



StorageGRID® Webscale 10.4

Audit Message Reference

April 2017 | 215-11697_A0
doccomments@netapp.com

Contents

Audit message overview	5
Audit message flow and retention	5
Audit log message levels	7
Changing audit message levels	8
Audit log file access	8
Audit log file and message formats	9
Audit log file format	9
Audit message format	9
Data types	10
Event-specific data	11
Common elements in audit messages	11
Audit message examples	12
Audit messages and the object lifecycle	14
Object ingest transactions	15
Example: S3 Object Ingest	15
Object delete transactions	16
Example: S3 Object Deletion	17
Object retrieve transactions	17
Example: S3 Object Retrieval	18
Metadata update messages	18
Example: CDMI metadata update	19
Audit messages	20
Audit message categories	20
System audit messages	20
Object storage audit messages	22
HTTP protocol audit messages	23
Management audit message (MGAU)	24
Audit messages	24
APCT: Archive Purge from Cloud-Tier	24
ARCB: Archive Object Retrieve Begin	24
ARCE: Archive Object Retrieve End	25
ARCT: Archive Retrieve from Cloud-Tier	25
AREM: Archive Object Remove	26
ASCE: Archive Object Store End	26
ASCT: Archive Store Cloud-Tier	27
ATCE: Archive Object Store Begin	27
AVCC: Archive Validate Cloud-Tier Configuration	28
CBRE: Object Receive End	28
CBSE: Object Send End	29
CDMD: CDMI Delete Transaction	30
CDMG: CDMI GET Transaction	31

CDMP: CDMI PUT or POST Transaction to Create Object	31
CDMU: CDMI PUT Transaction to Update Object	32
ECOC: Corrupt Erasure Coded Data Fragment	33
ETAF: Security Authentication Failed	33
ETCA: TCP/IP Connection Establish	34
ETCC: TCP/IP Connection Close	34
ETCF: TCP/IP Connection Fail	35
ETCR: TCP/IP Connection Refused	36
GNRG: GNDS Registration	37
GNUR: GNDS Unregistration	37
GTED: Grid Task Ended	37
GTST: Grid Task Started	38
GTSU: Grid Task Submitted	39
HTSC: HTTP Session Close	39
HTSE: HTTP Session Establish	40
LLST: Location Lost	40
MGAU: Management audit message	41
OLST: System Detected Lost Object	42
ORLM: Object Rules Met	42
SADD: Security Audit Disable	44
SADE: Security Audit Enable	44
SCMT: Object Store Commit	44
SDEL: S3 DELETE	45
SGET: S3 GET	45
SHEA: S3 HEAD	46
SPUT: S3 PUT	47
SREM: Object Store Remove	48
SUPD: S3 Metadata Updated	48
SVRF: Object Store Verify Fail	49
SVRU: Object Store Verify Unknown	50
SYSD: Node Stop	50
SYST: Node Stopping	50
SYSU: Node Start	51
VLST: User Initiated Volume Lost	51
WDEL: Swift DELETE	51
WGET: Swift GET	52
WHEA: Swift HEAD	53
WPUT: Swift PUT	53
Determining the security partition for an object	54
Glossary	55
Copyright information	62
Trademark information	63
How to send comments about documentation and receive update notifications	64
Index	65

Audit message overview

This guide contains information about the structure and content of StorageGRID Webscale system's audit logs, providing you with the information necessary to read and analyze the audit trail of system activity.

The guide is intended for administrators responsible for producing reports of system activity and usage that require analysis of the StorageGRID Webscale system's audit messages.

You are assumed to have a sound understanding of the nature of audited activities within the StorageGRID Webscale system. To use the text log file, you must have access to the configured audit share on the Admin Node.

Audit message flow and retention

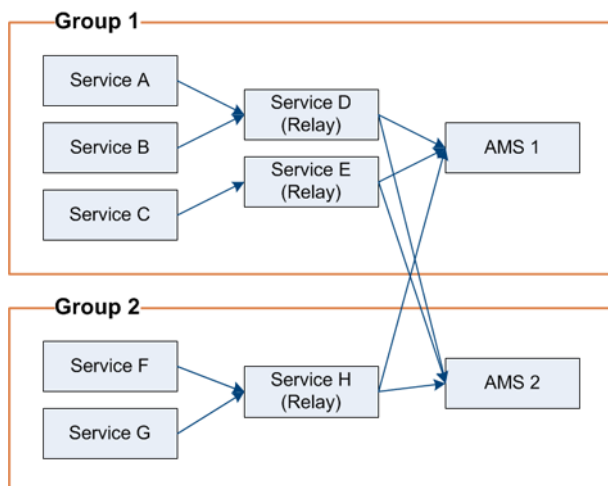
As StorageGRID Webscale services perform their various activities and process events, audit messages are generated to retain a record of this activity.

Audit messages are processed by the Audit Management System (AMS) service, which is hosted by the Admin Node, and they are stored in the form of text log files.

Audit message flow

Audit messages are generated internally by each service. All services generate audit messages during normal system operation. These messages are sent to all connected AMS services for processing and storage, so that each AMS service maintains a complete record of system activity.

Some services can be designated as audit message relay services. They act as collection points to reduce the need for every service to send its audit messages to all connected AMS services. As shown in the audit message flow diagram, each relay service must send messages to all AMS service destinations, whereas services can send messages to just one relay service.

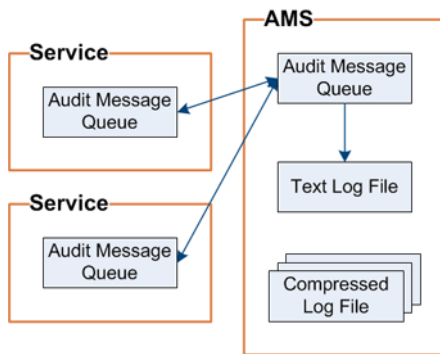


Relay services are designated at the time the topology of the StorageGRID Webscale deployment is configured. In a StorageGRID Webscale system, the ADC service is designated as the audit message relay.

Message retention

After an audit message is generated, it is stored on the grid node of the originating service until it has been committed to all connected AMS services, or a designated audit relay service. The relays in turn

store the message until it is committed at all AMS services. This process includes a confirmation (positive acknowledgment) to ensure that no messages are lost.



Messages arrive at the AMS service and are stored in a queue pending a confirmed write to the audit log file (`audit.log`). Confirmation of the arrival of messages is sent to the originating service (or audit relay) to permit the originator to delete its copy of the message.

Only after a message has been committed to storage at the AMS service can it be removed from the queue. If the backlog becomes unusually large, the local message buffer at the audit relay service (ADC) and the AMS service each have an alarm (AMQS) associated with it. During peak activity, the rate at which audit messages arrive can be faster than they can be relayed to the audit repository on the AMS service or committed to storage in the audit log file, causing a temporary backlog that clears itself when system activity declines.

Once a day the active audit log is saved to a file named for the date the file is saved (in the format `YYYY-MM-DD.txt`) and a new audit log file is started. If more than one audit log is created in a single day, each log is saved to a file named using the date when the file is saved and is appended with a number (in the format `YYYY-MM-DD.txt.#`): for example, `2010-04-23.txt.1`. Subsequent audit messages generated on the same day are saved to a new audit log. This new audit log is saved with the same date as the other, but with the appended number incremented by one: for example, `2010-04-23.txt.2`.

Audit logs are compressed after one day and are renamed `YYYY-MMDD.txt.gz` (where the original date is preserved). Audit logs files are saved to the Admin Node's `/var/local/audit/export` directory. Over time, this results in the consumption of storage allocated for audit logs on the Admin Node. A script monitors the audit log space consumption and deletes log files as necessary to free space in the `/var/local/audit/export` directory. Audit logs are deleted based on the date they were created, with the oldest being deleted first.

Duplicate messages

Audit messages are queued for storage by the AMS service. If system communications are interrupted (for example, because of service failures or network interruptions), the write status of some audit messages might be in doubt. The StorageGRID Webscale system takes a conservative approach in this case: all queued audit messages are resubmitted to the AMS service. This can result in duplicate messages in the audit log.

If duplicate messages are a cause for concern (for example, if the audit log is used for billing applications), you must detect and discard duplicate audit messages manually. To detect duplicate audit messages, you use the audit sequence count number (ASQN). Duplicate messages have the same ASQN.

Audit log message levels

You can adjust the message levels that are recorded in the audit log.

The AMS service and the HTTP audit feed filter incoming audit messages based on settings made through the **Configuration > Audit** page.

The following categories of messages are generated:

- **System** — By default, this level is set to Normal.
- **Object Storage** — By default, this level is set to Error.
- **Protocol - HTTP** — By default, this level is set to Normal.
- **Management** — By default, this level is set to Normal.

Note: For new installations starting at 10.3 and beyond, the above defaults are in effect. For upgrades of systems implemented prior to 10.3, the default for all categories is set to Normal.

Note: During upgrades, audit level configurations will not be effective immediately.

Audit

Audit Levels

System	Normal	▼
Storage	Error	▼
Protocol	Normal	▼
Management	Normal	▼

Audit Protocol Headers

Header Name 1	X-Forwarded-For	×
Header Name 2	x-amz-*	+ ×

Save

Related references

[System audit messages](#) on page 20

[Object storage audit messages](#) on page 22

[HTTP protocol audit messages](#) on page 23

[Management audit message \(MGAU\)](#) on page 24

Changing audit message levels

You can adjust the audit message levels to increase or decrease the number of messages recorded for each audit message category to match the detail level you require.

Before you begin

- You must be signed in to the Grid Management Interface using a supported browser.
- To perform this task, you need specific access permissions. For details, see information about controlling system access with administration user accounts and groups.

Steps

1. Select **Configuration > Audit**.
2. Select an audit level from the drop-down list:

Option	Description
Off	No audit messages from the category are logged.
Error	Only error messages are logged—audit messages for which the result code was not “successful” (SUCS).
Normal	Standard transactional messages are logged—the messages listed in this guide for the category.
Debug	Trace messages are logged; for troubleshooting only.

The messages included for any particular level include those that would be logged at the higher levels. Therefore, the Normal level includes all of the Error messages.

3. Under **Audit Protocol Headers**, enter the name of the HTTP request headers to be included in Protocol audit messages. Use an asterisk “*” as a wildcard, or use the escape sequence “*” as a literal asterisk. Use the plus sign to create a list of header name fields for entry.

Note: This field applies to S3/Swift requests only.

When such HTTP headers are found in a request, they are included in the audit message under the field HTRH.

Note: Because they are used for Protocol audit messages, request headers are logged only if the audit level for **Protocol** is not **Off**.

4. Click **Save**.

Related information

[StorageGRID Webscale 10.4 Administrator Guide](#)

Audit log file access

The audit share contains the active `audit.log` file and any compressed audit log files. For easy access to audit logs, you can configure client access to audit shares for both CIFS and NFS. You can also access audit log files directly from the command line of the Admin Node.

For more information about configuring audit shares, see the [StorageGRID Webscale 10.4 Administrator Guide](#).

Audit log file and message formats

You need to understand how audit log files and audit messages are formatted before you can use the audit log to gather information about your system and troubleshoot issues.

Audit log file format

The audit log file is found on every Admin Node and contains a collection of individual audit messages.

Each audit message contains the following:

- The Coordinated Universal Time (UTC) of the event that triggered the audit message (ATIM) in ISO 8601 format (that is, YYYY-MM-DDTHH:MM:SS.UUUUUU, where UUUUUU are microseconds), followed by a space.
- The audit message itself, enclosed within square brackets and beginning with “AUDT.”

The following is part of a sample log file:

```
2014-07-17T21:16:49.634203
[AUDT:[CNID(UI64):1405569153677780][INIE(FC32):LOCL][RSLT(FC
32):SUCS][AVER(UI32):10][ATIM(UI64):1405631809634203][ATYP(F
C32):ETCC][ANID(UI32):16259461][AMID(FC32):CONL][ATID(UI64):
14665900343074544330][ASQN(UI64):263][ASES(UI64):14055691536
33087]]
2014-07-17T21:16:49.863879
[AUDT:[CNID(UI64):1405569153677782][INIE(FC32):LOCL][RSLT(FC
32):SUCS][AVER(UI32):10][ATIM(UI64):1405631809863879][ATYP(F
C32):ETCC][ANID(UI32):16259461][AMID(FC32):CONL][ATID(UI64):
89577087142912923][ASQN(UI64):264][ASES(UI64):14055691536330
87]]
2014-07-17T21:16:52.606919
[AUDT:[SEID(FC32):CONS][CNDR(FC32):INBO][SVIP(UI32):1406][DA
IP(IPAD):"127.0.0.1"][SAIP(IPAD):"127.0.0.1"][CNID(UI64):140
5569167131673][RSLT(FC32):SUCS][AVER(UI32):10][ATIM(UI64):14
05631812606919][ATYP(FC32):ETCA][ANID(UI32):16039415][AMID(F
C32):PSVR][ATID(UI64):18353570832908893687][ASQN(UI64):258][
ASES(UI64):1405569167108360]]
```

Audit message format

Audit messages exchanged within the StorageGRID Webscale system include standard information common to all messages and specific content describing the event or activity being reported.

The following is a sample audit message as it might appear in the `audit.log` file:

```
2014-07-17T03:50:47.484627
[AUDT:[RSLT(FC32):VRGN][AVER(UI32):10][ATIM(UI64):1405569047484627]
[ATYP(FC32):SYSU][ANID(UI32):11627225][AMID(FC32):ARNI][ATID(UI64):
9445736326500603516][ASQN(UI64):0][ASES(UI64):1405569047484791]]
```

Each audit message is a string of attribute elements that has the following characteristics:

- Enclosed in square brackets []
- Introduced by the string AUDT, which indicates an audit message

- Without delimiters (no commas or spaces) between attributes
- Terminated by a line feed character \n

Each element includes an attribute code, a data type, and a value that are reported in this format:

```
[ATTR(type):value][ATTR(type):value]...
[ATTR(type):value]\n
```

The number of attribute elements in the message depends on the event type of the message.

The following list describes the attribute elements:

- **ATTR** is a four-character code for the attribute being reported. There are some attributes that are common to all audit messages and others that are event-specific.
- **type** is a four-character identifier of the programming data type of the value, such as UI64, FC32, and so on. The type is enclosed in parentheses ().
- **value** is the content of the attribute, typically a numeric or text value. Values always follow a colon (:). Values of data type CSTR are surrounded by double quotes “ ”.

Related concepts

[Audit messages](#) on page 20

[Audit message examples](#) on page 12

Related references

[Common elements in audit messages](#) on page 11

[Data types](#) on page 10

Data types

Different data types are used to store information in audit messages.

Type	Description
UI32	Unsigned long integer (32 bits); it can store the numbers 0 to 4,294,967,295.
UI64	Unsigned double long integer (64 bits); it can store the numbers 0 to 18,446,744,073,709,551,615.
FC32	Four-character constant; a 32-bit unsigned integer value represented as four ASCII characters such as “ABCD.”
IPAD	Used for IP addresses.
CSTR	<p>A variable-length array of UTF-8 characters. Characters can be escaped with the following conventions:</p> <ul style="list-style-type: none"> • Backslash is \. • Carriage return is \r. • Double quotes is \". • Line feed (new line) is \n. • Characters can be replaced by their hexadecimal equivalents (in the format \xHH, where HH is the hexadecimal value representing the character).

Event-specific data

The audit log records data that is specific to a system event in each audit message.

Following the opening “[AUDT:” container that identifies the message itself, the next set of attributes are items related to the event or action described by the audit message. These attributes are highlighted in the following example:

```
2014-06-20T00:14:20.424035 [AUDT:[HSID(UI64):1027401556]
[OBNS(CSTR):"UUID"][OBPA(CSTR):"/"][OBNA(CSTR):"DDE2
5220-7049-403D-8B71-B9D884A00864"][CBID(UI64):0x210C9
CFC55EACDC6][UUID(CSTR):"DDE25220-7049-403D-8B71-
B9D884A00864"][RSLT(FC32):SUCS][AVER(UI32):9][ATIM(UI64):
1213920860424035][ATYP(FC32):HHEA][ANID(UI32):12885257][AMID
(FC32):HTGM][ATID(UI64):9771581922913861059][ASQN(UI64):7374
859][ASES(UI64):1213662052895969]]
```

The event that these attributes describe is identified using the common ATYP element.

Related concepts

[Audit messages](#) on page 20

Related references

[Common elements in audit messages](#) on page 11

Common elements in audit messages

There is a set of elements that are common to all audit messages.

Code	Type	Description
AMID	FC32	Module ID: A four-character identifier of the module ID that generated the message. This indicates the code segment within which the audit message was generated.
ANID	UI32	Node ID: The grid node ID assigned to the service that generated the message. Each service is allocated a unique identifier at the time the StorageGRID Webscale system is configured and installed. This ID cannot be changed.
ASES	UI64	Audit Session Identifier: Indicates the time at which the audit system was initialized after the service started up. This time value is measured in microseconds since the operating system epoch (00:00:00 UTC on 1 January, 1970). It can be used to identify which messages were generated during a given runtime session.
ASQN	UI64	Sequence Count: A counter that is incremented for each generated audit message on the grid node (ANID). This counter is reset to zero at service restart. It can be used for consistency checks to ensure that no audit messages have been lost.
ATID	UI64	Trace ID: An identifier that is shared by the set of messages that were triggered by a single event.

Code	Type	Description
ATIM	UI64	<p>Timestamp: The time the event was generated that triggered the audit message, measured in microseconds since the operating system epoch (00:00:00 UTC on 1 January, 1970). Note that most available tools for converting the timestamp to local date and time are based on milliseconds.</p> <p>Rounding or truncation of the logged timestamp might be required. The human-readable time that appears at the beginning of the audit message in the <code>audit.log</code> file is the ATIM attribute in ISO 8601 format. (That is, the date and time is represented as YYYY-MMDDTHH:MM:SS.UUUUUU, where the T is a literal string character indicating the beginning of the time segment of the date. UUUUUU are microseconds).</p>
ATYP	FC32	<p>Event Type: A four-character identifier of the event being logged. This governs the “payload” content of the message: the attributes that are included.</p>
AVER	UI32	<p>Version: The version of the audit message. As the StorageGRID Webscale software evolves, new versions of services might incorporate new features in audit reporting. This field enables backward compatibility in the AMS service to process messages from older versions of services.</p>
RSLT	FC32	<p>Result: The result of event, process, or transaction. If is not relevant for a message, NONE is used rather than SUCS so that the message is not accidentally filtered.</p>

Audit message examples

You can find detailed information in each audit message. All audit messages use the same format.

The following is a sample audit message as it might appear in the `audit.log` file:

```
2014-07-17T21:17:58.959669
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):246979][S3AI(CSTR):"bc644d
381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"]
S3AK(CSTR):"UJXDKKQOXB7YARDS71Q2"]S3BK(CSTR):"s3small11"]S3K
Y(CSTR):"hello1"]CBID(UI64):0x50C4F7AC2BC8EDF7][CSIZ(UI64):0
][AVER(UI32):10][ATIM(UI64):1405631878959669][ATYP(FC32):SPUT
][ANID(UI32):12872812][AMID(FC32):S3RQ][ATID(UI64):1579224144
102530435][ASQN(UI64):45][ASES(UI64):1405569049324630]]
```

The audit message contains information about the event being recorded, as well as information about the audit message itself.

To identify which event is recorded by the audit message, look for the ATYP attribute (highlighted below):

```
2014-07-17T21:17:58.959669
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):246979][S3AI(CSTR):"bc644d
381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"]
S3AK(CSTR):"UJXDKKQOXB7YARDS71Q2"]S3BK(CSTR):"s3small11"]S3K
Y(CSTR):"hello1"]CBID(UI64):0x50C4F7AC2BC8EDF7][CSIZ(UI64):0
][AVER(UI32):10][ATIM(UI64):1405631878959669][ATYP(FC32):SP
UT][ANID(UI32):12872812][AMID(FC32):S3RQ][ATID(UI64):1579224
144102530435][ASQN(UI64):45][ASES(UI64):1405569049324630]]
```

The value of the ATYP attribute is SPUT. SPUT represents an S3 PUT transaction, which logs the ingest of an object to a bucket.

The following audit message also shows the bucket to which the object is associated:

```
2014-07-17T21:17:58.959669
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):246979][S3AI(CSTR):"bc644d
381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"]
S3AK(CSTR):"UJXDKKQOXB7YARDS71Q2"][S3BK(CSTR):"s3small11"][S3
KY(CSTR):"hello1"]][CBID(UI64):0x50C4F7AC2BC8EDF7][CSIZ(UI64):
0][AVER(UI32):10][ATIM(UI64):1405631878959669][ATYP(FC32):SPU
T][ANID(UI32):12872812][AMID(FC32):S3RQ][ATID(UI64):157922414
4102530435][ASQN(UI64):45][ASES(UI64):1405569049324630]]
```

To discover when the PUT event occurred, you should note the Universal Coordinated Time (UTC) timestamp at the beginning of the audit message. This value is a human-readable version of the ATIM attribute of the audit message itself:

```
2014-07-17T21:17:58.959669
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):246979][S3AI(CSTR):"bc644d
381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"]
S3AK(CSTR):"UJXDKKQOXB7YARDS71Q2"][S3BK(CSTR):"s3small11"][S3K
Y(CSTR):"hello1"]][CBID(UI64):0x50C4F7AC2BC8EDF7][CSIZ(UI64):0
][AVER(UI32):10][ATIM(UI64):1405631878959669][ATYP(FC32):SP
UT][ANID(UI32):12872812][AMID(FC32):S3RQ][ATID(UI64):15792241
44102530435][ASQN(UI64):45][ASES(UI64):1405569049324630]]
```

ATIM records the time, in microseconds, since the beginning of the UNIX epoch. For example, the value 1213920860718929 translates to Fri, 20 Jun 2008 00:14:20 UTC.

Related concepts

[*SPUT: S3 PUT*](#) on page 47

Related references

[*Common elements in audit messages*](#) on page 11

Audit messages and the object lifecycle

Each time an object is ingested, retrieved, or deleted, audit messages are generated. Audit messages are linked through identifiers specific to each protocol.

Protocol	Code
Linking S3 operations	S3BK (S3 Bucket) and/or S3KY (S3 Key)
Linking Swift operations	WCON (Swift Container) and/or WOBJ (Swift Object)
Linking CDMI operations	COID (CDMI Object Identifier)
Linking internal operations	CBID (Object's Internal Identifier)

Timing of audit messages

Because of factors such as timing differences between grid nodes, object size, and network delays, the order of audit messages generated by the different services can vary from that shown in the examples in this section.

Dual commit

Dual commit forces two copies of the object to be stored on two LDR services at the time of initial ingest. These initial copies are made before the object is evaluated against the active information lifecycle management (ILM) policy. At the same time that these initial copies are made, objects are queued for ILM evaluation. Depending on the configuration of ILM rules, additional copies might be made in different locations and the initial dual “commit” copies deleted. If the ILM rule is configured with only one content placement instruction, when ILM rules are evaluated, the StorageGRID Webscale system deletes one instance of the object so that only one instance of object data is stored.

For more information about dual commit, see the [StorageGRID Webscale 10.4 Administrator Guide](#).

Information lifecycle management policy configuration

With the default ILM policy (Baseline 2 Copy Rule v1.0), object data is copied once for a total of two copies. If the ILM policy requires more than two copies, there will be an additional set of CBRE, CBSE, and SCMT messages for each extra copy. For more information about ILM policies, see the [StorageGRID Webscale 10.4 Administrator Guide](#).

Archive Nodes

The series of audit messages generated when an Archive Node sends object data to an external archival storage system is similar to that for Storage Nodes except that there is no SCMT (Store Object Commit) message, and the ATCE (Archive Object Store Begin) and ASCE (Archive Object Store End) messages are generated for each archived copy of object data.

The series of audit messages generated when an Archive Node retrieves object data from an external archival storage system is similar to that for Storage Nodes except that the ARCB (Archive Object Retrieve Begin) and ARCE (Archive Object Retrieve End) messages are generated for each retrieved copy of object data.

The series of audit messages generated when an Archive Node deletes object data from an external archival storage system is similar to that for Storage Nodes except that there is no SREM (Object Store Remove) message, and there is an AREM (Archive Object Remove) message for each delete request.

Object ingest transactions

You can identify client ingest transactions in the audit log by locating API-specific (S3, Swift, and CDMI) audit messages.

Not all audit messages generated during an ingest transaction are listed in the following tables. Only the messages required to trace the ingest transaction are included.

S3 ingest audit messages

Code	Name	Description	Trace	See
SPUT	HTTP PUT Transaction Start	A request made to create a new object in a bucket.	CBID, S3BK, S3KY	<i>SPUT: S3 PUT</i> on page 47
ORLM	Object Rules Met	The ILM policy has been satisfied for this object.	CBID	<i>ORLM: Object Rules Met</i> on page 42

Swift ingest audit messages

Code	Name	Description	Trace	See
WPUT	Swift PUT transaction	A Swift PUT ingest transaction has successfully completed.	CBID, WCON, WOBJ	<i>WPUT: Swift PUT</i> on page 53
ORLM	Object Rules Met	The ILM policy has been satisfied for this object.	CBID	<i>ORLM: Object Rules Met</i> on page 42

CDMI ingest audit messages

Code	Name	Description	Trace	See
CDMP	CDMI PUT or POST Transaction to Create Object	A CDMI PUT or POST ingest transaction has successfully completed.	CBID, COID	<i>CDMP: CDMI PUT or POST Transaction to Create Object</i> on page 31
ORLM	Object Rules Met	The ILM policy has been satisfied for this object.	CBID	<i>ORLM: Object Rules Met</i> on page 42

Example: S3 Object Ingest

The series of audit messages below is an example of the audit messages generated and saved to the audit log when an S3 client ingests an object to a Storage Node (LDR service).

In this example, the active ILM policy is the default Baseline 2 Copy Rule.

Note: Not all audit messages generated during a transaction are listed in the example below. Only those related to the S3 ingest transaction (SPUT) are listed.

This example assumes that an S3 bucket has been previously created.

SPUT: S3 PUT

The SPUT message is generated to indicate that an S3 PUT transaction has been issued to create an object in a specific bucket.

```
2014-07-17T21:17:58.959669[AUDT:[RSLT(FC32):SUCS][TIME(UI64):25771]
[SAIP(IPAD):"10.96.112.29"]][S3AI(CSTR):"70899244468554783528"]
[S3AK(CSTR):"SGKHyalRU_5cLflqajtaFmxJn946lAWRJfBF33gAOg==" ]
[SUSR(CSTR):"urn:sgws:identity:70899244468554783528:root"]
[SACC(CSTR):"test"]][S3BK(CSTR):"example"]][S3KY(CSTR):"testobject-0-3"]
[CBID(UI64):0x8EF52DF8025E63A8][CSIZ(UI64):30720][AVER(UI32):10]
[ATIM(UI64):1462395667595443][ATYP(FC32):SPUT][ANID(UI32):12086324]
[AMID(FC32):S3RQ][ATID(UI64):14399932238768197038][ASQN(UI64):98]
[ASES(UI64):1461975217756408]]
```

ORLM: Object Rules Met

The ORLM message indicates that the ILM policy has been satisfied for this object. The message includes the object's CBID and the name of the ILM rule that was applied.

```
2014-07-17T21:18:31.230669[AUDT:[CBID(UI64):0x50C4F7AC2BC8EDF7]
[RULE(CSTR):"Make 2 Copies"]][STAT(FC32):DONE][CSIZ(UI64):0][SPAR(UI64):0]
[UUID(CSTR):"0B344E18-98ED-4F22-A6C8-A93ED68F8D3F"]][LOCS(CSTR):"CLDI
12872812, CLDI 12119796"]][RSLT(FC32):SUCS][AVER(UI32):10]
[ATYP(FC32):ORLM][ATIM(UI64):1405631911230669][ATID(UI64):
15494889725796157557][ANID(UI32):13100453][AMID(FC32):BCMS][ASQN(UI64):
14][ASES(UI64):1405569248205144]]
```

Object delete transactions

You can identify object delete transactions in the audit log by locating API-specific (S3 and Swift) audit messages.

Not all audit messages generated during a delete transaction are listed in the following tables. Only messages required to trace the delete transaction are included.

S3 delete audit messages

Code	Name	Description	Trace	See
SDEL	S3 Delete	Request made to delete the object from a bucket.	CBID, S3KY	<i>SDEL: S3 DELETE</i> on page 45

Swift delete audit messages

Code	Name	Description	Trace	See
WDEL	Swift Delete	Request made to delete the object from a container, or the container.	CBID, WOBJ	<i>WDEL: Swift DELETE</i> on page 51

CDMI delete audit messages

Code	Name	Description	Trace	See
CDMP	CDMI Delete Transaction	Request is made to delete a data object.	CBID, COID	<i>CDMP: CDMI PUT or POST Transaction to Create Object</i> on page 31

Example: S3 Object Deletion

When an S3 client deletes an object from a Storage Node (LDR service), an audit message is generated and saved to the audit log.

Note: Not all audit messages generated during a transaction are listed in the example below. Only those related to the S3 ingest transaction (SDEL) are listed.

SDEL: S3 Delete

Object deletion begins when the client sends a DELETE Object request to an LDR service. The message contains the bucket from which to delete the object and the object's S3 Key, which is used to identify the object.

```
A2014-07-17T21:17:58.959669[AUDT:[RSLT(FC32):SUCS][TIME(UI64):14316]
[SAIP(IPAD):"10.96.112.29"]][S3AI(CSTR):"70899244468554783528"]
[S3AK(CSTR):"SGKHyalRU_5cLflqajtaFmxJn946lAWRJfBF33gAOg==" ]
[SUSR(CSTR):"urn:sgws:identity::70899244468554783528:root"]
[SACC(CSTR):"test"]][S3BK(CSTR):"example"]][S3KY(CSTR):"testobject-0-7"]
[CBID(UI64):0x339F21C5A6964D89][CSIZ(UI64):30720][AVER(UI32):10]
[ATIM(UI64):1462395677674894][ATYP(FC32):SDEL][ANID(UI32):12086324]
[AMID(FC32):S3RQ][ATID(UI64):4727861330952970593][ASQN(UI64):132]
[ASES(UI64):1461975217756408]]
```

Object retrieve transactions

You can identify object retrieve transactions in the audit log by locating API-specific (S3 and Swift) audit messages.

Not all audit messages generated during a retrieve transaction are listed in the following tables. Only messages required to trace the retrieve transaction are included.

S3 retrieval audit messages

Code	Name	Description	Trace	See
SGET	S3 GET	Request made to retrieve an object from a bucket.	CBID, S3BK, S3KY	<i>SGET: S3 GET</i> on page 45

Swift retrieval audit messages

Code	Name	Description	Trace	See
WGET	Swift GET	Request made to retrieve an object from a container.	CBID, WCON, WOBJ	<i>WGET: Swift GET</i> on page 52

CDMI retrieval audit messages

Code	Name	Description	Trace	See
CDMG	CDMI GET Transaction	Request is made to read a data object.	CBID, COID	<i>CDMG: CDMI GET Transaction</i> on page 31

Example: S3 Object Retrieval

When an S3 client retrieves an object from a Storage Node (LDR service), an audit message is generated and saved to the audit log.

Note that not all audit messages generated during a transaction are listed in the example below. Only those related to the S3 retrieval transaction (SGET) are listed.

SGET: S3 GET

Object retrieval begins when the client sends a GET Object request to an LDR service. The message contains the bucket from which to retrieve the object and the object's S3 Key, which is used to identify the object.

```
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):4826][SAIP(IPAD):"10.96.112.29"]
[S3AI(CSTR):"70899244468554783528"]
[S3AK(CSTR):"SGKHyalRU_5cLflqajtaFmxJn946lAWRJfBF33gAOg==" ]
[SUSR(CSTR):"urn:sgws:identity:70899244468554783528:root"]
[SACC(CSTR):"test"]][S3BK(CSTR):"example"][S3KY(CSTR):"testobject-0-2"]
[CBID(UI64):0xD94CB36D1F4C3B82][CSIZ(UI64):30720][AVER(UI32):10]
[ATIM(UI64):1462395672856696][ATYP(FC32):SGET][ANID(UI32):12086324]
[AMID(FC32):S3RQ][ATID(UI64):642093248925982115][ASQN(UI64):116]
[ASES(UI64):1461975217756408]]
```

Metadata update messages

Audit messages are generated when a client (S3 or CDMI) updates an object's metadata.

S3 metadata update audit messages

Code	Name	Description	Trace	See
SUPD	S3 Metadata Updated	The S3 client makes a request to update metadata for a data or container object.	CBID	<i>SUPD: S3 Metadata Updated</i> on page 48

CDMI metadata update audit messages

Code	Name	Description	Trace	See
CDMU	CDMI PUT Transaction	The CDMI client makes a request to update a data or container object.	COID, CBID	<i>CDMU: CDMI PUT Transaction to Update Object</i> on page 32

Example: CDMI metadata update

The example below lists the events that take place when a CDMI client updates an object's metadata.

CDMU: CDMI PUT

The CDMI client makes a request to update a data or container object.

```
2016-02-03T19:31:33.985481 [AUDT:[RSLT(FC32):SUCS][TIME(UI64):30850]
[HSID(UI64):2313303137][OBSP(CSTR):""][SPAR(UI64):0][CURI(CSTR):"/
cdmi_objectid/00006FFD00194D3F00488E81086CA4419D8FB6CFCA54F32FF4"]
[COID(CSTR):"00006FFD00194D3F00488E81086CA4419D8FB6CFCA54F32FF4"]
[CBID(UI64):0x9FD403E54D3C6F8B][CSIZ(UI64):5][AVER(UI32):10][ATIM(UI64):
1454527893985481][ATYP(FC32):CDMU][ANID(UI32):12107434][AMID(FC32):CDRQ]
[ATID(UI64):2392204792264697975][ASQN(UI64):178][ASES(UI64):
1454375943224080]]
```

Audit messages

Detailed descriptions of audit messages returned by the system are listed in the following sections. Each audit message is first listed in a table that groups related messages by the class of activity that the message represents. These groupings are useful both for understanding the types of activities that are audited, and for selecting the desired type of audit message filtering.

The audit messages are also listed alphabetically by their four-character codes. This alphabetic listing enables you to find information about specific messages.

The four-character codes used throughout this chapter are the ATYP values found in the audit messages as shown in the following sample message:

```
2014-07-17T03:51:07.526110
[AUDT:[SEID(FC32):RCON][CNDR(FC32):OUTB][SVIP(UI32):1501]
[DAIP(IPAD):"10.96.99.41"][SAIP(IPAD):"10.96.99.40"][CNID(UI64):0]
[RSLT(FC32):CRFU][AVER(UI32):10][ATIM(UI64):1405569067526110]
[ATYP(FC32):ETCF][ANID(UI32):11627225][AMID(FC32):RCON][ATID(UI64):
15065141165638247697][ASQN(UI64):1][ASES(UI64):1405569047484791]
```

Related concepts

[Audit messages](#) on page 24

Related tasks

[Changing audit message levels](#) on page 8

Audit message categories

You should be familiar with the various categories within which audit messages are grouped. These groups are organized based on the class of activity that the message represents.

The categories of audit messages are system audit messages, object storage audit messages, HTTP protocol audit messages, and management audit messages.

System audit messages

You should be familiar with audit messages belonging to the system audit category. These are events related to the auditing system itself, grid node states, system-wide task activity (grid tasks), and service backup operations, so that you can address potential issues.

Code	Message title and description	See
ECOC	Corrupt Erasure Coded Data Fragment: Indicates that a corrupt erasure coded data fragment has been detected.	ECOC: Corrupt Erasure Coded Data Fragment on page 33
ETAF	Security Authentication Failed: A connection attempt using Transport Layer Security (TLS) failed.	ETAF: Security Authentication Failed on page 33
ETCA	TCP/IP Connection Establish: An incoming or outgoing TCP/IP connection was successfully established.	ETCA: TCP/IP Connection Establish on page 34
ETCC	TCP/IP Connection Close: An established connection was closed by either side of the connection (normally or abnormally).	ETCC: TCP/IP Connection Close on page 34

Code	Message title and description	See
ETCF	TCP/IP Connection Fail: An outgoing connection attempt failed at the lowest level, due to communication problems.	<i>ETCF: TCP/IP Connection Fail</i> on page 35
ETCR	TCP/IP Connection Refused: An incoming TCP/IP connection attempt was not allowed.	<i>ETCR: TCP/IP Connection Refused</i> on page 36
GNRG	GNDS Registration: A service updated or registered information about itself in the StorageGRID Webscale system.	<i>GNRG: GNDS Registration</i> on page 37
GNUR	GNDS Unregistration: A service contains unregistered information about itself from the StorageGRID Webscale system.	<i>GNUR: GNDS Unregistration</i> on page 37
GTED	Grid Task Ended: The CMN service finished processing the grid task.	<i>GTED: Grid Task Ended</i> on page 37
GTST	Grid Task Started: The CMN service started to process the grid task.	<i>GTST: Grid Task Started</i> on page 38
GTSU	Grid Task Submitted: A grid task was submitted to the CMN service.	<i>GTSU: Grid Task Submitted</i> on page 39
LLST	Location Lost: This is the audit message that is generated when a location is lost.	<i>LLST: Location Lost</i> on page 40
OLST	Object Lost: A requested object cannot be located within the StorageGRID Webscale system.	<i>OLST: System Detected Lost Object</i> on page 42
ORLM	Object Rules Met: Object data is stored as specified by the ILM rules.	<i>ORLM: Object Rules Met</i> on page 42
SADD	Security Audit Disable: Audit message logging was turned off.	<i>SADD: Security Audit Disable</i> on page 44
SADE	Node Stop: A service was gracefully stopped.	<i>SADE: Security Audit Enable</i> on page 44
SVRF	Object Store Verify Fail: A content block failed verification checks.	<i>SVRF: Object Store Verify Fail</i> on page 49
SVRU	Object Store Verify Unknown: Unexpected object data detected in the object store.	<i>SVRU: Object Store Verify Unknown</i> on page 50
SYSD	Node Stop: A shutdown was requested.	<i>SYSD: Node Stop</i> on page 50
SYST	Node Stopping: A service initiated a graceful stop.	<i>SYST: Node Stopping</i> on page 50
SYSU	Node Start: A service started; the nature of the previous shutdown is indicated in the message.	<i>SYSU: Node Start</i> on page 51
VLST	User Initiated Volume Lost: The <code>/proc/CMSI/Volume_Lost</code> command was run.	<i>VLST: User Initiated Volume Lost</i> on page 51

Object storage audit messages

You should be familiar with audit messages belonging to the object storage audit category. These are events related to the storage and management of objects within the StorageGRID Webscale system. These include object storage and retrievals, grid-node to grid-node transfers, and verifications.

Code	Description	See
APCT	Archive Purge from Cloud-Tier: Archived object data is deleted from an external archival storage system, which connects to the StorageGRID Webscale through the S3 API.	<i>APCT: Archive Purge from Cloud-Tier</i> on page 24
ARCB	Archive Object Retrieve Begin: The ARC service begins the retrieval of object data from the external archival storage system.	<i>ARCB: Archive Object Retrieve Begin</i> on page 24
ARCE	Archive Object Retrieve End: Object data has been retrieved from an external archival storage system, and the ARC service reports the status of the retrieval operation.	<i>ARCE: Archive Object Retrieve End</i> on page 25
ARCT	Archive Retrieve from Cloud-Tier: Archived object data is retrieved from an external archival storage system, which connects to the StorageGRID Webscale through the S3 API.	<i>ARCT: Archive Retrieve from Cloud-Tier</i> on page 25
AREM	Archive Object Remove: A content block was successfully or unsuccessfully deleted from the external archival storage system.	<i>AREM: Archive Object Remove</i> on page 26
ASCE	Archive Object Store End: A content block has been written to the external archival storage system, and the ARC service reports the status of the write operation.	<i>ASCE: Archive Object Store End</i> on page 26
ASCT	Archive Store Cloud-Tier: Object data is stored to an external archival storage system, which connects to the StorageGRID Webscale through the S3 API.	<i>ASCT: Archive Store Cloud-Tier</i> on page 27
ATCE	Archive Object Store Begin: Writing a content block to an external archival storage has started.	<i>ATCE: Archive Object Store Begin</i> on page 27
AVCC	Archive Validate Cloud-Tier Configuration: The account and bucket settings provided were successfully or unsuccessfully validated.	<i>AVCC: Archive Validate Cloud-Tier Configuration</i> on page 28
CBSE	Object Send End: The source entity completed a grid-node to grid-node data transfer operation.	<i>CBSE: Object Send End</i> on page 29
CBRE	Object Receive End: The destination entity completed a grid-node to grid-node data transfer operation.	<i>CBRE: Object Receive End</i> on page 28
SCMT	Object Store Commit: A content block was completely stored and verified, and can now be requested.	<i>SCMT: Object Store Commit</i> on page 44
SREM	Object Store Remove: A content block was deleted from a grid node, and can no longer be requested directly.	<i>SREM: Object Store Remove</i> on page 48

HTTP protocol audit messages

You should be familiar with audit messages belonging to the HTTP protocol category. These are events related to interactions with internal and external system components using the HTTP protocol.

Code	Description	Used by	See
CDMD	CDMI DELETE Transaction: Logs a successful transaction to delete a data or container object.	CDMI client	<i>CDMD: CDMI Delete Transaction</i> on page 30
CDMG	CDMI GET Transaction: Logs a successful transaction to read from a data or container object.	CDMI client	<i>CDMG: CDMI GET Transaction</i> on page 31
CDMP	CDMI PUT or POST Transaction to Create an Object: Logs a successful transaction to create a new data or container object.	CDMI client	<i>CDMP: CDMI PUT or POST Transaction to Create Object</i> on page 31
CDMU	CDMI PUT Transaction to Update an Object: Logs a successful transaction to update a data or container object.	CDMI client	<i>CDMU: CDMI PUT Transaction to Update Object</i> on page 32
HTSC	HTTP Session Close: Logs a successful close of a previously established HTTP session.	CDMI client	<i>HTSC: HTTP Session Close</i> on page 39
HTSE	HTTP Session Establish: Logs a successful establishment of an HTTP session with a grid node.	CDMI client	<i>HTSE: HTTP Session Establish</i> on page 40
SDEL	S3 DELETE: Logs a successful transaction to delete an object or bucket.	S3 client	<i>SDEL: S3 DELETE</i> on page 45
SGET	S3 GET: Logs a successful transaction to retrieve an object or list the objects in a bucket.	S3 client	<i>SGET: S3 GET</i> on page 45
SHEA	S3 HEAD: Logs a successful transaction to check for the existence of an object or bucket.	S3 client	<i>SHEA: S3 HEAD</i> on page 46
SPUT	S3 PUT: Logs a successful transaction to create a new object or bucket.	S3 client	<i>SPUT: S3 PUT</i> on page 47
SUPD	S3 Metadata Updated: Logs a successful transaction to update the metadata for an existing object or bucket.	S3 client	<i>SUPD: S3 Metadata Updated</i> on page 48
WDEL	Swift DELETE: Logs a successful transaction to delete an object or container.	Swift client	<i>WDEL: Swift DELETE</i> on page 51
WGET	Swift GET: Logs a successful transaction to retrieve an object or list the objects in a container.	Swift client	<i>WGET: Swift GET</i> on page 52

Code	Description	Used by	See
WHEA	Swift HEAD: Logs a successful transaction to check for the existence of an object or container.	Swift client	WHEA: Swift HEAD on page 53
WPUT	Swift PUT: Logs a successful transaction to create a new object or container.	Swift client	WPUT: Swift PUT on page 53

Management audit message (MGAU)

The Management category logs user requests to the Management API.

Code	Message title and description	See
MGAU	Management API audit message: A log of user requests.	MGAU: Management audit message on page 41

Audit messages

When system events occur, the StorageGRID Webscale system generates audit messages and records them in the audit log.

APCT: Archive Purge from Cloud-Tier

This message is generated by when archived object data is deleted from an external archival storage system, which connects to the StorageGRID Webscale through the S3 API.

Code	Field	Description
CBID	Content Block ID	The unique identifier for the content block that was deleted.
CSIZ	Content Size	The size of the object in bytes. Always returns 0.
RSLT	Result Code	Returns successful (SUCS) or the error reported by the backend.
SUID	Storage Unique Identifier	Unique identifier (UUID) of the cloud-tier from which the object was deleted.

ARCB: Archive Object Retrieve Begin

This message is generated when a request is made to retrieve archived object data and the retrieval process begins. Retrieval requests are processed immediately, but can be reordered to improve efficiency of retrieval from linear media such as tape.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the Content Block to be retrieved from the external archival storage system.

Code	Field	Description
RSLT	Result	Indicates the result of starting the archive retrieval process. Currently defined value is: SUCS: The content request was received and queued for retrieval.

This audit message marks the time of an archive retrieval. It allows you to match the message with a corresponding ARCE end message to determine the duration of archive retrieval, and whether the operation was successful.

ARCE: Archive Object Retrieve End

This message is generated when an attempt by the Archive Node to retrieve object data from an external archival storage system completes. If successful, the message indicates that the requested object data has been completely read from the archive location, and was successfully verified. After the object data has been retrieved and verified, it is delivered to the requesting service.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the Content Block to be retrieved from the external archival storage system.
VLID	Volume Identifier	The identifier of the volume on which the data was archived. If an archive location for the content is not found, a Volume ID of 0 is returned.
RSLT	Retrieval Result	The completion status of the archive retrieval process: <ul style="list-style-type: none"> • SUCS: successful • VRFL: failed (object verification failure) • ARUN: failed (external archival storage system unavailable) • CANC: failed (retrieval operation canceled) • GERR: failed (general error)

Matching this message with the corresponding ARCB message can indicate the time taken to perform the archive retrieval. This message indicates whether the retrieval was successful, and in the case of failure, the cause of the failure to retrieve the content block.

ARCT: Archive Retrieve from Cloud-Tier

This message is generated when archived object data is retrieved from an external archival storage system, which connects to the StorageGRID Webscale through the S3 API.

Code	Field	Description
CBID	Content Block ID	The unique identifier for the content block that was retrieved.
CSIZ	Content Size	The size of the object in bytes. The value is only accurate for successful retrieves.
RSLT	Result Code	Returns successful (SUCS) or the error reported by the backend.

Code	Field	Description
SUID	Storage Unique Identifier	Unique identifier (UUID) of the external archival storage system.
TIME	Time	Total processing time for the request in microseconds.

AREM: Archive Object Remove

The Archive Object Remove audit message indicates that a content block was successfully or unsuccessfully deleted from an Archive Node. If the result is successful, the Archive Node has successfully informed the external archival storage system that StorageGRID Webscale has released an object location. Whether the object is removed from the external archive storage system depends on the type of system and its configuration.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the Content Block to be retrieved from the external archival media system.
VLID	Volume Identifier	The identifier of the volume on which the object data was archived.
RSLT	Result	The completion status of the archive removal process: <ul style="list-style-type: none"> • SUCS: successful • ARUN: failed (external archival storage system unavailable) • GERR: failed (general error)

ASCE: Archive Object Store End

This message indicates that writing a content block to an external archival storage system has ended.

Code	Field	Description
CBID	Content Block Identifier	The identifier of the content block stored on the external archival storage system.
VLID	Volume Identifier	The unique identifier of the archive volume to which the object data is written.
VREN	Verification Enabled	Indicates if verification is performed for content blocks. Currently defined values are: <ul style="list-style-type: none"> • VENA: verification is enabled • VDSA: verification is disabled
MCLS	Management Class	A string identifying the TSM Management Class to which the content block is assigned if applicable.

Code	Field	Description
RSLT	Result	<p>Indicates the result of the archive process. Currently defined values are:</p> <ul style="list-style-type: none"> • SUCS: successful (archiving process succeeded) • OFFL: failed (archiving is offline) • VRFL: failed (object verification failed) • ARUN: failed (external archival storage system unavailable) • GERR: failed (general error)

This audit message means that the specified content block has been written to the external archival storage system. If the write fails, the result provides basic troubleshooting information about where the failure occurred. More detailed information about archive failures can be found by examining Archive Node attributes in the StorageGRID Webscale system.

ASCT: Archive Store Cloud-Tier

This message is generated when archived object data is stored to an external archival storage system, which connects to StorageGRID Webscale through the S3 API.

Code	Field	Description
CBID	Content Block ID	The unique identifier for the content block that was retrieved.
CSIZ	Content Size	The size of the object in bytes.
RSLT	Result Code	Returns successful (SUCS) or the error reported by the backend.
SUID	Storage Unique Identifier	Unique identifier (UUID) of the cloud-tier the content was stored to.
TIME	Time	Total processing time for the request in microseconds.

ATCE: Archive Object Store Begin

This message indicates that writing a content block to an external archival storage has started.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the content block to be archived.
VLID	Volume Identifier	The unique identifier of the volume to which the content block is written. If the operation fails, a volume ID of 0 is returned.

Code	Field	Description
RSLT	Result	<p>Indicates the result of the transfer of the content block. Currently defined values are:</p> <ul style="list-style-type: none"> • SUCS: success (content block stored successfully) • EXIS: ignored (content block was already stored) • ISFD: failed (insufficient disk space) • STER: failed (error storing the CBID) • OFFL: failed (archiving is offline) • GERR: failed (general error)

AVCC: Archive Validate Cloud-Tier Configuration

This message is generated when the configuration settings are validated for a Cloud Tiering - Simple Storage Service (S3) target type.

Code	Field	Description
RSLT	Result Code	Returns successful (SUCS) or the error reported by the backend.
SUID	Storage Unique Identifier	UUID associated with the external archival storage system being validated.

CBRE: Object Receive End

When transfer of a content block from one node to another is completed, this message is issued by the destination entity.

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node session/connection.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	<p>Indicates if the CBID transfer was push-initiated or pull-initiated:</p> <p>PUSH: The transfer operation was requested by the sending entity.</p> <p>PULL: The transfer operation was requested by the receiving entity.</p>
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the sequence count on which the transfer started.
CTAS	Actual End Sequence Count	Indicates the last sequence count successfully transferred. If the Actual End Sequence Count is the same as the Start Sequence Count, and the Transfer Result was not successful, no data was exchanged.

Code	Field	Description
RSLT	Transfer Result	<p>The result of the transfer operation (from the perspective of the sending entity):</p> <p>SUCS: transfer successfully completed; all requested sequence counts were sent.</p> <p>CONL: connection lost during transfer</p> <p>CTMO: connection timed-out during establishment or transfer</p> <p>UNRE: destination node ID unreachable</p> <p>CRPT: transfer ended due to reception of corrupt or invalid data (might indicate tampering)</p>

This audit message means a node-to-node data transfer operation was completed. If the Transfer Result was successful, the operation transferred data from “Start Sequence Count” to “Actual End Sequence Count”. Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow and to locate, tabulate, and analyze errors. When combined with storage audit messages, it can also be used to verify replica counts.

CBSE: Object Send End

When transfer of a content block from one node to another is completed, this message is issued by the source entity.

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node session/connection.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	<p>Indicates if the CBID transfer was push-initiated or pull-initiated:</p> <p>PUSH: The transfer operation was requested by the sending entity.</p> <p>PULL: The transfer operation was requested by the receiving entity.</p>
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the sequence count on which the transfer started.
CTAS	Actual End Sequence Count	Indicates the last sequence count successfully transferred. If the Actual End Sequence Count is the same as the Start Sequence Count, and the Transfer Result was not successful, no data was exchanged.

Code	Field	Description
RSLT	Transfer Result	<p>The result of the transfer operation (from the perspective of the sending entity):</p> <p>SUCS: Transfer successfully completed; all requested sequence counts were sent.</p> <p>CONL: connection lost during transfer</p> <p>CTMO: connection timed-out during establishment or transfer</p> <p>UNRE: destination node ID unreachable</p> <p>CRPT: transfer ended due to reception of corrupt or invalid data (might indicate tampering)</p>

This audit message means a node-to-node data transfer operation was completed. If the Transfer Result was successful, the operation transferred data from “Start Sequence Count” to “Actual End Sequence Count”. Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow and to locate, tabulate, and analyze errors. When combined with storage audit messages, it can also be used to verify replica counts.

CDMD: CDMI Delete Transaction

When a CDMI client issues a DELETE transaction, a request is made to remove the specified data or container object. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the corresponding content block requested. If the CBID is unknown, this field is set to 0.
CSIZ	Content Size	The size of the data object in bytes. Container objects do not include this field.
CURI	CDMI URI	The URI of the data or container object. For named objects, this is the named path or COID. For nameless objects, it is the COID only (/cdmi_objectid/COID). Does not contain the /CDMI root.
HSID	Session Identifier	The unique identifier assigned to the HTTP session.
OBSP	Security Partition Name	The name of the security partition assigned to the data or container object. An empty string is returned if the object is not associated with a security partition.
COID	CDMI Object Identifier	A unique value assigned at creation time to identify an object.
SPAR	Security Partition ID	The unique identifier of the security partition assigned to the object. Zero is returned if the object is not associated with a security partition. SPAR can be used to determine the security partition to which an object is associated. See Determining the security partition for an object on page 54.
RSLT	Result Code	Result of the DELETE transaction. Result is always: SUCS: successful
TIME	Time	Total processing time for the request in microseconds.

CDMG: CDMI GET Transaction

When a CDMI client issues a GET transaction, a request is made to read a data or container object. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the corresponding content block requested. If the CBID is unknown, this field is set to 0.
CSIZ	Content Size	The size of the retrieved object in bytes.
HSID	Session Identifier	The unique identifier assigned to the HTTP session.
CURI	CDMI URI	The URI of the data or container object. For named objects, this is the named path or COID. For nameless objects, it is the COID only (<code>/cdmi_objectid/COID</code>). Does not contain the <code>/CDMI</code> root.
OBSP	Security Partition Name	The name of the security partition assigned to the data or container object. An empty string is returned if the object is not associated with a security partition.
COID	CDMI Object Identifier	A unique value assigned at creation time to identify an object.
SPAR	Security Partition ID	The unique identifier of the security partition assigned to the object. Zero is returned if the object is not associated with a security partition. SPAR can be used to determine the security partition to which an object is associated. For more information, see Determining the security partition for an object on page 54.
RSLT	Result Code	Result of the GET transaction. Result is always: SUCS: successful
TIME	Time	Total processing time for the request in microseconds.

CDMP: CDMI PUT or POST Transaction to Create Object

When a CDMI client issues a PUT or POST request to create a new data or container object, this message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the corresponding content block requested. If the CBID is unknown, this field is set to 0.
CSIZ	Content Size	The size of the original content stored in bytes. Container objects do not include this field.
HSID	Session Identifier	The unique identifier assigned to the HTTP session.
CURI	CDMI URI	The URI of the data or container object. For named objects, this is the named path or COID. For nameless objects, it is the COID only (<code>/cdmi_objectid/COID</code>). Does not contain the <code>/CDMI</code> root.
OBSP	Security Partition Name	The name of the security partition assigned to the data or container object. An empty string is returned if the object is not associated with a security partition.

Code	Field	Description
COID	CDMI Object Identifier	A unique value assigned at creation time to identify an object.
SPAR	Security Partition ID	The unique identifier of the security partition assigned to the object. Zero is returned if the object is not associated with a security partition. SPAR can be used to determine the security partition to which an object is associated. For more information, see Determining the security partition for an object on page 54.
RSLT	Result Code	Result of the POST transaction. Result is always: SUCS: successful
TIME	Time	Total processing time for the request in microseconds.

CDMU: CDMI PUT Transaction to Update Object

When a CDMI client issues a PUT request to update a data or container object, this message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the corresponding content block requested. If the CBID is unknown, this field is set to 0.
CSIZ	Content Size	The size of the original content stored in bytes. Container objects do not include this field.
HSID	Session Identifier	The unique identifier assigned to the HTTP session.
CURI	CDMI URI	The URI of the data or container object. For named objects this is the named path or COID. For nameless objects, it is the COID only (<code>/cdmi_objectid/COID</code>). Does not contain the <code>/</code> CDMI root.
OBSP	Security Partition Name	The name of the security partition assigned to the data or container object. An empty string is returned if the object is not associated with a security partition.
COID	CDMI Object Identifier	A unique value assigned at creation time to identify an object.
SPAR	Security Partition ID	The unique identifier of the security partition assigned to the target object. Zero is returned if the object is not associated with a security partition. SPAR can be used to determine the security partition to which an object is associated. For more information, see Determining the security partition for an object on page 54.
RSLT	Result Code	Result of the PUT transaction. Result is always: SUCS: successful
TIME	TIME	Total processing time for the request in microseconds.

ECOC: Corrupt Erasure Coded Data Fragment

This audit message indicates that the system has detected a corrupt erasure-coded data fragment.

Code	Field	Description
VCCO	VCS ID	The name of the VCS that contains the corrupt chunk.
VLID	Volume ID	The RangeDB Volume that contains the corrupt erasure-coded fragment.
CCID	Chunk ID	The identifier of the corrupt erasure-coded fragment.
RSLT	Result	This field has the value 'NONE'. RSLT is a mandatory message field, but is not relevant for this particular message. 'NONE' is used rather than 'SUCS' so that this message is not filtered.

ETAF: Security Authentication Failed

A connection attempt using Transport Layer Security (TLS) has failed.

Code	Field	Description
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection over which the authentication failed.
RUID	User Identity	A service dependent identifier representing the identity of the remote user.
RSLT	Reason Code	The reason for the failure: SCNI: Secure connection establishment failed. CERM: Certificate was missing. CERT: Certificate was invalid. CERE: Certificate was expired. CERR: Certificate was revoked. CSGN: Certificate signature was invalid. CSGU: Certificate signer was unknown. UCRM: User credentials were missing. UCRI: User credentials were invalid. UCRU: User credentials were disallowed. TOUT: Authentication timed out.

When a connection is established to a secure service that uses TLS, the credentials of the remote entity are verified using the TLS profile and additional logic built into the service. If this authentication fails due to invalid, unexpected, or disallowed certificates or credentials, an audit message is logged. This enables queries for unauthorized access attempts and other security-related connection problems.

The message could result from a remote entity having an incorrect configuration, or from attempts to present invalid or disallowed credentials to the system. This audit message should be monitored to detect attempts to gain unauthorized access to the system.

ETCA: TCP/IP Connection Establish

When a connection to a service running on a grid node is permitted, this message is generated.

Code	Field	Description
SEID	Service Identifier	The unique identifier of the service to which the connection was established. Values of interest include: <ul style="list-style-type: none"> HING: HTTP Ingest Service HCLN: HTTP Query/Retrieve Service NCON: Neighbor Connection Service
CNDR	Connection Direction	Indicates whether the connection was opened by the grid node or by a remote host: INBO: Connection initiated by a remote host, which connected to the grid node. OUTB: Connection initiated by the grid node, which connected to a remote host.
SVIP	Destination Service Port	The port to which the connection was established.
DAIP	Destination IP Address	The IP address to which the connection was established.
SAIP	Source IP Address	The IP address from which the connection was established.
CNID	Connection Identifier	The unique identifier of the connection.
RSLT	Result Code	Connection status: SUCS: connection successfully established

This audit message means an incoming or outgoing TCP/IP connection was successfully established. This does not indicate the corresponding user was permitted to use the service — only that they were not rejected. Typically, each service implements additional authentication mechanisms specific to the service type (HTTP).

This message can be used to report on external hosts communicating with the system, and to correlate higher-level protocol messages back to the IP address initiating the activity. The “Connection Identifier” field allows correlation of audit messages related to actions performed during a session.

ETCC: TCP/IP Connection Close

When the system on either side of an established connection closes the connection (either normally or abnormally), this message is generated.

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the connection.
INIE	Initiating Entity	The entity causing the connection to be closed: LOCL: the grid node closed the connection RMOT: the remote entity closed the connection

Code	Field	Description
RSLT	Result Code	<p>Why the connection was closed:</p> <p>SUCS: connection closed at an expected point</p> <p>LOST: connection closed by the remote entity at an unexpected point</p> <p>UNEX: connection closed by the remote entity at an unexpected point</p> <p>TOUT: connection timed-out and was closed</p>

This audit message means a TCP/IP connection was closed. When this message is generated, the corresponding connection ID no longer exists, and the associated TCP/IP connection is no longer established.

This message can be used to detect problems within the system, such as network issues over a WAN, or interoperability problems between systems. The “Connection Identifier” field allows correlation of audit messages related to actions performed during a session.

ETCF: TCP/IP Connection Fail

When an attempt to establish a connection to a remote service fails during establishment, this message is generated.

Code	Field	Description
SEID	Service Identifier	<p>The unique identifier of the service to which the connection was attempted. Values of interest include:</p> <ul style="list-style-type: none"> • HING: HTTP Ingest Service • HCLN: HTTP Query/Retrieve Service • HCLN: HTTP Query/Retrieve Service
CNDR	Connection Direction	<p>Indicates whether the connection was opened by the grid node or by a remote host:</p> <p>INBO: connection initiated by a remote host connecting to the node</p> <p>OUTB: connection initiated by the grid node, attempting a connection to a remote host</p>
SVIP	Destination Service Port	The port to which the connection attempt was made.
DAIP	Destination IP Address	The IP address to which the connection attempt was made.
SAIP	Source IP Address	The IP address from which the connection attempt was made.
CNID	Connection Identifier	The unique identifier of the attempted connection.
RSLT	Result Code	<p>Why the attempted connection failed:</p> <p>IPAR: inbound IP address was not from allowed range</p> <p>CRFU: outgoing connection refused by remote host</p> <p>UNRE: destination (remote host) unreachable</p> <p>ATHF: TCP/IP connection level authentication failure</p>

This audit message means an outgoing or incoming connection attempt failed at the lowest level, due to communication problems — the corresponding service was unable to access the remote host, and the TCP/IP connection was not established.

This message can be used to detect system problems such as configuration errors where content is being pushed to unreachable hosts, or where routing problems result in inaccessibility of hosts. The message can also be used to report on the hosts to which content was pushed.

The “Connection Identifier” field allows correlation of audit messages related to actions performed during a session.

ETCR: TCP/IP Connection Refused

The Connection Refused Audit Message indicates that an incoming TCP/IP connection attempt was not allowed.

If the node refuses a connection, this message is generated. Failures of inbound connections can result from a variety of reasons, which are described in the entry below for the Result field.

Code	Field	Description
SEID	Service Identifier	The unique identifier of the service to which the connection was attempted. Values of interest include: <ul style="list-style-type: none"> • HING: HTTP Ingest Service • HCLN: HTTP Query/Retrieve Service • NCON: Neighbor Connection Service
CNDR	Connection Direction	Indicates that the connection was opened by a remote host: INBO: connection initiated by a remote host connecting to the node
SVIP	Destination Service Port	The port to which the connection attempt was made.
DAIP	Destination IP Address	The IP address to which the connection attempt was made (remote IP address).
SAIP	Source IP Address	The IP address from which the connection attempt was made (local IP address).
CNID	Connection Identifier	The unique identifier of the attempted connection.
RSLT	Result Code	Why the attempted connection was refused: IPAR: inbound IP address was not from allowed range ATHF: TCP/IP connection-level authentication failure

For incoming connections, this audit message means that a connection was not successfully established at the lowest level due to a security violation. When this message is received, the corresponding user was not able to access the service and the TCP/IP Connection was closed. The most common reporting use of this message is to detect unauthorized attempts to access services running on the system from foreign IP addresses that have not been explicitly given access to the service.

GNRG: GNDS Registration

The CMN service generates this audit message when a service has updated or registered information about itself in the StorageGRID Webscale system.

Code	Field	Description
RSLT	Result	The result of the update request: <ul style="list-style-type: none"> SUCS: Successful SUNV: Service Unavailable GERR: Other failure
GNID	Node ID	The node ID of the service that initiated the update request.
GNTTP	Device Type	The grid node's device type (for example, BLDR for an LDR service).
GNDV	Device Model version	The string identifying the grid node's device model version in the DMDL bundle.
GNGP	Group	The group to which the grid node belongs (in the context of link costs and service-query ranking).
GNIA	IP Address	The grid node's IP address.

This message is generated whenever a grid node updates its entry in the Grid Nodes Bundle.

GNUR: GNDS Unregistration

The CMN service generates this audit message when a service has unregistered information about itself from the StorageGRID Webscale system.

Code	Field	Description
RSLT	Result	The result of the update request: <ul style="list-style-type: none"> SUCS: Successful SUNV: Service Unavailable GERR: Other failure
GNID	Node ID	The node ID of the service that initiated the update request.

GTED: Grid Task Ended

This audit message indicates that the CMN service has finished processing the specified grid task and has moved the task to the Historical table. If the result is SUCS, ABRT, or ROLF, there will be a

corresponding Grid Task Started audit message. The other results indicate that processing of this grid task never started.

Code	Field	Description
TSID	Task ID	<p>This field uniquely identifies a generated grid task and allows the grid task to be managed over its lifecycle.</p> <p>Note: The Task ID is assigned at the time that a grid task is generated, not the time that it is submitted. It is possible for a given grid task to be submitted multiple times, and in this case the Task ID field is not sufficient to uniquely link the Submitted, Started, and Ended audit messages.</p>
RSLT	Result	<p>The final status result of the grid task:</p> <ul style="list-style-type: none"> • SUCS: The grid task completed successfully. • ABRT: The grid task was aborted without a rollback error. • ROLF: The grid task was aborted and was unable to complete the rollback process. • CANC: The grid task was canceled by the user before it was started. • EXPR: The grid task expired before it was started. • IVLD: The grid task was invalid. • AUTH: The grid task was unauthorized. • DUPL: The grid task was rejected as a duplicate.

GTST: Grid Task Started

This audit message indicates that the CMN service has started to process the specified grid task. The audit message immediately follows the Grid Task Submitted message for grid tasks initiated by the internal Grid Task Submission service and selected for automatic activation. For grid tasks submitted into the Pending table, this message is generated when the user starts the grid task.

Code	Field	Description
TSID	Task ID	<p>This field uniquely identifies a generated grid task and allows the task to be managed over its lifecycle.</p> <p>Note: The Task ID is assigned at the time that a grid task is generated, not the time that it is submitted. It is possible for a given grid task to be submitted multiple times, and in this case the Task ID field is not sufficient to uniquely link the Submitted, Started, and Ended audit messages.</p>
RSLT	Result	<p>The result. This field has only one value:</p> <ul style="list-style-type: none"> • SUCS: The grid task was started successfully.

GTSU: Grid Task Submitted

This audit message indicates that a grid task has been submitted to the CMN service.

Code	Field	Description
TSID	Task ID	Uniquely identifies a generated grid task and allows the task to be managed over its lifecycle. Note: The Task ID is assigned at the time that a grid task is generated, not the time that it is submitted. It is possible for a given grid task to be submitted multiple times, and in this case the Task ID field is not sufficient to uniquely link the Submitted, Started, and Ended audit messages.
TTYP	Task Type	The type of grid task.
TVER	Task Version	A number indicating the version of the grid task.
TDSC	Task Description	A human-readable description of the grid task.
VATS	Valid After Timestamp	The earliest time (UINT64 microseconds from January 1, 1970 - UNIX time) at which the grid task is valid.
VBTS	Valid Before Timestamp	The latest time (UINT64 microseconds from January 1, 1970 - UNIX time) at which the grid task is valid.
TSRC	Source	The source of the task: <ul style="list-style-type: none"> • TXTB: The grid task was submitted through the StorageGRID Webscale system as a signed text block. • GRID: The grid task was submitted through the internal Grid Task Submission Service.
ACTV	Activation Type	The type of activation: <ul style="list-style-type: none"> • AUTO: The grid task was submitted for automatic activation. • PEND: The grid task was submitted into the pending table. This is the only possibility for the TXTB source.
RSLT	Result	The result of the submission: <ul style="list-style-type: none"> • SUCS: The grid task was submitted successfully. • FAIL: The task has been moved directly to the historical table.

HTSC: HTTP Session Close

When a client finishes communicating with a remote host and closes the previously established HTTP session, this message is issued.

Code	Field	Description
HSID	Session Identifier	The unique identifier assigned to the HTTP session established to the node.

Code	Field	Description
RSLT	Result Code	<p>Why the session was closed:</p> <p>SUCS: session closed normally, without errors</p> <p>TOUT: timed-out by the node, due to inactivity</p> <p>ERRC: lost the connection over which the session was established</p> <p>ERRT: session terminated due to an error occurring on a transaction</p> <p>AUTH: session terminated due to a failed transaction authorization</p> <p>GERR: a general error occurred, causing the session to close</p>

“HTTP Session Close” always corresponds with a previously issued “HTTP Session Establish” message.

This message should be monitored to determine if there are any repetitive or excessive problems in attempting to establish a session. This can indicate potential communications or interoperability problems related to a client or server.

HTSE: HTTP Session Establish

When a client establishes an HTTP session, this message is issued.

Code	Field	Description
CNID	Connection Identifier	The unique identifier for the connection over which the HTTP session was established.
HSID	Session Identifier	The unique identifier assigned to the HTTP session established to the node.
OBCL	Client Name	The user-defined security partition client name assigned to the client certificate. An empty string is returned if the client has not been defined.
RSLT	Result Code	<p>Status at the time the session was established:</p> <p>SUCS: session successfully established.</p>

LLST: Location Lost

The LLST message is generated whenever a location for an object copy (replicated or erasure coded) cannot be found.

Code	Field	Description
CBIL	CBID	The affected CBID.
NOID	Source Node ID	The node ID on which the locations were lost.
ECPR	Erasure Coding Profile	For erasure-coded object data. The ID of the Erasure Coding Profile used.
LTYP	Location Type	<p>CLDI (Online): For replicated object data</p> <p>CLEC (Online): For erasure-coded object data</p> <p>CLNL (Nearline): For archived replicated object data</p>

Code	Field	Description
RSLT	Result	Always NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.
TSRC	Triggering Source	USER: User triggered SYST: System triggered

MGAU: Management audit message

The Management category logs user requests to the Management API. Every request that is not a GET or HEAD request to the API logs a response with the username, IP, and type of request to the API.

Code	Field	Description
MDIP	Destination IP Address	The server (destination) IP address.
MDNA	Domain name	The host domain name.
MPAT	Request PATH	The request path.
MRBD	Request body	<p>The content of the request body. While the response body is logged by default, the request body is logged in certain cases when the response body is empty. Because the following information is not available in the response body, it is taken from the request body for the following POST methods:</p> <ul style="list-style-type: none"> Username and account ID in POST authorize New subnets configuration in POST /grid/grid-networks/update New NTP servers in POST /grid/ntp-servers/update Decommissioned server IDs in POST /grid/servers/decommission <p>Note: Sensitive information is either deleted (for example, an S3 access key) or masked with asterisks (for example, a password).</p>
MRMD	Request method	<p>The HTTP request method:</p> <ul style="list-style-type: none"> POST PUT DELETE PATCH
MRSC	Response code	The response code.

Code	Field	Description
MRSP	Response body	The content of the response (the response body) is logged by default. Note: Sensitive information is either deleted (for example, an S3 access key) or masked with asterisks (for example, a password).
MSIP	Source IP address	The client (source) IP address.
MUUN	User URN	The URN (uniform resource name) of the user who sent the request.
RSLT	Result	Returns successful (SUCS) or the error reported by the backend.

OLST: System Detected Lost Object

This message is generated when the DDS service cannot locate any copies of an object within the StorageGRID Webscale system.

Code	Field	Description
CBID	Content Block Identifier	The CBID of the lost object.
UUID	Universal Unique ID	The identifier of the lost object within the StorageGRID Webscale system.
NOID	Node ID	The last known direct or nearline location of the lost object. It is possible to have no last known location if the location information is not available. It is possible to have just the Node ID without a Volume ID if the volume information is not available.
VOLI	Volume ID	The Volume ID of the Storage Node or Archive Node for the last known location of the lost object.
RSLT	Result	This field has the value NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.

ORLM: Object Rules Met

This message is generated when the object is successfully stored and copied as specified by the ILM rules.

Code	Field	Description
CBID	Content Block Identifier	The CBID of the object.
RULE	Rules Label	The human-readable label given to the ILM rule applied to this object.
FSIZ	File size	Indicates the size of the object in bytes.
UUID	Associated UUID	The identifier of the object within the StorageGRID Webscale system. The value for UUID is "" if the object does not have a UUID, content handle has been released.

Code	Field	Description
SGCB	Container CBID	CBID of the container for the segmented object. This value is available only if the object is segmented.
SEGC	Container UUID	UUID of the container for the segmented object. This value is available only if the object is segmented.
SPAR	Security Partition ID	The unique identifier of the security partition assigned to the target object. Zero is returned if the object is not associated with an HTTP security partition. SPAR can be used to determine the HTTP security partition to which an object is associated. For more information, see Determining the security partition for an object on page 54.
LOCS	Locations	The storage location of object data within the StorageGRID Webscale system. The value for LOCS is "" if the object has no locations (for example, it has been deleted). CLEC: ID of the Erasure Coding profile and group applied to the object's data if object data is erasure coded. CLDI: LDR and volume ID of the object's location if object data is replicated. CLNL: ARC ID of the object's location if the object data is archived.
STAT	Status	The status of ILM operation. DONE: ILM operations against the object have completed. DFER: The object has been marked for future ILM re-evaluation. PRGD: The object has been deleted from the StorageGRID Webscale system. NLOC: The object was deleted from the StorageGRID Webscale system without the CMS service being involved. This typically happens when ingest fails.
RSLT	Result	The result of the ILM operation. SUCS: The ILM operation was successful.

The ORLM audit message can be issued a number of times for a single object. For instance, it is issued whenever one of the following events take place:

- ILM rules for the object are satisfied forever.
- ILM rules for the object are satisfied for this epoch.
- ILM rules have deleted the object.
- The background verification process detects that a copy of replicated object data is corrupt. The CMS service then preforms an ILM evaluation to replace the corrupt object.

Related references

[Object ingest transactions](#) on page 15

[Object delete transactions](#) on page 16

SADD: Security Audit Disable

This message indicates that the originating service (node ID) has turned off audit message logging; audit messages are no longer being collected or delivered.

Code	Field	Description
AETM	Enable Method	The method used to disable the audit.
AEUN	User Name	The user name that executed the command to disable audit logging.
RSLT	Result	This field has the value NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.

The message implies that logging was previously enabled, but has now been disabled. This is typically used only during bulk ingest to improve system performance. Following the bulk activity, auditing is restored (SADE) and the capability to disable auditing is then permanently blocked.

SADE: Security Audit Enable

This message indicates that the originating service (node ID) has restored audit message logging; audit messages are again being collected and delivered.

Code	Field	Description
AETM	Enable Method	The method used to enable the audit.
AEUN	User Name	The user name that executed the command to enable audit logging.
RSLT	Result	This field has the value NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.

The message implies that logging was previously disabled (SADD), but has now been restored. This is typically only used during bulk ingest to improve system performance. Following the bulk activity, auditing is restored and the capability to disable auditing is then permanently blocked.

SCMT: Object Store Commit

Grid content is not made available or recognized as stored until it has been committed (meaning it has been stored persistently). Persistently stored content has been completely written to disk, and has passed related integrity checks. When a content block is committed to storage, this message is issued.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block committed to permanent storage.
RSLT	Result Code	Status at the time the object was stored to disk: SUCS: Object successfully stored.

This message means a given content block has been completely stored and verified, and can now be requested. It can be used to track data flow within the system.

SDEL: S3 DELETE

When an S3 client issues a DELETE transaction, a request is made to remove the specified object or bucket. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the corresponding content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets do not include this field.
CSIZ	Content Size	The size of the deleted object in bytes. Operations on buckets do not include this field.
DMRK	Delete Marker Version ID	The version ID of the delete marker created when deleting an object from a versioned bucket. Operations on buckets do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.
S3AK	S3 Access Key ID	The S3 access key ID as specified by the StorageGRID Webscale management API that uniquely identifies the client performing the transaction. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets do not include this field.
SACC	S3 Account ID	The S3 account user ID.
SAIP	IP address of requesting client	The IP address of the client application that made the request.
SUSR	S3 User Name	The S3 account and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: urn:sgws:identity::03393893651506583485:root
RSLT	Result Code	Result of the GET transaction. Result is always: SUCS: Successful
TIME	Time	Total processing time for the request in microseconds.
VSID	Version ID	The version ID of the specific version of an object that was deleted. Operations on buckets and objects in unversioned buckets do not include this field.

SGET: S3 GET

When an S3 client issues a GET transaction, a request is made to retrieve an object or list the objects in a bucket. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the corresponding content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets do not include this field.

Code	Field	Description
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on buckets do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.
S3AI	S3 Account ID	The unique account ID as specified by the StorageGRID Webscale system. An empty value indicates anonymous access.
S3AK	S3 Access Key ID	The S3 access key ID as specified by the StorageGRID Webscale management API that uniquely identifies the client performing the transaction. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets do not include this field.
SACC	S3 Account ID	The S3 account user ID.
SAIP	IP address of requesting client	The IP address of the client application that made the request.
SUSR	S3 User Name	The S3 account and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: <code>urn:sgws:identity::03393893651506583485:root</code>
RSLT	Result Code	Result of the GET transaction. Result is always: SUCS: Successful
TIME	Time	Total processing time for the request in microseconds.
VSID	Version ID	The version ID of the specific version of an object that was requested. Operations on buckets and objects in unversioned buckets do not include this field.

SHEA: S3 HEAD

When an S3 client issues a HEAD transaction, a request is made to check for the existence of an object or bucket and retrieve the metadata about an object. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the corresponding content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets do not include this field.
CSIZ	Content Size	The size of the checked object in bytes. Operations on buckets do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.
S3AI	S3 Account ID	The unique account ID as specified by the StorageGRID Webscale system. An empty value indicates anonymous access.

Code	Field	Description
S3AK	S3 Access Key ID	The S3 access key ID as specified by the StorageGRID Webscale management API that uniquely identifies the client performing the transaction. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets do not include this field.
SACC	S3 Account ID	The S3 account user ID.
SAIP	IP address of requesting client	The IP address of the client application that made the request.
SUSR	S3 User Name	The S3 account and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: <code>urn:sgws:identity::03393893651506583485:root</code>
RSLT	Result Code	Result of the GET transaction. Result is always: SUCS: Successful
TIME	Time	Total processing time for the request in microseconds.
VSID	Version ID	The version ID of the specific version of an object that was requested. Operations on buckets and objects in unversioned buckets do not include this field.

SPUT: S3 PUT

When an S3 client issues a PUT transaction, a request is made to create a new object or bucket. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the corresponding content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets do not include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on buckets do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.
S3AI	S3 Account ID	The unique account ID as specified by the StorageGRID Webscale system. An empty value indicates anonymous access.
S3AK	S3 Access Key ID	The S3 access key ID as specified by the StorageGRID Webscale management API that uniquely identifies the client performing the transaction. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3KY	The S3 key name, not including the bucket name. Operations on buckets do not include this field.
SACC	S3 Account ID	The S3 account user ID.

Code	Field	Description
SAIP	IP address of requesting client	The IP address of the client application that made the request.
SUSR	S3 User Name	The S3 account and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: <code>urn:sgws:identity::03393893651506583485:root</code>
RSLT	Result Code	Result of the GET transaction. Result is always: SUCS: Successful
TIME	Time	Total processing time for the request in microseconds.
VSID	Version ID	The version ID of a new object created in a versioned bucket. Operations on buckets and objects in unversioned buckets do not include this field.
VSST	Versioning State	The new versioning state of a bucket. Two states are used: "enabled" or "suspended". Operations on objects do not include this field.

SREM: Object Store Remove

This message is issued when content is removed from persistent storage and is no longer accessible through regular APIs.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block deleted from permanent storage.
RSLT	Result Code	Indicates the result of the content removal operations. The only defined value is: SUCS: Content removed from persistent storage

This audit message means a given content block has been deleted from a node and can no longer be requested directly. The message can be used to track the flow of deleted content within the system.

SUPD: S3 Metadata Updated

This message is generated by the S3 API after processing a metadata update for an ingested object.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the corresponding content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets do not include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on buckets do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.
S3AI	S3 Account ID	The unique account ID as specified by the StorageGRID Webscale system. An empty value indicates anonymous access.

Code	Field	Description
S3AK	S3 Access Key ID	The S3 access key ID as specified by the StorageGRID Webscale management API that uniquely identifies the client performing the transaction. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets do not include this field.
SACC	S3 Account ID	The S3 account user ID.
SAIP	IP address of requesting client	The IP address of the client application that made the request.
SUSR	S3 User Name	The S3 account and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: <code>urn:sgws:identity::03393893651506583485:root</code>
RSLT	Result Code	Result of the GET transaction. Result is always: SUCS: successful
TIME	Time	Total processing time for the request in microseconds.

SVRF: Object Store Verify Fail

This message is issued whenever a content block fails the verification process. Each time replicated object data is read from or written to disk, several verification and integrity checks are performed to ensure the data sent to the requesting user is identical to the data originally ingested into the system. If any of these checks fail, the system automatically quarantines the corrupt replicated object data to prevent it from being retrieved again.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block which failed verification.
RSLT	Result Code	Verification failure type: CRCF: Cyclic redundancy check (CRC) failed. HMAC: Hash-based message authentication code (HMAC) check failed. EHSB: Unexpected encrypted content hash. PHSH: Unexpected original content hash. SEQC: Incorrect data sequence on disk. PERR: Invalid structure of disk file. DERR: Disk error. FNAM: Bad file name.

Note: This message should be monitored closely. Content verification failures can indicate attempts to tamper with content or impending hardware failures.

To determine what operation triggered the message, see the value of the AMID (Module ID) field. For example, an SVFY value indicates that the message was generated by the Storage Verifier module, that is, background verification, and STOR indicates that the message was triggered by content retrieval.

SVRU: Object Store Verify Unknown

The LDR service's Storage component continuously scans all copies of replicated object data in the object store. This message is issued when an unknown or unexpected copy of replicated object data is detected in the object store and moved to the quarantine directory.

Code	Field	Description
RSLT	Result	This field has the value 'NONE'. RSLT is a mandatory message field, but is not relevant for this message. 'NONE' is used rather than 'SUCS' so that this message is not filtered.

Note: The SVRU: Object Store Verify Unknown audit message should be monitored closely. It means unexpected copies of object data were detected in the object store. This situation should be investigated immediately to determine how these copies were created, because it can indicate attempts to tamper with content or impending hardware failures.

SYSD: Node Stop

When a service is stopped gracefully, this message is generated to indicate the shutdown was requested. Typically this message is sent only after a subsequent restart, because the audit message queue is not cleared prior to shutdown. Look for the SYST message, sent at the beginning of the shutdown sequence, if the service has not restarted.

Code	Field	Description
RSLT	Clean Shutdown	The nature of the shutdown: SUCS: System was cleanly shutdown.

The message does not indicate if the host server is being stopped, only the reporting service. The RSLT of a SYSD cannot indicate a “dirty” shutdown, because the message is generated only by “clean” shutdowns.

SYST: Node Stopping

When a service is gracefully stopped, this message is generated to indicate the shutdown was requested and that the service has initiated its shutdown sequence. SYST can be used to determine if the shutdown was requested, before the service is restarted (unlike SYSD, which is typically sent after the service restarts.)

Code	Field	Description
RSLT	Clean Shutdown	The nature of the shutdown: SUCS: System was cleanly shutdown.

The message does not indicate if the host server is being stopped, only the reporting service. The RSLT code of a SYST message cannot indicate a “dirty” shutdown, because the message is generated only by “clean” shutdowns.

SYSU: Node Start

When a service is restarted, this message is generated to indicate if the previous shutdown was clean (commanded) or disorderly (unexpected).

Code	Field	Description
RSLT	Clean Shutdown	The nature of the shutdown: SUCS: System was cleanly shut down. DSDN: System was not cleanly shut down. VRGN: System was started for the first time after server installation (or re-installation).

The message does not indicate if the host server was started, only the reporting service. This message can be used to:

- Detect discontinuity in the audit trail.
- Determine if a service is failing during operation (as the distributed nature of the StorageGRID Webscale system can mask these failures). Server Manager restarts a failed service automatically.

VLST: User Initiated Volume Lost

This message is issued whenever the `/proc/CMSI/Volume_Lost` command is run.

Code	Field	Description
VOLL	Volume Identifier Lower	The lower end of the affected volume range or a single volume.
VOLU	Volume Identifier Upper	The upper end of the affected volume range. Equal to VOLL if a single volume.
NOID	Source Node ID	The node ID on which the locations were lost.
LTYP	Location Type	'CLDI' (Online) or 'CLNL' (Nearline). If not specified, defaults to 'CLDI'.
RSLT	Result	Always 'NONE'. RSLT is a mandatory message field, but is not relevant for this message. 'NONE' is used rather than 'SUCS' so that this message is not filtered.

WDEL: Swift DELETE

When a Swift client issues a DELETE transaction, a request is made to remove the specified object or container. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the corresponding content block requested. If the CBID is unknown, this field is set to 0. Operations on containers do not include this field.
CSIZ	Content Size	The size of the deleted object in bytes. Operations on containers do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.
SAIP	IP address of requesting client	The IP address of the client application that made the request.

Code	Field	Description
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID Webscale system.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.
WCON	Swift Container	The Swift container name.
WOBJ	Swift Object	The Swift object identifier. Operations on containers do not include this field.
RSLT	Result Code	Result of the DELETE transaction. Result is always: SUCS: Successful
TIME	Time	Total processing time for the request in microseconds.

WGET: Swift GET

When a Swift client issues a GET transaction, a request is made to retrieve an object, list the objects in a container, or list the containers in an account. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the corresponding content block requested. If the CBID is unknown, this field is set to 0. Operations on accounts and containers do not include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on accounts and containers do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.
SAIP	IP address of requesting client	The IP address of the client application that made the request.
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID Webscale system.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.
WCON	Swift Container	The Swift container name. Operations on accounts do not include this field.
WOBJ	Swift Object	The Swift object identifier. Operations on accounts and containers do not include this field.
RSLT	Result Code	Result of the GET transaction. Result is always: SUCS: successful
TIME	Time	Total processing time for the request in microseconds.

WHEA: Swift HEAD

When a Swift client issues a HEAD transaction, a request is made to check for the existence of an account, container, or object, and retrieve any relevant metadata. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the corresponding content block requested. If the CBID is unknown, this field is set to 0. Operations on accounts and containers do not include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on accounts and containers do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.
SAIP	IP address of requesting client	The IP address of the client application that made the request.
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID Webscale system.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.
WCON	Swift Container	The Swift container name. Operations on accounts do not include this field.
WOBJ	Swift Object	The Swift object identifier. Operations on accounts and containers do not include this field.
RSLT	Result Code	Result of the HEAD transaction. Result is always: SUCS: successful
TIME	Time	Total processing time for the request in microseconds.

WPUT: Swift PUT

When a Swift client issues a PUT transaction, a request is made to create a new object or container. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the corresponding content block requested. If the CBID is unknown, this field is set to 0. Operations on containers do not include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on containers do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.
SAIP	IP address of requesting client	The IP address of the client application that made the request.
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID Webscale system.

Code	Field	Description
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.
WCON	Swift Container	The Swift container name.
WOBJ	Swift Object	The Swift object identifier. Operations on containers do not include this field.
RSLT	Result Code	Result of the PUT transaction. Result is always: SUCS: successful
TIME	Time	Total processing time for the request in microseconds.

Determining the security partition for an object

You can use the security partition ID number listed in the SPAR field to determine the security partition with which an object is associated. This helps to ensure that the object is placed in the correct security partition when the object is evaluated against information lifecycle management (ILM) rules. Security partitions are ignored for objects ingested through either the S3 or Swift APIs.

Before you begin

- You must be signed in to the Grid Management Interface using a supported browser.
- To perform this task, you need specific access permissions. For details, see information about controlling system access with administration user accounts and groups.

Steps

1. Obtain the security partition ID from the SPAR field:

Example

```
[ SPAR(UI64) : 5064106809552273409 ]
```

If security partitioning is disabled for the StorageGRID Webscale system, SPAR is zero.

2. Convert the security partition ID to a hexadecimal value.

For example, 5064106809552273409 becomes 0x4647525000000001.

3. Interpret the result of the conversion to a hexadecimal value.

The first eight numbers determine the character code: HTTP (HTTP Security Partition) would be 0x48545450. The remainder of the hexadecimal number is the partition identifier; for example, 0x4854545000000001 means that security partition identifier 1 is associated with the security partition ID 325742710709288961.

4. Select **Configuration > CDMI**.
5. Select **Security Partitions**.
6. Under **Partitions**, determine the partition identifier associated with the partition number as detailed in Step 3.

Related information

[StorageGRID Webscale 10.4 Administrator Guide](#)

Glossary

ACL

Access control list. Specifies which users or groups of users are allowed to access an object and what operations are permitted, for example, read, write, and execute.

active-backup mode

A method for bonding two physical ports together for redundancy.

ADC service

Administrative Domain Controller. The ADC service maintains topology information, provides authentication services, and responds to queries from the LDR, CMN, and CLB services. The ADC service is present on each of the first three Storage Nodes installed at a site.

ADE

Asynchronous Distributed Environment. Proprietary development environment used as a framework for services within the StorageGRID Webscale system.

Admin Node

The Admin Node provides services for the web interface, system configuration, and audit logs. See also, *primary Admin Node*.

Amazon S3

Proprietary web service from Amazon for the storage and retrieval of data.

AMS service

Audit Management System. The AMS service monitors and logs all audited system events and transactions to a text log file. The AMS service is present on the Admin Node.

API Gateway Node

An API Gateway Node provides load balancing functionality to the StorageGRID Webscale system and is used to distribute the workload when multiple client applications are performing ingest and retrieval operations. API Gateway Nodes include a Connection Load Balancer (CLB) service.

ARC service

Archive. The ARC service provides the management interface with which you configure connections to external archival storage such as the cloud through an S3 interface or tape through TSM middleware. The ARC service is present on the Archive Node.

Archive Node

The Archive Node manages the archiving of object data to an external archival storage system.

atom

Atoms are the lowest level component of the container data structure, and generally encode a single piece of information.

audit message

Information about an event occurring in the StorageGRID Webscale system that is captured and logged to a file.

Base64

A standardized data encoding algorithm that enables 8-bit data to be converted into a format that uses a smaller character set, enabling it to safely pass through legacy systems

that can process only basic (low order) ASCII text excluding control characters. See RFC 2045 for more details.

bundle

A structured collection of configuration information used internally by various components of the StorageGRID Webscale system. Bundles are structured in container format.

Cassandra

An open-source database that is scalable and distributed, provides high availability, and handles large amounts of data across multiple servers.

CBID

Content Block Identifier. A unique internal identifier of a piece of content within the StorageGRID Webscale system.

CDMI

Cloud Data Management Interface. An industry-standard defined by SNIA that includes a RESTful interface for object storage. For more information, see www.snia.org/cdm.

CIDR

Classless Inter-Domain Routing. A notation used to compactly describe a subnet mask used to define a range of IP addresses. In CIDR notation, the subnet mask is expressed as an IP address in dotted decimal notation, followed by a slash and the number of bits in the subnet. For example, 192.0.2.0/24.

CLB service

Connection Load Balancer. The CLB service provides a gateway into the StorageGRID Webscale system for client applications connecting through HTTP. The CLB service is part of the API Gateway Node.

Cloud Data Management Interface

See *CDMI*.

CMN service

Configuration Management Node. The CMN service manages system-wide configurations and grid tasks. The CMN service is present on the primary Admin Node.

CMS service

Content Management System. The CMS service carries out the operations of the active ILM policy's ILM rules, determining how object data is protected over time. The CMS service is present on the Storage Node.

command

In HTTP, an instruction in the request header such as GET, HEAD, DELETE, OPTIONS, POST, or PUT. Also known as an HTTP method.

container

Created when an object is split into segments. A container object lists the header information for all segments of the split object and is used by the LDR service to assemble the segmented object when it is retrieved by a client application.

content block ID

See *CBID*.

content handle

See *UUID*.

CSTR

Null-terminated, variable-length string.

DC

Data Center site.

DDS service

Distributed Data Store. The DDS service interfaces with the distributed key-value store and manages object metadata. It distributes metadata copies to multiple instances of the distributed key-value store so that metadata is always protected against loss.

distributed key value store

Data storage and retrieval that unlike a traditional relational database manages data across grid nodes.

DNS

Domain Name System.

enablement layer

Used during installation to customize the Linux operating system installed on each grid node. Only the packages needed to support the services hosted on the grid node are retained, which minimizes the overall footprint occupied by the operating system and maximizes the security of each grid node.

Fibre Channel

A networking technology primarily used for storage.

Grid ID signed text block

A Base64 encoded block of cryptographically signed data that contains the grid ID. See also, *provisioning*.

grid node

The basic software building block for the StorageGRID Webscale system, for example, Admin Node or Storage Node. Each grid node type consists of a set of services that perform a specialized set of tasks.

grid task

System-wide scripts used to trigger various actions that implement specific changes to the StorageGRID Webscale system. For example, most maintenance and expansion procedures involve running grid tasks. Grid tasks are typically long-term operations that span many entities within the StorageGRID Webscale system. See also, *Task Signed Text Block*.

ILM

Information Lifecycle Management. A process of managing content storage location and duration based on content value, cost of storage, performance access, regulatory compliance, and other factors. See also, *Admin Node* and *storage pool*.

LACP

Link Aggregation Control Protocol. A method for bundling two or more physical ports together to form a single logical channel.

LAN

Local Area Network. A network of interconnected computers that is restricted to a small area, such as a building or campus. A LAN can be considered a node to the Internet or other wide area network.

latency

Time duration for processing a transaction or transmitting a unit of data from end to end. When evaluating system performance, both throughput and latency need to be considered. See also, *throughput*.

LDR service

Local Distribution Router. The LDR service manages the storage and transfer of content within the StorageGRID Webscale system. The LDR service is present on the Storage Node.

LUN

See *object store*.

mDNS

Multicast Domain Name System. A system for resolving IP addresses in a small network where no DNS server has been installed.

metadata

Information related to or describing an object stored in the StorageGRID Webscale system; for example, ingest time.

MLAG

Multi-Chassis Link Aggregation Group. A type of link aggregation group that uses two (and sometimes more) switches to provide redundancy in case one of the switches fails.

MTU

Maximum transmission unit. The largest size packet or frame that can be sent in any transmission.

namespace

A set whose elements are unique names. There is no guarantee that a name in one namespace is not repeated in a different namespace.

nearline

A term describing data storage that is neither “online” (implying that it is instantly available, like spinning disk) nor “offline” (which can include offsite storage media). An example of a nearline data storage location is a tape that is loaded in a tape library, but is not mounted.

NFS

Network File System. A protocol (developed by SUN Microsystems) that enables access to network files as if they were on local disks.

NMS service

Network Management System. The NMS service provides a web-based interface for managing and monitoring the StorageGRID Webscale system. The NMS service is present on the Admin Node. See also, *Admin Node*.

node ID

An identification number assigned to a service within the StorageGRID Webscale system. Each service (such as an NMS service or ADC service) must have a unique node ID. The number is set during system configuration and tied to authentication certificates.

NTP

Network Time Protocol. A protocol used to synchronize distributed clocks over a variable latency network, such as the Internet.

object

An artificial construct used to describe a system that divides content into data and metadata.

object segmentation

A StorageGRID Webscale process that splits a large object into a collection of small objects (segments) and creates a segment container to track the collection. The segment container contains the UUID for the collection of small objects as well as the header

information for each small object in the collection. All of the small objects in the collection are the same size. See also, *segment container*.

object storage

An approach to storing data where the data is accessed by unique identifiers and not by a user-defined hierarchy of directories and files. Each object has both data (for example, a picture) and metadata (for example, the date the picture was taken). Object storage operations act on entire objects as opposed to reading and writing bytes as is commonly done with files, and provided via APIs or HTTP instead of NAS (CIFS/NFS) or block protocols (iSCSI/ FC/FCOE).

object store

A configured file system on a disk volume. The configuration includes a specific directory structure and resources initialized at system installation.

OID

Object Identifier. The unique identifier of an object.

primary Admin Node

Admin Node that hosts the CMN service. Each StorageGRID Webscale system has only one primary Admin Node. See also, *Admin Node*.

provisioning

The process of generating a new or updated Recovery Package and GPT repository. See also, *SAID*.

quorum

A simple majority: $50\% + 1$. Some system functionality requires a quorum of the total number of a particular service type.

Recovery Package

A .zip file containing deployment-specific files and software needed to install, expand, upgrade, and maintain a StorageGRID Webscale system. The package also contains system-specific configuration and integration information, including server hostnames and IP addresses, and highly confidential passwords needed during system maintenance, upgrade, and expansion. See also, *SAID*.

SAID

Software Activation and Integration Data. The component in the Recovery Package that includes the `Passwords.txt` file.

SATA

Serial Advanced Technology Attachment. A connection technology used to connect server and storage devices.

SCSI

Small Computer System Interface. A connection technology used to connect servers and peripheral devices, such as storage systems.

segment container

An object created by the StorageGRID Webscale system during the segmentation process. Object segmentation splits a large object into a collection of small objects (segments) and creates a segment container to track the collection. A segment container contains the UUID for the collection of segmented objects as well as the header information for each segment in the collection. When assembled, the collection of segments creates the original object. See also, *object segmentation*.

server

Used when specifically referring to hardware. Might also refer to a virtual machine.

service

A unit of the StorageGRID Webscale system, such as the ADC service, NMS service, or SSM service. Each service performs unique tasks critical to the normal operations of a StorageGRID Webscale system.

SQL

Structured Query Language. An industry-standard interface language for managing relational databases. An SQL database is one that supports the SQL interface.

ssh

Secure Shell. A UNIX shell program and supporting protocols used to log in to a remote computer and run commands over an authenticated and encrypted channel.

SSL

Secure Socket Layer. The original cryptographic protocol used to enable secure communications over the Internet. See also, *TLS*.

SSM service

Server Status Monitor. A component of the StorageGRID Webscale software that monitors hardware conditions and reports to the NMS service. Every grid node runs an instance of the SSM service.

Storage Node

The Storage Node provides storage capacity and services to store, move, verify, and retrieve objects stored on disks.

storage pool

The element of an ILM rule that determines the location where an object is stored.

storage volume

See *object store*

StorageGRID

A registered trademark of NetApp, Inc., used for an object storage grid architecture and software system.

Task Signed Text Block

A Base64 encoded block of cryptographically signed data that provides the set of instructions that define a grid task.

TCP/IP

Transmission Control Protocol/Internet Protocol. A process of encapsulating and transmitting packet data over a network. It includes positive acknowledgment of transmissions.

throughput

The amount of data that can be transmitted or the number of transactions that can be processed by a system or subsystem in a given period of time. See also, *latency*.

Tivoli Storage Manager

IBM storage middleware product that manages storage and retrieval of data from removable storage resources.

TLS

Transport Layer Security. A cryptographic protocol used to enable secure communications over the Internet. See RFC 2246 for more details.

transfer syntax

The parameters, such as the byte order and compression method, needed to exchange data between systems.

URI

Universal Resource Identifier. A generic set of all names or addresses used to refer to resources that can be served from a computer system. These addresses are represented as short text strings.

UTC

A language-independent international abbreviation, UTC is neither English nor French. It means both “Coordinated Universal Time” and “Temps Universel Coordonné.” UTC refers to the standard time common to every place in the world.

UUID

Universally Unique Identifier. Unique identifier for each piece of content in the StorageGRID Webscale system. UUIDs provide client applications with a content handle that permits them to access content in a way that does not interfere with the StorageGRID Webscale system’s management of that same content. A 128-bit number that is guaranteed to be unique. See RFC 4122 for more details.

virtual machine (VM)

A software platform that enables the installation of an operating system and software, substituting for a physical server and permitting the sharing of physical server resources among several virtual servers.

VLAN

Virtual local area network (or virtual LAN). A group of devices that are located on different LAN segments but are configured to communicate as if they were attached to the same network switch.

WAN

Wide area network. A network of interconnected computers that covers a large geographic area, such as a country.

XFS

A scalable, high-performance journaled file system originally developed by Silicon Graphics.

XML

Extensible Markup Language. A text format for the extensible representation of structured information; classified by type and managed like a database. XML has the advantages of being verifiable, human readable, and easily interchangeable between different systems.

Copyright information

Copyright © 1994–2017 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

Active IQ, AltaVault, Arch Design, ASUP, AutoSupport, Campaign Express, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Element, Fitness, Flash Accel, Flash Cache, Flash Pool, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, Fueled by SolidFire, GetSuccessful, Helix Design, LockVault, Manage ONTAP, MetroCluster, MultiStore, NetApp, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANscreen, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, SolidFire Helix, StorageGRID, SyncMirror, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[*doccomments@netapp.com*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

ACPT
 See Archive Purge from Cloud-Tier

Archive Object Remove (AREM)
 description of [26](#)

Archive Object Retrieve Begin (ARCB)
 description of [24](#)

Archive Object Retrieve End (ARCE)
 description of [25](#)

Archive Object Store Begin (ATCE)
 description of [27](#)

Archive Object Store End (ASCE)
 description of [26](#)

Archive Purge from Cloud-Tier (APCT)
 description of [24](#)

Archive Retrieve from Cloud-Tier (ARCT)
 description of [25](#)

Archive Store Cloud-Tier (ASCT)
 description of [27](#)

Archive Validate Cloud-Tier Configuration (AVCC)
 description of [28](#)

audit log files
 access [8](#)
 description of [8](#)
 format of [9](#)
 names of [8](#)

audit log, messages
 system events and [11](#)

audit logs
 message levels of [7](#)

audit message levels
 changing [8](#)

audit messages
 Archive Purge from Cloud-Tier (APCT) [24](#)
 character codes for [20](#)
 common elements [11](#)
 data types [10](#)
 duplicates [5](#)
 examples of [12](#)
 flow [5](#)
 format of [9](#)
 listing of object storage [22](#)
 log file format of [9](#)
 metadata updates [18](#)
 object lifecycle and [14](#)
 overview [5](#)
 relay service [5](#)
 result [11](#)
 retention [5](#)
 timestamp [11](#)

audit messages, categories [20](#)

audit messages, HTTP protocol
 listing of [23](#)

audit messages, object deletion
 list of [16](#)

audit messages, object ingest
 list of [15](#)

audit messages, object retrieval

 list of [17](#)

audit messages, system
 table of [20, 24](#)

C

categories, audit messages [20](#)

CDMI Delete Transaction (CDMD)
 description of [30](#)

CDMI GET Transaction (CDMG)
 description of [31](#)

CDMI PUT or POST Transaction to Create Object (CDMP)
 description of [31](#)

CDMI PUT Transaction to Update Object (CDMU)
 description of [32](#)

character codes
 for audit messages [20](#)

client delete transactions
 list of audit messages generated by [16](#)

client ingest transactions
 list of audit messages generated by [15](#)

client retrieval transactions
 list of audit messages generated by [17](#)

comments
 how to send feedback about documentation [64](#)

Corrupt Erasure Coded Data Fragment (ECOC)
 description of [33](#)

D

data types
 description of [10](#)

delete audit messages
 list of [16](#)

delete transactions
 list of audit messages generated by [16](#)

documentation
 how to receive automatic notification of changes to [64](#)
 how to send feedback about [64](#)

F

feedback
 how to send comments about documentation [64](#)

formats
 of audit log files [9](#)
 of audit messages [9](#)
 of log files [9](#)

G

GNDS Registration (GNRG)
 description of [37](#)

GNDS Unregistration (GNUR)
 description of [37](#)

Grid Task Ended (GTED)

description of [37](#)
 Grid Task Started (GTST)
 description of [38](#)
 Grid Task Submitted (GTSU)
 description of [39](#)

H

HTTP protocol audit messages
 listing of [23](#)
 HTTP Session Close (HTSC)
 description of [39](#)
 HTTP Session Establish (HTSE)
 description of [40](#)

I

information
 how to send feedback about improving
 documentation [64](#)
 ingest audit messages
 list of [15](#)
 ingest transactions
 list of audit messages generated by [15](#)

L

levels
 changing for audit messages [8](#)
 Location Lost (LLST)
 description of [40](#)
 log files
 formats of [9](#)
 log files, audit
 access [8](#)
 format of [9](#)
 names of [8](#)
 logs, audit
 adjustment of message levels [7](#)

M

Management audit message (MGAU)
 description of [41](#)
 message levels
 adjustment for audit logs [7](#)
 messages, audit
 Archive Purge from Cloud-Tier (APCT) [24](#)
 changing levels for [8](#)
 examples of [12](#)
 format of [9](#)
 log file format of [9](#)
 messages, audit log
 system events and [11](#)
 messages, HTTP protocol audit
 listing of [23](#)
 messages, system audit
 table of [20](#), [24](#)
 metadata updates
 example [19](#)
 for audit messages [18](#)

N

Node Start (SYSU)
 description of [51](#)
 Node Stop (SYSD)
 description of [50](#)
 Node Stopping (SYST)
 description of [50](#)

O

object deletion
 example [17](#)
 object deletion transactions
 list of audit messages generated by [16](#)
 object ingest
 example [15](#)
 object ingest transactions
 list of audit messages generated by [15](#)
 object lifecycle
 audit messages and [14](#)
 Object Receive End (CBRE)
 description of [28](#)
 object retrieval
 example [18](#)
 object retrieval transactions
 list of audit messages generated by [17](#)
 Object Rules Met (ORLM)
 description of [42](#)
 Object Send End (CBSE)
 description of [29](#)
 object storage audit messages
 listing of [22](#)
 Object Store Commit (SCMT)
 description of [44](#)
 Object Store Remove (SREM)
 description of [48](#)
 Object Store Verify Fail (SVRF)
 description of [49](#)
 Object Store Verify Unknown (SVRU)
 description of [50](#)
 objects
 determining the security partition for [54](#)

P

protocol audit messages, HTTP
 listing of [23](#)

R

result
 audit messages [11](#)
 retrieval transactions
 list of audit messages generated by [17](#)
 RSLT [11](#)

S

S3 Delete (SDEL)
 description of [45](#)
 S3 Get (SGET)

- description of [45](#)
- S3 Head (SHEA)
 - description of [46](#)
- S3 Metadata Updated (SUPD)
 - description of [48](#)
- S3 PUT (SPUT)
 - description of [47](#)
- Security Audit Disable (SADD)
 - description of [44](#)
- Security Audit Enable (SADE)
 - description of [44](#)
- Security Authentication Failed (ETAF)
 - description of [33](#)
- security partitions
 - determining for objects [54](#)
- suggestions
 - how to send feedback about documentation [64](#)
- Swift DELETE (WDEL)
 - description of [51](#)
- Swift GET (WGET)
 - description of [52](#)
- Swift HEAD (WHEA)
 - description of [53](#)
- Swift PUT (WPUT)
 - description of [53](#)
- system audit messages

- table of [20](#), [24](#)
- System Detected Lost Object (OLST)
 - description of [42](#)

T

- TCP/IP Connection Close (ETCC)
 - description of [34](#)
- TCP/IP Connection Establish (ETCA)
 - description of [34](#)
- TCP/IP Connection Fail (ETCF)
 - description of [35](#)
- TCP/IP Connection Refused (ETCR)
 - description of [36](#)
- timestamp
 - audit messages [11](#)
- Twitter
 - how to receive automatic notification of documentation changes [64](#)

U

- User Initiated Volume Lost (VLST)
 - description of [51](#)