Hitachi Virtual Storage Software

1.12

Block System Administrator Guide

This guide describes and provides instructions for performing system administration tasks for Hitachi Virtual Storage Software block (VSS block), including managing users and event logs and using SNMP and CHAP authentication.

MK-91VSSB001-03
February 2023

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.

Contents

Preface

This guide describes and provides instructions for performing system administration tasks for Hitachi Virtual Storage Software block (VSS block), including managing users and event logs and using SNMP and CHAP authentication.

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

- 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 how much time it takes for password changes to be reflected in the console interface.
- Revised the description of teaming information for management ports and interstorage node ports.
- Added descriptions about email notification target and notification time.
- Revised the default values fror some user authentication settings.
- Added information about multi-tenancy.
- Corrected the description about the whitelist setting when using the VMware vCenter Server Plugin.
- Added an explanation of session management.
- Added a description of collecting VMware ESXi logs.
- Revised the description of SSL communication.

Document conventions

This document uses the following typographic conventions:

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

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

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

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

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

Comments

Please send 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: Overview of Hitachi Virtual Storage Software block

About Hitachi Virtual Storage Software block

Virtual Storage Software block is a storage software product that builds and sets up a virtual storage system from multiple general-purpose servers.

The operation and management features of the Virtual Storage Software block system are as follows. The system offers a high-performance, high-capacity block storage service with high reliability.

- The initial cost is low because general-purpose servers (x86 servers) are used.
- You can quickly introduce the system, add storage, or reduce storage.
- You can centrally manage multiple storage nodes as a single storage system.
- You can easily determine the operation status of resources, the maximum volume capacity that can be created, and I/O performance at the time of checking.
- The REST APIs for Virtual Storage Software block are compatible with the REST APIs for Hitachi Storage Advisor Embedded. Administrators with experience in managing Hitachi storage products can manage storage using Virtual Storage Software block without extra time on learning Virtual Storage Software block.

User data protection methods

Virtual Storage Software block supports Hitachi Polyphase Erasure Coding (HPEC) and Mirroring as user data protection functions. HPEC is a proprietary data protection method developed by Hitachi for SDS systems and uses narrow internode network bandwidth for outstanding capacity efficiency. If HPEC is used, the user data is stored on a local drive. Mirroring is a data protection method that stores a copy of user data on another storage node.

For HPEC, the 4D+1P or 4D+2P method can be selected. For Mirroring, only Duplication can be selected. The method to be used is specified during setup.

HPEC 4D+1P (4 data areas + 1 parity area)

This method is suitable when capacity efficiency and performance are important.

- User data and its parities are stored on five or more different storage nodes for redundancy.

- At least five storage nodes are required.

- The maximum capacity available to users is 60 to 75% of the physical capacity.

  However, if the rebuild capacity policy (rebuildCapacityPolicy) is set to "Fixed" (default), users can use a maximum of 60 to 75% of the physical capacity excluding the rebuild capacity on each storage node. For details about the rebuild capacity, see *Rebuild capacity of a storage pool*.

- The read performance is the same as HPEC 4D+2P, but the write performance is better than HPEC 4D+2P.

- The number of storage node or drive failures allowed is one.

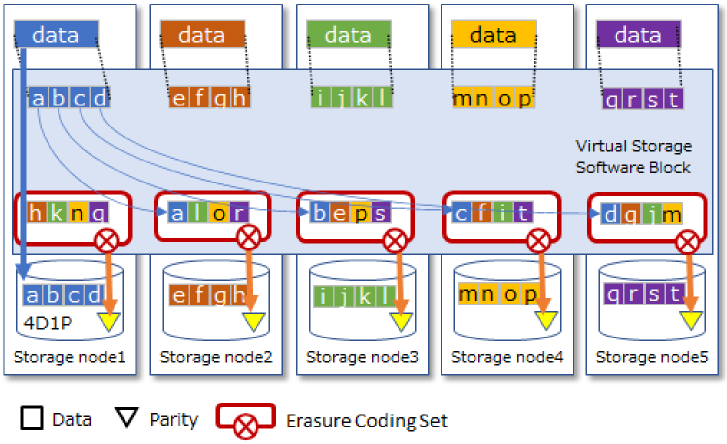

HPEC 4D+2P (4 data areas + 2 parity area)

This method is suitable when the number of failures that can be tolerated is important.

- User data and its parities are stored on six or more different storage nodes for redundancy.

- At least six storage nodes are required.

- The maximum capacity available to users is 50 to 65% of the physical capacity.

  However, if the rebuild capacity policy (rebuildCapacityPolicy) is set to "Fixed" (default), users can use a maximum of 50 to 65% of the physical capacity excluding the rebuild capacity on each storage node. For details about the rebuild capacity, see *Rebuild capacity of a storage pool*.

- The read performance is the same as HPEC 4D+1P, but the write performance is about 60% of HPEC 4D+1P.

- The number of storage node or drive failures allowed is two.

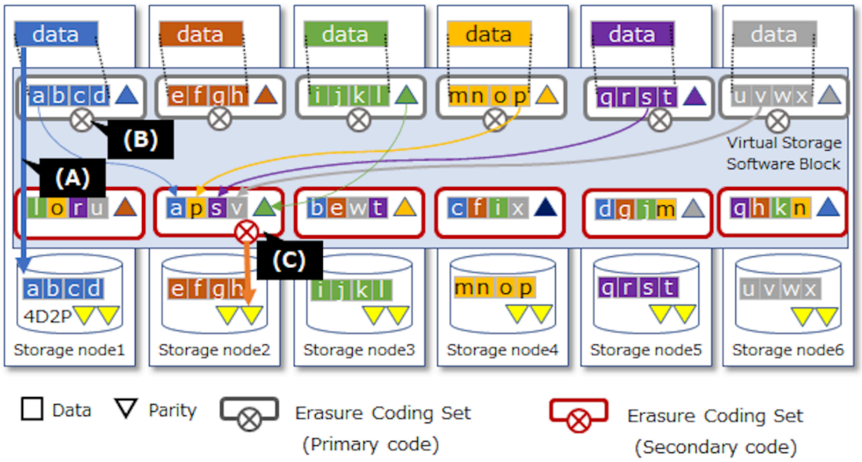

(A) Store data locally and reduce network communication during read

(B) Primary coding: Coding reduces data volume for two redundancies

(C) Secondary coding: Data storage capacity is reduced to achieve capacity efficiency that is equivalent to EC (Erasure Coding).

Mirroring Duplication (1 data area + 1 data copy)

This method is suitable if performance is priority.

- User data and its copies are stored redundantly on two different storage nodes.
- At least three storage nodes are required.
- The maximum capacity available to users is 40 to 48% of the physical capacity.

  However, if the rebuild capacity policy (rebuildCapacityPolicy) is set to "Fixed" (default), users can use a maximum of 40 to 48% of the physical capacity excluding the rebuild capacity on each storage node. For details about the rebuild capacity, see *Rebuild capacity of a storage pool*.

- The read performance of this method is equivalent to the HPEC 4D+1P and HPEC 4D+2P methods but the write performance is superior to the HPEC 4D+1P method. The fault tolerance against storage node or drive failures is also superior to the HPEC methods.
- The allowable number of defective storage nodes or drives is 1. However, two or more defective storage nodes or drives are tolerated except in the following cases:
  - Condition 1: Case where storage node or drive failures occur on both storage nodes that belong to redundant storage controllers

    For details about storage controllers, see *Capacity management of storage nodes by the storage controllers* in this manual.

    Caution:

    Failures might not be tolerated during the following periods, regardless of Condition 1:

    - After storage node addition until completion of drive data relocation
    - When a storage node is being removed
  - Condition 2: Case where failures occur on two or more cluster master nodes.

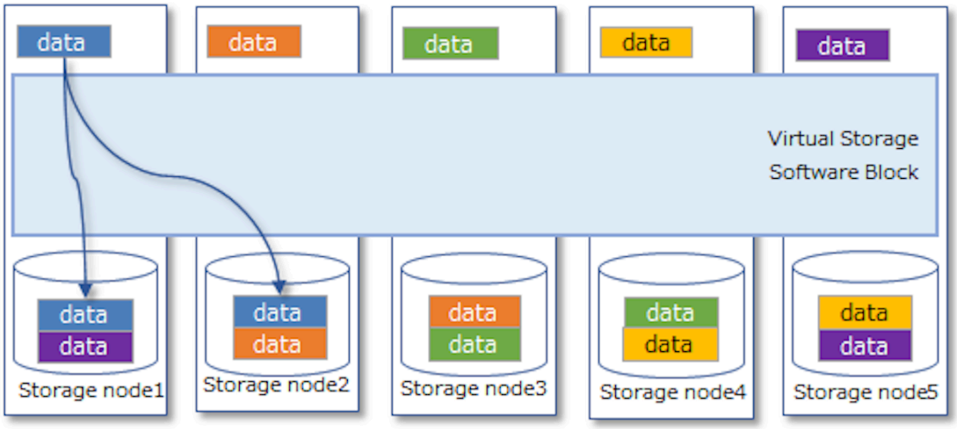

Note:

For how to design the capacity for the HPEC 4D+1P, HPEC 4D+2P, or Mirroring Duplication method, see *Capacity design (for HPEC 4D+1P)*, *Capacity design (for HPEC 4D+2P)*, or *Capacity design (for Mirroring)* in this manual.

Caution:

If the number of failures exceeds the allowable limit, you need to re-install Virtual Storage Software block or restore the configuration from the configuration backup file. In case more failures occur than are allowed, back up user data to other media and obtain the configuration backup file of Virtual Storage Software block. User data cannot be restored by re-installing Virtual Storage Software block or restoring the configuration from the configuration backup file.

System configuration

The system built by Virtual Storage Software block is a separate Software Defined Storage (SDS) that separates the storage node from the compute node.

Virtual Storage Software block's storage system consists of multiple storage nodes. A storage node consists of multiple drives.

Operations require the registration of compute nodes running by applications accessing the storage system and the installation of the controller node for management operations on the storage system.

Virtual machine

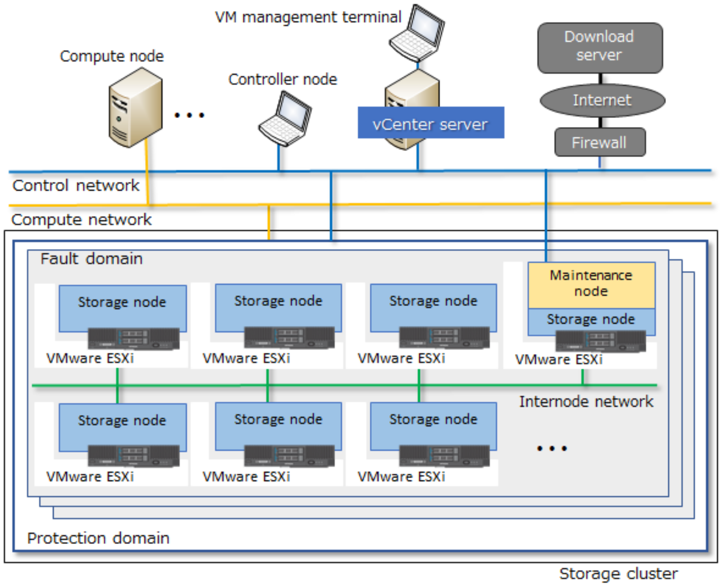

Bare metal

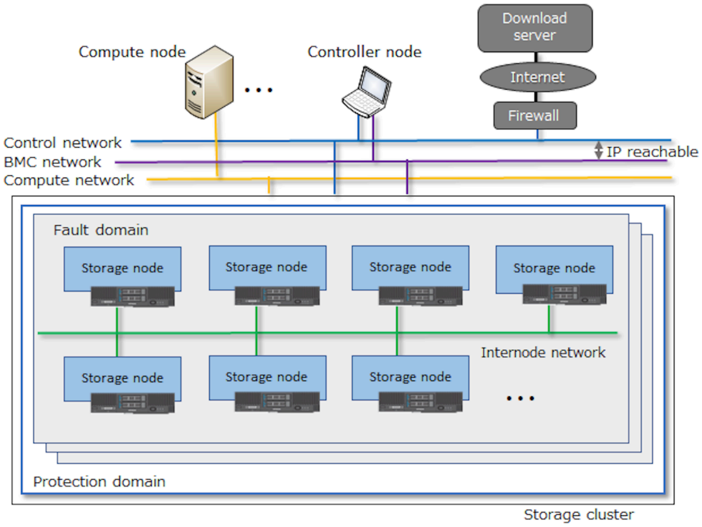

The Virtual Storage Software block storage system uses the following networks.

Control network

(Virtual machine) Network that connects the controller nodes to storage nodes and maintenance nodes. Use this network to send management commands to Virtual Storage Software block running on storage nodes and communicate with external services (such as SNMP managers, and NTP servers). The storage nodes are connected to this network through their control ports. This network is also used for communication with VMware vCenter Server. In the virtual machine model, you can configure a redundant configuration with multiple ports by using NIC teaming on VMware ESXi hosts on storage nodes.

(Bare metal) Network that connects the controller nodes to storage nodes. Use this network to send management commands to Virtual Storage Software block running on storage nodes and communicate with external services (such as SNMP managers and NTP servers). The storage nodes are connected to this network through their control ports. In the bare metal model, configure a redundant configuration with multiple ports by using teaming on storage nodes.

Compute network

Network that connects compute nodes to storage nodes. Use this network to send user data. The storage nodes are connected to this network through their compute ports.

(Virtual machine) This network uses the iSCSI or FC protocol.

(Bare metal) This network uses the iSCSI protocol.

Internode network

Network of storage nodes. Use this network to protect data by storing the same data in multiple storage nodes, exchange user data when capacity is balanced between storage nodes, and transmit management information between storage nodes. The storage nodes are connected to this network through their internode ports.

(Virtual machine) You can configure a redundant configuration with multiple ports by using NIC teaming on VMware ESXi hosts on storage nodes.

(Bare metal) Configure a redundant configuration with multiple ports by using teaming on storage nodes.

BMC network (Bare metal)

Network that connects the storage node BMC and the controller node. Use this network to operate the BMC from the controller node. The storage nodes are connected to this network through their BMC ports.

The BMC network must be able to communicate with the control network.

The storage nodes in the storage cluster are named depending on their assigned roles.

Cluster master node (primary)

This storage node supervises the entire storage cluster. The representative IP address of the storage cluster is assigned.

- The node has an interface for communicating with users and the management application.
- The node centrally controls all the storage nodes in the storage cluster. Users and the management application can control each storage node by simply issuing commands to the cluster master node (primary).

Cluster master node (secondary)

This storage node takes over the cluster master node (primary) if the primary node fails.

Cluster worker node

Any storage node that is controlled by the cluster master node (primary).

Protection domain and fault domain

Protection domain

A protection domain is a unit that configures the data assurance format and internal processing I/O resource utilization. There is one storage pool in the protection domain, and one or more fault domains in the protection domain.

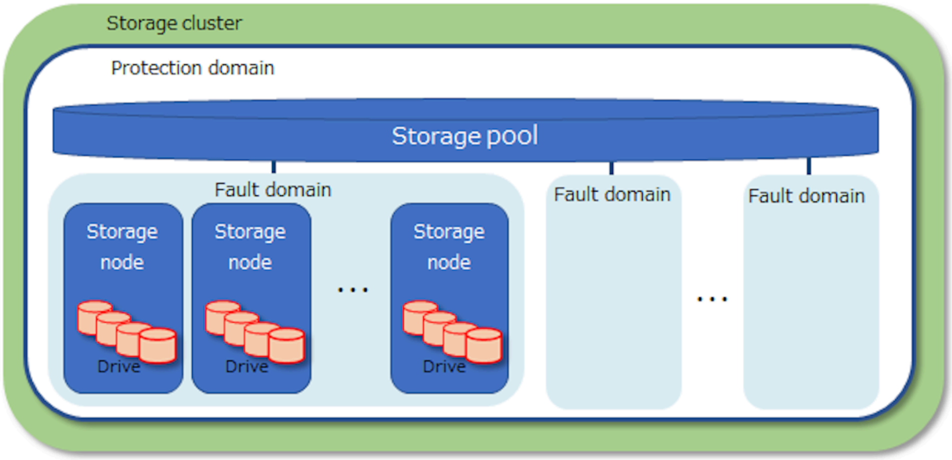

Fault domain

A fault domain is a group of storage nodes that share hardware, such as the same power supply system or network switches.

By setting up multiple fault domains and isolating hardware such as power supply system and network switches for each fault domain, the storage system can continue to operate if another fault domain is successful, even if there is a hardware failure within one fault domain.

Virtual Storage Software block automatically locates user data and storage controllers so that they are redundantly located across fault domains. However, cluster master nodes must be located manually at the time of mounting so that they are distributed evenly among fault domains.

However, even in a multiple fault domain configuration, if there is at least piece of hardware, such as a power supply system or network switch, that shares with other fault domains, the hardware failure can cause multiple fault domains to fail at the same time and cause operations to stop.

If hardware such as power supplies and network switches cannot be isolated for each fault domain, redundant configurations should be in place in case the non-isolated hardware fails.

The fault domain is set during The Virtual Storage Software block setup.

(Bare metal) For the bare metal model, isolate hardware for each fault domain in the BMC network in addition to other networks shown in the following figure.

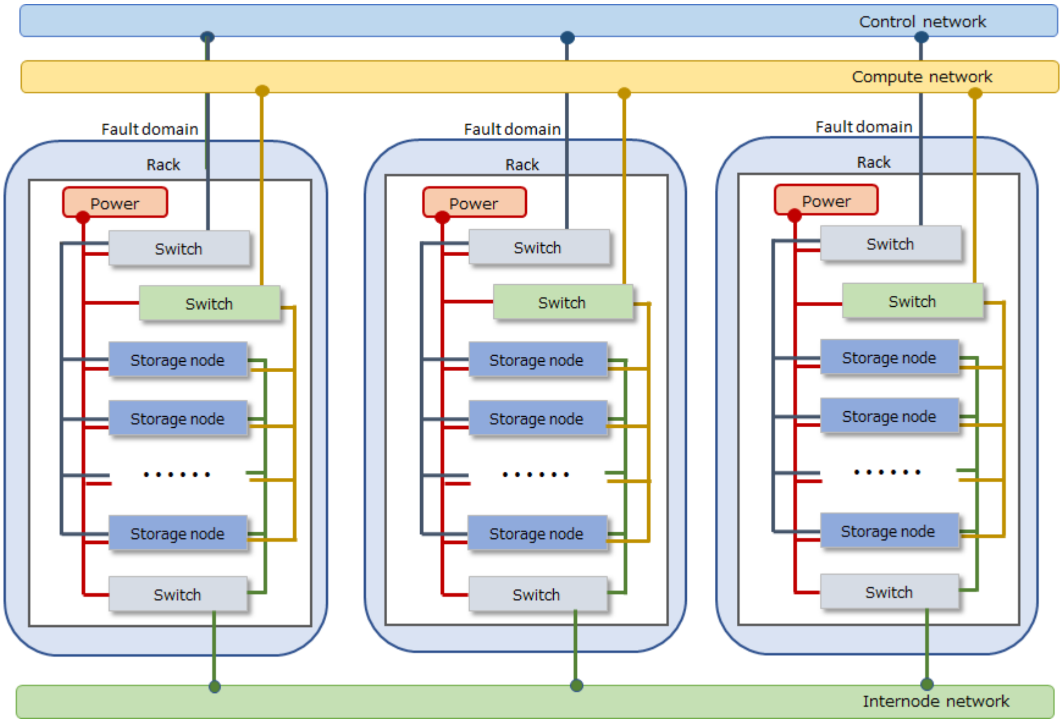

Maintenance node (Virtual machine)

A maintenance node is used to install and maintain Virtual Storage Software block (for example, when adding or replacing storage nodes, changing and setting the configuration file). Also, to manage a maintenance node itself (for example, when managing users or updating the maintenance node), you can use the commands and utilities provided for maintenance nodes.

Operation performed using a maintenance node	See:
Installing Virtual Storage Software block	Contact customer support.

Operation performed using a maintenance node	See:
Adding storage nodes	*Preparation and procedure for adding storage nodes*
Replacing storage nodes	*Replacing storage nodes (Virtual machine)*
Changing or setting configuration information for the storage cluster	*Changing or setting configuration information for the storage cluster*
Importing and exporting the configuration file	*Importing and exporting the configuration file*
Managing maintenance nodes	*Managing maintenance nodes (Virtual machine)*
Backing up and restoring the configuration	*Backing up the configuration information*

Configuration file

Virtual Storage Software block has the following configuration file:

- VSSB configuration file (SystemConfigurationFile.csv)

  This file defines the Virtual Storage Software block configuration itself.

- (Virtual machine) VM configuration file (vagrant_setup.yml)

  This file defines the configuration of the virtual machine built in the storage node of the Virtual Storage Software block.

These files define the configuration of the Virtual Storage Software block. These files are used when building storage clusters, adding storage nodes, replacing storage nodes, and changing and setting configuration information.

(Virtual machine) Also, when you change the configuration of your VMware environment, use the Import Configuration File feature to update the information in the Virtual Storage Software block. For more information, see *Importing and exporting the configuration file*.

Storage pools and volumes

A storage pool is a logical area for storing user data. Each pool is composed of multiple drives. You only need to manage the storage capacity of storage pools. You do not need to manage each drive or verify the physical boundaries of drives.

You can expand a storage pool by adding drives to an existing storage node or adding a storage node.

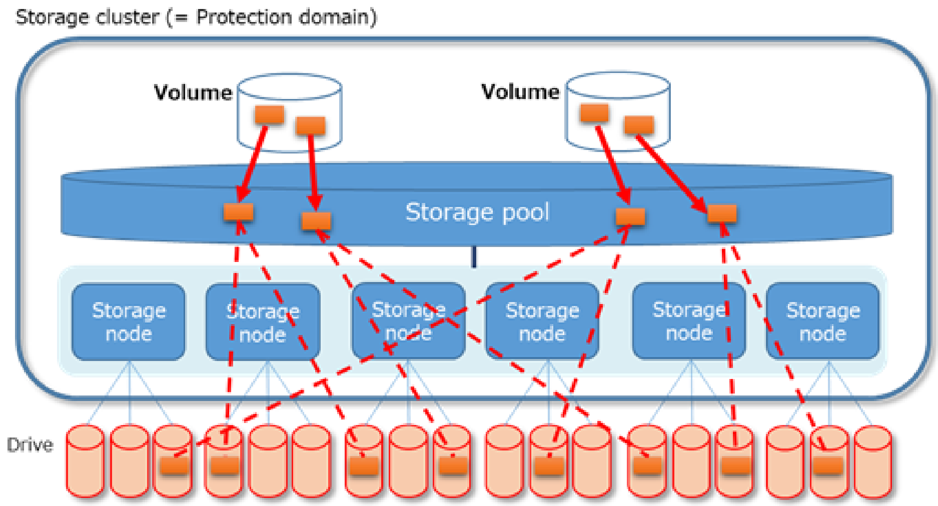

Volumes are logical devices to which user data is written or from which user data is read. Volumes are managed by the storage controller. When you write user data to a volume, it is written to the corresponding storage pool.

Capacity management of storage nodes by the storage controllers

Storage controllers are part of Virtual Storage Software block processes that manage storage node capacities and volumes.

The capacity of each storage node is managed by the storage controllers. The storage controllers also monitor the system configuration and usage.

The total storage pool capacity is the sum of the individual capacities managed by the storage controllers.

As many storage controllers as storage nodes are provided. A storage controller is deployed across multiple storage nodes and is redundant to tolerate storage node failures (the figure shows an example of redundancy degree 3).

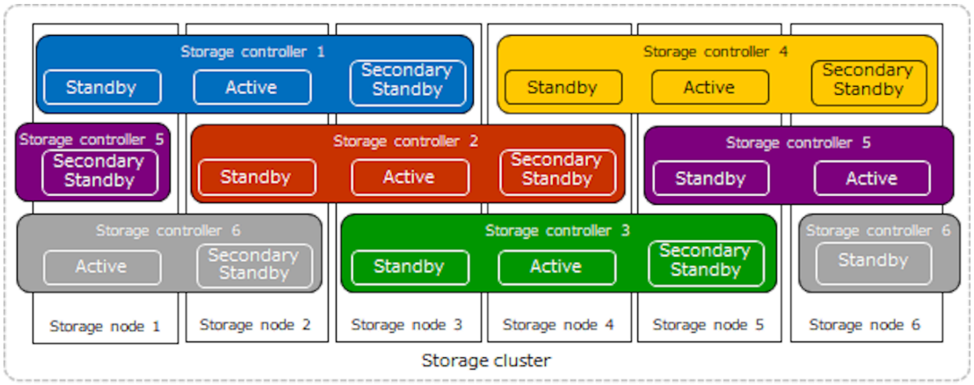

There are two redundancy degree settings for the storage controller: OneRedundantStorageNode (degree = 2) and TwoRedundantStorageNodes (degree = 3). One of these settings is automatically selected according to the setting of the user data protection method.

For details about management of storage controllers, see *Managing storage controllers* in the *Hitachi Virtual Storage Software Block Storage Administrator Guide*.

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.

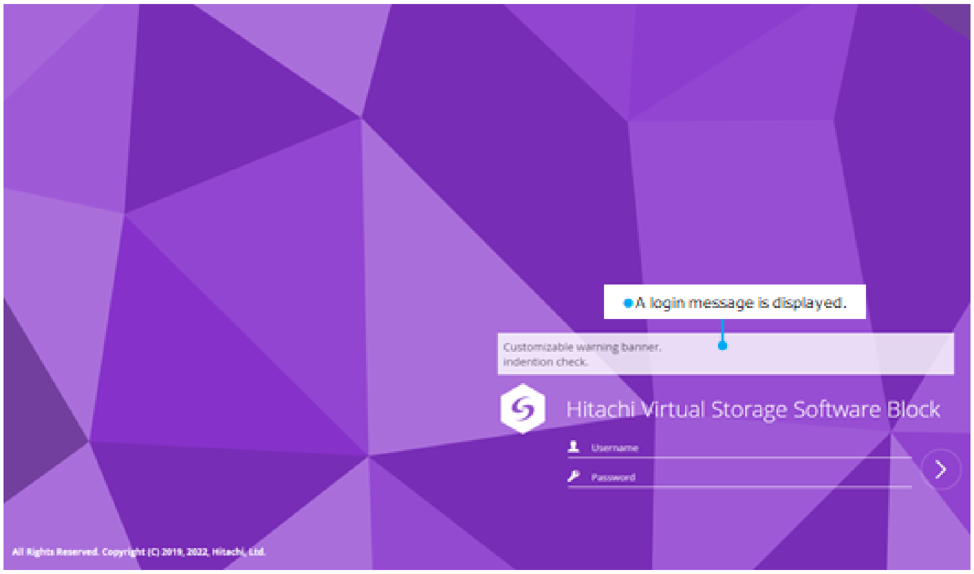

A login message is displayed.

Customizable warning banner.
indention check.

Hitachi Virtual Storage Software Block
Username
Password
All Rights Reserved. Copyright (C) 2019, 2022, Hitachi, Ltd.

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

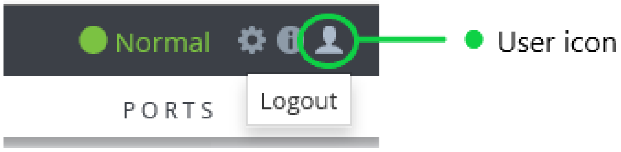

Dashboard window

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

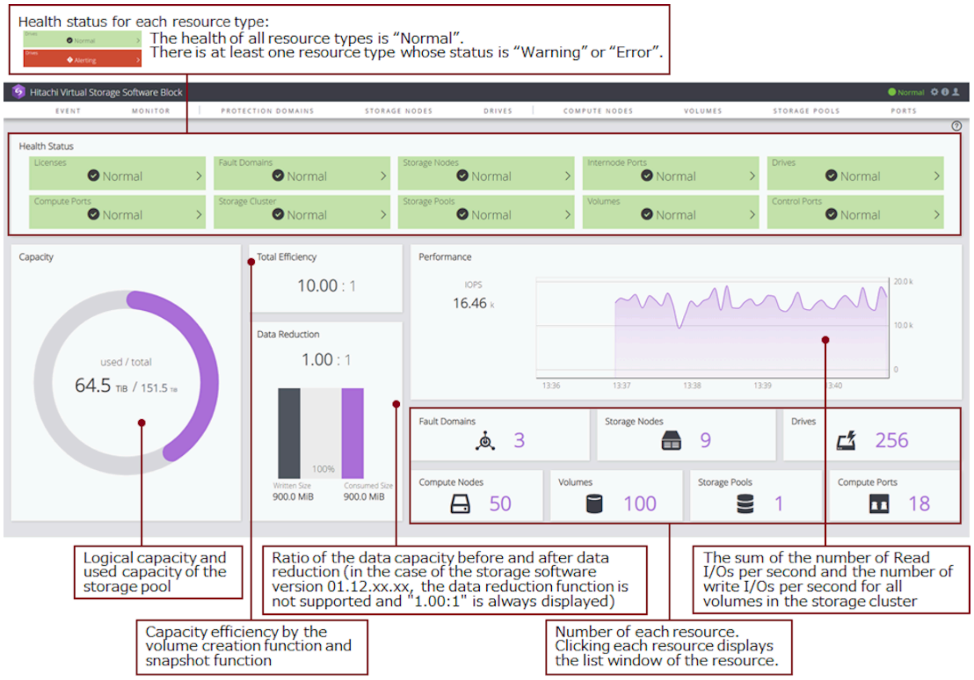

Health status for each resource type:
The health of all resource types is "Normal".
There is at least one resource type whose status is "Warning" or "Error".

Logical capacity and used capacity of the storage pool

Capacity efficiency by the volume creation function and snapshot function

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)

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

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

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.

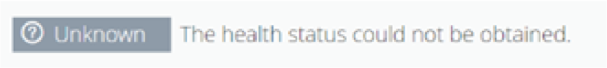

Unknown The health status could not be obtained.

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.

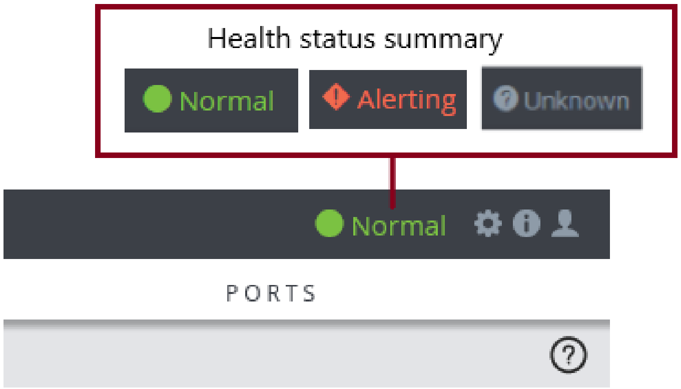

Clicking the health status summary opens the dashboard window.

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.

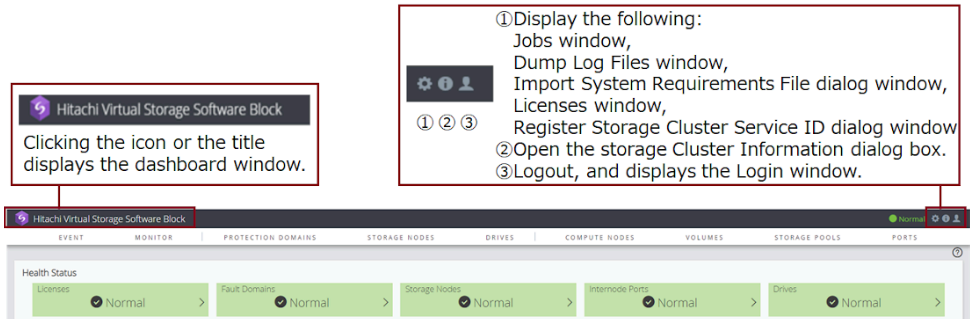

Clicking the icon or the title displays the dashboard window.

①Display the following:
Jobs window,
Dump Log Files window,
Import System Requirements File dialog window,
Licenses window,
Register Storage Cluster Service ID dialog window
②Open the storage Cluster Information dialog box.
③Logout, and displays the Login 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

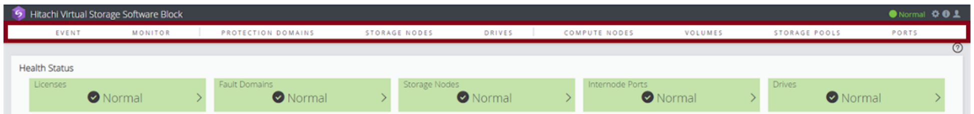

Clicking each item displays the following information.

- 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.

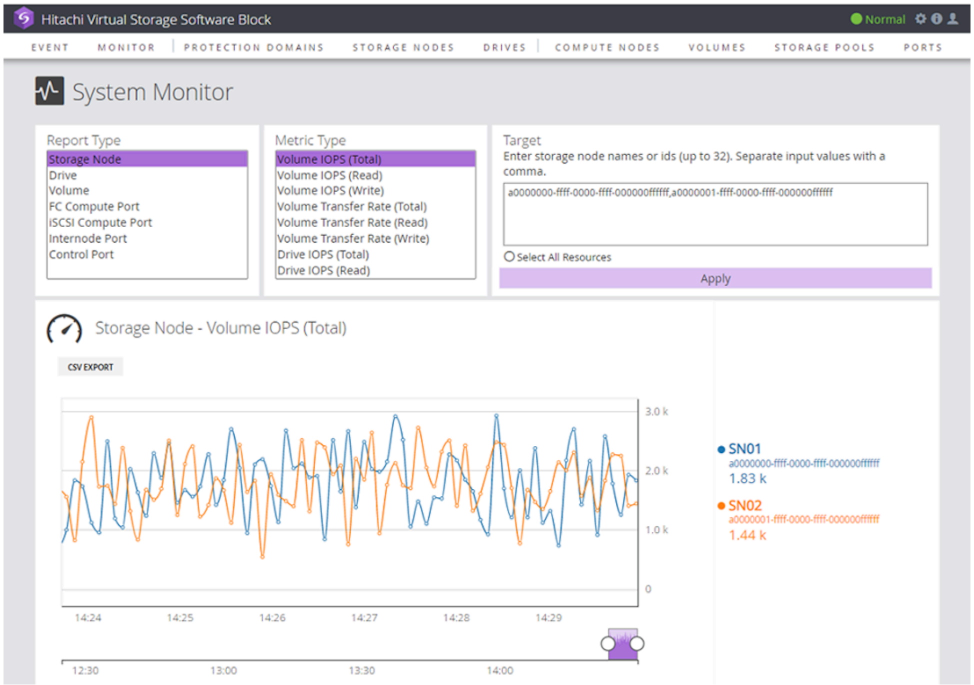

Report Type: Storage Node

Metric Type: Volume IOPS (Total)

Target

Enter storage node names or ids (up to 32). Separate input values with a comma.

a0000000-ffff-0000-ffff-000000ffffff,a0000001-ffff-0000-ffff-000000ffffff

○ Select All Resources

Storage Node - Volume IOPS (Total)

● SN01
a0000000-ffff-0000-ffff-000000ffffff
1.83 k

● SN02
a0000001-ffff-0000-ffff-000000ffffff
1.44 k

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.

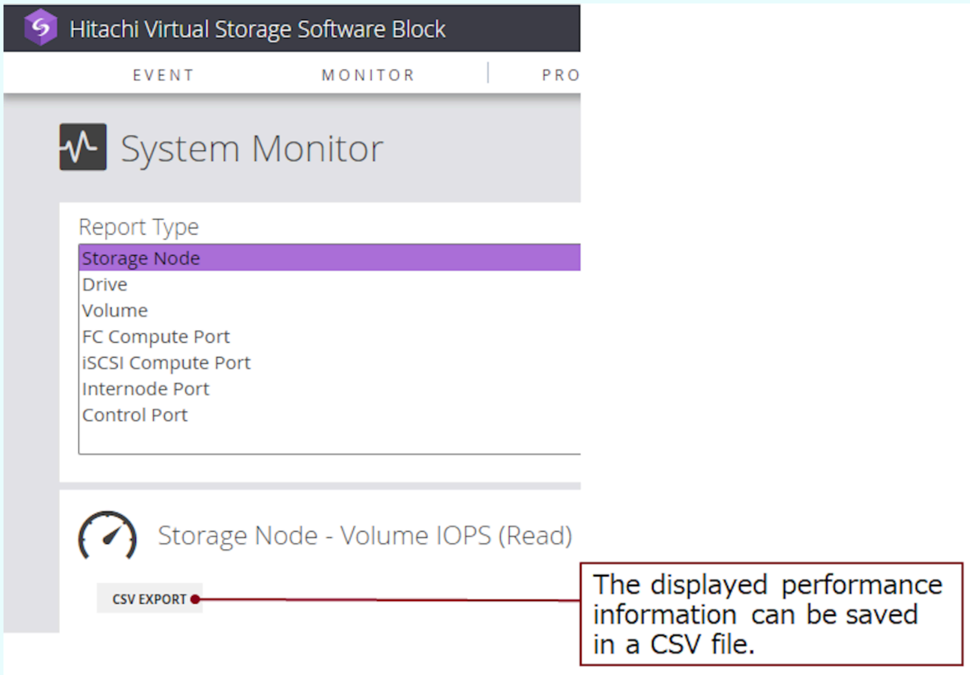

The displayed performance information can be saved in a CSV file.

- 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.

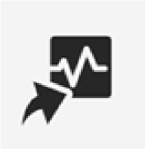

Clicking the activated monitor link icon selects the resource ID having this icon and displays the System Monitor window with the report type and resource ID set.

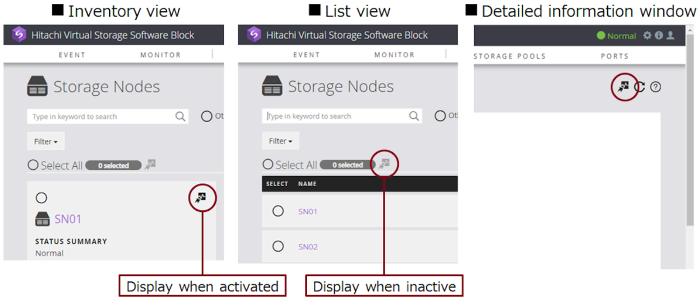

You can select a maximum of 32 resources. To select multiple resources, click the checkbox in each resource. The Monitor link icon to the right of the [Select All] will be activated.

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

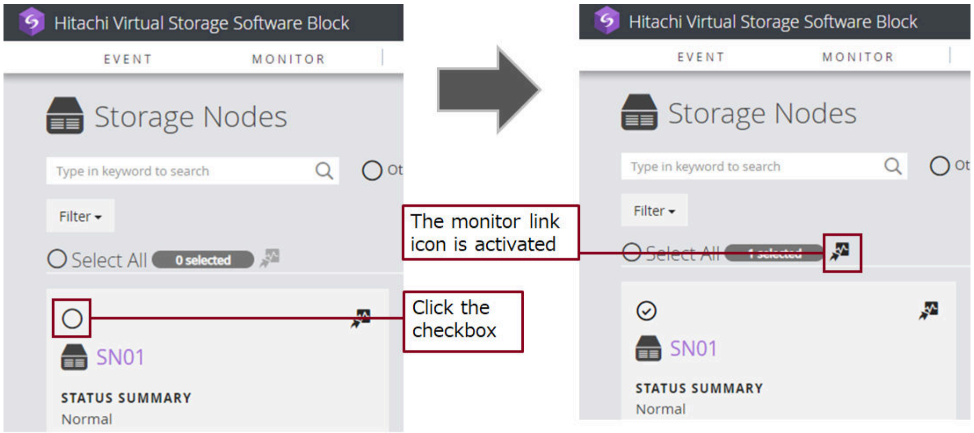

To select all resources, click the checkbox to the left of [Select All]. If 33 or more resources are selected, the Monitor link to the right of [Select All] is not activated because up to 32 resources can be selected. Clear the selection by clicking the checkbox in each resource and set the number of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

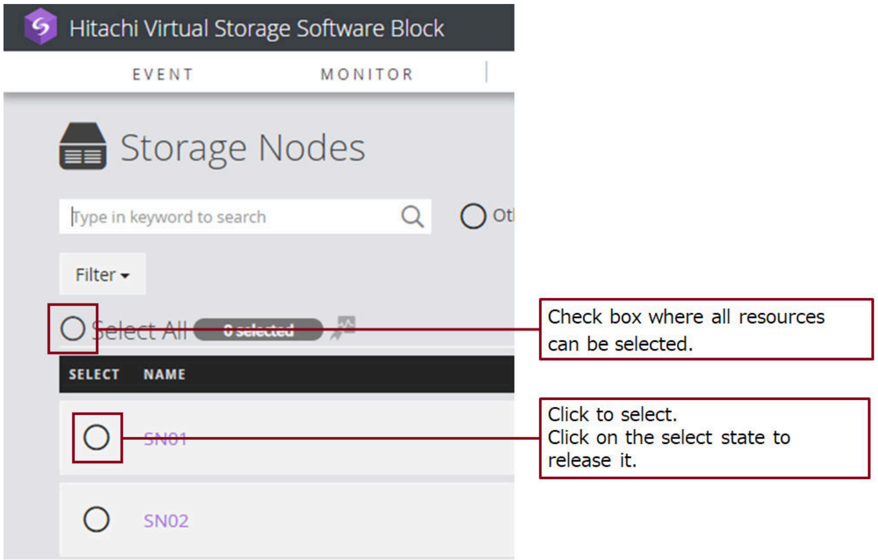

Check box where all resources can be selected.

Click to select.
Click on the select state to release it.

- 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.

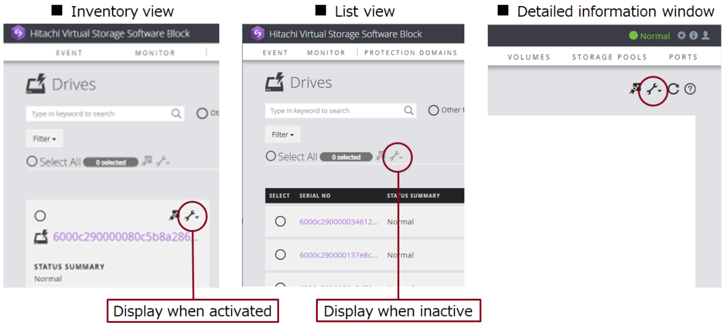

How many resources can be selected varies depending on the displayed menu. To select multiple resources, click the check box in each resource. Doing this activates the Maintenance icon to the right of the Monitor link icon.

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.

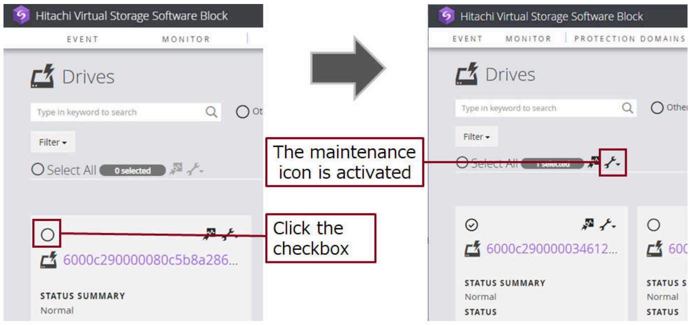

The maintenance icon is activated

Click the checkbox

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.

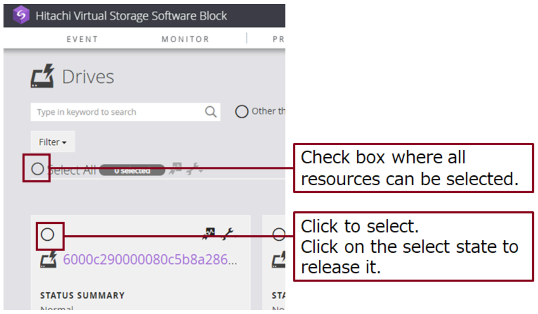
Hitachi Virtual Storage Software Block
EVENT
MONITOR
Drives
Type in keyword to search
Other th
Filter
Select All
0 selected
6000c290000080c5b8a286...
STATUS SUMMARY

Check box where all resources can be selected.

Click to select.
Click on the select state to release it.

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.

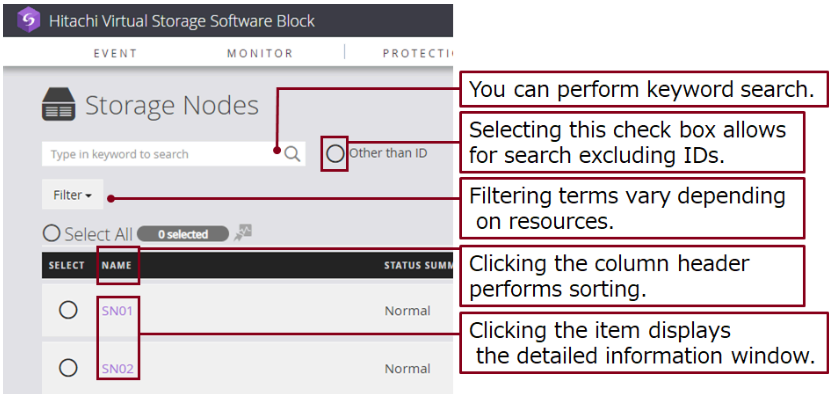

You can perform keyword search.

Selecting this check box allows for search excluding IDs.

Filtering terms vary depending on resources.

Clicking the column header performs sorting.

Clicking the item displays the detailed information window.

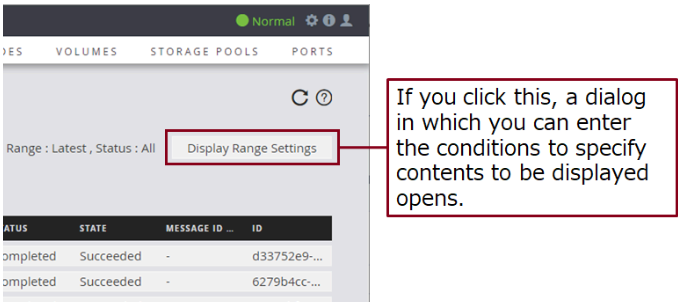

If you click this, a dialog in which you can enter the conditions to specify contents to be displayed opens.

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

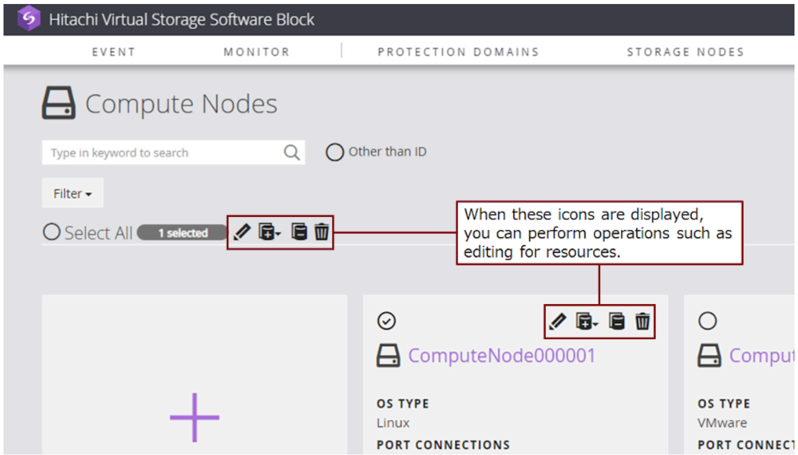

When these icons are displayed, you can perform operations such as editing for resources.

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.

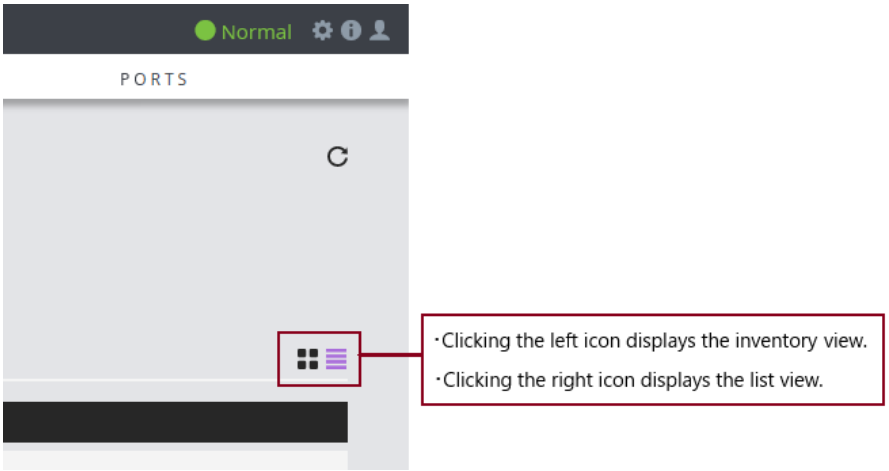

・Clicking the left icon displays the inventory view.

・Clicking the right icon displays the list view.

■List view

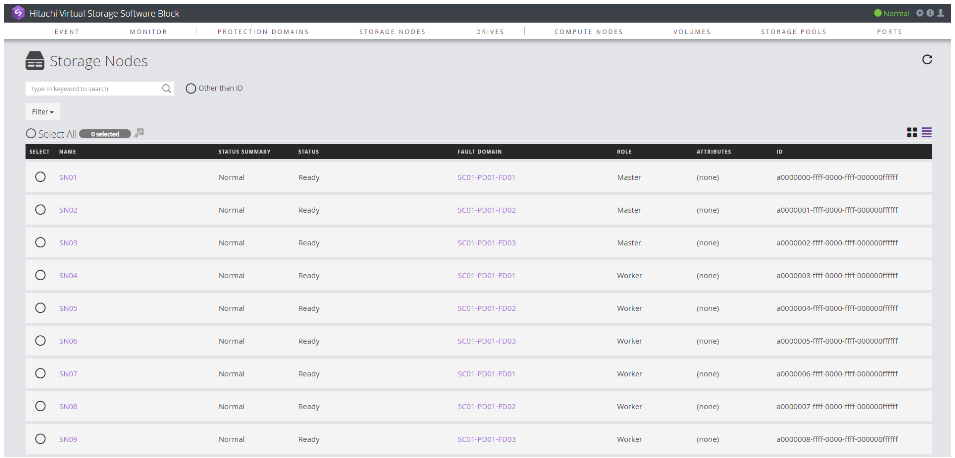

SELECT	NAME	STATUS SUMMARY	STATUS	FAULT DOMAIN	ROLE	ATTRIBUTES	ID
○	SN01	Normal	Ready	SC01-PD01-FD01	Master	(none)	a0000000-ffff-0000-ffff-000000ffffff
○	SN02	Normal	Ready	SC01-PD01-FD02	Master	(none)	a0000001-ffff-0000-ffff-000000ffffff
○	SN03	Normal	Ready	SC01-PD01-FD03	Master	(none)	a0000002-ffff-0000-ffff-000000ffffff
○	SN04	Normal	Ready	SC01-PD01-FD01	Worker	(none)	a0000003-ffff-0000-ffff-000000ffffff
○	SN05	Normal	Ready	SC01-PD01-FD02	Worker	(none)	a0000004-ffff-0000-ffff-000000ffffff
○	SN06	Normal	Ready	SC01-PD01-FD03	Worker	(none)	a0000005-ffff-0000-ffff-000000ffffff
○	SN07	Normal	Ready	SC01-PD01-FD01	Worker	(none)	a0000006-ffff-0000-ffff-000000ffffff
○	SN08	Normal	Ready	SC01-PD01-FD02	Worker	(none)	a0000007-ffff-0000-ffff-000000ffffff
○	SN09	Normal	Ready	SC01-PD01-FD03	Worker	(none)	a0000008-ffff-0000-ffff-000000ffffff

■Inventory view

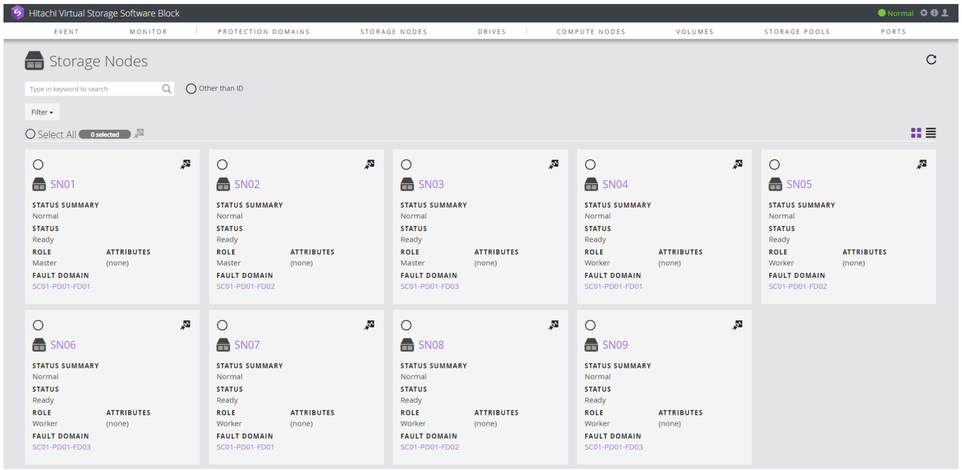

Storage Nodes

○ Other than ID

○ Select All

	STATUS SUMMARY	STATUS	ROLE	ATTRIBUTES	FAULT DOMAIN
SN01	Normal	Ready	Master	(none)	SC01-PD01-FD01
SN02	Normal	Ready	Master	(none)	SC01-PD01-FD02
SN03	Normal	Ready	Master	(none)	SC01-PD01-FD03
SN04	Normal	Ready	Worker	(none)	SC01-PD01-FD01
SN05	Normal	Ready	Worker	(none)	SC01-PD01-FD02
SN06	Normal	Ready	Worker	(none)	SC01-PD01-FD03
SN07	Normal	Ready	Worker	(none)	SC01-PD01-FD01
SN08	Normal	Ready	Worker	(none)	SC01-PD01-FD02
SN09	Normal	Ready	Worker	(none)	SC01-PD01-FD03

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.

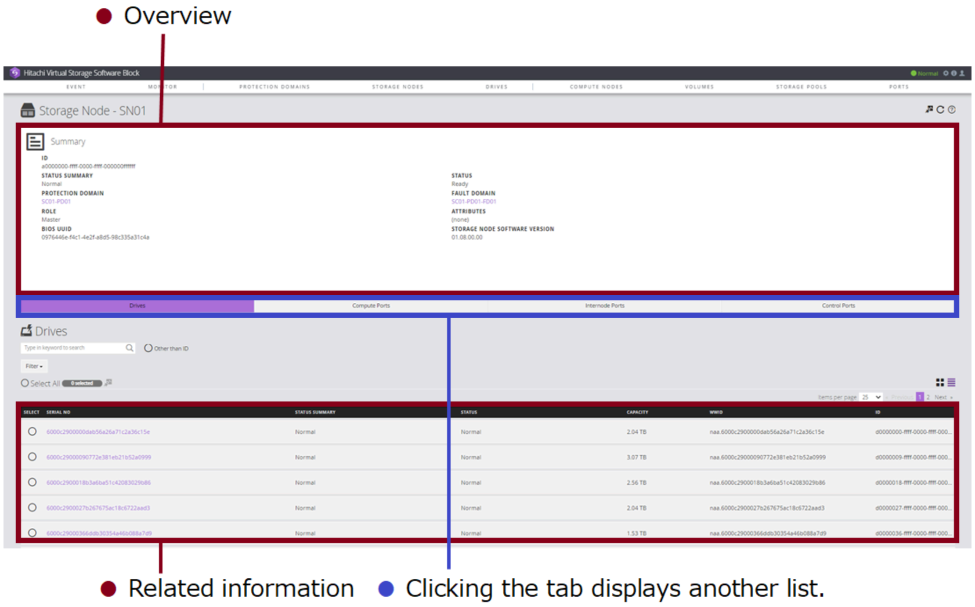

Storage Node - SN01

Summary

ID
a0000000-ffff-0000-ffff-000000ffffff

STATUS SUMMARY
Normal

PROTECTION DOMAIN
SC01-PD01

ROLE
Master

BIOS UUID
0976446e-f4c1-4e2f-a8d5-98c335a31c4a

STATUS
Ready

FAULT DOMAIN
SC01-PD01-FD01

ATTRIBUTES
(none)

STORAGE NODE SOFTWARE VERSION
01.08.00.00

Drives

○ Other than ID

○ Select All

SELECT	SERIAL NO	STATUS SUMMARY	STATUS	CAPACITY	WWID	ID
○	6000c2900000dab56a26a71c2a36c15e	Normal	Normal	2.04 TB	naa.6000c2900000dab56a26a71c2a36c15e	d0000000-ffff-0000-ffff-000...
○	6000c29000090772e381eb21b52a0999	Normal	Normal	3.07 TB	naa.6000c29000090772e381eb21b52a0999	d0000009-ffff-0000-ffff-000...
○	6000c2900018b3a6ba51c42083029b86	Normal	Normal	2.56 TB	naa.6000c2900018b3a6ba51c42083029b86	d0000018-ffff-0000-ffff-000...
○	6000c2900027b267675ac18c6722aad3	Normal	Normal	2.04 TB	naa.6000c2900027b267675ac18c6722aad3	d0000027-ffff-0000-ffff-000...
○	6000c29000366ddb30354a46b088a7d9	Normal	Normal	1.53 TB	naa.6000c29000366ddb30354a46b088a7d9	d0000036-ffff-0000-ffff-000...

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

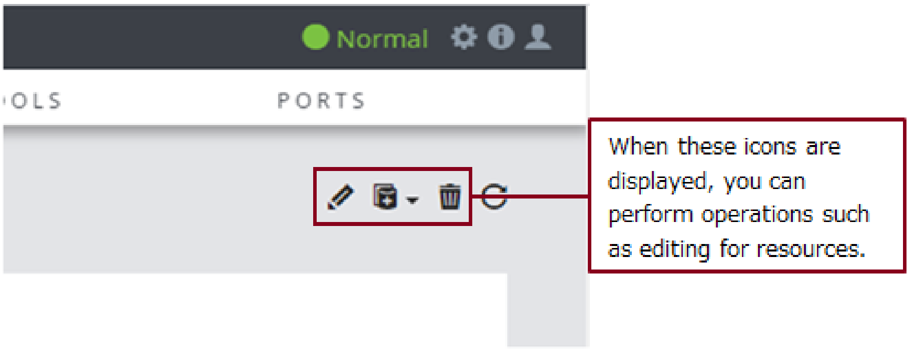

When these icons are displayed, you can perform operations such as editing for resources.

Refresh display

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

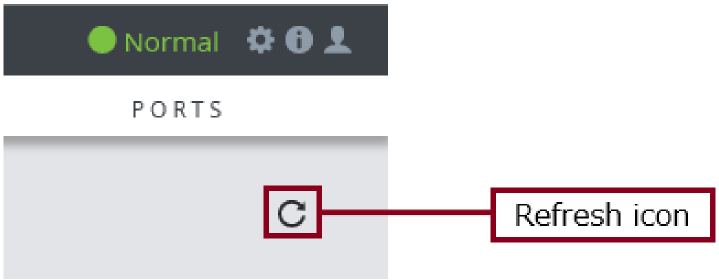

Help display

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

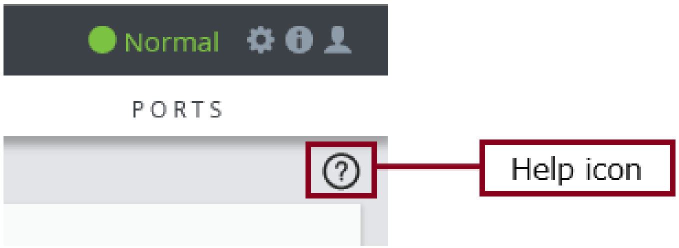

■Help display example

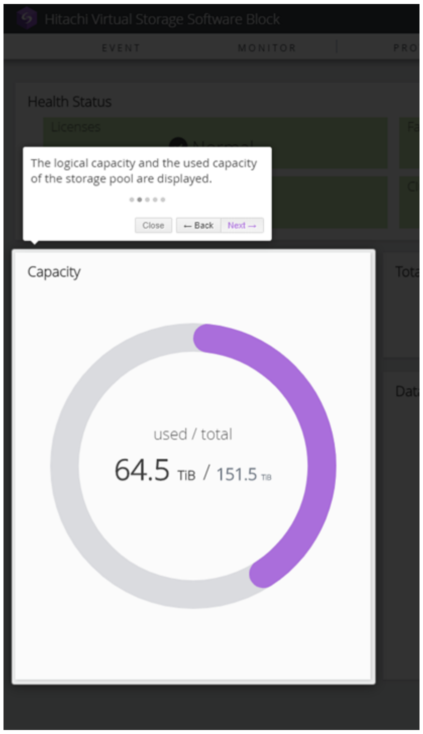

The logical capacity and the used capacity of the storage pool are displayed.

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

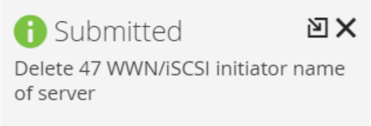
Submitted

Delete 47 WWN/iSCSI initiator name of server

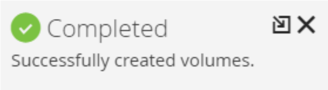
Completed

Successfully created volumes.

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

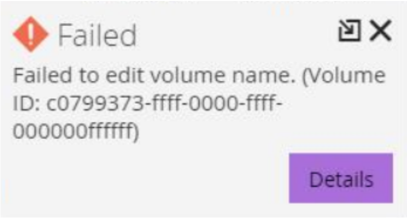
Failed

Failed to edit volume name. (Volume ID: c0799373-ffff-0000-ffff-000000ffffff)

Chapter 2: System configuration

Managing time settings

Synchronizing time

Configure NTP servers and the time zone for the storage cluster at setup or as described in *Changing or setting configuration information for the storage cluster*. The NTP servers synchronize the time within the storage cluster based on these settings.

Caution:

Synchronize the time on any storage cluster components* by using NTP or other time-synchronization methods. When you use multiple NTP servers, also synchronize the time on NTP servers. By synchronizing the time, you can avoid the risk of problems due to the time gap, and identifying the cause of a problem (such as a failure) becomes easier.

* The system components include VMware vCenter Server, VMware ESXi hosts, BMCs (storage nodes, (Bare metal) spare nodes), storage nodes, compute nodes, controller nodes, maintenance node, and network devices.

Changing the time of NTP servers

If you want to change the time of NTP servers, follow the procedure and make sure that you perform stopping the storage cluster beforehand.

1. Perform stopping the storage cluster.

2. Change the time of the NTP servers.

3. Configure the settings so that the time on other storage cluster components* is synchronized with the new set time.

4. Start the storage cluster.

* The system components include VMware vCenter Server, VMware ESXi hosts, BMCs, storage nodes, compute nodes, controller nodes, maintenance node, and network devices.

Using Windows Time Service as an NTP server

To use Windows Time Service on Windows Server as an NTP server for a storage system, you must specify the settings so that time can be synchronized with other NTP servers in the network.

Obtaining time settings of the storage cluster

The following information can be obtained.

Role-based execution is not subject to restriction for obtaining the time settings of the storage cluster.

- systemTime: UTC time of the storage cluster
- ntpServerNames: List of NTP servers (in order of precedence)
- timezone: Time zone of the storage cluster

Procedure

1. Obtain time settings.

   REST API: GET /v1/objects/storage-time-setting

   CLI: storage_time_setting_show

Setting a message to be displayed at login and during CLI Basic authentication (CLI or REST API)

(Virtual machine) You can set a message to be displayed on the GUI login window and in the warning banner during CLI basic authentication.

(Bare metal) You can set a message to be displayed in the GUI login window, at the console interface login, and in the warning banner during CLI Basic authentication.

To set such a message, you have to have the Security role. Role-based execution is not subject to restriction for obtaining the set message.

Before you begin

Required role: Security

Procedure

1. Set a message.

   Run either of the following commands with the message body specified. If you want to delete the existing message, specify an empty string.

   The maximum number of characters that can be specified is 6144. Specify a message in ASCII printable characters.

   REST API: PATCH /configuration/login-message

   CLI: login_message_set

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

> **Note:**
>
> Bare metal:
>
> - The screen size of the console interface is 80 x 24 character units. Set the message considering the screen size. If there are too many characters or line breaks, the message might be cut off.
> - The specified message is applied to the console interface by the internal processing that runs in one-minute cycle, and event log KARS20069-I is output. Therefore, it takes a certain amount of time until the message is applied.
> - No event log is output when you perform the following operation:
>   - You set a message to be displayed in the login window or in the warning banner during CLI basic authentication. However, there is no change in the content of the message.

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. Verify that the specified message was set correctly.

   REST API: GET /configuration/login-message

   CLI: login_message_show

   The set message will be displayed. If nothing is displayed, no message is set.

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.

> **Note:**
>
> If the CLI execution does not require authentication or it performs session authentication, nothing is displayed even if a message is set.

Setting a whitelist (CLI or REST API)

(Virtual machine) To prevent unauthorized management operations, you can set the IP address of a controller node or maintenance node in a whitelist. The maximum number of IP addresses that can be set is 10.

(Bare metal) To prevent unauthorized management operations, you can set the IP address of a controller node in a whitelist. The maximum number of IP addresses that can be set is 10.

Caution:

- Immediately after you set or change a whitelist, a REST API, CLI, or GUI operation cannot be performed temporarily. Wait for approximately 30 seconds before you perform an operation.
- To use a VMware vCenter plugin, set the IP address of VMware vCenter in the whitelist. If you have not set VMware vCenter in the whitelist, information cannot be referenced from VMware vCenter.
- (Bare metal) Operations on the console interface cannot be restricted by whitelist settings. If you want to restrict console interface operations, use the security settings provided by iLO for each storage node. For details, refer to the iLO User Guide provided by the vendor of the physical server to be used as the storage node.

Before you begin

Required role: Security

Procedure

1. Edit the whitelist setting.

   (Virtual machine) Run either of the following commands by enabling or disabling the whitelist and specifying the IP address (IPv4) of a controller node or maintenance node to be set in the whitelist. The set content is overwritten.

   (Bare metal) Run the command with the following specified: whether the whitelist is enabled and the IP address (IPv4) of the controller node to be set in the whitelist. The set content is overwritten.

   REST API: PATCH /v1/objects/web-server-access-setting

   CLI: web_server_access_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. Verify that the whitelist is set correctly.

   REST API: GET /v1/objects/web-server-access-setting

   CLI: web_server_access_setting_show

   After running the command, you receive a response indicating the set content.

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 3: Establishing an SSL/TLS session for operation and management

TLS

TLS (Transport Layer Security) is a protocol for safely transferring data over the Internet. Two peers (devices) with TLS enabled can establish a safe session using a private key and a public key. Both peers (devices) use a randomly created symmetric key to encrypt the data to be transferred.

The following terms are used in the explanation of TLS in this section.

Key pair

A combination of a private key and a public key. These two cryptographic keys are determined mathematically.

Server certificate

Also called digital certificate. A server certificate binds a server (Virtual Storage Software block storage system) to a key pair. With a server certificate, Virtual Storage Software block proves it is a server to a client. This way, Virtual Storage Software block and the client can communicate with each other using TLS. Two types of server certificates are available.

- Reliable certificate signed by a trusted Certificate Authority

  You can create a certificate issue request, send it to a trusted Certificate Authority, and receive a signed certificate. Notable Certificate Authorities include VeriSign, Symantec, and the Certificate Authority in your company. When you use a certificate signed by a Certificate Authority, reliability of the certificate improves despite the cost and the criteria you need to meet. The websites of Certificate Authorities explain the procedures for issuing certificates. The procedure for obtaining a certificate signed by a Certificate Authority described in this manual is just an example. For details, see the website of the Certificate Authority you want to apply for or contact the department in charge of the Certificate Authority in your company.

- Self-signed certificate

  You can create a certificate for yourself. In this case, the certified party is the same as the issuer of the certificate. This type of certificate might offer enough security if the controller node and the Virtual Storage Software block storage system communicate on an internal LAN protected by a firewall. However, best practice is to use this certificate only for testing encrypted connections.

Note:

The REST API or GUI becomes SSL/TLS communication by communicating over https.

Flow of preparation for establishing an SSL/TLS session

The preparation for establishing an SSL/TLS session is as follows.

Operations on the settings required for SSL/TLS communication should be performed on the controller node.

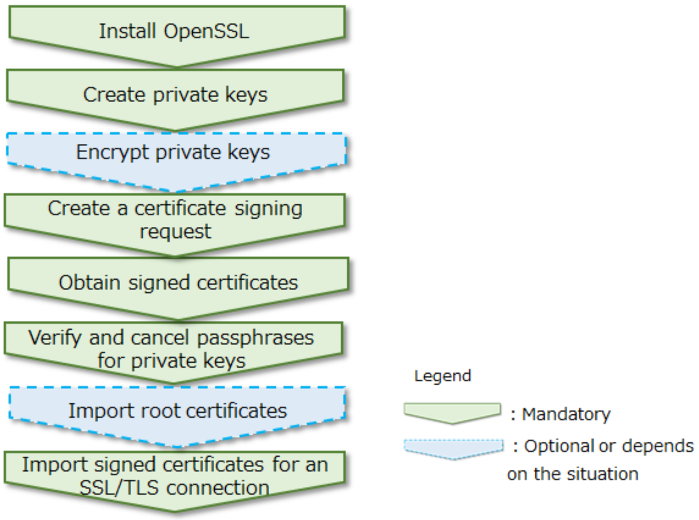
Install OpenSSL
Create private keys
Encrypt private keys
Create a certificate signing request
Obtain signed certificates
Verify and cancel passphrases for private keys
Import root certificates
Import signed certificates for an SSL/TLS connection
Legend
: Mandatory
: Optional or depends on the situation

Caution:

To enable an SSL/TLS session, make sure that your server certificate has not expired. If your server certificate has expired, a warning message appears when you try to make an SSL/TLS connection. Moreover, security of the communication is not guaranteed. If your server certificate has expired, repeat the steps in the above flowchart from “Create private keys” to “Import signed certificates for an SSL/TLS connection”.

Caution on updating server certificates

Note the following points when you update your server certificate.

In this section, "Subject Alternative Name" is abbreviated to "SAN".

- Immediately after Virtual Storage Software block is installed, a self-signed certificate is imported. When you update your server certificate for the first time, secure connection is not guaranteed.

  Also note that if you use a self-signed certificate, a warning appears when you try to establish an SSL/TLS session. Take action according to *Action to be taken when a warning message about a server certificate appears*.

- While a server certificate is being updated, the configuration of the storage system cannot be changed using the REST APIs or CLI. (Configuration change refers to management jobs including creation, update, and deletion. For example, creation of a volume is a configuration change.)

- While storage system configuration is being changed, the server certificate cannot be updated. Before you update the server certificate, make sure that management jobs are not being performed.

- Server certificates are updated asynchronously.

- Update of a server certificate has considerable impact on the system and should be approached with caution. Update might cause a failure in the Virtual Storage Software block storage system. To avoid any problem, thoroughly verify the content of server certificates and private keys.

- Be careful when you change the control network by changing or specifying the configuration settings if an IP address is set for the SAN in a server certificate. In this case, server certificate management becomes complex because you need to update the server certificate several times.

  When you create a server certificate, we recommend that you specify a FQDN for the SAN as shown below. For details, see *Creating a certificate signing request*.

  - Recommendation 1:

    Use a wildcard (*) to specify the subdomain for the FQDN that corresponds to the representative IP address of the cluster and the FQDN that corresponds to the IP address of each storage node control port.

    (Example: *.example.com)

  - Recommendation 2:

    - Specify the FQDN that corresponds to the representative IP address of the cluster.

      (Example: storage.example.com)

    - Specify the FQDN that corresponds to the IP address of each storage node control port.

      (Example: storage-node1.example.com, storage-node2.example.com, ...)

Installing OpenSSL

Use OpenSSL to create private keys and public keys that are necessary for establishing an SSL/TLS session.

In Linux, you do not need to install OpenSSL because it is pre-installed by default. When you enter "openssl version" from a terminal, if version information appears, OpenSSL is already installed.

In Windows, obtain the OpenSSL source code from the OpenSSL official website (http://www.openssl.org/) and compile the source code. Alternatively, use binary code provided by a third party.

Creating private keys

Use an OpenSSL command to create private keys.

You can use the same OpenSSL command for both Windows and Linux. Use a console where you can run commands. In Windows, open a command prompt and run the command. In Linux, open a terminal and run the command.

Procedure

1. Run the following command.

```
openssl genrsa -out server.key 2048
```

As a private key, the server.key file is created in the folder where you run the command.

Option	Description
-out<*private-key-file-name*>	A private key file is created with the specified name. The file name can be any.
2048	Key length. Specify 2048 or longer to increase security.

Encrypting private keys

If your private key is exposed to third parties, security of servers might be compromised. To prevent such a problem, you can encrypt private keys to protect them from third parties. Use an OpenSSL command to encrypt private keys.

You can use the same OpenSSL command for both Windows and Linux. Use a console where you can run commands. In Windows, open a command prompt and run the command. In Linux, open a terminal and run the command.

Procedure

1. Run the following command.

```
openssl genrsa -aes256 -out server.key 2048
```

   When you run the command, you are prompted to enter a passphrase. Passphrases are passwords for encrypting keys. Do not forget your passphrase. In addition, do not let others know your passphrase.

Option	Description
-aes256	Encryption algorithm. Other types of algorithms are available in addition to AES-256. To find out available encryption algorithms, run the following command. openssl genrsa -h
-out<*private-key-file-name*>	A private key file is created with the specified name.
2048	Key length.

Creating a certificate signing request

Use an OpenSSL command to create a certificate signing request.

You can use the same OpenSSL command for both Windows and Linux. Use a console where you can run commands. In Windows, open a command prompt and run the command. In Linux, open a terminal and run the command.

In this section, "Subject Alternative Name" is abbreviated to "SAN".

Procedure

1. Copy the OpenSSL settings file to the working folder.

   > **Note:**
   >
   > The location of the settings file differs depending on the environment.
   >
   > Linux: In many cases, the file is stored in /etc/pki/tls/openssl.cnf.
   >
   > Windows: The location of the file differs depending on the settings specified when OpenSSL is installed.

2. Open the file created in step 1 and add or modify the information according to the following.

Section name	Parameter name	Value	Description
req	req_extensions	v3_req	This parameter adds the v3_req section to the certificate signing request settings.
v3_req	subjectAltName	@alt_names	The system reads the "alt_names" section as the value of the SAN.
alt_names	DNS.<number>	<*FQDN-corresponding-to-the-IP address-of-a-storage-node-control-port-and-the-representative-IP address-of-the-storage-cluster*>	The system issues a certificate to the FQDN that corresponds to the IP address of a storage node control port and the representative IP address of the storage cluster. For <*number*> in the parameter name, specify a unique value beginning with 1 in ascending order. If the "alt_names" section does not exist, create it in the settings file.

(Example of setting): Items other than the above are omitted.

```
[ req ]
req_extensions = v3_req
[ v3_req ]
subjectAltName = @alt_names
[ alt_names ]
DNS.1 = storage.example.com
DNS.2 = storage-node2.example.com
DNS.3 = storage-node2.example.com
```

To ensure safe SSL/TLS communications with all storage nodes, we recommend that you use a certificate in which both of the following two items are specified for the SAN, or the subdomain of the following two FQDNs is specified by using a wildcard (*) for the SAN rather than the common name (CN). (Example: *.example.com)

- FQDN corresponding to the representative IP address (when using the representative IP address)

- FQDN corresponding to the IP address of a storage node control port

From version 58 onward, Google Chrome does not support the CN. Therefore, we recommend that you specify the FQDN for the SAN if you use the GUI.

For the SAN, we recommend that you register an FQDN because, when you change the IP address of a control port, you only have to change the relevant record in the DNS. If you use the FQDN, you also need to set up a DNS server.

If you use a wildcard character for the subdomain in each FQDN specified in a server certificate, you do not need to re-create the server certificate when adding or removing storage nodes.

3. Run the following command.

```
openssl req -sha256 -new -key server.key -out server.csr -config<file-created-in-step-2>
```

Option	Description
req	Requests creation of a certificate signing request (csr).
-sha256	Signature hash algorithm. Use an algorithm that is equivalent to SHA-2.
-key<*private-key-file-name*>	Specifies the name of the private key for creating a certificate signing request.
-out<*public-key-file-name*>	The command outputs a certificate signing request file with the specified name. You can assign any name to the certificate signing request file. Usually, use ".csr" as the extension.
-config<*file-created-in-step-2*>	Specifies the name of the settings file that was created in step 2. The information to be registered as the SAN in the certificate is specified in this file.

⚠ **Caution:**

Use SHA-256 as the hash algorithm. Do not use MD5 or SHA-1 because they might create security problems.

4. Enter the information, which will be written on the server certificate.

- Country Name (2 letter code) [AU]: Use two characters to enter the country name (example: JP).
- State or Province Name (full name) [Some-State]: Specify the name of your prefecture (example: Kanagawa).
- Locality Name (eg, city) []: Specify the name of your city, ward, town, village, or region (example: Odawara).

- Organization Name (eg, company) [Internet Widgits Pty Ltd]: Specify the name of your organization (example: Hitachi).

- Organization Unit Name (eg, section) []: Specify the name of your department in the organization (example: ITPD).

- Common Name (eg, YOUR name) []: Enter any value (not entered in the example).

- Email Address []: Enter your email address (not entered in the example).

- A challenge password []: Entry is not necessary.

- An optional company name []: Entry is not necessary.

Example

Input example

```
$ openssl req -sha256 -new -key server.key -out server.csr
Enter pass phrase for server.key:
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]:JP
State or Province Name (full name) [Some-State]:Kanagawa
Locality Name (eg, city) []:Odawara
Organization Name (eg, company) [Internet Widgits Pty Ltd]:Hitachi
Organizational Unit Name (eg, section) []:ITPD
Common Name (eg, server FQDN or YOUR name) []:
Email Address []:
Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []:
An optional company name []:
```

Obtaining signed certificates

After you create a private key and a public key, obtain a signed certificate file for the public key. Three methods are available to obtain signed certificate files.

- Obtain certificates from the Certificate Authority in your company. (Recommended)
- Obtain official certificates from a public Certificate Authority such as VeriSign. (Recommended)
- Create self-signed certificates.

Obtaining certificates from the Certificate Authority in your company

Contact the relevant department in your company for how to obtain server certificates. For the server certificates signed by the Certificate Authority in your company (in-house server certificates), root certificates exist to prove the validity of the server certificates. To enable a secure connection with the Virtual Storage Software block storage system, a root certificate must be imported to the client machine that communicates with Virtual Storage Software block. Obtain an in-house server certificate and a root certificate. For details about how to import root certificates, see *Importing root certificates*.

Obtaining official certificates from a public Certificate Authority

To obtain a signed, reliable certificate, send a certificate issue request file (csr file) to a public Certificate Authority such as VeriSign and obtain a signed certificate for a public key (crt file). For details about the procedure for requesting a certificate from a public Certificate Authority, go to the website of the Certificate Authority you want to apply for.

Obtaining this more reliable certificate requires additional cost and more criteria to be met.

When you request a public Certificate Authority to issue a certificate, specify the host name of Virtual Storage Software block in the Common Name field.

Obtaining self-signed certificates

You can create a signed certificate for a public key by signing the certificate yourself without asking any Certificate Authority to sign it for you. Use of self-signed certificates is preferable only for testing encrypted connections.

To create a self-signed certificate, run the following command.

Example:

```
openssl x509 -req -sha256 -days 10000 -in server.csr -signkey server.key -out
server.crt
```

In this example, the valid period is set to 10,000 days. In addition, when you run the command, the SHA-256 hash algorithm is used.

A server.crt file is created in the folder where you run the command. This server.crt file is the self-signed certificate for the public key.

⚠ **Caution:**

Make sure that you use SHA-256 for the hash algorithm. Do not use MD5 or SHA-1 to avoid security problems.

If you use a self-signed certificate, a warning message appears when you try to make an SSL/TLS connection. To make SSL/TLS connections using this certificate, add the option for ignoring warnings.

Note:

(Virtual machine) When importing a self-signed certificate with X.509 v3 extended attributes added as a trusted root certificate to a browser, CLI, or maintenance node, you need to set the following attribute values to make the certificate valid for an authority.

(Bare metal) When importing a self-signed certificate with X.509 v3 extended attributes added as a trusted root certificate to a browser or CLI, you need to set the following attribute values to make the certificate valid for an authority.

- X509v3 Basic Constraints must be CA:TRUE.
- The following values must be set (when setting X509v3 Key Usage).
  - Digital Signature
  - Certificate Sign

Verifying and canceling passphrases for private keys

If a passphrase is set for a private key, the private key cannot be used for the Virtual Storage Software block storage system. If a passphrase is set, cancel it.

You can use the same OpenSSL command for both Windows and Linux. Use a console where you can run commands. In Windows, open a command prompt and run the command. In Linux, open a terminal and run the command.

Before you begin

- The intended private key must be created beforehand.

Procedure

1. Move to the folder containing the private key file.
2. Run the following command.

If a passphrase is set, you are prompted to enter the passphrase when you run the command. In that case, go to step 3. If no character string is displayed, no passphrase is set. You do not need to perform step 3. The private key can be used for Virtual Storage Software block.

```
openssl rsa -in <name-of-private-key-file-for-input> -out<name-of-private-key-
file-for-output>
```

Note:

If you use the same private key file name for input and output, the private key file is overwritten. If you do not want the file to be overwritten, specify different file names for input and output or make a backup file, and then run the command.

3. If you are prompted to enter a passphrase, enter it.

   When you enter the correct passphrase, it is canceled. You can now use the private key for Virtual Storage Software block.

   Example:

```
openssl rsa -in server.key -out server.key Enter pass phrase for server.key:
```

Importing root certificates

(Virtual machine) In the browser, CLI, and maintenance node, root certificates that prove the validity of the certificates issued by the default trusted certification authority have been imported. In contrast, if a root certificate is issued by the Certificate Authority in your organization, you need to import the root certificate manually. The handling of the root certificates is different among the REST API, CLI, GUI (browser) and maintenance node.

(Bare metal) In the browser, and CLI, root certificates that prove the validity of the certificates issued by the default trusted certification authority have been imported. In contrast, if a root certificate is issued by the Certificate Authority in your organization, you need to import the root certificate manually. The handling of the root certificates is different among the REST API, CLI, and GUI (browser).

The following procedure describes the procedure for importing root certificates into the controller node.

Importing root certificates for REST APIs

When you use REST APIs as the management interface, how root certificates are used differs depending on the program that issues REST API requests. Import root certificates to satisfy the requirements of the program you are using.

Importing root certificates for CLI

CLI uses a Python library named certifi to manage root certificates. certifi has a list of root certificates. You can add new root certificate information to the list. Perform the following procedure.

Procedure

1. Start the Python 3 interpreter.

   ```
   $ python3
   ```

   Note:

   Depending on how Python 3 is installed, you need to type only "python" or type the detailed version of Python (python3.4).

2. Enter as follows to show the path for the certificate list file (cacert.pem) in certifi.

   ```
   >>>import certifi
   >>>certifi.where()
   ```

   Example: When running the command on Linux

   ```
   $ python3
   >>>import certifi
   >>>certifi.where()
   /usr/lib/python3.4/site-package/certifi/cacert.pem
   ```

3. Use a text editor to open the intended root certificate and copy the content.

   Copy the information in the following range.

   ```
   -----BEGIN CERTIFICATE----- //From here
            :
            :
            :
   -----END CERTIFICATE----- //Until here
   ```

4. Add the copied information to the certificate list file (cacert.pem) in certifi.

   Use a text editor to open the file at the location indicated in step 2. Add the information copied in step 3 at the end of the file.

Example

Caution:

- If you update the certifi library, you need to recopy the root certificate.
- To remove the installed root certificate, delete the contents added to the certificate list file (cacert.pem) in certifi in step 4.

Importing root certificates to a web browser

The import method differs depending on the web browser. Go to the official website of your web browser, confirm the import method, and import root certificates.

> **Caution:**
>
> If you import a root certificate that is not issued by a trusted Certificate Authority to a web browser, a serious security problem might occur. Carefully read the guidelines of your organization and use caution when you import such root certificates.

Importing signed certificates for an SSL/TLS connection (CLI or REST API)

To enable an SSL/TLS connection to Virtual Storage Software block using a certificate, import the applicable private key and a signed certificate for the applicable public key to Virtual Storage Software block and update the certificate.

Before you begin

- Required role: Security
- The applicable private key must be created beforehand.
- A signed certificate for the applicable public key must be obtained beforehand.
- The private key format must be PEM or DER.
- The format of the signed certificate for the public key must be X509.
- The passphrase for the private key must be canceled.
- Server certificates that can be imported are paired with private keys.
- If you import a server certificate in RSA format, a key length in the range from 1024 to 8192 bits is supported. The recommended key length is 2048 bits or longer. If you import a server certificate in ECC format, you can use one of the following Elliptic Curves: prime256v1, secp384r1, and secp521r1.

Procedure

1. Import the server certificate.

   You can perform this for the cluster master node (primary) only.

   Run either of the following commands with the server certificate file (public key) and the server certificate file (private key) to be transferred to the storage cluster specified.

   REST API: POST /v1/objects/server-certificate/actions/import/invoke

   CLI: server_certificate_import

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

2. 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.

3. Verify the server certificate.

   Confirm the requirements in *Client requirements for SSL communication*. Then, run a REST API or CLI or display a GUI by using a browser, and then confirm that a security warning* does not appear. If a security warning appears, contact customer support.

   * A security warning is an error message that contains any of the following text.

   "SSL", "TLS", "security certificate", "not protected", "not safe", "server certificate"

   If you are using a GUI, verify the message (including detailed error information), in the error window of the browser.

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.

   > ⚠ **Caution:**
   >
   > If you want to perform the import again due to an operation error such as importing the wrong server certificate, allow at least 60 seconds between imports to prevent operation conflicts.

Client requirements for SSL/TLS communication

Requirements for connection with server certificate validation enabled

(Virtual machine) To establish a connection with server certificate verification enabled, the Subject Alternative Name(SAN) or CN information in the server certificate must include the connection destinations to be specified when operations (for example, adding/replacing storage nodes, changing/setting configuration information, and importing/exporting configuration files) is performed from the REST API, CLI, GUI (browser), or maintenance node.

(Bare metal) To establish a connection with server certificate verification enabled, the Subject Alternative Name(SAN) or CN information in the server certificate must include the connection destinations to be specified when operation is performed from the REST API, CLI, or GUI (browser).

In this section, "Subject Alternative Name" is abbreviated to "SAN".

Note:

In server certificate verification of SSL/TLS communication, the connection source verifies whether the IP address or FQDN specified by the connection source as the connection destination is included in the SAN or CN information in the server certificate for the connection destination.

The following table shows examples of the connection destinations to be specified in SAN or CN in the server certificate when server certificate verification is performed.

If you create a server certificate as described in *Creating a certificate signing request*, the created server certificate content corresponds to item 1 or 2 of the following table.

No.	SAN	CN	Connection destination to be specified when server certificate verification is performed
1	▪ FQDN corresponding to the representative cluster IP address (Example: storage.example.com) ▪ FQDN corresponding to the IP address of the control port for each storage node (Example: storage-node1.example.com, storage-node2.example.com, ...)	-	Specify either of the following items: ▪ FQDN corresponding to the representative cluster IP address ▪ FQDN corresponding to the IP address of the control port for each storage node Specifying an IP address results in a server certificate verification error.
2	▪ Generic FQDN that uses a wildcard character to represent the FQDN that corresponds to the representative cluster IP address and the FQDNs that correspond to the IP addresses of the control ports for the storage nodes (Example: *.example.com)	-	

No.	SAN	CN	Connection destination to be specified when server certificate verification is performed
3	-	▪ Generic FQDN that uses a wildcard character to represent the FQDN that corresponds to the representative cluster IP address and the FQDNs that correspond to the IP addresses of the control ports for the storage nodes (Example: *.example.com)	
4	▪ Representative cluster IP address (Example: 192.0.2.100) ▪ IP address of the control port for each storage node (Example: 192.0.2.101, 192.0.2.102, ...)	-	Specify either of the following items: ▪ Representative cluster IP address ▪ IP address of the control port for the storage node Specifying an FQDN results in a server certificate verification error.

Cipher suite requirements

The client software from which REST APIs are used and the browser from which the GUI is used must meet the following requirements for SSL/TLS cipher suites.

Type of server certificates imported to Virtual Storage Software block	Cipher suite requirements
RSA certificate	At least one of the following is supported. ▪ TLS_DHE_RSA_WITH_AES_256_GCM_SHA384 ▪ TLS_DHE_RSA_WITH_AES_128_GCM_SHA256 ▪ TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 ▪ TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
ECC certificate	At least one of the following is supported. ▪ TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 ▪ TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256

Note:

- Cipher suite requirements differ depending on whether an RSA server certificate or ECC server certificate is imported to Virtual Storage Software block. By default, an RSA server certificate is imported. If you have created and imported a server certificate according to this manual, an RSA certificate is imported.
- If the cipher suite requirements are not met, SSL/TLS communication cannot be established.
- CLIs running on the Linux* controller nodes support the preceding cipher suites.
- CLIs running on the Windows* controller nodes support the preceding cipher suites only when cipher suites are enabled. Cipher suites are enabled by default on the OS. For details, see the Microsoft documentation.
- For whether other than CLIs on the controller nodes support the preceding cipher suites, see the documentation of the client software from which you use REST APIs or the browser from which you use the GUI.
- (Virtual machine)

  The preceding cipher suites are supported for operations on the maintenance node (adding and replacing nodes, changing and setting configuration information, importing and exporting configuration files, CLIs, and curl).

* For OSs supported on the controller nodes, contact customer support.

Action to be taken when a warning message about a server certificate appears

(Virtual machine) When you connect to Virtual Storage Software block via SSL/TLS communication to perform operations (such as adding/replacing devices, changing/setting configuration information, and importing/exporting configuration files) from the REST API, CLI, or GUI (browser) or maintenance node, if the server certificate is not one issued by a trusted certification authority, a warning message about the server certificate appears. To resolve this issue, import a server certificate issued by a trusted Certificate Authority to Virtual Storage Software block.

(Bare metal) When you connect to Virtual Storage Software block via SSL/TLS communication to perform operation from the REST API, CLI, or GUI (browser), if the server certificate is not one issued by a trusted certification authority, a warning message about the server certificate appears. To resolve this issue, import a server certificate issued by a trusted Certificate Authority to Virtual Storage Software block.

Also, confirm whether the requirements shown in *Client requirements for SSL/TLS communication* are met.

You can still make an SSL/TLS connection in this case. However, it will not be secure. (The connection is encrypted but the destination is not authenticated.)

- If you see a warning in the REST API, add or optionally ignore certificate warnings according to the program or command that issues the REST API.
- (Virtual machine) If a warning message appears when you perform an operation (such as adding/replacing devices, changing/setting configuration information, or importing/exporting configuration files) from the CLI or maintenance node, add the --ignore_certificate_errors option to ignore the warning message.
- (Bare metal) If a warning message appears when you perform an operation from the CLI, add the --ignore_certificate_errors option to ignore the warning message.
- If this warning appears while you are using a web browser, click "Continue to this website".

  The operation differs slightly depending on the web browser you use.

SSL/TLS communication with VMware vCenter Server (Virtual machine)

The SSL/TLS communications with VMware vCenter Server are established to perform a storage node installation, addition, and replacement, the configuration information change and setting, the configuration file import and export, and the configuration backup and restore on the maintenance node. Server certificate verification that is performed in those operations differs from server certificate verification described earlier in the following points.

Verification target	How to set verification	Valid period of the setting
Storage node	--ignore_certificate_errors option for each command	During the command execution
VMware vCenter Server	change_certificate_action command	Recorded as an attribute for each user. This setting is valid even if the user is logged out.

When server certificate verification for VMware vCenter Server is unsuccessful, an error message that contains "SSL connection", "TLS connection", or "certificate" is output.

- For how to change the server certificate for VMware vCenter Server, see the VMware documentation.
- If you have changed the server certificate, see *Add your own root CA certificate to the maintenance node (Virtual machine)* in this manual to perform necessary operation.

- For details about the change_certificate_action command, see *Managing maintenance nodes (Virtual machine)* in this manual.

- For how to specify the server certificate verification setting, contact customer support.

Chapter 4: Using CHAP authentication

Creating a CHAP user and configuring CHAP authentication (CLI or REST API)

Challenge-Handshake Authentication Protocol (CHAP) authentication can be used to verify if a connection request to the storage system comes from a valid compute node.

You can use CHAP authentication only if the compute node uses iSCSI connection.

For each compute port, you can set whether CHAP authentication is used.

The procedure for creating a CHAP user and set CHAP authentication is as follows.

The following table lists the system requirements for CHAP authentication.

Item	Requirement	Remarks
Maximum number of CHAP users	1024 per protection domain	Same number as the maximum number of compute nodes
Combination of a CHAP user name and a CHAP secret	The combination of a CHAP user name and a CHAP secret must be unique in the system.	
CHAP user name	Number of characters: 1 to 223 Allowed character types: Numbers (0 to 9), upper-case alphabet (A to Z), lower-case alphabet (a to z), space, symbols (. - + @ _ = : [ ] ~)	The conventions apply to the following parameter settings: ▪ TargetChapUserName(CLI: --target_chap_user_name) ▪ initiatorChapUserName(CLI: --initiator_chap_user_name)

Item	Requirement	Remarks
CHAP secret	Number of characters: 12 to 32 Allowed character types: Numbers (0 to 9), upper-case alphabet (A to Z), lower-case alphabet (a to z), space, symbols (. - + @ _ = : / [ ] ~)	The conventions apply to the following parameter settings: ▪ TargetChapSecret(CLI: --target_chap_secret) ▪ initiatorChapSecret(CLI: --initiator_chap_secret)

Caution:

- When changing the CHAP authentication setting, Virtual Storage Software block forcibly disconnects iSCSI connection between the compute node and the compute port to discard the connection before the setting change for safety. It is recommended to disconnect the iSCSI connection between the compute node and the compute port according to the disconnection procedure of each OS in advance. After changing the CHAP authentication setting, establish the iSCSI connection according to the changed setting.

- When a VPS is created, if you configure CHAP authentication, CHAP authentication must be performed for all connection requests to storage systems, including the connection requests from the compute node in the VPS to storage systems. For this reason, if you configure CHAP authentication, make sure that you notify the VPS administrator.

Before you begin

Required role: Security

Procedure

1. Create a CHAP user.

   Run either of the following commands with a CHAP user name and a CHAP secret specified.

   As required, specify a CHAP user name and a CHAP secret for mutual CHAP authentication.

   REST API: POST /v1/objects/chap-users

   CLI: chap_user_create

   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. **Obtain a list of compute ports to verify the ID of the compute port to be specified.**

   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

4. **Edit the authentication settings of the intended compute port.**

   Run the command with the following specified: ID of the compute port, authentication scheme of the compute port, whether CHAP authentication is enabled at the time of discovery in iSCSI connection, and whether mutual CHAP authentication is enabled.

   If you use the CLI, you can specify the WWN or iSCSI name instead of the compute port ID.

   The same CHAP user name cannot be created twice on the same compute port.

   REST API: PATCH /v1/objects/port-auth-settings/<*id*>

   CLI: port_auth_setting_set

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

5. **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.

6. **Allow the CHAP user to access the compute port.**

   Run either of the following commands with the ID of the compute port and the IDs of the CHAP users who are allowed to access the compute port with CHAP authentication specified.

   If you use the CLI, you can specify the CHAP user name instead of the CHAP user's ID.

   REST API: POST /v1/objects/port-auth-settings/<*id*>/chap-users

   CLI: port_auth_setting_chap_user_create

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

7. **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.

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 CHAP user information (CLI or REST API)

Obtain a list of CHAP user information as follows. The following information can be obtained.

- id: CHAP user ID (uuid)
- targetChapUserName: CHAP user name used for CHAP authentication on the compute port (i.e., target side)
- initiatorChapUserName: CHAP user name used for CHAP authentication on the initiator side of the compute node

Before you begin

Required role: Security

Procedure

1. Run either of the following commands to obtain a list of CHAP user information.

   REST API: GET /v1/objects/chap-users

   CLI: chap_user_list

Obtaining CHAP user information (CLI or REST API)

The following information can be obtained.

- portIds: List of compute port IDs (uuid) which the CHAP user is allowed to access with CHAP authentication
- id: CHAP user ID (uuid)
- targetChapUserName: CHAP user name used for CHAP authentication on the compute port (i.e., target side)
- initiatorChapUserName: CHAP user name used for CHAP authentication on the initiator side of the compute node

Before you begin

Required role: Security

Procedure

1. Verify the ID of the CHAP user.

   If you use the CLI to specify a CHAP user with a CHAP user name, check the CHAP user name.

   REST API: GET /v1/objects/chap-users

   CLI: chap_user_list

2. Obtain CHAP user information with the ID of the CHAP user specified.

   If you use the CLI, you can specify the CHAP user name instead of the CHAP user's ID.

   REST API: GET /v1/objects/chap-users/<*chapUserId*>

   CLI: chap_user_show

Editing CHAP user information (CLI or REST API)

Edit CHAP user information as follows.

Before you begin

Required role: Security

Procedure

1. Verify the ID of the CHAP user.

   If you use the CLI to specify a CHAP user with a CHAP user name, check the CHAP user name.

   REST API: GET /v1/objects/chap-users

   CLI: chap_user_list

2. Edit CHAP user information.

   Run either of the following commands with the ID of the CHAP user, CHAP user name, and a CHAP secret specified. As required, specify a CHAP user name and a CHAP secret for mutual CHAP authentication.

   If you use the CLI, you can specify the CHAP user name instead of the CHAP user's ID.

   REST API: PATCH /v1/objects/chap-users/<*chapUserId*>

   CLI: chap_user_set

   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.

Deleting a CHAP user (CLI or REST API)

Delete a CHAP user as follows.

Before you begin

Required role: Security

Procedure

1. Verify the ID of the CHAP user to be deleted.

   If you use the CLI to specify a CHAP user with a CHAP user name, check the CHAP user name.

   REST API: GET /v1/objects/chap-users

   CLI: chap_user_list

2. Delete a CHAP user.

   Run either of the following commands with the ID of the CHAP user specified.

   If you use the CLI, you can specify the CHAP user name instead of the CHAP user's ID.

   REST API: DELETE /v1/objects/chap-users/<*chapUserId*>

   CLI: chap_user_delete

   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 authentication settings for the compute port for the target operation (CLI or REST API)

The following information can be obtained.

- id: Compute port ID (uuid)

- authMode: Authentication scheme of the compute port

- isDiscoveryChapAuth: Whether CHAP authentication is enabled at the time of discovery in iSCSI connection
- isMutualChapAuth: Whether mutual CHAP authentication is enabled

Before you begin

Required role: Security

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 the authentication 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/port-auth-settings/<*id*>

   CLI: port_auth_setting_show

Editing the authentication settings for the compute port for the target operation (CLI or REST API)

Edit the authentication settings for the compute port for the target operation as follows.

Caution:

- When changing the CHAP authentication setting, Virtual Storage Software block forcibly disconnects iSCSI connection between the compute node and the compute port to discard the connection before the setting change for safety. It is recommended to disconnect the iSCSI connection between the compute node and the compute port according to the disconnection procedure of each OS in advance. After changing the CHAP authentication setting, establish the iSCSI connection according to the changed setting.
- When a VPS is created, editing the authentication settings affects the status of connection from the compute node in the VPS to a storage system. For this reason, after you edit the authentication settings, make sure that you notify the VPS administrator.

Before you begin

Required role: Security

Procedure

1. Verify the ID of the 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. Edit the authentication settings of the intended compute port.

   Run either of the following commands with the compute port ID and parameters for setting compute port authentication specified.

   If you use the CLI, you can specify the WWN or iSCSI name instead of the compute port ID.

   REST API: PATCH /v1/objects/port-auth-settings/<*id*>

   CLI: port_auth_setting_set

   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.

Allowing a CHAP user to access the compute port (CLI or REST API)

Allow the CHAP user to access the compute port through CHAP authentication as follows.

> ⚠ **Caution:**
>
> If you add a storage node after you grant access to the compute port of an existing storage node based on CHAP authentication, also grant CHAP users access to the compute port of the added storage node.

Before you begin

Required role: Security

Procedure

1. **Verify the ID of the applicable 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. Verify the ID of the CHAP user.

   If you use the CLI to specify a CHAP user with a CHAP user name, check the CHAP user name.

   REST API: GET /v1/objects/chap-users

   CLI: chap_user_list

3. Allow the CHAP user to access the compute port.

   Run either of the following commands with the ID of the compute port and the ID of the CHAP user specified.

   If you use the CLI, you can specify the WWN or iSCSI name instead of the compute port ID, or the CHAP user name instead of the CHAP user's ID.

   REST API: POST /v1/objects/port-auth-settings/<*id*>/chap-users

   CLI: port_auth_setting_chap_user_create

   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. 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 CHAP users who are allowed to access a compute port (CLI or REST API)

Obtain a list of information about CHAP users who are allowed to access a compute port as follows. The following information can be obtained.

- id: ID (uuid) of each CHAP user who is allowed to access the compute port
- targetChapUserName: CHAP user name used for CHAP authentication on the compute port (i.e., target side)
- initiatorChapUserName: CHAP user name used for CHAP authentication on the initiator side of the compute node

Before you begin

Required role: Security

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 a list of information about CHAP users who are allowed to access 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/port-auth-settings/<*id*>/chap-users

   CLI: port_auth_setting_chap_user_list

Obtaining information about individual CHAP users who are allowed to access a compute port (CLI or REST API)

Obtain information about a CHAP user who is allowed to access a compute port as follows. The following information can be obtained.

- id: ID (uuid) of a CHAP user who is allowed to access the compute port
- targetChapUserName: CHAP user name used for CHAP authentication on the compute port (i.e., target side)
- initiatorChapUserName: CHAP user name used for CHAP authentication on the initiator side of the compute node

Before you begin

Required role: Security

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. Verify the ID of the CHAP user.

   If you use the CLI to specify a CHAP user with a CHAP user name, check the CHAP user name.

   REST API: GET /v1/objects/chap-users

   CLI: chap_user_list

3. Obtain information about the CHAP user who is allowed to access the compute port.

   Run either of the following commands with the ID of the compute port and the ID of the CHAP user specified.

   If you use the CLI, you can specify the WWN or iSCSI name instead of the compute port ID, or the CHAP user name instead of the CHAP user's ID.

   REST API: GET /v1/objects/port-auth-settings/<*id*>/chap-users/<*chapUserId*>

   CLI: port_auth_setting_chap_user_show

Canceling permission for a CHAP user to access a compute port (CLI or REST API)

Cancel permission for a CHAP user to access a compute port for target operation in CHAP authentication as follows.

⚠ **Caution:**

- When a VPS is created, editing the authentication settings affects the status of connection from the compute node in the VPS to a storage system. For this reason, after you edit the authentication settings, make sure that you notify the VPS administrator.

Before you begin

Required role: Security

Procedure

1. Verify the ID of the applicable 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. Verify the ID of the CHAP user whose access permission is to be canceled.

   If you use the CLI to specify a CHAP user with a CHAP user name, check the CHAP user name.

   REST API: GET /v1/objects/chap-users

   CLI: chap_user_list

3. Cancel permission for a CHAP user to access a compute port as follows.

   Run either of the following commands with the ID of the compute port and the ID of the CHAP user specified.

   If you use the CLI, you can specify the WWN or iSCSI name instead of the compute port ID, or the CHAP user name instead of the CHAP user's ID.

   REST API: DELETE /v1/objects/port-auth-settings/<*id*>/chap-users/<*chapUserId*>

   CLI: port_auth_setting_chap_user_delete

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

4. 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.

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 5: Managing event logs

Overview of event logs

Each time an event that needs to be reported to the user occurs while the storage cluster is operating, Virtual Storage Software block collects such an event and generates an event log. You can verify event logs to know what happened in the storage cluster.

An event log can contain up to 864,000 events. If this maximum limit is exceeded, the log is overwritten in first-in-first-out basis.

In addition, you can set the transfer destination of event logs to the Syslog server or SMTP server. Note that, transfer to the Syslog server and SMTP server requires you to run different respective commands. Event logs that are transferred to the SMTP server are sent as emails.

Event logs that are created after event log transfer setting are transferred to the Syslog server or SMTP server. Event logs created before event log transfer setting are not transferred.

Event logs transferred to the SMTP server are those whose severity is "Critical", "Error", or "Warning". The time from when an applicable event log is generated until an email is sent is approximately 1 minute and 20 seconds.

Note:

When the storage cluster is restarted or a storage node is recovered from maintenance, event logs that were issued before a failure might be reissued.

An event log is issued when the storage cluster is performing some processing. However, if a failure occurs during the processing, the processing is passed to a normal storage node and is retried. As a result, the same event log might be issued.

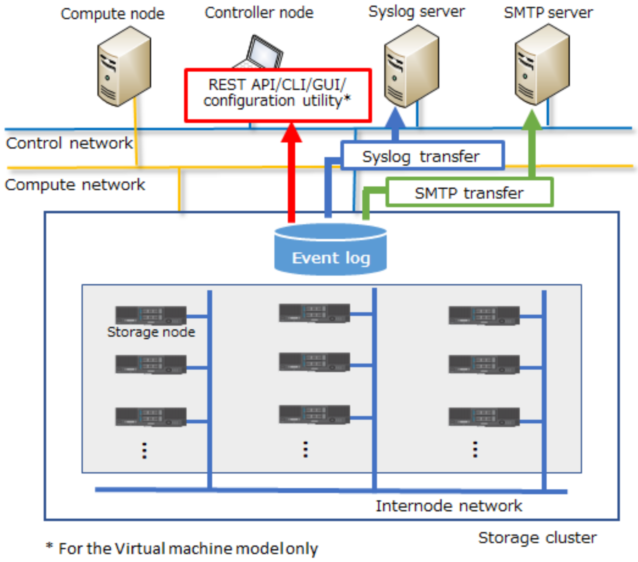
Compute node
Controller node
Syslog server
SMTP server
REST API/CLI/GUI/ configuration utility*
Control network
Compute network
Syslog transfer
SMTP transfer
Event log
Storage node
Internode network
Storage cluster
* For the Virtual machine model only

Requirements for the Syslog server and SMTP server

- Syslog server: Rsyslog 8 is supported.
- SMTP server: The requirements for the SMTP server are as follows:
  - STARTTLS is supported.
  - SMTP authentication is supported.
  - At least one of CRAM-MD5, PLAIN, or LOGIN is supported as the SMTP authentication method.
  - TLS 1.2 is supported.

    The SMTP client that is used by the storage system for connection with the SMTP server uses only TLS1.2. However, you should disable vulnerable protocol versions (SSL2.0, SSL3.0, TLS1.0, and TLS1.1) in the SMTP server settings.
  - At least one of the following is supported as a TLS cipher suite.
    - (1) TLS_DHE_RSA_WITH_AES_256_GCM_SHA384
    - (2) TLS_DHE_RSA_WITH_AES_128_GCM_SHA256
    - (3) TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
    - (4) TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
    - (5) TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
    - (6) TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256

    - If the cipher suite requirements are not met, TLS connection to the SMTP server cannot be established.

    - For whether the SMTP server supports any of the preceding cipher suites, see the documentation for the SMTP server to be used.

    - You might need to specify the settings for key exchange (DHE or ECDHE) on the SMTP server according to the key exchange method (DHE or ECDHE) of the cipher suite to be used. For details about the settings to be specified, see the documentation for the SMTP server to be used.

    - The cipher suites that can be used differ depending on whether the server certificate set on the SMTP server is an RSA or ECC certificate. For an RSA certificate, the first to fourth cipher suites in the above can be used. For an ECC certificate, the fifth and sixth cipher suites in the above can be used. For details about how to set a server certificate on the SMTP server, see the documentation for the SMTP server to be used.

    - The SMTP client that is used by the storage system for connection with the SMTP server uses only the preceding cipher suites. However, you should disable the other vulnerable cipher suites on the SMTP server settings. Vulnerable cipher suites refer to the cipher suites shown in *Appendix A. TLS 1.2 Cipher Suite Black List* of RFC 7540.

- If multiple server certificates are to be set on the SMTP server, they must be issued so that their authenticity can be proved by only one root certificate imported into the storage system.
- SMTP client provided by Virtual Storage Software block allows insecure renegotiation for backward compatibility.

  For TLS connection with SMTP server, it is recommended to use an SMTP server that complies with RFC5746.

Obtaining a list of event logs (CLI or REST API)

The following information can be obtained.

- id: IDs (uuid) of events
- time: Time and dates when events are detected
- timeInMicroseconds: Time (in microseconds) elapsed from January 1, 1970, 00:00:00:00 until when the event was detected
- category: Categories of events (such as drive, storage pool)
- eventName: Unique names of events
- messageId: Message ID
- severity: Severity levels of events
- message: Descriptions of events
- solution: Recommended countermeasures for events
- nodeLocation: Information about the storage nodes where events occurred
- eventType: Event type. This item is reserved for future use. "Informational" is always output.
- severityLevel: Degree of urgency of the event. This item is reserved for future use. "1" is always output.

Before you begin

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

Procedure

1. Obtain a list of event logs.

   REST API: GET /v1/objects/event-logs

   CLI: event_log_list

Note:

Although the maximum number of event logs that can be retained is 864,000, the maximum number of event logs that you can obtain at one time by running a command is 1,000. If you run the command without the parameter for specifying the range (startTime or endTime), the latest 1,000 logs are obtained. To obtain the older event logs, run the command with startTime or endTime specified. For details, see the *Hitachi Virtual Storage Software Block REST API Reference* or the *Hitachi Virtual Storage Software Block CLI Reference*.

Obtaining individual event logs (CLI or REST API)

The following information can be obtained for the event logs of the specified ID.

- id: ID (uuid) of the target event
- time: Time and date when the target event is detected
- timeInMicroseconds: Time (in microseconds) elapsed from January 1, 1970, 00:00:00:00 until when the event was detected
- category: Category of the event (such as drive, storage pool)
- eventName: Unique name of the event
- messageId: Message ID
- severity: Severity level of the event
- message: Description of the event
- solution: Recommended countermeasures for the event
- nodeLocation: Information about the storage node where the event occurred
- eventType: Event type. This item is reserved for future use. "Informational" is always output.
- severityLevel: Degree of urgency of the event. This item is reserved for future use. "1" is always output.

Before you begin

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

Procedure

1. Verify the ID of the target event.

   REST API: GET /v1/objects/event-logs

   CLI: event_log_list

2. Obtain the target event log.

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

   REST API: GET /v1/objects/event-logs/<*id*>

   CLI: event_log_show

Editing Syslog transfer settings of event logs (CLI or REST API)

Edit Syslog transfer of event logs as follows. You can set a maximum of two Syslog servers. For the structure of event logs transferred to the Syslog server, see *Structure of event logs transferred to the Syslog server*.

> ⚠ **Caution:**
>
> When using a DNS server, a storage node caches DNS inquiry results for the time (DNS TTL) set in the DNS server. For this reason, if the content registered in the DNS server (correspondence between the host name and IP address) is changed, the storage node might access an old address during DNS TTL. Therefore, if you have changed the content registered in the DNS server (correspondence between the host name and IP address), wait until the time specified for DNS TTL has passed, and then set Syslog transfer.

> **Note:**
>
> When syslog transfer of event logs is set, Virtual Storage Software block periodically sends ICMP echo requests to the set Syslog server to establish network paths.

Before you begin

Required role: Security

Procedure

1. Set Syslog transfer of event logs as follows.

   Run either of the following commands with the parameters for setting event log Syslog transfer specified.

   REST API: PATCH /v1/objects/event-log-setting

   CLI: event_log_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.

4. **On the Syslog server, set whether event logs can be received from Virtual Storage Software block.**

   See the manual of the Syslog server in use, and make the following settings as required.

   - IP address of Syslog transmission source: In Syslog transfer, the source IP address of the IP packets sent from Virtual Storage Software block is as follows.

     - If the representative IP address of the storage cluster is not specified:

       Control network IP address of the cluster master node (primary)

     - If the representative IP address of the storage cluster is specified:

       Representative IP address of the storage cluster or control network IP address of the cluster master node (primary)

     The cluster master node (primary) might be switched to a different storage node (due to a storage node failure, for example). Therefore, we recommend that you configure the Syslog server so that it can receive Syslog messages from the representative IP address of the storage cluster and the control network IP addresses of all the storage nodes.

   - Port number of Syslog transmission source: Port number set in step 1

   - Communications protocol: Communications protocol set in step 1

5. **Verify that event logs are correctly transferred to the Syslog server as follows.**

   Make the Syslog transfer setting of event logs again by using the same input values as step 1.

   If the Syslog transfer setting in Virtual Storage Software block and the reception setting in the Syslog server are made properly, the following event log is transferred to the Syslog server, indicating that the job for setting event log Syslog transfer has been started. If the log is not transferred, review the setting and the network.

   - Event log name: Start of job

   - Event description: The job has started. JobId=[jobId],Operation=event_log_setting_set

Editing SMTP transfer settings of event logs (CLI or REST API)

Virtual Storage Software block acts as an SMTP client and transfers an e-mail containing event log information to the outgoing SMTP server. A user is responsible for preparing an SMTP server and POP/IMAP server required for mail services. The following figure shows an example configuration required for mail services. Depending on operations, the configuration might include additional components, such as an SMTP server for relaying e-mails.

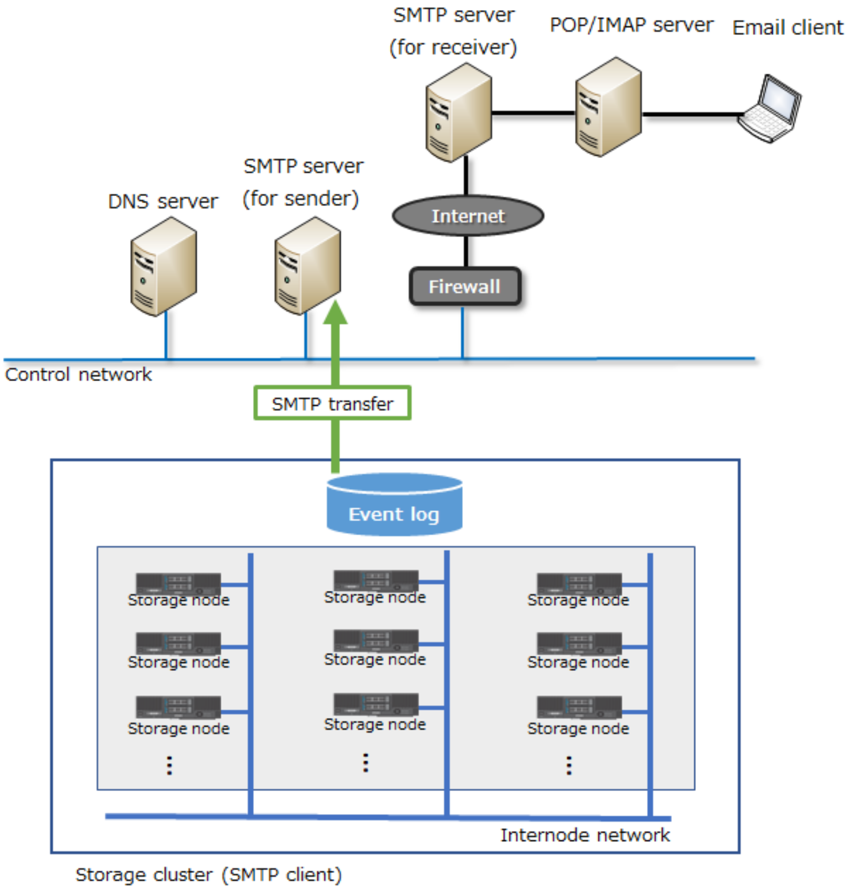

To use SMTP transfer, edit SMTP transfer settings of event logs as follows. Import a root certificate that proves the authenticity of the server certificates set on the outgoing SMTP server to the storage cluster. This enables TLS communication with the outgoing SMTP server. PEM- and DER-format certificate files are supported.

For the format of emails sent via the outgoing SMTP server, see *Email format*.

Event logs transferred to the SMTP server are those whose severity is "Critical", "Error", or "Warning". The time from when an applicable event log is generated until an email is sent is approximately 1 minute and 20 seconds.

Note:

In SMTP transfer, the source IP address of the IP packets sent from Virtual Storage Software block is as follows.

- If the representative IP address of the storage cluster is not set:

  Control network IP address of the cluster master node (primary)

- If the representative IP address of the storage cluster is set:

  Representative IP address of the storage cluster or control network IP address of the cluster master node (primary)

The cluster master node (primary) might be switched to a different storage node (due to a storage node failure, for example). Therefore, we recommend that you configure the SMTP server so that it can receive emails sent from the representative IP address of the storage cluster and the control network IP addresses of all the storage nodes.

Before you begin

Required role: Security

Procedure

1. Import a root certificate that proves the authenticity of the server certificates set on the outgoing SMTP server to the storage cluster.

   You can perform this for the cluster master node (primary) only.

   Run either of the following commands with the certificate file and the outgoing SMTP server ID specified.

   REST API: POST /v1/objects/smtp-server-root-certificates/<*targetServer*>/actions/import/invoke

   CLI: smtp_server_root_certificate_import

   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.

   If TLS communication with the outgoing SMTP server is enabled, the root certificate is applied immediately.

3. Set SMTP transfer of event logs.

   Run either of the following commands with the parameters for setting email notification for the event log transfer destination specified.

   REST API: PATCH /v1/objects/event-log-setting

   CLI: event_log_setting_set

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

> **Caution:**
>
> When using a DNS server, a storage node caches DNS inquiry results for the time (DNS TTL) set in the DNS server. For this reason, if the content registered in the DNS server (correspondence between the host name and IP address) is changed, the storage node might access an old address during DNS TTL. Therefore, if you have changed the content registered in the DNS server (correspondence between the host name and IP address), wait until the time specified for DNS TTL has passed, and then set SMTP transfer.

4. 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.

5. Verify that SMTP transfer of event logs is set correctly.

   If the SMTP transfer setting and the setting on the outgoing SMTP server side are made properly, an event log is notified via an email (Subject = VSSB-Report KARS10603-I), indicating that the job for setting event log SMTP transfer has been started. Although the email is sent immediately, it might be delayed depending on the network condition. If the email client could not receive the email, go to the next step.

6. Obtain a list of event logs.

   REST API: GET /v1/objects/event-logs

   CLI: event_log_list

7. If one of the following event logs is output, take the specified action.

messageId	Action to be taken
KARS10650-W	See the specific error cause described in the Detailed Information section of the event log, verify and correct the SMTP transfer settings and outgoing SMTP server settings, and then set SMTP transfer.
KARS10652-W	Import a root certificate, and then set SMTP transfer.
KARS10651-W	Verify whether the root certificate has expired or is invalid, import the correct root certificate, and then set SMTP transfer.

If none of the above event logs exists, Virtual Storage Software block has successfully transferred the e-mail to the outgoing SMTP server. This implies that Virtual Storage Software block as an SMTP client cannot detect the e-mail for the following reasons:

- In the outgoing SMTP server settings specified in the SMTP transfer settings, there is a problem in the settings for relaying the email to another SMTP server.
- In the outgoing SMTP server settings specified in the SMTP transfer settings, there is a problem in the settings for transferring the email to the POP/IMAP server.
- Setting for POP/IMAP server, SMTP server for relay, email client (or another component) is incorrect.
- A network problem exists between components.

The following table lists the major actions performed by Virtual Storage Software block as an SMTP client and event logs notified when an error occurs. In step 7, if there is no event log that indicates SMTP transfer error, the following operations were successfully performed. If the email client cannot receive the email, refer to this table for identifying the cause.

Action	Event log upon error
If a host name is specified for smtpServerName in the SMTP transfer settings, resolve the host name.	KARS10650-W (Detailed information = Connection to the SMTP server is not established.)
Establish a TCP connection with the outgoing SMTP server.	KARS10650-W (Detailed information = Connection to the SMTP server is not established.)
Send the SMTP protocol EHLO command over the TCP protocol, and check the response.	KARS10650-W (Detailed information = Connection to the SMTP server is not established.)
Use the SMTP protocol STARTTLS command to verify that the outgoing SMTP server allows TLS communication.	KARS10650-W (Detailed information = STARTTLS feature is not available in the SMTP server.)
Make sure the outgoing SMTP server supports TLS 1.2.	KARS10650-W (Detailed information = TLS 1.2 is not available in the SMTP server.)
Make sure the outgoing SMTP server supports a cipher suite that meets the requirements.	KARS10650-W (Detailed information = Connection to the SMTP server is not established.)

Action	Event log upon error
[Supported cipher suites] - TLS_DHE_RSA_WITH_AES_256_GCM_SHA384 - TLS_DHE_RSA_WITH_AES_128_GCM_SHA256 - TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 - TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 - TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 - TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	
Validate the server certificate of the outgoing SMTP server using the root certificate imported into the Virtual Storage Software block.	KARS10651-W
Send the SMTP protocol EHLO command over the TLS protocol, and verify the response.	KARS10650-W (Detailed information = Connection to the SMTP server is not established.)
Use the SMTP protocol AUTH command to send the smtpAuthAccount and smtpAuthPassword SMTP transfer settings and be authenticated by the outgoing SMTP server.	KARS10650-W (Detailed information = Authentication feature is not available in the SMTP server.) Or, KARS10650-W (Detailed information = Authentication is failed.)
Use the SMTP protocol MAIL FROM command to send the fromAddress SMTP transfer setting and inform the source address to the outgoing SMTP server.	KARS10650-W (Detailed information = Sender address is not accepted in the SMTP server.)
Use the SMTP protocol RCPT TO command to send the toAddress1, toAddress2, and toAddress3 SMTP transfer settings and inform the destination address to the outgoing SMTP server.	KARS10650-W (Detailed information = Recipient address is not accepted in the SMTP server.)

Action	Event log upon error
Use the SMTP protocol DATA command to send the email title and body information to the outgoing SMTP server.	KARS10650-W (Detailed information = Sent data is not accepted in the SMTP server.)
Send the SMTP protocol QUIT command.	None

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.

Caution:

After notifying an event log indicating an error (KARS10650-W, KARS10651-W, or KARS10652-W), the Virtual Storage Software block SMTP transfer function suppresses further notification of an event log with the same messageId until the problem is resolved. To enable notification of an event log with the same messageId again to identify the cause when the email client cannot receive the email, perform the following procedure.

a. In the event log SMTP transfer settings, set isEnabled to False. For any required parameters other than isEnabled, set a specifiable value.

REST API: PATCH /v1/objects/event-log-setting

CLI: event_log_setting_set

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

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

CLI: job_show

If you receive a response indicating "Succeeded" as the state, information configured in the SMTP transfer settings has been changed or set. Wait for approximately 20 seconds until Virtual Storage Software block operates as an SMTP client (according to the configuration information), and then set SMTP transfer of event logs by repeating the procedure from step 3.

> **Note:**
>
> - When the cluster master node (primary) is blocked, an email that was already sent might be sent again, depending on the timing.
> - If the same address is set for destination email addresses 1 to 3, the same email might be sent more than once, depending on the SMTP server setting.
> - You can obtain a root certificate that proves the authenticity of the server certificates set on the outgoing SMTP server by running either of the following commands.
>
>   You can perform this for the cluster master node (primary) only.
>
>   Run the command with the ID of the outgoing SMTP server specified.
>
>   A root certificate is obtained as a DER file.
>
>   REST API: GET /v1/objects/smtp-server-root-certificates/<*targetServer*>/download
>
>   CLI: smtp_server_root_certificate_download

Obtaining event log settings (CLI or REST API)

Obtain the settings for transferring event log data to syslog and the settings for reporting event log data by email. The following information can be obtained.

- syslogForwardingSetting: Settings for transferring event log data to syslog
- emailReportSetting: Settings for reporting event log data by email

Before you begin

Required role: Security

Procedure

1. Obtain the event log settings by running either of the following commands:

   REST API: GET /v1/objects/event-log-setting

   CLI: event_log_setting_show

Monitoring the event logs transferred to the Syslog server

Sometimes, event logs are not transferred to the Syslog server due to network failures. In such a case, verify the event log in Virtual Storage Software block.

Procedure

1. Verify the *message serial numbers* of the event logs received by the Syslog server.

   Message serial numbers are serial numbers allocated to events. By verifying *message serial numbers*, you can make sure that all event logs are transferred. If some message serial numbers are missing, some event logs were not transferred to the Syslog server due to network failure or other causes. In such a case, verify the event log by the following procedure.

2. Obtain a list of event logs (see *Obtaining a list of event logs*). Verify the event logs whose message serial numbers are missing in the Syslog server.

Structure of event logs transferred to the Syslog server

Each piece of event log information transferred to the Syslog server is structured as follows. No line feed code is inserted.

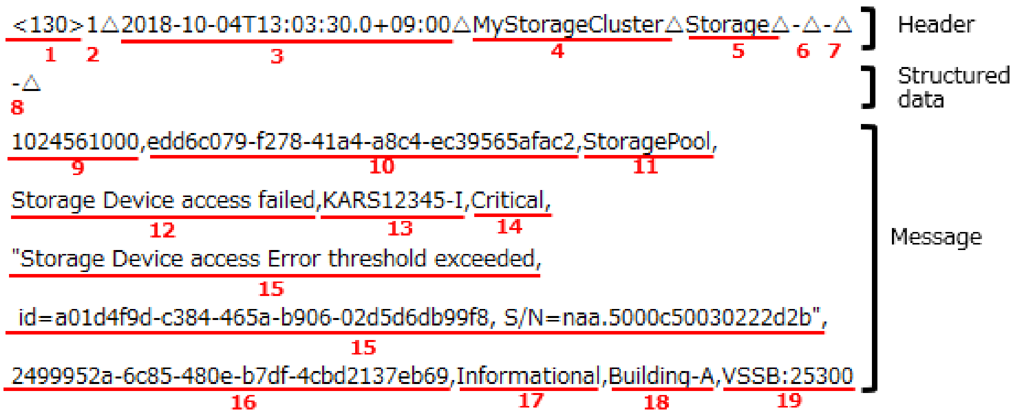

```
<130>1△2018-10-04T13:03:30.0+09:00△MyStorageCluster△Storage△-△-△
-△
1024561000,edd6c079-f278-41a4-a8c4-ec39565afac2,StoragePool,
Storage Device access failed,KARS12345-I,Critical,
"Storage Device access Error threshold exceeded,
 id=a01d4f9d-c384-465a-b906-02d5d6db99f8, S/N=naa.5000c50030222d2b",
2499952a-6c85-480e-b7df-4cbd2137eb69,Informational,Building-A,VSSB:25300
```

Header part

The following is an example of the header part. Each triangle (△) in the example indicates a space. Each number in the example corresponds to the equivalent in the following table.

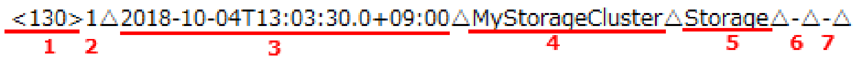

```
<130>1△2018-10-04T13:03:30.0+09:00△MyStorageCluster△Storage△-△-△
```

Number	Item	Description
1	Priority	▪ <130>: When the severity of the event is "Critical" ▪ <131>: When the severity of the event is "Error" ▪ <132>: When the severity of the event is "Warning" ▪ <134>: When the severity of the event is "Info"

Number	Item	Description
2	Version	Fixed to "1".
3	Date and time	The output format is "YYYY-MM-DDThh:mm:ss.s ±hh:mm" (YYYY: Year, MM: Month, DD: Date, hh: Hour, mm: Minute, ss.s: Second). The date and time are in coordinated universal time (UTC). ▪ "±hh:mm" indicates the time difference between the UTC and the time when the event occurred. "+" means that the time when the event occurred is hh:mm faster than the UTC. "-" means that the time when the event occurred is hh:mm slower than the UTC. ▪ "Z" attached to the date and time as in 2019-03-20T23:06:58.0Z means that the date and time is the same as the UTC.
4	Storage cluster name	Storage cluster name.
5	Program name	Fixed to "Storage".
6	Process name	Fixed to "-".
7	Message name	Fixed to "-".

Structured data part

The following is an example of the structured data part. Each triangle (△) in the example indicates a space. Each number in the example corresponds to the equivalent in the following table.

-△

Number	Item	Description
8	Structured data	Fixed to "-".

Message part

The following is an example of the message part. Each triangle (△) in the example indicates a space. Each number in the example corresponds to the equivalent in the following table.

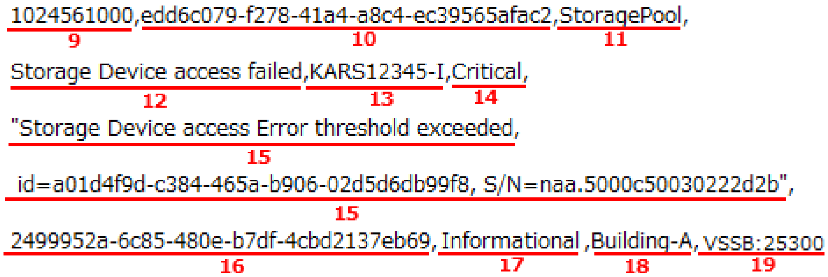

```
1024561000,edd6c079-f278-41a4-a8c4-ec39565afac2,StoragePool,
Storage Device access failed,KARS12345-I,Critical,
"Storage Device access Error threshold exceeded,
 id=a01d4f9d-c384-465a-b906-02d5d6db99f8, S/N=naa.5000c50030222d2b",
2499952a-6c85-480e-b7df-4cbd2137eb69,Informational ,Building-A,VSSB:25300
```

9
10
11
12
13
14
15
15
16
17
18
19

Note:

"%xEF.BB.BF" is added as the byte order mark (BOM) to the beginning of the message part.

Number	Item	Description
9	Message serial number	Serial number assigned to each event (Range: 0000000001 to 9999999999).
10	Event ID	Event ID.
11	Event category	Indicates an event category (such as drive and storage pool).
12	Unique event name	Event name uniquely given.
13	Message ID	Message ID.
14	Severity of the event	Severity of the event. ▪ Info: Notified for your information. ▪ Warning: Warning level. Immediate action is not required, but this event might cause a major failure in the future. ▪ Error: Error level. Corrective action is required. ▪ Critical: Fatal failure. Immediate action is required.
15	Event description	Event description.
16	Information about a storage node on which the event occurred	Information about a storage node on which the event occurred.
17	Event type	Type of the event. This item is reserved for future use. "Informational" is always output.
18	Location name	Location of the event output source, such as the installation location of equipment.

Number	Item	Description
19	Model name : Serial number	Model name and serial number.

Email format

The format of emails sent via the SMTP server is as follows.

Item	Format and description
Email title	VSSB-Report <*Message-ID*>

Header information	`----------------------------------------------------------------------`
	`This is a notification report of the VSSB. This is automatically sent to`
	`the registered addresses.`
	`Please refer to the reported event and take action if required.`
	`----------------------------------------------------------------------`
	Storage cluster name: Nickname of the storage cluster
	Storage cluster version: Version of the storage cluster
	Timezone: Time zone
	Date and time: Collection date and time (yyyy/MM/dd hh:mm:ss.fff)
	Severity: Severity of the event
	Message ID: Message ID
	Description: Event log message body
	Action: Action to be taken
	Event ID: Event ID
	`----------------------------------------------------------------------`
	`The followings show the events that occurred before and after the`
	`reported event.`
	`Note that * is marked at the beginning of the line corresponding to the`
	`reported event.`
	`----------------------------------------------------------------------`

List of event logs	A list showing the specific event log transferred by email and any previous and subsequent event logs is displayed. A maximum of 50 event logs (generated 30 minutes before the specific event log transferred by email) are displayed.

Item	Format and description
	A maximum of 50 event logs (generated 1 minute after the specific event log transferred by email) are displayed. Event logs displayed after the event log transferred by email are no longer notified by email even if they meet email transfer conditions. Format: <*collection-date-and-time*> <*severity*> <*message-ID*> <*event-log-message-body*> <*event-log-ID*> Example: 20xx/12/10 17:40:33.145 Info KARS08100-I The storage cluster started. aaee3c10-9aa6-45ef-84d9-b2b4e32a242c An asterisk (*) is indicated before the event log. Example: * 20xx/12/10 17:40:33.238 Error KARS08103-E The storage node, storage_node_uuid_1234567890 was blocked. fc1f2299-d0fc-42ed-96f9-c310b029443a

Using SNMP

Virtual Storage Software block has an SNMP agent function that issues failure information (pertaining to the Virtual Storage Software block storage cluster) as event logs (using the SNMP protocol) to an SNMP manager connected to a control network prepared by the user.

Also, you can verify the system configuration and operating state by using the SNMP protocol.

Virtual Storage Software block, which has detected failure information regarding the devices to be managed, sends a message called a trap to the SNMP manager without being requested by the SNMP manager.

To use SNMP, set items such as SNMP enable or disable, SNMP trap issuance destination, SNMP-based request permission, and system group information according to *Setting SNMP access control*.

For the SNMP protocol, a standard structure is defined for management information and database, which is called "Management Information Base (MIB)". Information about the hardware that has a failure is defined by MIB. In addition to a standard MIB, an extended MIB is provided.

The failure information issued to the SNMP manager includes information from event logs whose severity is "Critical", "Error", or "Warning".

Failure information issued to the SNMP manager is based on the event logs created after setting SNMP. Event logs created before setting SNMP are not applicable.

SNMP manager and environment requirements

To use each SNMP function, the following requirements for SNMP manager and environment must be prepared by the user:

Item		Requirement
Communication protocol		UDP must be supported.
Communication port	SNMP command	Must be able to send a Get request to destination port 161.
	SNMP trap	Must be able to receive a packet destined for port 162.
SNMP version		v2c must be supported.
Setting a source IP address for a trap to be received		In the event log SNMP transfer settings, the source IP address is as follows: ▪ If the representative IP address of the storage cluster is not set: Control network IP address of the cluster master node (primary) ▪ If the representative IP address of the storage cluster is set: Representative IP address of the storage cluster or control network IP address of the cluster master node (primary) The cluster master node (primary) might be switched to a different storage node (due to a storage node failure, for example). Therefore, we recommend that you configure SNMP server so that it can receive a trap from the representative IP address of the storage cluster and the control network IP addresses of all the storage nodes.

Setting SNMP access control (CLI or REST API)

Set SNMP access control as follows.

When you specify the community, storageSystemName (CLI: storage_system_name), contact, or location parameters for editing SNMP settings, observe the following:

- Number of characters: Maximum of 180
- Characters that can be used: Numbers (0 to 9), uppercase alphabet (A to Z), lowercase alphabet (a to z), symbols (! # $ ' ( ) + -.@ _ ` { } ~)

  Alternatively, spaces can also be used except for the starting and ending positions, where spaces are not allowed.

> **Caution:**
>
> When using a DNS server, a storage node caches DNS inquiry results for the time (DNS TTL) set in the DNS server. For this reason, if the content registered in the DNS server (correspondence between the host name and IP address) is changed, the storage node might access an old address during DNS TTL. Therefore, if you have changed the content registered in the DNS server (correspondence between the host name and IP address), wait until the time specified for DNS TTL has passed, and then edit SNMP settings.

Before you begin

Required role: Security

Procedure

1. Set SNMP access control.

   Run either of the following commands with the parameters for editing SNMP settings specified.

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

   CLI: snmp_setting_set

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

2. 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.

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.

4. If the settings in step 1 are correct and a trap transmission destination was specified, a trap for verifying the settings is issued. If the following event log has not been transferred, verify the settings and the network again.

   - messageId: KARS10607-I

   - SNMP setting complete. JobId=[JobId]

Viewing SNMP access control (CLI or REST API)

Obtain SNMP access control as follows. You can obtain the following information.

- isSNMPAgentEnabled: Whether SNMP is enabled or disabled
- snmpVersion: SNMP version
- sendingTrapSetting: SNMP trap transmission destination setting
- requestAuthenticationSetting: Whether to permit SNMP-based requests
- systemGroupInformation: System group information

Before you begin

Required role: Security

Procedure

1. Obtain SNMP access control.

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

   CLI: snmp_setting_show

Standard MIB

The structure of the standard MIB in Virtual Storage Software block is as follows. The standard MIB supports only the System group.

For the extended MIB in Virtual Storage Software block, see *Installing an extended MIB*.

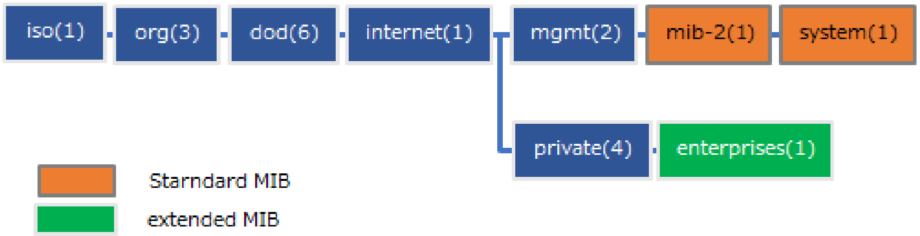

The following table shows the structure and contents of the System group.

OID	Name	Description	R/W attribute	Remarks
1	sysDescr	Product name	R	
2	sysObjectID	Start OID of the extended MIB	R	
3	sysUpTime	Operating time of the SNMP agent*	R	

OID	Name	Description	R/W attribute	Remarks
4	sysContact	Contact	R/W	
5	sysName	System name	R/W	
6	sysLocation	Installation location	R/W	
7	sysServices	Supported services	R	
8	sysORLastChange	-	-	Not supported
9	sysORTable	-	-	Not supported
* Reset by a restart of the SNMP agent				

Installing an extended MIB

To reference the data received by an SNMP command (Get, etc.) or a trap, an extended MIB file of Virtual Storage Software block must be included in the SNMP manager.

The name of the extended MIB file is as follows:

- File name: sdsExMib -<*version*>-<*number*>.mib

You can obtain the extended MIB file in the product media.

If you updated the storage software, obtain the file for which the <*version*> part of the name is the same as the version of the updated storage software from customer center.

The structure of the extended MIB of Virtual Storage Software block is as follows:

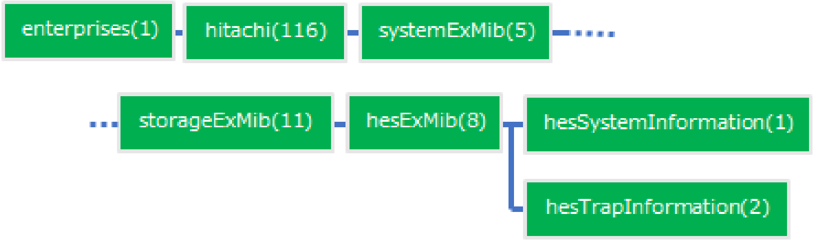

OID	Name	Description	R/W attribute	Remarks
1	sdsSystemInformation(1)	System information of Hitachi Virtual Storage Software block	R	

OID	Name	Description	R/W attribute	Remarks
2	sdsTrapInformation(2)	SNMP trap information	-	Neither writable nor readable

sdsSystemInformation

OID	Name	Description	Remarks
1	sdsProductName	Product name	Fixed to: Hitachi Virtual Storage Software block
2	sdsSoftwareVersion	Operating software version	Storage software version
3	sdsExMibVersion	Version of the extended MIB supported by the operating software	Version of the extended MIB of the storage node This version is different from the version of storage software.
4	sdsStorageNodeIdTable	List of IDs of the installed storage nodes	Array of storage node IDs
5	sdsStorageNodeInformationTable	Hardware information of the installed storage nodes	See the following table.

sdsStorageNodeInformationTable

OID	Name	Description	Setting
1	sdsStorageNodeNickName	Node name of the storage node	Node name of the storage node
2	sdsStorageNodeServerModel	Server model of the storage node	Server model of the storage node
3	sdsStorageNodeServerSerialNumber	Serial number of the storage node	Serial number of the storage node

sdsTrapInformation

OID	Name	Description	Setting
1	sdsInternalID	Device manufacturer's serial number	System ID
2	sdsEventLog	Event log	See the following table.

sdsEventLog

OID	Name	Description	Setting
1	sdsDateAndTime	Date and time when the event was detected	YYYY/MM/DD hh:mm:ss.msec
2	sdsSeverity	Severity of the event	"Critical", "Error", or "Warning"
3	sdsMessageId	Message ID	Message ID of the event
4	sdsEventName	Event name	Unique event name
5	sdsEventMessage	Detailed message of event content	Event description
6	sdsCategory	Failure category	Event category
7	sdsID	Individual ID of the event	Event ID
8	sdsNodeLocation	ID of the storage node which registered the event log	ID of the storage node

Identifying faulty hardware after receiving a trap

The procedure for identifying the hardware in which a failure occurred after receiving a trap is as follows.

Procedure

1. From the information received via a trap, verify the ID of the storage node which registered the event log.
2. Obtain a list of storage node IDs by running an SNMP command.
3. In the list of storage node IDs, find a storage node ID that matches the ID verified in step 1.

4. Obtain information about the storage node by using the entry number (X) of the matched storage node.

   Example:

   OID of the node name of the storage node: 1.3.6.1.4.1.116.3.11.8.1.5.X.1

   OID of the server model of the storage node: 1.3.6.1.4.1.116.3.11.8.1.5.X.2

   OID of the server serial number of the storage node: 1.3.6.1.4.1.116.3.11.8.1.5.X.3

5. From the results in step 4, you can identify the server in which a failure occurred.

Chapter 6: Managing users

Overview of user management

To operate and set up the Virtual Storage Software block storage cluster, you must register with Virtual Storage Software block as a user.

For example, in REST APIs, you must specify your user ID and password ("<*user-ID*>:<*password*>") in the Authorization header for the request header encoded in Base64.

In the case of CLI, specify the user Id in the --user option and enter your password interactively.

(Bare metal) In the console interface, enter your user ID and password interactively.

Built-in user groups and built-in user

Six user groups (built-in user groups), such as SecurityAdministrators and ServiceAdministrators, are pre-registered in Virtual Storage Software block. admin is registered as the built-in user.

Item	Description
User ID	admin
User object ID	admin
Password	hsds-admin You will be requested to change the password when you first log in to the system.
User group ID to which the user belongs	SystemAdministrators
Object ID of the user group to which the user belongs	SystemAdministrators

The following figure illustrates the initial states of users and user groups. More than one role defining an operation privilege (Storage, Service, Security, etc.) is set for each user group.

A built-in user group cannot be deleted. No set roles can be changed. Also, a built-in user cannot be deleted.

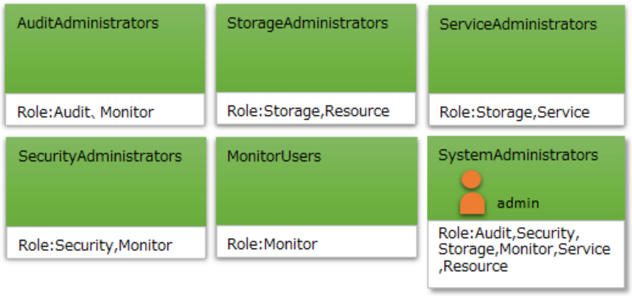

(Bare metal) Built-in users can use the console interface.

States after setup is completed

After setup is completed, a user required for system operation is created each for SecurityAdministrators and ServiceAdministrators groups. For the user ID, user object ID, and password of each created user, ask the administrator who created these users. Then, change the password as required.

Also, the admin user is disabled after setup is completed. You can enable it if required.

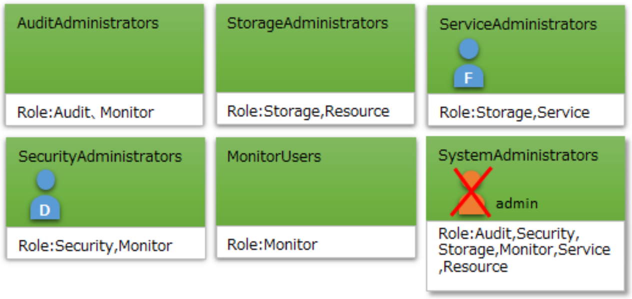

Creating users and user groups

A user's operation privilege is determined by the roles set for the user group to which the user belongs. For example, only a user who belongs to a user group having the Security role can create users.

Also, when the multi-tenancy function is used, it is possible to limit the targets that a user can operate by a scope that is set to the user group to which the user belongs. For details about a scope, see *Configuring multi-tenancy* in this manual.

The only operation a user who is created can perform initially is to change the password. After changing the password, the user can perform any operations allowed for the given role. Note that the password of a user who is created during setup has been changed.

A user can be registered for more than one user group.

You can create new user groups.

Roles and available operations

The following table lists the roles and available operations. Create users according to the system operation guidelines.

When multi-tenancy configuration is used, roles for VPS administrators (VpsSecurity, VpsStorage, or VpsMonitor) are provided in addition to the roles shown in the table further on in this section. Users that have only those roles cannot log in to the GUI. For details about multi-tenancy configuration, see *Configuring multi-tenancy* in this manual.

Caution:

Be careful not to lose the password of the valid user having the Security role. If the passwords of all valid users having the Security role are lost, account management and other operations required for system operation cannot be performed.

Role	Available operations
Security	Audit log file creation and downloading, CHAP authentication, user management, acquisition of a list of session information, event log setting, audit log setting, external authentication server linking, whitelist setting, server certificate importing, setting of a message to be displayed in the GUI login window and in CLI basic authentication, and authentication setting for the compute port for the target operation
Storage	License management, volume management, snapshot management, compute node management, compute node initiator information and path information registration and deletion, volume and compute node connection and disconnection, compute port setting, storage pool capacity expansion, storage node capacity management, drive management, performance and capacity information acquisition, drive data relocation suspension and resumption, etc.
Monitor	Performance and capacity information acquisition, storage node capacity management information acquisition, license information acquisition, etc.
Audit	Audit log file creation and downloading (To log in to the GUI, roles other than Audit and Resource are required.)
Service	Storage node management (maintenance, addition, removal, etc.), storage cluster stop, storage software update, etc.
Resource	VPS creation, edition, and deletion For users with the Resource role to use the GUI, allocate the Monitor role. This allows for login to the GUI and also allows for reference within the range that is allowed for the Resource role.

Role	Available operations
	(To log in to the GUI, roles other than Audit and Resource are required.)

- No role-based execution restriction is applied to the following operations:
  - Verifying, creating, and deleting your own session
  - Obtaining a message to be displayed in the GUI login window and in the warning banner during CLI basic authentication
  - Obtaining versions of APIs
  - Obtaining information about individual jobs
  - Obtaining information about storage cluster master (primary)
  - Obtaining information about control ports and internode ports
  - Network settings for the storage cluster
  - Storage cluster time settings
  - Obtaining your own user information
  - Changing your own password
- A user who has the Security, Storage, Monitor, Service, or Resource role can perform the following operations:
  - Obtaining storage pool information
  - Obtaining information about drives
  - Obtaining license information

- A user who has the Security, Storage, Monitor, Audit, Service, or Resource role can perform the following operations:
  - Obtaining the health status
  - Obtaining information about protection domains
  - Obtaining information about storage clusters
  - Obtaining information about storage nodes
  - Obtaining information about fault domains
- A user who has the Security, Storage, Monitor, Service, Resource, VpsSecurity, VpsStorage, or VpsMonitor role can perform the following operations:
  - Obtaining volume information
  - Obtaining S-VOL and P-VOL information
  - Obtaining compute node information
  - Obtaining compute node initiator information
  - Obtaining compute node path information
  - Obtaining volumes and compute node connection information
  - Obtaining compute port information
  - Obtaining storage node network settings

Basic authentication, session authentication, and ticket authentication

To perform a storage cluster operation through a REST API, for example, send an authentication request to Virtual Storage Software block with credentials specified in the Authorization header for the request header.

Virtual Storage Software block supports three authentication methods: basic authentication, session authentication, and ticket authentication.

In basic authentication, a user ID and a password are used as credentials. In basic authentication, authentication is performed for each request.

In session authentication, a token is used as credentials, and authentication can be omitted for a period of time. Therefore, session authentication is useful in application-based automatic operations. A token is obtained by running a REST API or CLI for generating a session. For how to generate a token, see *Generating a session*.

Ticket authentication is an alternative method used when basic authentication and session authentication cannot be performed due to storage node stoppage or a failure.

When performing ticket authentication, specify the user name and password when issuing the ticket at the same time. See *Security* of *Hitachi Virtual Storage Software Block REST API Reference* and *Master command options* of *Hitachi Virtual Storage Software Block CLI Reference* for how to specify the ticket, user name and password.

See *Managing authentication tickets* for the method of issuing and destroying discarding authentication tickets.

Note:

In case authentication is not possible with an authentication ticket, verify the following:

- Whether the user name and password is the same as the one used when the authentication ticket is issued.
- Whether the authentication ticket has expired.
- Whether the user who issued the authentication ticket had a role required for each operation to be performed by using the authentication ticket.
- Whether the authentication ticket is not the one issued by another storage cluster.
- Whether the authentication ticket is not discarded after being issued.

  If the authentication ticket has been discarded, issue an authentication ticket again. If an authentication ticket cannot be re-issued (due to storage cluster stoppage or other reason), contact customer support.

Upper limit of generated sessions and deleting them

The number of sessions that can be generated is limited. If this limit is exceeded, commands become unavailable. When the multi-tenancy function is not used, the upper limit of the number of sessions that can be generated is 64. When the multi-tenancy function is used, the upper limit of the number of sessions that can be generated is 64 for users that do not belong to a VPS, and 436 for users that belong to a VPS.

If this upper limit is already reached when requesting authentication, 503 Service Unavailable is returned.

A session is deleted by the user (manually), or when the token expires or the session has timed out (automatically). Also, the session of a user is deleted when the user is edited or deleted, the password is changed, the user is deleted from the user group, or the group to which the user belongs is edited.

When a session is deleted, the user must create a session.

User authentication settings and system requirements

The settings which are applied to user authentication are called user authentication settings. User authentication settings include password complexity, password expiration time, lockout, and session parameters. The setting values can be changed. See *Editing user authentication settings*.

The following table lists the system requirements.

Item	Requirement	Remarks
Maximum number of users	32	Including the built-in user and the users on the external authentication server.
Maximum number of user groups	32	Including the built-in groups and the groups on the external authorization server.
Maximum number of user groups to which a user can belong	8	
Number of characters and character types available for a user ID	▪ Number of characters: 6 to 255[1] ▪ Usable character types: Numbers (0 to 9), upper-case alphabet (A to Z), lower-case alphabet (a to z), symbols (! # $ & % ' - . @ ^ _ ` { } ~) (Bare metal) The following restrictions apply to users who are permitted to use the console interface: ▪ Number of characters: 6 to 28 ▪ Usable character types: Numbers (0 to 9), upper-case alphabet (A to Z), lower-case alphabet (a to z), symbols (- . _) ▪ Characters that can be used as the first character: upper-case alphabet (A to Z), lower-case alphabet (a to z), underscores (_)	
Number of characters and character types available for a password	Number of characters: 1 to 256	The minimum password length can be set in the user authentication setting.

Item	Requirement	Remarks
	Usable character types: Numbers (0 to 9), upper-case alphabet (A to Z), lower-case alphabet (a to z), symbols (! " # $ % & ' ( ) * + , - . / : ; < = > ? @ [¥] ^ _ ` { | } ~)	The user authentication default setting is 8.
Number of characters and character types available for a user group ID	Number of characters: 1 to 64[2] Usable character types: Numbers (0 to 9), upper-case alphabet (A to Z), lower-case alphabet (a to z), symbols (! # $ & % ' - . @ ^ _ ` { } ~)	

1. The maximum number of characters for a user name on the external authentication server that can be linked with Virtual Storage Software block is 64.

2. The maximum number of characters for a user group name on the external authentication server that can be linked with Virtual Storage Software block is 64.

About console interface users (Bare metal)

If you use a REST API or CLI to perform the following operations regarding the users who are allowed to use the console interface, the information is applied by the internal processing that runs in one-minute cycle, and event log KARS20067-I is output. Therefore, it takes a certain amount of time until the entire information is applied.

- Creating a user
- Editing the user information
- Deleting a user
- Changing your own password
- Adding a user to a user group
- Removing a user from a user group
- Editing the user group information

Note:

No event log will be output when you perform the following operations on users who are authorized to use the console interface:

- You edited information of a user, but the user is still disabled (isEnabled is "false").
- You edited information of a user and enabled the user, but requiresInitialPasswordReset in the user authentication settings is true.
- You added or removed a user to or from a user group, but the users role is not changed.
- You edited information of a user group, but the roles of the users belonging to the user group are not changed.

Using an external authentication server

When an external authentication server is linked, authentication can be performed by using the user information registered in the external authentication server. Only an OpenLDAP or Active Directory (AD) external authentication server can be linked.

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

A list of the following information about the registered users can be obtained.

- userId: User ID
- userObjectId: User object ID
- passwordExpirationTime: Expiration time of the password
- isEnabled: Whether the user is valid
- userGroups: List of user group IDs (user group ID and object ID of each user group)
- isBuiltIn: Whether the user is a built-in user
- authentication: Authentication type
- roleNames: Role of the user group
- isEnabledConsoleLogin:

  (Virtual machine) null

  (Bare metal) Whether the console interface can be used
- vpsId: ID of the VPS to which users belong
- privileges: List of the VPS information that users can access

Note:

If external authentication is enabled and "mappingMode" is set to "Group", information about the users on the external authentication server is not obtained and only information about the users whose authentication is "local" is obtained.

Before you begin

Required role: Security

Procedure

1. Run either of the following commands to obtain a list of user information:

   REST API: GET /v1/objects/users

   CLI: user_list

Obtaining detailed information about users (CLI or REST API)

The following information about the registered users can be obtained.

If external authentication is enabled and "mappingMode" is set to "Group", information about the user IDs on the external authentication server is not obtained.

- userId: User ID
- userObjectId: User object ID
- passwordExpirationTime: Expiration time of the password
- isEnabled: Whether the user is valid
- userGroups: List of IDs of user groups that the user belongs to (user group IDs and the object ID of each user group)
- isBuiltIn: Whether the user is a built-in user
- authentication: Authentication type
- roleNames: Role of the user group
- isEnabledConsoleLogin:

  (Virtual machine) null

  (Bare metal) Whether the console interface can be used
- vpsId: ID of the VPS to which the user belongs
- privileges: List of the VPS information that the user can access

Before you begin

- Role required to obtain detailed information about a user: Security

  Role-based execution is not subject to restriction for obtaining the user information about yourself.

Procedure

1. Verify the user ID.

   REST API: GET /v1/objects/users

   CLI: user_list

2. Obtain detailed information about users.

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

   REST API: GET /v1/objects/users/<*userId*>

   CLI: user_show

Creating a user (CLI or REST API)

Create a user as follows. When the multi-tenancy function is not used, you can register up to 32 users including built-in users. When the multi-tenancy function is used, the maximum number of users (including built-in users) that do not belong to the VPS is 32, and the maximum number of users that belong to the VPS is 256.

For a user who is authenticated externally, create the user with the same user name as the one on the external authentication server.

The only operation a user who is created can perform initially is to change the password. After changing the password, the user can run any commands allowed for the given role.

Before you begin

- Required role: Security

- When creating users in a VPS: Scope of the VPS

Procedure

1. Verify the user group ID to which the user is to belong.

   REST API: GET /v1/objects/user-groups

   CLI: user_group_list

2. Create a user.

   Run either of the following commands with the user ID, password, and user group ID (more than one ID (up to 8) can be specified), and authentication type specified.

   (Bare metal) To permit a user to use the console interface, specify true for isEnabledConsoleLogin.

Note:

- (Bare metal) If requiresInitialPasswordReset is set to "true" in the user authentication settings, the console interface is not available to the users created by this function even by granting them permission to use the console interface. Performing the procedure described in *Changing your own password* allows these users to use the console interface.
- (Bare metal) To permit a user to use the console interface, the following settings are required:
  - The user must belong to a user group to which the Security or Service role is assigned.
  - The authentication type (authentication) must be set to "local".

To authenticate the user with the external authentication server, specify "external" as the authentication type. (You do not have to specify a password.)

REST API: POST /v1/objects/users

CLI: user_create

After running the command, you receive a response indicating user information.

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.

Editing user information (CLI or REST API)

Edit the following registered information about a user as follows. You can also edit the user information of the built-in user. However, editing user information or disabling a user in such a way that results in there being no user in a user group to which the Security or Service role is assigned cannot be performed.

- (Bare metal) Any operation that will cause the number of users that are permitted to be used in the console interface to become 0 is not allowed.
- (Bare metal) If requiresInitialPasswordReset is set to "true" in the user authentication settings, the console interface is not available even by using this function to change the password. These users become able to use the console interface by performing the procedure described in *Changing your own password*.

If you edit user information, the session of the user is deleted.

- Password
- Enabling or disabling a user

> **Note:**
>
> When the password is changed for the function, if requiresInitialPasswordReset of the user authentication settings is true, the password of the user expires. The expired password can be recovered by using the password changing API of the local user.

Before you begin

- Required role: Security
- When creating users in a VPS: Scope of the VPS

Procedure

1. Verify the user ID.

   REST API: GET /v1/objects/users

   CLI: user_list

2. Edit information about the user.

   Run either of the following commands with the user ID, new password, and whether the user is enabled specified.

   REST API: PATCH /v1/objects/users/<*userId*>

   CLI: user_set

   After running the command, you receive a response indicating user information.

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.

Deleting a user (CLI or REST API)

A user cannot be deleted if no other user remains in the user group to which the Security or Service role is assigned. Also, a built-in user cannot be deleted.

(Bare metal) You cannot delete a user if all remaining users become unable to use the console interface as a result of deletion.

If you delete a user, the session of the user is deleted.

Before you begin

- Required role: Security
- When deleting users in a VPS: Scope of the VPS

Procedure

1. Verify the user ID.

   REST API: GET /v1/objects/users

   CLI: user_list

2. Delete a user.

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

   REST API: DELETE /v1/objects/users/<*userId*>

   CLI: user_delete

3. Obtain a list of user information and verify that the intended user is deleted.

   REST API: GET /v1/objects/users

   CLI: user_list

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.

Changing your own password (CLI or REST API)

Note that only a user for which authentication is set to local can change their password.

If you change your password, your session is deleted.

Role-based execution is not subject to restriction.

Procedure

1. Change your own password.

   Run either of the following commands with the user ID, current password, and new password specified.

   REST API: PATCH /v1/objects/users/<*userId*>/password

   CLI: user_password_set

   After running the command, you receive a response indicating user information.

2. 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.

   > ⚠ **Caution:**
   >
   > When changing the password of the user issuing the authentication ticket, issue an authentication ticket again after the change.

Note:

- (Bare metal) If requiresInitialPasswordReset is set to "true" in the user authentication settings, the console interface is not available to any users for which passwords were set by using the procedures described in *Creating a user* and *Editing user information*. The console interface becomes available by using this function.
- (Bare metal) Password changes are applied to the console interface by internal processing that is performed at 1-minute intervals and message KARS20067-I is output. Therefore, it takes some time for this operation to complete.

Adding a user to a user group (CLI or REST API)

Add a user to another user group as follows.

If you add a user to a user group, the session of the user is deleted.

Before you begin

- Required role: Security
- When adding users in a VPS to user groups: Scope of the VPS

Procedure

1. Verify the user group ID to which the user is added.

   REST API: GET /v1/objects/user-groups

   CLI: user_group_list
2. Run either of the following commands with the user ID and user group ID specified.

   You can specify more than one user group ID.

   REST API: POST /v1/objects/users/<*userId*>/actions/add-user-group/invoke

   CLI: user_add_user_group

   After running the command, you receive a response indicating user information.
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.

Example

User registration example

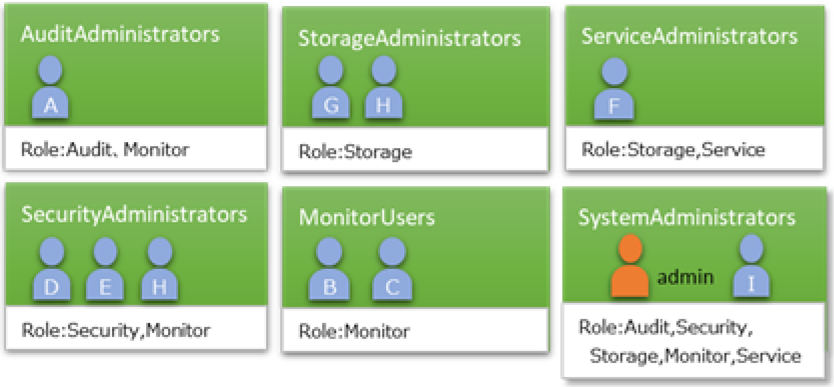

AuditAdministrators, StorageAdministrators, etc. are user group names.The operation privileges of users (A to I) are regulated by the role of the user group to which they belong.

Deleting a user from a user group (CLI or REST API)

If you delete a user from a user group, the session of the user is deleted.

(Bare metal) You cannot delete a user if all users remaining in the user group after the deletion have neither the Security role nor the Service role.

Before you begin

- Required role: Security
- When deleting users in a VPS from user groups: Scope of the VPS

Procedure

1. **Verify the user group ID from which the user is to be deleted.**

   REST API: GET /v1/objects/user-groups

   CLI: user_group_list

2. **Run either of the following commands with the user ID and user group ID specified.**

   You can specify more than one user group ID.

   REST API: POST /v1/objects/users/<*userId*>/actions/delete-user-group/invoke

   CLI: user_delete_user_group

   After running the command, you receive a response indicating user information.

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 a list of user groups (CLI or REST API)

Obtain a list of registered user groups as follows. The following information can be obtained.

- userGroupId: User group ID
- userGroupObjectId: User group object ID
- roleNames: Roles of user groups
- isBuiltIn: Whether the user is a built-in user
- externalGroupName: Name of the group registered with an external authorization server when the external authorization server is linked
- vpsId: ID of the VPS to which user groups belong
- scope: ID of the VPS that user groups can access

Before you begin

Required role: Security

Procedure

1. Obtain a list of user groups.

   REST API: GET /v1/objects/user-groups

   CLI: user_group_list

Obtaining detailed information about user groups (CLI or REST API)

The following information about a registered user group can be obtained.

- memberUsers: List of member users (user IDs and user object IDs)
- userGroupId: User group ID
- userGroupObjectId: User group object ID
- roleNames: Roles of user groups

- isBuiltIn: Whether the group is a built-in user group
- externalGroupName: Name of the group registered with an external authorization server when the external authorization server is linked
- vpsId: ID of the VPS to which the user group belongs
- scope: ID of the VPS that the user group can access

Before you begin

Required role: Security

Procedure

1. Verify the user group ID.

   REST API: GET /v1/objects/user-groups

   CLI: user_group_list

2. Obtain detailed information about the user group.

   Run either of the following commands with the user group ID specified.

   REST API: GET /v1/objects/user-groups/<*userGroupId*>

   CLI: user_group_show

Creating a user group (CLI or REST API)

When the multi-tenancy configuration is not used, the maximum number of user groups that can be registered is 32 including built-in user groups. When the multi-tenancy configuration is used, the maximum number of user groups that do not belong to a VPS is 32 including built-in users and the maximum number of user groups that belong to a VPS is 256.

When an external authorization server is linked, register the name of the group registered with the external authentication server. The type and number of characters used for a group name registered with the external authentication server shall follow the specifications for the external authentication server.

Before you begin

- Required role: Security
- When creating user groups in a VPS: Scope of the VPS

Procedure

1. Create a user group.

   Run either of the following commands with the user group ID, role (more than one role can be specified), and the name of the group registered with the external authentication server (when linked) specified.

   When you use the multi-tenancy function, also specify the VPS to which the user group belongs (for a system administrator, "system") and the scope.

   REST API: POST /v1/objects/user-groups

   CLI: user_group_create

   After running the command, you receive a response indicating user group information.

2. 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.

Editing user group information (CLI or REST API)

Edit the information (role) about a user group as follows. However, the role of the built-in user group cannot be changed.

Editing user group information is not possible if doing so would result in there being no user to which the Security Role or Service role is assigned.

(Bare metal) You cannot edit the information about a user group to which a user who can use the console interface belongs if the user has neither the Security role nor the Service role after editing the information.

If you edit user group information, the sessions of the users belonging to the user group are deleted. Also, if external authentication is enabled and users are managed on each group basis on the external authentication server, editing a local user group corresponding to the external group deletes all sessions.

Before you begin

- Required role: Security
- When editing information about user groups in a VPS: Scope of the VPS

Procedure

1. Verify the user group ID.

   REST API: GET /v1/objects/user-groups

   CLI: user_group_list

2. **2.** Edit information about the user group.

   Run either of the following commands with ID and role(s) of the user group specified.

   When the multi-tenancy function is used, also specify the scope. When you specify the scope, include “system” as well.

   REST API: PATCH /v1/objects/user-groups/<*userGroupId*>

   CLI: user_group_set

   After running the command, you receive a response indicating user group information.

3. **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.

Deleting a user group (CLI or REST API)

You can delete a user group which includes no user.

However, a built-in user group cannot be deleted.

If external authentication is enabled and users are managed on each group basis on the external authentication server, deleting a local user group corresponding to the external group deletes all sessions.

Before you begin

- Required role: Security
- When deleting user groups in a VPS: Scope of the VPS

Procedure

1. Verify the user group ID of the user group to be deleted.

   REST API: GET /v1/objects/user-groups

   CLI: user_group_list

2. Obtain detailed information about the user group.

   Run either of the following commands with the ID of the user group specified.

   REST API: GET /v1/objects/user-groups/<*userGroupId*>

   CLI: user_group_show

   Execute the following steps after verifying that there is no user in the group.

3. Delete the user group.

   Run either of the following commands with the ID of the user group specified.

   REST API: DELETE /v1/objects/user-groups/<*userGroupId*>

   CLI: user_group_delete

4. Obtain a list of user groups and verify that the intended user group is deleted.

   REST API: GET /v1/objects/user-groups

   CLI: user_group_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.

Using an external authentication server (CLI or REST API)

When an external authentication server is linked, authentication can be performed by using the user information registered in the external authentication server. Only an OpenLDAP or Active Directory (AD) external authentication server can be linked.

The following versions of external authentication servers are supported:

- Openldap: 2.4
- Windows Server (Active Directory): 2012 R2, 2016, 2019

TLS protocols available:

- LDAPS
- StartTLS

To use a TLS protocol, you must set isStartTlsEnabled, primaryLdapServerUrl, and secondaryLdapServerUrl in the external authentication server settings according to the following table:

TLS protocol to be used	External authentication server settings	
	isStartTlsEnabled	primaryLdapServerUrl secondaryLdapServerUrl
LDAPS	false	ldaps://example1.com
StartTLS	true	ldap://example2.com

Tip:

The following shows how you can specify the external authentication server settings in the cases of using the REST API and the CLI:

REST API: PATCH /v1/objects/external-auth-server-setting

CLI: external_auth_server_setting_set

Caution:

- To use a Windows Server machine as the external authentication server, enable the TLS protocol, and then import the root certificate of the external authentication server. In Windows Server, a setting requiring an LDAP server signature has been enabled by default since January 2020. Therefore, disabling the TLS protocol might result in unsuccessful connection.

- In an environment with multiple AD/LDAP servers specified, if the AD/LDAP server on the primary node goes down or becomes unreachable when STARTTLS is enabled, external authentication becomes unavailable. This is because STARTTLS does not support a redundant configuration of the AD/LDAP server. Note that external authentication is also unavailable if an AD/LDAP server is specified on only the secondary node.

  In the case that the external AD/LDAP server goes down if STARTTLS is enabled, log in to the system as a local user and change the designation of the external AD/LDAP server by editing the external authentication server settings.

- If an error occurs in a primary external authentication server in a redundant configuration with LDAPS, responses might be delayed up to about 10 minutes.

⚠ **Caution:**

To ensure the security of the connection between the storage system and the external authentication server, configure the external authentication server as follows.

- LDAP clients that the storage system uses to connect to external authentication servers allow unsafe renegotiating for backward compatibility. When connecting TLS to an external authentication server, we recommend that you use an external authentication server that is compatible with RFC 5746.
- The external authentication server is TLS1.2 supported.

  The LDAP client that is used by the storage system for connection with the external authentication server uses only TLS1.2. However, you should disable vulnerable protocol versions (SSL2.0, SSL3.0, TLS1.0, and TLS1.1) in the external authentication server settings.
- At least one of the following is supported by the external authentication server as a TLS cipher suite.
  - (1) TLS_DHE_RSA_WITH_AES_256_GCM_SHA384
  - (2) TLS_DHE_RSA_WITH_AES_128_GCM_SHA256
  - (3) TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
  - (4) TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
  - (5) TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
  - (6) TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256

  - If the cipher suite requirements are not met, TLS connection to the external authentication server cannot be established.

  - For whether the external authentication server supports any of the preceding cipher suites, see the documentation for the external authentication server to be used.

  - You might need to specify the settings for key exchange (DHE or ECDHE) on the external authentication server according to the key exchange method (DHE or ECDHE) of the cipher suite to be used. For details about the settings to be specified, see the documentation for the external authentication server to be used. Note that the OpenLDAP version that supports the ECDHE key exchange method is 2.4.48 or later.

  - The cipher suites that can be used differ depending on whether the server certificate set on the external authentication server is an RSA or ECC certificate. For an RSA certificate, the first to fourth cipher suites in the above can be used. For an ECC certificate, the fifth and sixth cipher suites in the above can be used. For details about how to set a server certificate on the external authentication server, see the documentation for the external authentication server to be used.

- The LDAP client that is used by the storage system for connection with the external authentication server uses only the preceding cipher suites. Disable the other vulnerable cipher suites in the settings on the external authentication server. Vulnerable cipher suites refer to the cipher suites shown in Appendix A. TLS 1.2 Cipher Suite Black List of RFC7540.

- If multiple server certificates are to be set on the external authentication server, they must be issued so that their authenticity can be proved by only one root certificate imported into the storage system.

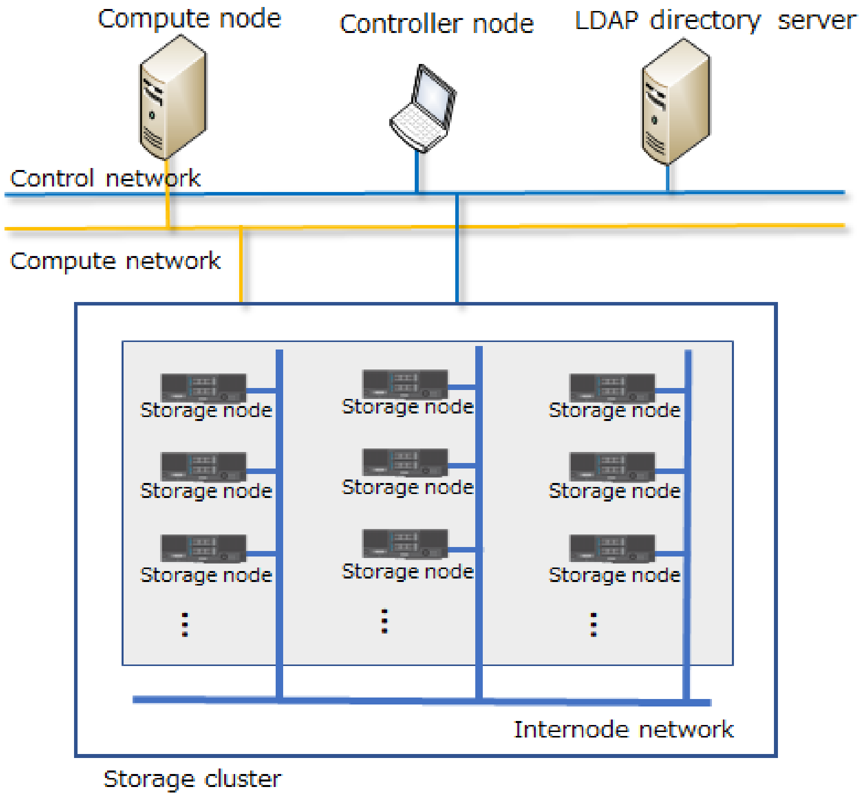

Two use cases are possible:

When users are not managed on a group basis on the external authentication server:

A user is authenticated by using the user group information registered in Virtual Storage Software block. The external authentication server is used as follows in Virtual Storage Software block. After setting a link with the external authentication server, create a user locally in Virtual Storage Software block by using the same user name (user ID) registered in the external authentication server. Hereinafter, these users might be called external users.

When users are managed on a group basis on the external authentication server:

This use case is supported when the external authentication server is Active Directory. A user is authenticated by using the user group information registered in the external authentication server. It is assumed that the external user belongs to a group of external authentication servers. In Virtual Storage Software block, set the link with the external authentication server, create a user group locally in Virtual Storage Software block, and then set the name of the external group in a parameter. In this case, you do not have to create a user individually. Hereinafter, these user groups might be called external groups.

Each time an external authentication server is used, an inquiry is made to the external authentication server (in basic authentication) to verify the latest information. In session authentication, no inquiry is made to the external authentication server, but authentication is performed just by verifying if the token is valid.

Note:

Precautions when using an external authentication server:

- The maximum number of characters that can be used for a user name and user group name on the external authentication server that can be linked with Virtual Storage Software block is 64.

- It might take approximately five minutes until a password change made on the external authentication server is reflected to Virtual Storage Software block. The time required to do this depends on the settings of the external authentication server.

- Regardless of the above use case or the setting value for mappingMode, when you specify settings on the external authentication server, you need the following six attributes of the target users:

  userIdAttribute, userTreeDn, userObjectClass, externalGroupNameAttribute, userGroupTreeDn, userGroupObjectClass

  For details about the setting values, see the *Hitachi Virtual Storage Software Block REST API Reference* or the *Hitachi Virtual Storage Software Block CLI Reference*.

Before you begin

- Required external service (if the TLS protocol is used): DNS server

  If the TLS protocol is used, the AD/LDAP server must be specified by using an FQDN. In Virtual Storage Software block, register a DNS server for which an AD/LDAP server is registered for the Address record.

  > **Caution:**
  >
  > When using a DNS server, a storage node caches DNS inquiry results for the time (DNS TTL) set in the DNS server. For this reason, if the content registered in the DNS server (correspondence between the host name and IP address) is changed, the storage node might access an old address during DNS TTL. If you have changed the content registered in the DNS server (correspondence between the host name and IP address), wait until the time specified for DNS TTL has passed, and then specify the external authentication server settings.

- Required role: Security

Procedure

1. Edit the settings of the external authentication server.

   You can perform this for the cluster master node (primary) only.

   For the communication parameters for connection to the external authentication server, timeoutSeconds (CLI: --timeout_seconds), retryInternalMilliseconds (CLI: --retry_interval_milliseconds), and maxRetries (CLI: --max_retries), it is recommended that you use the system defaults without changing them.

   The system defaults are as follows: timeoutSeconds (CLI: --timeout_seconds) is -1, retryIntervalMilliseconds (CLI: --retry_interval_milliseconds) is 100, and maxRetries (CLI: --max_retries) is 3.

   REST API: PATCH /v1/objects/external-auth-server-setting

   CLI: external_auth_server_setting_set

   After running the command, you receive a response indicating the setting information of the external authentication server.

2. Prepare a root certificate that proves the authenticity of the server certificates set on the external authentication server, and then import the root certificate into Virtual Storage Software block.

   You can perform this for the cluster master node (primary) only.

   Run either of the following commands with the external authentication server and the root certificate file specified.

   REST API: POST /v1/objects/external-auth-server-root-certificates/<*targetServer*>/actions/import/invoke

   CLI: external_auth_server_root_certificate_import

   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. **Verify the connection with the external authentication server.**

   You can perform this for the cluster master node (primary) only.

   REST API: POST /v1/objects/external-auth-server-setting/actions/verify-connectivity/invoke

   CLI: external_auth_server_setting_verify_connectivity

   After running the command, you will receive a response indicating the number of externally authenticable users and the number of externally authenticable user groups that could be searched on the LDAP server.

   If the connection with the external authentication server fails, the command response shows the message and error code for indicating the cause of the failure for each external authentication server set in step 1. Eliminate the cause of the failure based on the shown information and recheck the connection with the external authentication server.

5. **Register an external user or an external group.**

   To register an external user, create a user.

   REST API: POST /v1/objects/users

   CLI: user_create

   In this case, specify parameters as follows:

   - userId: User ID on the external authentication server
   - userGroupIds: ID of one or more user groups to which the user is to belong
   - authentication: "external"

   To register an external group, create a user group. In this case, specify parameters as follows:

   - userGroupId: ID of the user group
   - roleNames: Role or roles assigned to the user group
   - externalGroupName: Name of the group registered in the external authentication server

6. **Obtain a list of users or user groups to make sure that step 5 is carried out correctly.**
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 the settings of the external authentication server (CLI or REST API)

Obtain the settings of the external authentication server as follows. The following information can be obtained.

- isEnabled: Whether external authentication is enabled
- authProtocol: Authentication protocol used for external authentication
- ldapSetting: Setting information about LDAP authentication

Before you begin

Required role: Security

Procedure

1. Obtain the settings of the external authentication server.

   You can perform this for the cluster master node (primary) only.

   REST API: GET /v1/objects/external-auth-server-setting

   CLI: external_auth_server_setting_show

Obtaining a root certificate from the external authentication server (CLI or REST API)

The root certificate that was imported in *Using an external authentication server* is output as a file in DER format and is used by the storage system to verify the authenticity of the server certificates set on the external authentication server.

Before you begin

Required role: Security

Procedure

1. Obtain the root certificate that is used to verify the authenticity of the server certificates set on the external authentication server.

   You can perform this for the cluster master node (primary) only.

   Run either of the following commands with the external authentication server specified.

   REST API: GET /v1/objects/external-auth-server-root-certificates/<*targetServer*>/download

   CLI: external_auth_server_root_certificate_download

Chapter 7: Editing user authentication settings

Editing user authentication settings (CLI or REST API)

The following table lists the user authentication settings that you can perform.

Item	Description	System default	Allowed value range
minLength	Minimum password length.	8	1-256
minNumberOfUpperCaseChars	Minimum number of uppercase alphabetical characters contained in a password.	0	0-256
minNumberOfLowerCaseChars	Minimum number of lowercase alphabetical characters contained in a password.	1	0-256
minNumberOfNumerals	Minimum number of numerals (0 to 9) contained in a password.	1	0-256
minNumberOfSymbols	Minimum number of symbols (excluding alphanumeric characters) contained in a password.	0	0-256

Item	Description	System default	Allowed value range
numberOfPassword History	Number of generations from generation 1 (when the password was changed) for which use of a previous password is prohibited. For example, if this is 2, the previous password cannot be set. 1 means that this limit is disabled (the user can set a previously used password).	1	1-10
requiresInitialPasswordReset	Whether change of the initial password is requested when a new user runs a REST API or CLI or logs in to the GUI for the first time. If true, a new user is forced to change the default password before the initial operation.	true	Boolean
minAgeDays	Number of days after the password is changed until it can be changed. 0 means that the user can change the password immediately. This should be less than maxAgeDays.	0	0-10

Item	Description	System default	Allowed value range
maxAgeDays	Number of days you can use a password after it has been changed. The password is invalid if the specified number of days has passed. 0 means that this limit is disabled.	42	0-365
maxAttempts	Number of consecutive login failures until the account is disabled temporarily (account lock). 0 means that the account is not locked. (Bare metal) This setting is also applied at login to the console interface.	3	0-10
lockoutSeconds	Duration (in seconds) until the account is unlocked. (Bare metal) This setting is also applied at login to the console interface.	60	60-600
maxLifetimeSeconds	Token lifetime (in seconds).	86400	1800-604800
maxIdleSeconds	Time until a session times out (in seconds).	1800	300-86400

Item	Description	System default	Allowed value range
	When access is made during the session before the timeout, the timeout count starts from this time. If you specify the time until a session times out (aliveTime) at the time of session creation, the setting takes priority over maxIdleSeconds.		

Before you begin

Required role: Security

Procedure

1. Edit the user authentication settings.

   You can perform this for the cluster master node (primary) only.

   REST API: PATCH /v1/objects/user-auth-setting

   CLI: user_auth_setting_set

   After running the command, you receive a response indicating user authentication setting information.

> **Note:**
>
> (Bare metal)
>
> - The user authentication settings you edited are applied to the console interface by the internal processing that runs in one-minute cycle, and event log KARS20068-I is output. Therefore, it takes a certain amount of time until the entire setting is applied.
> - No event log is output when you perform the following operations:
>   - You edited the user authentication settings. However, there is no change in lockoutSetting.
>   - When requiresInitialPasswordReset is true in the user authentication settings, you disable a user who was enabled and for which the password has not yet been changed.
>   - When requiresInitialPasswordReset is true in the user authentication settings, you change the password of a user for which the password has not yet been changed.

2. 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 user authentication settings (CLI or REST API)

Role-based execution is not subject to restriction.

- passwordComplexitySetting: Specifies password complexity.
- passwordAgeSetting: Specifies the password expiration time.
- lockoutSetting: Specifies the lockout setting.
- sessionSetting: Specifies the session.

Procedure

1. Obtain the user authentication settings as follows.

   You can perform this for the cluster master node (primary) only.

   REST API: GET /v1/objects/user-auth-setting

   CLI: user_auth_setting_show

Chapter 8: Managing sessions

Overview of session management

A session is connection information between a user and a storage system.

A session can be generated by executing a REST API or CLI for generating a session. You can obtain a token by generating a session.

The number of sessions that can be generated in the entire storage system is limited. If this limit is exceeded, new sessions can no longer be generated.

When the multi-tenancy function is not used, the upper limit of the number of sessions that can be generated is 64. When the multi-tenancy function is used, the upper limit for the number of sessions that do not belong to a VPS is 64, and the number of sessions that belong to a VPS is 436.

After the upper limit of the number of sessions that can be generated is reached, you can use Basic authentication and ticket authentication. When the upper limit is reached, delete unnecessary sessions or consider using Basic authentication or ticket authentication.

A generated session is deleted when:

- The session (token) expired.
- The session has timed out without being used.
- The cluster master node (primary) has failed over.
- The session is deleted, the user is edited or deleted, the password is changed, the user is deleted from a user group, or the user group to which the user belongs is edited.

Token expiration time (default: 24 hours) applies to the entire storage system. For details about how to change the token expiration time, see *Editing user authentication settings*.

You can specify the session timeout time (default: 30 minutes) not only for each session but for the entire storage system. To specify the session timeout time for each session, specify the time when sessions are created. If you omit this specification, the session timeout time will be applied to the entire storage system. For details about how to change the session timeout time of the entire storage system, see *Editing user authentication settings*.

Caution:

If the token obtained by generating a session has leaked out, the Virtual Storage Software block storage system might be operated by an unintended third party. Properly manage the obtained token so that it won't be leaked out.

Obtaining a list of sessions (CLI or REST API)

Obtain a list of sessions that the user generated as follows. The following information can be obtained.

- sessionId: Session ID (uuid)
- userId: User ID
- userObjectId: User object ID
- expirationTime: Session expiration time
- createdTime: Date and time when the session is generated
- lastAccessTime: Date and time when the session was used last
- roleNames: List of roles assigned to the user who retains this session
- vpsId: ID of the VPS to which the user who generated sessions belongs
- privileges: List of the VPS information that the user who generated sessions can access

Before you begin

Required role: Security

Procedure

1. Obtain a list of sessions that the user generated.

   REST API: GET /v1/objects/sessions

   CLI: session_list

Obtaining session information (CLI or REST API)

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

- sessionId: Session ID (uuid)
- userId: User ID
- userObjectId: User object ID
- expirationTime: Session expiration time
- createdTime: Date and time when the session is generated
- lastAccessTime: Date and time when the session was used last

- roleNames: List of roles assigned to the user who retains this session
- vpsId: ID of the VPS to which the user who generated the session belongs
- privileges: List of the VPS information that the user who generated the session can access

Before you begin

Required role: Security

Procedure

1. **Verify the ID of the session.**

   REST API: GET /v1/objects/sessions

   CLI: session_list
2. Obtain information about the session with the session ID specified.

   REST API: GET /v1/objects/sessions/<*sessionId*>

   CLI: session_show

Generating a session (CLI or REST API)

A token for session authentication is generated and the following information about the session is displayed. Session authentication is possible with the generated token.

Role-based execution is not subject to restriction.

- token: Token
- sessionId: Session ID
- userId: User ID
- userObjectId: User object ID
- expirationTime: Session expiration time
- createdTime: Date and time when the session is generated
- lastAccessTime: Date and time when the session was used last
- roleNames: List of roles assigned to the user who retains this session
- vpsId: ID of the VPS to which the user who generated the session belongs
- privileges: List of the VPS information that the user who generated the session can access

Before you begin

- When adding sessions in a VPS: Scope of the VPS

Procedure

1. Generate a session.

   You can specify the time until a session times out.

   REST API: POST /v1/objects/sessions

   CLI: session_create

   After running the command, you receive a response indicating session information.

Deleting a session (CLI or REST API)

Delete a generated session as follows.

Before you begin

- Role required to specify all session IDs: Security

  A user who does not have the Security role can specify only the session IDs of the user.

- When deleting sessions in a VPS: Scope of the VPS

Procedure

1. Verify the ID of the session.

   REST API: GET /v1/objects/sessions

   CLI: session_list

2. Delete sessions with the session IDs specified.

   REST API: DELETE /v1/objects/sessions/<*sessionId*>

   CLI: session_delete

3. Obtain a list of sessions and verify that the session is deleted.

   REST API: GET /v1/objects/sessions

   CLI: session_list

Chapter 9: Managing authentication tickets

Issuing an authentication ticket (CLI or REST API)

Issue an authentication ticket. Authentication tickets are required to change and set configuration information and create dump files.

Before you begin

- Role required to issue the authentication ticket for changing and setting configuration information: Security

- Role required to issue the authentication ticket for generating a dump log file: Service

Procedure

1. Issue an authentication ticket.

   Run either of the following commands with expiration time of the ticket specified.

   REST API: POST /v1/objects/tickets

   CLI: ticket_create

   After running the command, you will receive a response indicating the authentication ticket and expiration time.

2. Create an authentication ticket file from the issued authentication ticket.

   In the response after executing the command in step 1, the validity period and authentication ticket are displayed as follows: <authentication ticket> the string of the part of the file to a file.

   Example of output in the REST API:

   ```
   {
       "ticket":"<authentication-ticket>"
       "expirationTime": "<expiration-time>",
   }
   ```

   Example of output in the CLI when the format option is not specified:

   ```
   Ticket: <authentication-ticket>
   ExpirationTime: <expiration-time>
   ```

Note:

- Note that the <*authentication ticket*> does not include line breaks.
- (Virtual machine) Use the nano editor to create an authentication ticket file on the maintenance node.

```
$ nano
```

3. (Virtual machine) The authentication ticket file that you created should be stored outside the maintenance node for the home directory and backup of the logged-in user of the maintenance node. Be careful not to leak the authentication ticket file to the outside. If you want to upload an authentication ticket created outside the maintenance node to the maintenance node, use SFTP to transfer it to the home directory of the logged-in user of the maintenance node.

```
$ sftp -P 10022 <user-name-of-the-mnservice-role>@<IP-address-of-the-maintenance-
node>
sftp> cd <user-name-of-the-mnservice-role>
sftp> put <authentication-ticket-file-name>
```

Make sure that authentication ticket files are not used for unauthorized purposes.

Discarding authentication tickets (CLI or REST API)

Discard all the authentication tickets that were previously issued.

Perform discarding of authentication tickets in the following cases:

- An authentication ticket is lost.
- A security incident such as theft has occurred.
- Customer support issued an authentication ticket.
- The configuration information was restored.

If discarding of the authentication tickets cannot be reflected on some of the storage nodes (due to a failure or other reason), an event log indicating such storage nodes is notified.

Before you begin

Required role: Security

Procedure

1. Discard the authentication tickets.

   REST API: POST /v1/objects/tickets/actions/revoke-all/invoke

   CLI: ticket_revoke_all

   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

   After running the command, if you receive a response indicating "Succeeded" as the state, the job is completed. Go to step 5.

   If the job is unsuccessful, go to steps 3 and 4.

3. If discarding of the authentication ticket cannot be successfully performed on some of the storage nodes, it is notified in an event log. Verify the event log and identify the unsuccessful storage node.
4. Verify the status of the identified storage node. If it is blocked, perform maintenance recovery or maintenance replacement. If the storage node is not blocked, perform maintenance blocking, and then perform maintenance recovery or maintenance replacement.
5. After the authentication tickets are successfully discarded, notify all the users having the Service role and Security role of discarding of the authentication tickets and request them to re-create an authentication ticket.

Chapter 10: Obtaining network setting information

Obtaining storage cluster network settings (CLI or REST API)

Obtain the following storage cluster network settings as follows. If a failure occurs in the storage node assigned with the representative IP address of the storage cluster, the IP address is passed to a different normal storage node.

Role-based execution is not subject to restriction.

- primaryDnsServerIpAddress: IP address of the DNS server for the first name resolution request destination
- secondaryDnsServerIpAddress: IP address of the DNS server for the second name resolution request destination
- virtualIpv4Address: Representative IP address of the storage cluster (IPv4)

Procedure

1. **Obtain the storage cluster network settings.**

   REST API: GET /v1/objects/storage-network-setting

   CLI: storage_network_setting_show

Obtaining a list of information about control ports (CLI or REST API)

The following information can be obtained.

Role-based execution is not subject to restriction.

- id: Control port IDs (uuid)
- storageNodeId: IDs (uuid) of the storage nodes with control ports
- macAddress: MAC addresses
- mtuSize: MTU size
- interfaceName: Interface names

- deviceName: Device name of the NIC
- configuredPortSpeed: Link speed setting of the physical port used for communication
- portSpeedDuplex:
  - (Virtual machine) Always “DependsOnHypervisor”
  - (Bare metal) Actual link speed and duplex settings of the physical port used for communication
- isTeamingEnabled:
  - (Virtual machine) Always “DependsOnHypervisor”
  - (Bare metal) Whether teaming is enabled/disabled
- ipv4Information: IPv4 settings of the control port
- teaming:
  - (Virtual machine) Always null
  - (Bare metal) Teaming information of the control port
- redundancy: Redundancy of the physical port
- status: Status of the control port
  - (Virtual machine) Always “Normal”
  - (Bare metal) "Normal", "Warning", or "Error"
- statusSummary: Status summary of the control port
  - (Virtual machine) Always “Normal”
  - (Bare metal) "Normal", "Warning", or "Error"

Procedure

1. Obtain a list of information about control ports.

   REST API: GET /v1/objects/control-ports

   CLI: control_port_list

Obtaining information about individual control ports (CLI or REST API)

The following information can be obtained.

Role-based execution is not subject to restriction.

- id: ID (uuid) of the intended control port
- storageNodeId: ID (uuid) of the storage node with the intended control port
- macAddress: MAC address
- mtuSize: MTU size
- interfaceName: Interface name
- deviceName: Device name of the NIC
- configuredPortSpeed: Link speed setting of the physical port used for communication
- portSpeedDuplex:
  - (Virtual machine) Always “DependsOnHypervisor”
  - (Bare metal) Actual link speed and duplex settings of the physical port used for communication
- isTeamingEnabled:
  - (Virtual machine) Always “DependsOnHypervisor”
  - (Bare metal) Whether teaming is enabled/disabled
- ipv4Information: IPv4 settings of the control port
- teaming:
  - (Virtual machine) Always null
  - (Bare metal) Whether teaming is enabled/disabled
- redundancy: Redundancy of the control port
- status: Status of the control port
  - (Virtual machine) Always “Normal”
  - (Bare metal) "Normal", "Warning", or "Error"
- statusSummary: Status summary of the control port
  - (Virtual machine) Always “Normal”
  - (Bare metal) "Normal", "Warning", or "Error"

Procedure

1. Verify the ID of the intended control port.

   REST API: GET /v1/objects/control-ports

   CLI: control_port_list

2. Obtain information about the intended control port.

   Run either of the following commands with the control port ID specified.

   REST API: GET /v1/objects/control-ports/<*id*>

   CLI: control_port_show

Obtaining a list of information about internode ports (CLI or REST API)

The following information can be obtained.

Role-based execution is not subject to restriction.

- id: Internode port IDs (uuid)
- storageNodeId: IDs (uuid) of the storage nodes with internode ports
- macAddress: MAC addresses
- mtuSize: MTU size
- interfaceName: Interface names
- deviceName: Device name of the NIC
- configuredPortSpeed: Link speed setting of the physical port used for communication
- portSpeedDuplex:
  - (Virtual machine) Always “DependsOnHypervisor”
  - (Bare metal) Actual link speed and duplex settings of the physical port used for communication
- isTeamingEnabled:
  - (Virtual machine) Always “DependsOnHypervisor”
  - (Bare metal) Whether teaming is enabled/disabled
- ipv4Information: IPv4 settings of the internode port

- teaming:
  - (Virtual machine) Always null
  - (Bare metal) Teaming information of the internode port
- redundancy: Redundancy of the physical port
- status: Status of the internode port
  - (Virtual machine) Always “Normal”
  - (Bare metal) "Normal", "Warning", or "Error"
- statusSummary: Status summary of the internode port
  - (Virtual machine) Always “Normal”
  - (Bare metal) "Normal", "Warning", or "Error"

Procedure

1. Obtain a list of information about internode ports.

   REST API: GET /v1/objects/internode-ports

   CLI: internode_port_list

Obtaining information about individual internode ports (CLI or REST API)

The following information can be obtained.

Role-based execution is not subject to restriction.

- id: ID (uuid) of the intended internode port
- storageNodeId: ID (uuid) of the storage node with the intended internode port
- macAddress: MAC address
- mtuSize: MTU size
- interfaceName: Interface name
- deviceName: Device name of the NIC
- configuredPortSpeed: Link speed setting of the physical port used for communication

- portSpeedDuplex:
  - (Virtual machine) Always “DependsOnHypervisor”
  - (Bare metal) Actual link speed and duplex settings of the physical port used for communication
- isTeamingEnabled:
  - (Virtual machine) Always “DependsOnHypervisor”
  - (Bare metal) Whether teaming is enabled/disabled
- ipv4Information: IPv4 settings of the internode port
- teaming:
  - (Virtual machine) Always null
  - (Bare metal) Teaming information for the internode port
- redundancy: Redundancy of the physical port
- status: Status of the internode port
  - (Virtual machine) Always “Normal”
  - (Bare metal) "Normal", "Warning", or "Error"
- statusSummary: Status summary of the internode port
  - (Virtual machine) Always “Normal”
  - (Bare metal) "Normal", "Warning", or "Error"

Procedure

1. Verify the ID of the intended internode port.

   REST API: GET /v1/objects/internode-ports

   CLI: internode_port_list

2. Obtain information about the intended internode port.

   Run either of the following commands with the internode port ID specified.

   REST API: GET /v1/objects/internode-ports/<*id*>

   CLI: internode_port_show

Obtaining a list of network settings of storage nodes (CLI or REST API)

The following information can be obtained.

- id: Storage node IDs (uuid)
- ipv4Route: Routing table (IPv4)

Before you begin

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

Procedure

1. Obtain a list of network settings of storage nodes.

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

   CLI: storage_node_network_setting_list

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

The following information can be obtained.

- id: ID (uuid) of the intended storage node
- ipv4Route: Routing table (IPv4)

Before you begin

Required role: Security, Storage, Monitor, Service, 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 network settings.

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

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

   CLI: storage_node_network_setting_show

Chapter 11: Obtaining common information

Obtaining versions and names of APIs (CLI or REST API)

The following information can be obtained.

Role-based execution is not subject to restriction.

- apiVersion: Versions of APIs
- productName: Names of APIs

Procedure

1. Obtain information about APIs such as version information.

   REST API: GET /configuration/version

   CLI: version_show

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

The following information can be obtained.

Role-based execution is not subject to restriction.

- jobId: IDs (uuid) of jobs
- self: URLs for accessing job information
- userId: IDs of the users who issued APIs that triggered creation of jobs
- status: Progress of jobs
- state: Status of jobs
- createdTime: Time when jobs were created
- updatedTime: Time when job status was updated
- completedTime: Time when jobs were completed
- request: Request information

- affectedResources: URLs for accessing the resources which are the operation targets of jobs
- error: Job error information

Procedure

1. Obtain a list of jobs.

   REST API: GET /v1/objects/jobs

   CLI: job_list

   > **Note:**
   >
   > - If you can't find job information, filter by the query parameter when you obtain a list of job information, or specify the job ID to get job information individually. For information about the query parameter, see *Hitachi Virtual Storage Software Block REST API Reference*.
   > - Up to 100,000 jobs can be retained in a storage system. Job information that exceeds the maximum number of retentions is deleted in the oldest order. However, because it also includes jobs that run automatically inside the storage system, the total number of job information that can be obtained is less than this number.

Obtaining information about individual jobs (CLI or REST API)

The following information can be obtained.

Role-based execution is not subject to restriction.

- jobId: ID (uuid) of the intended job
- self: URL for accessing the information about the intended job
- userId: ID of the user who issued an API that triggered creation of the intended job
- status: Progress of the intended job
- state: Status of the intended job
- createdTime: Time when the intended job was created
- updatedTime: Time when the job status was updated
- completedTime: Time when the job was completed
- request: Request information

- affectedResources: URL for accessing the resource which is the operation target of the intended job
- error: Job error information

Procedure

1. Obtain the ID of the intended job.

   REST API: GET /v1/objects/jobs

   CLI: job_list

2. Obtain information about the intended job.

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

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

   CLI: job_show

   Note:

   Up to 100,000 jobs can be retained in a storage system. Job information that exceeds the maximum number of retentions is deleted in the oldest order. However, because it also includes jobs that run automatically inside the storage system, the total number of job information that can be obtained is less than this number.

Verifying the model (CLI or REST API)

The model of the storage software that is used can be verified from the storage cluster information.

Role-based execution is not subject to restriction.

Procedure

1. Obtain information about the storage cluster.

   REST API: GET /v1/objects/storage

   CLI: storage_show

2. Verify the model from the value of "modelName".

Value of "modelName"	Model
VSSB	Virtual machine
VSSB1	Bare metal

Chapter 12: Obtaining system performance and capacity information

Performance and capacity information of the storage system

You can obtain performance and capacity information about each resource that Virtual Storage Software block is collecting.

The performance information the user can obtain is divided into two types: low-temporal-resolution (hereinafter, low-resolution) information that is collected at 1-minute intervals and high-temporal-resolution (hereinafter, high-resolution) information that is collected at 5-second intervals. Low-resolution information refers to the information obtained from archival records. High-resolution information refers to the latest information. (Some information is only available as low-resolution information.)

Low-resolution information is retained for up to two days. Low-resolution information that exceeds the maximum retention days is deleted in the oldest order.

When a problem occurs in the storage system, obtain and utilize high-resolution information about the specific faulty resource for troubleshooting.

Regularly obtain and utilize low-resolution information to verify whether the system is running normally and any abnormality has occurred.

Obtaining a list of low-resolution performance information of control ports (CLI or REST API)

The following information can be obtained.

- id: Control port IDs (uuid)
- receiveTransferRate: Amount of data received per second
- sendTransferRate: Amount of data sent per second

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Obtain a list of performance information about control ports.

   REST API: GET /v1/objects/performances/control-ports

   CLI: control_port_performance_list

Obtaining a list of high-resolution performance information of control ports (CLI or REST API)

The following information can be obtained.

- id: Control port IDs (uuid)
- receiveTransferRate: Amount of data received per second
- sendTransferRate: Amount of data sent per second

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

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

   REST API: GET /v1/objects/control-ports

   CLI: control_port_list

2. Obtain a list of performance information about control ports.

   Use the query parameter to specify the control port IDs and run one of the following commands. You can specify a maximum of 32 control port IDs.

   REST API: GET /v1/objects/detail-performances/control-ports

   CLI: control_port_detail_performance_list

Obtaining performance information about individual control ports (CLI or REST API)

The following information can be obtained.

- id: Control port IDs (uuid)
- receiveTransferRate: Amount of data received per second
- sendTransferRate: Amount of data sent per second

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

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

   REST API: GET /v1/objects/control-ports

   CLI: control_port_list

2. Obtain performance information.

   Run either of the following commands with the control port ID specified.

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

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

   CLI (low-resolution): control_port_performance_show

   CLI (high-resolution): control_port_detail_performance_show

Obtaining a list of low-resolution performance information about drives (CLI or REST API)

The following information can be obtained.

- id: Drive 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
- responseTime: Average response time
- usage: Ratio of I/O operating time of the drive against elapsed time

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Obtain a list of performance information.

   REST API: GET /v1/objects/performances/drives

   CLI: drive_performance_list

Next steps

Note:

When the following operations or failures occur, drive access may occur asynchronously with I/O from the compute node. If the system load is high, the performance may be temporarily affected.

- Expanding storage pool
- Volume operations
- Storage node maintenance
- Adding and removing storage nodes
- Adding and removing drives
- Storage node failure
- Drive failure

Obtaining a list of high-resolution performance information about drives (CLI or REST API)

The following information can be obtained.

- id: Drive 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
- responseTime: Average response time
- usage: Ratio of I/O operating time of the drive against elapsed time

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

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

   REST API: GET /v1/objects/drives

   CLI: drive_list

2. Obtain a list of performance information.

   Use the query parameter to specify the drive IDs and run either of the following commands. You can specify a maximum of 32 drive IDs.

   REST API: GET /v1/objects/detail-performances/drives

   CLI: drive_detail_performance_list

Next steps

> **Note:**
>
> When the following operations or failures occur, drive access may occur asynchronously with I/O from the compute node. If the system load is high, the performance may be temporarily affected.
>
> - Expanding storage pool
> - Volume operations
> - Storage node maintenance
> - Adding and removing storage nodes
> - Adding and removing drives
> - Storage node failure
> - Drive failure

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

The following information can be obtained.

- id: Drive 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
- responseTime: Average response time
- usage: Ratio of I/O operating time of the drive against elapsed time

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

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

   REST API: GET /v1/objects/drives

   CLI: drive_list

2. Obtain performance information.

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

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

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

   CLI (low-resolution): drive_performance_show

   CLI (high-resolution): drive_detail_performance_show

Next steps

Note:

When the following operations or failures occur, drive access may occur asynchronously with I/O from the compute node. If the system load is high, the performance may be temporarily affected.

- Expanding storage pool
- Volume operations
- Storage node maintenance
- Adding and removing storage nodes
- Adding and removing drives
- Storage node failure
- Drive failure

Obtaining a list of low-resolution performance information of internode ports (CLI or REST API)

The following information can be obtained.

- id: Internode port IDs (uuid)
- receiveTransferRate: Amount of data received per second
- sendTransferRate: Amount of data sent per second

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Obtain a list of performance information about internode ports.

   REST API: GET /v1/objects/performances/internode-ports

   CLI: internode_port_performance_list

Obtaining a list of high-resolution performance information of internode ports (CLI or REST API)

The following information can be obtained.

- id: Internode port IDs (uuid)
- receiveTransferRate: Amount of data received per second
- sendTransferRate: Amount of data sent per second

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

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

   REST API: GET /v1/objects/internode-ports

   CLI: internode_port_list

2. **Obtain a list of performance information about internode ports.**

   Use the query parameter to specify the IDs of the internode ports and run either of the following commands. You can specify a maximum of 32 internode port IDs.

   REST API: GET /v1/objects/detail-performances/internode-ports

   CLI: internode_port_detail_performance_list

Obtaining performance information about individual internode ports (CLI or REST API)

The following information can be obtained.

- id: Internode port IDs (uuid)
- receiveTransferRate: Amount of data received per second
- sendTransferRate: Amount of data sent per second

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

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

   REST API: GET /v1/objects/internode-ports

   CLI: internode_port_list

2. Obtain performance information.

   Run either of the following commands with the internode port ID specified.

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

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

   CLI (low-resolution): internode_port_performance_show

   CLI (high-resolution): internode_port_detail_performance_show

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

The following information can be obtained.

- id: Storage pool IDs (uuid)
- usedCapacity: Total consumed amount

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Obtain a list of capacity information about storage pools.

   REST API: GET /v1/objects/performances/pool-capacities

   CLI: pool_capacity_performance_list

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

The following information can be obtained.

- id: Storage pool IDs (uuid)
- usedCapacity: Total consumed amount

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Verify the ID of the storage pool for which you want to obtain capacity information.

   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 capacity information.

   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/performances/pool-capacities/<*id*>

   CLI: pool_capacity_performance_show

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

The following information can be obtained.

- id: Storage pool IDs (uuid)
- volumeReadIOPS: Number of read operations per second for all volumes in each storage pool
- volumeWriteIOPS: Number of write operations per second for all volumes in each storage pool
- volumeReadTransferRate: Read transfer amount per second for all volumes in each storage pool
- volumeWriteTransferRate: Write transfer amount per second for all volumes in each storage pool

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Obtain a list of performance information about storage pools.

   REST API: GET /v1/objects/performances/pools

   CLI: pool_performance_list

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

The following information can be obtained.

- id: Storage pool IDs (uuid)
- volumeReadIOPS: Number of read operations per second for all volumes in the specified storage pool
- volumeWriteIOPS: Number of write operations per second for all volumes in the specified storage pool
- volumeReadTransferRate: Read transfer amount per second for all volumes in the specified storage pool
- volumeWriteTransferRate: Write transfer amount per second for all volumes in the specified storage pool

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

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

   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 performance information.

   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/performances/pools/<*id*>

   CLI: pool_performance_show

Obtaining a list of low-resolution performance information of compute ports (CLI or REST API)

The following information can be obtained.

- id: Compute port IDs (uuid)
- fc:

  (Virtual machine) FC port performance information

  (Bare metal) null
- iscsi: iSCSI port performance information

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Obtain a list of performance information about compute ports.

   REST API: GET /v1/objects/performances/ports

   CLI: port_performance_list

Obtaining a list of high-resolution performance information of compute ports (CLI or REST API)

The following information can be obtained.

- id: Compute port IDs (uuid)
- fc:

  (Virtual machine) FC port performance information

  (Bare metal) null
- iscsi: iSCSI port performance information

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

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

   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 a list of performance information about compute ports.

   Use the query parameter to specify the compute port IDs and run one of the following commands. You can specify a maximum of 32 compute port IDs.

   If you use the CLI, you can specify the WWN or iSCSI name instead of the compute port ID.

   REST API: GET /v1/objects/detail-performances/ports

   CLI: port_detail_performance_list

Obtaining performance information about individual compute ports (CLI or REST API)

The following information can be obtained.

- id: Compute port ID (uuid)
- fc:

  (Virtual machine) FC port performance information

  (Bare metal) null
- iscsi: iSCSI port performance information

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. **Verify the ID of the compute port for which you want to obtain performance information.**

   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 performance information.

   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 (low-resolution): GET /v1/objects/performances/ports/<*id*>

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

   CLI (low-resolution): port_performance_show

   CLI (high-resolution): port_detail_performance_show

Obtaining performance information about the storage cluster (CLI or REST API)

The following information can be obtained.

- id: UUID of the storage cluster
- averageCpuUsage: Average CPU usage for all storage nodes
- averageMemoryUsage: Average memory usage for all storage nodes

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Obtain performance information about the storage cluster.

   REST API: GET /v1/objects/performances/storage

   CLI: storage_performance_show

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

The following information can be obtained.

- id: Storage node IDs (uuid)
- volumeReadIOPS: Number of read operations per second for volumes
- volumeWriteIOPS: Number of write operations per second for volumes
- volumeReadTransferRate: Read transfer amount per second for volumes
- volumeWriteTransferRate: Write transfer amount per second for volumes
- driveReadIOPS: Number of read operations per second for drives
- driveWriteIOPS: Number of write operations per second for drives
- driveReadTransferRate: Read transfer amount per second for drives
- driveWriteTransferRate: Write transfer amount per second for drives
- cpu: A list of CPU performance information of the storage node
- memory: Memory usage

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. Obtain a list of performance information about storage nodes.

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

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

   CLI (low-resolution): storage_node_performance_list

   CLI (high-resolution): storage_node_detail_performance_list

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

The following information can be obtained.

- id: Storage node IDs (uuid)
- volumeReadIOPS: Number of read operations per second for volumes
- volumeWriteIOPS: Number of write operations per second for volumes
- volumeReadTransferRate: Read transfer amount per second for volumes
- volumeWriteTransferRate: Write transfer amount per second for volumes
- driveReadIOPS: Number of read operations per second for drives
- driveWriteIOPS: Number of write operations per second for drives
- driveReadTransferRate: Read transfer amount per second for drives
- driveWriteTransferRate: Write transfer amount per second for drives
- cpu: A list of CPU performance information of the storage node
- memory: Memory usage

Before you begin

Required role: Storage, Monitor, or Resource

Procedure

1. **Verify the ID of the storage node for which you want to obtain performance information.**

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

   CLI: storage_node_list

2. **Obtain performance information.**

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

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

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

   CLI (low-resolution): storage_node_performance_show

   CLI (high-resolution): storage_node_detail_performance_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 13: Operating the console interface (Bare metal)

Overview of console interface (Bare metal)

Note that the console screen shown here appears after a storage cluster setup operation, storage node addition operation, or storage node replacement operation is completed on a storage node.

In the bare metal model of Virtual Storage Software block, the console interface provides the functionality necessary for troubleshooting Virtual Storage Software block in a situation in which the GUI or REST API is unavailable due to a storage node failure after setup or other reasons.

The console interface is operated using the iLO web interface and the iLO remote console. For details about how to connect to the iLO web interface and iLO remote console, see the iLO User Guide provided by the vendor of the physical server to be used as the storage node.

> ⚠ **Caution:**
>
> - Only the remote console (HTML5 Console) provided by iLO is supported for operations with the console interface. Operations with a connection other than a remote console (HTML5 Console) might not work properly.

Login and logout (Bare metal)

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

Before you begin

- Required role: Security or Service

- The user must be permitted to use the console interface (isEnabledConsoleLogin is enabled).

Procedure

1. Log in to the iLO web interface window on the storage node, and then connect to the iLO remote console (HTML5 Console).
2. The **Keyboard layout configuration** screen, which indicates the current keyboard layout, appears.

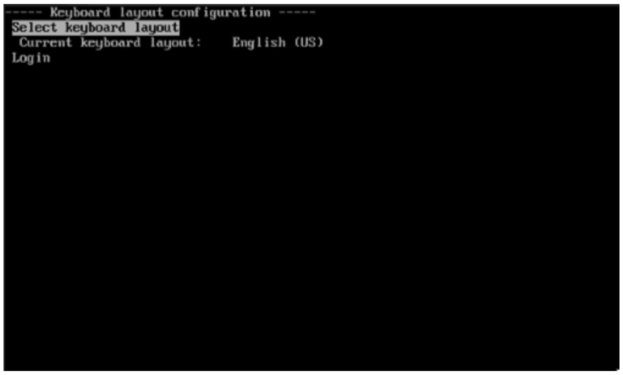

```
----- Keyboard layout configuration -----
Select keyboard layout
 Current keyboard layout:    English (US)
Login
```

If you accept the current setting, you can skip this step. Go to the next step.

To change the setting, select **Select keyboard layout**, and then press the **Enter** key.

The **Select keyboard layout** screen, which lists the supported keyboard layouts, appears. If [...] is displayed on the last line of the displayed keyboard layout list, there are other list items on the next page. You can view them by selecting [...] and pressing the **Enter** key. If [...] is displayed on the first line of the displayed keyboard layout list, there are other list items on the previous page. You can view them by selecting [...] and pressing the **Enter** key. Select the keyboard layout that is appropriate for the keyboard you are using, and then press the **Enter** key.

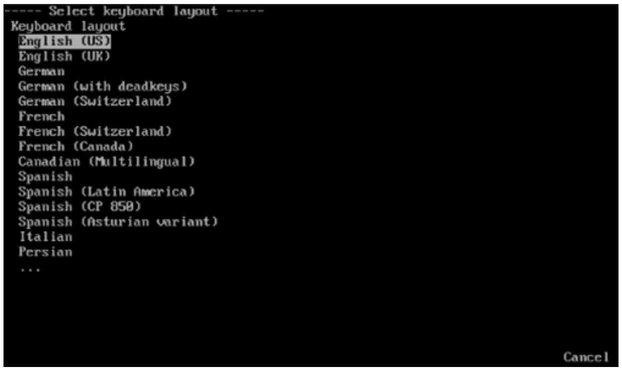

```
----- Select keyboard layout -----
Keyboard layout
 English (US)
 English (UK)
 German
 German (with deadkeys)
 German (Switzerland)
 French
 French (Switzerland)
 French (Canada)
 Canadian (Multilingual)
 Spanish
 Spanish (Latin America)
 Spanish (CP 850)
 Spanish (Asturian variant)
 Italian
 Persian
 ...

                                                                Cancel
```

Note:

- When you select a keyboard layout and press the **Enter** key, the screen color changes to blue for a moment. This phenomenon is not a problem.
- If you do not change the keyboard layout, you can go back to the **Keyboard layout configuration** screen by selecting **Cancel** and pressing the **Enter** key.

The **Select keyboard layout (Keyboard test)** screen appears. In the "Keyboard test" field, you can verify whether the selected keyboard layout works as expected. When verification is complete, select **OK** and press the **Enter** key.

When the **Keyboard layout configuration** screen appears, verify that the selected keyboard layout is indicated as the current keyboard layout.

3. In the **Keyboard layout configuration** screen, use the ↓ key to select **Login**, and then press the **Enter** key.

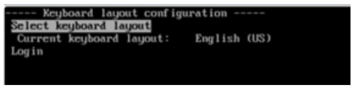

```
----- Keyboard layout configuration -----
Select keyboard layout
 Current keyboard layout:    English (US)
Login
```

4. When the login prompt is displayed, enter the user ID and password to log in.

   If login is successful, the **Storage node console top menu** screen is displayed.

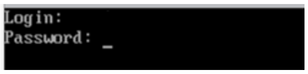

```
Login:
Password: _
```

5. To log out, on the **Storage node console top menu** screen, use the down arrow (↓) key to select **Logout**, and then press the **Enter** key.

   If logout is successful, the **Keyboard layout configuration** menu appears.

> **Caution:**
>
> - Do not forget to log out when you have completed your work in the console interface.
>
>   If you exit the iLO remote console (HTML5 Console) without logging out, you will remain logged in. If you access the iLO remote console (HTML5 Console) again in this state, the logged-in screen might be displayed.
>
> - If you leave the screen while operating the console interface, log out.
>
>   When multiple users make a connection to the iLO remote console (HTML5 Console), iLO displays a dialog notifying them of the simultaneous connections. You can use the dialog to deny connection from a subsequent user. However, if no action is taken, subsequent users will be allowed to connect, and the logged-in screen might become operable by them.
>
> - If no key operation is performed for 30 minutes, the user is automatically logged out.

About the Storage node console top menu screen (Bare metal)

The **Storage node console top menu** screen displays the following information:

- Time of the storage node (local time display)
- Storage node Name: Name of the storage node
- IP address: IP address for the control network of the storage node

In the **Storage node console top menu** screen , you can select from the following menus. Use the up ↑ and down ↓ arrow keys and **Enter** key to select a menu.

Displayed menu	Description	Transition-destination screen
Show storage cluster information	Displays the storage cluster information.	**Show storage cluster information** screen
Logout	Log out of the console interface.	**Keyboard layout configuration** screen

Log out when you have completed your work in the console interface.

About the Show storage cluster information screen (Bare metal)

The **Show storage cluster information** screen displays the information that the user views on demand from customer support.

The **Show storage cluster information** screen displays the following information:

- Time of the storage node (local time display)
- Storage node Name: Name of the storage node
- Storage cluster Id: ID of the storage cluster
- System random Id: Device-specific random ID

On the **Show storage cluster information** screen, the following menu item is available: Use the up ↑ and down ↓ arrow keys and **Enter** key to select a menu.

Displayed menu	Description	Transition-destination screen
Back to top	Return to the **Storage node console top menu** screen.	**Storage node console top menu** screen

After you have confirmed the information displayed on this screen, select **Back to top** to return to the **Storage node console top menu** screen.

Notes on the console interface (Bare metal)

This section describes how you can take action if you become unable to perform operations while using the console interface.

- Simultaneously pressing **Alt** and another key might cause the iLO remote console screen to black out, preventing you from performing operations. In such a case, press **Alt** + **F1** key.
- Pressing the **ScrollLock** key might lock the screen, preventing you from performing key entries. In such a case, press the **ScrollLock** key again to unlock the screen.
- The login prompt might not be displayed correctly. In such a case, press the **Enter** key to cause a login error temporarily, and then redisplay the login prompt.

Also, in input mode, do not enter any characters other than the following in the input field:

- Numbers (0-9)
- Uppercase letters (A-Z) and lowercase letters (a-z)
- Following symbols: ! " # $ % & ' ( ) * + , - . / : ; < = > ? @ [\ ] ^ _ ` { | } ~

Entering other characters or keys might result in the screen not displaying properly. In this case, display another screen and then display the original screen again.

Chapter 14: 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, you must restart the storage system after power supply 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 *Hitachi Virtual Storage Software Block CLI Reference*.	

Points to confirm	Remarks	Done
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.	▪ 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] 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.	
When using REST API/CLI, verify that the parameters and other options are correctly specified.	See the *Hitachi Virtual Storage Software Block REST API Reference* or *Hitachi Virtual Storage Software Block CLI Reference*.	
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.	For more information, see *Establishing an SSL/TLS session for operation and management*.	
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.	For more information, see *Establishing an SSL/TLS session for operation and management*.	

Points to confirm	Remarks	Done
Verify that the external servers accessed by the storage system are configured and operating normally.	▪ 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. ▪ 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 ▪ A user of the storage system might have to be registered with an external server, such as an LDAP server.	
Verify that the control network, compute network, and internode network that constitute Virtual Storage Software block are all running normally.	▪ 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
	▪ 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 (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.		

Collecting logs and making a query

Procedure for collecting Virtual Storage Software block dump log files

You can collect the dump log file of Virtual Storage Software block.

(Bare metal) If a failure occurs when you configure a storage cluster, add storage nodes, or replace storage nodes and you need to collect dump log files, you might need to collect them via the console interface rather than using the procedure described in this section. The following table shows cases requiring collection of dump log files by using the procedure described in this section, and cases requiring collection of dump log files by using the procedure described in *Procedure for collecting dump log files via the console interface (Bare metal)*.

Operation for which a failure occurred	Storage nodes with which a storage cluster was being configured[1]	Storage nodes being added	Storage nodes being replaced	Spare nodes	Other existing storage nodes
Configuring a storage cluster	Procedure in *Procedure for collecting dump log files via the console interface (Bare metal)*	-	-	-	-

Operation for which a failure occurred	Storage nodes with which a storage cluster was being configured[1]	Storage nodes being added	Storage nodes being replaced	Spare nodes	Other existing storage nodes
Adding storage nodes	-	Procedure in *Procedure for collecting dump log files via the console interface (Bare metal)*[2]	-	-	Procedure in this section
Replacing storage nodes	-	-	Procedure in *Procedure for collecting dump log files via the console interface (Bare metal)*[2]	-	Procedure in this section
Other than above	-	-	-	Procedure in *Procedure for collecting dump log files via the console interface (Bare metal)*[3]	Procedure in this section

1. There might be a case where you only have to collect dump log files for only specific storage nodes.

When you collect dump log files as described in *Action to be taken for a failure that occurred during setup of each storage node (Bare metal)*, you only have to collect dump log files for the storage node on which the failure occurred.

2. The procedure in this section applies depending on the progress of the processing for the target storage node. The procedure to be used depends on whether [Login:setup] is displayed on the first line of the console interface login screen.

You can choose the appropriate procedure according to the *Procedure for collecting dump log files via the console interface (Bare metal)*.

3. Perform this procedure only when instructed to do so in the manual or by customer support.

This section describes how to collect dump log files by using the REST API/CLI. For details about how to collect dump log files by using the GUI, see *Operation with dump log files* in the *Hitachi Virtual Storage Software Block GUI Guide*.

Dump log files are created in the following file name:

- When splitting of dump log files is not specified:

  hsds_log_<*YYYYMMDD*>_<*hhmmss*>_<*LABEL*>_<*HOSTNAME*>_<*MODE*>.tar.gz

  Example) hsds_log_20210120_204236_labelname1_storagenode1_Base.tar.gz

- When splitting of dump log files is specified:

  hsds_log_<*YYYYMMDD*>_<*hhmmss*>_<*LABEL*>_<*HOSTNAME*>_<*MODE*>.tar.gz.<*INDEX*>

  Example) hsds_log_20210120_204236_labelname1_storagenode1_Base.tar.gz.001

When executing each REST API/CLI in the procedure, authenticate with the issued authentication ticket and execute it. At the time of authentication, use the user name and password that were specified when the authentication ticket was issued. If you cannot issue an authentication ticket due to a failure of Virtual Storage Software block, contact customer support.

In Virtual Storage Software block, the system might automatically create dump log files in the event of a failure. Verify whether dump log files have been created according to the procedure below, and if they were created automatically, download the dump log files. Automatically-created dump log files include "AutoCollection" for <*MODE*>.

A dump log file created with the REST API/CLI includes a value set in the mode parameter for <*MODE*>. When the mode parameter was omitted, "Base" is set.

A maximum of one generation of automatically-created dump log files and a maximum of one generation of dump log files created with the REST API/CLI are stored in each storage node.

> **Note:**
>
> If you have not yet imported a server certificate into the Virtual Storage Software block, you may receive a warning message about the server certificate.

Procedure

1. **Verify the dump log file creation status.**

   Perform this step for all storage nodes.

   REST API: GET /v1/objects/dump-statuses

   CLI: dump_status_list

   If dump log files whose mode is "AutoCollection" exist in the storage nodes, verify the status of the dump log files.

   When the status is "Creating", wait until the status becomes "Created" or "Failed". " It will take a maximum of 60 minutes until the status becomes "Created" or "Failed".

   After the status of all the dump log files whose mode is "AutoCollection" becomes "Created" or "Failed", verify the size and fileName of only the dump log files whose status is "Created" and the number of files into which such dump log files can be split, and then proceed to step 2.

   The size of the dump log files are displayed (unit: MiB) in "size".

   If there is no storage node in which dump log files whose mode is "AutoCollection" exist, proceed to step 3.

2. Download the dump log files whose mode is "AutoCollection" and status is "Created" to the controller node.

   Perform this step for all storage nodes in which dump log files whose mode is "AutoCollection" and status is "Created" exist.

   Secure free space equal to or more than the size confirmed in step 1 before downloading.

   REST API: GET /v1/objects/dump-files/<*fileName*>/download

   CLI: dump_file_download

   Note:

   - If you specified an index parameter, you can download split files whose file size is 400 Mib. When you want to download the dump log files separately, perform download for the number of files into which dump log files can be split.

   - If a dump log file could not be downloaded using a GUI, collect it using a REST API or CLI.

3. Request the creation of a dump log file from the controller node to the storage node to which the dump log file is created.

   Perform this step for all storage nodes.

   Specify "Base" for the mode parameter.

   You can run the request for creating the dump log files in parallel.

   REST API: POST /v1/objects/dump-file/actions/create-file/invoke

   CLI: dump_file_create_file

Note:

When you specify <LABEL> of the file name in the label parameter, observe the following.

Number of characters: Maximum of 64

Characters that can be used: Numbers (0 to 9), uppercase alphabet (A to Z), lowercase alphabet (a to z), symbols (! # $ % & ' - . @ ^ _ ` { } ~)

Caution:

- When failure information is being generated in the storage node, an error message (KARS10665-E) is returned. When this message appears, wait for a while, and then retry executing the request for creating the dump log files. If the message is consecutively returned for 60 minutes or more, skip step 4 and subsequent steps for the storage node.
- If the information to be collected does not exist on a storage node, the KARS10666-E message is returned. If this message is returned, skip the operations in steps 4, 5, and 7 for the relevant storage node.

4. Check the creation status of the dump log file.

Perform this step for all storage nodes.

It takes a maximum of 60 minutes to create the dump log file.

REST API: GET /v1/objects/dump-statuses

CLI: dump_status_list

When the status of the dump log files for which you requested creation changes to "Created", verify the size, fileName, and the number of files into which the dump log files can be split, and then proceed to the next step.

5. Download the dump log files to the controller node. Secure free space equal to or more than the size confirmed in step 4 before downloading.

Perform this step for all storage nodes.

REST API: GET /v1/objects/dump-files/<*fileName*>/download

CLI: dump_file_download

Note:

- If you specified an index parameter, you can download split files whose file size is 400 Mib. When you want to download files separately, perform download for the number of files into which the dump log files can be split.
- If a dump log file could not be downloaded using a GUI, collect it using a REST API or CLI.

6. Verify that all the dump log files you downloaded in step 5 were created with names conforming to the naming convention.

7. Delete the dump log files from all the storage nodes from which you downloaded them in step 5.

   REST API: DELETE /v1/objects/dump-files/<*fileName*>

   CLI: dump_file_delete

8. If the CLI was also being run, collect the CLI log in the controller node as described in *Procedure for collecting CLI logs*.

   > **Note:**
   >
   > When sending dump log files to customer support, give priority to those files whose <MODE> in the file name is "AutoCollection".

Procedure for collecting CLI logs

If a failure occurs while using the CLI, collect the CLI logs as follows, and then contact customer support.

> **Note:**
>
> (Virtual machine) If a failure occurs when using the CLI on the maintenance node, execute the *Procedure for collecting maintenance node logs (Virtual machine)* instead of this procedure. By collecting the log of the maintenance node, the log of CLI usage on the maintenance node is also collected.

Procedure

1. Verify the "hsds" folder in the user folder of the controller node in which the CLI was executed. The user folder differs depending on the settings.

   The following are examples of the "hsds" folder for each OS. <*User*> indicates a user ID.

   - For Linux:

     Normal user: /home/<*User*>/hsds/cli

     Root user: /root/hsds/cli

   - For Windows:

     C:\Users\<*User*>\hsds\cli

2. Zip the "cli" folder in the "hsds" folder.

   Zip the file according to the environment of the user.

3. With the created zip file, contact customer support.

Procedure for collecting dump log files via the console interface (Bare metal)

This section describes the procedure for collecting dump log files via the console interface.

Caution:

Dump log files generated via the console interface can no longer be downloaded after configuring a storage cluster, after adding or replacing storage nodes in the cluster, or after switching spare nodes. Make sure that you download any necessary log files before performing these operations.

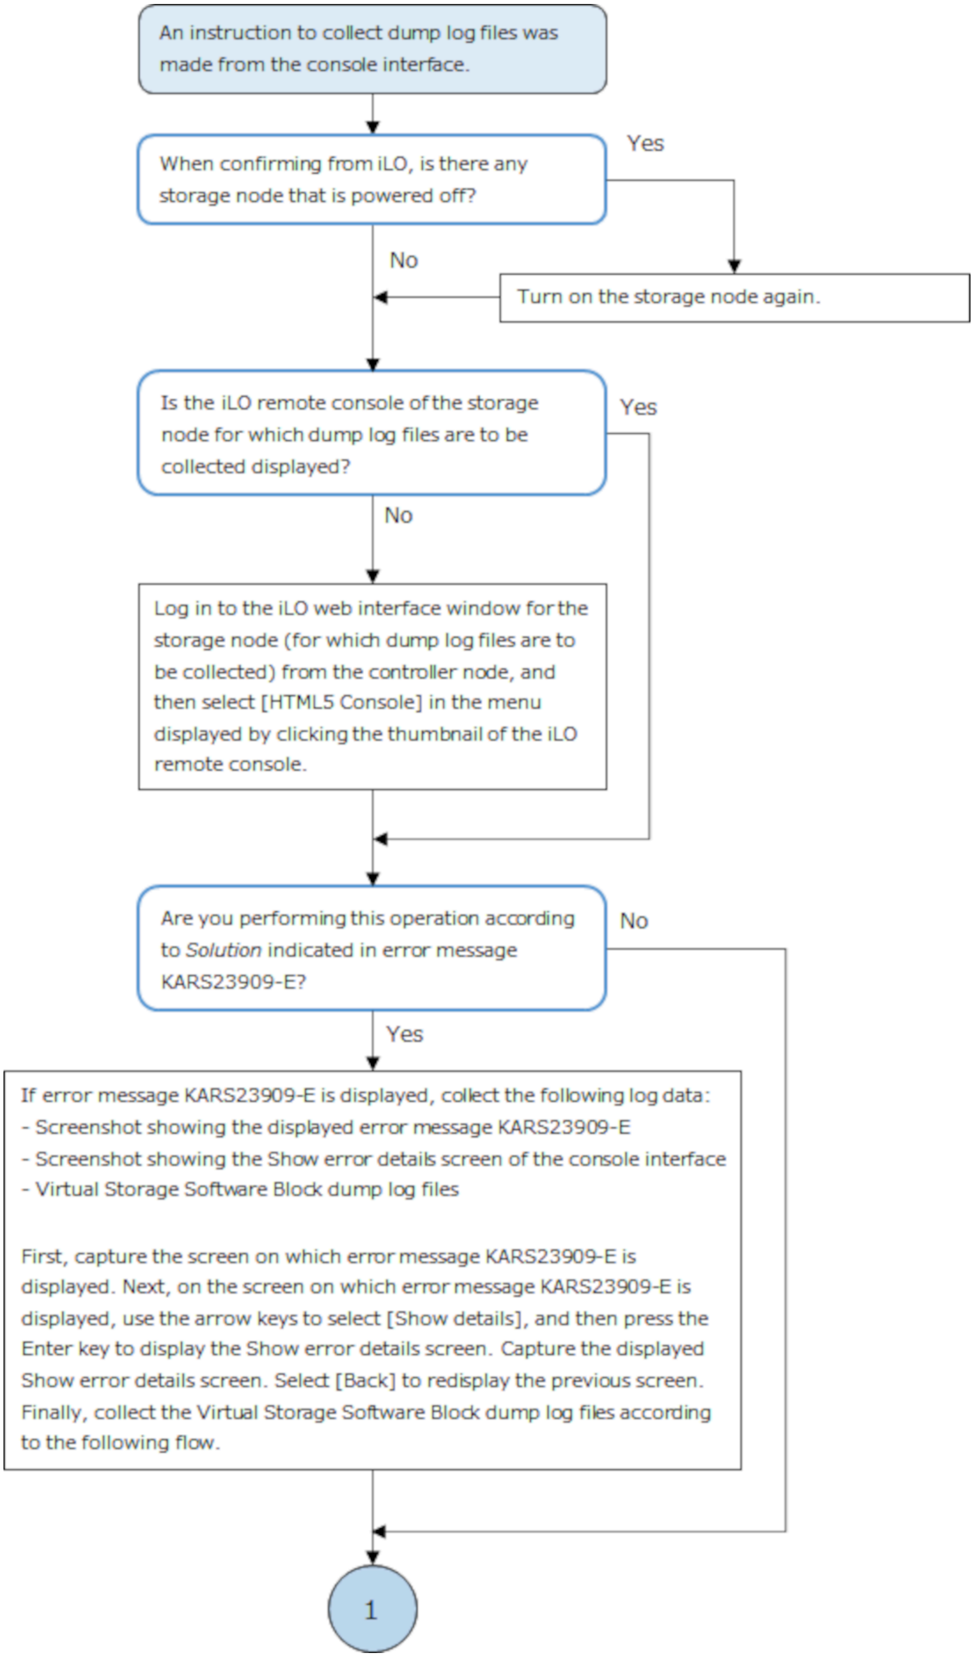
An instruction to collect dump log files was made from the console interface.
When confirming from iLO, is there any storage node that is powered off?
Yes
Turn on the storage node again.
No
Is the iLO remote console of the storage node for which dump log files are to be collected displayed?
Yes
No
Log in to the iLO web interface window for the storage node (for which dump log files are to be collected) from the controller node, and then select [HTML5 Console] in the menu displayed by clicking the thumbnail of the iLO remote console.
Are you performing this operation according to *Solution* indicated in error message KARS23909-E?
No
Yes
If error message KARS23909-E is displayed, collect the following log data:
- Screenshot showing the displayed error message KARS23909-E
- Screenshot showing the Show error details screen of the console interface
- Virtual Storage Software Block dump log files
First, capture the screen on which error message KARS23909-E is displayed. Next, on the screen on which error message KARS23909-E is displayed, use the arrow keys to select [Show details], and then press the Enter key to display the Show error details screen. Capture the displayed Show error details screen. Select [Back] to redisplay the previous screen. Finally, collect the Virtual Storage Software Block dump log files according to the following flow.
1

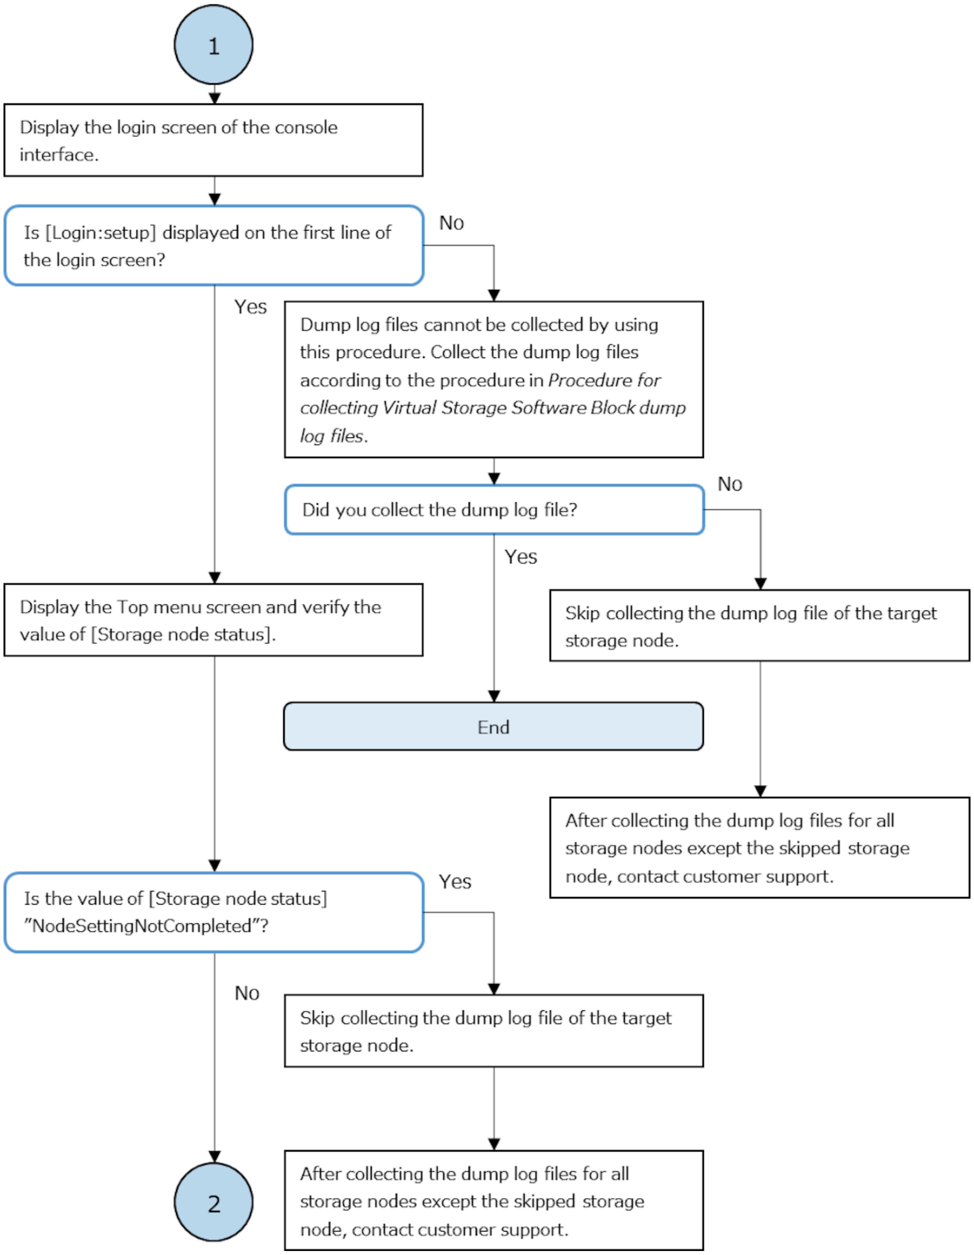
1
Display the login screen of the console interface.
Is [Login:setup] displayed on the first line of the login screen?
No
Yes
Dump log files cannot be collected by using this procedure. Collect the dump log files according to the procedure in *Procedure for collecting Virtual Storage Software Block dump log files*.
Did you collect the dump log file?
No
Yes
Display the Top menu screen and verify the value of [Storage node status].
Skip collecting the dump log file of the target storage node.
End
After collecting the dump log files for all storage nodes except the skipped storage node, contact customer support.
Is the value of [Storage node status] "NodeSettingNotCompleted"?
Yes
No
Skip collecting the dump log file of the target storage node.
2
After collecting the dump log files for all storage nodes except the skipped storage node, contact customer support.

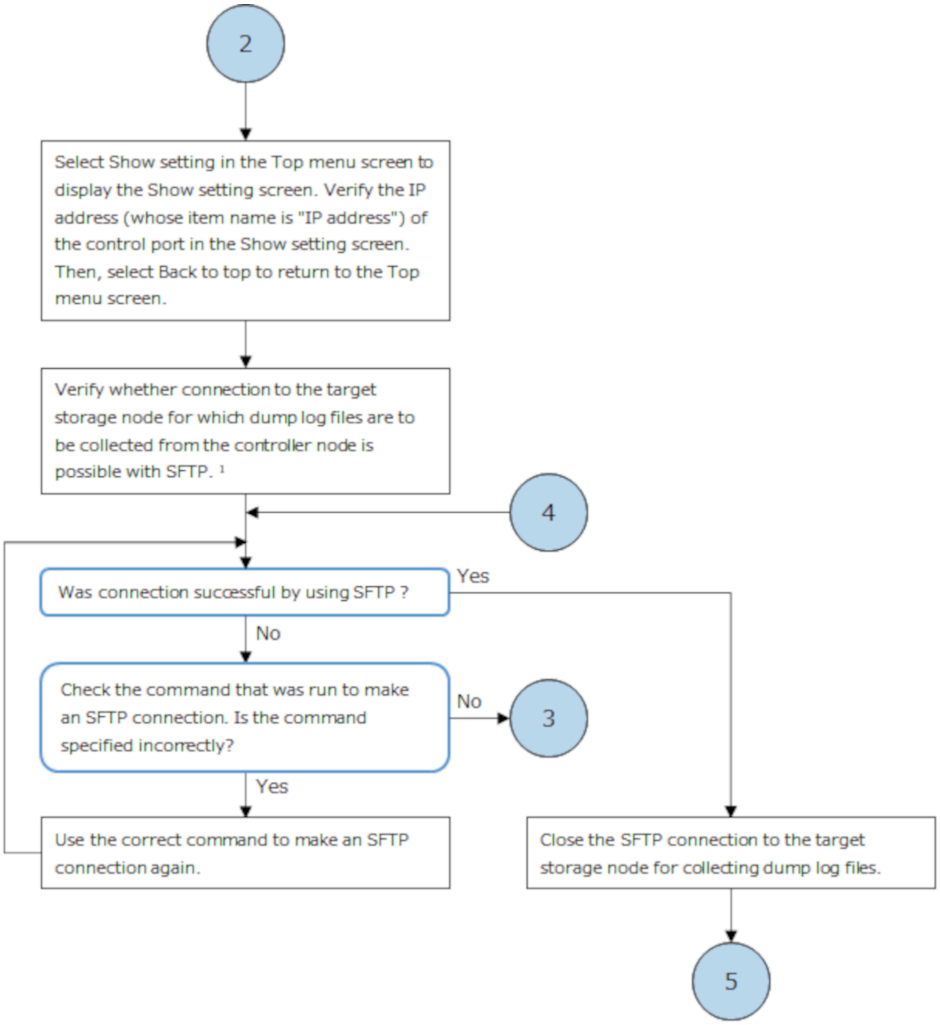
2
Select Show setting in the Top menu screen to display the Show setting screen. Verify the IP address (whose item name is "IP address") of the control port in the Show setting screen. Then, select Back to top to return to the Top menu screen.
Verify whether connection to the target storage node for which dump log files are to be collected from the controller node is possible with SFTP. [1]
4
Was connection successful by using SFTP ?
Yes
No
Check the command that was run to make an SFTP connection. Is the command specified incorrectly?
No
3
Yes
Use the correct command to make an SFTP connection again.
Close the SFTP connection to the target storage node for collecting dump log files.
5

1. See below for the information required to check the SFTP connection to the target storage node for collecting dump log files and a command example for checking the connection:
   User name: setup
   Password: Password of the setup user
   Port number: 10022
   Connection destination IP address: Control port IP address of the target storage node for collecting dump log files
   Command example

```
$ sftp -P 10022 setup@<control-port-IP-address-of-target-storage-node-for-collecting-dump-log-files>
sftp>
```

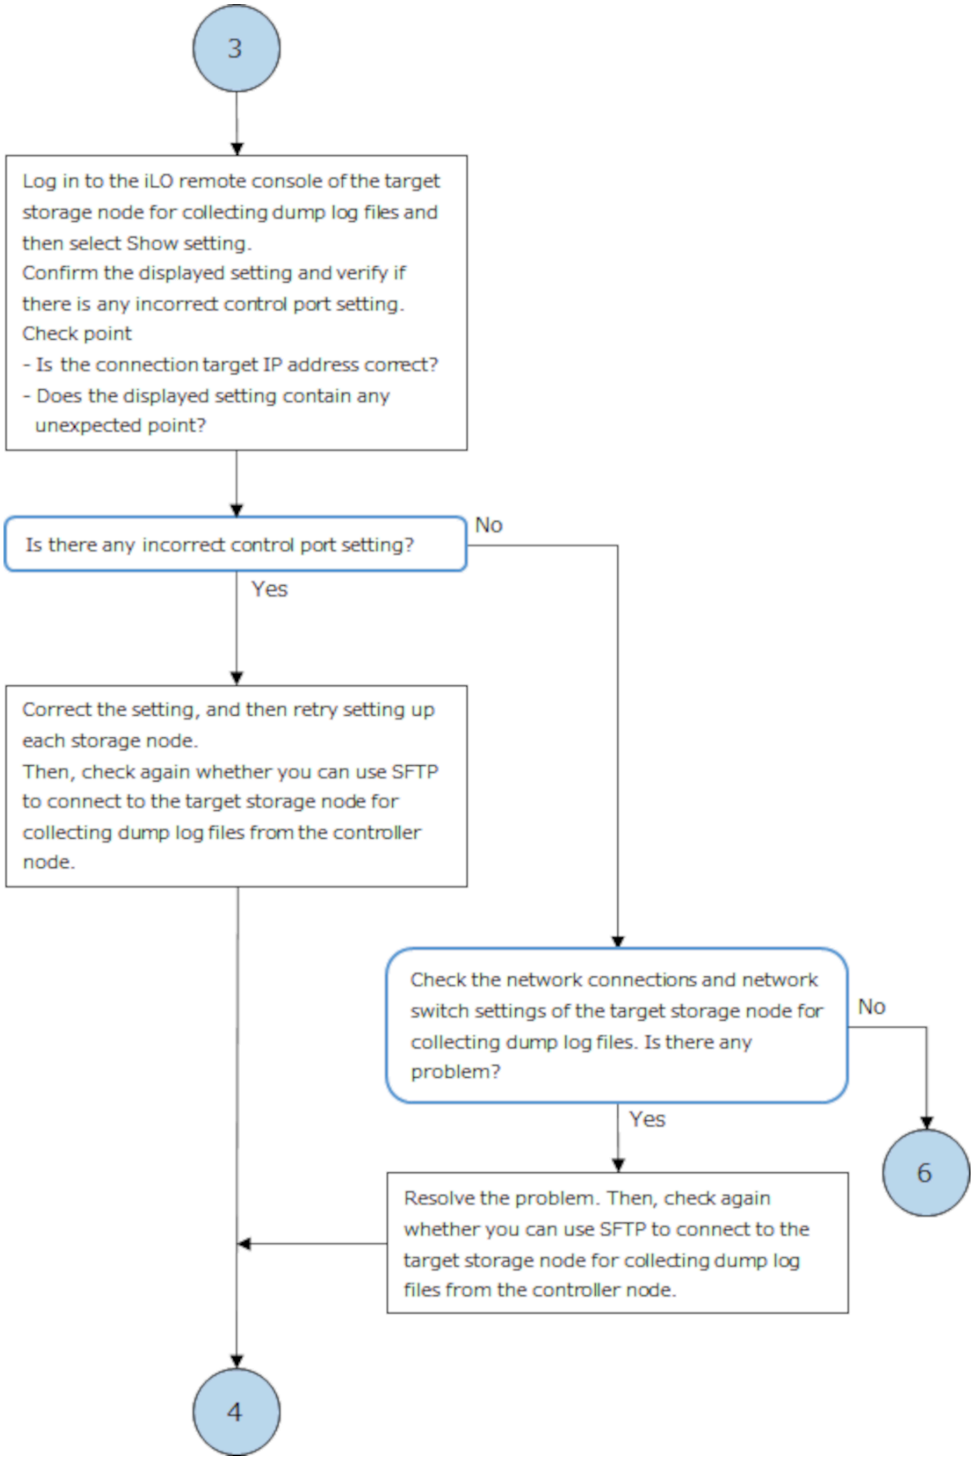
3
Log in to the iLO remote console of the target storage node for collecting dump log files and then select Show setting.
Confirm the displayed setting and verify if there is any incorrect control port setting.
Check point
- Is the connection target IP address correct?
- Does the displayed setting contain any unexpected point?
Is there any incorrect control port setting?
No
Yes
Correct the setting, and then retry setting up each storage node.
Then, check again whether you can use SFTP to connect to the target storage node for collecting dump log files from the controller node.
Check the network connections and network switch settings of the target storage node for collecting dump log files. Is there any problem?
No
Yes
6
Resolve the problem. Then, check again whether you can use SFTP to connect to the target storage node for collecting dump log files from the controller node.
4

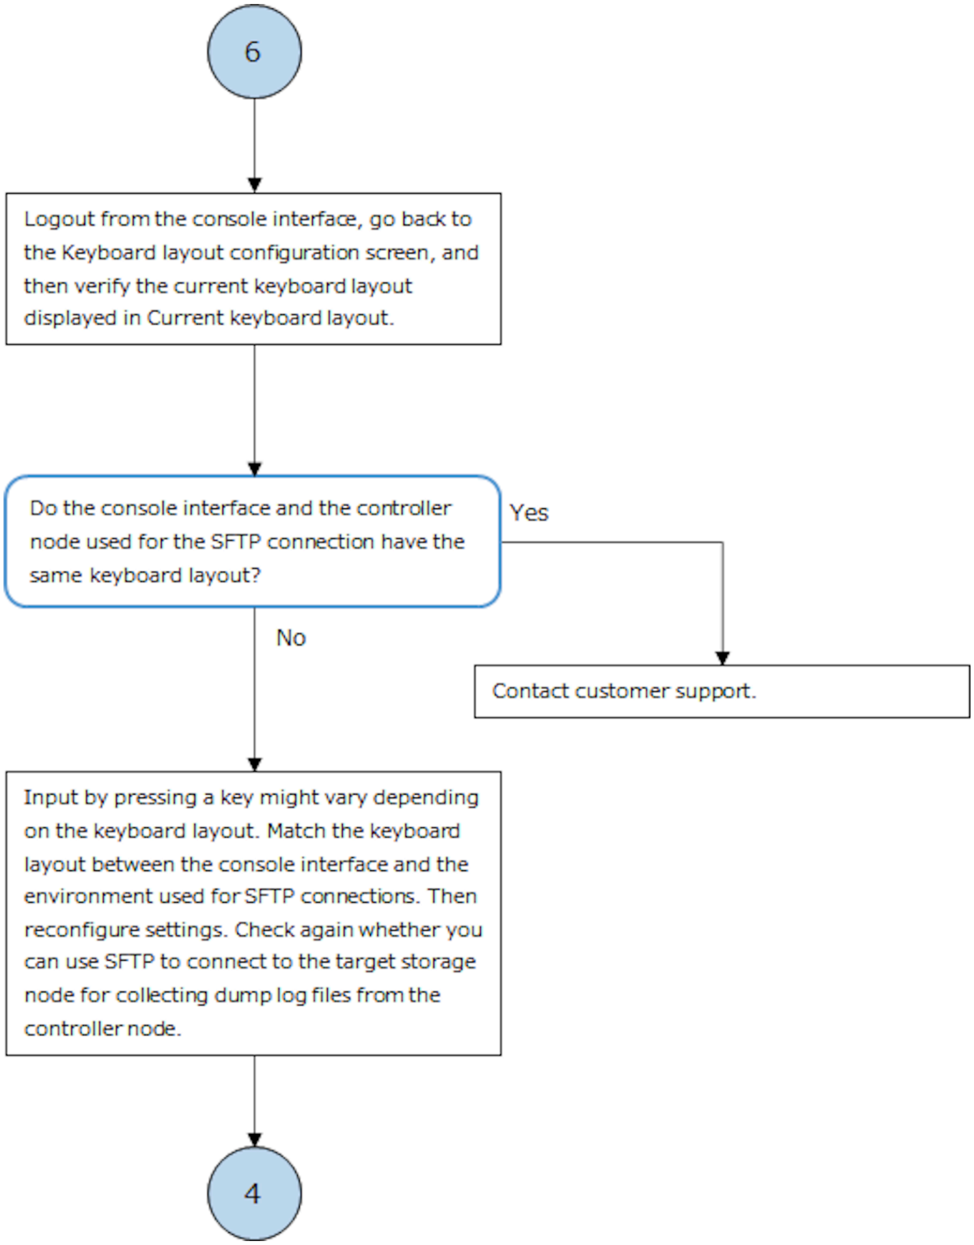

Logout from the console interface, go back to the Keyboard layout configuration screen, and then verify the current keyboard layout displayed in Current keyboard layout.

Do the console interface and the controller node used for the SFTP connection have the same keyboard layout?

Contact customer support.

Input by pressing a key might vary depending on the keyboard layout. Match the keyboard layout between the console interface and the environment used for SFTP connections. Then reconfigure settings. Check again whether you can use SFTP to connect to the target storage node for collecting dump log files from the controller node.

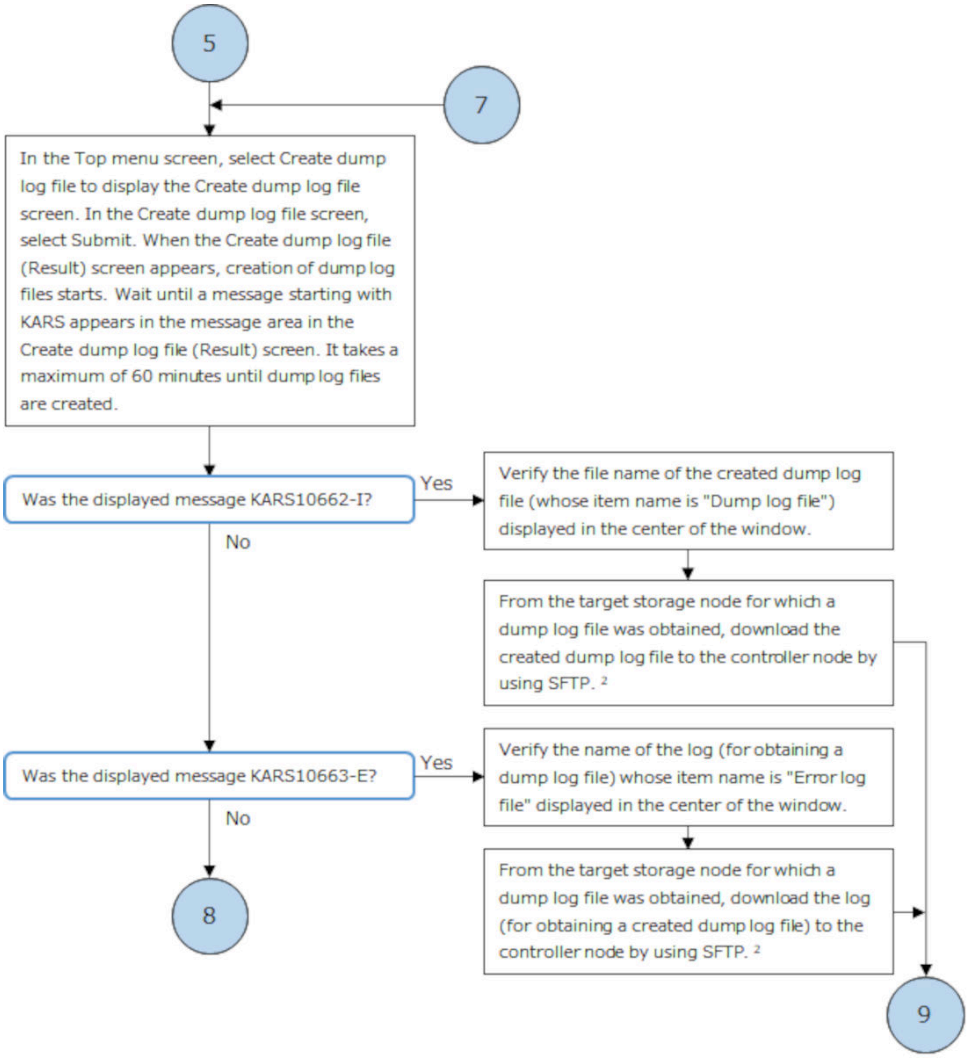
5
7
In the Top menu screen, select Create dump log file to display the Create dump log file screen. In the Create dump log file screen, select Submit. When the Create dump log file (Result) screen appears, creation of dump log files starts. Wait until a message starting with KARS appears in the message area in the Create dump log file (Result) screen. It takes a maximum of 60 minutes until dump log files are created.
Was the displayed message KARS10662-I?
Yes
No
Verify the file name of the created dump log file (whose item name is "Dump log file") displayed in the center of the window.
From the target storage node for which a dump log file was obtained, download the created dump log file to the controller node by using SFTP. 2
Was the displayed message KARS10663-E?
Yes
No
Verify the name of the log (for obtaining a dump log file) whose item name is "Error log file" displayed in the center of the window.
From the target storage node for which a dump log file was obtained, download the log (for obtaining a created dump log file) to the controller node by using SFTP. 2
8
9

2. See below for the information required to download the files to be downloaded from the target storage node for collecting dump log files using SFTP and a command example for downloading:

User name: setup
Password: Password of the setup user
Port number: 10022
Connection destination IP address: Control port IP address of the target storage node for collecting dump log files
Directory for storing the files to be downloaded: /dump
Name of the file to be downloaded: Name of the dump log file confirmed in the immediately previous step or the name of the log file for collecting dump log files

Command example

```
$ sftp -P 10022 setup@<control-port-IP-address-of-target-storage-node-for-collecting-dump-log-files>
sftp> cd dump
sftp> get <name-of-file-to-be-downloaded>
```

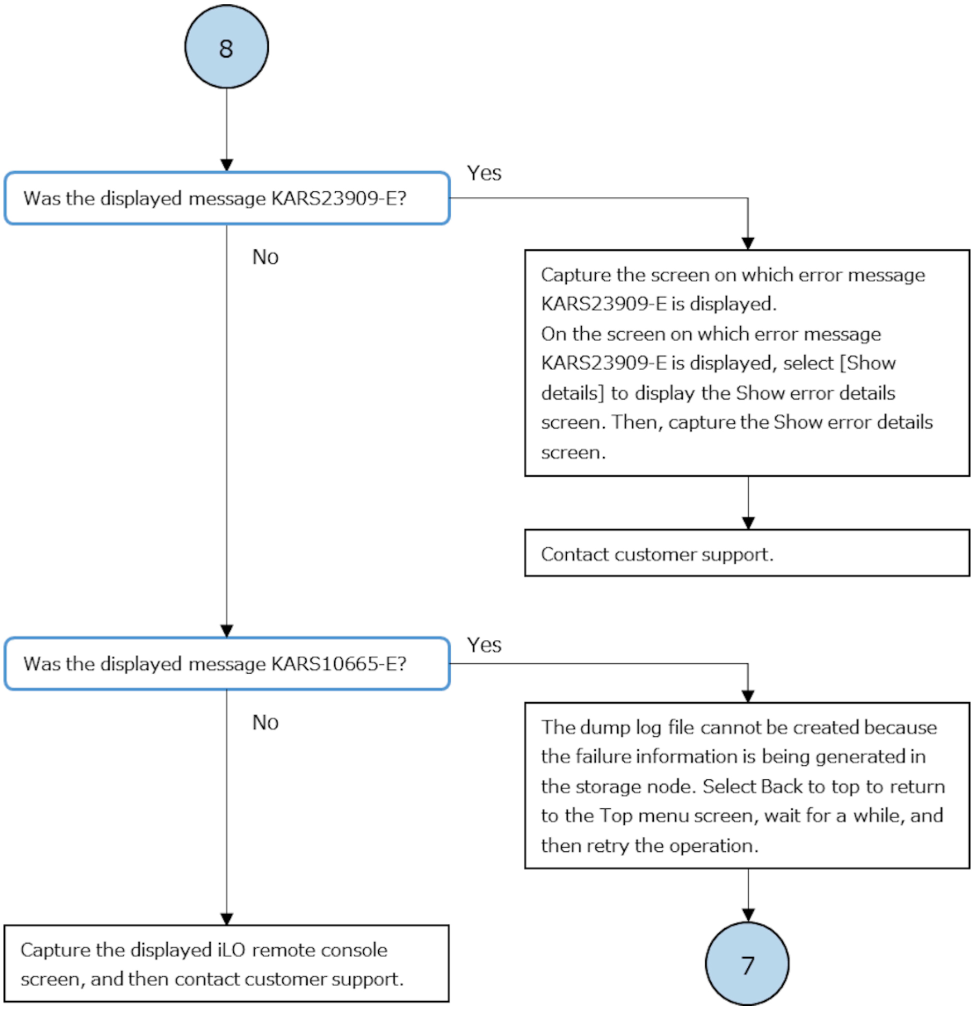

Was the displayed message KARS23909-E?

Capture the screen on which error message KARS23909-E is displayed.
On the screen on which error message KARS23909-E is displayed, select [Show details] to display the Show error details screen. Then, capture the Show error details screen.

Contact customer support.

Was the displayed message KARS10665-E?

The dump log file cannot be created because the failure information is being generated in the storage node. Select Back to top to return to the Top menu screen, wait for a while, and then retry the operation.

Capture the displayed iLO remote console screen, and then contact customer support.

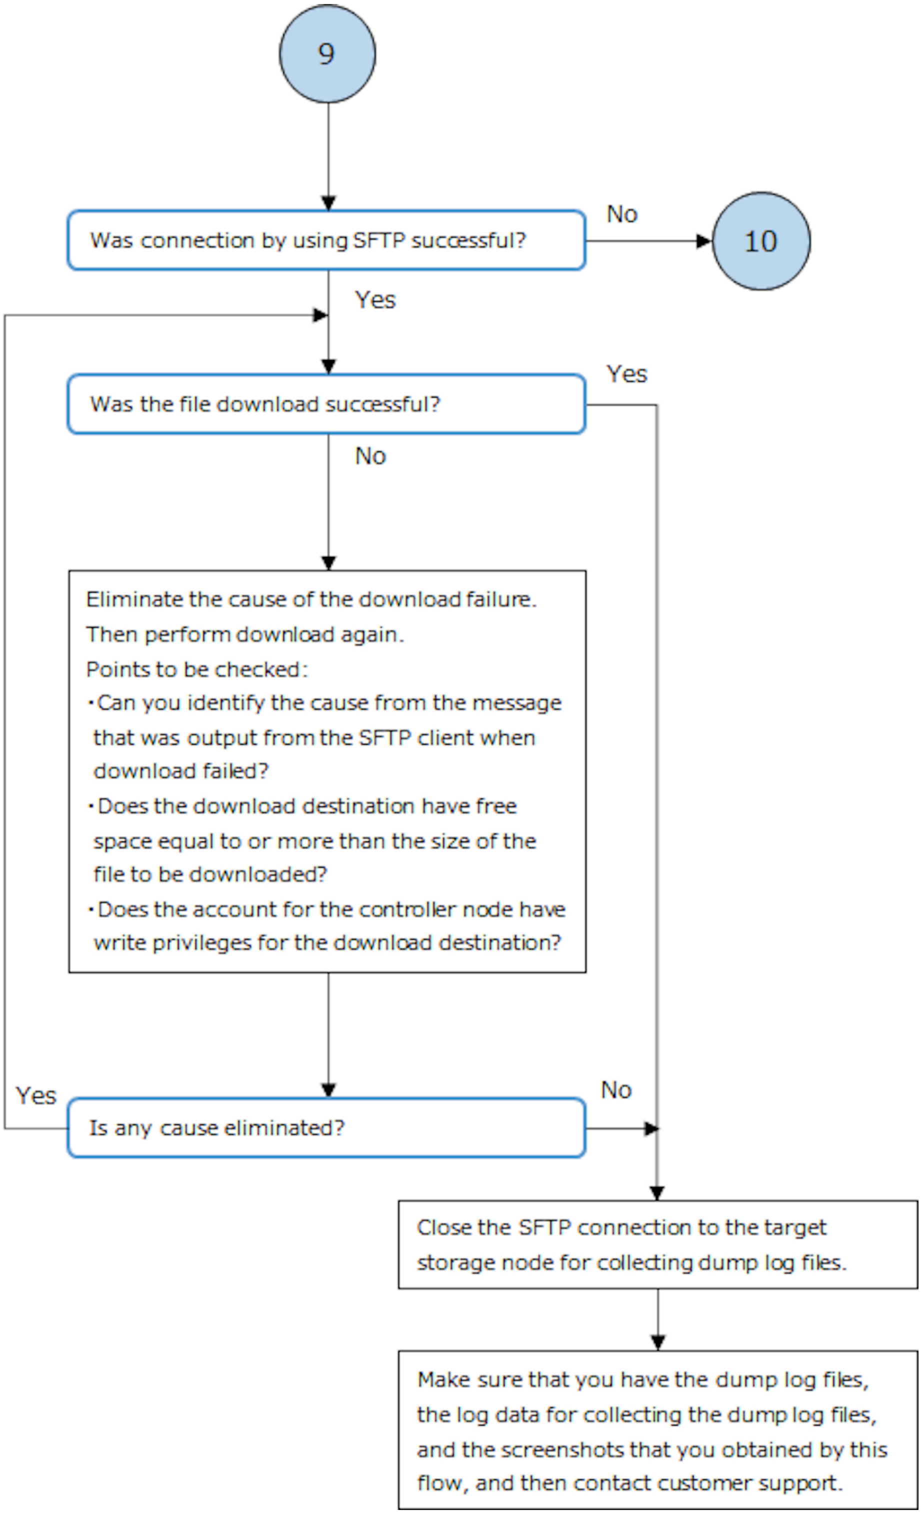
9
Was connection by using SFTP successful?
No
10
Yes
Was the file download successful?
Yes
No
Eliminate the cause of the download failure.
Then perform download again.
Points to be checked:
・Can you identify the cause from the message that was output from the SFTP client when download failed?
・Does the download destination have free space equal to or more than the size of the file to be downloaded?
・Does the account for the controller node have write privileges for the download destination?
Yes
Is any cause eliminated?
No
Close the SFTP connection to the target storage node for collecting dump log files.
Make sure that you have the dump log files, the log data for collecting the dump log files, and the screenshots that you obtained by this flow, and then contact customer support.

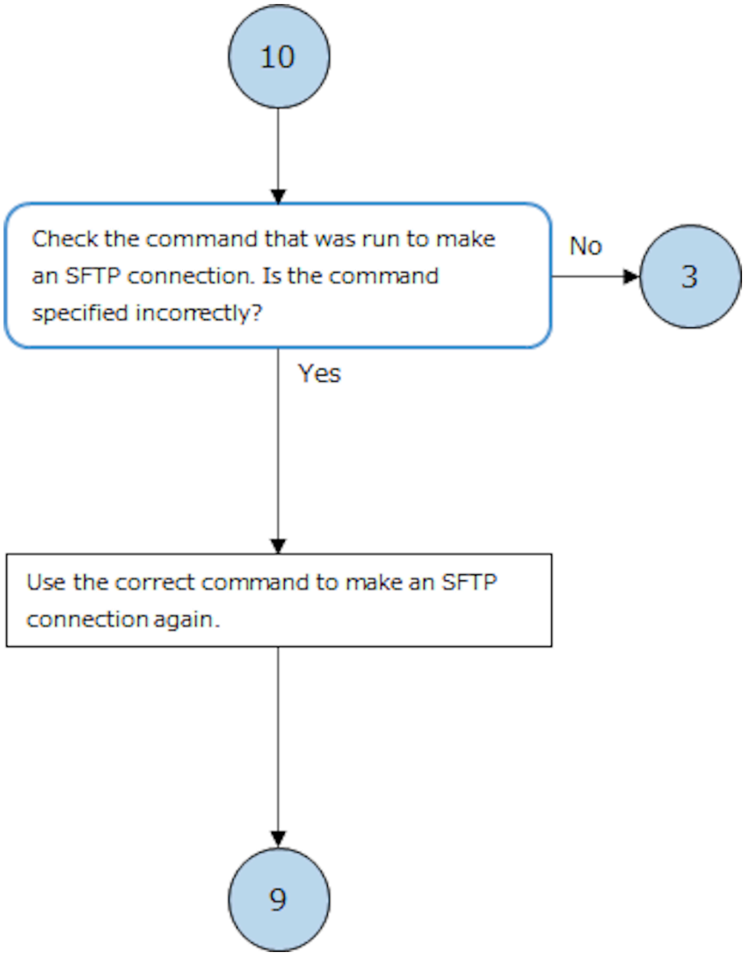
10
Check the command that was run to make an SFTP connection. Is the command specified incorrectly?
No
3
Yes
Use the correct command to make an SFTP connection again.
9

Action to be taken for an SSL/TLS certificate error

(Virtual machine) If a security warning is displayed when operating with the REST API/CLI or operating on the maintenance node (to add/replace storage nodes, change/set configuration information, or import/export a configuration file) or when the GUI is displayed, troubleshoot it according to the following flowchart.

(Bare metal) If a security warning is displayed when operating with the REST API/CLI or when the GUI is displayed, troubleshoot it according to the following flowchart.

In this section, "Subject Alternative Name" is abbreviated to "SAN".

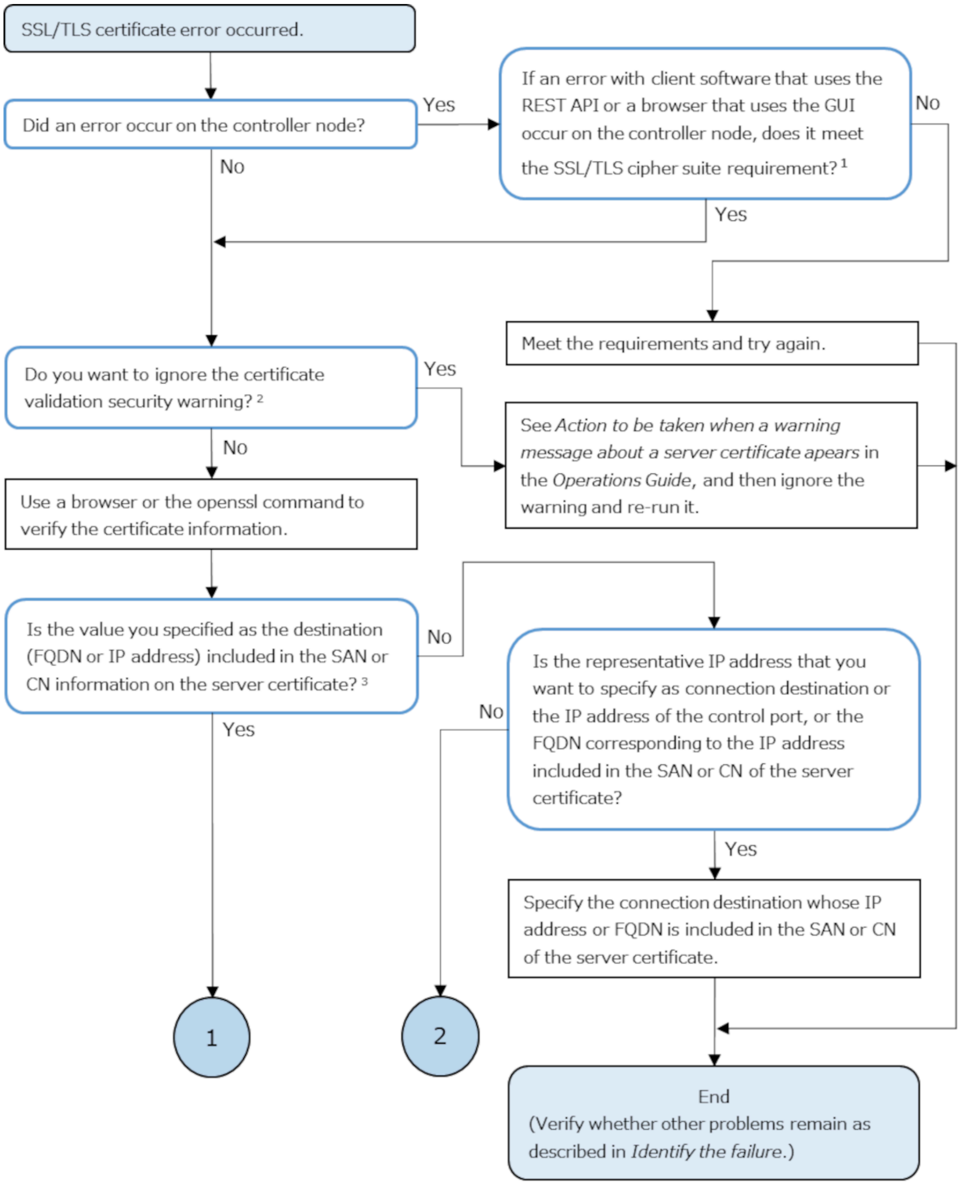
SSL/TLS certificate error occurred.
Did an error occur on the controller node?
Yes
No
If an error with client software that uses the REST API or a browser that uses the GUI occur on the controller node, does it meet the SSL/TLS cipher suite requirement? 1
No
Yes
Meet the requirements and try again.
Do you want to ignore the certificate validation security warning? 2
Yes
No
See Action to be taken when a warning message about a server certificate apears in the Operations Guide, and then ignore the warning and re-run it.
Use a browser or the openssl command to verify the certificate information.
Is the value you specified as the destination (FQDN or IP address) included in the SAN or CN information on the server certificate? 3
No
Yes
Is the representative IP address that you want to specify as connection destination or the IP address of the control port, or the FQDN corresponding to the IP address included in the SAN or CN of the server certificate?
No
Yes
Specify the connection destination whose IP address or FQDN is included in the SAN or CN of the server certificate.
1
2
End
(Verify whether other problems remain as described in Identify the failure.)

1. See *Client requirements for SSL/TLS communication* in the *Operations Guide*.
2. If you ignore the warning, you can still do SSL/TLS communication, which is not secure communication. (The communication is encrypted, but the other party is not authenticated.)
3. For information about the destinations to be specified for each SAN or CN information for server certificates, see *Client requirements for SSL/TLS communication* in the *Operations Guide*.

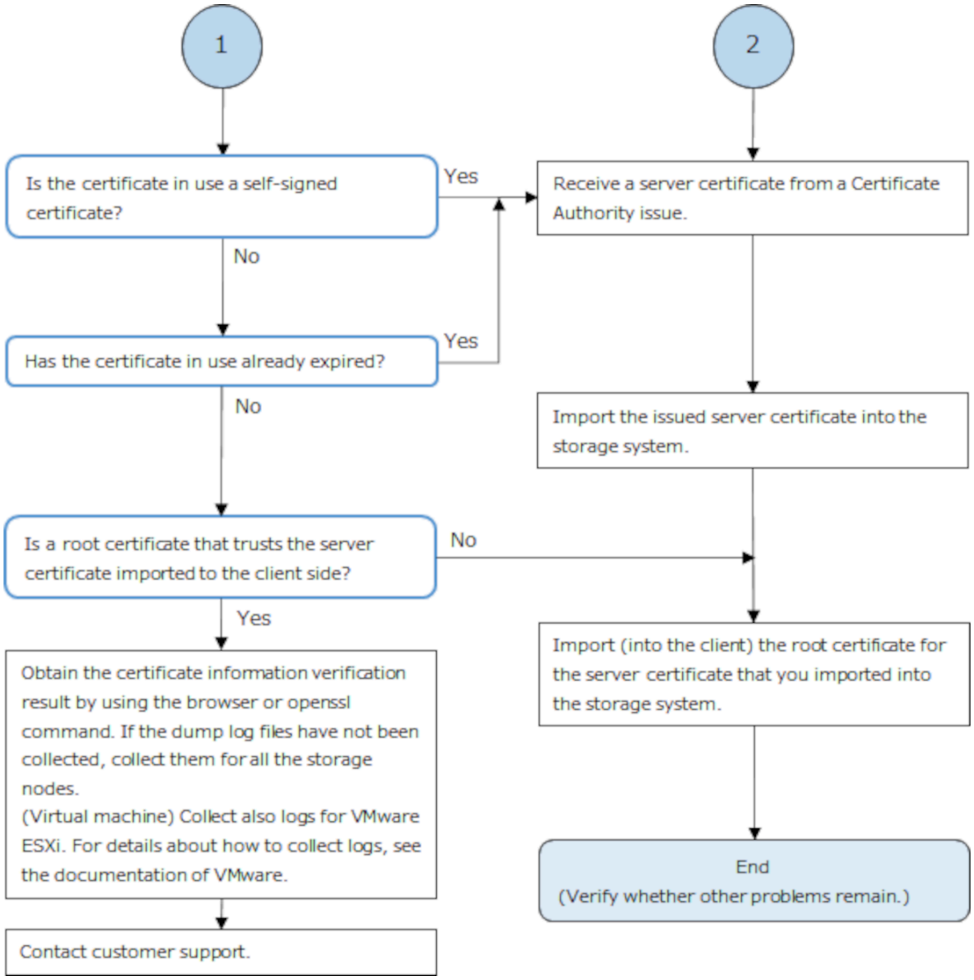
1
2
Is the certificate in use a self-signed certificate?
Yes
No
Receive a server certificate from a Certificate Authority issue.
Has the certificate in use already expired?
Yes
No
Import the issued server certificate into the storage system.
Is a root certificate that trusts the server certificate imported to the client side?
No
Yes
Import (into the client) the root certificate for the server certificate that you imported into the storage system.
Obtain the certificate information verification result by using the browser or openssl command. If the dump log files have not been collected, collect them for all the storage nodes.
(Virtual machine) Collect also logs for VMware ESXi. For details about how to collect logs, see the documentation of VMware.
End
(Verify whether other problems remain.)
Contact customer support.

Action to be taken for a user authentication error

If a failure occurred during user authentication, troubleshoot it according to the following flowchart.

- Could not log in to GUI and an error message was displayed.
- HTTP status code 401 or 403 was returned when a REST API/CLI command is run.
- When running the REST API/CLI, a solution indicating that the authentication error is to be handled according to troubleshooting guidelines was displayed. For more information, contact customer support.

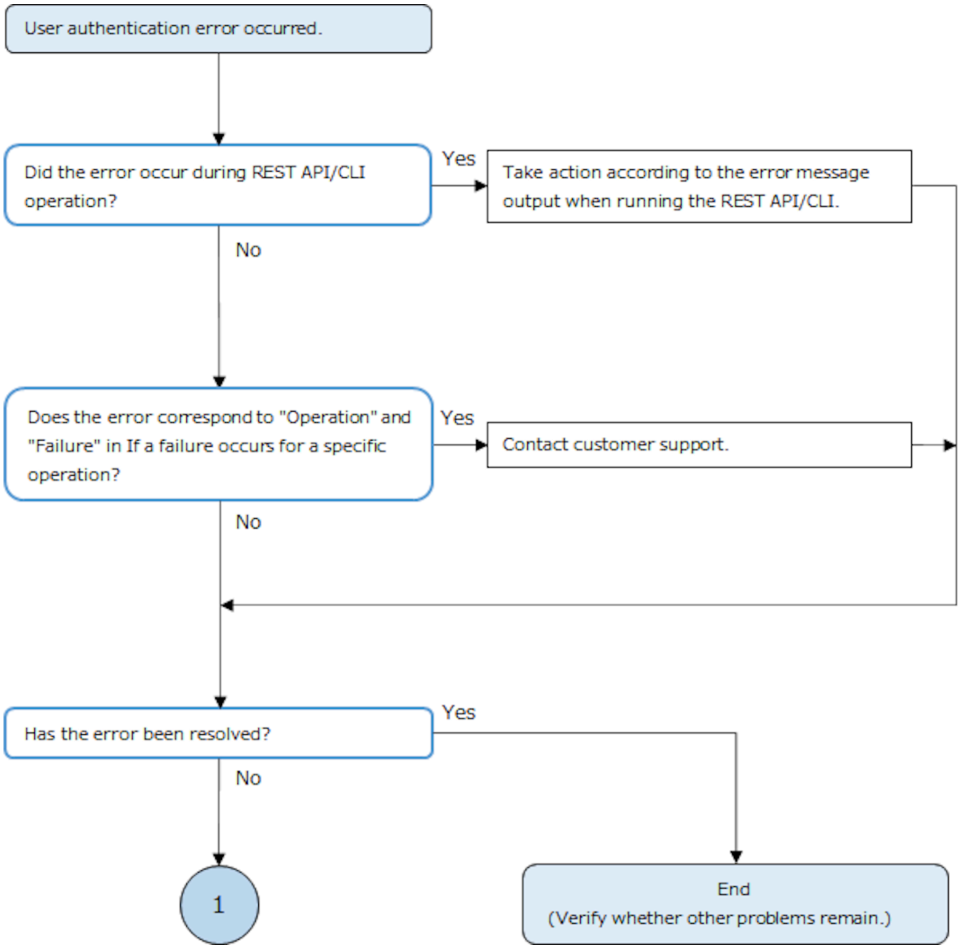
User authentication error occurred.
Did the error occur during REST API/CLI operation?
Yes
Take action according to the error message output when running the REST API/CLI.
No
Does the error correspond to "Operation" and "Failure" in If a failure occurs for a specific operation?
Yes
Contact customer support.
No
Has the error been resolved?
Yes
No
1
End
(Verify whether other problems remain.)

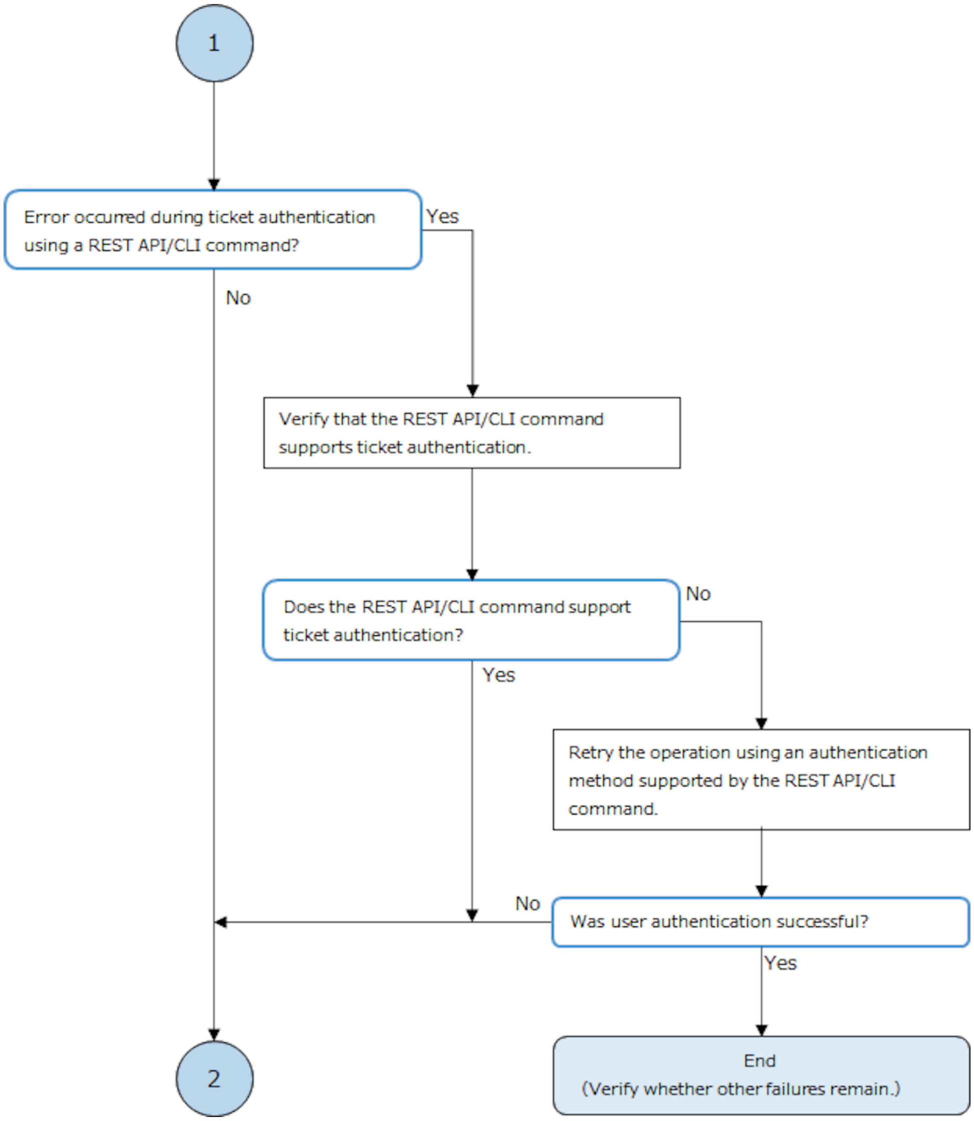
1
Error occurred during ticket authentication using a REST API/CLI command?
Yes
No
Verify that the REST API/CLI command supports ticket authentication.
Does the REST API/CLI command support ticket authentication?
No
Yes
Retry the operation using an authentication method supported by the REST API/CLI command.
No
Was user authentication successful?
Yes
2
End
(Verify whether other failures remain.)

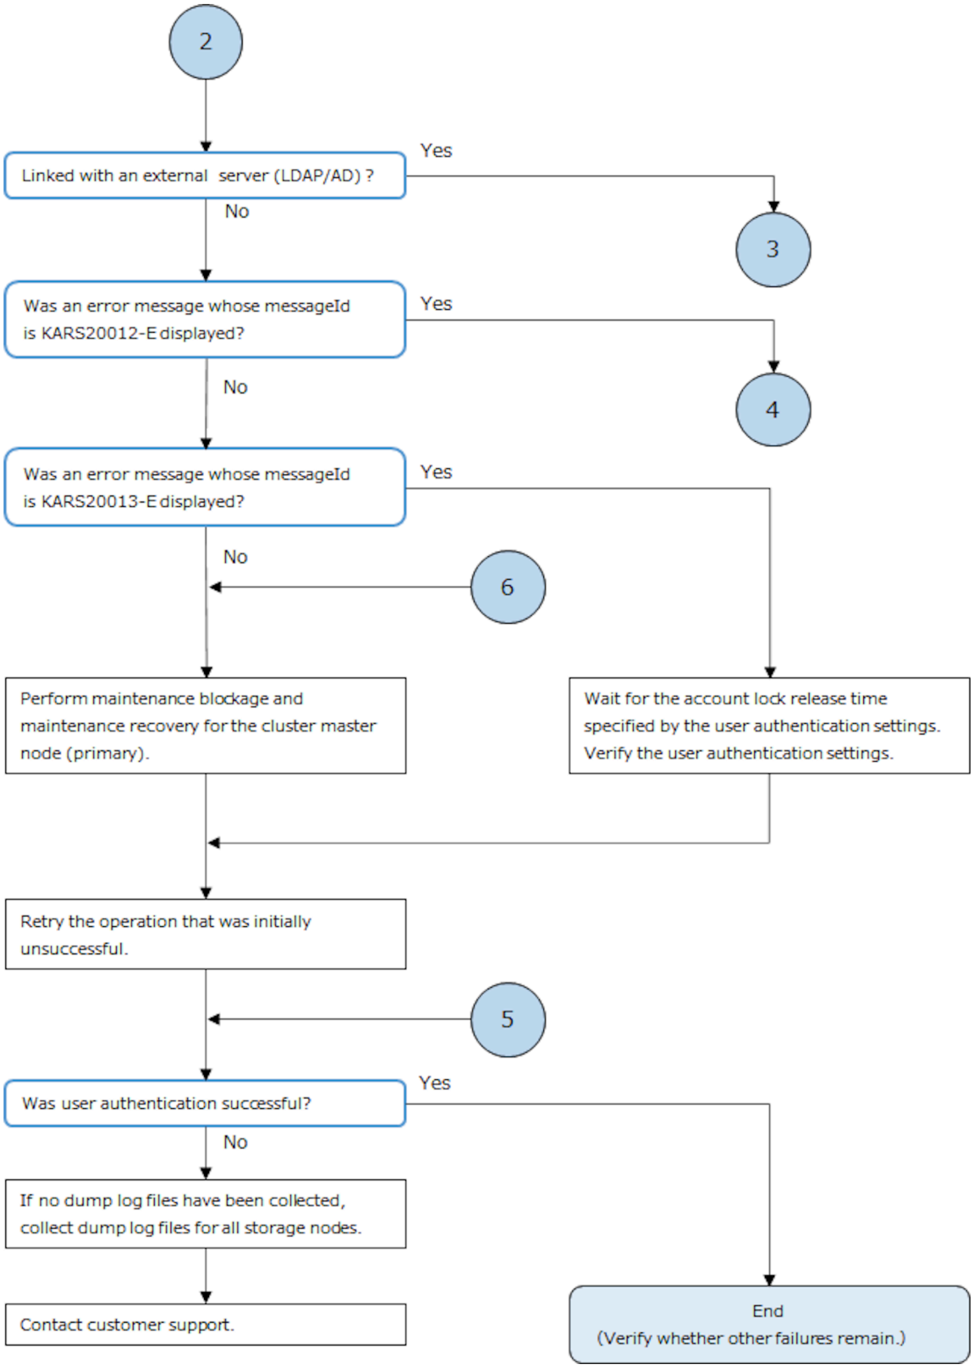
2
Linked with an external server (LDAP/AD) ?
Yes
3
No
Was an error message whose messageId is KARS20012-E displayed?
Yes
4
No
Was an error message whose messageId is KARS20013-E displayed?
Yes
No
6
Perform maintenance blockage and maintenance recovery for the cluster master node (primary).
Wait for the account lock release time specified by the user authentication settings. Verify the user authentication settings.
Retry the operation that was initially unsuccessful.
5
Was user authentication successful?
Yes
No
If no dump log files have been collected, collect dump log files for all storage nodes.
Contact customer support.
End
(Verify whether other failures remain.)

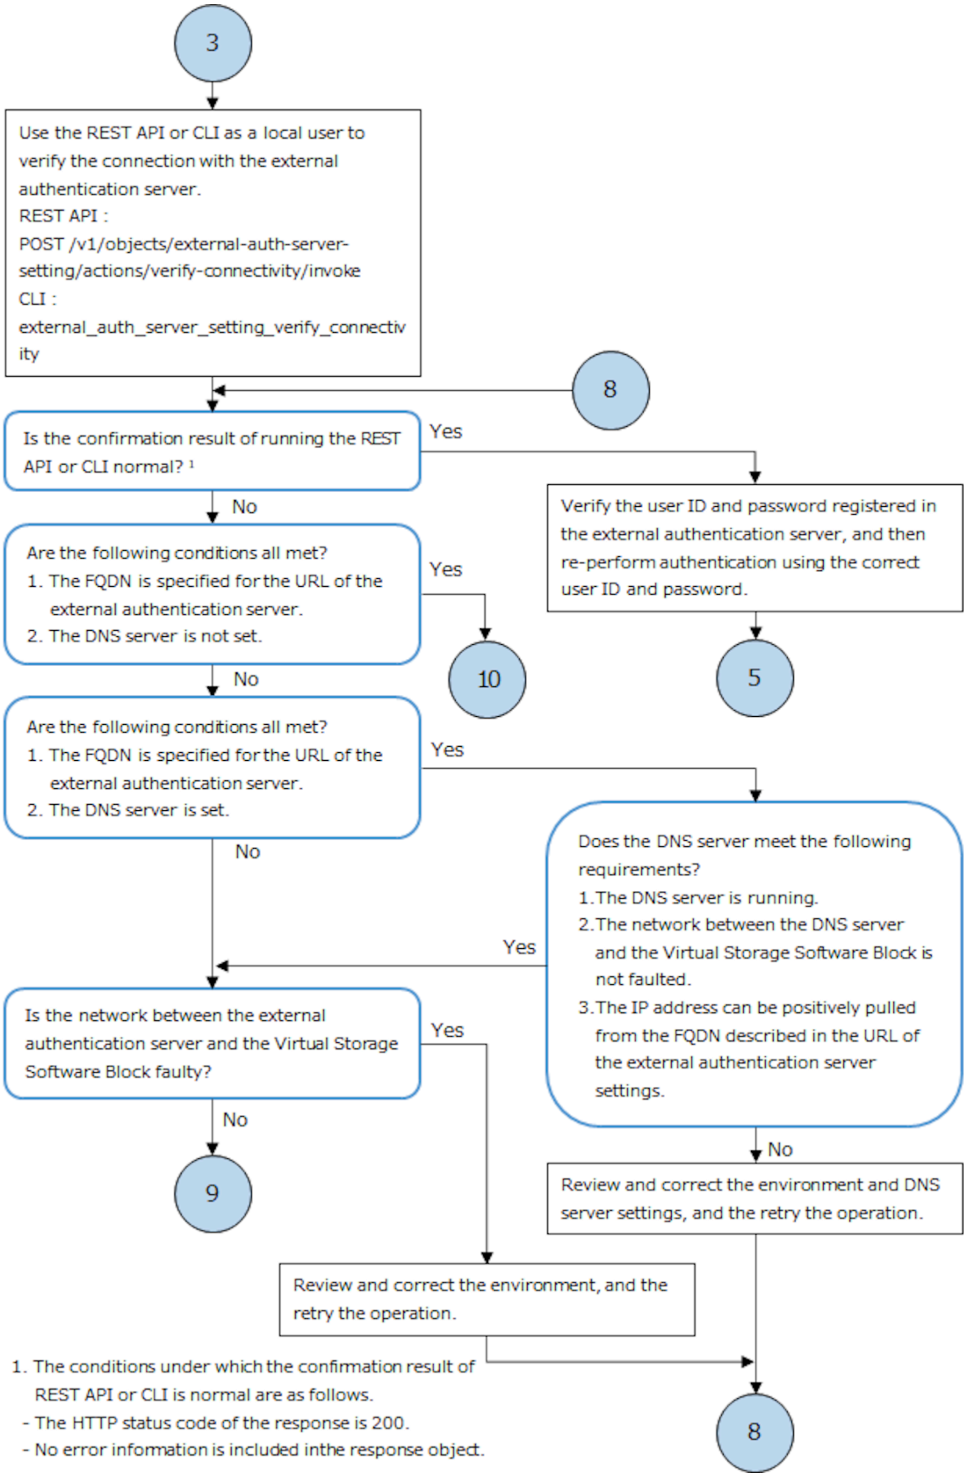
3
Use the REST API or CLI as a local user to verify the connection with the external authentication server.
REST API :
POST /v1/objects/external-auth-server-setting/actions/verify-connectivity/invoke
CLI :
external_auth_server_setting_verify_connectivity
8
Is the confirmation result of running the REST API or CLI normal? 1
Yes
Verify the user ID and password registered in the external authentication server, and then re-perform authentication using the correct user ID and password.
5
No
Are the following conditions all met?
1. The FQDN is specified for the URL of the external authentication server.
2. The DNS server is not set.
Yes
10
No
Are the following conditions all met?
1. The FQDN is specified for the URL of the external authentication server.
2. The DNS server is set.
Yes
Does the DNS server meet the following requirements?
1.The DNS server is running.
2.The network between the DNS server and the Virtual Storage Software Block is not faulted.
3.The IP address can be positively pulled from the FQDN described in the URL of the external authentication server settings.
No
Yes
Is the network between the external authentication server and the Virtual Storage Software Block faulty?
Yes
No
9
No
Review and correct the environment and DNS server settings, and the retry the operation.
Review and correct the environment, and the retry the operation.
8

1. The conditions under which the confirmation result of REST API or CLI is normal are as follows.
- The HTTP status code of the response is 200.
- No error information is included inthe response object.

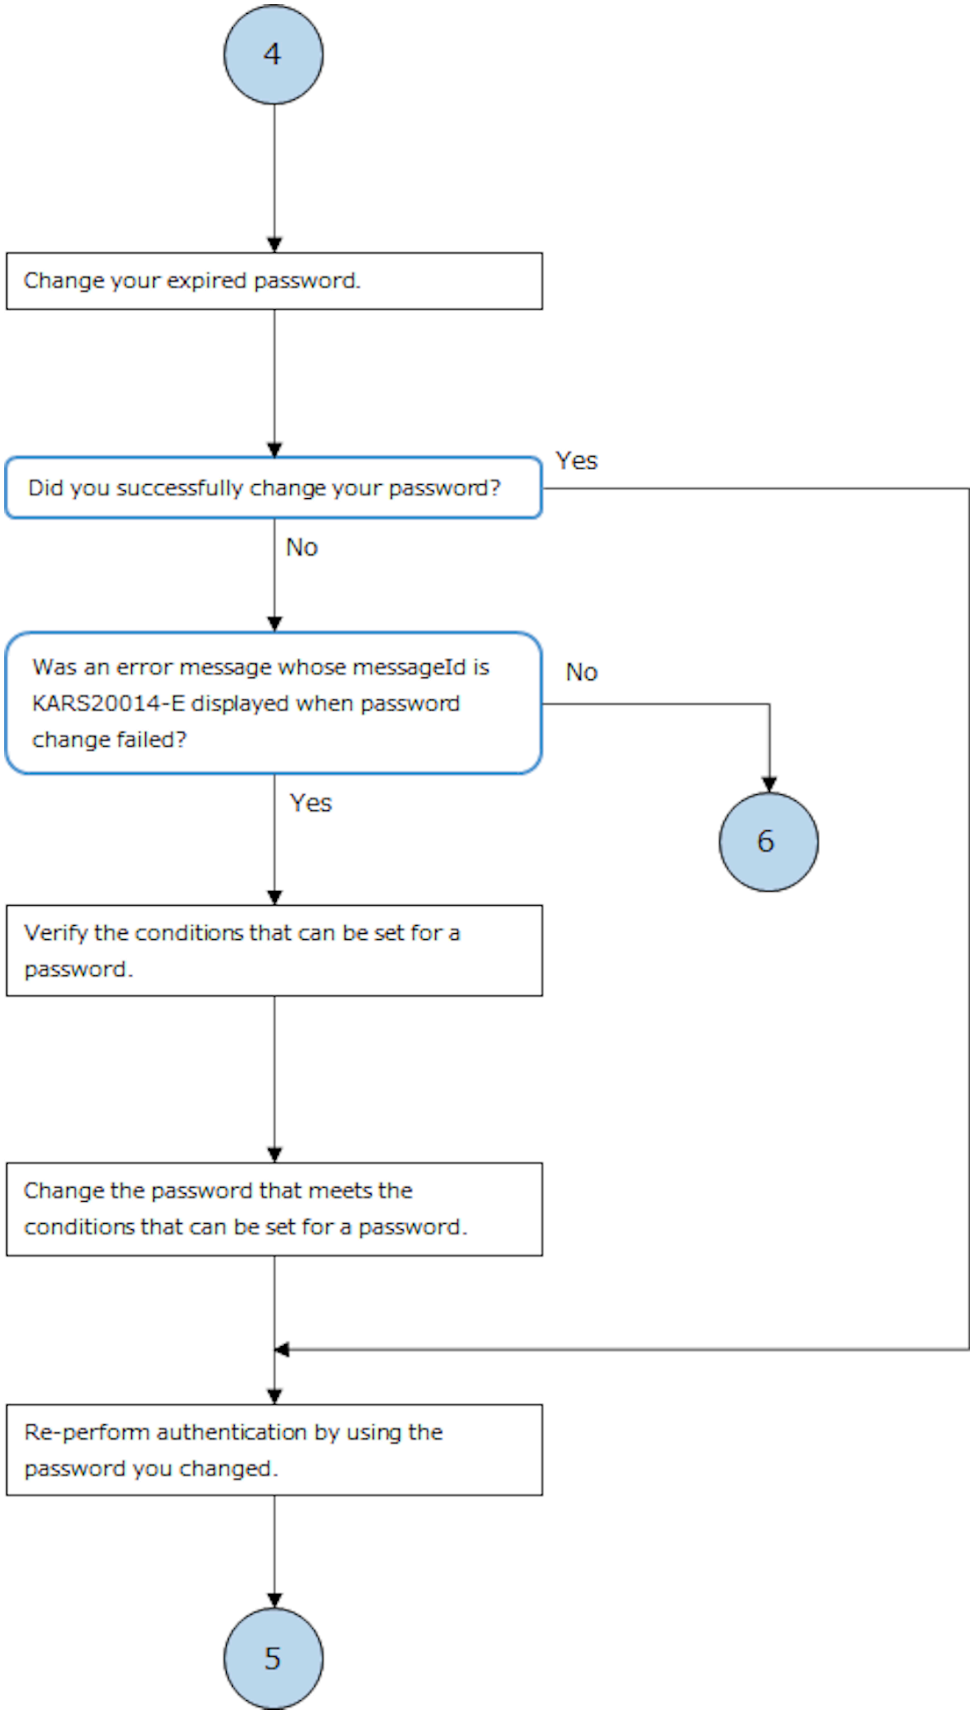
4
Change your expired password.
Did you successfully change your password?
Yes
No
Was an error message whose messageId is KARS20014-E displayed when password change failed?
No
6
Yes
Verify the conditions that can be set for a password.
Change the password that meets the conditions that can be set for a password.
Re-perform authentication by using the password you changed.
5

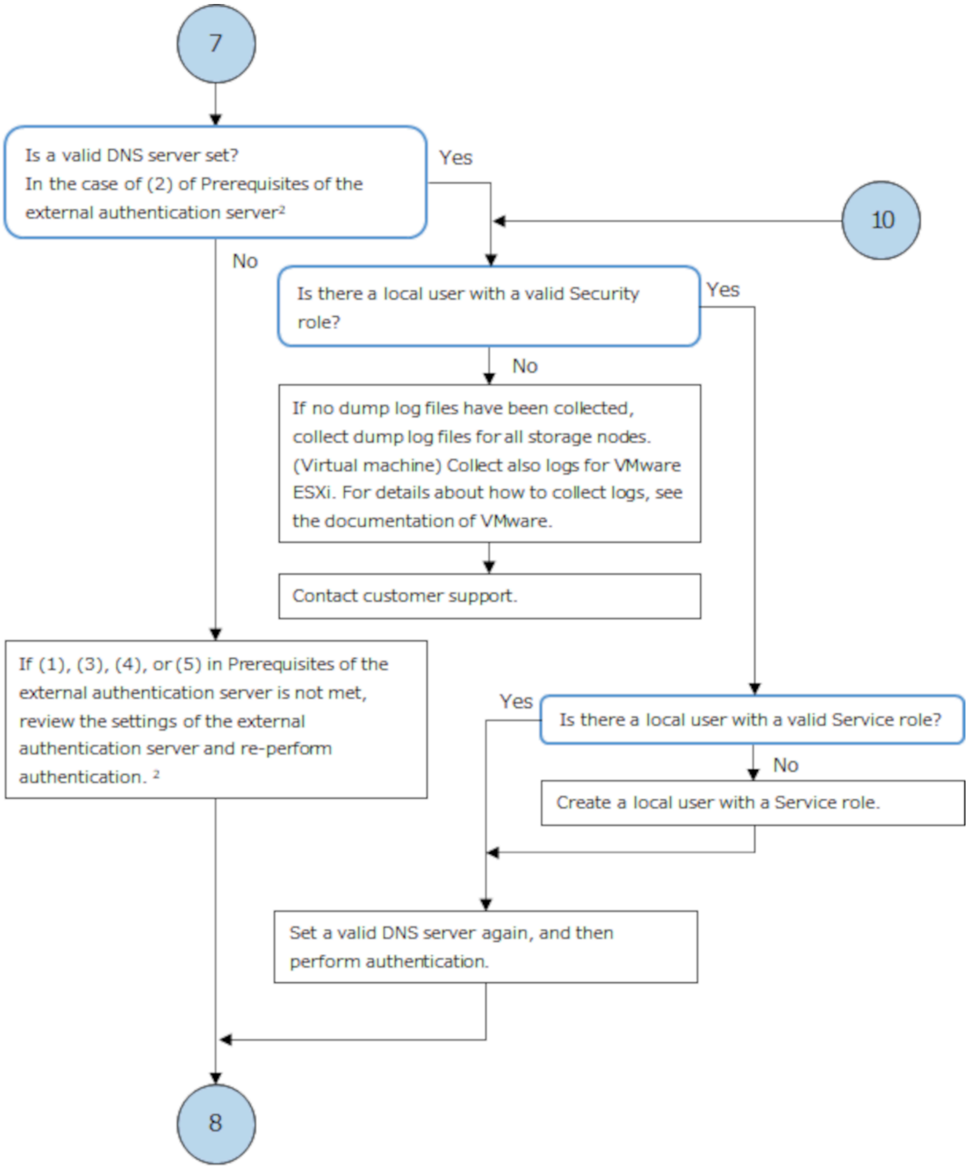
7
Is a valid DNS server set?
In the case of (2) of Prerequisites of the external authentication server2
Yes
10
No
Is there a local user with a valid Security role?
Yes
No
If no dump log files have been collected, collect dump log files for all storage nodes. (Virtual machine) Collect also logs for VMware ESXi. For details about how to collect logs, see the documentation of VMware.
Contact customer support.
If (1), (3), (4), or (5) in Prerequisites of the external authentication server is not met, review the settings of the external authentication server and re-perform authentication. 2
Yes
Is there a local user with a valid Service role?
No
Create a local user with a Service role.
Set a valid DNS server again, and then perform authentication.
8

2. Prerequisites of the external authentication server

(1) The URL of the external authentication server is specified with the FQDN.

(2) The DNS server is set.

(3) The cipher suites that can be used on the external authentication server meet the requirements.

(4) The root certificate that proves the authenticity of the server certificates set on the externalauthentication server has been imported into the storage.

(5) The FQDN in the URL of the external authentication server set to Virtual Storage Software Block is included in Common Name or Subject Alternative Name of the certificate imported into the external authentication server.

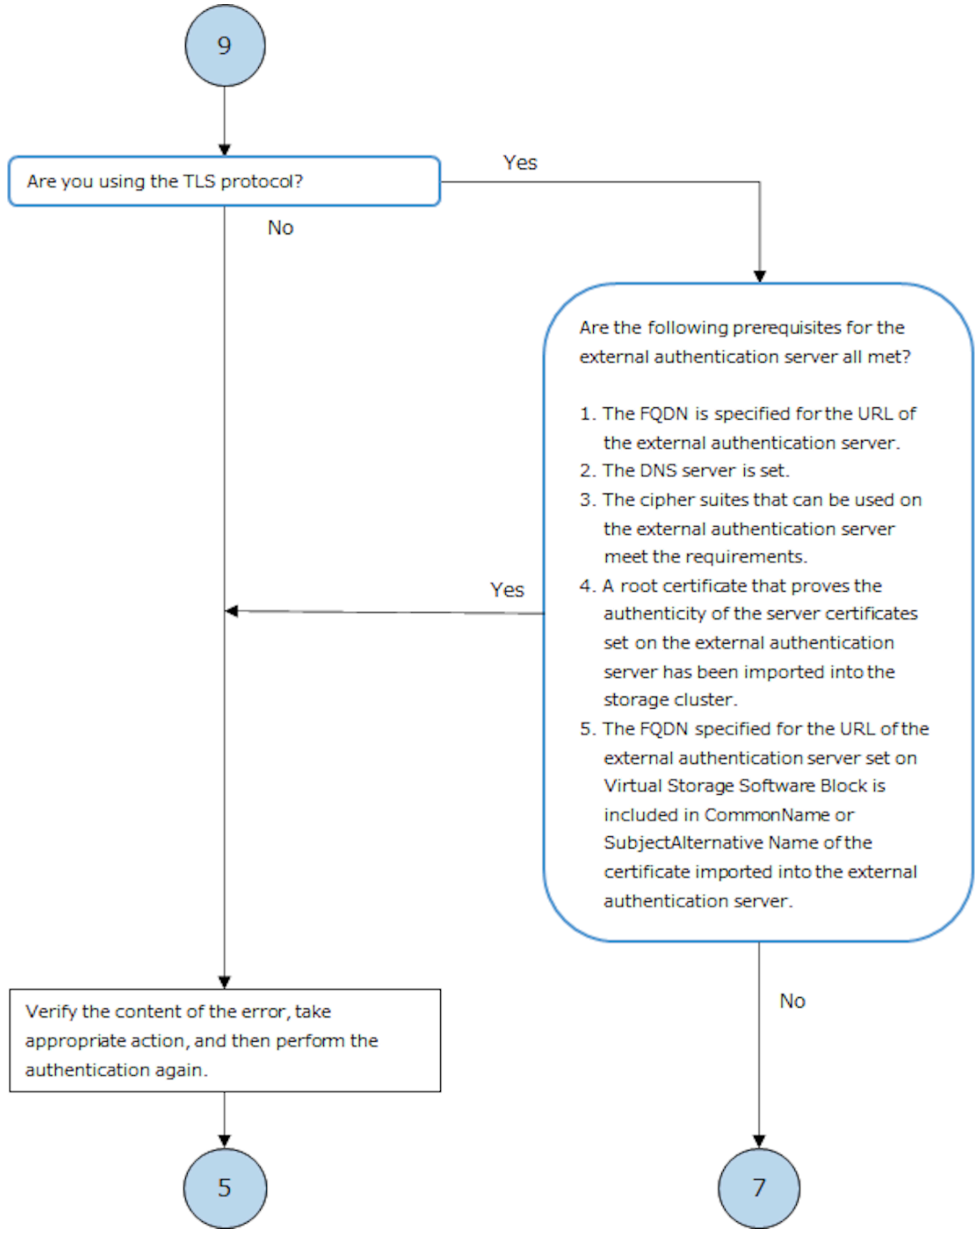
9
Are you using the TLS protocol?
Yes
No
Are the following prerequisites for the external authentication server all met?
1. The FQDN is specified for the URL of the external authentication server.
2. The DNS server is set.
3. The cipher suites that can be used on the external authentication server meet the requirements.
4. A root certificate that proves the authenticity of the server certificates set on the external authentication server has been imported into the storage cluster.
5. The FQDN specified for the URL of the external authentication server set on Virtual Storage Software Block is included in CommonName or SubjectAlternative Name of the certificate imported into the external authentication server.
Yes
No
Verify the content of the error, take appropriate action, and then perform the authentication again.
5
7

Action to be taken when REST API or CLI is unavailable

If REST API/CLI execution is unsuccessful, troubleshoot the failure according to the following flowchart.

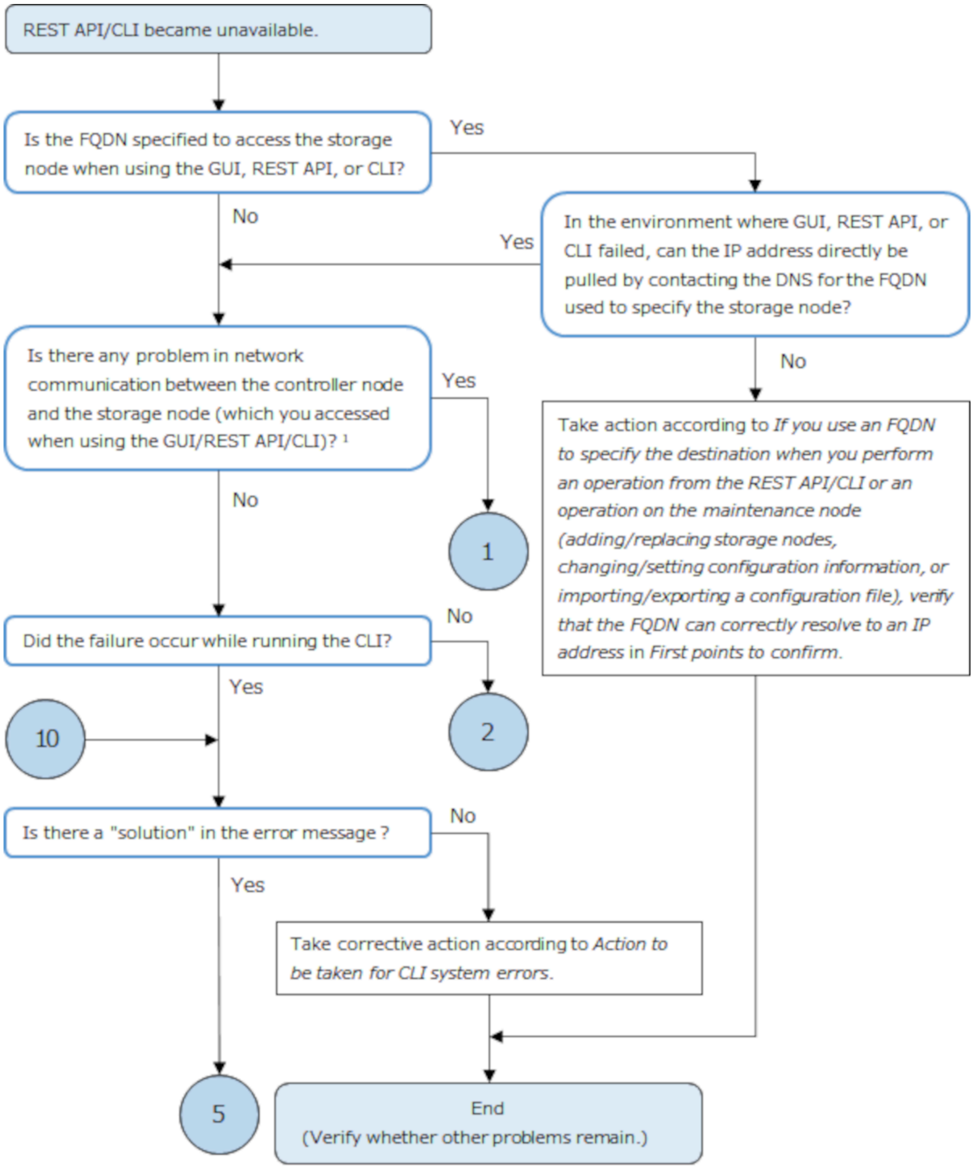
REST API/CLI became unavailable.
Is the FQDN specified to access the storage node when using the GUI, REST API, or CLI?
Yes
No
In the environment where GUI, REST API, or CLI failed, can the IP address directly be pulled by contacting the DNS for the FQDN used to specify the storage node?
Yes
No
Is there any problem in network communication between the controller node and the storage node (which you accessed when using the GUI/REST API/CLI)? 1
Yes
No
1
Take action according to If you use an FQDN to specify the destination when you perform an operation from the REST API/CLI or an operation on the maintenance node (adding/replacing storage nodes, changing/setting configuration information, or importing/exporting a configuration file), verify that the FQDN can correctly resolve to an IP address in First points to confirm.
Did the failure occur while running the CLI?
No
Yes
2
10
Is there a "solution" in the error message ?
No
Yes
Take corrective action according to Action to be taken for CLI system errors.
5
End
(Verify whether other problems remain.)

1. You can verify the network connectivity by running the following command from a console such as the command prompt:
   - In Windows (run the following commands on PowerShell)

   ```
   > Test-NetConnection <target-storage-node-IPaddress> -Port 443
   ```

   - In Linux

   ```
   $ curl -k -I https://<target-storage-node-IPaddress>/hsds/
   ```

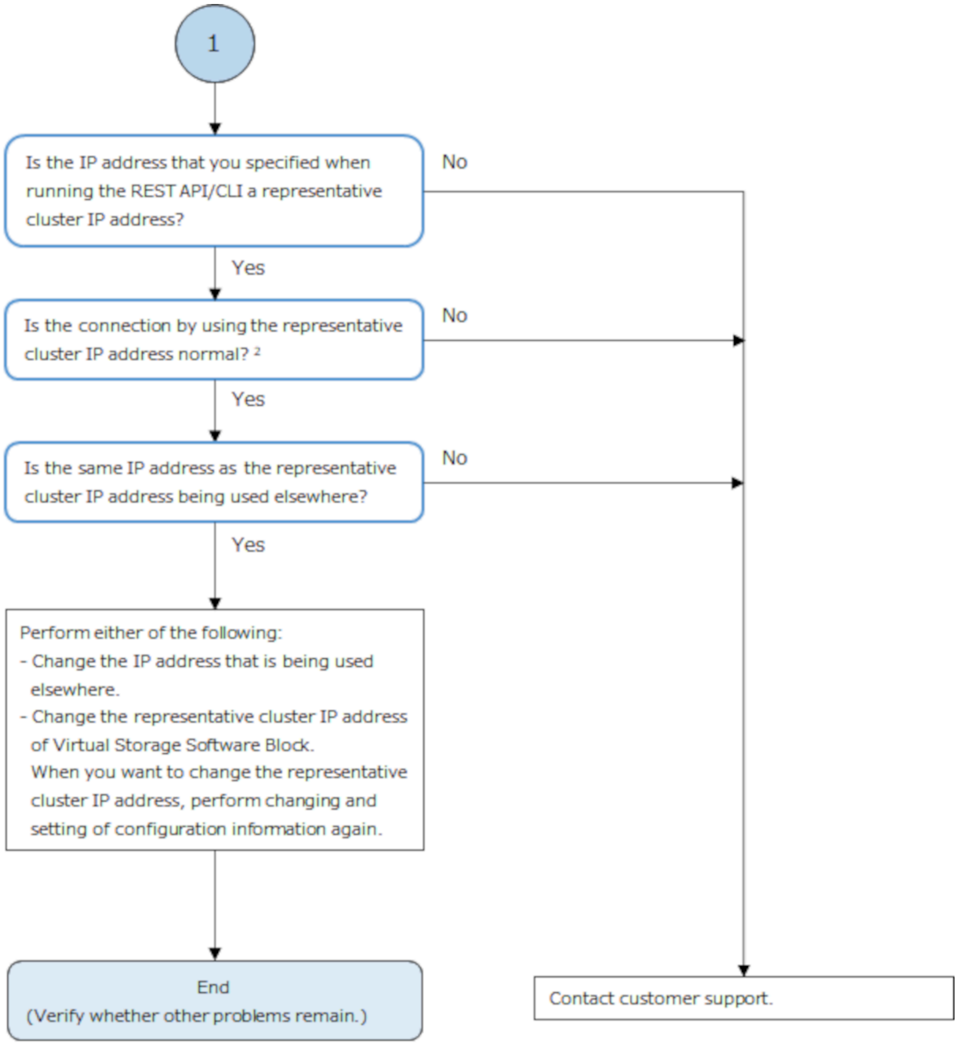
1
Is the IP address that you specified when running the REST API/CLI a representative cluster IP address?
No
Yes
Is the connection by using the representative cluster IP address normal? 2
No
Yes
Is the same IP address as the representative cluster IP address being used elsewhere?
No
Yes
Perform either of the following:
- Change the IP address that is being used elsewhere.
- Change the representative cluster IP address of Virtual Storage Software Block.
When you want to change the representative cluster IP address, perform changing and setting of configuration information again.
End
(Verify whether other problems remain.)
Contact customer support.

2. You can verify the connection with the representative IP address by running the following command.

```
> ping <representative-cluster-IPaddress>
```

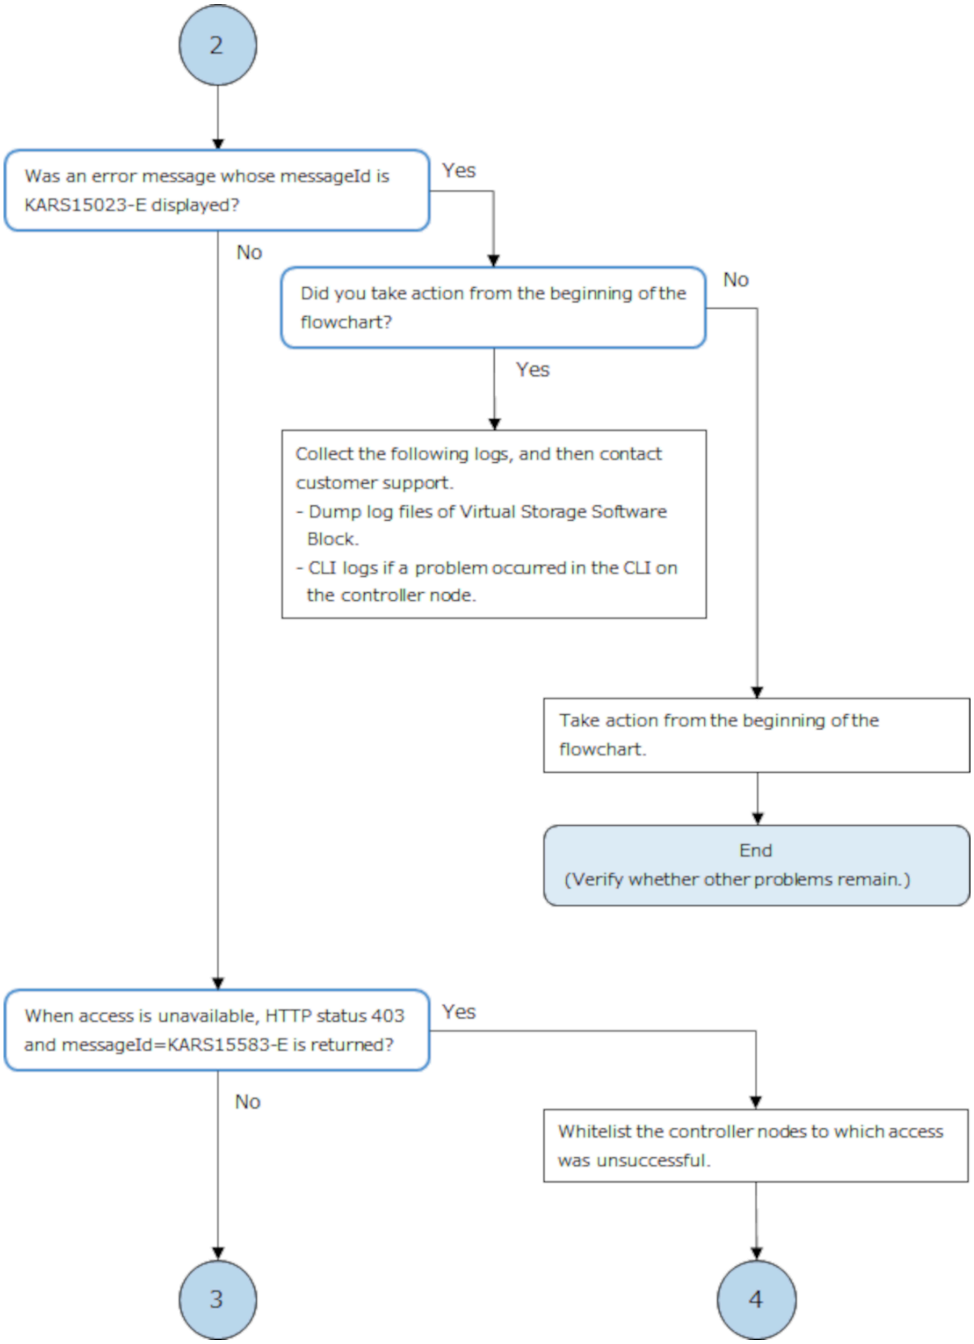
2
Was an error message whose messageId is KARS15023-E displayed?
Yes
No
Did you take action from the beginning of the flowchart?
No
Yes
Collect the following logs, and then contact customer support.
- Dump log files of Virtual Storage Software Block.
- CLI logs if a problem occurred in the CLI on the controller node.
Take action from the beginning of the flowchart.
End
(Verify whether other problems remain.)
When access is unavailable, HTTP status 403 and messageId=KARS15583-E is returned?
Yes
No
Whitelist the controller nodes to which access was unsuccessful.
3
4

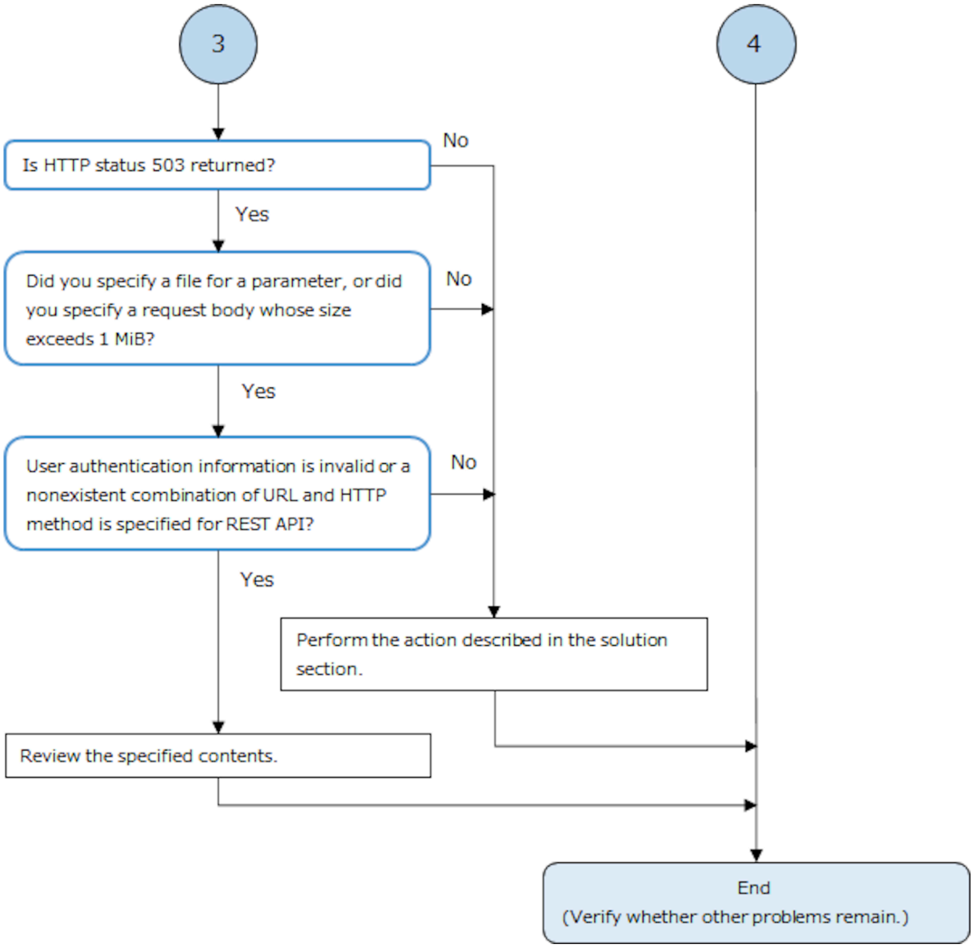
3
4
Is HTTP status 503 returned?
No
Yes
Did you specify a file for a parameter, or did you specify a request body whose size exceeds 1 MiB?
No
Yes
User authentication information is invalid or a nonexistent combination of URL and HTTP method is specified for REST API?
No
Yes
Perform the action described in the solution section.
Review the specified contents.
End
(Verify whether other problems remain.)

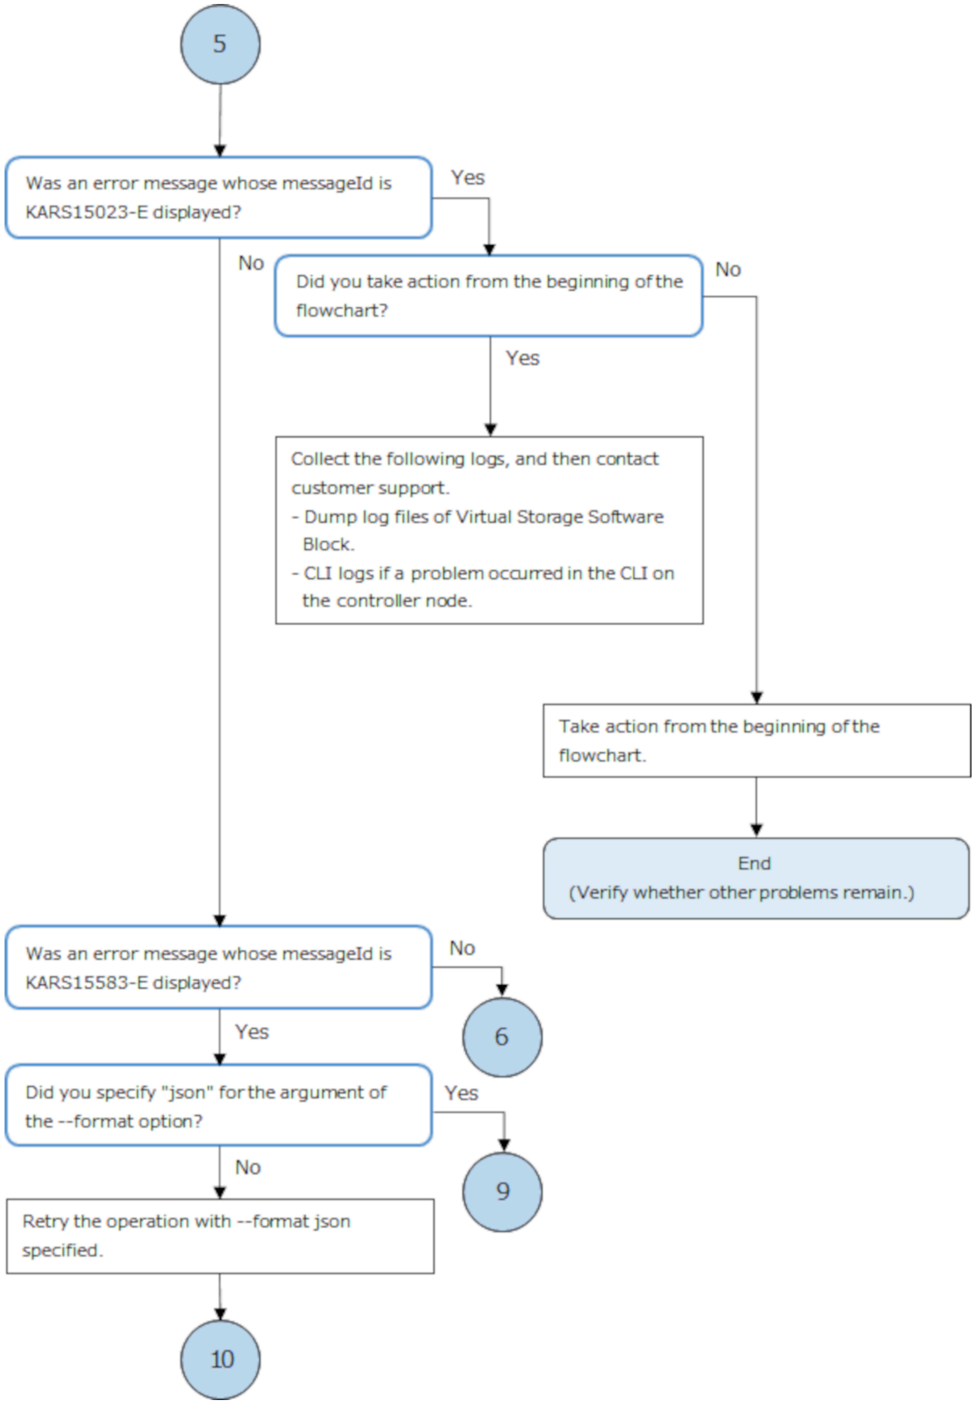
5
Was an error message whose messageId is KARS15023-E displayed?
Yes
No
Did you take action from the beginning of the flowchart?
No
Yes
Collect the following logs, and then contact customer support.
- Dump log files of Virtual Storage Software Block.
- CLI logs if a problem occurred in the CLI on the controller node.
Take action from the beginning of the flowchart.
End
(Verify whether other problems remain.)
Was an error message whose messageId is KARS15583-E displayed?
No
6
Yes
Did you specify "json" for the argument of the --format option?
Yes
9
No
Retry the operation with --format json specified.
10

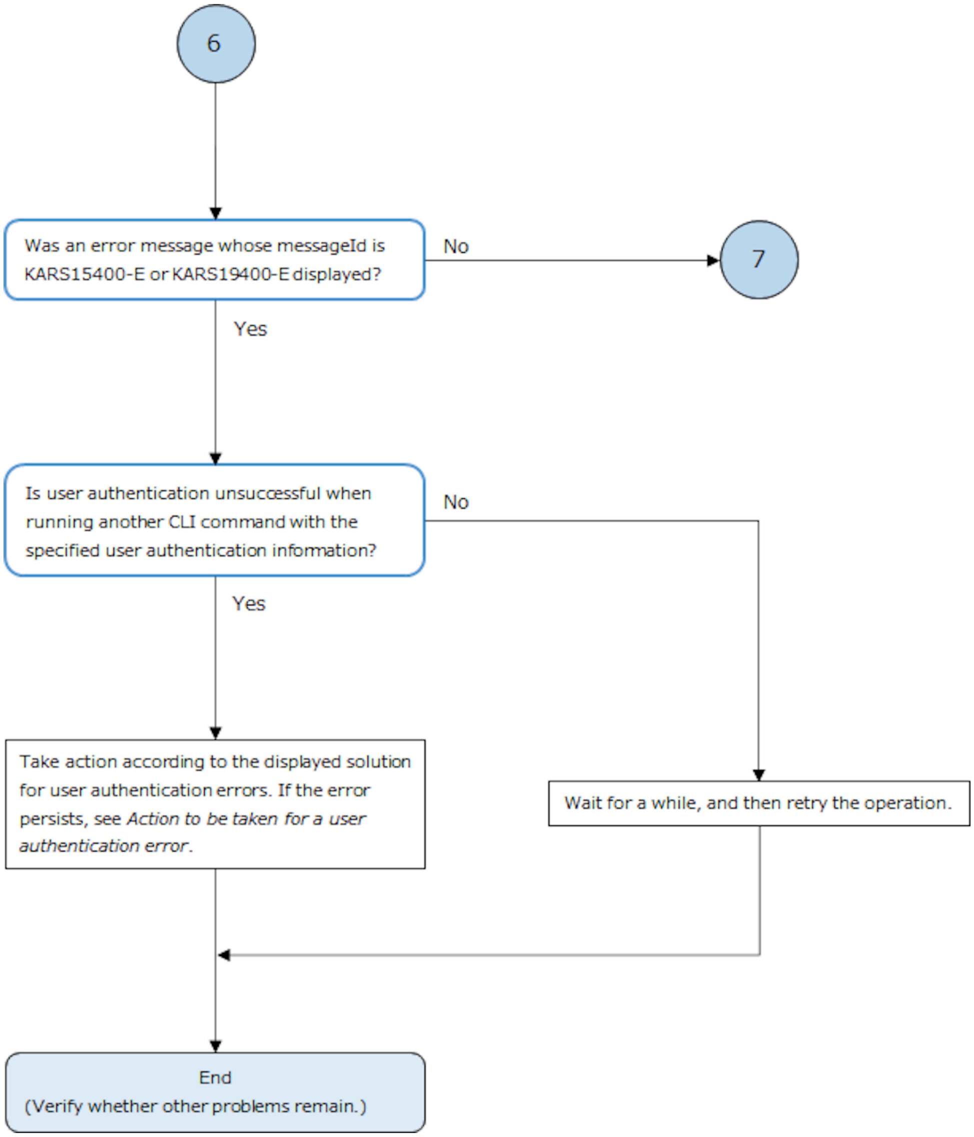
6
Was an error message whose messageId is KARS15400-E or KARS19400-E displayed?
No
7
Yes
Is user authentication unsuccessful when running another CLI command with the specified user authentication information?
No
Yes
Take action according to the displayed solution for user authentication errors. If the error persists, see *Action to be taken for a user authentication error*.
Wait for a while, and then retry the operation.
End
(Verify whether other problems remain.)

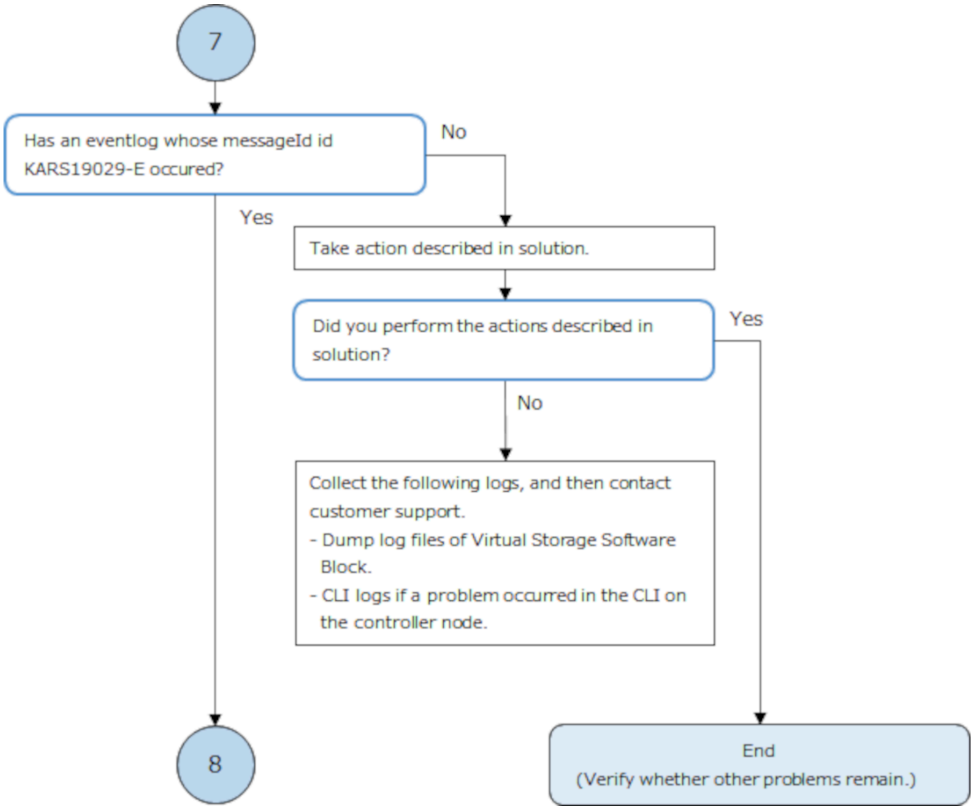
7
Has an eventlog whose messageId id KARS19029-E occured?
No
Yes
Take action described in solution.
Did you perform the actions described in solution?
Yes
No
Collect the following logs, and then contact customer support.
- Dump log files of Virtual Storage Software Block.
- CLI logs if a problem occurred in the CLI on the controller node.
8
End
(Verify whether other problems remain.)

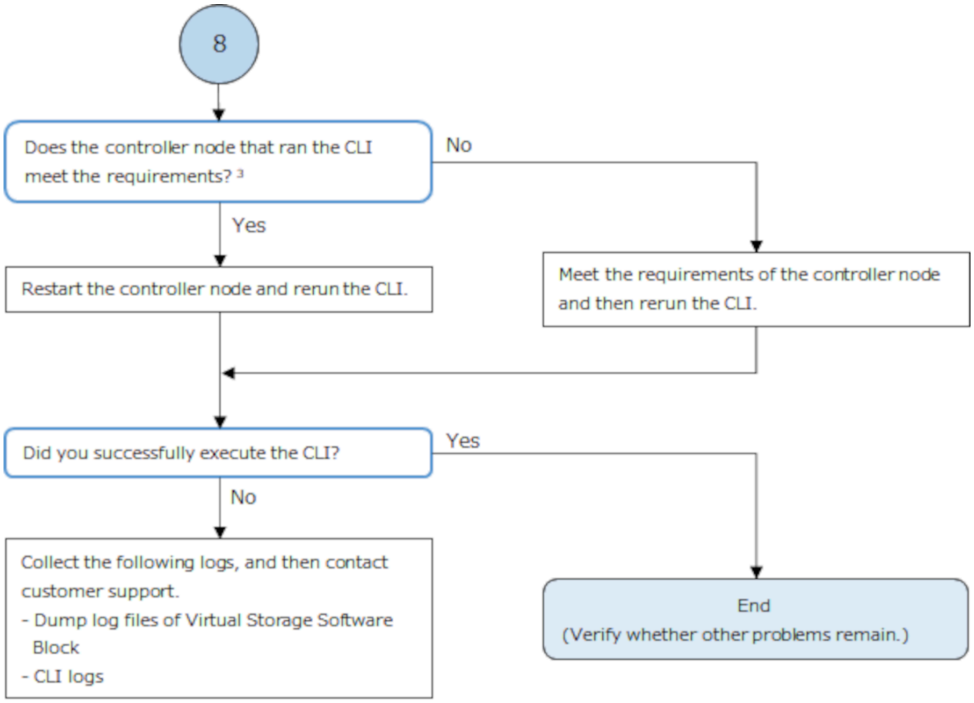
8
Does the controller node that ran the CLI meet the requirements? 3
No
Yes
Restart the controller node and rerun the CLI.
Meet the requirements of the controller node and then rerun the CLI.
Did you successfully execute the CLI?
Yes
No
Collect the following logs, and then contact customer support.
- Dump log files of Virtual Storage Software Block
- CLI logs
End
(Verify whether other problems remain.)

3. Contact customer support for controller node requirements.

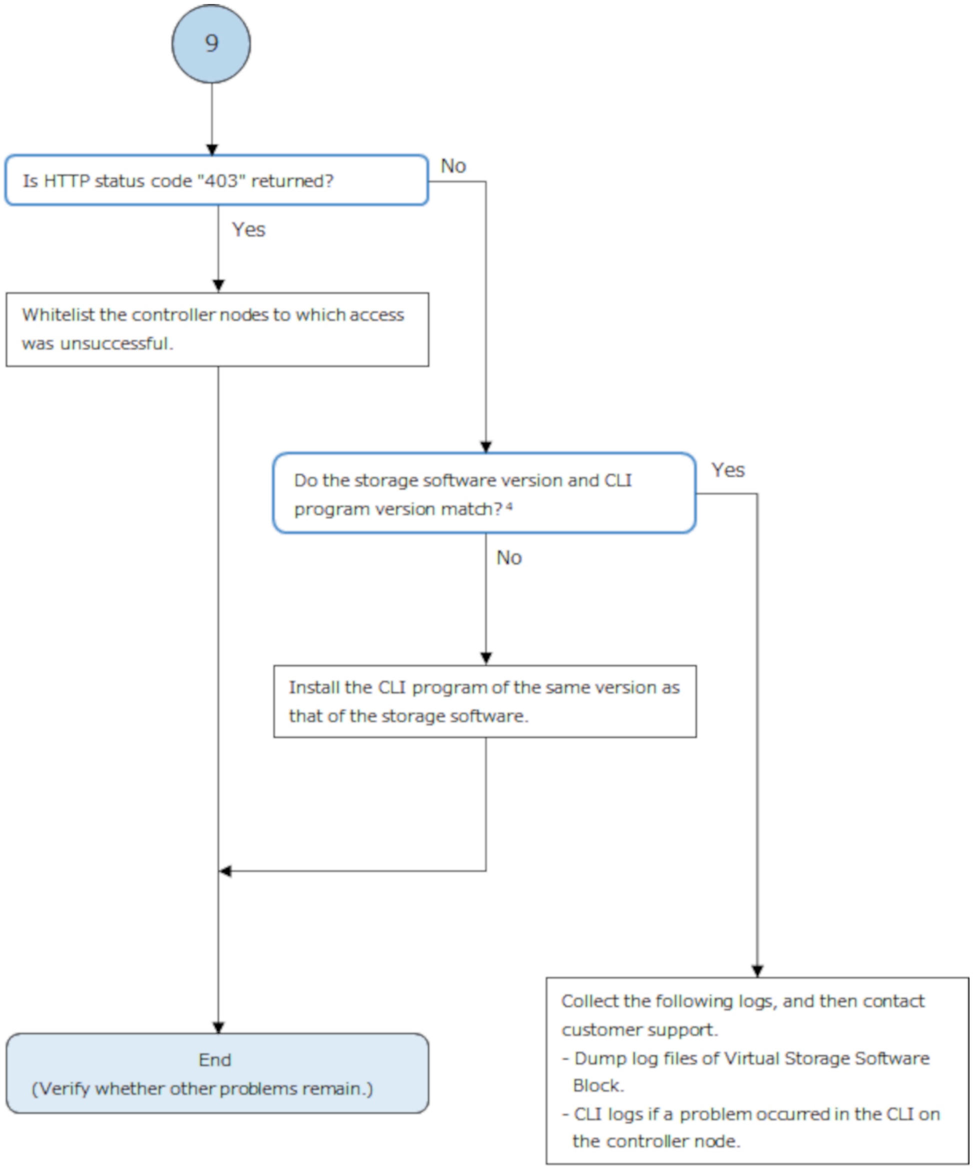
9
Is HTTP status code "403" returned?
No
Yes
Whitelist the controller nodes to which access was unsuccessful.
Do the storage software version and CLI program version match?4
Yes
No
Install the CLI program of the same version as that of the storage software.
End
(Verify whether other problems remain.)
Collect the following logs, and then contact customer support.
- Dump log files of Virtual Storage Software Block.
- CLI logs if a problem occurred in the CLI on the controller node.

4. To verify the storage software version, contact customer support. You can also verify the CLI program version by using the master command option --version.

Actions to be taken for CLI system errors

If an error occurs before the CLI process finishes starting when you run the CLI command, you may receive an error message from the system. If this is the case, check the last line of the error message.

If any of the reviews match the *Output string* in the table below, follow the *Action* section.

Output string	Cause	Action
KeyboardInterrupt	Keyboard input interrupted the process.	Rerun the command.
ImportError: No module named 'xxx'	The required package is not installed.	Reinstall the CLI program.
MemoryError	Processing has been interrupted due to low memory.	(Virtual machine) Restart the controller node or the maintenance node to free up memory, and then rerun the command. If the problem persists, collect the maintenance node logs if a problem occurred in the CLI on the maintenance node and CLI logs if a problem occurred in the CLI on the controller node. Then, contact customer support. (Bare metal) Restart the controller node to free up memory and then rerun the command. If the problem persists, collect the CLI logs from the controller node and contact customer support.
Messages other than above	There is an internal error.	(Virtual machine) Collect the maintenance node logs if a problem occurred in the CLI on the maintenance node and CLI logs if a problem occurred in the CLI on the controller node. Then, contact customer support. (Bare metal) Collect CLI logs on the controller node, and then contact customer support.

Action to be taken when the same InitiatorName (iSCSI name) is used in two compute nodes

When an initiator name (iSCSI name) is duplicated in the iSCSI connection, no volumes are visible from a compute node having the corresponding initiator name. Take the following actions according to the OS you are using.

In Windows:

Procedure

1. In the compute node that has an initiator with a duplicated initiator name (iSCSI name), open [Control Panel], [Administrative tools], and then the [iSCSIInitiator] settings window.

   In the [Targets] tab, select the target which is connected to Virtual Storage Software block, and disconnect the connection by clicking [Disconnect].

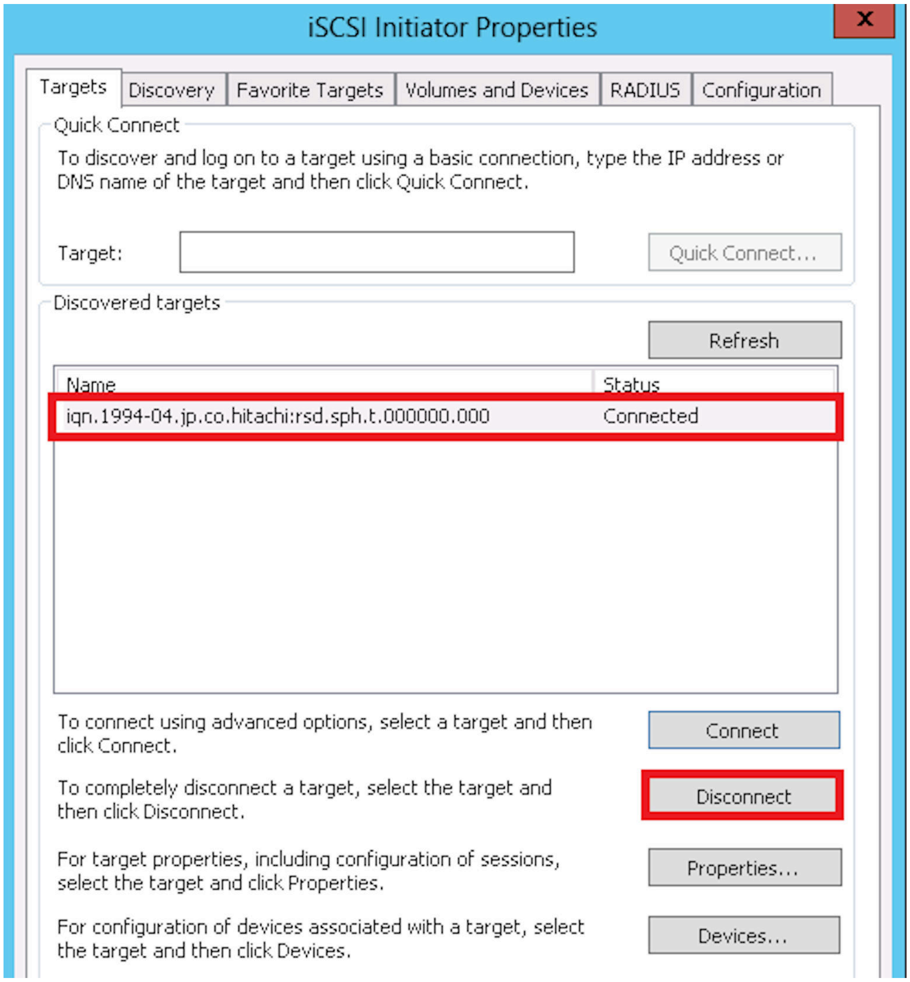

To discover and log on to a target using a basic connection, type the IP address or DNS name of the target and then click Quick Connect.

Name	Status
iqn.1994-04.jp.co.hitachi:rsd.sph.t.000000.000	Connected

To connect using advanced options, select a target and then click Connect.

To completely disconnect a target, select the target and then click Disconnect.

For target properties, including configuration of sessions, select the target and click Properties.

For configuration of devices associated with a target, select the target and then click Devices.

2. In the [Configuration] tab, click [Change] to change the initiator name (iSCSI name).

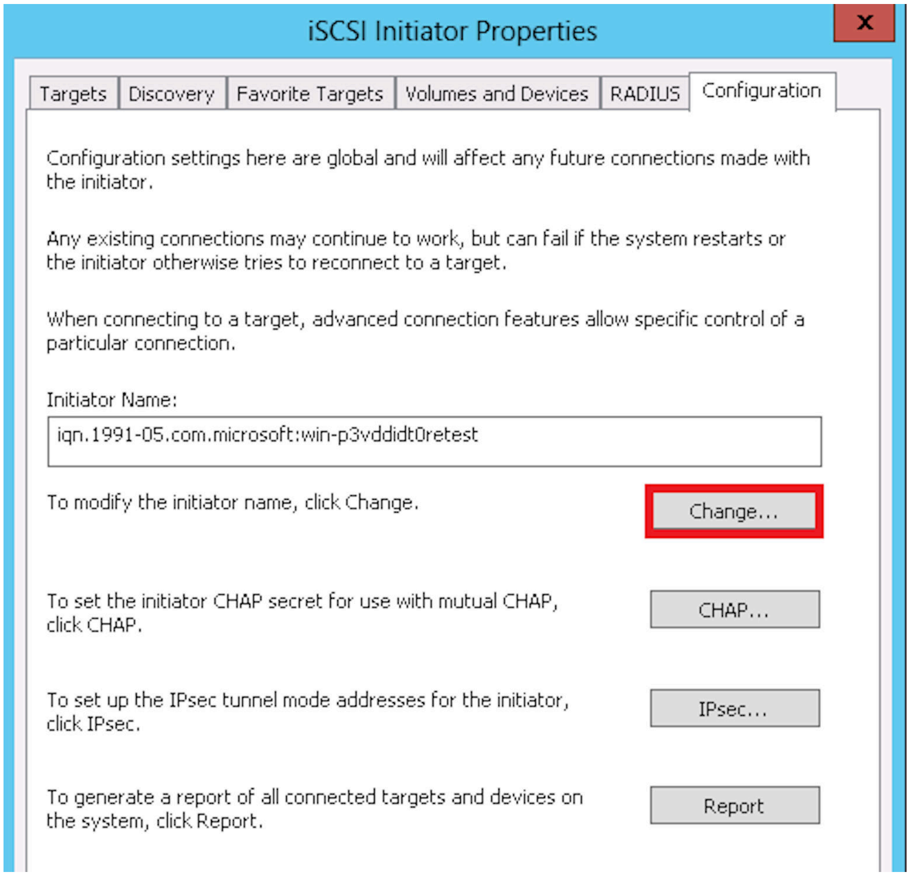

Configuration settings here are global and will affect any future connections made with the initiator.

Any existing connections may continue to work, but can fail if the system restarts or the initiator otherwise tries to reconnect to a target.

When connecting to a target, advanced connection features allow specific control of a particular connection.

Initiator Name:

iqn.1991-05.com.microsoft:win-p3vddidt0retest

To modify the initiator name, click Change. Change...

To set the initiator CHAP secret for use with mutual CHAP, click CHAP. CHAP...

To set up the IPsec tunnel mode addresses for the initiator, click IPsec. IPsec...

To generate a report of all connected targets and devices on the system, click Report. Report

3. Enter a unique initiator name (iSCSI name) in the New initiator name field, and then click [OK].

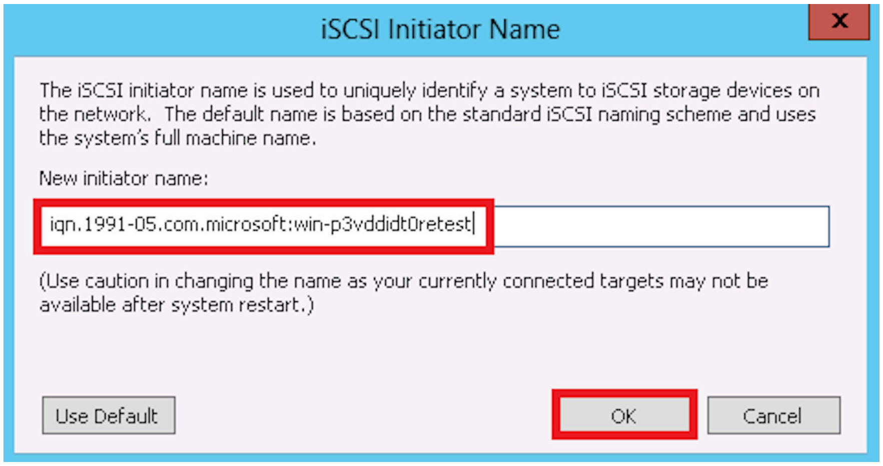

The iSCSI initiator name is used to uniquely identify a system to iSCSI storage devices on the network. The default name is based on the standard iSCSI naming scheme and uses the system's full machine name.

New initiator name:

iqn.1991-05.com.microsoft:win-p3vddidt0retest

(Use caution in changing the name as your currently connected targets may not be available after system restart.)

4. In the [Targets] tab, select the target which is connected to Virtual Storage Software block, and reconnect to Virtual Storage Software block by clicking [Connect].

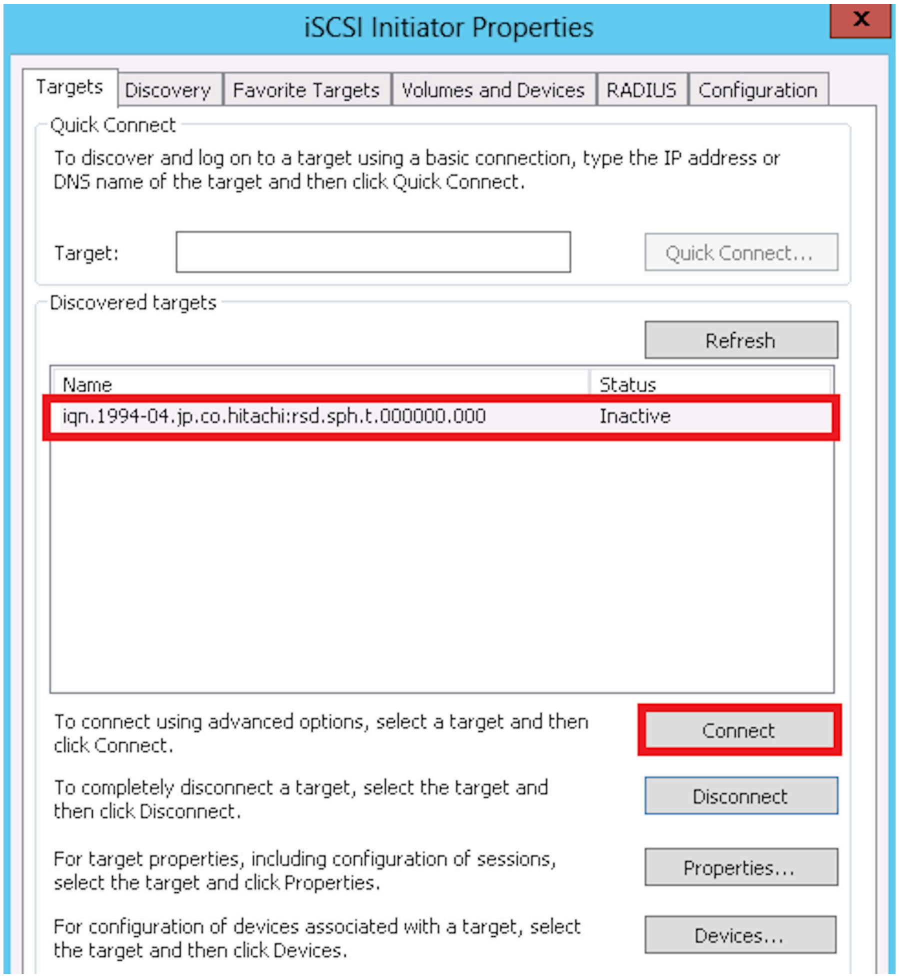

To discover and log on to a target using a basic connection, type the IP address or DNS name of the target and then click Quick Connect.

Name	Status
iqn.1994-04.jp.co.hitachi:rsd.sph.t.000000.000	Inactive

To connect using advanced options, select a target and then click Connect.

To completely disconnect a target, select the target and then click Disconnect.

For target properties, including configuration of sessions, select the target and click Properties.

For configuration of devices associated with a target, select the target and then click Devices.

5. After reconnection, verify that volumes are visible from the compute node.

   If volumes are still not visible after reconnection, check for other failures by following *If a failure status is detected in the storage cluster*.

In Linux:

If the target is connected to Virtual Storage Software block via a means other than openiscsi, contact customer support which handles the connection method.

Procedure

1. Verify the connection information on the compute node that has a duplicated initiator name (iSCSI name).

   ```
   # iscsiadm -m node
   ```

   Of the displayed information, record the targetname and portalname.

   After the command is run, a targetname and portalname are indicated as shown in the following example display.

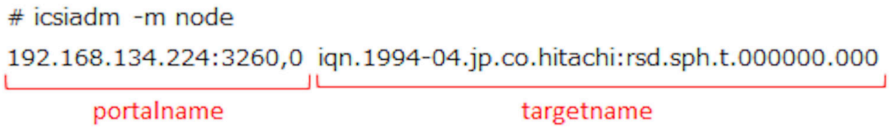

```
# icsiadm -m node
192.168.134.224:3260,0 iqn.1994-04.jp.co.hitachi:rsd.sph.t.000000.000
```

2. Disconnect from the initiator that has a duplicated initiator name (iSCSI name). For <*targetname*> and <*portalname*>, specify the targetname and portalname recorded in step 1.

   ```
   # iscsiadm -m node -T <targetname> --portal <portalname> --logout
   ```

3. Edit the /etc/iscsi/initiatorname.iscsi file using an editor such as vi commands.

   ```
   # vi /etc/iscsi/initiatorname.iscsi
   ```

4. In the file, change the "InitiatorName = <*duplicated-initiator-name* (*iSCSI-name*)>" line to a unique initiator name (iSCSI name)

   ```
   InitiatorName= <ubique-initiator-name (iSCSI-name)>
   ```

5. Reconnect the initiator. For <*targetname*> and <*portalname*>, specify the targetname and portalname recorded in step 1.

   ```
   # iscsiadm -m node -T <targetname> --portal <portalname> --login
   ```

6. After reconnection, verify that volumes are visible from the compute node.

   If volumes are still not visible after reconnection, check for other failures by following *If a failure status is detected in the storage cluster*.

In ESXi:

Procedure

1. In VMware vCenter Server or VMware ESXi of the compute node (VMware ESXi) with a duplicated initiator name (iSCSI name), open [Storage], [Adapters], and then the [Configure iSCSI] settings window.

   Record the target information displayed in [Static targets] and [Dynamic targets], and delete all of them.

In ESXi:

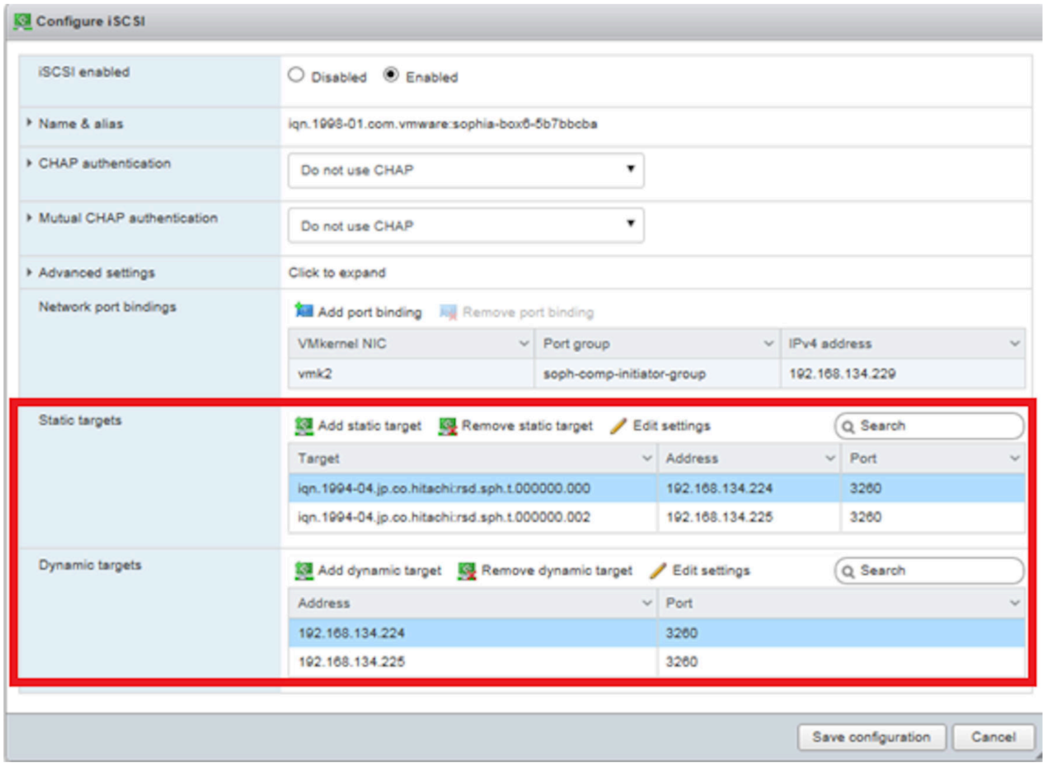

Configure iSCSI

iSCSI enabled	○ Disabled ● Enabled
Name & alias	iqn.1998-01.com.vmware:sophia-box6-5b7bbcba
CHAP authentication	Do not use CHAP
Mutual CHAP authentication	Do not use CHAP
Advanced settings	Click to expand
Network port bindings	Add port binding, Remove port binding

VMkernel NIC	Port group	IPv4 address
vmk2	soph-comp-initiator-group	192.168.134.229

Static targets: Add static target, Remove static target, Edit settings, Search

Target	Address	Port
iqn.1994-04.jp.co.hitachi:rsd.sph.t.000000.000	192.168.134.224	3260
iqn.1994-04.jp.co.hitachi:rsd.sph.t.000000.002	192.168.134.225	3260

Dynamic targets: Add dynamic target, Remove dynamic target, Edit settings, Search

Address	Port
192.168.134.224	3260
192.168.134.225	3260

2. By using an ESXCLI command, get the Software iSCSI Adapter information and record the adapter name.

```
# esxcli iscsi adapter list
```

After the command is run, an adapter name is indicated as shown in the following example display.

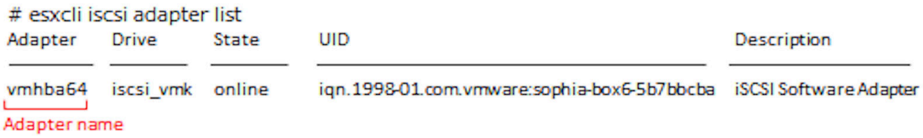

```
# esxcli iscsi adapter list
Adapter   Drive      State    UID                                          Description
--------  ---------  -------  -------------------------------------------  ----------------------
vmhba64   iscsi_vmk  online   iqn.1998-01.com.vmware:sophia-box6-5b7bbcba  iSCSI Software Adapter
```

3. By using an ESXCLI command, change the initiator name (iSCSI name) to a unique one. For <*AdapterName*>, specify the adapter name recorded in step 2.

```
#esxcli iscsi adapter set -n <unique-InitiatorName> -A <AdapterName>
```

4. In VMware vCenter Server or VMware ESXi, open [Storage], [Adapters], and then the [Configure iSCSI] settings window, and then register target information in [Static targets] and [Dynamic targets]. As the target information to be registered, re-specify the information recorded in step 1.

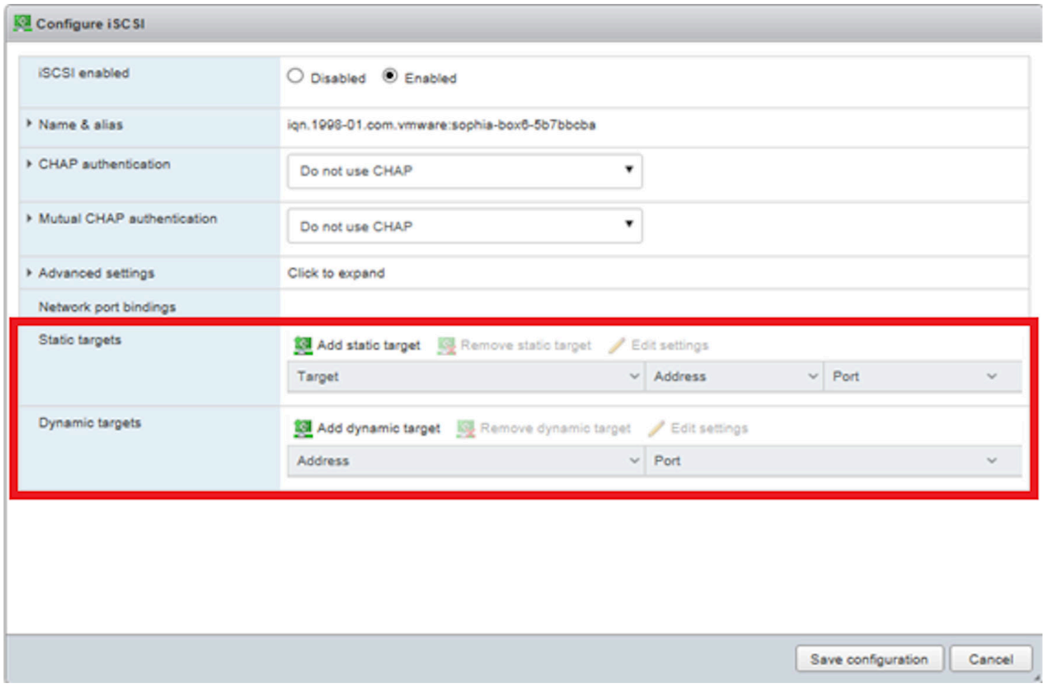

Configure iSCSI

iSCSI enabled	○ Disabled ● Enabled
Name & alias	iqn.1998-01.com.vmware:sophia-box6-5b7bbcba
CHAP authentication	Do not use CHAP
Mutual CHAP authentication	Do not use CHAP
Advanced settings	Click to expand
Network port bindings	
Static targets	Target / Address / Port
Dynamic targets	Address / Port

5. After reconnection, verify that volumes are visible from the compute node.

If volumes are still not visible after reconnection, check for other failures by following *If a failure status is detected in the storage cluster*.

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.

Action to be taken when "Alerting" is shown in the license health status

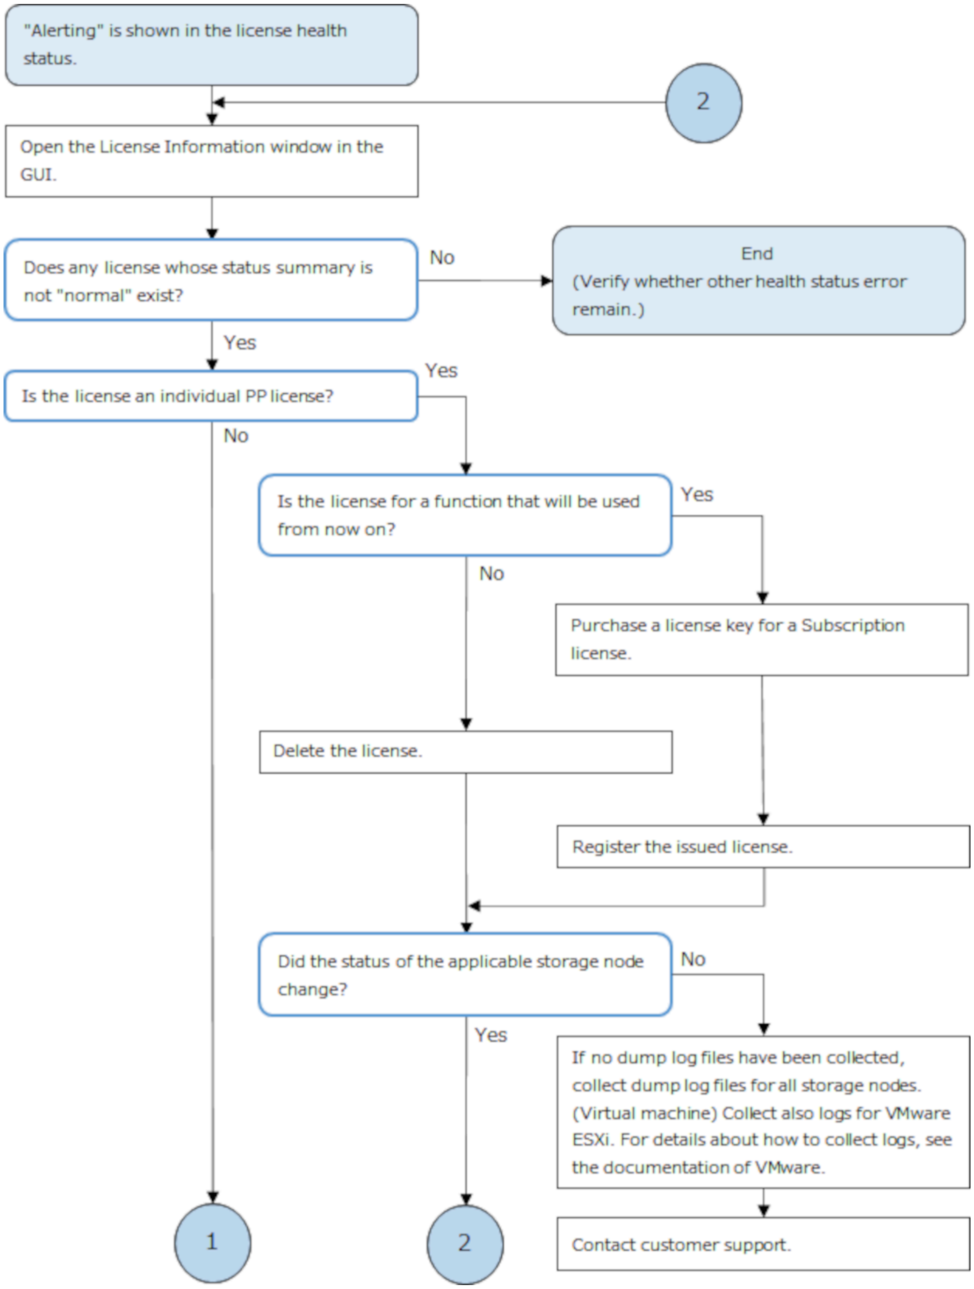
"Alerting" is shown in the license health status.
2
Open the License Information window in the GUI.
Does any license whose status summary is not "normal" exist?
No
End (Verify whether other health status error remain.)
Yes
Is the license an individual PP license?
Yes
No
Is the license for a function that will be used from now on?
Yes
No
Purchase a license key for a Subscription license.
Delete the license.
Register the issued license.
Did the status of the applicable storage node change?
No
Yes
If no dump log files have been collected, collect dump log files for all storage nodes. (Virtual machine) Collect also logs for VMware ESXi. For details about how to collect logs, see the documentation of VMware.
Contact customer support.
1
2

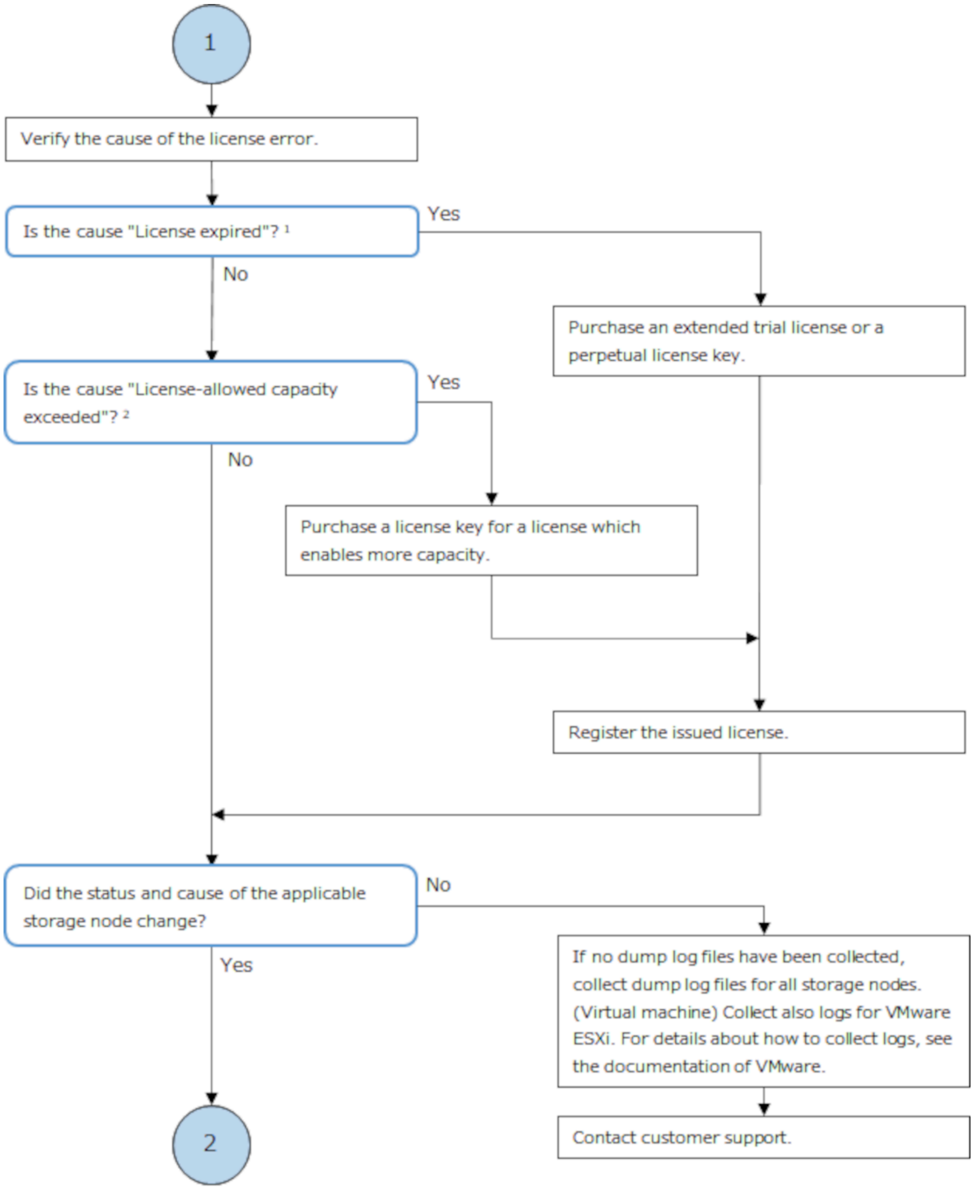
1
Verify the cause of the license error.
Is the cause "License expired"? 1
Yes
No
Purchase an extended trial license or a perpetual license key.
Is the cause "License-allowed capacity exceeded"? 2
Yes
No
Purchase a license key for a license which enables more capacity.
Register the issued license.
Did the status and cause of the applicable storage node change?
No
Yes
If no dump log files have been collected, collect dump log files for all storage nodes. (Virtual machine) Collect also logs for VMware ESXi. For details about how to collect logs, see the documentation of VMware.
Contact customer support.
2

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.	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. ▪ 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 status[1] 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 status[1] 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 status[1] 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"

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.
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.
Expansion Failed	All statuses	Volume expansion ended abnormally.	Normal	Expand the volume which ended abnormally again without specifying a value for the addictionalCapacity.

STATUS	SNAPSHOT STATUS	Description	VOLUME TYPE	Action[*]
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.				

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