocation is not being performed.
- Running: Status in which Drive data relocation is being performed. When you want to suspend the Drive data relocation processing, contact customer support.
- Error: Status in which Drive data relocation resulted in an error, or other conditions are not satisfied, therefore, Drive data relocation cannot be performed.
- Suspended: Status in which Drive data relocation processing is suspended, contact customer support.

- DRIVE DATA RELOCATION PROGRESS RATE

Progress rate (%) is displayed whenever data is transferred for relocation.



Note:

When Drive data relocation is suspended, the progress rate is reset to 0 when resuming. The progress rate will resume from 0. Only data for which processing did not complete before suspension will be processed again.

In addition, when Drive data relocation is suspended, event log KARS07012-I or KARS07013-I is output.

If Drive data relocation is not being performed, go to the next step.

If Drive data relocation is being performed, wait until the operation completes or interrupt Drive data relocation, and then go to the next step.

4. Confirm the user data protection type and user data redundancy in the Storage Pool window.

Verify whether the user data redundancy (DATA REDUNDANCY) satisfies the conditions shown in the following table.

User data protection type (REDUNDANT POLICY)	Condition
4D+1P	The value of DATA REDUNDANCY is 1.
4D+2P	The value of DATA REDUNDANCY is 2.
Duplication	The value of DATA REDUNDANCY is 1.

- If the conditions are met, the storage node can be blocked for maintenance, and the confirmation procedure is complete.

However, if a storage node exists that does not contain a drive for which a storage pool has been expanded, even if it satisfies conditions, you need to go to the next step to confirm.

- If the conditions are not met, go to the next step.

5. Verify the status of the failure in the Storage Node list window and Drives list window.

Verify STATUS of the storage node and drive to see if the applicable conditions are met.

User data protection type (REDUNDANT POLICY)	Number of fault domains	Condition
4D+1P	1	The total number of failure-status storage nodes and failure-status drives is 0. ^{1,2}

User data protection type (REDUNDANT POLICY)	Number of fault domains	Condition
4D+2P	1	The total number of failure-status storage nodes and failure-status drives is 1 or less. ^{1,2}
	3	Either of the following must be met: <ul style="list-style-type: none"> <li data-bbox="1094 562 1414 695">▪ The total number of failure-status storage nodes and failure-status drives is 1 or less.^{1,2} <li data-bbox="1094 730 1414 968">▪ The failure-status storage nodes, failure-status drives, and storage nodes to be blocked for maintenance are all in the same fault domain.¹

User data protection type (REDUNDANT POLICY)	Number of fault domains	Condition
Duplication	1	<p>Either of the following must be met:</p> <ul style="list-style-type: none"> ▪ The total number of failure-status storage nodes and failure-status drives is 0.^{1,2} ▪ The failure-status storage nodes, failure-status drives, and storage nodes to be blocked for maintenance do not span both storage nodes that the redundant storage controllers belong to. Also, the failure-status storage nodes and storage nodes to be blocked for maintenance do not include two or more cluster master nodes in total.^{1,3}

User data protection type (REDUNDANT POLICY)	Number of fault domains	Condition
	3	<p>Either of the following must be met:</p> <ul style="list-style-type: none"> ▪ The total number of failure-status storage nodes and failure-status drives is 0.^{1,2} ▪ The failure-status storage nodes, failure-status drives, and storage nodes to be blocked for maintenance are all in the same fault domain.¹ ▪ The failure-status storage nodes, failure-status drives, and storage nodes to be blocked for maintenance do not span both storage nodes that the redundant storage controllers belong to. Also, the failure-status storage nodes and storage nodes to be blocked for maintenance do not include two or more cluster master nodes in total.^{1,3}
Notes:		

User data protection type (REDUNDANT POLICY)	Number of fault domains	Condition
		<p>1. A failure status refers to the following storage nodes and drives.</p> <ul style="list-style-type: none"> ■ Failure status of storage nodes: STATUS including the following character strings: "TemporaryBlockage", "MaintenanceBlockage", "TemporaryBlockageFailed", "MaintenanceBlockageFailed", "InstallationFailed", "PersistentBlockage", or "RemovalFailed" ■ Failure status of drives: "Blockage" <p>2. The number of failures is counted as one for the following cases:</p> <ul style="list-style-type: none"> ■ A failure-status drive exists in a failure-status storage node. ■ Multiple failure-status drives exist in the same storage node. <p>3. You can use the following command to view the information about the storage controllers:</p> <p>REST API: GET /v1/objects/storage-controllers</p> <p>CLI: storage_controller_list</p>

If the conditions described in the preceding table are met, the procedure for verifying the conditions for performing storage node maintenance blocking is complete.

If the conditions are not met, take the following action. Taking the following action completes the procedure for verifying the conditions for performing storage node maintenance blocking.

- If the GUI displays "Alerting" for "Health status" of a storage node, contact customer support.
- If there are storage nodes whose STATUS is MaintenanceBlockage, perform maintenance recovery for the nodes according to the procedure in *Performing maintenance recovery for storage nodes*.
- In the Protection Domain window, verify the following statuses to make sure that Rebuild operation is not being performed and no error occurred during the Rebuild.

If Rebuild operation is being performed or an error occurred during the Rebuild operation, confirm the following statuses, and then take appropriate action.

- REBUILD STATUS

- Stopped: Status in which Rebuild operation is not being performed.
- Running: Status in which Rebuild operation is being performed. The Rebuild operation cannot be stopped. Wait until the Rebuild operation is completed, and then verify the Rebuild status in the Protection Domain window again.
- Error: Status in which Rebuild processing cannot be performed due to an error. Verify the event logs and perform troubleshooting.

- REBUILD PROGRESS RATE

Shows the progress rate (%) of Rebuild operation. The progress rate is updated when it fluctuates by 1 point or more. (When progress is made in a short period of time, such as when Fast Rebuild is performed, the progress rate might be updated by several points rather than by 1 point.)

Performing storage node maintenance blocking (GUI)

Perform maintenance blocking for storage nodes as follows.

Before performing maintenance blocking by using the GUI, verify that maintenance blocking is possible according to *Verifying the conditions for performing storage node maintenance blocking*.

Before you begin

- Required role: Service

Procedure

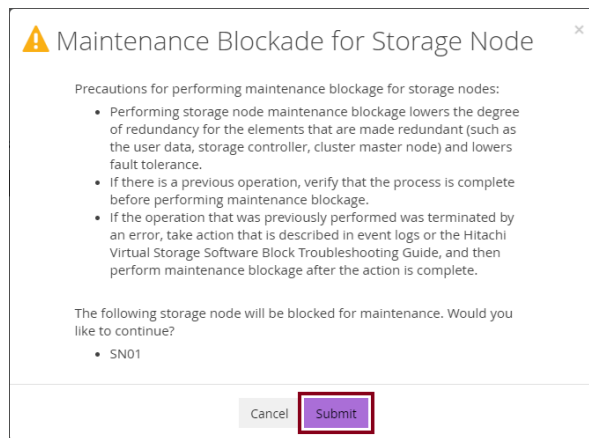
1. Blockade the storage node in the Storage Node list window or Storage Node detailed information window by either of the following means.



- Select a storage node to be blocked in the Storage Node list window (one node), click the preceding icon to the right of the Monitor link icon, and then select **Maintenance Blockade**.
- Click the preceding icon displayed for the storage node to be blocked in the Storage Node list window (displayed as inventory view), and then select **Maintenance Blockade**.
- Click the preceding icon in the Storage node detailed information window for the storage node to be edited, and then select **Maintenance Blockade**.

The storage node maintenance blocking dialog is displayed.

2. Verify the displayed information, and then click **Submit**.



The storage node maintenance blocking dialog is closed and storage node maintenance blocking is performed.

3. When the following "Completed" message is displayed, processing is completed.
 - Successfully maintenance blockaded for storage node.
 - (Virtual machine) The VM in the storage node is turned off.
 - (Bare metal) The storage node is powered off.
4. If you are logged out during blockage processing, verify the Jobs window.

In some cases, storage node maintenance blocking takes a long time, which might cause logout.

In the Jobs window, verify a job whose PATH is "/v1/objects/storage-nodes/<id>/actions/block-for-maintenance/invoke". If the STATUS is "Completed", storage node maintenance blocking is successfully completed.

* <id> represents the ID of the target storage node.



Caution:

The Jobs window only shows jobs performed by the user. Therefore, once you are logged out, log in as the same user as when performing storage node maintenance blocking.

Chapter 6: Verifying the status of drive data relocation

Overview of drive data relocation

If the capacity becomes unbalanced among storage nodes due to storage node addition or removal, drive data relocation runs to optimize the capacity efficiency of each storage node by enabling capacity balancing among storage nodes.

Suspend drive data relocation in the following cases. For how to suspend drive data relocation, see *Suspending drive data relocation*. For how to resume drive data relocation from a suspended state, see *Resuming drive data relocation*.

- To prevent performance degradation due to drive data relocation (resume drive data relocation at night when system load is light)
- To perform other maintenance operations first (resume drive data relocation after other maintenance operations are finished)

Approximate time until drive data relocation is completed

The approximate processing time until drive data relocation is completed can be calculated as follows:

For HPEC 4D + 1P:

Processing time = ((Total physical capacity of drives of expanded storage pool installed in the added or removed storage nodes - 625 [GiB]) / 625 [GiB]) × unit processing time for drive data relocation [min]

For HPEC 4D + 2P:

Processing time = ((Total physical capacity of drives of expanded storage pool installed in the added or removed storage nodes - 900 [GiB]) / 900 [GiB]) × unit processing time for drive data relocation [min]

For Mirroring Duplication:

Processing time = ((A - 200 [GiB]) / 200 [GiB]) × B [min]

(A: Total physical capacity of drives that are added to an expanded storage pool and are installed on the storage node to be added or removed, B: Unit processing time for drive data relocation)

**Caution:**

If physical capacity of a storage node exceeds the capacity that can be allocated to a storage controller, the calculated approximate processing time might be longer. For details about the maximum capacity that can be allocated to a storage controller, contact customer support

Depending on the user data protection method (`redundantType`) settings and resource usage rate of internal processing I/O (`asyncProcessingResourceUsageRate`), unit processing time for drive data relocation is as follows.

The `asyncProcessingResourceUsageRate` setting determines the usage of internal processing I/O (rebuild and drive data relocation) resources. To change the setting, see *Changing usage of internal processing I/O resources*.

<code>redundantType</code>	<code>asyncProcessingResourceUsageRate</code>	Unit processing time for drive data relocation (approximate time)
4D+1P	VeryHigh/High/Middle	100 [min]
	Low	200 [min]
4D+2P	VeryHigh/High/Middle	300 [min]
	Low	500 [min]
Duplication	VeryHigh/High/Middle	11 [min]
	Low	24 [min]

**Note:**

- Setting `asyncProcessingResourceUsageRate` to VeryHigh, High, or Middle makes no difference in the processing time of drive data relocation.
- Processing time for drive data relocation is based on measurement results in a configuration that meets the following prerequisites:
 - With I/O load
 - Number of installed user data drives per storage node: 6
 - The MTU size of the network switch is set to 9000

**Caution:**

Drive data relocation might take longer than the approximate time due to the following factors.

- Network switch performance and I/O load conditions.
- Case in which the number of installed user data drives per storage node is less than 6.

(Due to increased I/O load on each user data drive)

- When the user data protection method is HPEC and storage nodes are removed.

No less than the approximate processing time and no more than approximately the following time are required.

- 4D+1P: approximate processing time × (number of storage nodes before removal / 5)
- 4D+2P: approximate processing time × (number of storage nodes before removal / 6)

Verifying the drive data relocation status (CLI or REST API)

You can verify the status and progress of drive data relocation by obtaining information about the protection domain.

Role-based execution restriction is not subject to restriction.

Procedure

1. Obtain a list of information about the protection domain, and then verify `driveDataRelocationStatus` and `driveDataRelocationProgressRate`.

REST API: GET /v1/objects/protection-domains

CLI: `protection_domain_list`

2. `driveDataRelocationStatus` shows the status of drive data relocation, and `driveDataRelocationProgressRate` shows the progress of drive data relocation.

driveDataRelocationStatus

- Stopped: Status in which drive data relocation is not being performed
- Running: Status in which drive data relocation is being performed. For how to suspend drive data relocation, see *Suspending drive data relocation*.
- Error: Drive data relocation resulted in an error, or the conditions described in *Trigger and operation conditions of drive data relocation* are not satisfied, and therefore, drive data relocation cannot be performed.

- Suspended: Drive data relocation is suspended by the procedure described in *Suspending drive data relocation*.

driveDataRelocationProgressRate

Progress is displayed as a percentage (%) whenever data is transferred for relocation.



Note:

When drive data relocation is canceled, the progress rate is reset to 0 when resuming. The progress rate will resume from 0. Only data for which processing did not complete before suspension will be processed again.

In addition, when drive data relocation is suspended, event log KARS07012-I or KARS07013-I is output.

Suspending drive data relocation (CLI or REST API)

You can suspend drive data relocation as follows.

If drive data relocation is suspended while it is running or stopped, the following occurs:

- If drive data relocation is running, the processing is suspended and stopped.
- If drive data relocation is stopped, drive data relocation after storage node addition is not performed.



Caution:

- If drive data relocation triggered by storage node removal runs, the suspended state of drive data relocation is cleared.
- To suspend drive data relocation triggered by storage node removal, stop storage node removal.
- Do not interrupt drive data relocation triggered by storage node removal.
- If you interrupt drive data relocation triggered by storage node removal, then stop storage node removal.

For details, see *Stopping storage node removal*.

Before you begin

Required role: Storage or Service

Procedure

1. Obtain a list of information about the protection domain, and then verify driveDataRelocationStatus.

REST API: GET /v1/objects/protection-domains

CLI: protection_domain_list

When driveDataRelocationStatus is "Stopped" or "Running", go to the next step.

2. Suspend drive data relocation.

Run either of the following commands with the protection domain ID specified.

REST API: POST /v1/objects/protection-domains/<id>/actions/suspend-drive-data-relocation/invoke

CLI: protection_domain_suspend_drive_data_relocation

Verify the job ID which is displayed after the command is run.

3. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

4. Obtain a list of information about the protection domain, and verify driveDataRelocationStatus.

REST API: GET /v1/objects/protection-domains

CLI: protection_domain_list

Verify that driveDataRelocationStatus is "Suspended".

**Caution:**

If any of the following failures occurs after a request to suspend drive data relocation is executed, drive data relocation might not be suspended even when the job is successful.

- Storage cluster failure
- Cluster master node (primary) failure

In such a case, try to suspend drive data relocation again.

Resuming drive data relocation (CLI or REST API)

You can resume suspended drive data relocation as follows.

Before you begin

Required role: Storage or Service

Procedure

1. Obtain a list of information about the protection domain, and verify the ID of the protection domain and driveDataRelocationStatus.

REST API: GET /v1/objects/protection-domains

CLI: protection_domain_list

When driveDataRelocationStatus is "Suspended", go to the next step.

2. Resume drive data relocation.

Run either of the following commands with the protection domain ID specified.

REST API: POST /v1/objects/protection-domains/<id>/actions/resume-drive-data-relocation/invoke

CLI: protection_domain_resume_drive_data_relocation

Verify the job ID which is displayed after the command is run.

3. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

4. Obtain a list of information about the protection domain, and verify driveDataRelocationStatus.

REST API: GET /v1/objects/protection-domains

CLI: protection_domain_list

Verify that driveDataRelocationStatus is "Stopped" or "Running".

Chapter 7: Managing compute ports

Obtaining a list of information about compute ports (CLI or REST API)

The following information can be obtained.

- id: IDs (uuid) of compute ports
- protocol: Connection protocols of compute ports
- type: Types of compute ports
- nickname: Nicknames of compute ports
- name: WWNs or iSCSI names of compute ports
- configuredPortSpeed: Link speed setting
- portSpeed:
 - (Virtual machine) Always "DependsOnHypervisor"
 - (Bare metal) Actual link speed
- portSpeedDuplex:
 - (Virtual machine) Link speed only (for FC connection configuration) or always "DependsOnHypervisor" (for iSCSI connection configuration)
 - (Bare metal) Actual link speed and duplex settings of the physical port used for communication
- protectionDomainId: IDs (uuid) of protection domains containing compute ports
- storageNodeId: IDs (uuid) of storage nodes containing compute ports
- interfaceName: Interface name
- statusSummary: Summary of statuses of compute ports
- status: Status of each compute port

- fclInformation :
(Virtual machine) FC port information
(Bare metal) Null
- iscsiInformation: iSCSI port information

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Obtain a list of information about compute ports.
REST API: GET /v1/objects/ports
CLI: port_list

Obtaining information about individual compute ports (CLI or REST API)

The following information can be obtained for the compute port with the ID specified.

- id: ID (uuid) of the intended compute port
- protocol: Connection protocol of the intended compute port
- type: Type of the intended compute port
- nickname: Nickname of the intended compute port
- name: WWN or iSCSI name of the intended compute port
- configuredPortSpeed: Link speed setting
- portSpeed:
(Virtual machine) Always "DependsOnHypervisor"
(Bare metal) Actual link speed
- portSpeedDuplex:
(Virtual machine) Link speed only (for FC connection configuration) or always "DependsOnHypervisor" (for iSCSI connection configuration)
(Bare metal) Actual link speed and duplex settings of the physical port used for communication
- protectionDomainId: ID (uuid) of the protection domain containing the intended compute port

- storageNodeId: ID (uuid) of the storage node containing the intended compute port
- interfaceName: Interface name
- statusSummary: Summary of statuses of compute ports
- status: Status of each compute port
- fclInformation:
 - (Virtual machine) FC port information
 - (Bare metal) Null
- iscsiInformation: iSCSI port information

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Verify the ID of the intended compute port.

If you use the CLI to specify a compute port by WWN or iSCSI name, check the WWN or iSCSI name of the compute port.

REST API: GET /v1/objects/ports

CLI: port_list

2. Obtain information about the intended compute port.

Run either of the following commands with the compute port ID specified.

If you use the CLI, you can specify the WWN or iSCSI name instead of the compute port ID.

REST API: GET /v1/objects/ports/<id>


CLI: port_show

Modifying settings of compute ports (CLI or REST API)

Modify the settings of the intended compute port.

Use of an iSNS server

To use an iSNS server, make settings for the iSNS server.



Internet Storage Name Service (iSNS) is a naming service for iSCSI. By using iSNS, information about the initiator and the target (IP address, TCP port number, iSCSI name, etc.) can be discovered without any settings.



Caution:

- For the server address specified when registering an iSNS server, specify a network in the same segment that can be connected to the compute port.
- If two compute ports exist on a storage node, both ports must have the same settings for the iSNS client function (such as enable/disable of the iSNS client function and information of the connection-target iSNS server). If different settings are used, information might not be correctly applied to the iSNS server.
- Even when correct information is set, not all iSCSI targets might be correctly recognized in combinations of control ports and some iSNS servers. In such a case, add individual iSCSI targets manually.

With iSNS enabled in Virtual Storage Software block, the following are supported by connecting with the iSNS server:

(1) Naming service

iSNS registers the initiator/target information on the network with the iSNS database and makes inquiries to the iSNS database at the request of the initiator or target.

(2) SCN (State Change Notification)

An SCN is a message to notify the initiator or target of changes in the iSNS database. When the initiator or target registers an SCN notification request, iSNS sends an SCN to the registered peer.

(3) ESI (Entity Status Inquiry)

iSNS supports ESI (Entity Status Inquiry), a heartbeat feature for the initiator and target. More specifically, iSNS sends an "ESI" to the initiator and target periodically. If no response is received for a certain number of ESIs due to a network failure or other failure, iSNS unregisters the initiator and target from the iSNS database.

When the type of compute ports is "target", all compute ports are registered with the iSNS server.

The following table lists the attributes registered in the iSNS server in the naming service.

Attribute name	Registration required	Attribute value type	Attribute value length	Description
Delimiter	-	Value	0	Delimiter. Used in a delimiter attribute.
Entity Identifier (EID)	Key	Character string	256	ID specifying an entity uniquely. Generated from a MAC address. Unique value for each physical port. The setting cannot be changed. The following defaults are set. iqn.1994-04.jp.co.hitachi:rsd.sph.t.XX XXX.YYY XXXXX: Serial number (0x0 to 0xFFFFF) YYY: Port number (0x0 to 0xFFFF)
Entity Protocol	A	Value	4	Specify "2".
IP address	Key	IP address	16	Specify the IP address of the port.
Portal TCP/UDP Port	Key	Value	4	Specify the port number of the port.
ESI Interval	A	Value	4	ESI issuance interval (in seconds). Set 300 seconds.
ESIPort	Key	Value	4	ESI receive port number. Specify the value obtained when enabling the iSNS service for ISP4022.

Attribute name	Registration required	Attribute value type	Attribute value length	Description
SCNPort	Key	Value	4	SCN receive port number. Specify the value obtained when enabling the iSNS service for ISP4022.
iSCSI Name	Key	Character string	224	Specify the iSCSI name of the iSCSI target.
iSCSI Node Type	A	Value	4	Specify "Target" (0x00000001).
iSCSI Node Alias	Key	Character string	256	Specify the iSCSI alias of the iSCSI target.
iSCSI Node SCN Bitmap	Key	Value	4	Information received via an SCN. Set only bit 25 to ON to receive only the target information.
Key: Registered key attribute, A: Registered attribute				

Before you begin

Required role: Storage

Procedure

1. Verify the ID of the intended compute port.

If you use the CLI to specify a compute port by WWN or iSCSI name, check the WWN or iSCSI name of the compute port.

REST API: GET /v1/objects/ports

CLI: port_list

2. Modify the settings of the intended compute port.

Run either of the following commands with the compute node ID specified.

If you use the CLI, you can specify the WWN or iSCSI name instead of the compute port ID.

Set a nickname and iSCSI name (for iSCSI connection), and then perform settings for FC (network connection) or iSCSI.

In iSNS settings, enable or disable the iSNS client function and perform settings for the iSNS server.

Nickname

- Each compute port must have a unique nickname.
- Conventions to be followed when setting a nickname
 - Number of characters: 1 to 32
 - Characters that can be used: Numbers (0 to 9), uppercase alphabet (A to Z), lowercase alphabet (a to z), symbols (! # \$ % & ' + - . = @ ^ _ { } ~ () [] :)

iSCSI name

- The same iSCSI name as another compute port cannot be set.
- The default name of the compute port added when a storage node is added cannot be set.

Example: iqn.1994-04.jp.co.hitachi:rsd.sph.t.<ID-of-the-storage-cluster>.XXX

Where, <ID-of-the-storage-cluster> is the hexadecimal notation of internalId obtained in *Obtaining information about storage clusters*.

REST API: PATCH /v1/objects/ports/<id>

CLI: port_set

Verify the job ID which is displayed after the command is run.

3. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

4. Obtain information about the intended compute port and verify the settings of the compute port.

Run either of the following commands with the compute port ID specified.

If you use the CLI, you can specify the WWN or iSCSI name instead of the compute port ID.

REST API: GET /v1/objects/ports/<id>

CLI: port_show

5. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Chapter 8: Managing drives

Overview of managing drives

For the drives in a storage node, you can perform the following operations including obtaining information about the drives and changing the configuration of the drives.

Item	Description
Obtaining information about drives	You can obtain information about drives by obtaining a list of drive information or by specifying a specific drive ID.
Adding drives	You can add drives to increase the capacity of storage pools. See <i>Adding drives</i> .
Removing drives	You can remove failed drives.
Replacing drives	You can replace failed drives with other drives.



Note:

To update the firmware on drives, follow the instructions in the manuals of your server vendors.

Obtaining a list of information about drives (CLI or REST API)

The following information can be obtained.

- id: Drive IDs (uuid)
- wwid: Drive WWIDs (WWN)
- statusSummary: Status summary of drives
- status: Status of drives
- typeCode: Codes (model names of drives) that indicate the types of drives
- serialNumber: Serial numbers of drives

- storageNodeId: Storage node IDs (uuid)
- deviceFileName: Names of device files corresponding to drives
- vendorName: Vendor names
- firmwareRevision: Firmware revision numbers
- locatorLedStatus: On/off status of the locator LED
- driveType: Drive types
- driveCapacity: Drive capacity



Note:

The output results of typeCode, serialNumber, and vendorName might include spaces depending on the drive type.

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Obtain a list of information about drives.

REST API: GET /v1/objects/drives

CLI: drive_list

Obtaining information about individual drives (CLI or REST API)

The following information can be obtained.

- id: Drive ID (uuid)
- wwid: Drive WWID (WWN)
- statusSummary: Status summary of the intended drive
- status: Status of the intended drive
- typeCode: Code (model name) that indicates the type of the intended drive
- serialNumber: Serial number of the intended drive
- storageNodeId: Storage node ID (uuid)
- deviceFileName: Name of the device file corresponding to the intended drive

- vendorName: Vendor name
- firmwareRevision: Firmware revision number
- locatorLedStatus: On/off status of the locator LED
- driveType: Type of the intended drive
- driveCapacity: Drive capacity



Note:

The output results of typeCode, serialNumber, and vendorName might include spaces depending on the drive type.

Before you begin

Required role: Security, Storage, Monitor, Service or Resource

Procedure

1. Verify the ID of the intended drive.
REST API: GET /v1/objects/drives
CLI: drive_list
2. Obtain information about the intended drive.
Run either of the following commands with the drive ID specified.
REST API: GET /v1/objects/drives/<id>
CLI: drive_show

Removing drives (CLI or REST API)

Remove a failed drive as follows.

Before you begin

Required role: Storage

Procedure

1. Verify the ID of the faulty drive to be removed and the ID of the storage node that has the faulty drive.

Also record the WWID of the faulty drive to be removed. WWID is used to remove the faulty drive from the server.

Run either of the following commands with status "Blockage" specified in the query parameter.

REST API: GET /v1/objects/drives

CLI: drive_list

2. Verify the status of the storage node containing the faulty drive.

Run either of the following commands with the storage node ID containing the faulty drive specified.

REST API: GET /v1/objects/storage-nodes/<id>

CLI: storage_node_show

If the status of the storage node is "Ready" or "RemovalFailed", go to the next step.

3. Turn on the locator LED of the drive to be removed.

Run either of the following commands with "TurnOn" specified for the operationType parameter (operation_type in the case of CLI).

REST API: POST /v1/objects/drives/<id>/actions/control-locator-led/invoke

CLI: drive_control_locator_led

Verify the job ID which is displayed after the command is run.

**Caution:**

If a storage node failure occurs during a drive removal operation, the locator LED on/off status shown by using the REST API, CLI, or GUI might become different from the on/off status of the locator LED on the physical drive. The locator LED on/off status shown by using the REST API, CLI, or GUI is updated and corrected after the storage node is recovered from the failure.

4. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobID>

CLI: job_show

If the job state is "Succeeded", the job is completed.

5. On the server, find the drive whose locator LED is lit and confirm its mounting position.

Then, remove the faulty drive from the server.

For details, see the documentation of your server vendor.

**Caution:**

If a failure occurs during a drive removal operation, the locator LED might be turned off. In such a case, resume from step 3.

If you interrupt a drive removal operation and perform a maintenance operation that requires the storage node to be restarted, the locator LED might be turned off. In such a case, resume from step 3.

**Note:**

- If the locator LED is not lit, confirm the mounting position of the drive to be removed using the following method. Find the drive that matches the WWID of the failed drive recorded in step 1 and the WWN value of the drive recorded at the time of expansion. Confirm the location in which the drive recorded in association with the WWN is installed.
- The WWN of the drive recorded at the time of addition might differ from the last 1 to 3 digits of the right-side 16-digit part of the WWID recorded in step 1.

6. See *Verifying Rebuild status* and determine whether the Rebuild is being performed or whether an error has occurred during the Rebuild.

If the Rebuild is not being performed and no error has occurred, go to the next step.

If the Rebuild is being performed or an error has occurred during the Rebuild, take appropriate action (see *Verifying Rebuild status*).

7. Remove the faulty drive.

Run either of the following commands with the drive ID obtained in step 1 specified.

REST API: POST /v1/objects/drives/<id>/actions/remove/invoke

CLI: drive_remove

Verify the job ID which is displayed after the command is run.

8. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

9. Obtain a list of drives and make sure that the target faulty drive has been removed.

After step 7, removal of the drive might take approximately one minute.

REST API: GET /v1/objects/drives

CLI: drive_list

10. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Replacing drives (CLI or REST API)

Replace the faulty drive with another drive.

Before you begin

Required role: Storage

Procedure

1. Verify the ID of the faulty drive to be removed and the ID of the storage node that has the faulty drive.

Also record the WWID of the faulty drive to be removed. WWID is used to remove the faulty drive from the server.

Run either of the following commands with status "Blockage" specified in the query parameter.

REST API: GET /v1/objects/drives

CLI: drive_list

2. Verify the status of the storage node containing the faulty drive.

Run either of the following commands with the storage node ID containing the faulty drive specified.

REST API: GET /v1/objects/storage-nodes/<id>

CLI: storage_node_show

Go to the next step when the status of the storage node is "Ready" or "RemovalFailed".

3. Turn on the locator LED of the drive to be removed.

Run either of the following commands with "TurnOn" specified for the operationType parameter (operation_type in the case of CLI).

REST API: POST /v1/objects/drives/<id>/actions/control-locator-led/invoke

CLI: drive_control_locator_led

Verify the job ID which is displayed after the command is run.

**Caution:**

If a storage node failure occurs during a drive replacement operation, the locator LED on/off status shown by using the REST API, CLI, or GUI might become different from the on/off status of the locator LED on the physical drive. The locator LED on/off status shown by using the REST API, CLI, or GUI is updated and corrected after the storage node is recovered from the failure.

4. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobID>

CLI: job_show

If the job state is "Succeeded", the job is completed.

5. On the server, find the drive whose locator LED is lit and confirm its mounting position.

Then, remove the faulty drive from the server.

For details, see the documentation of your server vendor.

**Caution:**

If a failure occurs during a drive replacement operation, the locator LED might be turned off. In such a case, resume from step 3.

If you interrupt a drive replacement operation and perform a maintenance operation that requires the storage node to be restarted, the locator LED might be turned off. In such a case, resume from step 3.

**Note:**

- If the locator LED is not lit, confirm the mounting position of the drive to be removed using the following method. Find the drive that matches the WWID of the failed drive recorded in step 1 and the WWN value of the drive recorded at the time of expansion. Confirm the location in which the drive recorded in association with the WWN is installed.
- The WWN of the drive recorded at the time of addition might differ from the last 1 to 3 digits of the right-side 16-digit part of the WWID recorded in step 1.

6. Perform the steps from *Inserting drives* in *Adding drives* to step 4 of *Expanding storage pools*.

Note that the drive you are adding should be a new drive, not the failed drive that you removed in step 5.

**Caution:**

If you are replacing multiple drives at the same time, perform steps 1 through 6 (physical drive reduction and expansion operations) one drive at a time. When you have completed step 6 on all drives, complete step 7 or later.

7. See *Verifying Rebuild status* and determine whether the Rebuild is being performed or whether an error has occurred during the Rebuild.

If the Rebuild is not being performed and no error has occurred, go to the next step.

If the Rebuild is being performed or an error has occurred during the Rebuild, take appropriate action (see *Verifying Rebuild status*).

8. Remove the faulty drive.

Run either of the following commands with the faulty drive ID obtained in step 1 specified.

REST API: POST /v1/objects/drives/<id>/actions/remove/invoke

CLI: drive_remove

Verify the job ID which is displayed after the command is run.

9. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

10. Obtain a list of drives and make sure that the target faulty drive has been removed.

After step 8, removal of the drive might take approximately one minute.

REST API: GET /v1/objects/drives

CLI: drive_list

11. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Adding drives

Overview of adding drives

You can add drives to increase the capacity of storage pools.

To add storage pool capacity, you need to add more drives to all storage nodes. The amount of drive space required depends on the configuration. See *Capacity design when adding storage nodes* to prepare the required capacity.

Choose a drive to be added. Additional requirements for added drives are as follows.

Item	Requirement	Remarks
Status of the storage node where drives are to be added	"Ready" or "RemovalFailed"	
Maximum number of user data drives (per storage node)	24	It differs depending on the server used as a storage node.
Drive type	SSD (SAS)	
Drive capacity	The required capacity per drive must be satisfied. In HPEC 4D+1P: 249 GiB or more In HPEC 4D+2P: 298 GiB or more In Mirroring Duplication: 201 GiB or more	Even if drive capacity in excess of the following capacity is installed, it is not used by the system. In HPEC 4D+1P: 124 TiB In HPEC 4D+2P: 148 TiB In Mirroring Duplication: 198 TiB




Caution:

- Existing data for the drives you are adding will be lost. Please be careful.
- After drives are added, user data might be automatically moved to the added drives. While user data is being moved, host I/O performance might decrease.
- When spare nodes have already been registered, also add drives for the spare nodes.

The workflow for adding drives is as follows.

Now, if you want to add more drives to multiple storage nodes, repeat *Inserting drives* for the number of storage nodes, review all the drive information that you want to add, and then perform *Expanding storage pool*.



Inserting drives (CLI or REST API)

Insert the drives you want to add into the intended storage node.

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Make a note of the correlation among the installation location, drive SKU, and WWN written on the label of all the actual drives to be added. Refer to this note when removing or replacing drives.
2. Insert the drives you want to add into the intended storage node.

For details, see the server vendor documentation.



Note:

- Mount a user data drive in any drive slot of your choice, making sure that the user data drive is connected to a disk controller that is connected to only user data drives. At this time, if there are multiple disk controllers to which user data drives are connected, connect the same number of user data drives to each disk controller for failure point distribution and load sharing.

3. Obtain a list of drives and verify the added drives.

Use the query parameter to specify the ID of the storage node with the added drives and run either of the following commands.

REST API: GET /v1/objects/drives

CLI: drive_list

It may take up to a minute for the added drives to be verified.

If the added drives cannot be verified, contact customer support.

Expanding a storage pool (CLI or REST API)

Before you begin

Required role: Storage

Procedure

1. Before you add drives, verify the logical capacity, free capacity, and ID of the intended storage pool.

If you use the CLI to specify a storage pool by name, check the name of the storage pool.

REST API: GET /v1/objects/pools

CLI: pool_list

2. If you want to change the rebuild capacity policy and the tolerable number of drive failures for the storage pool, see *Editing storage pool settings* to perform necessary operation.

**Caution:**

- To be able to edit storage pool settings, the degree of redundancy of user data must not be low. If the degree of redundancy of user data is low, perform the recovery procedure, and then edit storage pool settings.
- If you change the rebuild capacity policy after expanding storage pools, sufficient rebuild capacity might not be secured. Therefore, if you want to change the rebuild capacity policy, do so before expanding storage pools.

3. Expand the storage pool.

Run either of the following commands with the ID of the intended storage pool and the IDs of all drives to be added specified.

If you use the CLI, you can specify a name instead of the ID of the storage pool.

REST API: POST /v1/objects/pools/<id>/actions/expand/invoke

CLI: pool_expand

Verify the job ID which is displayed after the command is run.

**Caution:**

To perform other operations, make sure that the storage pool has been expanded as described in step 5 before doing so.

4. Verify the state of the job by specifying the job ID.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

**Note:**

If rebuildStatus is "Error" due to insufficient capacity of the storage controller on the storage node with the new drive, the rebuild process is executed before the capacity allocation process to the controller that is executed next. If capacity to be allocated to a storage controller becomes depleted as a result of this Rebuild processing, capacity will not be allocated to the storage controller.

Also, if the storage pool does not have sufficient rebuild capacity and the added drive capacity is replenished as the rebuild capacity, the internal processing of allocating capacity to the storage controller will not be performed. As a result, storage pool capacity will not be increased, an event log in step 5 will therefore not be output, and the logical and free capacity of the storage pool will be the same before and after this procedure.

However, if the storage pool rebuild capacity policy is set to "Variable", the logical capacity and free capacity of the storage pool might increase even if the internal processing of allocating capacity to the storage controller is not performed.

Subsequently, the capacity allocation process to storage controller, which is an internal process, is executed. Upon completion of these processes, capacity added by storage pool expansion becomes available.

The approximate processing time required for allocating capacity to storage controllers can be calculated as follows:

The processing is executed in each storage controller in parallel. Therefore, the total processing time is almost the same as the processing time for a single storage controller.

- For HPEC 4D+1P:
30 [sec] × (total physical capacity of added drives per storage node / 121 [GiB])
- For HPEC 4D+2P:
30 [sec] × (total physical capacity of added drives per storage node / 145 [GiB])
- For Mirroring Duplication:
30 [sec] × (total physical capacity of added drives per storage node / 194 [GiB])

**Caution:**

If physical capacity of a storage node exceeds the capacity that can be allocated to a storage controller, the calculated approximate processing time might be longer. For details about the maximum capacity that can be allocated to a storage controller, contact customer support

5. Obtain an event log to make sure that the drive specified in step 3 is now available as storage pool capacity.

Wait for the event log KARS16022-I or KARS16081-I to output as many existing storage nodes as possible. Until then, you may not be able to perform other operations.

If as many event logs KARS16022-I or KARS16081-I as the number of storage nodes have been output, the storage pool has been expanded as specified.

However, if a storage node failure occurs until as many KARS16022-I or KARS16081-I event logs as the number of the existing storage nodes are output, the storage pool might not be expanded. In such a case, perform troubleshooting.

If the event log KARS16046-I is output, the logical capacity (user-available capacity) could not be expanded because the capacity of the specified drive is insufficient. To expand the logical capacity, calculate the required capacity according to *Capacity design when adding storage nodes*, and then expand the storage pool according to *Adding drives*.

6. Obtain information about the storage pool that has the newly-added drives.

REST API: GET /v1/objects/pools

CLI: pool_list

Verify that the logical capacity and free capacity of the storage pool has increased as a result of adding the drives.

7. Verify data movement subject to the storage controller.

After expansion of the storage pool, user data is automatically moved so that the user data is stored evenly.

After data moving starts, verify that the status of data moving (dataRebalanceStatus) becomes "Running" for a while, then changes to "Stopped" finally.

You can verify the status of moving data (dataRebalanceStatus) and the progress rate (dataRebalanceProgressRate) by using the following commands:

Verify the status for all the storage controllers.

REST API: GET /v1/objects/storage-controllers

CLI: storage_controller_list

You can skip this confirmation and perform the following step.

**Note:**

- The progress rate (dataRebalanceProgressRate) is the ratio of the number of volumes whose data moving is completed against the total number of volumes subject to the storage controller. Updating the progress rate might take a long time depending on the volume capacity.

You can estimate how long the data moving takes by using the following formula.

$(\text{Used capacity of the storage pool managed by this storage controller}^1 \times \text{total logical capacity before adding storage nodes}^2 / \text{total logical capacity after adding drives}^3) / 42 \text{ [MiB]} \times 2 \text{ [sec]}$

1. The value of 1 is the following that is described in *Managing storage controllers*.

usedCapacity: Consumed capacity of storage pools managed by this storage controller

2. 3. For the values of 2 and 3, see *Capacity design (for HPEC 4D+1P)*, *Capacity design (for HPEC 4D+2P)*, and *Capacity design (for Mirroring)* to calculate the total logical capacity after adding drives.

The required time varies according to the actual storage pool usage.

- If one of the following events occurs during data movement, the processing status (dataRebalanceStatus) might become "Waiting".
 - The storage pool was expanded.
 - The volume was deleted.
 - The capacity of the storage pool managed by the storage controller was depleted.

When the status of the data movement processing becomes "Waiting", the data movement is suppressed until the cause of the delay is resolved and the progress rate (dataRebalanceProgressRate) becomes null. Once the cause of the delay is resolved, the data movement resumes. Progress rate (dataRebalanceProgressRate) is recalculated and displayed from 0 again.

8. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

**Tip:**

A drive whose status is "Offline" is not incorporated into the storage pool. You can verify the drive status from drive information.




Operating drives

Overview

The following table describes the drive operations that you can perform.

For details about the procedure flow and prerequisites for operating drives, contact customer support.

This manual describes the procedure that can be performed by using the GUI.

Operation	Window to operate on	Operation icon	Dialog
Removing drives	Drives list window Drive detailed information window	 Select the [Remove Drives] menu.	Remove Drives
Adding drives	Drives list window Drive detailed information window	 Select the [Add Drives] menu.	Add Drives
Replacing drives	You can replace drives by performing drive removal and addition operations.		
Turning on/off the locator LED	Drives list window Drive detailed information window	 Select the [Turn On/Off Locator LEDs] menu.	Turn On/Off Locator LEDs

Removing drives (GUI)

Drive removal can be performed to remove a faulty drive.

Before you begin

- Required role: Storage

Procedure

1. Note down the WWID of the faulty drive to be removed. This WWID is used to remove the faulty drive from the server.

2. Verify STATUS of the storage node containing the faulty drive in the storage node detailed information window.

If STATUS of the storage node is "Ready" or "RemovalFailed", go to the next step.

3. Turn on the locator LED of the drive to be removed.

For details about the procedure, see *Turning on/off the locator LED*.

4. From the server, locate the drive whose locator LED is turned on, and then confirm the installation location of the drive to be removed.

Remove the faulty drive from the server.

For details, see the server vendor's documentation.



Note:

- If the locator LED cannot be turned on, confirm the installation location of the drive to be removed as follows: Find a drive that has the same WWID as that of the faulty drive recorded in step 1 and the same WWN as that of the drive recorded at the time of addition. Confirm the location in which the drive recorded in association with the WWN is installed.
- The WWN of the drive recorded at the time of addition might differ from the last 1 to 3 digits of the right-side 16-digit part of the WWID recorded in step 1.

5. Determine whether Rebuild operation is being performed or whether an error has occurred during Rebuild operation by confirming the Protection Domain window.

If Rebuild operation is not being performed and no error has occurred, go to the next step.

If Rebuild operation is being performed or an error has occurred during the Rebuild operation, confirm the following status, and then take appropriate action.

- **REBUILD STATUS**

- Stopped: Status in which Rebuild operation is not being performed.
- Running: Status in which Rebuild operation is being performed. The Rebuild operation cannot be stopped. Wait until the Rebuild operation is completed, and then verify the Rebuild status in the Protection Domain window again.
- Error: Status in which Rebuild processing cannot be performed due to an error. Verify the event logs and perform troubleshooting.

- **REBUILD PROGRESS RATE**

Shows the progress rate (%) of Rebuild operation. The progress rate is updated when it fluctuates by 1 point or more. (When progress is made in a short period of time, such as when Fast Rebuild is performed, the progress rate might be updated by several points rather than by 1 point.)

6. Remove the drive in the Drives list window or Drive detailed information window by either of the following means.

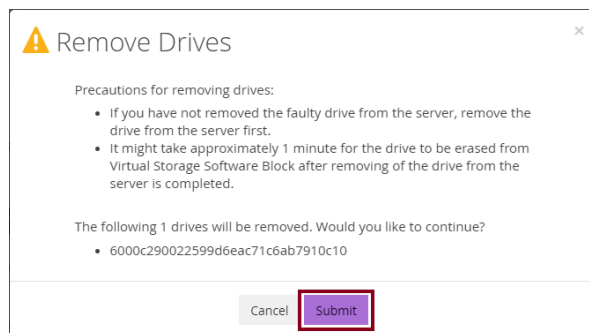
A drive can be removed when STATUS of the drive to be removed is "Blockage".



- In the Drives list window, select the drive to be edited (1 to 25 drives), click the preceding icon to the right of the Monitor link icon, and then select **Remove Drives**.
- Click the preceding icon displayed for the drive to be edited in the Drives list window (displayed as inventory view), and then select **Remove Drives**.
- Click the preceding icon in the Drive detailed information window for the drive to be edited, and then select **Remove Drives**.

The drive removal dialog is displayed.

7. Verify the displayed information, and then click **Submit**.



The drive removal dialog is closed and drive removal is performed.

8. When the following "Completed" message is displayed, processing is completed.
 - Successfully removed drives.
9. In the Drives list window, verify STATUS of the faulty drive to be removed. If STATUS of the drive is "Offline", go to the next step.
10. Back up the configuration information.

For details about the procedure, contact customer support.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.



Note:

The information of a drive removed according to the steps described earlier is deleted from Virtual Storage Software block.

Deletion of drive information from Virtual Storage Software block might take approximately one minute.

Adding drives (GUI)

You can add drives to increase the capacity of storage pools.

Contact customer support of a list of compatible hardware.

Before you begin

- Required role: Storage

Procedure

1. Make a note of the correlation between the location of all the actual drives (to be added), model name of the actual drives, and WWNs written on the labels attached to the actual drives. Refer to this note when removing or replacing drives.
2. Insert the drives you want to add into the intended storage node.

For details, see the server vendor's documentation.

**Note:**

Mount a user data drive in any drive slot of your choice, making sure that the user data drive is connected to a disk controller that is connected to only user data drives. At this time, if there are multiple disk controllers to which user data drives are connected, make sure the number of user data drives connected to each disk controller is equal for failure point distribution and load sharing.

3. Add drives in the Drives list window or Drive detailed information window by either of the following means.


A drive can be added when STATUS of the drive to be added is "Offline".



- In the Drives list window, select the drive to be edited (1 to 1,024 drives), click the preceding icon to the right of the Monitor link icon, and then select **Add Drives**.
- Click the preceding icon displayed for the drive to be edited in the Drives list window (displayed as inventory view), and then select **Add Drives**.
- Click the preceding icon in the Drive detailed information window for the drive to be edited, and then select **Add Drives**.

The drive addition dialog is displayed.

4. Verify the displayed information, and then click **Submit**.



The drive addition dialog is closed and drive addition is performed.

5. When the following "Completed" message is displayed, processing is completed.
 - Successfully added drives.

Confirmation of the added drives might take approximately one minute. If the added drives cannot be verified, contact customer support.

Replacing drives (GUI)

Replace the faulty drive.

Before you begin

- Required role: Storage

Procedure

1. Perform step 1 to step 4 in *Removing drives*.
2. Perform the operation described in *Adding drives*.



Caution:

If you replace multiple drives at the same time, perform the operation for each drive. When step 2 is completed for all the drives, continue the procedure from step 3.

Make sure that the drive to be added is a new drive, not the faulty drive detached in step 1.

3. Perform step 5 and the subsequent steps in *Removing drives*.

Turning on/off the locator LED (GUI)

By turning on the locator LED, the drive to be replaced can be easily recognized.

Before you begin

- Required role: Storage or Service

Procedure

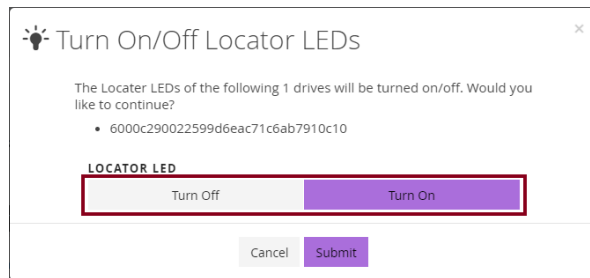
1. Specify the drive in the Drives list window or Drive detailed information window by either of the following means.



- In the Drives list window, select the drive to be edited (1 to 25 drives), click the preceding icon to the right of the Monitor link icon, and then select **Turn On/Off Locator LEDs**.
- In the Drives list window (displayed as inventory view), click the preceding icon displayed for the drive to be edited, and then select **Turn On/Off Locator LEDs**.
- Click the preceding icon in the Drive detailed information window for the drive to be edited, and then select **Turn On/Off Locator LEDs**.

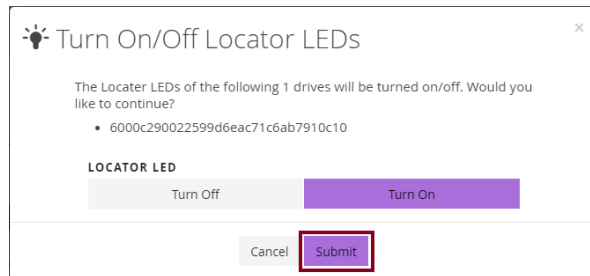
The locator LED on/off dialog is displayed.

2. Specify on or off by clicking **Turn On** or **Turn Off**, respectively.



- TurnOn: Turns on the LED.
- TurnOff: Turns off the LED.

3. Verify the displayed information, and then click **Submit**.



The locator LED on/off dialog is closed and the locator LED is turned on or off.

4. When the following "Completed" message is displayed, processing is completed.
 - Successfully turned on/off locator LEDs.

The Locator LED status is displayed in "LOCATOR LED STATUS" in the Drives list window or Drive detailed information window.

Importing a system requirements file

Overview

For details about the procedure flow and prerequisites for importing a system requirements file, contact customer support.

Importing a system requirements file (GUI)

Import a system requirements file that describes system requirements.

Before you begin

- Required role: Service


Procedure

1. Click the gear icon on the navigation bar, and then select "Import System Requirements File".



The window to import the system requirements file is displayed.

2. Click the **Select** button.




The file selection dialog is displayed.

3. Specify a system requirements file.




Tip:

You can also drag and drop the system requirements file that you want to import into the import dialog.



4. Click **Submit**.



The window to import the system requirements file is closed and import of the system requirements file is performed.

5. When the following "Completed" message is displayed, processing is completed.
 - Successfully imported system requirements file.

Chapter 9: Managing volumes

Creating volumes (CLI or REST API)

In the command used to create volumes, set "capacity" to a value smaller than the value calculated by the following expression:

$\text{logicalLimit [MiB]} - ((168 \text{ [MiB]} \times n) + 42 \text{ [MiB]})$

- n = Logical capacity of a volume to be created [MiB] / 3,145,548 [MiB] (A decimal is rounded up to the nearest integer.)
- To verify the value of logicalLimit, run either of the following commands.

REST API: GET /v1/objects/storage-controllers

CLI: storage_controller_list

A volume requires control information in addition to the logical capacity. For how to calculate the maximum capacity of the storage pool consumed by a volume, see *Maximum storage pool capacity consumed by a volume*.



Note:

The data reduction function is not supported.

Before you begin

- Required role: Storage
- When creating volumes in a VPS: Scope of the VPS

Procedure

1. When you are not using the multi-tenancy function or when you are using the multi-tenancy function and you create volumes in volumes within a system scope, verify the ID of the storage pool in which you want to create volumes.

If you want to specify the storage pool by its name in the CLI, verify the name of the intended storage pool.

REST API: GET /v1/objects/pools

CLI: pool_list

2. When you create volumes in a VPS, verify the creation-destination VPS ID and conditions set for the VPS (upper limit for the number of volumes, upper limit for volume

capacity, upper limit for a single volume capacity, or QoS settings) instead of a storage pool ID.

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list



Note:

A system administrator is not limited to the QoS settings of the VPS, and can configure QoS settings for individual volumes. If QoS parameters are not specified, volumes are created with the QoS parameters set to a VPS.

3. Create volumes.

Run either of the following commands with the parameters for creating volumes specified.

REST API: POST /v1/objects/volumes

CLI: volume_create

Verify the job ID which is displayed after the command is run.



Tip:

- You can set the name and nickname for a volume. The name must be unique in the storage cluster. The nickname does not have to be unique (can also be used for other volumes). For this reason, a name can be used for identifying each individual volume, and a nickname can be used for create a group of volumes.
- By setting the storageControllerId (CLI: --storage_controller_id) parameter, you can create a volume by specifying the storage controller that manages the volume.

4. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

5. Obtain a list of volumes and verify that the volumes are created.

REST API: GET /v1/objects/volumes

CLI: volume_list

6. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

**Note:**

In the following cases, use the function for creating a volume by specifying the storage controller that manages the volume.

- Create a volume by avoiding storage controllers on storage nodes with high I/O loads.
- Create a volume preferentially on the storage controller running on an added storage node.

Also note the following points when creating a volume by specifying the storage controller that will manage the volume.

- When the usage rate of a storage controller becomes high, the capacity balancing process might move the volume to a different storage controller.
- If you want to prevent a volume from being moved from the current controller by Capacity balance, confirm the status of the volumes that are already created on that storage controller, and design the capacity of the storage controller so that its capacity usage is kept less than 70%. If the capacity usage of the storage controller might exceed 70%, consider disabling Capacity balance of the storage node on which the storage controller runs.

However, if the capacity balancing setting is disabled, even if the usage rate of the storage controller becomes high, the capacity will not be balanced and the storage controller capacity might be exhausted. Monitor the usage rate of storage controllers to ensure that they have the necessary free space.

- If you create a maximum number of volumes for a storage controller when total capacity of the volumes managed by the storage controller is smaller than the storage controller capacity, storage pool capacity cannot be fully utilized.

Deleting volumes (CLI or REST API)

Delete unnecessary volumes as follows.

Volumes cannot be deleted in the following cases.

- If a path is set between a compute node and a volume, the volume cannot be deleted.

Before you delete a volume, cancel the connection between the volume to be deleted and the associated compute node.

- Volumes with "P-VOL", "P/S-VOL", or "S-VOL" snapshotAttribute cannot be deleted using the procedure described in this section. For details about how to delete these volumes, see *Deleting snapshots*.
- A volume whose status is "ExpansionFailed" cannot be deleted. If you want to delete a volume whose status is "ExpansionFailed", expand the volume, verify that the status becomes "Normal", and then run the volume deletion command again.

Before you begin

- Required role: Storage
- When deleting volumes in a VPS: Scope of the VPS

Procedure

1. When deleting volumes in a VPS, verify the VPS ID.
If you want to specify a VPS by its name in the CLI, verify the VPS name.
REST API: GET /v1/objects/virtual-private-storages
CLI: vps_list
2. Verify the ID of the intended volume.
If you use the CLI to specify a volume by name, check the name of the volume.
REST API: GET /v1/objects/volumes
CLI: volume_list
3. Delete the intended volume.
Run either of the following commands with the volume ID specified.
If you use the CLI, you can specify a name instead of the ID of the volume.
REST API: DELETE /v1/objects/volumes/<id>
CLI: volume_delete
Verify the job ID which is displayed after the command is run.
4. Verify the state of the job.
Run either of the following commands with the job ID specified.
REST API: GET /v1/objects/jobs/<jobId>
CLI: job_show
If the job state is "Succeeded", the job is completed.
5. Obtain a list of volumes and verify that the volumes are deleted.
REST API: GET /v1/objects/volumes
CLI: volume_list
6. Back up the configuration information.
Perform this step by referring to *Backing up the configuration information*.
If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Expanding volumes (CLI or REST API)

You can only expand a volume whose status is "Normal" or "ExpansionFailed", and "volumeType" is "Normal".

- If you expand a volume whose "status" is "Normal", the volume is expanded by the capacity specified in the parameter "additionalCapacity".
- If you expand a volume whose status is "ExpansionFailed", the volume expansion is performed again.

In a command used to expand volumes, set "additionalCapacity" to a value smaller than the value calculated by the following expression:

logicalLimit [MiB] - ((168 [MiB] × n) + 42 [MiB])

- n = Logical capacity of a volume to be added to the expanded volume [MiB] / 3,145,548 [MiB] (A decimal is rounded up to the nearest integer.)



Tip:

To verify the value of logicalLimit, run either of the following commands.

REST API: GET /v1/objects/storage-controllers

CLI: storage_controller_list



Caution:

If the event log KARS06170-C is issued and no actions are taken, make sure to take actions for KARS06170-C before the operations on the volume.

Before you begin

- Required role: Storage
- When expanding volumes in a VPS: Scope of the VPS

Procedure

1. When expanding volumes in a VPS, verify the VPS ID and conditions set for the VPS (upper limit for volume capacity and upper limit for a single volume capacity).

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list

2. Check the ID of the volume you want to extend.

If you use the CLI to specify a volume by name, check the name of the volume.

It also checks the status and volumeType of the volume to expand the capacity of the volume.

REST API: GET /v1/objects/volumes

CLI: volume_list

When "status" of the volume to be expanded is "Normal" or "ExpansionFailed", and "volumeType" is "Normal", go to the next step.

3. Expand the volume capacity.

Run one of the following commands with the capacity to be added to the volume specified for additionalCapacity. Run the command with a value specified if the status of the volume is "Normal", or with no value specified if the status is "ExpansionFailed".

Run either of the following commands with the volume ID specified.

If you use the CLI, you can specify a name instead of the ID of the volume.

REST API: POST /v1/objects/volumes/<id>/actions/expand/invoke

CLI: volume_expand

Verify the job ID which is displayed after the command is run.

4. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

5. Obtain volume information to verify that the volume capacity is expanded.

Run either of the following commands with the volume ID specified.

If you use the CLI, you can specify the name of the volume instead of its ID.

REST API: GET /v1/objects/volumes/<id>

CLI: volume_show

6. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Obtaining a list of volumes (CLI or REST API)

The following information can be obtained.

**Note:**

The data reduction function is not supported.

- savingEffects: Effect of the data reduction function
- id: IDs (uuid) of volumes
- name: Names of volumes
- nickname: Nicknames of volumes
- volumeNumber: Numbers of volumes
- poolId: IDs (uuid) of storage pools
- poolName: Names of storage pools
- totalCapacity: Total logical capacity
- usedCapacity: Consumed logical capacity
- numberOfConnectingServers: Number of connected compute nodes
- numberOfSnapshots: Number of snapshots
- protectionDomainId: IDs (uuid) of protection domains containing volumes
- fullAllocated: Whether all the area where user data is written is allocated in advance
- volumeType: List of volume types (attributes)
- statusSummary: Summary of statuses of volumes
- status: Status of each volume
- storageControllerId: IDs (uuid) of storage controllers managing volumes
- snapshotAttribute: Attributes of snapshots
- snapshotStatus: Statuses of snapshots
- savingSetting: Setting of the data reduction function
- savingMode: Processing mode of the data reduction function
- dataReductionStatus: Status of the data reduction function
- dataReductionProgressRate: Progress rate of the data reduction function
- vpsId: ID of the VPS to which volumes belong

- vpsName: Name of the VPS to which volumes belong
- qosParam: QoS-related parameter
 - upperLimitForIops: Upper limit for volume performance (IOPS)
 - upperLimitForTransferRate: Upper limit for volume performance (MiB/s)
 - upperAlertAllowableTime: Alert threshold value for the upper limit of volume performance (second)
 - upperAlertTime: The last time the upper limit of volume performance was continuously exceeded and the conditions for the alert threshold of the performance upper limit were met (UTC)
- naaid: NAA ID of the VPS to which volumes belong

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Obtain a list of volumes.
 - REST API: GET /v1/objects/volumes
 - CLI: volume_list

Obtaining information about individual volumes (CLI or REST API)

The following information can be obtained for the volume with the ID specified.



Note:

The data reduction function is not supported.

- reservedCapacity: Reserved logical capacity
- freeCapacity: Free logical capacity
- luns: List of volume LUNs
- snapshotProgressRate: Progress rate of preparation for, deleting, and restoring snapshot volumes
- snapshotTimestamp: For S-VOL or P/S-VOL, time when recording of difference data started for the P-VOL at the time of snapshot creation
- snapshotType: Snapshot type ("Snapshot" for S-VOL or P/S-VOL, null for others)

- `savingEffects`: Effect of the data reduction function
- `snapshotConcordanceRate`: Match rate between the target volume and one newer generation of S-VOL or copy-source volume
- `isWrittenInSvol`: Whether the volume can be written from the controller node
- `id`: ID (uuid) of the intended volume
- `name`: Name of the intended volume
- `nickname`: Nickname of the intended volume
- `volumeNumber`: Number of the intended volume
- `poolId`: ID (uuid) of the storage pool
- `poolName`: Name of the storage pool
- `totalCapacity`: Total logical capacity
- `usedCapacity`: Consumed logical capacity
- `numberOfConnectingServers`: Number of connected compute nodes
- `numberOfSnapshots`: Number of snapshots
- `protectionDomainId`: ID (uuid) of the protection domain containing the intended volume
- `fullAllocated`: Whether all the area where user data is written is allocated in advance
- `volumeType`: List of volume types (attributes)
- `statusSummary`: Summary of statuses of volumes
- `status`: Status of each volume
- `storageControllerId`: ID (uuid) of the storage controller that manages this volume.
- `snapshotAttribute`: Attributes of snapshots
- `snapshotStatus`: Statuses of snapshots
- `savingSetting`: Setting of the data reduction function
- `savingMode`: Processing mode of the data reduction function
- `dataReductionStatus`: Status of the data reduction function
- `dataReductionProgressRate`: Progress rate of the data reduction function
- `vpsId`: ID of the VPS to which the intended volume belongs

- vpsName: Name of the VPS to which the intended volume belongs
- qosParam: QoS-related parameter
 - upperLimitForIops: Upper limit for volume performance (IOPS)
 - upperLimitForTransferRate: Upper limit for volume performance (MiB/s)
 - upperAlertAllowableTime: Alert threshold value for the upper limit of volume performance (second)
 - upperAlertTime: The last time the upper limit of volume performance was continuously exceeded and the conditions for the alert threshold of the performance upper limit were met (UTC)
- naaid: NAA ID of the VPS to which the intended volume belongs

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Verify the ID of the intended volume.
If you use the CLI to specify a volume by name, check the name of the volume.
REST API: GET /v1/objects/volumes
CLI: volume_list
2. Obtain information about the intended volume.
Run either of the following commands with the volume ID specified.
If you use the CLI, you can specify a name instead of the ID of the volume.
REST API: GET /v1/objects/volumes/<id>
CLI: volume_show

Modifying settings of individual volumes (CLI or REST API)

You can edit volume names, nicknames, and QoS-related parameters (upperLimitForIops, upperLimitForTransferRate, and upperAlertAllowableTime).

You can only edit settings of volumes whose status is "Normal" or "UpdateFailed". For the volumes with the "UpdateFailed" status, re-edit volume settings without specifying parameters corresponding to the editing target.

It is possible to use the GUI for editing names or nicknames only of normal volumes.

When you perform editing, note the following points:

- Name can only be set to a unique value in the storage cluster.
- Nickname can have duplicate values between volumes.
- If you specify multiple parameters in step 3, you can specify them only in the following combinations: If you want to make a configuration change that cannot be combined, run the command again.
 - Combination of name and nickname
 - Combination of upperLimitForlops, upperLimitForTransferRate, and upperAlertAllowableTime
- You can edit settings of snapshot volumes similarly.



Caution:

If the event log KARS06170-C is issued and no actions are taken, make sure to take actions for KARS06170-C before the operations on the volume.

Before you begin

- Required role: Storage
- When editing volumes in a VPS: Scope of the VPS

Procedure

1. When editing settings of volumes in a VPS, verify the VPS ID and conditions set for the VPS (QoS settings).

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list



Note:

A system administrator can make QoS settings for individual volumes without limitation for the QoS settings of volumes.

2. Verify the ID of the intended volume.

If you use the CLI to specify a volume by name, check the name of the volume.

REST API: GET /v1/objects/volumes

CLI: volume_list

3. Change the settings of the intended volume.

Run either of the following commands with the volume ID and parameters for setting volumes specified.

If you use the CLI, you can specify a name instead of the ID of the volume.

REST API: PATCH /v1/objects/volumes/<id>

CLI: volume_set

Verify the job ID which is displayed after the command is run.

4. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

5. Obtain information about the intended volume and verify that the settings of the volume are modified.

Run either of the following commands with the volume ID specified.

If you use the CLI, you can specify a name instead of the ID of the volume.

REST API: GET /v1/objects/volumes/<id>

CLI: volume_show

6. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Managing volumes

Creating volumes (GUI)

Create volumes as follows.

Before you begin

- Required role: Storage


Procedure

1. Open the Volumes window or Storage Pool window, and click either of the icons shown below.

For List view:



For Inventory view:



The following dialog appears.

2. Enter each of the following parameters:
 - **CAPACITY:** Logical capacity of the volume and its unit
 - **NUMBER OF VOLUMES:** The number of volumes to be created
 - **VOLUME NAME:** The name of the volume
 - **SUFFIX START NUMBER:** The first sequential number suffixed to a volume name or nickname, in the case multiple volumes are created with the same volume name or nickname. If omitted, no number is added.
 - **NUMBER OF DIGITS:** The number of digits of a number suffixed to a name or nickname
 - **VOLUME NICKNAME:** The nickname of the volume. If omitted, **VOLUME NAME** is used.
3. Click **Submit**.
4. When the following "Completed" message is displayed, processing is completed.
 - Successfully created volumes.

Modifying settings of individual volumes (GUI)

Edit the settings of a volume as follows. You can edit the name and nickname of the volume.

Before you begin

- Required role: Storage

Procedure

1. In the Volumes window, the Volume detailed information window, the Compute Node detailed information window, or the Storage Pool window, edit the volume by using one of the following methods:



- In the Volumes window, select the editing-target volume (one volume), and then click the preceding icon shown to the right of Select All.
- Change the view mode of the Volumes window to Inventory view, and then click the preceding icon shown for the editing-target volume.
- In the Volume detailed information window for the editing-target volume, click the preceding icon.
- In the Compute Node detailed information window, select the editing-target volume (one volume), and then click the preceding icon shown to the right of Select All.
- Change the view mode of the Compute Node detailed information window to Inventory view, and then, in the list of volumes, click the preceding icon shown for the editing-target volume.
- In the Storage Pool window, select the editing-target volume (one volume), and then click the preceding icon shown to the right of Select All.
- Change the view mode of the Storage Pool window to Inventory view, and then, in the list of volumes, click the preceding icon shown for the editing-target volume.

The following dialog appears.

2. Enter each of the following parameters:
 - VOLUME NAME: A new name to be set for a volume
 - VOLUME NICKNAME: A new nickname to be set for a volume
3. Click **Submit**.
4. When the following "Completed" message is displayed, processing is completed.
 - Successfully edited volume name.

Expanding volumes (GUI)

Expand the volume capacity as follows.



Caution:

When you expand volumes by specifying [Total Capacity], make sure that no other user expands volumes by specifying [Total Capacity] at the same time.

Otherwise, the volume might not be expanded to the specified volume.

Before you begin

- Required role: Storage

Procedure

1. In the Volumes window, the Volume detailed information window, or the Compute Node detailed information window, expand volumes by using one of the following methods:



- In the Volumes window, select the expansion-target volumes (1 to 25 volumes), and then click the preceding icon shown to the right of Select All.
- Change the view mode of the Volumes window to Inventory view, and then click the preceding icon shown for each expansion-target volume.
- In the Volume detailed information window for each expansion-target volume, click the preceding icon.
- In the Compute Node detailed information window, select the expansion-target volumes (1 to 25 volumes), and then click the preceding icon shown to the right of Select All.
- Change the view mode of the Compute Node detailed information window to Inventory view, and then, in the list of volumes, click the preceding icon shown for each expansion-target volume.

The following dialogs appears.

When [SPECIFY CAPACITY BY] is set to "Additional Capacity":

The screenshot shows the 'Expand Volumes' dialog box. Under 'SPECIFY CAPACITY BY', the 'Additional Capacity' option is selected. The 'ADDITIONAL CAPACITY' field is empty, and the unit is set to 'GIB'. The 'PREVIEW' section shows a volume named 'VolumeName000001' with 42% usage and a total capacity of 2.50 TiB + 0 B = 2.50 TiB. Buttons for 'Cancel' and 'Submit' are at the bottom.

When [SPECIFY CAPACITY BY] is set to "Total Capacity":

The screenshot shows the 'Expand Volumes' dialog box. Under 'SPECIFY CAPACITY BY', the 'Total Capacity' option is selected. The 'TOTAL CAPACITY' field is empty, and the unit is set to 'GIB'. The 'PREVIEW' section shows a volume named 'VolumeName000001' with 42% usage and a total capacity of 2.50 TiB + 0 B = 2.50 TiB. Buttons for 'Cancel' and 'Submit' are at the bottom.

2. Enter each of the following parameters:
 - SPECIFY CAPACITY BY: Select how you specify capacity expansion.
 - ADDITIONAL CAPACITY: When [SPECIFY CAPACITY BY] is set to "Additional Capacity", specify the logical capacity to be added to the volume.
 - TOTAL CAPACITY: When [SPECIFY CAPACITY BY] is set to "Total Capacity", specify the capacity after expansion.
3. Click **Submit**.
4. When the following "Completed" message is displayed, processing is completed.
 - Successfully expanded volumes.

Deleting volumes (GUI)

Delete unnecessary volumes as follows.

Before you begin

- Required role: Storage

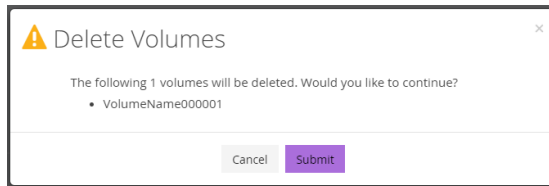
Procedure

1. In the Volumes window, the Volume detailed information window, or the Storage Pool window, delete volumes by using one of the following methods:



- In the Volumes window, select the deletion-target volumes (1 to 25 volumes), and then click the preceding icon shown to the right of Select All.
- Change the view mode of the Volumes window to Inventory view, and then click the preceding icon shown for each deletion-target volume.
- In the Volume detailed information window for each deletion-target volume, click the preceding icon.
- In the Storage Pool window, select the deletion-target volumes (1 to 25 volumes), and then click the preceding icon shown to the right of Select All.
- Change the view mode of the Storage Pool window to Inventory view, and then, in the list of volumes, click the preceding icon shown for each deletion-target volume.

The following dialog appears.



2. Click **Submit**.
3. When the following "Completed" message is displayed, processing is completed.
 - Successfully deleted volumes.

Chapter 10: Backing up and restoring volumes by using snapshots

Overview of snapshots

A duplicate image of a volume at a certain point in time is referred to as a snapshot. If you unintentionally update copy-source volume data, you can restore the data from a previously taken snapshot. The copy-source volume is referred to as the primary volume (P-VOL), and the copy-destination volume is referred to as the secondary volume (S-VOL). When the P-VOL is updated, only the data existing before updating in the update-target area is copied to the storage pool as data in snapshot volumes.


By storing only difference data between P-VOL and S-VOL, this scheme reduces capacity and allows a cost-efficient copy to be created. Furthermore, because data is shared between S-VOLs, less capacity is used even if multiple generations of snapshots are taken.

Because the storage pool is shared, there is no need to create a dedicated storage pool for snapshots.

**Caution:**


- You should back up any necessary data to another device (for example, via the host) with the following in mind. Also, monitor the used capacity of the storage pool to secure sufficient free space.
 - The S-VOL uses the data in the P-VOL. Therefore, if a failure occurs in the P-VOL and the data of the P-VOL becomes unavailable, the data of the S-VOL also becomes unavailable.
 - Difference data is stored in the storage pool. If the storage pool becomes unavailable due to a failure or the storage pool becoming full, data in snapshot volumes will be lost.
 - If data in snapshot volumes is lost during restore (that is, during execution of an operation to restore P-VOL data from S-VOL data), data in the restore-destination volume (P-VOL) also becomes unavailable.
 - If event log KARS06170-C is issued, the snapshot volume status related to the target volume (P-VOL) status changes to "IOSuppressed" and I/O operations to the volume are suppressed. Delete all the snapshot volumes created from the target volume (P-VOL). Then, delete the target volume (normal volume at this point). I/O operations are not possible even when the target volume becomes normal. (The volume status remains "IOSuppressed".)
- If snapshot volumes exist when creating configuration backup files according to *Performing the configuration backup operation*, a system administrator or users need to delete snapshot volumes and volume paths whose configuration was restored by restoring storage cluster configuration.

A snapshot configuration consisting of the copy source and copy destinations as follows is referred to as a snapshot tree. A different generation of the S-VOL is created by the time at which a snapshot is taken.



Another snapshot can be created by using an already created snapshot as the copy source. Another snapshot can be created by using an already created snapshot as the copy source. This snapshot configuration is called a cascade configuration. A cascade configuration can comprise of a maximum of 64 layers. In a cascade configuration, a volume can have both the P-VOL and S-VOL attributes. Such a volume is called a P/S-VOL.

A cascade configuration with a maximum of 64 layers can be created under the P-VOL layer. For example, the cascade configuration in the following figure has three layers.



The following table lists the operations on snapshots.

Operation	Description	See:
Prepare to take a snapshot	Create a volume (S-VOL) for taking snapshots. When you prepare to take a snapshot, a volume (S-VOL) will be created. A snapshot cannot be taken without creating a volume (S-VOL).	<i>Preparing to take a snapshot</i>
Take a snapshot	Take a snapshot for the prepared S-VOL.	<i>Taking a snapshot</i>
Delete snapshots	Delete snapshot volumes that are no longer needed.	<i>Deleting snapshots</i>
Restore P-VOLs from a snapshot	Restore P-VOL data from S-VOL data.	<i>Restoring P-VOLs from a snapshot</i>
Obtain P-VOL information for the volume	Obtain information about the P-VOL from which S-VOL was created.	<i>Obtaining P-VOL information for the volume</i>
Obtain a list of S-VOL information for the volume	Obtain a list of S-VOL information created from P-VOL.	<i>Obtaining a list of S-VOL information for the volume</i>

Operation	Description	See:
Modify settings of individual volumes	Edit the settings for the volume.	<i>Modifying settings of individual volumes</i>

Snapshot specifications

The following table lists the specifications for snapshots.

Item	Specifications
Number of S-VOLs (per P-VOL)	1,023
Number of volumes (per storage node)	Total number of P-VOLs, P/S-VOLs, S-VOLs, and ordinary volumes 1,024
Number of volumes (per system)	<ul style="list-style-type: none"> ▪ Plus-1 redundant configuration: Total number of P-VOLs, P/S-VOLs, S-VOLs, and ordinary volumes $1,024 \times n$ (number of storage nodes) ▪ Plus-2 redundant configuration: When the number of P-VOLs, number of P/S-VOLs, number of S-VOLs, and number of ordinary volumes are summed up, the smaller of the following applies: <ul style="list-style-type: none"> ▪ $1,024 \times n$ (number of storage nodes) ▪ 16,384
Maximum total volume capacity when snapshots are used (per storage controller)	Approx. 4 PiB If you prepare to take snapshots of approximately 4 PiB or larger, the operation might fail with KARS06002-E.
Maximum number of simultaneous operations of volume operations and snapshot operations	Approx. 20 If you perform more than 20 operations at the same time, operations might not be performed successfully due to a timeout. Volume operations refer to volume creation, deletion, expansion, and setting edition. Snapshot operations refer to preparation to take snapshots, taking, deleting and restoring snapshots.

Item	Specifications
	<p>A single operation handling multiple volumes is counted as one operation. For example, a volume creation operation with 20 or more volumes specified is counted as one operation.</p> <p>Completion of the job means completion of the operation. When the job state is "Succeeded", the operation is complete.</p>
Snapshot volume name	<p>Number of characters: 1 to 32</p> <p>Characters that can be used: Numbers (0 to 9), uppercase alphabet (A to Z), lowercase alphabet (a to z), symbols (-, ., : @ _)</p>
Maximum number of layers in a cascade configuration	64

Snapshot-specific information obtained from volume information

You can obtain the following snapshot-specific information by obtaining a list of snapshot information.

Type of information	Description
snapshotStatus	<p>Snapshot status indication: Normal, Empty (empty status without snapshot meta data)*, Deleting, Restoring, Preparing (to take snapshots), Prepared, or Error An ordinary volume is indicated by "-".</p> <p>* This is a temporary status when creating or deleting S-VOLs. When you create S-VOLs, the status transitions to Preparing after a while, and the applicable volumes are deleted when you delete S-VOLs.</p> <p>You can also obtain snapshotStatus from a list of volume information.</p>
snapshotProgressRate	<p>Progress rate of snapshot operation. Indicated only when snapshotStatus is Preparing, Restoring or Deleting. In the case of other statuses, null is output.</p> <p>snapshotProgressRate is not displayed at the same time as the match rate (snapshotConcordanceRate).</p>

Type of information	Description
	You can also obtain snapshotProgressRate from volume information.
snapshotTimestamp	<p>Time when the snapshot is taken.</p> <p>This is displayed when snapshotStatus is "Normal". In the case of other statuses, null is output.</p> <p>You can also obtain snapshotTimestamp from volume information.</p>
snapshotType	<p>"Snapshot" is displayed for a snapshot volume.</p> <p>You can also obtain snapshotType from volume information.</p>
snapshotConcordanceRate	<p>Match rate (%) between this volume and the comparison-target volume.</p> <p>When this volume is the latest generation S-VOL, the volume to be compared will be the P-VOL. When this volume is not the latest generation S-VOL, the volume to be compared will be an S-VOL of a newer generation. For example, if the P-VOL is written in a multi-generation configuration (1:N configuration for a P-VOL and S-VOL), only the latest generation of S-VOL will have a lower match rate. If the S-VOL is written, the match rate will decrease only for the S-VOL of the written generation.</p> <p>The match rate is an approximate value and might have a large margin of error.</p> <p>Differential data indicated by the match rate is managed in units of 256 KiB.</p> <p>When restoring and deleting snapshots, the larger the amount of data difference between P-VOL and S-VOL is, the longer the process will take.</p> <p>This information can be used to estimate the processing time for snapshot restoration and deletion in advance. *</p> <p>You can also obtain snapshotConcordanceRate from volume information.</p>
isWrittenInSvol	This information shows whether the host has written to the S-VOL.

Type of information	Description
	This is displayed when snapshotStatus is "Normal". In the case of other statuses, null is output. You can also obtain isWrittenInSvol from volume information.
* See <i>Processing time for snapshot operations</i> in this manual.	

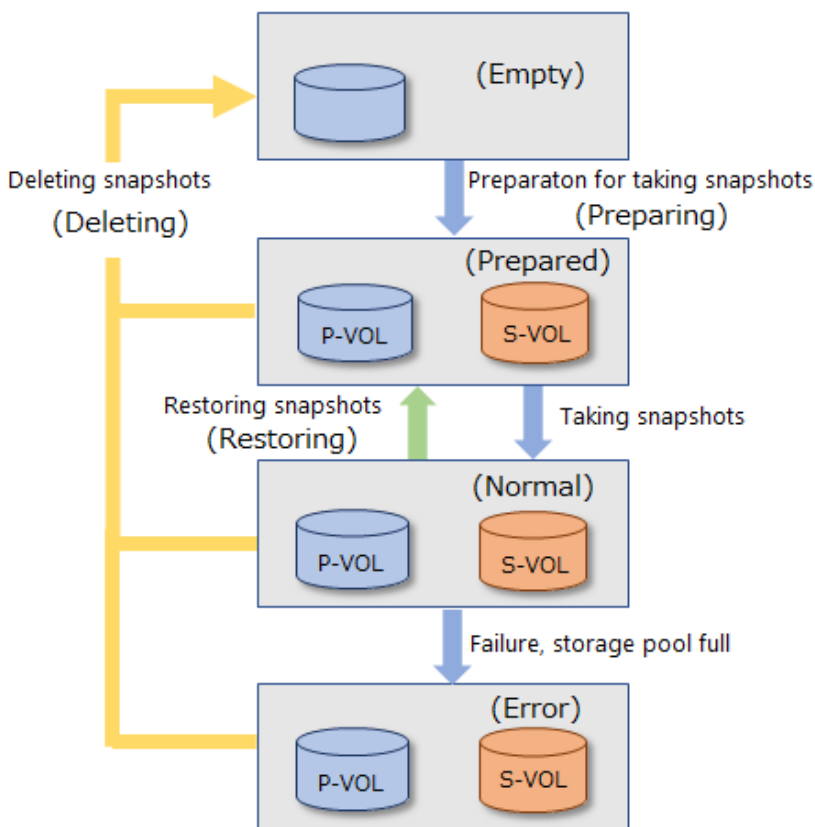
Status transition (snapshot operation and snapshotStatus)

The snapshot Status transits depending on the snapshot operation.



Note:

It might take about one minute for snapshotStatus to transit. To confirm the transition of snapshotStatus, wait about one minute and then check again.



**Tip:**

Whether the volume is a snapshot volume can be found in snapshotAttribute of volume information. "P-VOL", "P/S-VOL", or "S-VOL" is indicated. An ordinary volume is indicated by "-".

Also, the number of S-VOLs in the P-VOL or P/S-VOL can be found in numberOfSnapshots in volume information.

Whether the P-VOL and S-VOL can be written or read

The P-VOL is readable/writable regardless of snapshotStatus of the copy-destination volume.

The S-VOL is readable/writable only when snapshotStatus is "Normal".

Maximum pool capacity consumed by snapshots

You can calculate the maximum pool capacity that is consumed by snapshots for each P-VOL by using the following formula. To calculate the total maximum capacity, calculate and sum the capacity for all P-VOLs.

Maximum consumed capacity of a storage pool for each P-VOL [MiB] = (1) × 42 + (2) × 175,434

- (1): Result of the following calculation rounded up to the nearest integer
(P-VOL capacity [MiB] × number of snapshots × 2 / 42)
- (2): Result of the following calculation rounded up to the nearest integer
(P-VOL capacity [MiB] × number of snapshots × 2 / 2,921,688)

The number of snapshots in the above formula is not the number of current snapshots but the maximum number of snapshots that have been created.

Notes on storage pool capacity

The storage pool capacity minus the virtual volume management area is the capacity that can store data in that storage pool. The virtual volume management area is approximately 3% of the storage pool capacity.

The capacity of the virtual volume management area varies depending on the capacity of the P-VOL, the number of snapshots, and other factors. The estimated capacity for each P-VOL is calculated using the following formula. The capacity for each P-VOL is calculated using the following formula. To calculate the capacity of all the virtual volume management areas, calculate and sum the capacity of all P-VOLs.

Capacity of the virtual volume management area for each P-VOL [MB] = $(4 + (1)) \times 42$

- (1): Result of the following calculation rounded up to the nearest integer
 $(2 \times \text{update type} \times \text{number of snapshots} + 64) \times (\text{volume capacity [MiB]} / 3145728 \text{ [MiB]})$
- Update type: Specify 1 or 2 depending on the volume to be updated.
 - Specify 1 to update only the P-VOL.
 - Specify 1 to update only the S-VOL.
 - Specify 2 to update both the P-VOL and S-VOL.
- Number of snapshots: Number of S-VOLs per P-VOL
- Volume capacity: Capacity of the P-VOL [MiB]

If the capacity of the storage pool is insufficient, the storage pool might become full and the data in snapshot volumes might be lost. Therefore, when creating a storage pool, it is necessary to estimate the copy capacity (capacity of the differential data of the snapshots to be taken) and set the storage pool capacity to be sufficient to accommodate the copy capacity. The capacity of the storage pool is always consumed in units of 42 [MiB]. Therefore, even if the capacity of the difference data is less than 42 [MiB], the storage pool will still need 42 [MiB] of space.



Note:

Pay attention to the free capacity of the storage pool during operation, and increase the capacity of the storage pool before it exceeds the threshold or becomes full.

- Deleting some snapshots in the snapshot tree will not reduce storage pool usage because the space allocated as data in snapshot volumes in the storage pool will not be released. The unreleased space will be used again when storing new data in snapshot volumes in the same snapshot tree.

If you want to reduce usage of the storage pool, you need to delete all snapshots created in the corresponding snapshot tree.

- Note that if the storage pool becomes full, and a snapshot failure occurs (snapshotStatus becomes "Error"), you cannot restore the snapshot even if you resolve the insufficient capacity by deleting volumes or adding drives.



Caution:

If a storage cluster failure occurs because power supply to the storage cluster stops or for other reasons, data in snapshot volumes might be lost.

Data in the snapshot volume (S-VOL or P/S-VOL) cannot be read or written if it is lost. Data in all snapshot volumes might be lost depending on the failure type. You have no choice but to delete snapshot volumes that cannot be read or written.

We recommend that you back up your data to another device (e.g., via the host) as necessary.

Whether or not the operations can be performed for each snapshotStatus

The snapshot tree configuration and whether or not the operations can be performed for each snapshotStatus are as follows:




Note:

It might take about one minute for snapshotStatus to transit. To confirm the transition of snapshotStatus, wait about one minute and then check again.


■ Non-cascade configuration

"Snapshot operation for A" in the table indicates preparation for taking a snapshot of volume A or taking a snapshot of volume A. To prepare for taking a snapshot of volume B under volume A or take a snapshot of volume B under volume A, see "Snapshot operation for B" in "■ Cascade configuration".



snapshot Type	snapshot Status of A	Snapshot operation for A				
		Prepare and take a snapshot (PrepareAndFinalize)	Prepare	Take a snapshot (Finalize)	Restore	Delete
None	None (A does not exist)	○	○	-	-	-
Snapshot	Empty	-	-	×	×	× *
	Preparing	-	-	×	×	×
	Prepared	-	-	○	△	○
	Normal	-	-	△	○	○
	Restoring	-	-	×	×	×
	Error	-	-	×	×	○
	Deleting	-	-	×	×	△
<p>* If removal processing is interrupted because the storage cluster has stopped, the processing can be completed by retrying the operation.</p> <p>○: Succeeds △: Succeeds but nothing is performed (The job ends normally without performing processing.) ×: Does not succeed -: The operation is not available</p>						

■ Cascade configuration



You can take a snapshot of volume B only if the volume is operable.

snapshot Type	snapshot Status of A	Snapshot operation for B				
		Prepare and take a snapshot (PrepareAndFinalize)	Prepare	Take a snapshot (Finalize)	Restore	Delete
Snapshot	Empty	×	×	-	-	-
	Preparing	×	×	-	-	-
	Prepared	×	○	×	△	○
	Normal	○	○	○	○	○
	Restoring	×	×	×	△	○
	Error	×	×	×	×	○
	Deleting	×	×	-	-	-

○: Succeeds △: Succeeds but nothing is performed (The job ends normally without performing processing.) ×: Does not succeed -: The operation is not available


You can take a snapshot of volume A only if the volume is operable.

snapshot Type	snapshot Status of B	Snapshot operation for A				
		Prepare and take a snapshot (PrepareAndFinalize)	Prepare	Take a snapshot (Finalize)	Restore	Delete
Snapshot	Empty	-	-	○	○	×
	Preparing	-	-	○	○	×
	Prepared	-	-	○	○	×
	Normal	-	-	△	×	×
	Restoring	-	-	△	×	×

snapshot Type	snapshot Status of B	Snapshot operation for A				
		Prepare and take a snapshot (PrepareAndFinalize)	Prepare	Take a snapshot (Finalize)	Restore	Delete
	Error	-	-	△	×	×
	Deleting	-	-	△	×	×

○: Succeeds △: Succeeds but nothing is performed (The job ends normally without performing processing.) ×: Does not succeed -: The operation is not available

■ When sharing the copy-source VOL among multiple snapshots



You can take a snapshot of volume B only if the volume is operable.

snapshot Type	snapshot Status of A	Snapshot operation for B				
		Prepare and take a snapshot (PrepareAndFinalize)	Prepare	Take a snapshot (Finalize)	Restore	Delete
Snapshot	Empty	○	○	○	○	○
	Preparing	○	○	○	○	○
	Prepared	○	○	○	○	○
	Normal	○	○	○	○	○
	Restoring	×	○	×	×	○
	Error	×	×	○	○	○
	Deleting	×	×	○	○	○

○: Succeeds △: Succeeds but nothing is performed (The job ends normally without performing processing.) ×: Does not succeed

* Successful when the KARS06132-C message is output to the event log.

- QoS settings of snapshots

QoS settings can be applied to a P-VOL, S-VOL, and P/S-VOL.

For details about the QoS function, see *About the Quality of Service (QoS) function*.



Caution:

Snapshots do not inherit QoS settings of the copy-source volumes.

Preparing to take a snapshot

Run either of the following commands with the copy-source volume specified to create a volume (S-VOL) to take a snapshot of.

- To prepare to take a snapshot, the following calculation results must be less than 80.

$$\text{usedCapacity[MiB]} / \text{allocatableCapacity[MiB]} \times 100$$

You can verify the `usedCapacity` and `allocatableCapacity` by running either of the following commands:

REST API: GET /v1/objects/storage-controllers

CLI: `storage_controller_list`

- Preparation for taking a snapshot cannot be performed if there is a snapshot whose status is "Deleting" or "Error" on the copy-source volume.
- If the copy-source volume is a snapshot, the system can be prepared for taking a snapshot only when the copy-source snapshot is in Normal state.



Note:

Preparing to take a snapshot and taking a snapshot can be performed in one step. Batch processing simplifies the operation, but if the snapshot is to be taken after I/O is stopped, the I/O stop time will be longer because I/O must also be stopped for the time it takes to prepare to take the snapshot.

Therefore, this manual separately describes the operation for preparing to take a snapshot and the operation for taking a snapshot.

To sequentially perform the operation for preparing to take a snapshot and the operation for taking a snapshot, specify `PrepareAndFinalize` for the `operationType` parameter (in the case of CLI, `--operation_type`) when you perform step 3.

Note that the operation for preparing to take a snapshot and the operation for taking a snapshot are sequentially performed if the `operationType` parameter (in the case of CLI, `--operation_type`) is omitted.

Before you begin

- Required role: Storage
- When preparing for obtaining snapshots for volumes in a VPS: Scope of the VPS

Procedure

1. When you are not using the multi-tenancy function or when you are using the multi-tenancy function and you prepare obtaining snapshots for volumes in a system scope, you can skip this task. Go to the next step.

When you prepare obtaining snapshots for volumes in a VPS, verify the VPS ID and conditions set for the VPS (QoS settings).

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list

**Note:**

A system administrator can make QoS settings for individual volumes without limitation for the QoS settings of volumes. If QoS parameters are not specified, volumes (S-VOL) for obtaining snapshots are created with the QoS parameters set to a VPS.

2. Verify the ID of the volume for which you want to create a snapshot.

If you want to specify a volume by its name in the CLI, verify the name of the intended volume.

REST API: GET /v1/objects/volumes

CLI: volume_list

3. Obtain a list of S-VOL information by specifying the ID of the copy-source volume, and then confirm that there are no volumes for which "snapshotStatus" is "Deleting" or "Error".

If you use the CLI, you can specify the name of the volume instead of its ID.

REST API: GET /v1/objects/volumes/<id>/snapshot-volumes

CLI: snapshot_volume_list

If there are any volumes with snapshotStatus of "Deleting", wait for the respective processes to complete. If there are any volumes with "Error" snapshotStatus, delete the associated S-VOLs.

4. Prepare to take a snapshot.

Run either of the following commands with the ID of the copy-source volume specified and "Prepare" specified for the operationType parameter.

If you use the CLI, you can specify the name of the volume instead of its ID. The parameter is --operation_type and the specified value is "Prepare".

REST API: POST /v1/objects/volumes/actions/create-snapshot/invoke

CLI: volume_create_snapshot

Verify the job ID which is displayed after the command is run.

5. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

6. Obtain a list of S-VOL information to verify that an S-VOL whose snapshot status is "Prepared" has been created.

Run one of the following commands with the ID of the copy-source volume specified.

If you use the CLI, you can specify a name instead of the ID of the volume.

REST API: GET /v1/objects/volumes/<id>/snapshot-volumes

CLI: snapshot_volume_list

After performing step 5, it might take approximately one minute for the snapshot status to change. When you want to verify that the snapshot status has changed, wait for approximately one minute before verification.

Taking a snapshot (CLI or REST API)

Run either of the following commands with the copy-destination volume created in the preparatory step specified.

- If a snapshot in the Restoring state exists under the copy-source volume, a snapshot cannot be taken.
- You can take a snapshot only when the copy-destination snapshot is in Prepared state.
- If the copy-source volume is a snapshot, you can take a snapshot only when the copy-source snapshot is in Normal state.



Note:

- After a snapshot is taken, a write request for the P-VOL is processed with the CoW (Copy-on-Write) method. In the CoW method, the data existing before a write is performed is copied to the storage pool, and then the write is performed.
- Immediately after creating a snapshot volume, whenever the data in the P-VOL is updated, the difference data is copied. If many write I/O operations are performed from the host continuously, a large amount of difference data will be copied. This will increase the load on the storage node, causing degradation in IOPS performance of write and read I/O operations. If IOPS performance is degraded, reduce the number of write I/O operations from the host.

**Caution:**

To keep data consistent, stop the host I/O until a snapshot has been taken. If the volume is for the OS system disk, take a snapshot after shutting down the OS. For a data disk volume, forcibly write file cache to disk and stop writing data to the drive (for example, with the sync command and fsfreeze command) or unmount the disk, and then take a snapshot. After a snapshot has been obtained, resume the host I/O operations. The steps for stopping the host I/O operations depend on the OS. Stop the host I/O operations according to the steps for the OS you are using.

Before you begin

- Required role: Storage
- When obtaining snapshots for volumes in a VPS: Scope of the VPS

Procedure

1. When you are not using the multi-tenancy function or when you are using the multi-tenancy function and you obtain snapshots for volumes in a system scope, you can skip this task. Go to the next step.

When obtaining snapshots for volumes in a VPS, verify the VPS ID.

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list

2. Verify the ID of the volume for which you want to create a snapshot.

If you want to specify a volume by its name in the CLI, verify the name of the intended volume.

REST API: GET /v1/objects/volumes

CLI: volume_list

3. Obtain a list of S-VOL information by specifying the ID of the copy-source volume, and verify that there is no volume with snapshotStatus of "Restoring". Also, verify the ID of the copy-destination volume and make sure that snapshotStatus is "Prepared".

If you use the CLI, you can specify a name instead of the ID of the volume.

REST API: GET /v1/objects/volumes/<id>/snapshot-volumes

CLI: snapshot_volume_list

After verification, go to the next step.

4. Take a snapshot.

Run either of the following commands with the ID of the copy-source volume specified and "Finalize" specified for the operationType parameter. When you specified "Finalize" for the operationType parameter, QoS parameters cannot be specified.

If you want to use the CLI, you can specify the name of the volume instead of its ID. The parameter is operation_type and the specified value is "finalize".

REST API: POST /v1/objects/volumes/actions/create-snapshot/invoke

CLI: volume_create_snapshot

Verify the job ID which is displayed after the command is run.

5. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

6. Obtain information about replication-destination volumes to verify that the snapshot status has changed to "Normal".

REST API: GET /v1/objects/volumes/<id>

CLI: volume_show

After performing step 5, it might take approximately one minute for the snapshot status to change. When you want to verify that the snapshot status has changed, wait for approximately one minute before verification.

Deleting snapshots (CLI or REST API)

When snapshot volumes are no longer necessary, you can delete them as follows.

It is impossible to delete a snapshot volume from the GUI.

When you perform editing, note the following points:

- If the path (volume path) between the snapshot volume and the compute node is set, the snapshot volume cannot be deleted. Disconnect the snapshot volume from the compute node before deleting it.
- You can delete a snapshot only when the snapshot is in Prepared, Normal, or Error state.
- A snapshot that is the copy source of another existing snapshot cannot be deleted.
- To delete the S-VOL, run the snapshot deletion command.

- To delete the P-VOL, run the snapshot deletion command to delete all the snapshot volumes (S-VOL) that were created from the P-VOL to be deleted. Then, run the volume deletion command to delete the P-VOL (which is now a normal volume due to S-VOL deletion).
- To delete the P/S-VOL, run the snapshot deletion command to delete all the snapshot volumes (S-VOL) that were created from the P/S-VOL to be deleted. Then, run the volume deletion command to delete the P/S-VOL (which is now a S-VOL due to S-VOL deletion).

**Caution:**

- If snapshot volumes exist when creating configuration backup files according to *Performing the configuration backup operation*, a system administrator or users need to delete snapshot volumes and volume paths whose configuration was restored by restoring storage cluster configuration.
- Deleting some snapshots in the snapshot tree will not reduce the storage pool usage because the space allocated for data in snapshot volumes in the storage pool will not be released. The unreleased space will be used again when storing new data in snapshot volumes. If you want to reduce the usage of the storage pool, you need to delete all snapshots created in the corresponding snapshot tree.

Before you begin

- Required role: Storage
- When deleting snapshot volumes in a VPS: Scope of the VPS

Procedure

1. When you are not using the multi-tenancy function or when you are using the multi-tenancy function and you delete snapshot volumes in a system scope, you can skip this task. Go to the next step.

When deleting snapshot volumes in a VPS, verify the VPS ID.

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list

2. Verify the ID of the volume (P-VOL or S-VOL).

If you use the CLI to specify P-VOL and S-VOL by name, check the names of P-VOL and S-VOL.

REST API: GET /v1/objects/volumes

CLI: volume_list

3. Obtain information of the snapshot volume to be deleted, and verify that snapshotStatus is "Prepared", "Normal" or "Error".

Run either of the following commands with the volume ID specified.

If you want to use the CLI, you can specify the name of the volume instead of its ID.

REST API: GET /v1/objects/volumes/<id>

CLI: volume_show

After verification, go to the next step.

4. Delete the snapshot volume.

Run either of the following commands with the S-VOL ID of the volume to be deleted specified.

If you use the CLI, you can specify the name of the S-VOL instead of its ID.

REST API: POST /v1/objects/volumes/actions/delete-snapshot/invoke

CLI: volume_delete_snapshot

Verify the job ID which is displayed after the command is run.

5. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

6. Obtain a list of information about S-VOLs and verify that S-VOLs are deleted.

Run either of the following commands with the P-VOL volume ID specified.

If you use the CLI, you can specify a name instead of the ID of the volume.

REST API: GET /v1/objects/volumes/<id>/snapshot-volumes

CLI: snapshot_volume_list



Note:

If the snapshotStatus of an S-VOL becomes "Error" and the snapshot deletion command is unsuccessful even if the intended S-VOL is specified, you can delete all S-VOLs created from the P-VOL simultaneously by running the snapshot deletion command with the volume ID of the P-VOL and true specified for masterVolumeld and snapshotTree (for the REST API) and for --master_volume_id and --snapshot_tree (for the CLI), respectively. In performing this operation, if a command that performs an operation on a snapshot volume existing in the tree on the P-VOL is already being run by a prior job, the job might become unsuccessful. The job that runs the snapshot deletion command might also result in failure. If the job results in failure, rerun the command. For detailed operations, see *POST/v1/objects/volumes/actions/delete-snapshot/invoke* in the *Hitachi Virtual Storage Software Block REST API Reference* and *volume_delete_snapshot* in the *Hitachi Virtual Storage Software Block CLI Reference*.

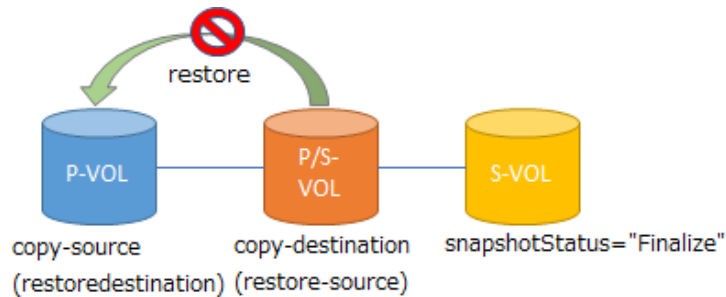
Restoring P-VOLs from a snapshot (CLI or REST API)

The operation of restoring the copy-source volume data from the copy-destination volume data is called snapshot restore.

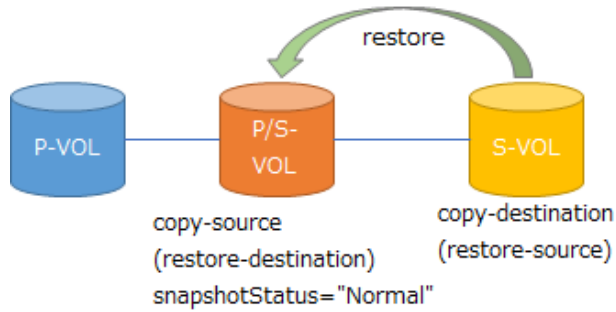
When restore is started and snapshotStatus becomes Restoring, the data in the restore-destination volume will virtually reflect the data in the restore-source volume, even if the restore copy is not yet complete, and the restored data can be accessed immediately.

When restore starts, access to the restore-source volume is suppressed. When the restore is complete, the restore-source volume will transition to the Prepared state. The volume that has transitioned to the Prepared state can take the next snapshot without the need to perform preparation for taking a snapshot.

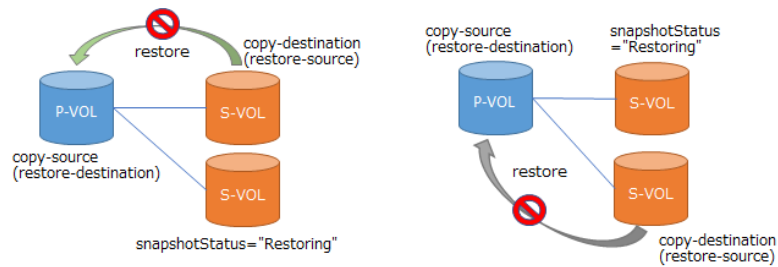
- You can restore a volume only when snapshotStatus of the copy-destination (restore-source) volume is "Normal".
- You cannot restore a volume if a snapshot has been taken by using the copy-destination (restore-source) volume as the copy-source volume in a cascade configuration.



- If the copy-source (restore-destination) volume is a snapshot in a cascade configuration, a volume can be restored only when snapshotStatus of the copy-source (restore-destination) volume is "Normal".



- When the copy-source volume is shared, if there is already a snapshot whose snapshotStatus is "Restoring" under the copy-source (restore-destination) volume, a new restore operation cannot be started.



**Note:**

- When a restore operation is performed, the data in the restore-source volume is written to the restore-destination volume. Because the data in the restore-destination volume before the data in the restore-source volume is written is copied to the storage pool here, the storage pool usage will increase. If data in snapshot volumes is lost during restore due to the storage pool being full, data in the restore-destination volume will also become unavailable. When performing a restore operation, check if the capacity is sufficient for the restore operation by following the steps 1 to 6 in the procedure below. If it is insufficient, secure necessary free capacity.

- When restoring a volume, do so when IOPS to the restore-destination volume is low.

If the data in the restore-destination volume is updated by a write I/O, differential data copy will run. Furthermore, copying data from the restore-source volume to the restore-destination volume by restore will also run. These copy operations increase the load in the storage node and deteriorate not only the IOPS performance of write I/O but also the IOPS performance of read I/O.

- If usage of internal resources in the storage node increases due to the load becoming higher, restore might temporarily stop. In this case, an event log indicating that the restore has temporarily stopped (KARS06160-I) is output. When usage of internal resources in the storage node decreases, the restore processing automatically resumes.
- If more failures than the redundant configuration allows occur during restoration, restoration processing might be unsuccessful. However, processing automatically restarts after the node is restored. In such cases, obtain the volume information, and then verify the progress of the restoration from "snapshotStatus". If "snapshotStatus" is "Restoring", restoration is in progress. Wait until the processing finishes. If "snapshotStatus" is "Prepared", restoration has completed.

**Caution:**

- If the storage cluster becomes abnormal during restoration due to failures exceeding redundancy (for example, a network switch failure, power failure, or an unintentional operation to shut down the storage nodes on multiple storage nodes), data in snapshot volumes will be lost and the data on the restoration-destination volume will become unavailable. Before performing restoration, back up data on the restoration-destination volume to another normal volume via hosts or other means, and record information of the restoration-destination volume by performing the operations described in *Obtaining information about individual volumes*.
- To keep the data consistent, stop the host I/O operations until a snapshot has been restored. If the volume is for the OS system disk, restore a snapshot from the restore-source volume (S-VOL) after shutting down the OS. For a data disk volume, restore a snapshot after unmounting the disk. In this case, the `fsfreeze` command cannot be used. After a snapshot restoration is completed, mount the file system, and then resume the host I/O operations. The steps for stopping the host I/O operations depend on the OS. Perform the process according to the steps for the OS you are using.

Before you begin

- Required role: Storage
- When restoring volumes in a VPS from snapshots: Scope of the VPS

Procedure

1. When you are not using the multi-tenancy function or when you are using the multi-tenancy function and you restore volumes from snapshots in a system scope, you can skip this task. Go to the next step.

When restoring volumes in a VPS from snapshots, verify the VPS ID.

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: `GET /v1/objects/virtual-private-storages`

CLI: `vps_list`

2. Obtain a list of volumes and check the ID of the restore-source volume and the logical capacity (`totalCapacity`).

If you use the CLI to specify volumes by name in the following steps, check the names of the volumes.

Verify that `snapshotStatus` of the restore-source volume is "Normal".

REST API: `GET /v1/objects/volumes`

CLI: `volume_list`

After verifying that `snapshotStatus` of the restore-source volume is "Normal", go to the next step.

3. Obtain the ID of the restore-destination volume (P-VOL).

See *Obtaining P-VOL information for the volume*.

4. Obtain the information of the restore-destination volume and check the ID of the storage controller that manages the restore-destination volume.

REST API: GET /v1/objects/volumes/<id>

CLI: volume_show

5. Obtain the information of the storage controller of the ID that is checked in step 4 and check the used capacity (usedCapacity) and the maximum logical capacity (allocatableCapacity) of the storage pool that are managed on the storage controller.

REST API: GET /v1/objects/storage-controllers/<id>

CLI: storage_controller_show

6. Verify that the difference between allocatableCapacity and usedCapacity that you confirmed in step 5 is no less than totalCapacity that you confirmed in step 2.

After verification, go to the next step.

7. Obtain a list of S-VOL information by specifying the restore-destination volume, and verify that there is no volume with snapshotStatus of "Restoring".

REST API: GET /v1/objects/volumes/<id>/snapshot-volumes

CLI: snapshot_volume_list

After verification, go to the next step.

8. Restore the intended P-VOL with data in snapshot volumes.

Run either of the following commands with the restore-source volume (S-VOL) ID specified.

If you use the CLI, you can specify a name instead of the ID of the volume.

REST API: POST /v1/objects/volumes/actions/restore-snapshot/invoke

CLI: volume_restore_snapshot

Verify the job ID which is displayed after the command is run.

You can check the progress of the restoration by obtaining the volume information, snapshotProgressRate. For how to obtain the volume information, see *Obtaining information about individual volumes*. For the volume ID, specify the ID of the restore-source volume.

9. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

10. Obtain information about restore-source volumes (S-VOLs) to verify that the snapshot status has changed to "Prepared".

REST API: GET /v1/objects/volumes/<id>

CLI: volume_show

After performing step 9, it might take approximately one minute for the snapshot status to change. When you want to verify that the snapshot status has changed, wait for approximately one minute before verification.

Obtaining P-VOL information for the volume (CLI or REST API)

The following information can be obtained.

- masterVolumeld: Volume ID (uuid) of the P-VOL
- vpsId: ID of the VPS to which the P-VOL belongs
- vpsName: Name of the VPS to which the P-VOL belongs
- qosParam: QoS-related parameter
 - upperLimitForIops: Upper limit for volume performance (IOPS)
 - upperLimitForTransferRate: Upper limit for volume performance (MiB/s)
 - upperAlertAllowableTime: Alert threshold value for the upper limit of volume performance (second)
 - upperAlertTime: The last time the upper limit of volume performance was continuously exceeded and the conditions for the alert threshold of the performance upper limit were met (UTC)

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Verify the volume ID of the S-VOL that is created from the intended P-VOL.

If you use the CLI to specify a volume by name, check the name of the volume.

REST API: GET /v1/objects/volumes

CLI: volume_list

2. Obtain P-VOL information.

Run either of the following commands with the S-VOL volume ID specified.

If you use the CLI, you can specify a name instead of the ID of the volume.

REST API: GET /v1/objects/volumes/<id>/master-volume

CLI: master_volume_show

Obtaining a list of S-VOL information for the volume (CLI or REST API)

The following information can be obtained.

- snapshotVolumeld: Volume ID (uuid) of the S-VOL
- snapshotVolumeName: Name of the S-VOL
- snapshotVolumeNickname: Nickname of the S-VOL
- statusSummary: Status summary of the S-VOL
- status: Volume status
- snapshotStatus: Snapshot status

It might take about one minute for snapshotStatus to transit. To confirm the transition of snapshotStatus, wait about one minute and then check again.

- snapshotProgressRate: Progress rate of preparation for taking snapshots of, deleting, and restoring the S-VOL
- snapshotTimestamp: Time when recording of difference data started for the P-VOL at the time of snapshot creation
- snapshotType: Snapshot type
- snapshotConcordanceRate: Match rate between the volume and one newer generation of S-VOL or copy-source volume (P-VOL)
- isWrittenInSvol: Whether the volume has been written from the compute node (host)
- vpsId: ID of the VPS to which the S-VOL belongs

- vpsName: Name of the VPS to which the S-VOL belongs
- qosParam: QoS-related parameter
 - upperLimitForIops: Upper limit for volume performance (IOPS)
 - upperLimitForTransferRate: Upper limit for volume performance (MiB/s)
 - upperAlertAllowableTime: Alert threshold value for the upper limit of volume performance (second)
 - upperAlertTime: The last time the upper limit of volume performance was continuously exceeded and the conditions for the alert threshold of the performance upper limit were met (UTC)

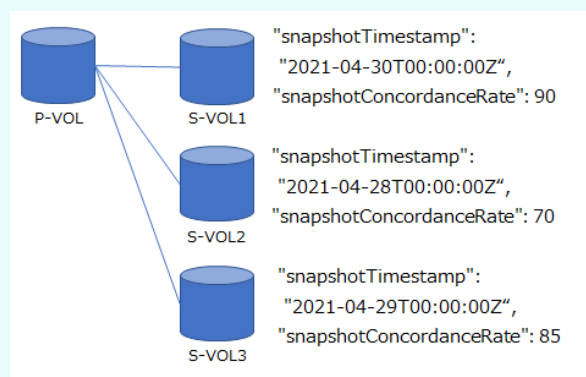


Note:

The matching rate of the S-VOL is handled as follows. Whether the generation is older or newer can be verified from "snapshotTimestamp".

- "snapshotConcordanceRate" for the newest generation (S-VOL1) is the rate of matching with the P-VOL.
- "snapshotConcordanceRate" for the S-VOL2 is the rate of matching with the S-VOL3, which is one generation newer than the S-VOL2.
- "snapshotConcordanceRate" for the S-VOL3 is the rate of matching with the S-VOL1, which is one generation newer than the S-VOL3.

When data is written to the P-VOL, the matching rate of the S-VOL1 (newest generation) seemingly decreases because the data written to the P-VOL is assumed to be differential data. If data is written to an S-VOL, the matching rate of only the S-VOL decreases.



Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. When you are not using the multi-tenancy function or when you are using the multi-tenancy function and you obtain a list of information about S-VOLs for volumes in a system scope, you can skip this task. Go to the next step.

When obtaining a list of information about S-VOLs for volumes in a system scope, verify the VPS ID.

If you want to specify a VPS by its name in the CLI, verify the VPS name.

REST API: GET /v1/objects/virtual-private-storages

CLI: vps_list

2. Verify the volume ID of the P-VOL that is the source of the intended S-VOL.

If you use the CLI to specify a volume by name, check the name of the volume.

REST API: GET /v1/objects/volumes

CLI: volume_list

3. Obtain a list of S-VOL information.

Run either of the following commands with the P-VOL volume ID specified.

If you use the CLI, you can specify a name instead of the ID of the volume.

REST API: GET /v1/objects/volumes/<id>/snapshot-volumes

CLI: snapshot_volume_list



Note:

If the storage cluster stops while snapshot volumes are being deleted or restored, the delete or restore job disappears due to interruption or failure. As a result, the storage cluster stops. When you restart the storage cluster, the delete or restore operation will be automatically resumed.

To verify the status of deleting or restoring snapshot volumes, obtain the volume information and verify snapshotStatus. When the delete operation ends, "snapshotStatus" changes to "Empty". When the restore operation ends, it changes to "Prepared".

Note that the S-VOLs remain even if the snapshot volume is deleted and "snapshotStatus" is changed to "Empty" by the resumed delete operation (Automatic deletion does not take place if the storage cluster stops.). To delete the S-VOLs, run the snapshot deletion command again.

Chapter 11: Obtaining information about storage clusters

Obtaining information about storage clusters (CLI or REST API)

The following information can be obtained.

Role-based execution is not subject to restriction.

- storageDeviceId: Type ID of storage devices
- id: UUID of the storage cluster
- modelName: Product model name
- internalId: ID used inside the storage cluster
- nickname: Nickname of the storage cluster
- numberOfTotalVolumes: Number of created volumes
- numberOfTotalServers: Number of registered compute nodes
- numberOfTotalStorageNodes: Total number of storage nodes that compose the storage cluster
- numberOfReadyStorageNodes: Number of normally operating storage nodes in the storage cluster
- numberOfFaultDomains: Total number of fault domains that comprise the storage cluster.
- totalPoolRawCapacity: Total valid physical capacity of all storage pools that comprise the storage cluster
- totalPoolPhysicalCapacity: Total physical capacity of all storage pools that comprise the storage cluster
- totalPoolCapacity: Total logical capacity of all storage pools that comprise the storage cluster
- usedPoolCapacity: Total consumed capacity of all storage pools that comprise the storage cluster

- `freePoolCapacity`: Total free capacity of all storage pools that comprise the storage cluster
- `savingEffects`: Effect of the data reduction function on the entire storage cluster
- `softwareVersion`: Storage software version
- `statusSummary`: Status summary of the storage cluster
- `status`: Management status of the storage cluster
- `systemRequirementsFileVersion`: Version of the system requirements file registered in the storage cluster
- `serviceId`: Service ID of the storage cluster. The value set by service personnel or maintenance personnel is displayed. When no value is set, null is output.



Note:

- `totalEfficiency` indicates the used capacity reduction effect achieved by the volume creation function and snapshot function.

This is an average ratio (unit: %), where the ratio of `usedCapacity` (used capacity of the storage pool) to `totalVolumeCapacity` (total capacity of created volumes) is calculated for each storage controller, and then weighted according to the total capacity of the created volumes for each storage controller. The larger the total capacity of created volumes of a storage controller is, the more the capacity reduction effect of that storage controller are reflected to the `totalEfficiency` value.

For example, if the total capacity of created volumes is 1,000 for each of the storage controllers and the used capacity of the storage pool is 50, `totalEfficiency` of 2,000 (unit: %) is output by using a REST API or CLI.

In GUI, the value corresponding to `totalEfficiency` is displayed as "20.00:1" in the dashboard window (Total Efficiency) and the Storage Pool window (TOTAL EFFICIENCY). Also, in the dashboard window of the GUI, if the ratio is greater than 9,999,999:100, ">99999.99:1" is displayed.

When snapshot volumes are used, the used storage pool capacity can be reduced by storing only the differential data between P-VOL and S-VOL. Therefore, the value of `totalEfficiency` is larger than that when snapshot volumes are not included, even if the total capacity of created volumes is the same.

When storage pool expansion processing is being performed and KARS16017-I, KARS16020-I, KARS16022-I, or KARS16081-I is output, the value of `totalEfficiency` might become larger even when snapshot operations are not being performed. Note that snapshot operations refer to preparation (for obtaining snapshots) and obtaining, deleting, and restoring snapshots.

- (Virtual machine) The number of created volumes that can be obtained using `numberOfTotalVolumes` does not include the number of volumes whose `volumeType` is `ExternalMigrationOrigin`.

Procedure

1. Obtain information about the storage cluster.

REST API: GET /v1/objects/storage

CLI: storage_show

Obtaining a list of fault domain information (CLI or REST API)

Obtain a list of information for each fault domain. The following information can be obtained.

Role-based execution is not subject to restriction.

- id: ID (uuid) of the fault domain
- name: Name of the fault domain
- statusSummary: Status summary of the fault domain
- status: Status of the fault domain
- numberOfStorageNodes: Number of storage nodes that belong to the fault domain

Procedure

1. Obtain a list of fault domain information.

REST API: GET /v1/objects/fault-domains

CLI: fault_domain_list

Obtaining information about individual fault domain (CLI or REST API)

Obtain the following information about the fault domain for the specified ID.

Role-based execution is not subject to restriction.

- storageNodeIds: List of IDs (uuid) of storage nodes that belong to the fault domain
- id: ID (uuid) of the fault domain
- name: Name of the fault domain
- statusSummary: Status summary of the fault domain
- status: Status of the fault domain
- numberOfStorageNodes: Number of storage nodes that belong to the fault domain

Procedure

1. Obtain a list of fault domain information.
REST API: GET /v1/objects/fault-domains
CLI: fault_domain_list
2. Obtain fault domain information.
Run either of the following commands with the fault domain ID specified.
REST API: GET /v1/objects/fault-domains/<id>
CLI: fault_domain_show

Obtaining health status (CLI or REST API)

Obtain the health status of each resource.

Role-based execution is not subject to restriction.

- type: Resource type
- status: Health status
"status" is displayed as "Normal" or "Alerting". For storage nodes, storage pools, compute ports, drives, volumes, control ports, and internode ports, if the statusSummary of at least one resource among the resource types belonging to the target protection domain contains "Warning" or "Error", the health status will be "Alerting".
- protectionDomainId: ID (uuid) of the protection domain
For the resources (storage clusters, licenses, and fault domains) to manage across the storage cluster, the ID of the protection domain is null.

Procedure

1. Obtain the health status of each resource.
REST API: GET /v1/objects/health-status
CLI: health_status_show

Verifying the cluster master node (primary) (CLI or REST API)

Obtain the following information that indicates whether the request-destination storage node is the cluster master node (primary).

Role-based execution is not subject to restriction.

- `isStorageMasterNodePrimary`: Information indicating whether the request-destination storage node is the cluster master node (primary)

Procedure

1. Verify that the node in question is the cluster master node (primary).

REST API: GET `/configuration/storage-master-node-primary-flag`

CLI: `storage_master_node_primary_flag_show`

If "true" is returned for `isStorageMasterNodePrimary`, the request-destination storage node is the cluster master node (primary). In this case, HTTP status code 200 (OK) is returned for the REST API. However, if error KARS15590-E is returned, the request-destination storage node is not a cluster master node (primary). In this case, HTTP status code 404 (Not Found) is returned for the REST API.

Chapter 12: Managing protection domains

Obtaining a list of information about protection domains (CLI or REST API)

The following information can be obtained.

Role-based execution is not subject to restriction.

- id: IDs (uuid) of protection domains
- name: Names of protection domains
- redundantPolicy: User data protection method
- redundantType: User data protection type
- driveDataRelocationStatus: Status of drive data relocation
- driveDataRelocationProgressRate: Progress rate of drive data relocation
- rebuildStatus: Status of Rebuild
- rebuildProgressRate: Progress rate of Rebuild
- memoryMode: Settings for the memory in use
- asyncProcessingResourceUsageRate: Usage of internal processing I/O resources
- numberOfFaultDomains: Total number of fault domains in the protection domain
- storageControllerClusteringPolicy: Storage controller cluster policy
- minimumMemorySize: Minimum memory value among storage nodes belonging to the protection domain

Procedure

1. Obtain a list of information about protection domains.

REST API: GET /v1/objects/protection-domains

CLI: protection_domain_list

Obtaining information about individual protection domains (CLI or REST API)

The following information can be obtained.

Role-based execution is not subject to restriction.

- totalPhysicalCapacity: Total drive capacity in a protection domain
- isFastRebuildEnabled: Whether Fast Rebuild is enabled or disabled
- id: ID (uuid) of the protection domain
- name: Name of the protection domain
- redundantPolicy: User data protection method
- redundantType: User data protection type
- driveDataRelocationStatus: Status of drive data relocation
- driveDataRelocationProgressRate: Progress rate of drive data relocation
- rebuildStatus: Status of Rebuild
- rebuildProgressRate: Progress rate of Rebuild
- memoryMode: Settings for the memory in use
- asyncProcessingResourceUsageRate: Usage of internal processing I/O resources
- numberOfFaultDomains: Total number of fault domains in the protection domain
- storageControllerClusteringPolicy: Storage controller cluster policy
- minimumMemorySize: Minimum memory value among storage nodes belonging to the protection domain

Procedure

1. Verify the ID of the intended protection domain.

REST API: GET /v1/objects/protection-domains

CLI: protection_domain_list

2. Obtain information about the intended protection domain.

Run either of the following commands with the protection domain ID specified.

REST API: GET /v1/objects/protection-domains/<id>

CLI: protection_domain_show

Changing usage of internal processing I/O resources (CLI or REST API)

Set a ratio of internode network bandwidth allocation to internal processing I/O operations in cases where host I/O operations run simultaneously with internal processing I/O operations (Rebuild, drive data relocation).

When internal processing I/O operations are not being performed, the entire internode network bandwidth is allocated to host I/O operations irrespective of this setting.

Set value	Description	Remarks
VeryHigh	<p>Gives the highest priority to internal processing I/O.</p> <p>A maximum of approximately 80% of the internode network bandwidth is allocated to internal processing I/O operations.</p> <p>The condition for "VeryHigh" to work effectively is that at least 25-Gbps bandwidth should be available for the internode network.</p>	<p>The allocation ratios shown on the left are based on the assumption that 10 Gbps or more (25 Gbps or more for VeryHigh) bandwidth is available in the internode network.</p>
High	<p>Prioritizes internal processing I/O operations.</p> <p>A maximum of approximately 80% of the internode network bandwidth is allocated to internal processing I/O operations.</p>	
Middle (Default)	<p>A maximum of approximately 60% of the internode network bandwidth is allocated to internal processing I/O operations.</p>	
Low	<p>Prioritizes the host I/O operations.</p> <p>A maximum of approximately 20% of the internode network bandwidth is allocated to internal processing I/O operations.</p>	



Caution:

To operate the internal processing I/O with the VeryHigh setting, at least 25-Gbps bandwidth should be available for the internode network. Before setting, review the configuration of the internode network to ensure that it meets the requirements. If the conditions are not met, performance of both host I/O and internal processing I/O might be significantly degraded because all available internode network bandwidth is used up.

Before you begin

Required role: Storage or Service

Procedure

1. Verify the ID of the intended protection domain.

REST API: GET /v1/objects/protection-domains

CLI: protection_domain_list

2. Edit the protection domain settings

Run either of the following commands with the protection domain ID specified.

REST API: PATCH /v1/objects/protection-domains/<id>

CLI: protection_domain_set

Verify the job ID which is displayed after the command is run.

3. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

If the change is made while internal processing I/O operations are being performed, the specified value will be reflected within (approximately) an hour.

4. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Chapter 13: Managing storage controllers

Overview of the features

The storage controller is a Virtual Storage Software block process that manages the capacities and volumes of storage nodes. The capacity of each storage node is managed by the storage controller. The storage controller monitors the system configuration and usage and automatically performs capacity expansion.

As many storage controllers as storage nodes exist, and multiple storage controllers are operating. The total storage pool capacity is the sum of the individual capacities managed by the storage controllers.

Triggers for capacity expansion and approximate processing time are as follows:

Item	Description	Estimated processing time
Triggers for capacity expansion	<p>Make the capacity added by storage pool expansion or capacity reserved for Rebuild (data reconstruction) available as user data.</p> <p>Capacity expansion is mainly performed when:</p> <ul style="list-style-type: none">▪ A storage pool is expanded.▪ Capacity usage of the storage pool increases.	<ul style="list-style-type: none">▪ In HPEC 4D+1P: 30 [sec] × (total physical capacity of added drives per storage node / 121 [GiB])▪ In HPEC 4D+2P: 30 [sec] × (total physical capacity of added drives per storage node / 145 [GiB])▪ In Mirroring Duplication: 30 [sec] × (total physical capacity of added drives per storage node / 194 [GiB])

**Caution:**

If physical capacity of a storage node exceeds the capacity that can be allocated to a storage controller, the calculated approximate processing time might be longer. For details about the maximum capacity that can be allocated to a storage controller, contact customer support

Capacity status of storage controllers

Among the information obtained in *Obtaining a list of information about storage controllers* or *Obtaining information about individual storage controllers*, the table below shows the capacity status of the storage controllers.

Capacity status (capacityStatus) of storage controllers

Status	Description
Normal	<p>The usage rate of the capacity managed by the storage controller is not high and is normal.</p> <p>Specifically, the usage rate is as follows:</p> <ul style="list-style-type: none"> ▪ The usage rate for the maximum logical capacity that can be managed on the storage controller is less than the maximum (storage pool depletion threshold is less than 80%). ▪ The usage rate for the logical capacity managed on the storage controller is less than 80%.
Warning	<p>The usage rate for the capacity managed by the storage controller is high. Or, maximum capacity expansion of the storage controller, which operates when capacity usage is high, is not being performed.</p> <p>Specifically, the usage rate is as follows:</p> <ul style="list-style-type: none"> ▪ The usage rate for the maximum logical capacity that can be managed on the storage controller has exceeded the maximum (storage pool depletion threshold is 80%), but is below the storage controller depletion threshold. ▪ The usage rate for the logical capacity managed on the storage controller is 80% or more and is less than the storage controller depletion threshold, and the storage controller's capacity has not been expanded to the maximum (using the capacity for the Rebuild) at least for five minutes.
Error	<p>The usage rate for the capacity managed by the storage controller is high, and the capacity might be depleted.</p>

Status	Description
	<p>Specifically, the usage rate is as follows:</p> <ul style="list-style-type: none"> ▪ The usage rate for the maximum logical capacity that can be managed on the storage controller has exceeded the storage controller depletion threshold. ▪ The usage rate for the logical capacity that is managed on the storage controller has exceeded the storage controller depletion threshold.

Obtaining a list of information about storage controllers (CLI or REST API)

The following information can be obtained.

- id: IDs (uuid) of storage controllers
- allocatableCapacity: Maximum logical capacity that can be managed by storage controllers
- currentlyAllocatableCapacity: Logical capacity managed by storage controllers
- usedCapacity: Consumed capacity of storage pools managed by storage controllers
- logicalLimit: Maximum logical capacity that can be specified for volumes on storage controllers
- volumeMaximumCapacity: Maximum capacity of a single volume that can be created on the specified storage controller
- freeCapacity: Free capacity of storage pools managed by storage controllers
- status: Status of storage controllers
- activeStorageNodeid: ID (uuid) of the storage node in which the active storage controller is running
- standbyStorageNodeid: ID (uuid) of the storage node in which the standby storage controller is running
- secondaryStandbyStorageNodeid: ID (uuid) of the storage node in which the second standby storage controller is running
- isDetailedLoggingMode: Whether the detailed logging mode is enabled or disabled

- `allocatableCapacityUsageRate`: Usage rate for the maximum logical capacity that can be managed on the storage controller
- `currentlyAllocatableCapacityUsageRate`: Usage rate for the logical capacity managed on the storage controller
- `capacityStatus`: Status of the capacity status managed by the specified storage controller
- `dataRebalanceStatus`: Status of balancing processing of user data capacity managed by each storage controller
- `dataRebalanceProgressRate`: Progress rate of balancing processing of user data capacity managed by each storage controller
- `totalVolumeCapacity`: Total capacity of volumes that have been created on the applicable storage controller
- `provisionedVolumeCapacity`: Capacity of provisioned volumes that have been created on the applicable storage controller
- `otherVolumeCapacity`: Capacity of other volumes that have been created on the applicable storage controller
- `temporaryVolumeCapacity`: Temporary capacity of volumes that have been created on the applicable storage controller
- `capacitiesExcludingSystemData`: Size of storage pool space used by the storage controller

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Obtain a list of information about storage controllers.

REST API: `GET /v1/objects/storage-controllers`

CLI: `storage_controller_list`

Obtaining information about individual storage controllers (CLI or REST API)

The following information can be obtained.

- `id`: ID (uuid) of the intended storage controller
- `allocatableCapacity`: Maximum logical capacity that can be managed by storage controllers

- `currentlyAllocatableCapacity`: Logical capacity managed by storage controllers
- `usedCapacity`: Consumed capacity of storage pools managed by the intended storage controller
- `logicalLimit`: Maximum logical capacity that can be specified for volumes on the intended storage controller
- `volumeMaximumCapacity`: Maximum capacity of a single volume that can be created on the specified storage controller
- `freeCapacity`: Free capacity of storage pools managed by the intended storage controller
- `status`: Status of the intended storage controller
- `activeStorageNodeId`: ID (uuid) of the storage node in which the active storage controller is running
- `standbyStorageNodeId`: ID (uuid) of the storage node in which the standby storage controller is running
- `secondaryStandbyStorageNodeId`: ID (uuid) of the storage node in which the second standby storage controller is running
- `isDetailedLoggingMode`: Whether the detailed logging mode is enabled or disabled
- `allocatableCapacityUsageRate`: Usage rate for the maximum logical capacity that can be managed on the storage controller
- `currentlyAllocatableCapacityUsageRate`: Usage rate for the logical capacity managed on the specified storage controller
- `capacityStatus`: Status of the capacity status managed by the specified storage controller
- `dataRebalanceStatus`: Status of balancing processing of user data capacity managed by each storage controller
- `dataRebalanceProgressRate`: Progress rate of balancing processing of user data capacity managed by each storage controller
- `totalVolumeCapacity`: Total capacity of volumes that have been created on the applicable storage controller
- `provisionedVolumeCapacity`: Capacity of provisioned volumes that have been created on the applicable storage controller
- `otherVolumeCapacity`: Capacity of other volumes that have been created on the applicable storage controller

- `temporaryVolumeCapacity`: Temporary capacity of volumes that have been created on the applicable storage controller
- `capacitiesExcludingSystemData`: Size of storage pool space used by the storage controller

Before you begin

Required role: Security, Storage, Monitor, Service, or Resource

Procedure

1. Verify the ID of the intended storage controller.
REST API: `GET /v1/objects/storage-controllers`
CLI: `storage_controller_list`
2. Obtain information about the intended storage controller.
Run either of the following commands with the storage controller ID specified.
REST API: `GET /v1/objects/storage-controllers/<id>`
CLI: `storage_controller_show`

Editing the settings for the storage controller (CLI or REST API)

Edit the settings for the storage controller. You can enable or disable log detail mode. However, you should set the setting to true only when requested by customer support.

Before you begin

Required role: Service

Procedure

1. To determine the storage controller, run the following command to check the ID of the storage controller.
If you do not identify a storage controller, it will be set for all storage controllers.
REST API: `GET /v1/objects/storage-controllers`
CLI: `storage_controller_list`

2. Edit the settings for the storage controller.

If you want to identify a storage controller, run the command with the ID of the storage controller.

The setting parameters for the log detailing mode are required.

REST API: POST /v1/objects/storage-controllers/actions/configure/invoke

CLI: storage_controller_configure

Verify the job ID which is displayed after the command is run.

3. Verify the state of the job.

Run either of the following commands with the job ID specified.

REST API: GET /v1/objects/jobs/<jobId>

CLI: job_show

If the job state is "Succeeded", the job is completed.

4. Back up the configuration information.

Perform this step by referring to *Backing up the configuration information*.

If you continue operations with other procedures, you must back up the configuration information after you have completed all operations.

Chapter 14: Verifying the status of Rebuild (data reconstruction)

Overview of Rebuild

The Rebuild (data reconstruction) function automatically restores redundancy of data if the level of redundancy is reduced because of a drive failure or a storage node failure.

Two types of Rebuild are available:

- Normal Rebuild: Rebuilds all the data on the storage nodes and drives to be rebuilt. "Rebuild" refers to Normal Rebuild.
- Fast Rebuild: Rebuilds only the necessary data among the storage nodes and drives to be rebuilt. Fast Rebuild requires less rebuild time than Normal Rebuild.

The following table lists the types of rebuild for each rebuild trigger.

Rebuild trigger event		Rebuild type
Drive failure ¹		Normal Rebuild
Updating the software ²		Fast Rebuild
Performing storage node maintenance recovery	Manual maintenance recovery after occurrence of a storage node failure	Normal Rebuild
	Manual maintenance recovery after maintenance blockade	Fast Rebuild
	Storage node auto recovery	
Replacing storage nodes		Normal Rebuild
1. In the event of failure of a drive when the storage node containing that drive does not have enough free space, the Rebuild will result in a failure due to insufficient drive capacity.		
2. If software update is interrupted, Normal Rebuild may be initiated only for the storage nodes that were undergoing software update.		

**Note:**

- If, during the Rebuild, another Rebuild is triggered by performing operations such as maintenance recovery of storage nodes, the Rebuild is temporarily suspended and then restarted automatically. If the event log indicating that the Rebuild started (KARS07000-I) is output after the event log indicating that the Rebuild is suspended (KARS07002-I), the Rebuild has been successfully restarted.
- Even in cases where Fast Rebuild should apply, Normal Rebuild will be performed if the following conditions are met.

Also, Normal Rebuild is to be performed if one of the following conditions is detected after Fast Rebuild. In this case, event log KARS07102-I is output.

- When the Rebuild is interrupted¹
- When drives of a storage node are replaced for maintenance recovery
- When an error occurred in the internal Rebuild processing
- When user data redundancy (dataRedundancy) is 0 (no redundancy) in HEPC 4D+2P²
- If the storage cluster stops or power supply to the storage cluster is interrupted after a storage node is blocked for maintenance or after a storage node fails

1. When the Rebuild is interrupted due to storage node or drive blockage, power-off or shutdown of the storage cluster, suspension of the software update, or other cause. Normal Rebuild is to be performed if it starts after the interruption.

2. Normal Rebuild will continue until the redundancy level is restored to 2.

- When failure in drives and failure in storage nodes in which the drives are installed occur at the same time, the following phenomena might occur.
 - Multiple Rebuild processings are performed.
 - Event log KARS07005-E (indicating that the free space for Rebuild is insufficient) is output.

If multiple Rebuild processings are performed, event log KARS07002-I (indicating cancellation of Rebuild) is output. After a while, the Rebuild is resumed. When event log KARS07001-I (indicating Rebuild completion) is output finally, the operation is successfully completed.

Even when event log KARS07005-E (indicating that the free space for Rebuild is insufficient) is output, the situation might be temporary. If event log KARS07001-I (indicating Rebuild completion) is output finally, the capacity insufficiency is resolved.

Cases when the Rebuild is possible

- Free space required for the Rebuild can be secured.
If free drive space is insufficient, the Rebuild is performed only in areas where rebuilding is possible.
- The status of the storage cluster is "Ready" or "NondisruptiveUpdating".
If the status is other than "Ready" or "NondisruptiveUpdating", perform the following to make the Rebuild run.
 - When the status is "Stopping" or "Stopped", start the storage cluster according to the procedure in *Starting the storage cluster*.
 - When the status is other than the above, the processing to the storage cluster is in progress. Wait until the status changes and ensure that the status is "Ready" or "NondisruptiveUpdating".
- Another Rebuild is not running.
If another Rebuild is already running, the subsequent Rebuild might be delayed until the previous Rebuild completes.

Estimated time before completion of Rebuild

The estimated time before completion of Rebuild can be calculated as follows:

Data capacity to be rebuilt [TiB] × unit processing time for Rebuild [min]

The following table shows the data capacities to be rebuilt.

Category		Data capacities to be rebuilt
Normal Rebuild	For a drive failure	Physical capacity of the drive in which a failure occurred
	For storage node maintenance recovery	<ul style="list-style-type: none"> ▪ When the rebuild capacity policy is "Variable" Physical capacity of the storage node to be recovered ▪ When the rebuild capacity policy is "Fixed" Capacity obtained by subtracting the allocated rebuild capacity from the physical capacity on the storage node to be recovered

Category		Data capacities to be rebuilt
		<p>Here, the allocated rebuild capacity is as follows.</p> <p>Based on the number of rebuildable drives (numberOfDrives) of the rebuildable resources (rebuildableResources) and the drive configuration of each storage node to be recovered, the number of rebuildable drives is allocated to the drives in order of the drive capacity, starting from the drive with the largest capacity. The rebuild capacity will be the sum of all the capacities of the drives.</p> <p>For details about the allocated rebuild capacity, see the <i>Precautions on rebuild capacity management</i> described in <i>Rebuild capacity of a storage pool</i> in this manual.</p> <p>Sometimes, the number of rebuildable drives (numberOfDrives) of the rebuildable resources (rebuildableResources) for each storage node to be recovered cannot be referenced (for example, when the status of the storage node is "PersistentBlockage"). If you need estimation in such a case, refer to the number of times Rebuild can be performed (numberOfDrives) in the storage pool information.</p>
Fast Rebuild		<p>Data capacity updated by write I/O to the drive when the storage node is blocked. The capacity varies depending on write I/O load conditions. Write I/O to the drive might occur asynchronously with I/O from the compute node.</p> <p>However, even in areas having no write I/O, the processing time is approximately three seconds per 100 GiB.</p>

**Tip:**

For how to verify whether there is write I/O to the drive, see *Obtaining a list of low-resolution performance information about drives* or *Obtaining a list of high-resolution performance information about drives*.

Estimate the unit processing time for Rebuild as follows based on the storage pool usage (usedCapacityRate). The formula depends on whether user data (used areas) or no user data (unused areas) exists on the drive.

Rebuild type	Unit processing time for Rebuild
Normal Rebuild	Usage of the storage pool [%] × unit processing time for Rebuild of used areas [min] + (100% - usage of the storage pool [%]) × unit processing time for Rebuild of unused areas [min]
Fast Rebuild	Unit processing time for Rebuild of used areas [min]

The unit processing time for Rebuild of used areas and unused areas varies depending on the user data protection method (RedundantType) settings and the resource usage rate of the internal processing I/O (asyncProcessingResourceUsageRate).

Note that the AsyncProcessingResourceUsageRate setting defines the resource utilization for internal processing I/O (rebuild, drive data relocation). To change the settings, see *Changing usage of internal processing I/O resources*.

redundantType	asyncProcessingResourceUsageRate	Unit processing time for Rebuild of used areas	Unit processing time for Rebuild of unused areas
4D+1P	VeryHigh	30[min]	8[min]
	High	60[min]	12[min]
	Middle	80[min]	15[min]
	Low	200[min]	30[min]
4D+2P	VeryHigh	45[min]	10[min]
	High	90[min]	15[min]
	Middle	120[min]	20[min]
	Low	210[min]	30[min]
Duplication	VeryHigh	15[min]	8[min]
	High	25[min]	12[min]
	Middle	35[min]	15[min]
	Low	160[min]	80[min]

Processing time for Rebuild is based on measurement results in a configuration that meets the following prerequisites:

- With I/O load
- Number of installed user data drives per storage node: 8
- Number of storage nodes: 6
- The user data drives to be used are SAS SSDs.
- The MTU size of the network switch is set to 9000
- The internode network bandwidth is 25 Gbps when `asyncProcessingResourceUsageRate` is set to `VeryHigh`.
- The usage rate for the maximum logical capacity that can be managed on a storage controller (`allocatableCapacityUsageRate`) is basically uniform across storage controllers.

The rebuild processing time may increase or decrease according to the following conditions.

- Network switch performance and I/O load

- Number of installed user data drives per storage node

If the number of installed user data drives per storage node is less than eight, the Rebuild might take longer due to the increased I/O load per user data drive.

If the number of installed user data drives is nine or more, the rebuild processing time might be shorter than the estimated time due to the reduced I/O load per user data drive.

- Number of storage nodes in the case of HPEC

If the number of storage nodes is less than six, the rebuild processing time might be longer than the estimated time due to the increased I/O load per storage node.

If the number of storage nodes is seven or more, the rebuild processing time might be shorter than the estimated time due to the reduced I/O load per storage node.

- Whether the usage rate against the maximum logical capacity that can be managed on the storage controller (`allocatableCapacityUsageRate`) is balanced between storage controllers.

If `allocatableCapacityUsageRate` is unbalanced between storage controllers, the rebuild processing time might differ from the estimated time.

- Immediately after the operation that caused storage pool usage reduction or immediately after internal processing is performed *

The rebuild processing time might become longer than the estimated time immediately after the operation that caused storage pool usage reduction or immediately after internal processing is performed (due to used areas remaining on the drive).

* The storage pool usage rate will be reduced when the following operations or internal processing is performed.

- Deleting volumes
- Deleting snapshots
- Adding storage nodes
- Volume migration by the Capacity balance function
- Capacity balance
- I/O by the UNMAP command
- I/O by the WRITE SAME command



Note:

- When performing a storage node maintenance recovery or replacement for multiple storage nodes, so the rebuild takes as long as (the rebuild processing time per storage node) × (number of storage nodes to be recovered or replaced).
- If Rebuild is triggered by a drive failure, the Rebuild might take longer than the estimated time due to change of data layout among drives. To verify the time required to complete the Rebuild, confirm the following event log KARS07003-I.



Caution:

If physical capacity of a storage node exceeds the capacity that can be allocated to a storage controller, the calculated approximate processing time might be longer. For details about the maximum capacity that can be allocated to a storage controller, contact customer support

When Rebuild is being performed, you can use event log KARS07003-I to confirm the approximate remaining time to complete the Rebuild. However, the time required until the Rebuild is completed might vary depending on the following conditions:

- Network switch performance and I/O load
- When Rebuild was suspended due to failures, maintenance operation, or other causes, and then the Rebuild was performed again
- When the storage cluster was shut down during Rebuild
- When Fast Rebuild changed to Normal Rebuild

For how to obtain a list of event logs, see *Obtaining a list of event logs*.

**Note:**

- The progress rate and remaining time displayed in KARS07003-I are shown in the following information.
 - Progress rate: The progress rate of the rebuild process for the entire data to be rebuilt. If all rebuilt data is successfully processed, it will be 100%. If there is data that failed to rebuild, it will not be 100%.
 - Remaining time: The approximate time before the rebuild process ends (completes or suspends). 0 even if there is data that failed to rebuild.
- If any data fails in the rebuild process, the rebuild process is terminated once. After that, the rebuild process will be re-executed and will try to rebuild the failed data.
- Event log KARS07003-I is output when the progress rate fluctuates by 10 points or more*, or the progress rate becomes 100%.
 - * The progress rate might move backward due to the internal processing or storage system status. The remaining time displayed when the progress rate is moving backward might be longer than the actual remaining time.
- The remaining time to complete the Rebuild displayed in KARS07003-I tends to be less accurate during periods of low progress. As the rate of progress proceeds, the time is gradually corrected to match actual conditions.

Verifying Rebuild status (CLI or REST API)

You can verify the status and progress rate of Rebuild by obtaining information about the protection domain.

Role-based execution is not subject to restriction.

Procedure

1. Obtain a list of information about the protection domain, and verify rebuildStatus and rebuildProgressRate.

REST API: GET /v1/objects/protection-domains

CLI: protection_domain_list

2. rebuildStatus indicates the operation status of the Rebuild, and rebuildProgressRate indicates the progress rate of the Rebuild.
 - rebuildStatus
 - Stopped: Status in which Rebuild is not being performed.
 - Running: Status in which Rebuild is being performed. Rebuild cannot be stopped. Wait until Rebuild is completed, and then perform steps 1 and 2 again to verify the operation state.
 - Error: Status in which Rebuild cannot be performed due to an error. Verify the event logs and perform troubleshooting.
 - rebuildProgressRate

Shows the progress rate (%) of the Rebuild. The progress rate is updated when it fluctuates by 1 point or more. (When progress is made in a short period of time, such as when Fast Rebuild is performed, the progress rate might be updated by several points rather than by 1 point.)

**Note:**

If any of the following occurs, the progress rate of the rebuild process will restart from 0.

- Fast Rebuild is switched to Normal Rebuild.
- Rebuild is suspended and then resumed.
- A storage cluster failure or cluster master node (primary) failure occurs. The progress rate will start from 0, however, already processed data will not be processed again and only unprocessed data will be processed.
- The progress rate might move backward due to the internal processing or storage system status. However, this is normal. If failures occur during the Rebuild, areas for which processing is being performed are rolled back, and the progress rate moves backward.

Chapter 15: Troubleshooting

First points to confirm

A failure that occur has many possible causes, including simple errors (such as wrong cable connections and operational mistakes) and factors outside the Virtual Storage Software block system. The following table shows basic points to confirm. Confirm these points beforehand.

Points to confirm	Remarks	Done
Verify that the cables of the storage system and network peripheral devices are connected correctly.	Simple errors such as disconnection of network cables might often cause a failure.	
Verify that power is supplied to the storage system and network peripheral devices.	If power to the storage system stops, contact customer support to have the storage system restarted after power is restored.	
Verify that you performed the necessary tasks by using the correct procedure as described in the manual.	Verify again of the content of the notes and cautions described in the procedures.	
When operating with the REST API/CLI or maintenance node (adding and replacing nodes, changing and setting configuration information, importing and exporting configuration files), verify that the access-destination address is correctly specified.	Specify the correct IP address or host name for the storage cluster or individual storage node that you want to work with. For information on how to address the CLI, see <i>Hitachi Virtual Storage Software Block CLI Reference</i> .	

Points to confirm	Remarks	Done
<p>When operating with the REST API/CLI or maintenance node (adding and replacing nodes, changing and setting configuration information, importing and exporting configuration files), and the FQDN is specified for the destination, verify that the name resolution can be properly performed.</p>	<ul style="list-style-type: none"> ▪ Verify that the DNS setting is completed. ▪ Verify the path to the DNS server for any failures. ▪ Verify that the DNS server is running. ▪ Verify that the DNS record is correctly set for the FQDN. ¹ <p>1. When changing and setting configuration information, a changed IP address of the control network might be set as a DNS record for the FQDN.</p>	
<p>When using REST API/CLI, verify that the parameters and other options are correctly specified.</p>	<p>See the <i>Hitachi Virtual Storage Software Block REST API Reference</i> or <i>Hitachi Virtual Storage Software Block CLI Reference</i>.</p>	
<p>When operating with the REST API/CLI/GUI or maintenance node (adding and replacing nodes, changing and setting configuration information, importing and exporting configuration files), verify that the SSL communication setting is correctly set.</p>	<p>For more information, contact customer support.</p>	
<p>When operating with the maintenance node (installation, addition, replacement, changing and setting configuration information, importing and exporting configuration files, and creating the configuration backup file), verify that the SSL communication with VMware vCenter Server is correctly set.</p>	<p>For more information, contact customer support.</p>	

Points to confirm	Remarks	Done
<p>Verify that the external servers accessed by the storage system are configured and operating normally.</p>	<ul style="list-style-type: none"> ▪ The storage system works in cooperation with multiple external servers. ▪ If the external servers are not set correctly, a failure might occur. Verify that the following external servers are configured and operating normally. <ul style="list-style-type: none"> ▪ Maintenance node audit log transfer server (external Auditd server) ▪ External authentication (LDAP) server ▪ DNS server ▪ iSNS server ▪ NTP server ▪ SMTP server ▪ SNMP manager ▪ Syslog server <p>For information about how to verify each external server setting, see the <i>Hitachi Virtual Storage Software Block System Administrator Guide</i>.</p> <ul style="list-style-type: none"> ▪ A user of the storage system might have to be registered with an external server, such as an LDAP server. 	
<p>Verify that the control network, compute network, and internode network that constitute Virtual Storage Software block are all running normally.</p>	<ul style="list-style-type: none"> ▪ A failure that occurred in the network affects the storage system. ▪ Use a network monitoring system or survey terminal that is independent of the Virtual Storage Software block to check if there are any failures or performance degradation (switch failure, link failure, increased delay, increased packet loss, connection at unexpected link speed and duplex operating values, etc.) in the network. 	

Points to confirm	Remarks	Done
	<ul style="list-style-type: none"> ▪ If the network is redundant and recovery processing such as switching to the redundant route is occurring due to a failure, check the log of the switch or VMware ESXi, and verify that the time required for the recovery process meets the recovery time confirmed by customer support. ▪ If you have installed a firewall on each network and are filtering using TCP / UDP port numbers, contact customer support and ask them to check if the TCP / UDP port number required for each communication is set to be allowed. ▪ When temporarily connecting an investigation terminal to the internode network for fault isolation, be sure to disconnect after the fault isolation is completed, and verify that the network is closed with no connections other than the internode network port of the Virtual Storage Software block node and the network monitoring system. ▪ For details about how to verify operation of ESXi and vCenter Server, see the VMware documentation. ▪ For information on how to check the monitoring system and switch of the network independent of the Virtual Storage Software block, refer to the manual of each vendor. 	
Verify that the hardware of each storage node is operating normally.	For details about how to verify that hardware is operating normally, see the hardware vendor documentation.	
Verify that the hardware of each storage node is configured correctly.	If the serial number and model name are not set, the storage node will fail to build. Check your hardware vendor's manual for information on how to set up your hardware.	
Verify that the VM on each storage node is operating normally.	For details about how to verify normal VM operation and specify the VM settings, see the VMware documentation.	

Points to confirm	Remarks	Done
Verify that vNIC and other virtual settings for the VM on each storage node are specified correctly.	If all the VMs are suspended due to a commercial power failure or power unit failure, start the storage cluster. For more details, contact customer support.	
Verify that ESXi and vCenter Server are operating normally.	For details about how to verify operation of ESXi and vCenter Server, see the VMware documentation.	
Verify that ESXi and vCenter Server are set up correctly.		

Action to be taken for a volume failure

If you have a volume with STATUS listed in the table below on the Volumes window of the Virtual Storage Software block GUI, verify that the storage node's Health Status is not "Alerting".

If "Alerting" occurs in the Health Status of the storage node, see *If a health status error is detected in the GUI* and address the failure in turn. However, if the status of the storage node where "Alerting" is occurring is only "RemovalFailed", take action for a volume failure as shown in the following table without referring to *If a health status error is detected in the GUI*.

If no failure occurred in a storage node, take action for a volume failure as shown in the following table.



Caution:


If 20 or more volume management operations are being performed at the same time, a volume might transition to one of the following statuses. In this case, reduce the number of simultaneous volume management operations to less than 20, and then take the action described in the table.

STATUS	SNAPSHOT STATUS	Description	VOLUME TYPE	Action*
IOSuppressed	All statuses	The data of the volume or snapshot volume can no longer be guaranteed. I/O operations to the faulty volume or snapshot volume is not possible.	Common for normal and snapshot volumes	Contact customer support.

STATUS	SNAPSHOT STATUS	Description	VOLUME TYPE	Action*
CreationFailed	All statuses	Volume creation ended abnormally.	Normal	Delete the volume which ended abnormally and create a volume again.
			Snapshot	Delete the snapshot volume which ended abnormally and create a volume again.
DeletionFailed	All statuses	Volume deletion ended abnormally.	Normal	Delete the volume which ended abnormally again.
			Snapshot	Delete the snapshot volume which ended abnormally again.
ExpansionFailed	All statuses	Volume expansion ended abnormally.	Normal	Expand the volume which ended abnormally again without specifying a value for the additionalCapacity.
UpdateFailed	All statuses	An update of the volume settings ended abnormally.	Both normal and snapshot volume types	Update the settings of the volume which ended abnormally again. Or, delete the volume which ended abnormally.
* If a problem persists after taking the corresponding action, collect the dump log files of Virtual Storage Software block, and then contact customer support.				

Resolving other error conditions

The following table lists error conditions and provides recommended action for resolving the error. If the problem persists after taking each action, collect the dump log file of the Virtual Storage Software block including the dump log file of the storage controller, and contact the support center.

Error condition	Recommended action
<p>"Alerting" is shown in the license health status.</p>	<p>Check the health status of the licenses, and take appropriate action.</p> <ul style="list-style-type: none"> ▪ If a license is expired (description includes "Term"), request an extended trial license or a perpetual license key and then register the issued license. ▪ If the licensed capacity has been exceeded (description includes "Capacity"), request a license with more capacity and then register the issued license. ▪ If a license whose status is not normal is no longer needed, delete the license. ▪ If the status of all licenses is normal, collect dump log files for all storage nodes and then contact customer support.
<p>A fault set with the status Error is shown in the Protection Domains window.</p>	<p>One or more storage nodes in the target fault set have been blocked.</p> <p>Check the health status of the storage node for the target fault set. If all storage nodes are blocked, there might be an abnormality in the power supply or network. Check the following and take action:</p> <ul style="list-style-type: none"> ▪ Are you powered by a storage system or peripheral device? ▪ Are the cables for storage systems and network peripherals connected correctly? ▪ Is the network that accesses the storage system working properly?
<p>"Alerting" is shown for the drive health status.</p>	<p>Check the health status of the drives.</p> <ul style="list-style-type: none"> ▪ If "Blockage" is shown for a drive, collect dump log files for all storage nodes, and then contact customer support. ▪ If "Blockage" is not shown for a drive, check for other health status errors and take appropriate action. <div style="background-color: #e0f2f1; padding: 10px; margin-top: 10px;"> <p> Note:</p> <p>In the event of a drive failure, storage pool capacity might decrease. In this case, the storage pool capacity is restored by addressing the drive failure.</p> </div>

Error condition	Recommended action
<p>"Alerting" is shown for the storage pool health status.</p>	<p>Check the health status of the storage nodes and drives.</p> <ul style="list-style-type: none"> ▪ If "Alerting" is shown for drive health status, take action to resolve the error condition as described above. ▪ If "Alerting" is shown for storage node status, contact customer support. ▪ If "Maintenance Blockage" is shown for storage node status, wait for the maintenance to be completed. If maintenance is not in process, or if the status does not change to "Normal" after the maintenance, contact customer support. <p>Check the health status of the storage pool.</p> <ul style="list-style-type: none"> ▪ If the pool status is "Exceeded Threshold" or "Error", delete unnecessary volumes or expand capacity so the status becomes "Normal". ▪ If the "StorageControllerCapacitiesGeneralStatus" is "Warning" or "Error", a software update might be in progress. If an update is not in process, or if the status does not change to "Normal" after the update, contact customer support. <p>If you are not able to change the storage pool status to "Normal", collect the dump log files for all storage nodes, and then contact customer support.</p>

Action to be taken if KARS06132-C occurs when starting the storage cluster

If the KARS06132-C event log is output when starting the storage cluster, troubleshoot it according to the following flowchart.



Caution:

If a failure occurred in the storage cluster during its restore, an event log (KARS06170-C) is output. Delete the volume (P-VOL) indicated in the event log (KARS06170-C). After that, create a new volume, and then restore the data from the backup volume to the new volume.

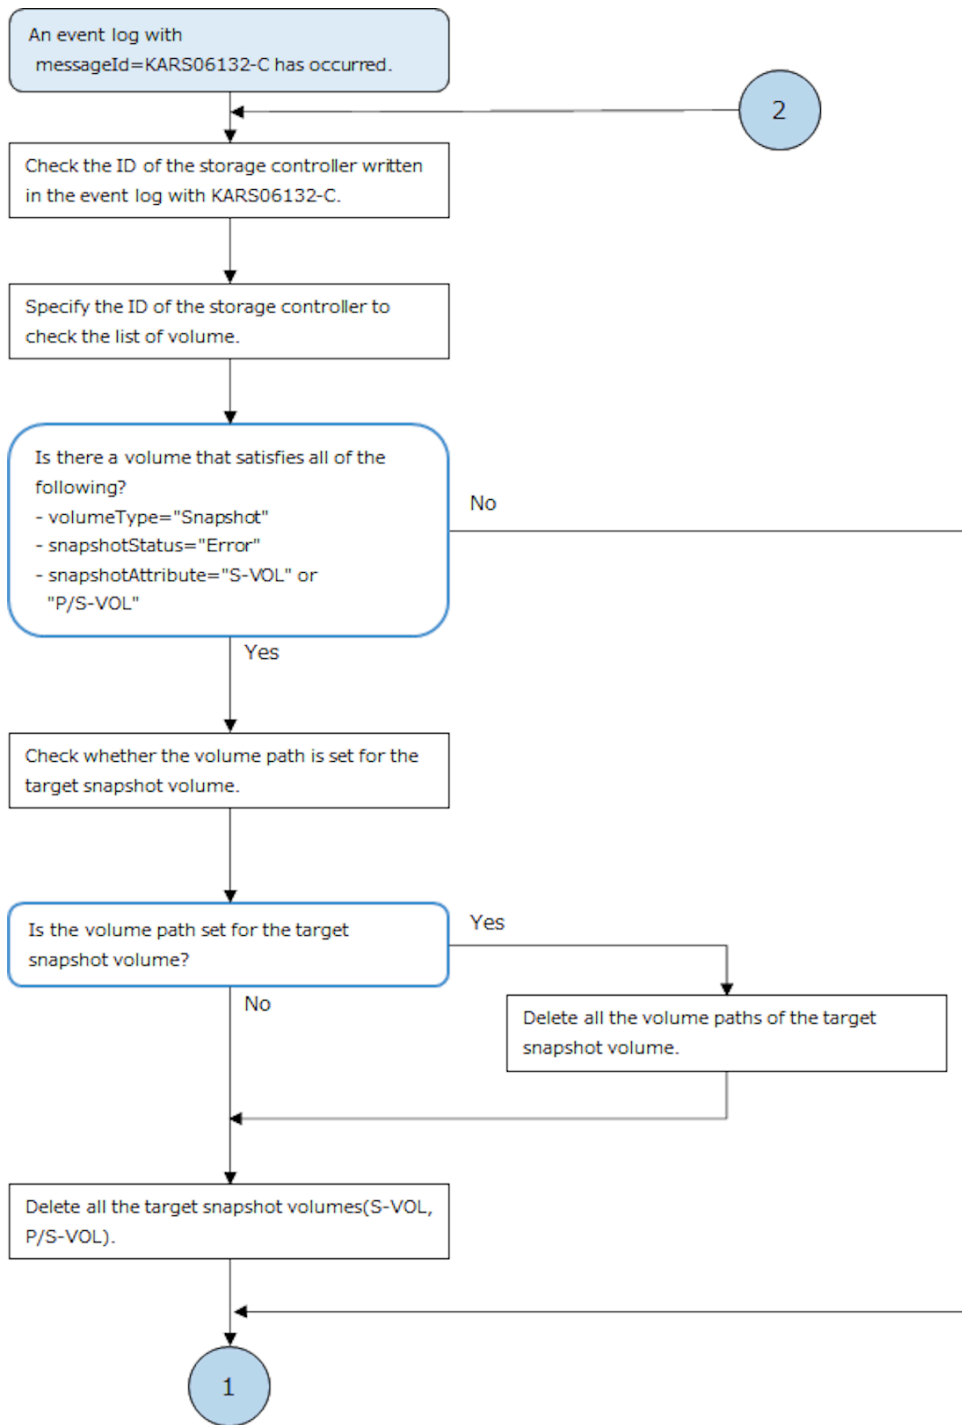


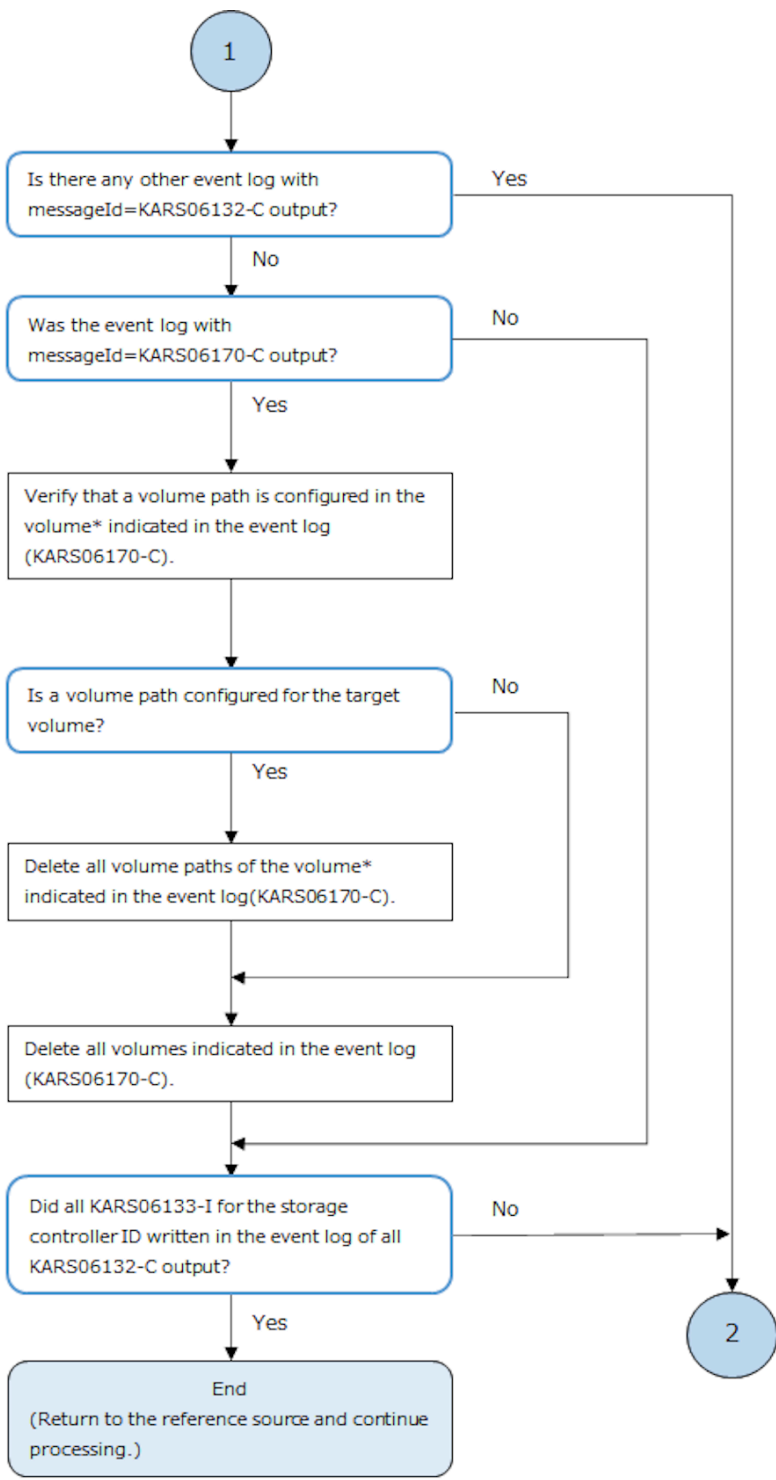
Note:

KARS06132-C is an event log that indicates an anomaly in a snapshot managed by a single storage controller. If more than one storage controller fails, multiple KARS06132-C are output.

Delete all the snapshot volumes whose snapshotStatus is Error among S-VOLs and P/S-VOLs managed by the storage controller corresponding to the event log (KARS06132-C), and then delete all the P-VOLs indicated in the event log (KARS06170-C). By doing so, an event log (KARS06132-I) is output.

In order to deal with it without overlooking it, it is recommended to list KARS06132-C using REST API or CLI, and to list KARS06133-I after it is addressed.

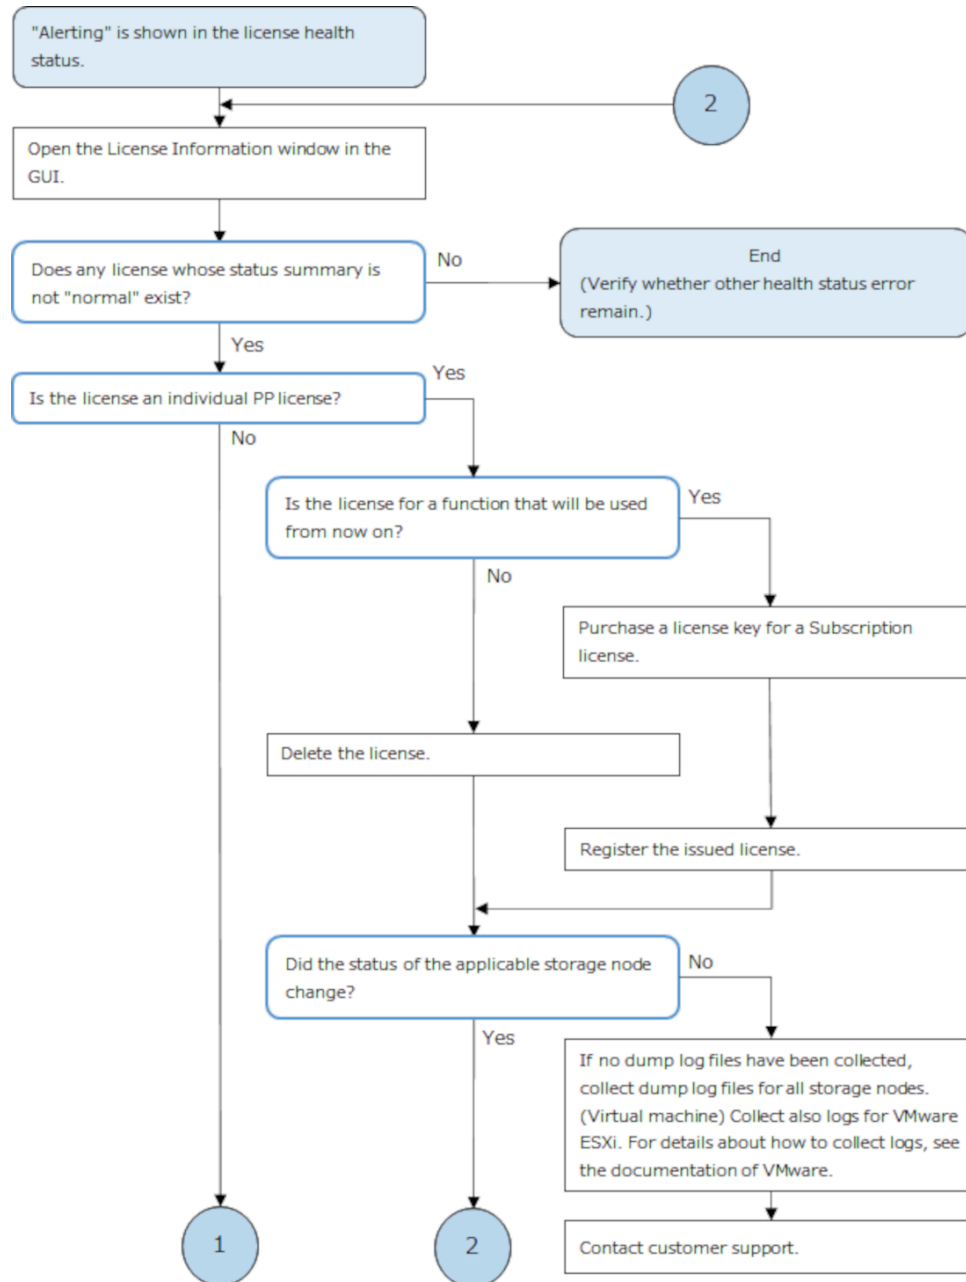


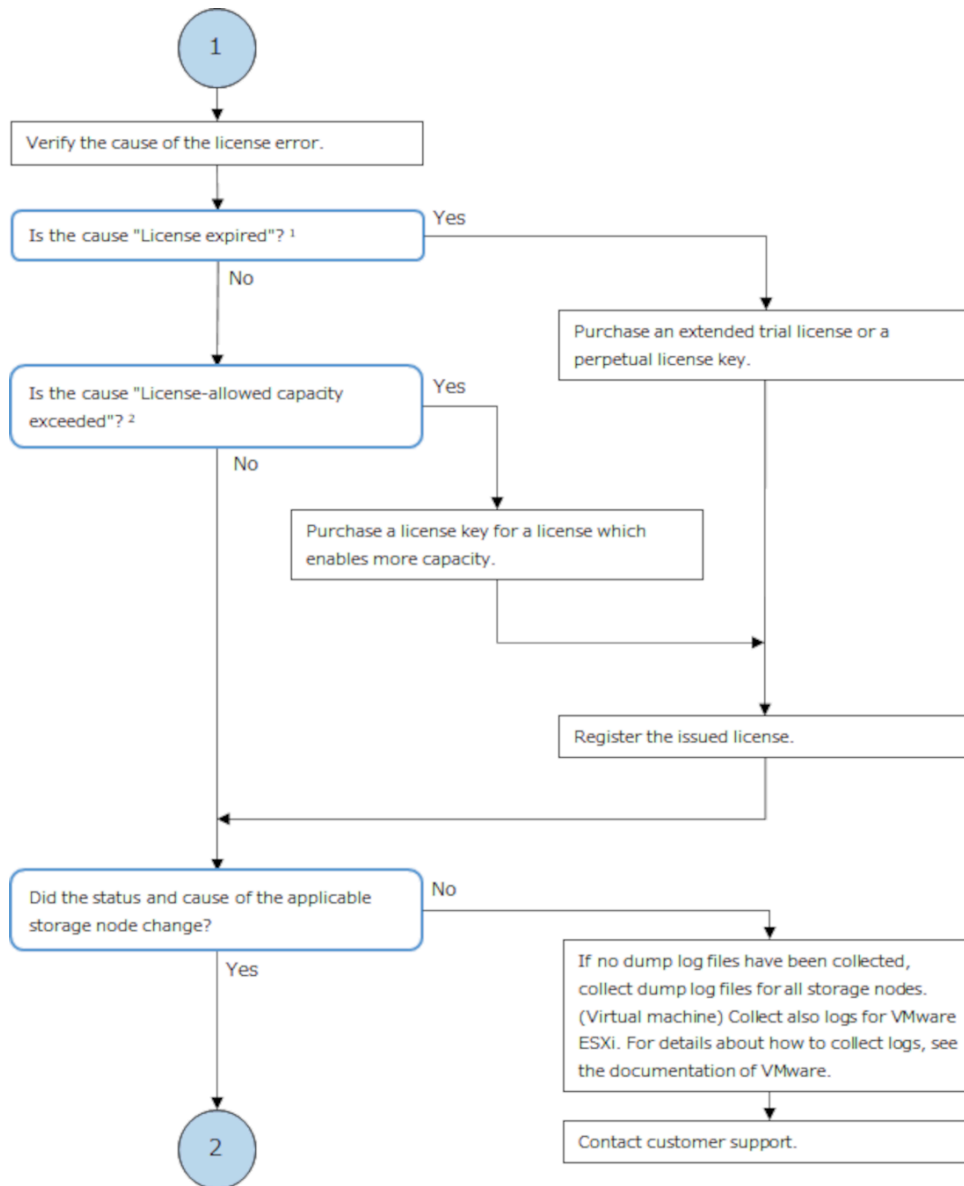


* The target volumes include not only the ones indicated in the event log but also the ones whose status is "IOSuppressed".

Action to be taken when "Alerting" is shown in the license health status

If "Alerting" is displayed in the license health status in the Virtual Storage Software block GUI, troubleshoot the failure according to the following flowchart.





1. If the license has expired, the description of the cause includes "Term".

2. If the license-allowed capacity was exceeded, the description of the cause includes "Capacity".

Action to be taken when "Alerting" is shown in the Fault Domains health status

If the fault domain with STATUS shown in the table below is displayed on the Protection Domain window of the Virtual Storage Software block GUI, follow the action below.

STATUS	Description	Action
Error	One or more storage nodes in the target fault domain have been blocked.	<p>Check the status of the storage node for the target fault domain. If all storage nodes are blocked, there may be an abnormality in the power supply or network. Check the following and take action.</p> <ul style="list-style-type: none"> ▪ Are you powered by a storage system or peripheral device? ▪ Are the cables for storage systems and network peripherals connected correctly? ▪ Is the network that accesses the storage system working properly?

Determining the maintenance priority of storage nodes

If more than one storage node in the following status exists in "the storage node group to be dealt with" determined before running the workflows in *Action to be taken when "Alerting" is shown in the storage node health status*, make sure to determine the order in which you troubleshoot the storage nodes:

Status:

- TemporaryBlockage
- MaintenanceBlockage
- InstallationFailed



Note:

If you are unable to address the target storage node immediately, such as when part procurement takes a long time, target the storage node with the highest priority.

When there are three cluster master nodes in the storage cluster

Take action starting for the storage node with the highest priority (the smallest number in the "Priority" column) in the table below. If there are two or more storage nodes with the highest priority, troubleshoot any of the storage nodes.

Type of storage node	Precedence
Cluster master node	1

Type of storage node	Precedence
Cluster worker node	2

When there are five cluster master nodes in the storage cluster

Perform the following procedure to verify the precedence.

1. Check the number of cluster master nodes in the blocked status¹ in the storage cluster.
2. See the list of storage controllers, and then record the ID of the storage node (hereinafter referred to as "storage node with storage controller in the status of TwoNodesDown") that is "standbyStorageNodeid" or "secondaryStandbyStorageNodeid" of the storage controller whose status is "TwoNodesDown" .
3. Check whether each target storage node in the status that troubleshoots the failure is a cluster master node or a cluster worker node.
4. Based on the check results in step 1 through step 3, make sure to take actions on the failure in the order from the storage node with the highest priority (the smallest number in the Precedence column) by following the tables. The storage nodes with the same priority have the same precedence.

- When the number of cluster master nodes in the blocked status¹ is 1 or less

Type of storage node	Precedence
"Storage node with storage controller in the status of TwoNodesDown" and cluster master node	1
"Storage node with storage controller in the status of TwoNodesDown" and cluster worker node	2
Not "storage node with storage controller in the status of TwoNodesDown" but cluster master node	3
Not "storage node with storage controller in the status of TwoNodesDown" but cluster worker node	4

- When the number of cluster master nodes in the blocked status¹ is 2

Type of storage node	Precedence
"Storage node with storage controller in the status of TwoNodesDown" and cluster master node	1
Not "storage node with storage controller in the status of TwoNodesDown" but cluster master node	2
"Storage node with storage controller in the status of TwoNodesDown" and cluster worker node	3
Not "storage node with storage controller in the status of TwoNodesDown" but cluster worker node	4

1. The status indicates the following:

TemporaryBlockage, MaintenanceBlockage, TemporaryBlockageFailed, MaintenanceBlockageFailed, InstallationFailed, PersistentBlockage, and the status containing the string "RemovalFailed"

Operations with dump log files (GUI)

Overview

The following table describes the operations you can perform for dump log files.

For details about dump log file operation, contact customer support.

**Caution:**

- Do not pause download on your web browser. If you pause download of multiple dump log files, the controller node on which the GUI is used consumes more memory than usual, and the download might fail.
- Do not close your browser while downloading dump log files. If you close the browser during the download, delete the dump log files already downloaded, and then retry the download.
- While downloading dump log files, do not download dump log files from other tabs or windows simultaneously.
- Delete the dump log files that have been downloaded.
- The GUI cannot be used to collect dump log files whose STATUS is "Unavailable". If you want to collect such files, perform the operation by using ticket authentication from a REST API or CLI.

Operation	Window to operate on	Operation icon	Dialog
Creating dump log files	Dump Log Files list window [Base] tab		Create Base Dump Log Files
	Dump Log Files list window [Monitor] tab		Create Monitor Dump Log File
Downloading dump log files	Dump Log Files list window [Auto Collection] tab		Download Auto Collection Dump Log Files
	Dump Log Files list window [Base] tab		Download Base Dump Log Files
	Dump Log Files list window [Monitor] tab		Download Monitor Dump Log File
Deleting dump log files	Dump Log Files list window [Auto Collection] tab		Delete Auto Collection Dump Log Files

Operation	Window to operate on	Operation icon	Dialog
	Dump Log Files list window [Base] tab		Delete Dump Log Files
	Dump Log Files list window [Monitor] tab		

Creating dump log files (GUI)

Create dump log files.

Before you begin

- Required role: Service

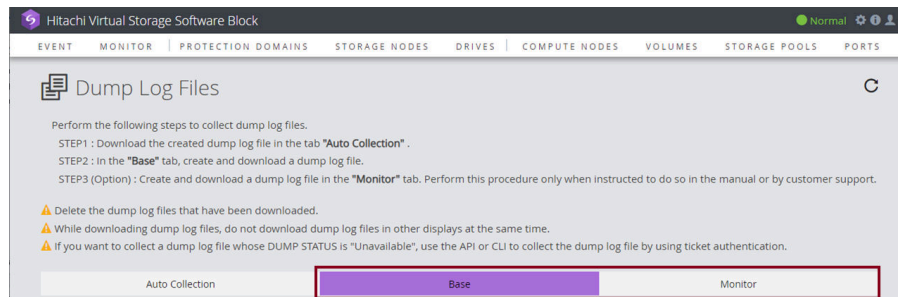
Procedure

- Click the gear icon on the navigation bar, and then select "Dump Log Files".



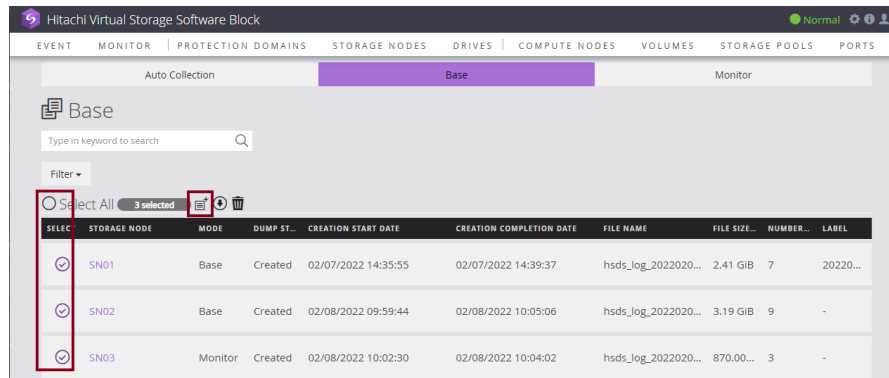
The Dump Log Files list window is displayed.

- Click **Base** or **Monitor**.



A list of dump log file creation statuses for each STORAGE NODE is displayed.

- Select the target storage node, and then click the creation icon.



The dump log file creation dialog is displayed.

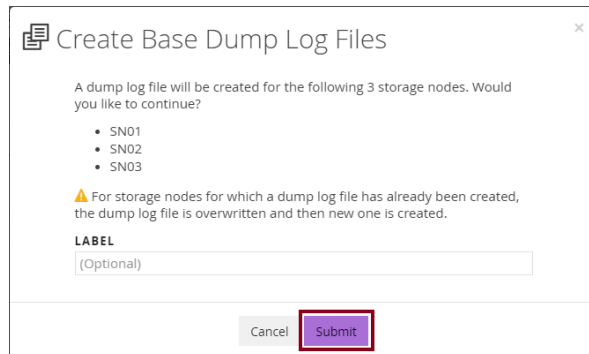


Note:

The dump log file creation icon is activated in the following cases:

- **Base** tab: When one or more STORAGE NODEs are selected (multiple nodes can be selected).
- **Monitor** tab: When one STORAGE NODE is selected (multiple nodes cannot be selected).

4. Verify the displayed information, and then click **Submit**.



The dump log file creation dialog is closed and a dump log file is created.



Note:

The dialog started from the **Base** tab has the title "Create Base Dump Log Files".

The dialog started from the **Monitor** tab has the title "Create Monitor Dump Log File".

5. When the following "Completed" message is displayed, processing is completed.
- Successfully created dump log files.

- If you are logged out during creation processing, verify the Dump Log File window. In some cases, creating dump log files takes a long time, which can cause logout. When the dump status of the storage node for which you requested dump log files is "Created", dump log file creation has been completed normally. If you recreated dump log files, also confirm the date of "CREATION COMPLETION DATA".

Downloading dump log files (GUI)

Download dump log files.

Before you begin

- Required role: Service

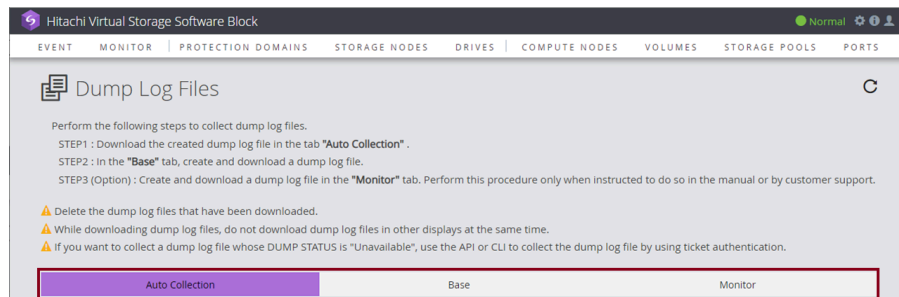
Procedure

- Click the gear icon on the navigation bar, and then select "Dump Log Files".

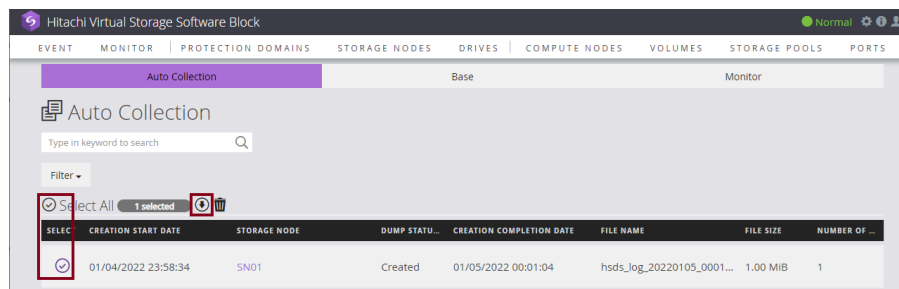


The Dump Log Files list window is displayed.

- Click **Auto Collection**, **Base**, or **Monitor**.



- Auto Collection** tab: A list of creation statuses for dump log files whose type is "Auto Collection" is displayed.
 - Base** tab, **Monitor** tab: A list of dump log file creation statuses for each STORAGE NODE is displayed.
- Select the target storage node, and then click the download icon.



The dump log file download dialog is displayed.



Note:

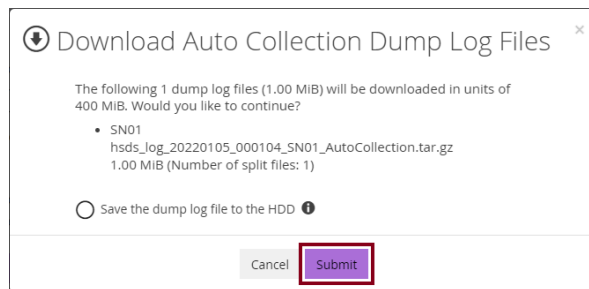
The download icon is activated in the following cases:

- **Auto Collection** tab: When one or more STORAGE NODEs are selected (multiple nodes can be selected).
- **Base** tab: When one STORAGE NODE is selected (multiple nodes can be selected).
- **Monitor** tab: When one STORAGE NODE is selected (multiple nodes cannot be selected).

4. Verify the displayed information, and then click **Submit**.

If an HDD is specified as the download destination drive, select the check box to the left of "Save the dump log file to the HDD", and then click **Submit**. Select this check box when the disk is slow to write to.

Before clicking **Submit**, confirm that there is free space equal to or larger than the file size displayed on the dialog box.



The dump log file download dialog is closed, and the dump log files are downloaded.

The processing status of the download is displayed in the bottom right of the window as a pop up.



Note:

- The storing location of dump log files is specified in your browser. You can confirm or change the storing location in the setting window of your browser.
- If a dump log file is not downloaded successfully, collect it using a REST API or CLI. For details, contact customer support.
- The dialog started from the **Auto Collection** tab has the title "Download Auto Collection Dump Log Files".
- The dialog started from the **Base** tab has the title "Download Base Dump Log Files".
- The dialog started from the **Monitor** tab has the title "Download Monitor Dump Log File".

5. When the following "Completed" message and download history of the browser are displayed, processing is completed.
 - Successfully downloaded dump log files.

**Caution:**

I/O performance might be affected while dump log files are being downloaded from the GUI.

Deleting dump log files (GUI)

Delete dump log files.

Before you begin

- Required role: Service

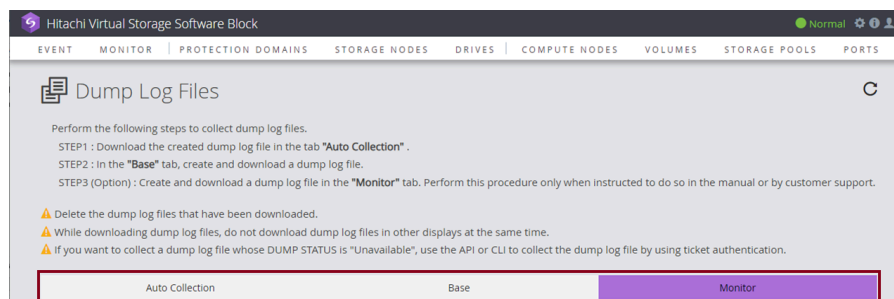
Procedure

1. Click the gear icon on the navigation bar, and then select "Dump Log Files".

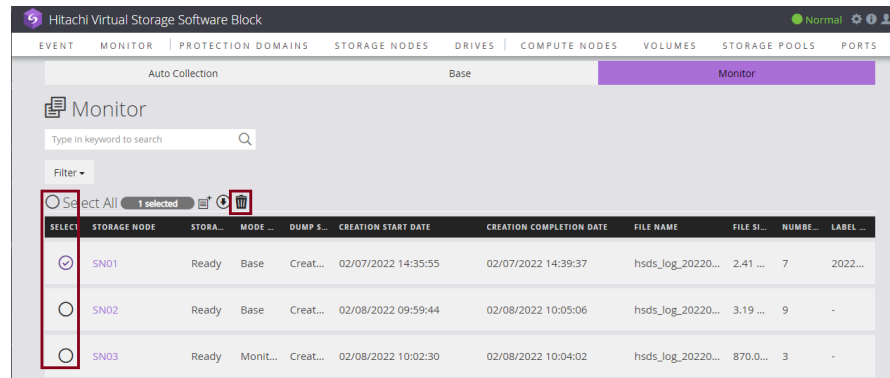


The Dump Log Files list window is displayed.

2. Click **Auto Collection**, **Base**, or **Monitor**.



- **Auto Collection** tab: A list of creation statuses for dump log files whose type is "Auto Collection" is displayed.
 - **Base** tab, **Monitor** tab: A list of dump log file creation statuses for each STORAGE NODE is displayed.
3. Select the target storage node, and then click the deletion icon.



The dump log file deletion dialog is displayed.

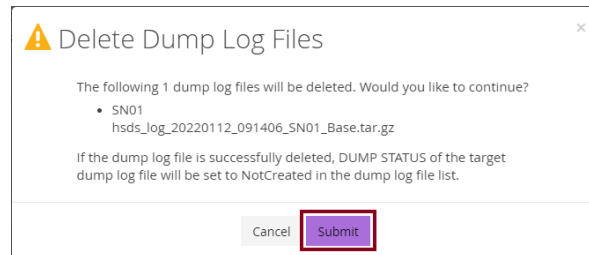


Note:

You can select multiple storage nodes.

The deletion icon is activated when one or more storage nodes are selected.

4. Verify the displayed information, and then click **Submit**.



The dump log file deletion dialog is closed, and the dump log files are deleted.

For the **Auto Collection** tab, the deleted file is cleared from the list, and for the **Base** and **Monitor** tabs, DUMP STATUS in the list becomes "No Created" (indicating that deletion was completed).



Note:

The dialog started from the **Auto Collection** tab has the title "Delete Auto Collection Dump Log Files".

The dialog started from the **Base** tab has the title "Delete Dump Log Files".

The dialog started from the **Monitor** tab has the title "Delete Dump Log Files".

Appendix A: Capacity design (for HPEC 4D+1P)

Capacity design principles (1) : Physical capacity (for HPEC 4D+1P)

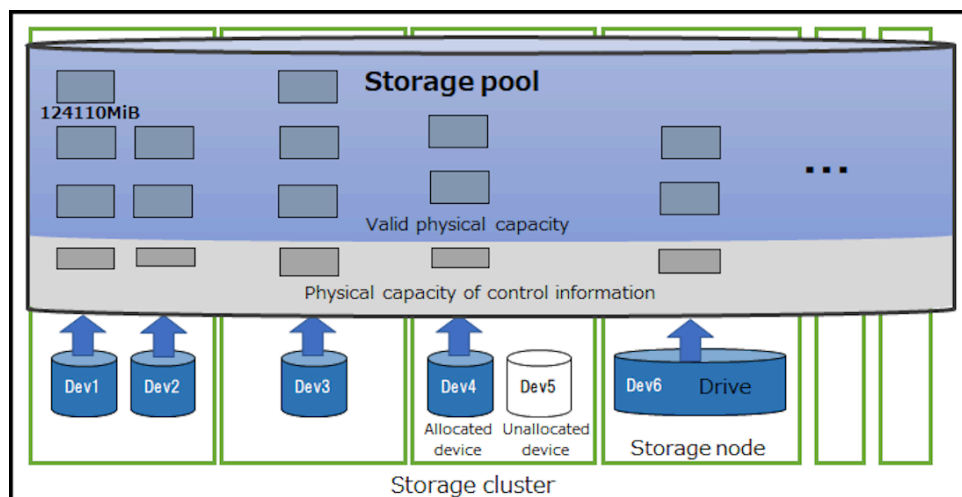
The physical capacity in drives is classified into the following two categories:

- Valid physical capacity (totalRawCapacity): The amount of space that has been available since the drive was added to the storage pool.
- Physical capacity of control information (metaDataPhysicalCapacity): The capacity for control information in the total capacity (total physical capacity) of the drives allocated to the storage pool.

For HPEC 4D+1P, the effective physical capacity ($C_{rawDevice}$)[MiB] is determined by the following calculations:

```
( $C_{rawDevice}$ ) = floor(( $C_{device} \times 0.9846$ ) - 2048[MiB] - 124110[MiB], 124110[MiB])
```

floor(value, reference value)
: The value is rounded off to a multiple close to the reference value.
 C_{device} : Drive capacity[MiB]



Capacity design principles (2) : Logical capacity (for HPEC 4D+1P)

The following assumptions are assumed:

- Each storage node (or drive) in the storage cluster has the same valid physical capacity.
- The cluster does not have faulty storage nodes or faulty drives.
- The configuration and the capacity of each storage node satisfy the product specifications.

Logical capacity

Logical capacity (TotalCapacity) is the amount of storage available to users in a storage pool. To allow Capacity balance to resolve imbalance among the storage controllers, configure logical capacity (totalCapacity) with adequate margin so that planned capacity usage is kept to 80% or below.

You can calculate logical capacity (totalCapacity) by using the formula shown further on in this section.

First, calculate physical capacity used as logical capacity (usablePhysicalCapacity).

Subtract the capacity reserved for configuration changes (620550 MiB) and the capacity allocated as the rebuild capacity ($C_{rawDevice} \times N_{rebuild}$) from the valid physical capacity of storage node. Out of the remaining capacity, capacity for multiples of 620550 MiB is used for physical capacity used as logical capacity (usablePhysicalCapacity).

```
physical capacity used as logical capacity [MiB/Node] = RoundDown((CrawNode-620550-
CrawDevice×Nrebuild) / 620550) × 620550
CrawNode: Valid physical capacity [MiB/Node]
CrawDevice: Valid physical capacity [MiB/Device]
Nrebuild: Tolerable number of drive failures set for the storage pool
(0 if the rebuild capacity policy is set to "Variable")
RoundDown (numerical value): A numerical value is rounded up to an integer.
```

Out of every 620550 MiB of physical capacity used as logical capacity, 496440 MiB is used for logical capacity.

In addition, each storage controller needs 4200MiB of managed capacity, which is excluded from the logical capacity.

```
Logical capacity [MiB] = ((CusableNode / 620550 × 496440) - 4200) × Nnode
CusableNode: Physical capacity used as logical capacity [MiB/Node]
Nnode: Number of storage nodes [number of storage nodes / cluster]
```

As a result of the calculation, if the logical capacity is 0 or less, the logical capacity will not be configured. However, when the calculated logical capacity is 0 or less while $N_{rebuild}$ is 1 (default value), if the logical capacity can be formed by setting $N_{rebuild}$ to 0 and then recalculating, the logical capacity will be formed without securing the rebuild capacity.

The capacity that a storage controller can manage has an upper limit. The capacity exceeding the following values cannot be used as logical capacity.

- Physical capacity used as logical capacity 125971650 [MiB]
- Logical capacity: 100773120 [MiB/Node]

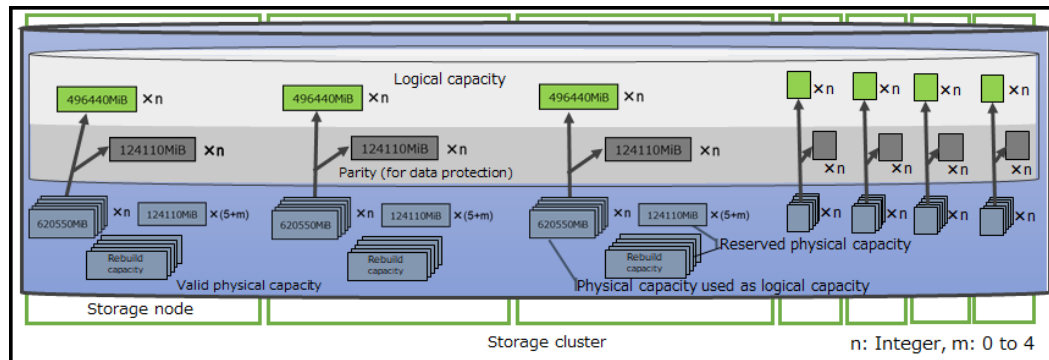
Reserved physical capacity

Physical reserved capacity (reservedPhysicalCapacity) is the amount of available physical capacity in the storage pool that is not used for logical capacity. The physical reserved capacity is determined by the following calculation:

$$\text{Reserved physical capacity [MiB]} = (C_{\text{rawNode}} - C_{\text{usableNode}}) \times N_{\text{node}}$$

C_{rawNode} : Valid physical capacity [MiB/Node]
 $C_{\text{usableNode}}$: Physical capacity used as logical capacity [MiB/Node]
 N_{node} : Number of storage nodes [number of storage nodes / cluster]

* Part truncated by RoundDown in the logical capacity equation $\times 620550$



Note:

The logicalized physical capacity (usablePhysicalCapacity) shown in the figure above is the amount of effective physical capacity in the storage pool that is available for the logical capacity.

Physical reserved capacity (reservedPhysicalCapacity) is the amount of effective physical capacity in the storage pool that is not used for logical capacity.

Capacity design method (1)-1 : Logical capacity from physical capacity (for HPEC 4D+1P)

The following is assumed:

- Each storage node (or drive) in the storage cluster has the same valid physical capacity.
- The storage cluster does not have faulty storage nodes or faulty drives.
- The configuration and the capacity of each storage node satisfy the product specifications.

The following example procedure assumes that the configuration consists of five storage nodes and each storage node has a capacity of 16 TB (1.6-TB drive × 10) with the rebuild capacity policy set to "Fixed" and the tolerable number of drive failures set to 1.

Procedure

1. Convert the unit for drive capacity to TiB:

$$1.6 \times 1000^4 / 1024^4 = 1.4552[\text{TiB}]$$

2. Calculate the valid physical capacity per drive (1.6 TB) by using the expression in section *Capacity design principles (1) : Physical capacity (for HPEC 4D+1P)*.

```
floor((1.4552 × 10242 × 0.9846) - 2048[MiB] - 124110[MiB], 124110[MiB])
= 1365210 [MiB/device]
floor(value, reference value)
: The value is rounded off to a multiple close to the reference
value.
```

3. Calculate the valid physical capacity per storage node.

$$1365210 \times 10 = 13652100[\text{MiB/node}]$$

4. Calculate the physical capacity used as logical capacity of the storage cluster by using the expression shown in *Capacity design principles (2) : Logical capacity (for HPEC 4D +1P)*.

```
RoundDown((13652100 - 620550 - 1365210) / 620550) × 496440 - 4200) × 5
= 44658600[MiB/Cluster]
RoundDown( numerical value )
: A numerical value is rounded down to an integer.
```

5. Calculate the logical capacity of the storage cluster by using the expression shown in *Capacity design principles (2) : Logical capacity (for HPEC 4D+1P)*.

$$(11169900 / 620550 \times 496440 - 4200) \times 5 = 44658600[\text{MiB/Cluster}]$$

Capacity design method (2) : Physical capacity from logical capacity (for HPEC 4D+1P)

The following is assumed:

- Each storage node (or drive) in the storage cluster has the same valid physical capacity.
- The storage cluster does not have faulty storage nodes or faulty drives.
- The configuration and the capacity of each storage node satisfy the product specifications.

In the following example procedure, you calculate the physical capacity required per storage node when a logical capacity of 50 TiB is allocated to five storage nodes (20 device slots), the rebuild capacity policy is set to "Fixed", and the tolerable number of drive failures is set to 1.

Procedure

1. Since the logical capacity can be increased in 496440-MiB increments (section *Capacity design principles (2) : Logical capacity (for HPEC 4D+1P)*), the valid physical capacity required in each storage node is:

```
RoundUp(((50 × 10242) / 5 + 4200) × (620550 / 496440)) + 620550
= 13733000[MiB/node]
RoundUp( numerical value )
: A numerical value is rounded up to an integer.
```

2. Calculate the valid physical capacity of the used drives.

For example, the valid physical capacity of a drive (1.6 TB) is 1365210 [MiB/device].

Since the capacity for one drive is allocated as the rebuild capacity, the number of necessary drives can be calculated by using the following expression:

```
RoundUp(13733000 / 1365210) + 1 =12[Drives/Node]
RoundUp( numerical value )
: A numerical value is rounded up to an integer.
```

Through the preceding calculation, it is found that 50-TiB logical capacity is available by constructing a storage cluster consisting of five storage nodes with twelve 1.6-TB drives.

Appendix B: Capacity design (for HPEC 4D+2P)


Capacity design principles (1) : Physical capacity (for HPEC 4D+2P)

The physical capacity in drives is classified into the following two categories:

- Valid physical capacity (totalRawCapacity): The amount of space that has been available since the drive was added to the storage pool.
- Physical capacity of control information (metaDataPhysicalCapacity): The capacity for control information in the total capacity (total physical capacity) of the drives allocated to the storage pool.

For HPEC 4D+2P, the effective physical capacity ($C_{rawDevice}$ [MiB]) is determined by the following calculations:

```
 $C_{rawDevice} = \text{floor}((C_{device} \times 0.9846) - 2048[\text{MiB}] - 148932[\text{MiB}], 148932[\text{MiB}])$   
 $\text{floor}(\text{value}, \text{reference value})$   
: The value is rounded off to a multiple close to the reference value.  
 $C_{device}$ : Drive capacity[MiB]
```



Capacity design principles (2) : Logical capacity (for HPEC 4D+2P)

The following assumptions are assumed:

- Each storage node (or drive) in the storage cluster has the same valid physical capacity.
- The cluster does not have faulty storage nodes or faulty drives.
- The configuration and the capacity of each storage node satisfy the product specifications.

Logical capacity

The logical capacity (totalCapacity) is the storage pool capacity available to users. To allow Capacity balance to resolve imbalance among the storage controllers, configure logical capacity (totalCapacity) with adequate margin so that planned capacity usage is kept to 80% or below.

You can calculate logical capacity (totalCapacity) by using the formula shown further on in this section.

First, calculate physical capacity used as logical capacity (usablePhysicalCapacity).

Subtract the capacity reserved for configuration changes (893592MiB MiB) and the capacity allocated as the rebuild capacity ($C_{rawDevice} \times N_{rebuild}$) from the valid physical capacity of each storage node. Out of the remaining capacity, capacity for multiples of 893592MiB is used for physical capacity used as logical capacity (usablePhysicalCapacity).

```
Physical capacity used as logical capacity [MiB/Node] = RoundDown((CrawNode - 893592 -
CrawDevice × Nrebuild) / 893592) × 893592
CrawNode: Valid physical capacity [MiB/Node]
CrawDevice: Valid physical capacity [MiB/Device]
Nrebuild: Tolerable number of drive failures set for the storage pool
(0 if the rebuild capacity policy is set to "Variable")
RoundDown (numerical value): A numerical value is rounded up to an integer.
```

Out of every 893592MiB of physical capacity used as logical capacity, 595728 MiB is used for logical capacity.

In addition, each storage controller needs 4200MiB of managed capacity, which is excluded from the logical capacity.

```
Logical capacity [MiB] = ((CusableNode / 893592 × 595728) - 4200) × Nnode
CusableNode: Physical capacity used as logical capacity [MiB/Node]
Nnode: Number of storage nodes [number of storage nodes / cluster]
```

As a result of the calculation, if the logical capacity is 0 or less, the logical capacity will not be configured. However, when the calculated logical capacity is 0 or less while $N_{rebuild}$ is 1 (default value), if the logical capacity can be formed by setting $N_{rebuild}$ to 0 and then recalculating, the logical capacity will be formed without securing the rebuild capacity.

The capacity that a storage controller can manage has an upper limit. The capacity exceeding the following values cannot be used as logical capacity.

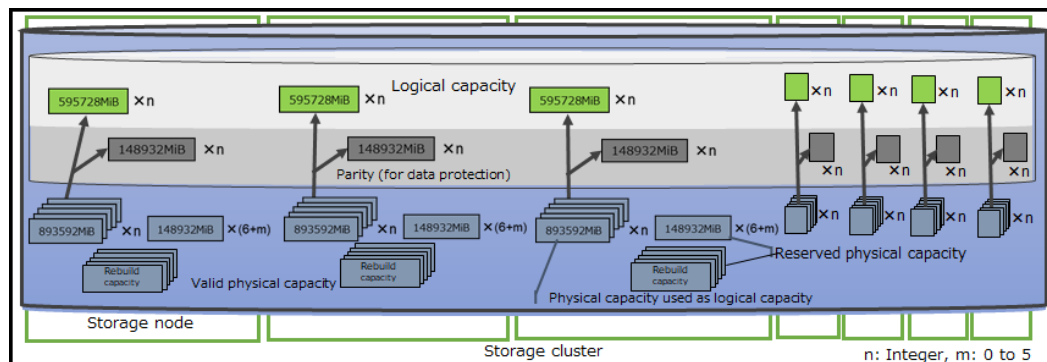
- Physical capacity used as logical capacity 151017048 [MiB/Node]
- Logical capacity: 100673832 [MiB/Node]

Reserved physical capacity

Physical reserved capacity (reservedPhysicalCapacity) is the amount of available physical capacity in the storage pool that is not used for logical capacity. The physical reserved capacity is determined by the following calculation:

$$\text{Reserved physical capacity [MiB]} = (C_{\text{rawNode}} - C_{\text{usableNode}}) \times N_{\text{node}}$$

C_{rawNode} : Valid physical capacity [MiB/Node]
 $C_{\text{usableNode}}$: Physical capacity used as logical capacity [MiB/Node]
 N_{node} : Number of storage nodes [number of storage nodes / cluster]



Tip:

The logicalized physical capacity (usablePhysicalCapacity) shown in the figure above is the amount of effective physical capacity in the storage pool that is available for the logical capacity.

Physical reserved capacity (reservedPhysicalCapacity) is the amount of effective physical capacity in the storage pool that is not used for logical capacity.

Capacity design method (1)-1 : Logical capacity from physical capacity (for HPEC 4D+2P)

The following is assumed:

- Each storage node (or drive) in the storage cluster has the same valid physical capacity.
- The storage cluster does not have faulty storage nodes or faulty drives.
- The configuration and the capacity of each storage node satisfy the product specifications.

The following example procedure assumes that the configuration consists of six storage nodes and each storage node has a capacity of 16 TB (1.6-TB drive × 10) with the rebuild capacity policy set to "Fixed" and the tolerable number of drive failures set to 1.

Procedure

1. Convert the unit for drive capacity to TiB:

$$1.6 \times 1000^4 / 1024^4 = 1.4552[\text{TiB}]$$

2. Calculate the valid physical capacity per drive (1.6 TB) by using the expression in section *Capacity design principles (1) : Physical capacity (for HPEC 4D+2P)*.

```
floor((1.4552 × 10242 × 0.9846) - 2048[MiB] - 148932[MiB], 148932[MiB])
= 1340388 [MiB/device]
floor(value, reference value)
: The value is rounded off to a multiple close to the reference
value.
```

3. Calculate the valid physical capacity per storage node.

$$1340388 \times 10 = 13403880[\text{MiB/node}]$$

4. Calculate the physical capacity used as logical capacity of the storage cluster by using the expression shown in *Capacity design principles (2) : Logical capacity (for HPEC 4D+2P)*.

```
(RoundDown((13403880 - 893592 - 1340388) / 893592) × 893592 = 10723104[MiB]
RoundDown( numerical value )
: A numerical value is rounded down to an integer.
```

5. Calculate the logical capacity of the storage cluster by using the expression shown in *Capacity design principles (2) : Logical capacity (for HPEC 4D+2P)*.

$$(10723104 / 893592 \times 595728 - 4200) \times 6 = 42867216[\text{MiB/Cluster}]$$

Capacity design method (2) : Physical capacity from logical capacity (for HPEC 4D+2P)

The following is assumed:

- Each storage node (or drive) in the storage cluster has the same valid physical capacity.
- The storage cluster does not have faulty storage nodes or faulty drives.
- The configuration and the capacity of each storage node satisfy the product specifications.

In the following example procedure, you calculate the physical capacity required per storage node when a logical capacity of 50 TiB is allocated to six storage nodes (20 device slots). The rebuild capacity policy is set to "Fixed", and the number of tolerable drive failures is set to 1.

Procedure

1. Since the logical capacity can be increased in 496440-MiB increments (section *Capacity design principles (2) : Logical capacity (for HPEC 4D+2P)*), the valid physical capacity required in each storage node is:

```
RoundUp(((50 × 10242) / 6 + 4200) × (893592 / 595728)) + 893592
= 14007092[MiB/node]
RoundUp( numerical value )
: A numerical value is rounded up to an integer.
```

2. Calculate the valid physical capacity of the used drives.

For example, the valid physical capacity of a drive (1.6 TB) is 1340388 [MiB/device].

Because the capacity for one drive is allocated as the rebuild capacity, the number of necessary drives can be calculated by using the following expression:

```
RoundUp(14007092 / 1340388) + 1 = 12[Drives/Node]
RoundUp(value): Result of rounding value up to an integer
```

Through the preceding calculation, it is found that 50-TiB logical capacity is available by constructing a storage cluster consisting of six storage nodes with twelve 1.6-TB drives.

Appendix C: Capacity design (for Mirroring)

Capacity design principles (1) : Physical capacity (for Mirroring)


The physical capacity of a drive consists of the valid physical capacity and the physical capacity of control information.

- Valid physical capacity (totalRawCapacity): Capacity that becomes available as a result of adding the drive to a storage pool
- Physical capacity of control information (metaDataPhysicalCapacity): The capacity for storing control information in the drive's total physical capacity allocated to the storage pool

For Mirroring, the valid physical capacity ($C_{rawDevice}$ [MiB]) can be calculated by using the following expression:

```
( $C_{rawDevice}$ ) = floor(( $C_{device} \times 0.9846$ ) - 3072[MiB] - 99288[MiB], 99288[MiB])
```

floor(value, reference value)
: The value is rounded off to a multiple close to the reference value.
 C_{device} : Drive capacity[MiB]



Capacity design principles (2) : Logical capacity (for Mirroring)

The following assumptions are assumed:

- Each storage node (or drive) in the storage cluster has the same valid physical capacity.
- The cluster does not have faulty storage nodes or faulty drives.
- The configuration and the capacity of each storage node satisfy the product specifications.

Logical capacity

Logical capacity (TotalCapacity) is the amount of storage available to users in a storage pool. To allow Capacity balance to resolve imbalance among the storage controllers, configure logical capacity (totalCapacity) with adequate margin so that planned capacity usage is kept to 80% or below.

You can calculate logical capacity (totalCapacity) by using the formula shown further on in this section.

First, calculate physical capacity used as logical capacity (usablePhysicalCapacity).

Subtract the capacity reserved for configuration changes (198576MiB MiB) and the capacity allocated as the rebuild capacity ($C_{rawDevice} \times N_{rebuild}$) from the valid physical capacity of each storage node. Out of the remaining capacity, capacity for multiples of 198576 MiB is used for physical capacity used as logical capacity (usablePhysicalCapacity).

```
Capacity used as logical capacity [MiB/Node] = RoundDown((CrawNode - 198576 - CrawDevice
× Nrebuild) / 198576) × 198576
CrawNode: Valid physical capacity [MiB/Node]
CrawDevice: Valid physical capacity [MiB/Device]
Nrebuild: Tolerable number of drive failures set for the storage pool
(0 if the rebuild capacity policy is set to "Variable")
RoundDown (numerical value): A numerical value is rounded up to an integer.
```

Out of every 198576 MiB of physical capacity used as logical capacity, 99288 MiB is used for logical capacity.

Each storage controller requires a capacity of 4200 MiB for management use. This capacity is excluded from the logical capacity.

```
Logical capacity [MiB] = ((CusableNode / 198576 × 99288) - 4200) × Nnode
CusableNode: Physical capacity used as logical capacity [MiB/Node]
Nnode: Number of storage nodes [number of storage nodes / cluster]
```

As a result of the calculation, if the logical capacity is 0 or less, the logical capacity will not be configured. However, when the calculated logical capacity is 0 or less while $N_{rebuild}$ is 1 (default value), if the logical capacity can be formed by setting $N_{rebuild}$ to 0 and then recalculating, the logical capacity will be formed without securing the rebuild capacity.

The capacity that a storage controller can manage has an upper limit. The capacity exceeding the following values cannot be used as logical capacity.

- Physical capacity used as logical capacity 203143248 [MiB/Node]
- Logical capacity: 101567424 [MiB/Node]

Reserved physical capacity

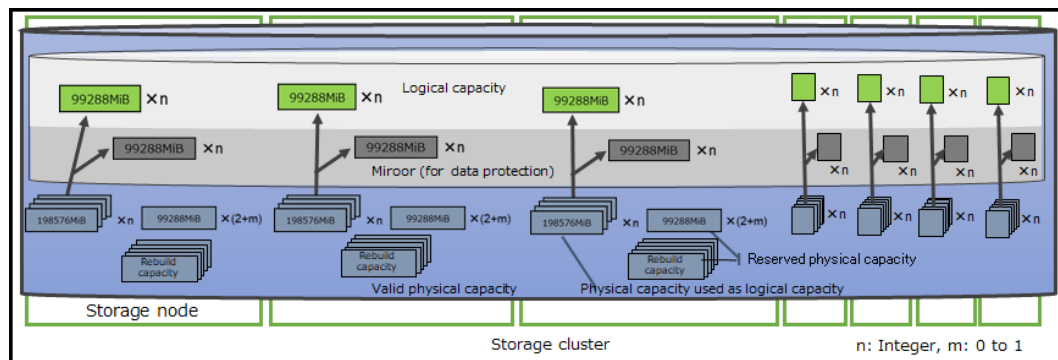
Physical reserved capacity (reservedPhysicalCapacity) is the amount of available physical capacity in the storage pool that is not used for logical capacity. The physical reserved capacity is determined by the following calculation:

$$\text{Reserved physical capacity [MiB]} = (C_{\text{rawNode}} - C_{\text{usableNode}}) \times N_{\text{node}}$$

C_{rawNode} : Valid physical capacity [MiB/Node]

$C_{\text{usableNode}}$: Physical capacity used as logical capacity [MiB/Node]

N_{node} : Number of storage nodes [number of storage nodes / cluster]



Tip:

The logicalized physical capacity (usablePhysicalCapacity) shown in the figure above is the amount of effective physical capacity in the storage pool that is available for the logical capacity.

Physical reserved capacity (reservedPhysicalCapacity) is the amount of effective physical capacity in the storage pool that is not used for logical capacity.

Capacity design method (1)-1 : Logical capacity from physical capacity (for Mirroring)

The following is assumed:

- Each storage node (or drive) in the storage cluster has the same valid physical capacity.
- The storage cluster does not have faulty storage nodes or faulty drives.
- The configuration and the capacity of each storage node satisfy the product specifications.

The following example procedure assumes that the configuration consists of six storage nodes and each storage node has a capacity of 16 TB (1.6-TB drive × 10) with the rebuild capacity policy set to "Fixed" and the tolerable number of drive failures set to 1.

Procedure

1. Convert the unit for drive capacity to TiB:

$$1.6 \times 1000^4 / 1024^4 = 1.4552[\text{TiB}]$$

2. Calculate the valid physical capacity per drive (1.6 TB) by using the expression in section *Capacity design principles (1) : Physical capacity (for Mirroring)*.

```
floor((1.4552 × 10242 × 0.9846) - 3072[MiB] - 99288[MiB], 99288[MiB])
= 1390032 [MiB/device]
floor(value, reference value)
: The value is rounded off to a multiple close to the reference
value.
```

3. Calculate the valid physical capacity per storage node.

$$1390032 \times 10 = 13900320[\text{MiB/node}]$$

4. Calculate the physical capacity used as logical capacity of the storage cluster by using the expression shown in *Capacity design principles (2) : Logical capacity (for Mirroring)*.

```
(RoundDown((13900320 - 198576 - 1390032) / 198576) × 99288 - 4200) × 6 =
36909936[MiB/Cluster]
RoundDown( numerical value )
: A numerical value is rounded down to an integer.
```

5. Calculate the logical capacity of the storage cluster by using the expression shown in *Capacity design principles (2) : Logical capacity (for Mirroring)*.

$$(12311712 / 198576 \times 99288 - 4200) \times 6 = 36909936[\text{MiB/Cluster}]$$

Capacity design method (2) : Physical capacity from logical capacity (for Mirroring)

The following is assumed:

- Each storage node (or drive) in the storage cluster has the same valid physical capacity.
- The storage cluster does not have faulty storage nodes or faulty drives.
- The configuration and the capacity of each storage node satisfy the product specifications.

In the following example procedure, you calculate the physical capacity required per storage node when a logical capacity of 50 TiB is allocated to six storage nodes (20 device slots). The rebuild capacity policy is set to "Fixed", and the number of tolerable drive failures is set to 1.

Procedure

1. Since the logical capacity can be increased in 99288-MiB increments (section *Capacity design principles (2) : Logical capacity (for Mirroring)*), the valid physical capacity required in each storage node is:

```
RoundUp(((50 × 10242) / 6 + 4200) × (198576 / 99288)) + 198576
= 17683243[MiB/node]
RoundUp( numerical value )
: A numerical value is rounded up to an integer.
```

2. Calculate the valid physical capacity of the used drives.

For example, the valid physical capacity of a drive (1.6 TB) is 1340388 [MiB/device].

Since the capacity for one drive is allocated as the rebuild capacity, the number of necessary drives can be calculated by using the following expression:

```
RoundUp(17683243 / 1340388) + 1 = 15 [drives/node]
RoundUp( numerical value )
: A numerical value is rounded up to an integer.
```

Through the preceding calculation, it is found that 50-TiB logical capacity is available by constructing a storage cluster consisting of six storage nodes with fifteen 1.6-TB drives.

Appendix D: Capacity design when adding storage nodes

Adding capacity

The following two ways can be used to add capacity to Virtual Storage Software block as follows.

Adding storage node

By adding storage nodes to the storage cluster, increase the total capacity of the storage cluster.

- Advantage: Adding storage nodes might improve the performance of a storage cluster, depending on the IO pattern.
- Disadvantage: It costs more than the case of adding drives. (Hardware cost, installation cost, etc.)
- Disadvantages: Before the added capacity is available, it takes time to complete the process of relocating drive data and allocating capacity to the storage controller.

Adding drives

By adding drives to the existing storage nodes, increase the total capacity of the storage cluster.

- Advantage: It costs less than the case of adding nodes because it only needs installation cost of drives.
- Disadvantage: The number of drives for storage node is limited.

Capacity design when adding storage nodes (for HPEC 4D +1P)

Describe about capacity design when adding storage nodes for HPEC 4D+1P.

The following is assumed:

- The effective physical capacity of the added storage node is uniform with the storage node in the pre-expansion configuration.
- The cluster does not have faulty storage nodes or faulty drives.
- The configuration and the capacity of each storage node satisfy the product specifications.

Procedure

1. Under the assumptions, the increased logical capacity of the expansion storage node is the same as the logical capacity of the storage node in the pre-expansion configuration, so the total logical capacity can be calculated using the following formula:


Total logical capacity after adding storage nodes = $C \times (N + A) / N$

C: Total logical capacity before adding storage nodes

N: Number of storage nodes before adding storage nodes

A: Number of added storage nodes

(storage nodes can be added on a per-node basis)



Capacity design when adding drives (for HPEC 4D+1P)

Describe about capacity design when adding drives for HPEC 4D+1P.

The following is assumed:

- The effective physical capacity of the added drives is uniform and is added to all storage nodes.
- The cluster does not have faulty storage nodes or faulty drives.
- The configuration and the capacity of each storage node satisfy the product specifications.

Procedure

1. Calculate the valid physical capacity of the added drive by using the method described in *Capacity design principles (1) : Physical capacity (for HPEC 4D+1P)*.

2. The increased logical capacity of the added drive can be calculated by:


```

Total logical capacity after adding drives[MiB]
= C + RoundDown((R + A - 620550 - CrawDevice × Nrebuild) / 620550) × N × 496440
C: Total logical capacity before adding drives
RoundDown( numerical value )
      : A numerical value is rounded down to an integer.
R: Reserved physical capacity of each storage node before adding
  drives
A: Valid physical capacity of each storage nodes before adding
  drives
N: Number of storage nodes
CrawDevice: Valid physical capacity [MiB/device]
Nrebuild: Tolerable number of drive failures set for the storage pool (0 if
the rebuild capacity policy is set to "Variable")

```

Note that the capacity that a storage controller can manage has an upper limit. The capacity exceeding the following values cannot be used as logical capacity.

```
100773120 × the number of storage nodes [MiB]
```



Capacity design when adding storage nodes (for HPEC 4D+2P)

Describe about capacity design when adding storage nodes for HPEC 4D+2P.

The following is assumed:

- The effective physical capacity of the added storage node is uniform with the storage node in the pre-expansion configuration.
- The cluster does not have faulty storage nodes or faulty drives.
- The configuration and the capacity of each storage node satisfy the product specifications.

Procedure

- Under the assumptions, the increased logical capacity of the expansion storage node is the same as the logical capacity of the storage node in the pre-expansion configuration, so the total logical capacity can be calculated using the following formula:

Total logical capacity after adding storage nodes = $C \times (N + A) / N$

C : Total logical capacity before adding storage nodes

N : Number of storage nodes before adding storage nodes

A : Number of added storage nodes

(storage nodes can be added on a per-node basis)

**Capacity design when adding drives (for HPEC 4D+2P)**

Describe about capacity design when adding drives for HPEC 4D+2P.

The following is assumed:

- The effective physical capacity of the added drives is uniform and is added to all storage nodes.
- The cluster does not have faulty storage nodes or faulty drives.
- The configuration and the capacity of each storage node satisfy the product specifications.

Procedure

- Calculate the valid physical capacity of the added drive by using the method described in *Capacity design principles (1) : Physical capacity (for HPEC 4D+2P)*.
- The increased logical capacity of the added drive can be calculated by:

Total logical capacity after adding drives[MiB]

= $C + \text{RoundDown}((R + A - 893592 - C_{\text{rawDevice}} \times N_{\text{rebuild}}) / 893592) \times N \times 595728$

C : Total logical capacity before adding drives

RoundDown(numerical value)

: A numerical value is rounded down to an integer.

R : Reserved physical capacity of each storage node before adding drives

A : Valid physical capacity of the added drives


N : Number of storage nodes

$C_{\text{rawDevice}}$: Valid physical capacity [MiB/device]

$N_{rebuild}$: Tolerable number of drive failures set for the storage pool (0 if the rebuild capacity policy is set to "Variable")

Note that the capacity that a storage controller can manage has an upper limit. The capacity exceeding the following values cannot be used as logical capacity.

$100673832 \times \text{the number of storage nodes [MiB]}$



Capacity design when adding storage nodes (for Mirroring)

Describe about capacity design when adding storage nodes for Mirroring.

The following is assumed:

- The effective physical capacity of the added storage node is uniform with the storage node in the pre-expansion configuration.
- The cluster does not have faulty storage nodes or faulty drives.
- The configuration and the capacity of each storage node satisfy the product specifications.

Procedure


1. Under the assumptions, the increased logical capacity of the expansion storage node is the same as the logical capacity of the storage node in the pre-expansion configuration, so the total logical capacity can be calculated using the following formula:

Total logical capacity after adding storage nodes = $C \times (N + A) / N$

C: Total logical capacity before adding storage nodes

N: Number of storage nodes before adding storage nodes

A: Number of added storage nodes
(storage nodes can be added on a per-node basis)



Capacity design when adding drives (for Mirroring)

Describe about capacity design when adding drives for Mirroring.

The following is assumed:

- The effective physical capacity of the added drives is uniform and is added to all storage nodes.
- The cluster does not have faulty storage nodes or faulty drives.
- The configuration and the capacity of each storage node satisfy the product specifications.

Procedure


1. Calculate the valid physical capacity of the added drive by using the method described in *Capacity design principles (1) : Physical capacity (for Mirroring)*.
2. The increased logical capacity of the added drive can be calculated by:

```
Total logical capacity after adding drives[MiB]
= C + RoundDown((R + A - 198576 - CrawDevice × Nrebuild) / 198576) × N × 99288
C: Total logical capacity before adding drives
RoundDown( numerical value )
    : A numerical value is rounded down to an integer.
R: Reserved physical capacity of each storage node before adding
    drives
A: Valid physical capacity of each storage nodes before adding
    drives
N: Number of storage nodes
CrawDevice: Valid physical capacity [MiB/device]
Nrebuild: Tolerable number of drive failures set for the storage pool (0 if
the rebuild capacity policy is set to "Variable")
```

Note that the capacity that a storage controller can manage has an upper limit. The capacity exceeding the following values cannot be used as logical capacity.

```
1001567424 × the number of storage nodes [MiB]
```

Capacity design when adding drives (for Mirroring)



Appendix E: ALUA configuration guidelines

Required preparations

VMware

Common settings

In the case of VMware ESXi 6.X (including 5.X), apply the following updates and KBs.

- VMware ESXi5.5 Update3 and KB2146717
- VMware ESXi6.0 Update2 and KB2145663

Note that you do not need to change the setting value for SCSI Timeout because the SCSI Timeout value is 120 seconds by default.

Alternate path software setting

No additional settings are required.

Inter-guest cluster setting

For a device in which the inter-guest cluster (MSCS, etc.) uses reserved volumes (SCSI-2, SCSI-3), change "perennially-reserved" to TRUE from each VMware ESXi (see each KB). Make this setting irrespective of the storage configuration.

- ESXi/ESX hosts with visibility to RDM LUNs being used by MSCS nodes with RDMS may take a long time to start or during LUN rescan (1016106)

<https://kb.vmware.com/s/article/1016106>

Windows

Common settings

Change the setting value of SCSI Timeout by referring to *Installing a compute node* in this manual.

Alternate path software setting (in the case of MPIO)

You must change the registry settings. Set the following registry values (see the following URL):

[http://technet.microsoft.com/en-us/library/ee619749\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/ee619749(v=ws.10).aspx)

- Set PathVerifyEnabled to 1.
- Set UseCustomPathRecoveryInterval to 1.
- Set PathRecoveryInterval (default: 40) to a value smaller than PDORemovePeriod (default: 20).

Edit registry values carefully (see the following URL):

<http://support.microsoft.com/kb/136393>

Cluster settings

No additional settings are required.

Linux

Change the setting value of SCSI Timeout and multipath settings by referring to *Installing a compute node* in this manual.

ALUA setting procedure

VMware (in the case of NMP)

This procedure allows the compute node to recognize whether the ALUA is enabled (specified on the storage side).

No procedure is available for allowing the compute node to recognize an ALUA priority change set on the storage side. ALUA priority changes are recognized automatically.

Procedure

1. Verify that a dedicated SATP has not been created yet.

If the following command outputs nothing, you can judge that an SATP has not been created. If an SATP has already been created, you do not need to perform the following steps.

```
# esxcli storage nmp satp rule list | grep Hi-SDS
```

2. Create a dedicated SATP.

```
# esxcli storage nmp satp rule add -V HITACHI -M "^Hi-SDS*" -s VMW_SATP_ALUA -c tpgs_on --psp="VMW_PSP_RR"
```

3. Reboot the compute node to allow the setting to be recognized.

Windows (in the case of MPIO)

No procedure is available for allowing the compute node to recognize whether ALUA is enabled (specified on the storage side).

To allow the compute node to recognize an ALUA priority change set on the storage side, scan hardware changes from the device manager. Run the following command for verification.

```
# mpclaim -s -d XXX
```

Priority paths and non-priority paths are indicated as follows:

Priority path: Active/Optimized

Non-priority path: Active/Unoptimized

Linux system (in the case of Device Mapper)

The following procedure allows the compute node to recognize whether the ALUA is enabled (specified on the storage side).

Procedure

1. Change the settings for the following items in /etc/multipath.conf:

Item	Setting value
vendor	"HITACHI"
product	"Hi-SDS*"
path_grouping_policy	group_by_prio
prio	alua
detect_path_checker/ detect_checker ¹	no
path_checker ²	readsector0
dev_loss_tmo	Refer to your distribution's documentation and set the maximum value.
<p>1. The parameter name differs depending on your distribution. Refer to your distribution's documentation and set the appropriate parameter name.</p> <p>2. If a failure occurs in a storage node of Virtual Storage Software block, when a value other than readsector0 is set, paths might be wrongly blocked. For this reason, make sure that you set readsector0.</p>	

2. Reload the alternate path daemon.

```
# service multipathd reload
```

Example

No procedure is available for letting the compute node recognize an ALUA priority change set on the storage side. ALUA priority changes are recognized automatically.

You can verify the ALUA priority set on the storage by running the following command:

```
# multipath -ll
```

Priority paths and non-priority paths are indicated as follows:

Priority path: prio=50 status=active

Non-priority path: prio=10 status=enable

Virtual Storage Software block does not support a mixed configuration of ALUA devices and non-ALUA devices.

Appendix F: Processing time for snapshot operations

Estimated processing time for snapshot operations

The processing time (in seconds) described here is estimated on the assumption that one snapshot operation is performed while no I/O is issued and no other operation is being performed. Estimate the processing time based on the table below.

- The processing time for restoration and deletion varies depending on totalCapacity and snapshotConcordanceRate of the volume.
- The processing time for PrepareAndFinalize and Prepare includes the time required for S-VOL creation.
- The processing time for restoration might be about twice as long as the time described in the table below depending on the target generation of restoration and the amount of differential data between the copy-source volume and the copy-destination volume of restoration. To restore data from a generation other than the latest one or from a generation of S-VOL to which data is written, estimate the processing time by doubling it.
- If the load of the copy processing of restoration has a high impact on I/O, the copy processing might be stopped by the storage controller. If the copy processing does not progress even if the storage node is not overloaded, collect the dump log file for the storage node and contact customer support.
- A maximum of 500 restore operations can be requested at the same time in a system. This upper limit is set because the time required for restoration may take long due to the match rate or suppression of copy processing. If more than 500 restore operations are requested, the KARS06143-E message will be returned, and the restore operations will fail.

totalCapacity [GiB]	snapshotConcordanceRate [%]	Prepare and take a snapshot (PrepareAndFinalize) [second]	Prepare(Prepare) [second]	Take a snapshot (Finalize) [second]	Restore[s econd]	Delete[sec ond]
51200	100	36	38	0	9727	1350
16384	100	14	13	0	3723	573
4096	100	8	8	0	961	162
1024	100	7	6	0	241	54
256	100	6	6	0	61	29
128	100	8	6	0	60	22
128	75	-	-	-	2583	44
128	50	-	-	-	5045	74
128	0	-	-	-	9729	101
-: Not available						

You can delete all snapshot volumes (S-VOL or P/S-VOL) created from the P-VOL by running the snapshot deletion command with the volume ID of the P-VOL and true specified for masterVolumeId and snapshotTree (for the REST API) and for --master_volume_id and --snapshot_tree (for the CLI).

The deletion time can be calculated as follows:

Number of snapshot volumes¹ x deletion time per volume + snapshot configuration deletion time + 600 seconds²

1. Number of all the S-VOLs or P/S-VOLs that were created from the specified P-VOL.
2. Waiting time for copy stop processing required only when conflict with snapshot copy processing within the system occurs.

The deletion time per volume and snapshot configuration deletion time vary depending on the size of the P-VOL as follows:

Size of the P-VOL [GiB]	Deletion time per volume [second]	Snapshot configuration deletion time [second]
128	10	4
256	10	4

Size of the P-VOL [GiB]	Deletion time per volume [second]	Snapshot configuration deletion time [second]
1024	10	4
4096	11	10
8192	11	15
16384	15	32
32768	19	64
51200	24	92
65536	29	107
98304	41	161

Glossary

Auto recovery

See *Storage node auto-recovery* in the Glossary.

base license

A license that provides basic functionality.

blocked, blocking, blockage

A state for a storage or resources that comprise a storage where I/O operations cannot be performed.

BMC network

Network that connects the storage node BMC and the controller node. This network is used to operate the BMC from the controller node.

BMC port

The port that is on a storage node and is used for connection to the BMC network.

capacity balancing

Function of moving volumes automatically from high capacity usage storage controllers to low capacity usage storage controllers when capacity usage is not balanced among storage controllers.

cluster master node (primary)

A storage node within the storage cluster that has the role of managing the entire storage cluster.

cluster master node (secondary)

A storage node in the storage cluster that is responsible for managing the entire storage cluster in the event of failure of the cluster master node (primary).

cluster worker node

A storage node in the storage cluster that does not have the role of managing the entire storage cluster.

compute network

A network between a compute node and a storage node. Used for input / output of user data.

compute node

A node that the application of the user operates and instructs input / output of user data to the storage node. A host connected to the compute port.

compute port

(Virtual machine) The virtual port that is on a storage node and connects to the compute network.
(Bare metal) The port that is on a storage node and connects to the compute network.

configuration backup file

Backup file of storage cluster configuration information.

Configuration file

(Virtual machine) Generic term for VSS block configuration file and VM configuration file.

(Bare metal) A synonym for the VSS block configuration file.

Console interface

The interface of a storage node console (such as a virtual console via BMC).

control network

(Virtual machine) The network between the controller node and the storage node or maintenance node. It is used for Virtual Storage Software block management operation and communication with external service such as SNMP and NTP.

(Bare metal) The network between the controller node and the storage node. It is used for Virtual Storage Software block management operation and communication with external service such as SNMP and NTP.

control port

(Virtual machine) The virtual port that is on a storage node and connects to the control network.

(Bare metal) The port that is on a storage node and connects to the control network.

controller node

A management node used to instruct Virtual Storage Software block's management function (volume creation, etc.).

data migration

A functionality to migrate data from an external storage system into Virtual Storage Software block in volume units.

disk controller

Hardware required to use a drive.

drive

A physical device that stores user data and the OS. Common name for SSDs and HDDs.

drive data relocation

Function of balancing data capacity among storage nodes (to optimize capacity efficiency of each storage node) when capacity becomes unbalanced among storage nodes due to storage node addition or removal.

event log

A file that records the operation of the system. In Virtual Storage Software block, it refers to the log for the purpose of fault notification.

Failover

Switching the cluster master (secondary) to the cluster master (primary) in the event of failure of the cluster master (primary).

fault domain

A group of storage nodes sharing power system and network switch. A configuration for making it possible to continue the operation of storage even if the storage nodes in a group collectively become abnormal.

initiator

An endpoint on the compute node side when accessing a volume from a compute node.

internode network

Network between storage nodes. Used for communication of user data and management information between storage nodes.

internode port

(Virtual machine) The virtual port that is on a storage node and connects to the internode network.

(Bare metal) The port that is on a storage node and connects to the internode network.

license key

Key to activate the corresponding license in Virtual Storage Software block.

maintenance blockage

See *Storage node maintenance blocking* in the Glossary.

maintenance node

VM that is configured inside some of the storage nodes, and which is used to configure and manage Virtual Storage Software block.

maintenance recovery

See *Storage node maintenance recovery* in the Glossary.

multi-tenancy function

Function to allow resources of a storage in a large storage system to be distributed to and shared by multiple tenants (companies and divisions). A storage distributed to each tenant is called VPS (Virtual Private Storage).

normal volume

Volume that is neither P-VOL, S-VOL, nor P/S-VOL.

other volume capacity

Total capacity of snapshot volumes (S-VOLs and P/S-VOLs).

OVA

An acronym for the Open Virtualization Appliance/Application. The following files are bundled into one tar ball.

1. OVF file contains the virtual machine attributes etc.
2. Disk image or ISO image created by certain Hypervisor software.
3. Manifest file contains hash value for each file (mf option).
4. Certification file for digital signage for Manifest files (cert option).

OVF

Acronym for Open Virtualization Format. OVF is a standard format designed to allow different virtualization software to exchange virtual machine image files with each other.

P-VOL

Volume of the copy source.

P/S-VOL

Volume having both the P-VOL and S-VOL attributes in a snapshot tree in cascade configuration.

physical node

In an environment where storage is used, a physical server that belongs to that environment.

program product license

A license provided on a per-function basis.

protection domain

Setting for limiting the range of failure if an error occurs in a storage node or the network between storage nodes.

provisioned volume capacity

Total capacity of normal volumes and snapshot volumes (P-VOLs).

rebuild

Function of automatically restoring redundancy of data whose redundancy was reduced due to a drive failure or storage node failure.

Rebuild capacity

Capacity in a storage pool secured for Data rebuild at the time of drive failure.

Representative storage node

A storage node that is used to configure a storage cluster in the setup procedure for the bare metal model. This node is different from a cluster master node (primary).

S-VOL

The copy destination volume.

scale out

A method of increasing the number of CPUs, memory capacity, and the number of drives by adding storage nodes to improve system performance and capacity.

scope

The range of resources that users can operate. A scope is set for a user group. A scope for a user is determined according to the user group to which the user belongs.

snapshot volume

Volume that is either a P-VOL, S-VOL, or P/S-VOL.

spare node

Standby storage node used for the spare node function.

spare node function

Function to allow restoration of redundancy by performing spare node switchover. Spare node switchover from a faulty storage node to a storage node that is registered as a standby storage node in the storage cluster is performed when the faulty storage node cannot be restored by the auto-recovery function.

storage cluster

A virtual storage system built from multiple storage nodes.

storage controller

Part of Virtual Storage Software block processes that manage storage node capacities and volumes.

storage controller relocation

Function of optimizing the number of the storage controllers of each storage node when the number of the storage controllers becomes unbalanced among storage nodes due to storage node addition or removal.

storage node

Physical server to which the CPU, memory, and drives that comprise Virtual Storage Software block are assigned. Alternatively, this term refers to a process group of Virtual Storage Software block software running on storage nodes.

storage node addition

A process of adding a storage node to a storage cluster.

Storage node auto-recovery

Function to execute self-diagnosis and self-recovery by a storage node to recover the storage node from server failures due to software factors (firmware, driver, and so on) or due to temporary network problems between storage nodes.

Storage node maintenance blocking

Process of separating a storage node from a storage cluster temporarily and placing the storage node in a status that allows for part replacement or other maintenance.

Storage node maintenance recovery

Process of returning a storage node to the available status again after it was blocked by manual operation or due to a failure.

storage node removal

A process of removing a storage node from a storage cluster.

storage node replacement

A functionality or process that manually recovers a blocked storage node.

Replace the following to recover the blocked storage node.

(Virtual machine) Storage node VM

(Bare metal) Physical node

storage pool

Logical user data storage area that combines multiple drives.

storage software

The Virtual Storage Software block software that realizes a storage cluster.

system administrator

Administrator who manages the entire system.

target

An endpoint on the storage cluster side when accessing a volume from a compute node.

temporary volume capacity

Total capacity of volumes created temporarily by Data migration and Capacity balance.

thin provisioning

Method of creating a virtual storage in which the minimum required capacity is initially secured, and then expanded as required.

virtual machine (VM)

Virtual machine.

virtual private storage

Virtual storage logically divided from a storage cluster in a multi-tenancy configuration.

volume

A logical device that mounts on a compute node to read or write user data.

volume migration

Moving volumes (existing on a storage node to be removed) to another storage node.

volume path

Connection information between a compute node and a volume. One of the setting information necessary for using a volume from a compute node.

VPS

Acronym for Virtual Private Storage. See *virtual private storage* in the Glossary.

VPS administrator

Administrator who manages a virtual private storage (VPS) in a multi-tenancy configuration.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA



HitachiVantara.com/contact