



StorageGRID® Webscale 10.3

Cloud Data Management Interface Implementation Guide

September 2016 | 215-10813_A0
doccomments@netapp.com

 **NetApp**®

Contents

Introduction to CDMI implementation	5
Who should read this guide	5
Revision history	5
Supported versions of CDMI and HTTP	5
How the StorageGRID Webscale system implements CDMI	6
CDMI specification sections supported by the StorageGRID Webscale system	6
How the StorageGRID Webscale system resolves conflicts	9
How the system's ILM rules and metadata manage data objects	9
How the StorageGRID Webscale system implements immediate redundancy	9
CDMI namespace permissions you can specify	11
Client application access permissions and HTTP methods	11
What the Read access type is	11
Use of GET method to read data objects and data object metadata	11
Use of GET method to retrieve data and container object metadata	11
Modify or Write access	12
How applications use the HTTP PUT and POST methods to create and store objects	12
How applications use HTTP PUT to update metadata	12
Delete access	12
Last access time metadata	12
Enabling or disabling last access time	13
Connecting client applications using CDMI	14
Configuring the StorageGRID Webscale system to accept client connections	14
Associating client application IP addresses with link costs	14
Creating HTTP profiles to set namespace permissions	15
Associating HTTP profiles with client applications	16
How client application authentication works	17
Identifying IP addresses for API Gateway Nodes and Storage Nodes	17
Finding the port number for API Gateway Nodes and Storage Nodes for CDMI	17
How the StorageGRID Webscale system implements security in CDMI	18
How client applications use certificates for security in CDMI	19
Supported hashing and encryption algorithms for TLS libraries	19
Choosing hash algorithms for data objects	19
How security partitions are used	20
Testing client application connections to the system	25
CDMI root URI	25
Testing HTTP connections by using Telnet	25
Testing HTTP connections by using OpenSSL	26
Retrieving CDMI capabilities with curl	26
Testing nameless data object storage and retrieval by using curl	28
Managing CDMI clients with the StorageGRID Webscale system	30

Managing HTTP connections	30
Viewing HTTP transactions for CDMI objects	30
Accessing and reviewing audit logs	31
Viewing information about data objects	32
How to retrieve objects through UUID	32
How the StorageGRID Webscale system uses UUIDs	32
How the StorageGRID Webscale system uses UUIDs and CDMI object IDs	33
Benefits of active, idle, and concurrent HTTP connections	34
Benefits of different types of HTTP connections	34
Benefits of keeping idle HTTP connections open	34
Benefits of active HTTP connections	34
Benefits of concurrent HTTP connections	35
Copyright information	36
Trademark information	37
How to send comments about documentation and receive update notifications	38
Index	39

Introduction to CDMI implementation

The StorageGRID Webscale system supports the storage and retrieval of objects by applications that interface to the StorageGRID Webscale system through the Cloud Data Management Interface (CDMI).

For basic information about the StorageGRID Webscale system, see the [StorageGRID Webscale 10.3 Grid Primer](#) and the [StorageGRID Webscale 10.3 Administrator Guide](#).

Who should read this guide

You should use the *Cloud Data Management Interface Implementation Guide* if you are creating applications that interface with the StorageGRID Webscale system through CDMI. You can also use the guide to gain a basic understanding of how the StorageGRID Webscale system supports CDMI.

Revision history

The revision history lists changes to the StorageGRID Webscale system's support for CDMI along with the accompanying changes to the documentation. You might find it helpful to identify those changes in preparing for CDMI implementation.

Date	Release	Comments
June 2016	10.3	Removed support for named objects.
December 2015	10.2	The use of named data objects is deprecated and no longer supported.
April 2015	10.1	No change.
September 2014	10.0	Added support for named objects and multipart MIME.

Supported versions of CDMI and HTTP

The StorageGRID Webscale system supports specific versions of CDMI and HTTP. When you are developing the interface with your client application, it is helpful to know the supported versions.

The following versions are supported:

Item	Version
CDMI specification	1.0.2 Published by the Storage Networking Industry Association (SNIA)
HTTP	1.1 For more information about HTTP, see HTTP/1.1 (RFC 2616).

Related information

[IETF RFC 2616: Hypertext Transfer Protocol \(HTTP/1.1\)](#)

[SNIA: SNIA Cloud Data Management Interface \(CDMI\) Version 1.0.2](#)

How the StorageGRID Webscale system implements CDMI

A client application can connect to the CLB service or LDR service, create object containers, and store and retrieve data objects. The StorageGRID Webscale system uses information lifecycle management (ILM) rules to manage these data objects ingested into the StorageGRID Webscale system by a client application that interfaces to the system through CDMI. When you are designing interfaces to the system, it might be helpful to know how the system processes your data objects.

For more information about ILM rules, see the *Administrator Guide*.

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)

CDMI specification sections supported by the StorageGRID Webscale system

The StorageGRID Webscale system supports the *Cloud Data Management Interface (CDMI)* specification published by the Storage Networking Industry Association (SNIA). You might find it helpful to understand how the system implements various sections of the CDMI specification.

Data object resource operations

The StorageGRID Webscale system supports a single CDMI domain.

The following table lists the supported “Data Resource Operations” sections of the *Cloud Data Management Interface (CDMI)* specification.

Section number	Section description
8.4	Read a data object (CDMI Content Type)
8.5	Read a data object (Non-CDMI Content Type)
8.6	Update a data object (CDMI Content Type)
8.7	Update a data object (Non-CDMI Content Type)
8.8	Delete a data object (CDMI Content Type)
8.9	Delete a data object (Non-CDMI Content Type)

Byte range read operations

The StorageGRID Webscale system supports byte range read operations using both CDMI and non-CDMI content types. For byte range read using non-CDMI content type, the following byte range is returned by the StorageGRID Webscale system:

- If a single contiguous byte range is requested, the system returns the byte range.
- If multiple byte ranges are requested that can be coalesced without holes, the system returns a single coalesced range.
- If multiple byte ranges are requested that cannot be coalesced without holes, the system returns the entire data object bytes.

The length of time the system takes to return requested portions of a data object is impacted by the following:

- If you enable compression (by using the **Configuration > Grid Options** option, selecting **Configuration**, and enabling stored object compression) or if the data object is retrieved from tape, the StorageGRID Webscale system locates and returns the requested portion of the data object by reading the data object, starting at the beginning of the segment containing the first byte of the requested range.
- If you disable compression and the data object is retrieved from disk, the StorageGRID Webscale system can begin reading the segment from the start of the requested byte range and not the beginning of the segment.

Thus, if compression is enabled or the data object is retrieved from tape, it takes the system longer to return the requested portion of a data object.

Domain object resource operations

The StorageGRID Webscale system supports a single CDMI domain. The creation of additional domains is not supported.

Capability object resource operations

The following table lists the supported “Capability Object Resource Operations” sections of the *Cloud Data Management Interface (CDMI)* specification:

Section	Section description	Notes
12.2	Read a capabilities object (CDMI Content Type)	<p>The StorageGRID Webscale system supports the following:</p> <ul style="list-style-type: none"> • System-wide capabilities • Storage system metadata capabilities • Data system metadata capabilities • Data object capabilities

Metadata

The following table lists the supported “Metadata” sections of the *Cloud Data Management Interface (CDMI)* specification:

Section	Section description	Notes
16.3	Storage system metadata	<p>The StorageGRID Webscale system supports the following:</p> <ul style="list-style-type: none"> • cdm_i_size • cdm_i_ctime • cdm_i_atime • cdm_i_hash

Section	Section description	Notes
16.4	Data system metadata	The StorageGRID Webscale system supports the following: <ul style="list-style-type: none"> cdmi_data_redundancy cdmi_immediate_redundancy cdmi_value_hash
16.5	Provided data system metadata	The StorageGRID Webscale system supports the following: <ul style="list-style-type: none"> cdmi_value_hash_provided

The StorageGRID Webscale system converts the following StorageGRID Webscale metadata to populate the values of some CDMI storage system metadata. The following table identifies which StorageGRID Webscale metadata is used to populate CDMI system metadata.

CDMI storage system metadata	StorageGRID Webscale metadata
cdmi_size	CSIZ
cdmi_ctime	CTME
cdmi_atime	LATM The StorageGRID Webscale system uses the value from cdmi_ctime when a data object lacks LATM (last access time) metadata.
cdmi_hash	The StorageGRID Webscale system returns the hash for the data object.
cdmi_value_hash_provided	The StorageGRID Webscale system returns the name of the hash algorithm selected in the StorageGRID Webscale system when the system stored the data object.

Extensions

The StorageGRID Webscale system supports the CDMI Multi-part MIME Extension 1.0g for the creation and retrieval of both named and nameless data objects, otherwise subject to the same limitations for data object operations within the StorageGRID Webscale system's standard CDMI implementation. This extension allows clients to read and write the CDMI value as binary in a multipart body part, which avoids incurring encoding (for example, base64) overhead on both the client and the server.

Related concepts

[How to retrieve objects through UUID](#) on page 32

Related tasks

[Retrieving CDMI capabilities with curl](#) on page 26

Related information

[SNIA: SNIA Cloud Data Management Interface \(CDMI\) Version 1.0.2](#)
[StorageGRID Webscale 10.3 Administrator Guide](#)

How the StorageGRID Webscale system resolves conflicts

If the StorageGRID Webscale system detects that two or more objects with same name exist in the same container, the system allows these objects to remain, keeping the objects unique through the use of each object's UUID.

However, conflicts might arise when a client application attempts to access an object by name and the action is not performed on the expected object. The StorageGRID Webscale system resolves this conflict by always performing actions on the most recently created object.

How the system's ILM rules and metadata manage data objects

The StorageGRID Webscale system's information lifecycle management (ILM) rules enable you to use metadata in rules to manage data objects automatically.

You can use the following metadata:

- You can use *object size*, *last access time*, and *CDMI user-defined* metadata in ILM rules for data objects stored in the StorageGRID Webscale system by client applications interfacing to the system through CDMI.
- You can use *object size* to ensure that small objects are stored to disk and not tape, which avoids the poor retrieval performance of tape.
- You can use *last access time* metadata to identify content that has not been retrieved for a set amount of time and have this content moved to a cheaper grade of storage.
- You can set the filter criteria to evaluate data objects against *CDMI user defined* metadata. The last access time is StorageGRID Webscale system metadata.

For more information about ILM rules, see the *StorageGRID Webscale Administration Guide*.

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)

How the StorageGRID Webscale system implements immediate redundancy

The StorageGRID Webscale system supports the CDMI Data System Metadata Capabilities `cdmi_data_redundancy` and `cdmi_immediate_redundancy`, which are used to enable Dual Commit.

When enabled, at ingest and before an object is evaluated against the active ILM policy, Dual Commit synchronously creates two copies of object data and distributes these copies to two Storage Nodes. To enable Dual Commit, a client application includes data system metadata `cdmi_data_redundancy` and `cdmi_immediate_redundancy` in a POST request. For more information about Dual Commit, see the *Administration Guide*.

If this data system metadata is not included in a CDMI content type request, the following defaults are assumed:

- `"cdmi_data_redundancy": "2"`

- "cdmi_immediate_redundancy": true

These defaults also apply to non-CDMI content type `POST` data object create requests. For CDMI content type requests, you can display redundancy by setting `cdmi_data_redundancy` to `false`.

The following table summarizes the responses for successful redundancy requests:

Client application request		Does the system use dual commit?	CDMI response for success	
cdmi_data_redundancy	cdmi_immediate_redundancy		cdmi_data_redundancy_provided	cdmi_immediate_redundancy_provided
not present	not present	Yes	2	true
0 or 1	true	No	1	true
2 or greater	true	Yes	2	true
not present	false	No	1	true because one data object was stored

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)

CDMI namespace permissions you can specify

You can specify permissions (such as read, modify, and delete permissions) for client applications in the CDMI namespace in the StorageGRID Webscale system. For each permission, it is helpful to know its implementation and which HTTP methods you should use.

Client application access permissions and HTTP methods

In the StorageGRID Webscale system, you can specify whether a client application has permission to read, write, modify, or delete data objects in the CDMI namespace. You can also specify whether to enable last access time metadata. For each permission, it is helpful to know which HTTP methods you should use.

You can grant or deny the following access to client applications through the StorageGRID Webscale system. The table also indicates the HTTP method used for each permission:

Permission type	HTTP method
Read	GET
Modify/Write	PUT for modify access POST for write access
Delete	DELETE
Last Access Time	GET You must enable both Last Access Time and Read . When a CDMI client application uses GET to retrieve a data object, the StorageGRID Webscale system stores the time that the CDMI application client retrieved the data object in internal object metadata called <i>last access time</i> metadata.

What the Read access type is

The Read access type determines whether a client application has permission to read and retrieve data objects and data object metadata from the CDMI namespace.

Use of GET method to read data objects and data object metadata

Client applications use the HTTP GET method and an object ID or name to read data objects and data object metadata in the CDMI namespace.

Use of GET method to retrieve data and container object metadata

Client applications use the HTTP GET method and an object ID or name to retrieve data object and container object metadata from the CDMI namespace.

For GET, when the requested field `childrenrange` is specified and the requested field `children` is not included, the complete range of objects within the container is returned. For example, GET / CDMI/foo/?childrenrange returns "children range" : " 0-8989899".

If no requested fields are specified or the requested field `children` is specified without a range or a range that includes more than 10,000 objects, GET is limited by a maximum range of 10,000 objects

(0 through 9999). If a container includes more than 10,000 objects, multiple GET operations must be run.

Modify or Write access

The Modify/Write access type determines whether a client application has permission to store data objects and update data object metadata in the namespace.

How applications use the HTTP PUT and POST methods to create and store objects

Client applications can create container objects and store (create) unnamed data objects in the StorageGRID Webscale system. Use the HTTP PUT and POST methods to create and store objects.

Client applications use the following methods in the CDMI namespace:

- HTTP POST method to store unnamed data objects
- HTTP PUT method to create container objects

By default, the StorageGRID Webscale system enables immediate redundancy for all POST requests for the CDMI content type and the non-CDMI content type. You can disable immediate redundancy for the CDMI content type, but not the non-CDMI content type. However, you should include `cdmi_immediate_redundancy` metadata set to true in all PUT and POST requests to enable immediate redundancy and protect against data loss.

The StorageGRID Webscale system supports user metadata for CDMI data objects; however, the system does not support user metadata for CDMI container objects.

Related concepts

[How the StorageGRID Webscale system implements immediate redundancy](#) on page 9

How applications use HTTP PUT to update metadata

Client applications use the HTTP PUT method to update data object user metadata in the CDMI namespace.

The following considerations apply with the PUT method:

- The StorageGRID Webscale system supports adding or updating only all user metadata for a data object. It does not support adding or updating an individual user metadata item (using the URI syntax `?metadata:<metadataname>`).
- Updating other fields (for example, `value` or `mimetype`) using the PUT method is not supported.

Delete access

The Delete access type determines whether a client application has permission to delete data objects and container objects from the namespace.

Last access time metadata

The last access time permission determines whether the StorageGRID Webscale system updates last access time metadata for a data object when a client application retrieves the object. You can create

information lifecycle management (ILM) rules to take action on data objects based on the last time that a client application retrieved the object.

When a client application that is assigned a CDMI profile with last access time enabled uses `GET` to retrieve a data object, the StorageGRID Webscale system saves the retrieval time in internal object metadata called last access time metadata.

Only the StorageGRID Webscale system can use internal metadata. For example, ILM policies can use last access time metadata to identify when an object was last retrieved.

Note: Because last access time metadata updates each time that a client application retrieves a data object, it can affect system performance. It is recommended that you disable **Last Access Time** in the StorageGRID Webscale system when no ILM policies use last access time metadata.

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)


Enabling or disabling last access time

You can enable and disable last access time by creating a CDMI profile. With last access time enabled, you can create an ILM rule that uses Last Access Time as the Reference Time.

Before you begin

- You must have signed in to the Grid Management Interface using a supported browser.
- To perform this task, you need specific access permissions.

Steps

1. Select **Configuration > CDMI**.
2. From the CDMI menu, select **Permissions**.
3. Under HTTP/ CDMI and UUID Namespaces, click **Edit**  next to the CDMI profile that you want to modify.
4. Perform one of the following actions:
 - To enable last access time, select the **Read** and the **Last Access Time** check boxes for each CDMI client.
 - To disable last access time, clear the **Last Access Time** check box for each CDMI client.
5. Click **Apply Changes**.

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)

Connecting client applications using CDMI

You must configure the StorageGRID Webscale system to accept HTTP connections from client applications. Client applications use HTTP connections to access and communicate with the StorageGRID Webscale system.

Note: IPv6 is only supported for client application connections through the API Gateway Node.

For more information about support for IPv6, see the *Administrator Guide*.

Connecting client applications to the StorageGRID Webscale system involves the following tasks:

- Configuring client connections to accept HTTP connections
- Identifying IP address for API Gateway Nodes and Storage Nodes
- Identifying port number for API Gateway Nodes and Storage Nodes
- Copying the system's certificate authority (CA) certificate for client applications that require server validation

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)

Configuring the StorageGRID Webscale system to accept client connections

Configuring the StorageGRID Webscale system to accept HTTP connections from client applications requires you to complete several steps.

Steps

1. [Associating client application IP addresses with link costs](#) on page 14
You can associate a link cost with the IP addresses that client applications use to connect with the StorageGRID Webscale system. A link cost is the relative cost of communicating between data center sites and is used by the StorageGRID Webscale system to determine which grid nodes can best provide a requested operation: for example, an object retrieval.
2. [Creating HTTP profiles to set namespace permissions](#) on page 15
You can create HTTP profiles that identify whether read, write, modify, or delete permissions are enabled or disabled in a namespace. You can create multiple HTTP profiles.
3. [Associating HTTP profiles with client applications](#) on page 16
HTTP profiles identify whether read, write, modify, or delete permissions are enabled in a namespace. You can associate HTTP profiles with individual client applications or with groups of client applications, based on IP addresses. The association gives client applications access to the StorageGRID Webscale namespace and identifies the HTTP permissions for the client application in the namespace.

Associating client application IP addresses with link costs

You can associate a link cost with the IP addresses that client applications use to connect with the StorageGRID Webscale system. A link cost is the relative cost of communicating between data center

sites and is used by the StorageGRID Webscale system to determine which grid nodes can best provide a requested operation: for example, an object retrieval.

Before you begin

- You must have 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 > Link Cost**.
2. In the **Client Site IP Ranges** table, perform one of the following actions:

When...	Then...
No entries exist	Click Edit .
One or more entries exist	Click Insert .

3. In the **IP Range Name** box, type a name for the IP address or the range of IP addresses.
You can use any name. The configuration does not reference the name elsewhere.
4. In the **IP Range** box, type the IP address or the range of IP addresses that the client uses to contact the StorageGRID Webscale system.
Use a hyphen or slash to indicate an inclusive range of IP addresses, as shown in the following examples:
 - 192.168.120.0/24 (CIDR format)
 - 192.168.142.20-192.168.142.28 (dotted decimal format)
 You can use an abbreviated format for masks in eight-bit steps. For example, 192.168.142.0 is equivalent to the CIDR notation 192.168.142.0/24, and you can extend it as follows: n.n.0.0 is equivalent to n.n.0.0/16.
5. In the **Site ID** list, select a site ID.
Site ID is the unique identification number assigned to a data center site.
6. Click **Apply Changes**.
7. Repeat this procedure for each range of IP addresses that client applications use to access the StorageGRID Webscale system.

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)

Creating HTTP profiles to set namespace permissions

You can create HTTP profiles that identify whether read, write, modify, or delete permissions are enabled or disabled in a namespace. You can create multiple HTTP profiles.

Before you begin

- You must have 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 > CDMI**.
2. In the **HTTP /CDMI and /UUID Namespace** table, perform one of the following actions:

When...	Then...
No entries exist	Click Edit .
One or more entries exist	Click Insert .

3. Check the boxes for the HTTP operations that you want to enable in the profile.
4. Click **Apply Changes**.
5. Create additional profiles as needed.

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)

Associating HTTP profiles with client applications

HTTP profiles identify whether read, write, modify, or delete permissions are enabled in a namespace. You can associate HTTP profiles with individual client applications or with groups of client applications, based on IP addresses. The association gives client applications access to the StorageGRID Webscale namespace and identifies the HTTP permissions for the client application in the namespace.

Before you begin

- You must have 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 > CDMI**.
2. From the CDMI menu, select **Clients**.
3. In the **HTTP Entities** table, perform one of the following actions:

When...	Then...
No HTTP entities exist	Click Edit .
One or more HTTP entities exist	Click Insert .

4. In the Description box, enter a description of the client.
5. In the IP Range box, enter the range of IP addresses that the client application can use to connect to the LDR service or the CLB service.
6. In the Profile Name list, select the name of the HTTP profile that you created.
7. Click **Apply Changes**.

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)

How client application authentication works

The StorageGRID Webscale system uses its HTTP management settings to authenticate client application requests for access to the StorageGRID Webscale system. Understanding client application authentication helps in establishing successful connections.

When a client application requests access to the StorageGRID Webscale system, the system authenticates the request against the HTTP management settings that you created for the client application and completes the following steps:

1. The StorageGRID Webscale system checks that the client application is using the same IP address or range of IP addresses that are defined in the HTTP management settings.
2. When the client application passes the authentication process, the StorageGRID Webscale system opens a TCP/IP connection.

Identifying IP addresses for API Gateway Nodes and Storage Nodes

You need the grid node's IP address to connect API client applications to StorageGRID Webscale.

Steps

1. Sign in to the Grid Management Interface using a supported browser.
2. Select **Grid**.
3. In the **Grid Topology** tree, locate and expand the Storage Node or API Gateway Node to which you want to connect.

The services for the selected grid node appear.

4. In the Storage Node or API Gateway Node, select **SSM > Resources**, and then scroll to the **Network Addresses** table.

You can establish HTTPS connections from API client applications to any of the listed IP addresses.

Related tasks

[Finding the port number for API Gateway Nodes and Storage Nodes for CDMI](#) on page 17

Finding the port number for API Gateway Nodes and Storage Nodes for CDMI

You can find the port number for API Gateway Nodes or Storage Nodes by using the Network Management Service (NMS) Management Interface (MI). You need this port number to create an HTTP connection from client applications to the StorageGRID Webscale system.

About this task

The following ports are used for client applications that interface to the StorageGRID Webscale through CDMI:

Grid node	Port number
API Gateway Node (CLB HTTP Port)	8080
Storage Node (LDR HTTP Port)	18080

Steps

1. Sign in to the Grid Management Interface using a supported browser.
2. Select **Configuration > Storage Options**.
3. In the **Ports** table, locate the port number for the API Gateway Node (CLB HTTP Port) or Storage Node (LDR HTTP Port).

How the StorageGRID Webscale system implements security in CDMI

The StorageGRID Webscale system employs the use of Transport Layer Security (TLS) connection security, server authentication, client authentication, client authorization, and security partitions. When considering security issues, you might find it helpful to understand how the StorageGRID Webscale system implements security, authentication, and authorization.

The StorageGRID Webscale system accepts only HTTP commands submitted over a network connection that uses TLS to provide connection security, application authentication and, optionally, transport encryption.

TLS enables the exchange of certificates as entity credentials and allows a negotiation that can use hashing and encryption algorithms.

Server authentication uses server certificates signed by the StorageGRID Webscale system's certificate authority (CA) certificate. The administrator might replace the system's CA certificate with a single, common server certificate applicable to all the API ports within the system.

The following table lists security issues when using CDMI:

Security issue	CDMI
Connection security	TLS
Server authentication	X.509 server certificate signed by system CA or server certificate supplied by administrator
Client authentication	Anonymous or client certificate (security partition) For CDMI, a client certificate (needed for security partitions) is supported only on nameless objects.
Client authorization	Client profile permissions and object ownership A system-wide command to disable the client's ability to delete can override client authorization.
Client origin permissions	IP range

Related concepts

[How security partitions are used](#) on page 20

How client applications use certificates for security in CDMI

When a client application establishes a TLS session with the StorageGRID Webscale system, the system sends a server certificate to the client application for verification, to ensure that the HTTP connection is secure. Understanding certificate use is important for system security.

You can verify the server certificate by using the StorageGRID Webscale system's CA certificate. The client application should load the system CA certificate and use it to verify that the client application is communicating with the expected StorageGRID Webscale system. This process protects against man-in-the-middle and impersonation attacks.

Client applications can send client certificates to the StorageGRID Webscale system as part of session establishment.

For information about copying the CA certificate, refer to the *StorageGRID Webscale Administration Guide*.

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)

Supported hashing and encryption algorithms for TLS libraries

Client applications use the HTTPS protocol to communicate with the StorageGRID Webscale system over a network connection that uses Transport Layer Security (TLS). The StorageGRID Webscale system supports a limited set of hashing and encryption algorithms from the TLS libraries that client applications can use when establishing a TLS session. When you are setting up the communication processes, it is important for you to know which security algorithms the system uses.

The StorageGRID Webscale system supports the following cipher suite security algorithms:

- AES128-SHA
- AES256-SHA
- AES128-GCM
- AES256-GCM

AES128-SHA and AES256-SHA provide secure encryption and efficient processing of objects. AES128-GCM and AES256-GCM provide secure encryption and more efficient processing of large objects. The TLS session negotiates the connection, using either AES128 or AES256 based on the client application requirements, and the need to balance performance with encryption security.

Choosing hash algorithms for data objects

The StorageGRID Webscale system can use either SHA-1 or SHA-2 256 secure hash algorithms to generate a hash for each data object stored in the StorageGRID Webscale system. You can choose the algorithm that best meets your security needs.

Before you begin

- You must have 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.

About this task

The following table maps the choices in the StorageGRID Webscale system to the names of the hash algorithms:

Algorithm choice	Name of algorithm
SHA-256	SHA-2 256 bits
SHA-1	SHA-1

Because you can change the algorithm, you might have data objects in the system with hashes generated by different algorithms. As a result, metadata for different data objects might include different algorithm names. The algorithm name associated with the data object depends on which algorithm was selected in the StorageGRID Webscale system when the data object was stored in the system.

Steps

1. Select **Configuration > Grid Options**.
2. From the Grid Options menu, select **Configuration**.
3. In the Stored Object Hashing option, select a hash algorithm.
4. Click **Apply Changes**.

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)

How security partitions are used

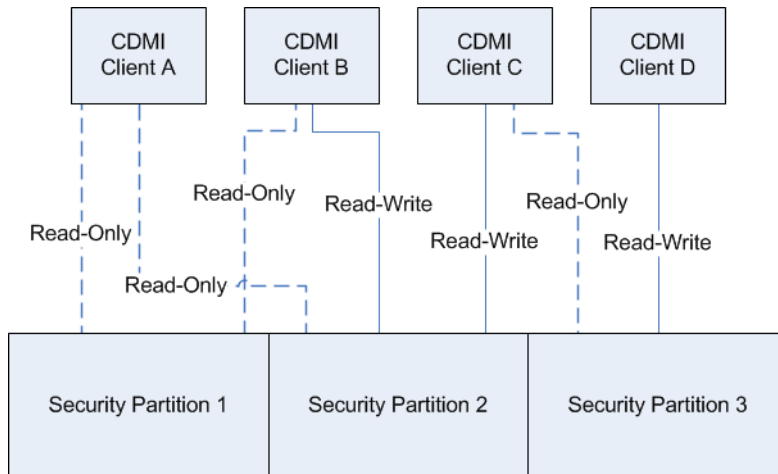
Security partitions are system-wide settings that provide you with a means to restrict access to objects for clients that connect to the StorageGRID Webscale system through CDMI. For instance, two client applications ingesting objects to the same system can be denied access to each other's objects.

Security partitions are supported only for nameless data objects, not for named data objects.

If the client application is assigned to a security partition or if the client application's assigned HTTP profile requires certificate authentication, a certificate is required. This certificate must be loaded to the StorageGRID Webscale system as part of the configuration process. If you use security partitions, consider how they are supported and how the StorageGRID Webscale system uses a certificate.

Security partition permissions

A client application can be configured with read-write permission to one security partition only. Multiple client applications can be configured with read-write permission to the same security partition. A client application can be configured with read-only permission to multiple security partitions.



Security partitions and CDMI transactions

When a client application assigned to a read-write security partition ingests an object into the StorageGRID Webscale system, the object is tagged with a metadata field that identifies the security partition assigned to that object.

CDMI POST/PUT requests are denied unless the security partition of the target object matches the read-write security partition associated with the client application.

CDMI GET requests are denied unless the security partition of the target object matches the read-write or read-only security partitions associated with the client application.

CDMI DELETE requests are denied unless the security partition of the target object matches the read-write security partition associated with the client application.

Objects stored by a client application that is not associated with a read-write security partition are not assigned to any security partition. Similarly, objects ingested into the StorageGRID Webscale system before security partitioning is enabled are not assigned to a partition when a security partition is enabled. Objects that are not assigned to a security partition can be retrieved, queried, and deleted by any client application.

Security partitions and certificates

Configuring a client application for security partitioning involves mapping a client certificate to a security partition.

When a client application establishes a session to the StorageGRID Webscale system, the client application presents a client certificate during TLS authentication. This certificate is verified against the certificates loaded into the StorageGRID Webscale system. Furthermore, if the client application presents a certificate signed by a certificate authority (CA), the StorageGRID Webscale system performs another layer of validation by not only validating the certificate against its list of permitted client certificates, but also against its list of CA certificates. After the certificate has been accepted, the StorageGRID Webscale system looks up the client application name associated with the certificate. If a matching client application name is found, the StorageGRID Webscale system looks up the security partitions associated with the client application name.

If the client application does not present a certificate or the certificate is not associated with a security partition, the client application can only access objects that are not assigned to a security partition.

Acquiring client certificates

The client application must be issued a certificate that uniquely identifies the client application to the StorageGRID Webscale system. There are two ways to acquire a client certificate: from a CA or by creating a self-signed certificate.

CA-issued certificates are signed certificates available from commercially available CAs. You can also host your own CA and issue your own certificates. Details on managing and generating certificates are beyond the scope of this guide. For more information, see third-party products.

Enable security partitions

To restrict client access to objects ingested through CDMI, enable security partitions.

Before you begin

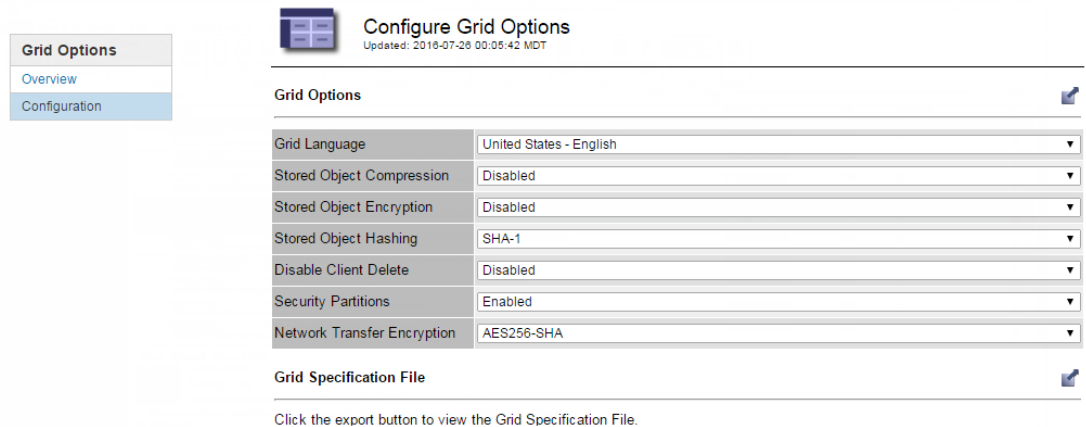
- You must have signed in to the Grid Management Interface using a supported browser.
- To perform this task, you need specific access permissions. For details, see information about controlling system access with administration user accounts and groups.

About this task

Objects ingested into the StorageGRID Webscale system before security partitioning is enabled is accessible to all applications. It is recommended that you configure client application access before enabling security partitioning. It is also recommended that you enable security partitioning before a StorageGRID Webscale system becomes operational so that objects are secure.

Steps

1. Select **Configuration > Grid Options**.
2. From the Grid Options menu, select **Configuration**.
3. Select **Security Partitions > Enabled**.



The screenshot shows the 'Configure Grid Options' page in the StorageGRID Webscale interface. On the left, a sidebar menu has 'Grid Options' selected, with 'Overview' and 'Configuration' as sub-items. The main content area is titled 'Configure Grid Options' and includes a timestamp 'Updated: 2016-07-26 00:05:42 MDT'. Below the title, there is a 'Grid Options' section with a table of settings:

Grid Options	
Grid Language	United States - English
Stored Object Compression	Disabled
Stored Object Encryption	Disabled
Stored Object Hashing	SHA-1
Disable Client Delete	Disabled
Security Partitions	Enabled
Network Transfer Encryption	AES256-SHA

Below the table is a 'Grid Specification File' section with a note: 'Click the export button to view the Grid Specification File.'

4. Click **OK**.

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)

Configure security partitions





You must manually associate client applications with a security partition and set permissions.


Before you begin


- You must have 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.
- You must have enabled security partitions.

Steps


1. Add the client certificate to the StorageGRID Webscale system:
 - a. Select **Configuration > CDMI**.
 - b. From the CDMI menu, select **Certificates**.
 - c. In the Certificate Authorities table, click **Edit**  (or **Insert**  if this is not the first certificate).
 - d. Enter a new value for **CA Name** and paste the CA certificate or, if the client certificate is self-signed, paste the client certificate into **CA Certificate**.
 - e. In the Clients table, click **Edit**  (or **Insert**  if this is not the first certificate).
 - f. Enter a new value for **Client Name** and paste the client certificate into **Client Certificate**.





 **Certificates**
Updated: 2016-07-26 00:18:54 MDT

Certificate Authorities (0 - 0 of 0) 


CA Name	CA Certificate	Actions
Client Cert 1	J6NZjSSwtCPkU1H505rx2xrRi9sh5FBvJDFMCwm8hfXPeo/wuA 9Gk4G7m6ij4ce YzMovOGbqdtKBjJIT7P2EhK64nkHQYp6NaFizLk9Ruiy/e8XJJtIM v4= -----END CERTIFICATE-----	   

Show Records Per Page Previous « » Next


Clients (0 - 0 of 0) 


Client Name	Client Certificate	Actions
Client Cert 1	J6NZjSSwtCPkU1H505rx2xrRi9sh5FBvJDFMCwm8hfXPeo/w uA9Gk4G7m6ij4ce YzMovOGbqdtKBjJIT7P2EhK64nkHQYp6NaFizLk9Ruiy/e8XJJtI Mv4= -----END CERTIFICATE-----	   







Show Records Per Page Previous « » Next




- g. Click **Apply Changes**.
2. Configure the security partition:
 - a. Select **Configuration > Grid Options**.
 - b. From the CDMI menu, select **Security Partitions**.





 **Security Partitions**
Updated: 2018-07-26 01:13:34 MDT

Partitions (1 - 2 of 2) 


Partition Name	Partition Identifier	Actions
Partition 1	1	  
Partition 2	2	  





Show Records Per Page Previous « 1 » Next

Client Partitions (1 - 1 of 1) 


Client Name	Partition Name	Actions
HTTP Client 1	Partition 1	   

Show Records Per Page Previous « 1 » Next

Read-Only Client Partitions (1 - 1 of 1) 

Client Name	Partition Name	Actions
HTTP Client Cert 1	Partition 2	   





Show Records Per Page Previous « 1 » Next



- c. In the Partitions table, click **Edit**  (or **Insert**  if this is not the first partition) and enter a value for **Partition Name**.
- d. Click **Apply Changes**.

A new security partition is created.

Note: Once created a security partition cannot be deleted from the StorageGRID Webscale system.

- e. Optionally, in the Client Partitions table, click **Edit**  (or **Insert**  if this is not the first mapping), select a **Client Name** and then a **Partition Name**.
This action assigns the client application to a read-write security partition.
- f. Optionally, in the Read-Only Client Partitions table, click **Edit**  (or **Insert**  if this is not the first mapping), select a **Client Name** and then a **Partition Name**.
This action grants the client application read-only access to the selected security partition.
- g. Click **Apply Changes**.

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)

Testing client application connections to the system

Using either Telnet or OpenSSL, you can test the HTTP connection between the client application and the StorageGRID Webscale system to ensure that the connection works. You can also test that the client application can store objects to and retrieve objects from the StorageGRID Webscale system.

If you copy a command from this section and paste the command into another application, the copy-and-paste process might remove dashes that appear between words near a line break. Ensure that the pasted command includes all dashes before you run the command.

CDMI root URI

The root URI for CDMI access to the StorageGRID Webscale system is `http://IP_address:port/CDMI`. You need the root URI when testing HTTP connections.

For `IP_address` and `port`, you must use the IP address and port for an API Gateway Node or Storage Node.

Testing HTTP connections by using Telnet

You can use a utility such as Telnet to test the HTTP connection between client applications and the StorageGRID Webscale system to ensure that the HTTP connection is correctly configured.

Before you begin

- You must have configured an IP address for the client application in the StorageGRID Webscale system.
- You must know the IP address and port number for an API Gateway Node or Storage Node.

About this task

You can connect the client application to an API Gateway Node or Storage Node.

Step

1. Use Telnet from a client application to connect to an API Gateway Node or Storage Node:

```
telnet IP_address port
```

For `IP_address` and `port`, use the IP address and port for an API Gateway Node or Storage Node.

If you correctly configured the IP address for the client application in the StorageGRID Webscale system, a delay of several seconds occurs, and then the API Gateway Node or Storage Node drops the connection.

If you incorrectly configured the IP address for the client application in the StorageGRID Webscale system, the connection closes immediately. If the API Gateway Node's CLB service or Storage Node's LDR service is not running or if a network error occurs, Telnet is unable to connect to the API Gateway Node or Storage Node.

Testing HTTP connections by using OpenSSL

You can use the `openssl` command to test the HTTP connection between client applications and the StorageGRID Webscale system to ensure that the HTTP connection is correctly configured.

Before you begin

- You must have configured an IP address for the client application in the StorageGRID Webscale system.
- You must know the IP address and port number for an API Gateway Node or Storage Node.

About this task

You can connect the client application to an API Gateway Node or Storage Node.

Step

1. From a client application, establish an HTTP connection to an API Gateway Node or Storage Node:

```
openssl s_client -tls1 -connect IP_address:port
```

For *IP_address* and *port*, you must use the IP address and port for an API Gateway Node or Storage Node.

If you correctly configured the IP address for the client application in the StorageGRID Webscale system, a connected response appears.

If you incorrectly configured the IP address for the client application in the StorageGRID Webscale system, an error response appears.

Retrieving CDMI capabilities with curl

You can retrieve the CDMI capabilities of the StorageGRID Webscale system to determine which CDMI functions the StorageGRID Webscale system supports. Knowing the CDMI capabilities helps you understand the functions that client applications can perform with the StorageGRID Webscale system.

Before you begin

You must know the IP address and port number for an API Gateway Node or the Storage Node.

About this task

In the following task, *IP_address* and *port* are for an API Gateway Node or the Storage Node.

Note: If you copy a command from this section and paste the command into another application, the copy-and-paste process might remove dashes that appear between words near a line break. You must ensure that the pasted command includes all dashes before you run the command.

Steps

1. From a client application, use `curl` to retrieve CDMI capabilities from the StorageGRID Webscale system:

```
curl -X GET --header 'Host: IP_address:port' --header 'Content-Type: application/cdmf-capability' --header 'X-CDMI-Specification-Version: 1.0.2' -k https://IP_address:port/CDMI/cdmf_capabilities/
```

A response that looks similar to the following appears:

```
{ "objectType": "application/cdmi-capability",
  "objectID": "00006FFD0009B74801", "objectName": "cdmi_capabilities/",
  "parentURI": "/CDMI/",
  "parentID": "00006FFD0019124105000000000000000000000000000000000000",
  "capabilities": { "cdmi_domains": "true",
    "cdmi_metadata_maxitems": "32",
    "cdmi_metadata_maxsize": "4096",
    "cdmi_metadata_maxtotalsize": "32768",
    "cdmi_multipart_mime": "true",
    "cdmi_object_access_by_ID": "true",
    "cdmi_post_dataobject_by_ID": "true",
    "cdmi_security_data_integrity": "true" }, "childrenrange": "0-2",
  "children": [ "container/", "dataobject/", "domain/" ] }
```

The supported CDMI capabilities are displayed after "capabilities": for example, "cdmi_domains": "true".

IP_address and *port* are for an API Gateway Node or the Storage Node.

- Optional: Retrieve CDMI capabilities for other CDMI object types, such as domain and data objects, by appending the required capability name to the URL with a trailing forward slash.

For more information about the different objects, see the CDMI specification.

Example

The following command retrieves the data system metadata capabilities:

```
curl -X GET --header 'Host: IP_address:port' --header 'Content-Type: application/cdmi-capability' --header 'X-CDMI-Specification-Version: 1.0.2' -k https://IP_address:port/CDMI/cdmi_capabilities/dataobject/
```

IP_address and *port* are for an API Gateway Node or Storage Node.

A response similar to the following appears:

```
{ "objectType": "application/cdmi-capability",
  "objectID": "00006FFD0009B60802", "objectName": "dataobject/",
  "parentURI": "/CDMI/cdmi_capabilities/",
  "parentID": "00006FFD0009B74801",
  "capabilities": {
    "cdmi_read_value": "true", "cdmi_read_value_range": "true",
    "cdmi_read_metadata": "true", "cdmi_delete_dataobject": "true",
    "cdmi_size": "true", "cdmi_ctime": "true",
    "cdmi_atime": "true",
    "cdmi_data_redundancy": "2", "cdmi_immediate_redundancy": "true",
    "cdmi_hash": "true", "cdmi_value_hash": [ "SHA256" ] },
  "childrenrange": "", "children": [ ] }
```

This response includes the "cdmi_value_hash" capability, indicating that the SHA-256 hash algorithm is supported.

Related references

[CDMI specification sections supported by the StorageGRID Webscale system](#) on page 6

Testing nameless data object storage and retrieval by using curl

You can store and retrieve a test nameless data object to ensure that these functions work.

Before you begin

You must know the IP address and port number for an API Gateway Node or Storage Node.

About this task

You can connect client applications to an API Gateway Node or Storage Node. In the following steps, *IP_address* and *port* are for an API Gateway Node or Storage Node.

Note: If you copy a command from this section and paste the command into another application, the copy-and-paste process might remove dashes that appear between words near a line break. You must ensure that the pasted command includes all dashes before you run the command.

Steps

1. From a client application, use curl to store nameless data object in the StorageGRID Webscale system:

```
curl -X POST --header 'Host: IP_address:port' 'Accept: application/cdmi-object' --header 'Content-Type: application/cdmi-object' --header 'X-CDMI-Specification-Version: 1.0.2' -d '{"domainURI":"/cdmi_domains/", "value":"Hello Big World"}' -k https://IP_address:port/CDMI/cdmi_objectid/
```

A response similar to the following example appears:

```
{ "capabilitiesURI": "/CDMI/cdmi_capabilities/dataobject/",
  "completionStatus": "Complete", "domainURI": "/CDMI/cdmi_domains/",
  "mimetype": "text/plain; charset=utf-",
  "objectID": "00006FFD0019F9EF00177C0596AE654ACF8E01C46395A0CD45",
  "objectType": "application/cdmi-object",
  "metadata": { "a": "A", "ab": "AB", "b": "B",
  "cdmi_atime": "2013-06-14T23:44:47.203728Z",
  "cdmi_ctime": "2013-06-14T23:44:47.203728Z",
  "cdmi_data_redundancy_provided": "2",
  "cdmi_hash": "185F8DB32271FE25F561A6FC938B2E264306EC304EDA518007D1764826381969",
  "cdmi_immediate_redundancy_provided": "true", "cdmi_size": "5",
  "cdmi_value_hash_provided": "SHA256" }
```

The response includes a "completionStatus": "Complete" text that indicates that you successfully created the nameless data object. The response also includes the "objectID" number for the nameless data object. In this example, the "objectID" is 00006FFD0019F9EF00177C0596AE654ACF8E01C46395A0CD45.

2. Copy the "objectID" number from the response.
3. Use curl to retrieve the test nameless data object from the StorageGRID Webscale system and include the "objectID" number:

```
curl -X GET --header 'Host: IP_address:port' --header 'Accept: application/cdmi-object' --header 'X-CDMI-Specification-Version: 1.0.2' -k https://IP_address:port/CDMI/cdmi_objectid/00006FFD0019F9EF00177C0596AE654ACF8E01C46395A0CD45
```

A response similar to the following example appears:

```
{ "capabilitiesURI": "/CDMI/cdmi_capabilities/dataobject/",
  "completionStatus": "Complete", "domainURI": "/CDMI/cdmi_domains/",
  "mimetype": "text/plain; charset=utf-",
  "objectID": "00006FFD0019F9EF00177C0596AE654ACF8E01C46395A0CD45",
  "objectType": "application/cdmi-object",
  "valuetransferencoding": "utf-8",
  "metadata": { "a": "A", "ab": "AB", "b": "B",
  "cdmi_atime": "2013-06-14T23:44:47.203728Z",
  "cdmi_ctime": "2013-06-14T23:44:47.203728Z",
  "cdmi_hash": "185F8DB32271FE25F561A6FC938B2E264306EC304EDA518007D1764826381969",
  "cdmi_size": "5", "cdmi_value_hash_provided": "SHA256" },
  "valuerange": "0-4", "value": "Hello" }
```

4. Optional: Retrieve the value of a specific field by replacing `/CDMI/cdmi_objectid/00006FFD0019692A00FAE83433343A47939555F14AA4D2F115` with `/CDMI/cdmi_objectid/00006FFD0019692A00FAE83433343A47939555F14AA4D2F115?field_name`.

The response includes the value for the requested field for the nameless data object.

5. Optional: Retrieve all metadata that begins with the prefix `cdmi` by replacing `/CDMI/cdmi_objectid/00006FFD0019692A00FAE83433343A47939555F14AA4D2F115` with `/CDMI/cdmi_objectid/00006FFD0019692A00FAE83433343A47939555F14AA4D2F115?metadata:cdmi_`.

The response includes all metadata for the nameless data object that begins with the `cdmi` prefix.

Managing CDMI clients with the StorageGRID Webscale system

You can use the StorageGRID Webscale system to manage CDMI client access to the StorageGRID Webscale system by changing the state of HTTP connections to the system. You can also use the StorageGRID Webscale system to view operations for CDMI clients and to look up CDMI object IDs.

You cannot use a CDMI client application to retrieve objects ingested through an S3 client application. You cannot use an S3 client application to retrieve objects ingested through a CDMI client.

Managing HTTP connections

In the StorageGRID Webscale system, you can change the state of an HTTP connection to the StorageGRID Webscale system to online, online (read-only), redirect, or offline to manage client application access to the StorageGRID Webscale system.

Before you begin

- You must have 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.

About this task

The state of the HTTP connection affects client applications. For example, when you change **HTTP/CDMI State** to **Offline**, client applications cannot access the StorageGRID Webscale system because the HTTP connection is closed.

Steps

1. Select **Grid**.
2. In the Grid Topology tree, select **Storage Node > LDR > ConfigurationMain**.
3. In the **HTTP/CDMI State** list, select a state.
4. Click **Apply Changes**.

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)

Viewing HTTP transactions for CDMI objects

You can view the number of successful and failed attempts by client applications to read, write, and modify CDMI objects in the StorageGRID Webscale system. You can view a summary of all

transactions for all LDR services, or you can view the transactions for a specific LDR service. Viewing successful and failed transactions might help you in troubleshooting.

Before you begin

- You must have 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 **Grid**.
2. To view information about individual LDR services, select **Storage Node > LDR > CDMI > Overview > Main**.
3. To reset the counter for the LDR service to zero, on the **Configuration** page, select **Reset CDMI Counts** and click **Apply Changes**.

The numbers on the **Main** page reset to zero and start incrementing again.

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)

Accessing and reviewing audit logs

The StorageGRID Webscale system securely and reliably transports audit messages from each service within the StorageGRID Webscale system to one or more audit repositories. API-specific (S3, Swift, and CDMI) audit messages provide critical security, operations, and performance monitoring data that can help you evaluate the health of your system.

About this task

The StorageGRID Webscale system compresses audit logs after one day and renames them using the format `YYYY-MM-DD.txt.gz` (where the original date is preserved).

Steps

1. Log in to the server using the user name and password as recorded in the `Passwords.txt` file.
2. Access the audit log directory through a command line of the server that hosts the AMS service.
3. Go to the `/var/local/audit/export/` directory.
4. View the `audit.log` file.

Related information

[StorageGRID Webscale 10.3 Audit Message Reference](#)

Viewing information about data objects

You can use an object ID in the StorageGRID Webscale system to view information about the data object. You can check on the current location of the object and obtain any metadata associated with the object.

Before you begin

- You must have 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.
- You must have the object ID from the client application. Object ID can be one of:
 - CBID (content block identifier)
 - UUID (universally unique identifier)
 - Object ID
 - Container/Object_Key

Steps

1. Select **Grid**.
2. In the Grid Topology tree, select **primary Admin Node > CMN > Object Lookup > Configuration**.
3. In the **Object Identifier** box, enter an object ID, and click **Apply Changes**:
Note: If you enter an invalid object ID, an error message appears.
4. Click the **Overview** tab to review the results.

Related information

[StorageGRID Webscale 10.3 Administrator Guide](#)

How to retrieve objects through UUID

You can use CDMI client applications to retrieve a data object by deriving an object ID from the UUID for the stored data objects. When you implement the retrieval of data objects, you cannot use object containers. Understanding how the StorageGRID Webscale system uses UUIDs and how CDMI uses object IDs helps you derive an object ID from a UUID and retrieve data objects stored by client applications.

How the StorageGRID Webscale system uses UUIDs

The StorageGRID Webscale system uses the LDR service to assign a universally unique identifier (UUID) to each object in the system.

UUIDs are 128 bits with an internal binary structure and a string representation in the form of A-B-C-D-E:

- A is eight hexadecimal digits.

- B is four hexadecimal digits.
- C is four hexadecimal digits.
- D is four hexadecimal digits.
- E is 12 hexadecimal digits.

Each hexadecimal digit can take the values from 0 to 9, *A* through *F* (or lowercase *a* through lowercase *f*). The following is an example of a UUID: F81D4fAE-7DEC-11D0-A765-00A0C91E6BF6.

The StorageGRID Webscale system randomly generates the UUID for each object, as described in IETF RFC 4122.

Related information

<http://www.ietf.org>

How the StorageGRID Webscale system uses UUIDs and CDMI object IDs

The StorageGRID Webscale system uses universally unique identifiers (UUIDs) to track and manage content, but CDMI client applications use object IDs to track data objects and container objects. Understanding how the system uses UUIDs and object IDs helps in tracking objects.

When a CDMI client application submits content to the StorageGRID Webscale system, the system stores the content and treats the stored content as one or more data objects. However, the LDR service creates a UUID, embeds the UUID in an object ID, and assigns the object ID to each data object that CDMI client applications submit to the system. CDMI client applications use the OID to retrieve, update, and delete these data objects.

Benefits of active, idle, and concurrent HTTP connections

How you configure HTTP connections can impact the performance of the StorageGRID Webscale system. Configurations differ depending on whether the HTTP connection is active or idle or you have concurrent multiple connections.

Benefits of different types of HTTP connections

The type of HTTP connection duration can impact the performance of the StorageGRID Webscale system.

You can identify the performance benefits for the following types of HTTP connections:

- Idle HTTP connections
- Active HTTP connections
- Concurrent HTTP connections

Benefits of keeping idle HTTP connections open

You should keep HTTP connections open even when client applications are idle to allow client applications to perform subsequent transactions over the open connection. Based on system measurements and integration experience, you should keep an HTTP connection open for a maximum of 10 minutes. The LDR service might automatically close an HTTP connection that is kept open and idle for longer than 10 minutes.

Open and idle HTTP connections provide the following benefits:

- Reduced latency from the time that the StorageGRID Webscale system determines it has to perform an HTTP transaction to the time that the StorageGRID Webscale system can perform the transaction
Reduced latency is the main advantage, especially for the amount of time required to establish TCP/IP and TLS connections.
- Increased data transfer rate by priming the TCP/IP slow-start algorithm with previously performed transfers
- Instantaneous notification of several classes of fault conditions that interrupt connectivity between the client application and the StorageGRID Webscale system

Determining how long to keep an idle connection open is a trade-off between the benefits of slow start that is associated with the existing connection and the ideal allocation of the connection to internal system resources.

Benefits of active HTTP connections

You should limit the duration of an active HTTP connection for a maximum of 10 minutes, even if the HTTP connection continuously performs transactions. Determining the maximum duration that a connection should be held open is a trade-off between the benefits of connection persistence and the ideal allocation of the connection to internal system resources.

Limited active HTTP connections provide the following benefits:

- Enables optimal load balancing across the StorageGRID Webscale system

To optimize load balancing across the StorageGRID Webscale system, you should prevent long-lived TCP/IP connections. You should configure client applications to track the duration of each HTTP connection and close the HTTP connection after a set time so that the HTTP connection can be reestablished and rebalanced.

The StorageGRID Webscale system balances its load when a client application establishes an HTTP connection. Over time, an HTTP connection that the StorageGRID Webscale system uses for a compute resource might no longer be optimal as load balancing requirements change. The system performs its best load balancing when client applications establish a separate HTTP connection for each transaction, but this negates the much more valuable gains associated with persistent connections.

- Allows maintenance procedures to start
Some maintenance procedures start only after all the in-progress HTTP connections are complete.
- Allows client applications to direct HTTP transactions to LDR services that have available space.

Benefits of concurrent HTTP connections

You must keep multiple TCP/IP connections to the StorageGRID Webscale system open to allow idle connections to perform transactions as required. The number of client applications also affects how you handle multiple TCP/IP connections.

Concurrent HTTP connections provide the following benefits:

- Reduced latency
Transactions can start immediately instead of waiting for other transactions to be completed.
- Increased throughput
The StorageGRID Webscale system can perform parallel transactions and increase aggregate transaction throughput.

Client applications should establish multiple HTTP connections, either on a client-by-client basis or on a connection-pool basis. When a client application has to perform a transaction, it can select and immediately use any established connection that is not currently processing a transaction.

Each StorageGRID Webscale system's topology has different peak throughput for concurrent transactions and connections before performance begins to degrade. Peak throughput depends on factors such as computing resources, network resources, storage resources, and WAN links. The number of servers and services and the number of applications that the StorageGRID Webscale system supports are also factors.

StorageGRID Webscale systems often support multiple client applications. You should keep this in mind when you determine the maximum number of concurrent connections used by a client application. If the client application consists of multiple software entities that each establish connections to the StorageGRID Webscale system, you should add up all the connections across the entities. You might have to adjust the maximum number of concurrent connections in the following situations:

- The StorageGRID Webscale system's topology affects the maximum number of concurrent transactions and connections that the system can support.
- Client applications that interact with the StorageGRID Webscale system over a network with limited bandwidth might have to reduce the degree of concurrency to ensure that individual transactions are completed in a reasonable time.
- When many client applications share the StorageGRID Webscale system, you might have to reduce the degree of concurrency to avoid exceeding the limits of the system.

Copyright information

Copyright © 1994–2016 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

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, 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

- access, CDMI
 - root URI [25](#)
- algorithms
 - encryption [19](#)
 - hash [19](#)
 - supported by TLS [19](#)
 - supported hash for object storage [19](#)
- API Gateway Node
 - finding the port number of [17](#)
- API Gateway Nodes
 - IP address of [17](#)
 - IP addresses on CLB service [17](#)
- APIs
 - CDMI implementation [6](#)
- applications, client
 - permissions [11](#)
 - testing connection to the system [25](#)
 - testing HTTP connections [26](#)
- audience
 - CDMI information [5](#)
- audit logs
 - reviewing [31](#)
- authentication
 - client application access [17](#)
 - HTTP connections [18](#)

B

- best practices
 - for active HTTP connections [34](#)
 - for concurrent HTTP connections [35](#)
 - for idle HTTP connections [34](#)

C

- CA certificate
 - security partitions [21](#)
- CDMI
 - determining available functions [26](#)
 - how the StorageGRID Webscale system implements
 - immediate redundancy [9](#)
 - implementation [6](#)
 - introduction [5](#)
 - root URI [25](#)
 - security partitions [21](#)
 - security partitions, configure [22](#)
 - security partitions, enable [22](#)
 - supported specification sections [6](#)
 - version supported [5](#)
- CDMI client applications
 - object IDs [33](#)
 - retrieving content with UUIDs [32](#)
 - UUIDs [33](#)
- CDMI documentation
 - revision history [5](#)
- CDMI namespace

- delete access [12](#)
- last access time metadata [12](#)
- modify access [12](#)
- permission for client application access [11](#)
- permissions [11](#)
- read access [11](#)
- read access using GET method [11](#)
- what the Read access type is [11](#)
- write access [12](#)
- certificate authority (CA) certificates
 - security partitions [20](#)
- certificates
 - security partitions [21](#)
- CLB service
 - IP addresses [17](#)
- client applications
 - accessing the CDMI namespace [11](#)
 - associating IP addresses with [16](#)
 - associating with link costs [14](#)
 - authenticating access [17](#)
 - configuring StorageGRID Webscale to accept HTTP connections for [14](#)
 - connecting to the StorageGRID Webscale system [14](#)
 - deleting objects [12](#)
 - HTTP certificates [19](#)
 - managing HTTP connection states for [30](#)
 - permissions for [11](#)
 - testing HTTP connections [25, 26](#)
 - viewing HTTP transactions for [30](#)
- code examples
 - retrieving CDMI capabilities with curl [26](#)
 - retrieving CDMI objects by using curl [28](#)
 - storing CDMI objects by using curl [28](#)
- comments
 - how to send feedback about documentation [38](#)
- connecting
 - managing HTTP [30](#)
 - testing between client applications and the system [25](#)
- connections
 - benefits for concurrent HTTP [35](#)
 - benefits for idle HTTP [34](#)
 - best practices for active HTTP [34](#)
- connections, HTTP
 - testing using Telnet [25](#)
 - testing with the openssl command [26](#)
- container objects
 - creating using HTTP PUT method [12](#)
- curl
 - retrieving CDMI capabilities using [26](#)
 - using to test nameless data object retrieval [28](#)
 - using to test nameless data object storage [28](#)

D

- data objects
 - how ILM manages [9](#)
 - permission to read and retrieve [11](#)
 - reading with GET method [11](#)
 - storing using HTTP POST method [12](#)

DELETE method
 permissions for [11](#), [12](#)
 documentation
 how to receive automatic notification of changes to [38](#)
 how to send feedback about [38](#)
 dual commit
 how the StorageGRID Webscale system implements immediate redundancy [9](#)

E

encryption algorithms
 supported by TLS [19](#)

F

feedback
 how to send comments about documentation [38](#)

G

GET method
 access permissions [11](#)
 reading objects and object data [11](#)
 retrieving data [11](#)
 grid nodes
 IP addresses for [17](#)

H

hash algorithms
 supported by TLS [19](#)
 supported for object storage [19](#)

HTTP

certificates for security [19](#)
 DELETE access method [12](#)
 DELETE method and permissions [11](#)
 GET method and permissions [11](#)
 GET method to read objects and object data [11](#)
 GET method to retrieve object metadata [11](#)
 POST method and permissions [11](#)
 POST method to create and store objects [12](#)
 PUT method [11](#)
 version supported [5](#)

HTTP connections

associating IP addresses with client applications [16](#)
 benefits for concurrent [35](#)
 benefits for idle [34](#)
 benefits for idle, active, and concurrent [34](#)
 benefits of different types [34](#)
 best practices for active [34](#)
 configuring StorageGRID Webscale to accept [14](#)
 creating between client applications and the system [14](#)
 managing state of [30](#)
 testing client application [25](#)
 testing with Telnet [25](#)
 testing with the openssl command [26](#)
 used by client applications to access the system [30](#)

HTTP ports

finding [17](#)

finding for API Gateway Node [17](#)
 HTTP ports |||
 CLB service
 finding the port number of [17](#)
 finding for Storage Node [17](#)
 LDR service
 finding the port number of [17](#)
 port number
 finding for a CLB service [17](#)
 finding for an LDR service [17](#)

HTTP profiles

associating with client application IP addresses [16](#)
 defining permissions for client applications [15](#)

HTTP PUT method

use to update data object user metadata [12](#)

HTTP transactions

generated by CDMI client applications [30](#)
 viewing for client applications [30](#)

HTTPS connections

IP address for grid nodes [17](#)

I

ILM

and managing data objects [9](#)
 last access time metadata [12](#)
See also information lifecycle management

immediate redundancy

how the StorageGRID Webscale system implements [9](#)

information

how to send feedback about improving documentation [38](#)

information lifecycle management

rules and CDMI [6](#)
See also ILM

IP addresses

associating with client applications [16](#)
 associating with link costs [14](#)
 for API Gateway Nodes [17](#)
 for Storage Nodes [17](#)

L

Last Access Time [13](#)

last access time metadata

permissions for [11](#)
 used for ILM [9](#), [12](#)

Last Access Time, disable [13](#)

Last Access Time, enable [13](#)

LDR service

IP addresses [17](#)

link costs

associating with clients [14](#)

logs

reviewing audit [31](#)

M

metadata

GET method to retrieve [11](#)
 how used in CDMI object management [9](#)

- last access time [9](#), [11](#), [12](#)
- permission to retrieve [11](#)
- support for [6](#)
- updating [12](#)
- updating using the HTTP PUT method [12](#)
- viewing for objects [32](#)

modify access [11](#)

N

namespace

- creating HTTP profiles for permissions [15](#)
- permissions [11](#)

O

object IDs

- deriving from UUIDs [32](#)
- looking up [30](#)

object storage

- supported hash algorithms [19](#)

objects

- creating using HTTP PUT method [12](#)
- deleting from the CDMI namespace [12](#)
- duplicate names [9](#)
- GET method to retrieve metadata [11](#)
- how immediate redundancy works [9](#)
- how the system assigns UUIDs [32](#)
- identified by object IDs [33](#)
- identified by UUIDs [33](#)
- last access time metadata [12](#)
- retrieving affects last access time [12](#)
- storing [12](#)
- storing using HTTP POST method [12](#)
- testing retrieval of [28](#)
- testing storage of [28](#)
- UUIDs and IDs [33](#)
- viewing HTTP transactions for [30](#)
- viewing information about [32](#)

objects, dual commit

- how the StorageGRID Webscale system implements immediate redundancy [9](#)

OpenSSL

- using to test client applications to the system [25](#)
- using to test HTTP connections [26](#)

operations

- metadata, support for [6](#)
- supported CDMI specification [6](#)

P

permissions

- defining for CDMI client applications [15](#)
- for client applications [11](#)
- in the CDMI namespace [11](#)
- last access time metadata [12](#)
- storing objects in the CDMI namespace [12](#)
- updating object metadata in the CDMI namespace [12](#)

port number

- finding for a Storage Node [17](#)
- finding for an API Gateway Node [17](#)

POST method

- creating and storing objects [12](#)
- immediate redundancy and [9](#)
- permissions for [11](#)

PUT method

- access permissions [11](#)

R

read access [11](#)

redundancy

- how the StorageGRID Webscale system implements immediate [9](#)

root URI

- for CDMI access [25](#)

S

security

- certificates [19](#)
- partitions [19](#)
- using partitions to achieve [20](#)
- using TLS [18](#)

security partitions

- CA certificate [21](#)
- CDMI transactions [21](#)
- certificate [21](#)
- configure [22](#)
- enable [22](#)

servers

- HTTP certificates [19](#)
- link costs for [14](#)

SHA-1 hash algorithm

- support for object storage [19](#)

SHA-2 256 bits hash algorithm

- support for object storage [19](#)

states

- managing for an HTTP connection [30](#)

Storage Node

- finding the port number of [17](#)

Storage Nodes

- IP address of [17](#)
- IP addresses on LDR service [17](#)

StorageGRID Webscale

- accepting HTTP connections from client applications [14](#)
- available CDMI functions [26](#)
- CDMI permissions [11](#)
- immediate redundancy overview [9](#)
- supported CDMI specification sections [6](#)

suggestions

- how to send feedback about documentation [38](#)

T

Telnet

- using to test client applications to the system [25](#)
- using to test HTTP connections [25](#)

testing

- HTTP connections [25](#)
- HTTP connections using OpenSSL [26](#)
- HTTP connections using Telnet [25](#)

TLS

- HTTP certificates [19](#)
- supported hashing algorithms [19](#)
- Transport Layer Security
 - See* TLS
- troubleshooting
 - using audit logs [31](#)
- Twitter
 - how to receive automatic notification of documentation changes [38](#)

U

- universally unique identifiers
 - See* UUIDs
- URI, root
 - for CDMI access [25](#)
- UUIDs

- assigning to objects [32](#)
- defined [32](#)
- extracting object IDs from [32](#)
- managing content [33](#)

V

- versions
 - CDMI [5](#)
 - HTTP [5](#)

W

- write access
 - for client applications in the CDMI namespace [12](#)