Hitachi Virtual Storage Software

1.12

Block Audit Log User Guide

Hitachi Virtual Storage Software block provides audit logs of events that occur in the storage nodes. The VSS block audit logs can be viewed on the syslog server and also downloaded using the CLI or the REST API.

MK-91VSSB004-03
February 2023

Contents

Preface....4
Intended audience....4
Product version....4
Release notes....4
Changes made in this revision....5
Document conventions....5
Accessing product documentation....6
Getting help....6
Comments....7
Chapter 1: Overview....8
Overview of audit logs....8
Viewing audit logs....9
Chapter 2: Format of audit logs....11
Structure of audit logs....11
Header part....14
Structured data part....15
Message part....15
Chapter 3: Information in the audit logs....19
Information that is output in audit logs....19
Chapter 4: Managing audit logs....25
Overview of audit logs....25
Downloading an audit log to the controller node (CLI or REST API)....26
Obtaining Syslog transfer settings of audit logs (CLI or REST API)....27
Editing Syslog transfer settings of audit logs (CLI or REST API)....28
Monitoring the audit logs transferred to the Syslog server....30
Glossary....31

Preface

Hitachi Virtual Storage Software block provides audit logs of events that occur in the storage nodes. The VSS block audit logs can be viewed on the syslog server and also downloaded using the CLI or the REST API.

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 operating Virtual Storage Software block.

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

- Knowledge of Windows and Linux
- Knowledge of REST API and CLI for Virtual Storage Software block

Product version

This document revision applies to Virtual Storage Software block version 1.12 (01.12.*0x*.*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

Updated the document for v1.12.

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`
< > 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

Overview of audit logs

Audit logs are records of the operations performed on the storage cluster. Audit logs allow you to verify "when" "who" did what" to see if each operation complies with audit standards such as laws, regulations, industrial standards, and in-house stipulations.

Audit logs are subject to the following capacity limits:

- Maximum number of characters per line (event): 8,192 bytes (including line feed code). If the maximum number of characters is exceeded, the excess characters are truncated.
- Maximum number of events: 750,000
- Maximum capacity: 5,859 MiB

If the maximum capacity or the maximum number of events for an audit log has been reached, old information will be overwritten with new information. To prevent old information from being lost, you can configure the system to transfer audit logs to the syslog server.

(Virtual machine) Hitachi Virtual Storage Software block (VSSB) provides an audit log of events that occur in the VMs of storage nodes. For events that occur in non-VM hardware or software (for example, VMware ESXi on which VMs run, physical servers, and switches), verify the audit log for the hardware or software.

(Bare metal) Virtual Storage Software block creates audit log data for only the events that occurred on storage nodes. For the audit log data for other events that occurred in software or hardware other than storage nodes (such as a physical server or switch), see the audit log for the relevant software or hardware.

Viewing audit logs

Audit logs are stored in the cluster primary node (primary). As illustrated in the following figure, you can obtain the logs by either of the following means:

- Configure transfer of audit logs to the syslog server and obtain the logs from the transfer destination syslog server.
  - Audit logs are transferred to the syslog server as text data.
  - Rsyslog 8 is supported as a syslog server.
  - Audit logs created after syslog transfer settings were made are transferred to the syslog server. Audit logs created before syslog transfer settings were made are not transferred to the syslog server.
  - For details about syslog transfer settings, see Editing Syslog transfer settings of audit logs (CLI or REST API) (on page 28).
- Download audit log files compiled by using the REST API or CLI.
  - Audit log files are in the csv format.
  - For the method to compile an audit log into a file by using the API or CLI and then download it, see Downloading an audit log to the controller node (CLI or REST API) (on page 26).

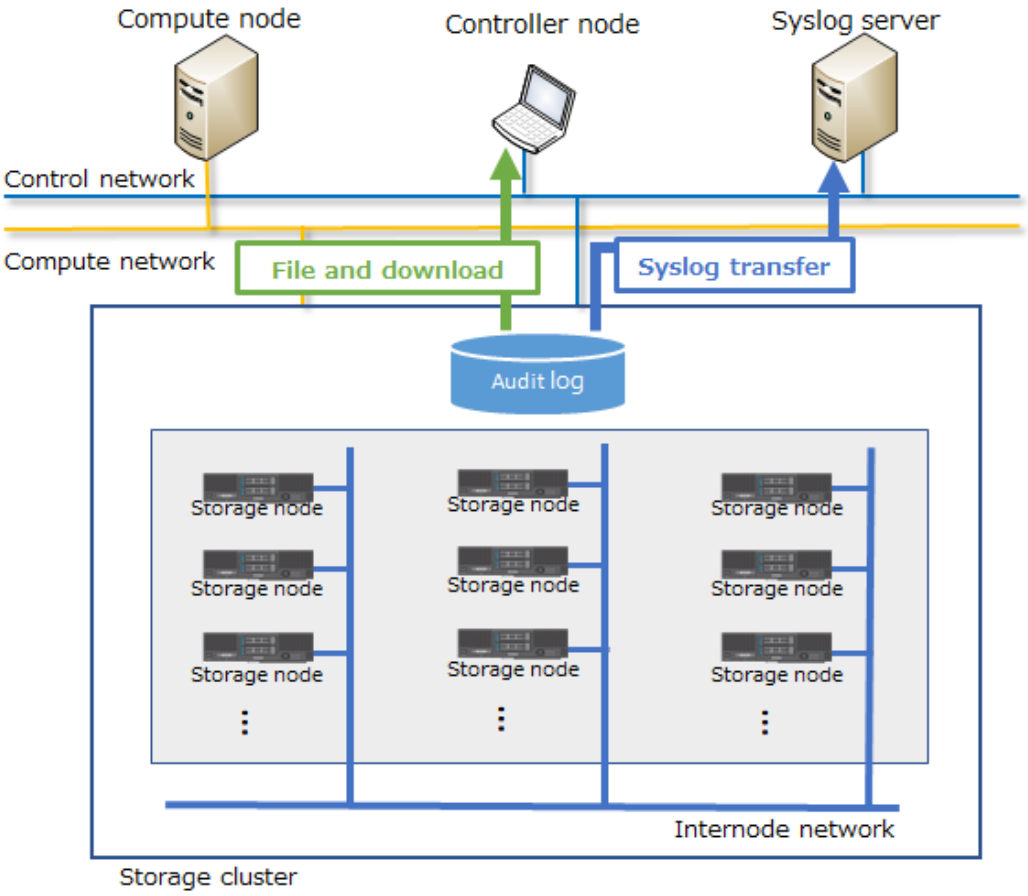
Compute node
Controller node
Syslog server
Control network
Compute network
File and download
Syslog transfer
Audit log
Storage node
Storage node
Storage node
Storage node
Storage node
Storage node
Storage node
Storage node
Storage node
Internode network
Storage cluster

Chapter 2: Format of audit logs

Structure of audit logs

Audit logs of Virtual Storage Software block are encoded in UTF-8.

The following two types of structures are used in audit logs:

- Text data transferred to the syslog server
- Audit log files in the csv format

Structure of text data transferred to the syslog server

The following shows an example of text data transferred to the syslog server. Each triangle (△) in the example indicates a space. The structure of text data transferred to the syslog server complies with RFC 5424.

Structure of audit log files in the csv format

The following shows an example of an audit log file in the csv format. Each triangle (△) in the example indicates a space, and each down arrow (↓) indicates a line feed.

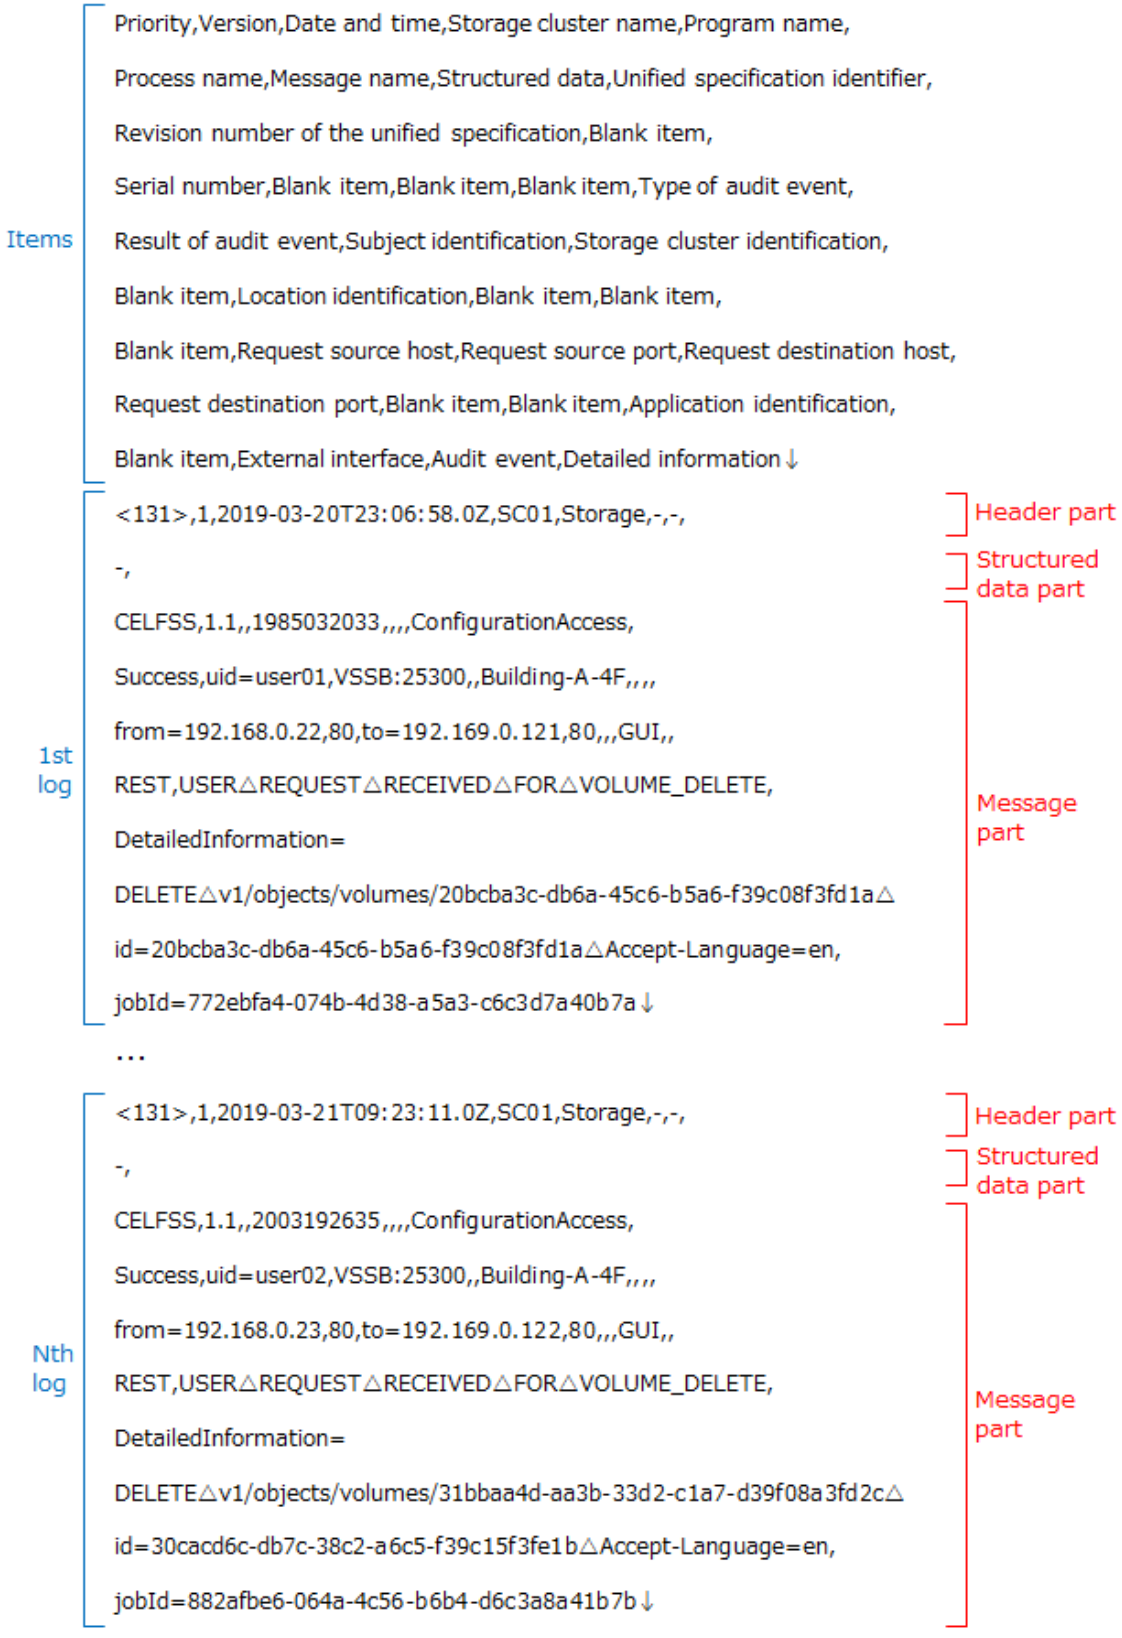
Items
Priority,Version,Date and time,Storage cluster name,Program name,
Process name,Message name,Structured data,Unified specification identifier,
Revision number of the unified specification,Blank item,
Serial number,Blank item,Blank item,Blank item,Type of audit event,
Result of audit event,Subject identification,Storage cluster identification,
Blank item,Location identification,Blank item,Blank item,
Blank item,Request source host,Request source port,Request destination host,
Request destination port,Blank item,Blank item,Application identification,
Blank item,External interface,Audit event,Detailed information↓
1st log
<131>,1,2019-03-20T23:06:58.0Z,SC01,Storage,-,-,
-,
CELFSS,1.1,,1985032033,,,,ConfigurationAccess,
Success,uid=user01,VSSB:25300,,Building-A-4F,,,,
from=192.168.0.22,80,to=192.169.0.121,80,,,GUI,,
REST,USER△REQUEST△RECEIVED△FOR△VOLUME_DELETE,
DetailedInformation=
DELETE△v1/objects/volumes/20bcba3c-db6a-45c6-b5a6-f39c08f3fd1a△
id=20bcba3c-db6a-45c6-b5a6-f39c08f3fd1a△Accept-Language=en,
jobId=772ebfa4-074b-4d38-a5a3-c6c3d7a40b7a↓
Header part
Structured data part
Message part
...
Nth log
<131>,1,2019-03-21T09:23:11.0Z,SC01,Storage,-,-,
-,
CELFSS,1.1,,2003192635,,,,ConfigurationAccess,
Success,uid=user02,VSSB:25300,,Building-A-4F,,,,
from=192.168.0.23,80,to=192.169.0.122,80,,,GUI,,
REST,USER△REQUEST△RECEIVED△FOR△VOLUME_DELETE,
DetailedInformation=
DELETE△v1/objects/volumes/31bbaa4d-aa3b-33d2-c1a7-d39f08a3fd2c△
id=30cacd6c-db7c-38c2-a6c5-f39c15f3fe1b△Accept-Language=en,
jobId=882afbe6-064a-4c56-b6b4-d6c3a8a41b7b↓
Header part
Structured data part
Message part

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 descriptions in the following table.

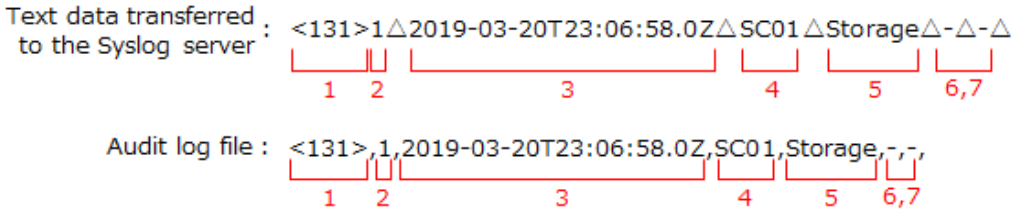

No.	Item	Description
1	Priority	▪ <131>: This number is shown if "Result of audit event" is "Failed", or if "Result of audit event" is "Occurred" and the user is requested to take action. ▪ <134>: This number is shown if "Result of audit event" is "Success", or if "Result of audit event" is "Occurred" and the user is not requested to take action.
2	Version	Fixed to "1".
3	Date and time	Date and time when an event under audit occurred. The output format is "YYYY-MM-DDThh:mm:ss.s±hh:mm". "YYYY-MM-DDThh:mm:ss.s" indicates a date and time in coordinated universal time (UTC). (YYYY: Year, MM: Month, DD: Date, hh: Hour, mm: Minute, ss.s: Second (up to the first decimal place)) "±hh:mm" indicates a time difference from UTC. ▪ "+hh:mm" indicates that the output date and time is hh:mm ahead of UTC. ▪ "-hh:mm" indicates that the output date and time is hh:mm behind UTC. ▪ "Z" means that the output date and time are the same as UTC.
4	Storage cluster name	Name of the storage cluster as the target of audit logs.
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.

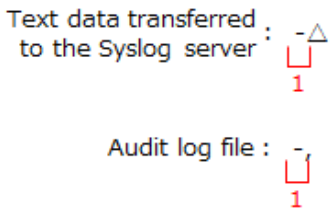

No.	Item	Description
1	Structured data	Fixed to "-".

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.

Depending on the type of audit log, the value of a specific item might be empty. (Items that become empty characters differ depending on the type of audit log.)

Note:

In syslog transfer, "%xEF.BB.BF" is added as the byte order mark (BOM) to the beginning of the message part (before "CELFSS" in the following figure).

No.	Item	Description
1	Unified specification identifier	Fixed to "CELFSS". This ID indicates that this audit log complies with the unified specification for Hitachi storage security products.
2	Revision number of the unified specification	Fixed to "1.1".
3	Blank item	Empty character.
4	Serial number	Serial number in the audit log (0000000001 to 9999999999) Some events might not be assigned serial numbers depending on the time the audit log was filed.
5	Blank item (for three items)	Empty character.
6	Type of audit event	Type of event under audit.
7	Result of audit event	Result of an event under audit: ▪ "Success": The event succeeded. ▪ "Failed": The event did not succeed. ▪ "Occurred": The event occurred.

No.	Item	Description
8	Subject identification	Subject that caused the event under audit. The output content differs depending on the cause of the event: ▪ User name (e.g., uid=userID): The event was generated through the REST API, CLI, GUI or console interface execution. ▪ "<*System*>": The event is generated through notification from the storage cluster. ▪ iSCSI name (e.g., iSN=*iSCSI-name*): The event was generated through access from the compute node.
9	Storage cluster identification	Product type identifier. The following format is used: <*model-name*>: <*Internal-ID-(internalId)-of-the-storage-cluster*> <*model-name*> is replaced by either of the following: ▪ "VSSB": Indicates the virtual machine model. ▪ "VSSBB1": Indicates the bare metal model.
10	Blank item	Empty character.
11	Location identification	Location of the audit log output source, such as the installation location of equipment.
12	Blank item (for three items)	Empty character.
13	Request source host	Host name or IP address of the request source of the event under audit.*
14	Request source port	Request source port of the event under audit.*
15	Request destination host	Host name or IP address of the request destination of the event under audit.*
16	Request destination port	Request destination port of the event under audit.*
17	Blank item (for two items)	Empty character.
18	Application identification	Name of the application that caused an event generating an audit log.
19	Blank item	Empty character.

No.	Item	Description
20	External interface	Name of the external interface that caused an event generating an audit log. ▪ "REST": Access through the REST API (including GUI and CLI) ▪ "CONSOLE": Access through the console interface ▪ "HOST": Access from the compute node
21	Audit event	Name of the executed operation or event.
22	Detailed information	Detailed information about the audit event.

* When you have created, downloaded, or deleted a dump log file by using the GUI, an audit log whose "Type of audit event" is ConfigurationAccess and that has USER REQUEST RECEIVED FOR followed by either of the following strings, and whose request source is localhost/127.0.0.1/0:0:0:0:0:0:0:1/::1, is displayed.

- DUMP_FILE_CREATE_FILE
- DUMP_FILE_DOWNLOAD
- DUMP_FILE_DELETE

If you want to confirm the request source for creation, downloading, and deletion of a dump log file by using the GUI, view an audit log that has USER REQUEST RECEIVED FOR followed by either of the following strings.

- Creating a dump log file by using the GUI: STORAGE_NODE_DUMP_FILE_CREATE_FILE
- Downloading a dump log file by using the GUI: STORAGE_NODE_DUMP_FILE_DOWNLOAD
- Deleting a dump log file by using the GUI: STORAGE_NODE_DUMP_FILE_DELETE

Chapter 3: Information in the audit logs

Information that is output in audit logs

"Type of audit event" and "Audit event" in the message part of an audit log provide summary information about the audit log. "Detailed information" is output depending on the event.

The following are the types of audit events:

- AnomalyEvent: Indicates that an abnormality occurred (for example, a threshold was exceeded).
- Authentication: Indicates that authentication or authorization processing was performed.
- ConfigurationAccess: Indicates the occurrence of a configuration change and the associated change in job status.
- Maintenance: Indicates that a maintenance operation was performed.
- StartStop: Indicates a start or stop of the storage cluster.

The following table shows audit events and their description and detailed information categorized by each type of audit event.

Type of audit event: AnomalyEvent

Audit event	Description	Detailed information
AUDIT LOG REACHED THRESHOLD	The number of untransferred audit logs exceeded 70% of the maximum limit.	None
AUDIT LOG REACHED UPPER LIMIT	The maximum number of untransferred audit logs has been reached.	None
CREATING AUDIT LOG FAILED	Indicates that the generated audit event includes information about violation of the unified specification for Hitachi storage security products. If this audit event name is output, contact customer support.	The following information about the generated audit event is output: ▪ Priority ▪ Type of audit event

Audit event	Description	Detailed information
		▪ Result of audit event ▪ Subject identification ▪ Request source host ▪ Request source port ▪ Request destination host ▪ Request destination port ▪ Application identification ▪ External interface ▪ Audit event ▪ Detailed information

Type of audit event: Authentication

Audit event	Description	Detailed information
AUTHORIZATION FAILED	Indicates that authentication did not succeed.	For ticket authentication, the upper 8 digits of the SHA-256 hash of the ticket are output.
PASSWORD AUTHENTICATION EXECUTED	Indicates that Basic authentication was performed.	(Virtual machine) None (Bare metal) If authentication to the console interface is performed, the ID of the storage node is output. Example: StorageNodeId=e839d554-56e6-40a6-929a-6a8c2a0d6e3d
SESSION AUTHENTICATION EXECUTED	Indicates that session authentication did not succeed. Not output when session authentication succeeded.	None

Audit event	Description	Detailed information
TICKET AUTHENTICATION EXECUTED	Indicates that ticket authentication was performed.	The upper 8 digits of the SHA-256 hash of the ticket are output.
CHAP AUTHENTICATION	Indicates that CHAP authentication was performed.	None

Type of audit event: ConfigurationAccess

Audit event	Description	Detailed information
JOB {STARTED| SUCCEEDED| FAILED| STOPPED} FOR <*CLI-subcommand-name*>	Indicates that the job is in one of the following states: ▪ STARTED: The job was started. ▪ SUCCEEDED: The job succeeded. ▪ FAILED: The job did not succeed. ▪ STOPPED: The job was stopped. Although <*CLI-subcommand-name*> indicates the operation that started the job, it does not necessarily indicate an operation through CLI. <*CLI-subcommand-name*> also covers any operations through REST API or GUI.[1]	The job ID (JobID) of the applicable event is output.[2] Example: JobID=772ebfa4-074b-4d38-a5a3-c6c3d7a40b7a
MODIFY CONFIGURATION EXECUTED	Indicates that the configuration information has been changed or set.	None
USER REQUEST RECEIVED FOR <*CLI-subcommand-name*>	Indicates that the storage cluster accepted a user operation. Although <*CLI-subcommand-name*> indicates the details of a user operation, it does not necessarily indicate an operation through CLI. <*CLI-subcommand-name*> also covers any operations through REST API or GUI.[1]	The input information (REST API name + input parameter) received by the REST server and job ID are output.[2, 3] Parameters that are not specified by the user are displayed with default values. Internal parameters may also be displayed. The input information and the job ID are delimited by a comma.

Audit event	Description	Detailed information
		Example:[4] DetailedInformation=DELETE v1/objects/volumes/20bcba3c-db6a-45c6-b5a6-f39c08f3fd1a id=20bcba3c-db6a-45c6-b5a6-f39c08f3fd1a Accept-Language=en, jobId=772ebfa4-074b-4d38-a5a3-c6c3d7a40b7a
CONTROL PORT SETTING EXECUTED	Indicates that control port configuration was performed through the console interface.	The ID of the target storage node is output. Example: StorageNodeId=e839d554-56e6-40a6-929a-6a8c2a0d6e3d

1. Depending on the operation, a CLI subcommand name that is not described in the *Hitachi Virtual Storage Software Block CLI Reference* might be displayed. The subcommand names and their meanings are as follows:

- STORAGE_ADD_NODE: Adding storage nodes
- CONFIGURATION_UPLOAD: Transferring configuration files
- CONFIGURATION_FILE_IMPORT: Importing configuration files
- PARTIAL_CONFIGURATION_FILE_CREATE: Exporting configuration files and creating configuration backup files
- CONFIGURATION_BACKUP_FILE_DOWNLOAD: Downloading configuration backup files
- STORAGE_SET_SERVICE_ID: Setting the service ID of the storage cluster
- STORAGE_NODE_DUMP_FILE_CREATE_FILE: Creating a dump log file by using the GUI
- STORAGE_NODE_DUMP_FILE_DOWNLOAD: Downloading a dump log file by using the GUI
- STORAGE_NODE_DUMP_FILE_DELETE: Deleting a dump log file by using the GUI
- STORAGE_NODE_CONFIGURATION_PARAMETER_PARAMETERS_SHOW: Obtaining configuration parameters

Audit event	Description	Detailed information
▪ STORAGE_NODE_CONFIGURATION_PARAMETER_PARAMETERS_SET: Setting configuration parameters ▪ STORAGE_NODE_CONFIGURATION_PARAMETER_POLLING_MODE_SHOW: Obtaining the configuration parameter setting mode ▪ STORAGE_MODIFY_CONFIGURATION: Changing and setting configuration information 2. Job information is not output to the audit log for the following operations among configuration change operations: ▪ User management ▪ Edition of user authentication settings ▪ Session management 3. When performing the following operations, a REST API that is not described in the *Hitachi Virtual Storage Software Block REST API Reference* might be generated: ▪ Adding storage nodes ▪ Transferring configuration files ▪ Importing configuration files ▪ Exporting configuration files ▪ Creating a configuration backup file ▪ Creating, downloading, and deleting a dump log file by using the GUI ▪ Changing and setting configuration information ▪ When service personnel or maintenance personnel set the service ID of the storage cluster 4. Some characters in "Detailed information" are replaced with certain types of characters: ▪ Confidential information such as a password is replaced with asterisks (*). ▪ 0x00 to 0x1F (NULL and other characters), 0x2C (comma), and 0x7F (DEL) in ASCII code are replaced with question marks ("?": 0x3F in ASCII code).		

Type of audit event: Maintenance

Audit event	Description	Detailed information
START MAINTENANCE MODE	Indicates that Maintenance mode is enabled.	None
START RESCUE MODE	Indicates that Rescue mode is enabled.	None

Type of audit event: StartStop

Audit event	Description	Detailed information
STORAGE CLUSTER STARTED	Indicates that the storage cluster started.	None

Chapter 4: Managing audit logs

Overview of audit logs

An audit log records the operations performed on the storage cluster. An audit log allows you to verify when and by whom operations have been performed and to see if each operation performed to the storage system or user data complies with audit standards such as laws, regulations, industrial standards, and in-house stipulations.

An audit log can contain a maximum of 750,000 activities. If this maximum limit is exceeded, the log is overwritten in first-in-first-out basis.

An audit log is stored in the cluster master node (primary) and you can obtain it in either of the following ways:

- Create an audit log file and download it to the controller node.
- Configure transfer of audit logs to the Syslog server and obtain them from the Syslog server.

Each audit log created after Syslog transfer is set is transferred to the Syslog server. An audit log created before Syslog transfer is set is not transferred to the Syslog server.

The following Syslog server is supported.

- Rsyslog 8

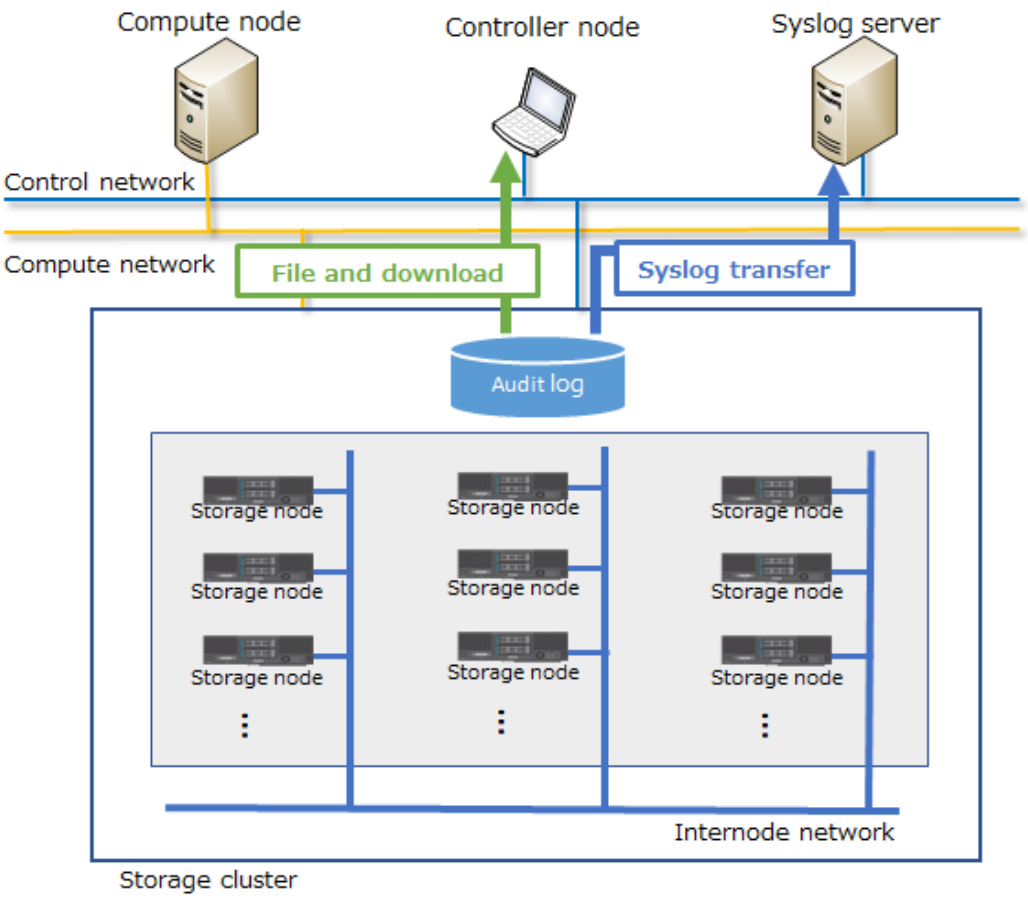

Downloading an audit log to the controller node (CLI or REST API)

Create an audit log file and download it to the controller node as follows.

Before you begin

Required role: Audit or Security

Procedure

1. Create an audit log file.

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

   REST API: POST /v1/objects/audit-logs/actions/create-file/invoke

   CLI: audit_log_create_file

   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. Download an audit log file.

   An audit log file is downloaded as a zip file to the current folder in which a REST API or CLI command has been run.

   REST API: GET /v1/objects/audit-logs/download

   CLI: audit_log_download

   The name of an audit log file when running a CLI command is "download-auditlog.zip".

Obtaining Syslog transfer settings of audit logs (CLI or REST API)

This section describes how to obtain the settings for the Syslog server to which audit logs are transferred. The following information can be obtained.

- locationName: Location information
- index: ID of the Syslog server
- isEnabled: Whether audit logs are transferred to the Syslog server
- serverName: Host name or IP address (IPv4) of the syslog server
- port: Port number of the Syslog server
- transportProtocol: Communication protocol

Before you begin

Required role: Security

Procedure

1. Run either of the following commands to obtain the settings for the Syslog server.

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

   CLI: audit_log_setting_show

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

Set Syslog transfer of audit logs as follows. You can set up to two Syslog servers.

An audit log can contain up to 750,000 activities. If transfer to the Syslog server is enabled, a user is notified by an event log and audit log when the number of untransferred activities reaches 70% of the maximum and when it reaches 100%.

Note:

- 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.
- When syslog transfer of audit logs is set, Virtual Storage Software block periodically sends ICMP echo requests to the set Syslog server to establish network paths.
- In the audit log Syslog 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 the Syslog server so that it can accept communication 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: Audit or Security

Procedure

1. Set Syslog transfer of audit logs as follows.

   Run either of the following commands with the parameters for setting audit logs specified.

When you specify the locationName (CLI: --location_name) parameter, observe the following:

- Number of characters: 1 to 180
- Characters that can be used: Numbers (0 to 9), uppercase alphabet (A to Z), lowercase alphabet (a to z), symbols (! # $ ' ( ) + -.@ _ ` { } ~)

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

CLI: audit_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. In the Syslog server, make settings so that audit 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: Control port IP address of the cluster master node (primary)
- Port number of Syslog transmission source: Port number set in step 1
- Communications protocol: Communications protocol set in step 1

4. Verify that audit logs are correctly transferred to the Syslog server.

Make the Syslog transfer setting of audit 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 audit log is transferred to the Syslog server, indicating that the job for audit log Syslog transfer setting has been started. If the log is not transferred, review the setting and the network.

- Audit event: JOB STARTED FOR audit_log_setting_set

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.

Monitoring the audit logs transferred to the Syslog server

Sometimes, audit logs are not transferred to the Syslog server due to network failures. In such a case, download and verify an audit log.

Procedure

1. Verify the *serial number* of each audit log received by the Syslog server.

   If no *serial number* is missing, all audit logs are correctly transferred to the Syslog server. Any missing *serial number* indicates that an audit log was not transferred to the Syslog server due to network failure. In such a case, verify the audit log by the following procedure.

2. Download the audit log (see *Downloading an audit log to the controller node*). Verify any audit log whose message *serial number* is missing in the Syslog server.

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