



**StorageGRID® 11.3**

# **Swift Implementation Guide**

November 2019 | 215-14207\_2019-11\_en-us  
doccomments@netapp.com

 **NetApp®**



# Contents

<b>OpenStack Swift API support in StorageGRID .....</b>	<b>4</b>
History of Swift API support in StorageGRID .....	4
How StorageGRID implements the Swift REST API .....	4
Recommendations for implementing the Swift REST API .....	5
<b>Configuring tenant accounts and connections .....</b>	<b>7</b>
Creating and configuring Swift tenant accounts .....	7
Configuring client connections .....	8
Identifying IP addresses for client connections .....	9
Deciding to use HTTPS or HTTP connections .....	10
Testing your connection in the Swift API configuration .....	10
<b>Swift REST API supported operations .....</b>	<b>12</b>
Supported Swift API endpoints .....	12
Account operations .....	14
Container operations .....	16
Object operations .....	19
OPTIONS request .....	23
Error responses to Swift API operations .....	24
<b>StorageGRID Swift REST API operations .....</b>	<b>25</b>
GET container consistency request .....	25
PUT container consistency request .....	27
<b>Configuring security for the REST API .....</b>	<b>29</b>
How StorageGRID provides security for the REST API .....	29
Supported hashing and encryption algorithms for TLS libraries .....	30
<b>Monitoring and auditing operations .....</b>	<b>32</b>
Monitoring object ingest and retrieval rates .....	32
Accessing and reviewing audit logs .....	34
Swift operations tracked in the audit logs .....	34
<b>Copyright .....</b>	<b>36</b>
<b>Trademark information .....</b>	<b>37</b>
<b>How to send comments about documentation and receive update     notifications .....</b>	<b>38</b>

## OpenStack Swift API support in StorageGRID

StorageGRID supports the following specific versions of Swift and HTTP.

Item	Version
Swift specification	OpenStack Swift Object Storage API v1 as of November 2015
HTTP	1.1 For more information about HTTP, see HTTP/1.1 (RFCs 7230-35). <b>Note:</b> StorageGRID does not support HTTP/1.1 pipelining.

### Related information

[OpenStack: Object Storage API](#)

## History of Swift API support in StorageGRID

You should be aware of changes to the StorageGRID system's support for the Swift REST API.

Release	Comments
11.3	Updated PUT Object operations to describe the impact of ILM rules that use synchronous placement at ingest (the Balanced and Strict options for Ingest Behavior). Added description of client connections that use load balancer endpoints or high availability groups. Updated list of supported TLS cipher suites.
11.2	Minor editorial changes to document.
11.1	Added support for using HTTP for Swift client connections to grid nodes. Updated the definitions of consistency controls.
11.0	Added support for 1,000 containers for each tenant account.
10.4	Added support for POST container and PUT, GET, and HEAD container ACL headers for Keystone configured grids. <b>Note:</b> Keystone is disabled by default. To enable Keystone, contact your NetApp representative.
10.3	Administrative updates and corrections to the document. Removed sections for configuring custom server certificates.
10.2	Initial support of the Swift API by the StorageGRID system. The currently supported version is OpenStack Swift Object Storage API v1.

## How StorageGRID implements the Swift REST API

A client application can use Swift REST API calls to connect to Storage Nodes and Gateway Nodes to create containers and to store and retrieve objects. This enables service-oriented applications

developed for OpenStack Swift to connect with on-premise object storage provided by the StorageGRID system.

### Swift object management

After Swift objects have been ingested in the StorageGRID system, they are managed by the information lifecycle management (ILM) rules in the system's active ILM policy. The ILM rules and policy determine how StorageGRID creates and distributes copies of object data and how it manages those copies over time. For example, an ILM rule might apply to objects in specific Swift containers and might specify that multiple object copies be saved to several data centers for a certain number of years.

Contact your StorageGRID administrator if you need to understand how the grid's ILM rules and policies will affect the objects in your Swift tenant account.

### Conflicting client requests

Conflicting client requests, such as a two clients writing to the same key, are resolved on a “latest-wins” basis. The timing for the “latest-wins” evaluation is based on when the StorageGRID system completes a given request, and not on when Swift clients begin an operation.

### Consistency guarantees and controls

By default, StorageGRID provides read-after-write consistency for newly created objects and eventual consistency for object updates and HEAD operations. Any GET following a successfully completed PUT will be able to read the newly written data. Overwrites of existing objects, metadata updates, and deletes are eventually consistent. Overwrites generally take seconds or minutes to propagate, but can take up to 15 days.

StorageGRID also allows you to control consistency on a per container basis. You can change the consistency control to make a trade-off between the availability of the objects and the consistency of those objects across different Storage Nodes and sites, as required by your application.

### Related references

[GET container consistency request](#) on page 25

[PUT container consistency request](#) on page 27

### Related information

[Administering StorageGRID](#)

## Recommendations for implementing the Swift REST API

You should follow these recommendations when implementing the Swift REST API for use with StorageGRID.

### Recommendations for HEADs to non-existent objects

If your application routinely checks to see if an object exists at a path where you do not expect the object to actually exist, you should use the “Available” consistency control. For example, you should use the “Available” consistency control if your application performs a HEAD operation to a location before performing a PUT operation to that location.

Otherwise, if the HEAD operation does not find the object, you might receive a high number of 500 Internal Server errors if one or more Storage Nodes are unavailable.

You can set the “Available” consistency control for each bucket using the PUT Bucket consistency request.

### Recommendations for object names

You should not use random values as the first four characters of object names. Instead, you should use non-random, non-unique prefixes, such as `image`.

If you do need to use random and unique characters in object name prefixes, you should prefix the object names with a directory name. That is, use this format:

```
mycontainer/mydir/f8e3-image3132.jpg
```

Instead of this format:

```
mycontainer/f8e3-image3132.jpg
```

### Recommendations for “range reads”

If the **Stored Object Compression** grid option is enabled for StorageGRID, Swift client applications should avoid performing GET object operations that specify a range of bytes be returned. These “range read” operations are inefficient because StorageGRID must effectively uncompress the objects to access the requested bytes. GET Object operations that request a small range of bytes from a very large object are especially inefficient; for example, it is very inefficient to read a 10 MB range from a 50 GB compressed object.

If ranges are read from compressed objects, client requests can time out.

**Note:** If you need to compress objects and your client application must use range reads, increase the read timeout for the application.

### Related references

[GET container consistency request](#) on page 25

[PUT container consistency request](#) on page 27

### Related information

[Administering StorageGRID](#)

# Configuring tenant accounts and connections

---

Configuring StorageGRID to accept connections from client applications requires creating one or more tenant accounts and setting up the connections.

## Steps

1. [Creating and configuring Swift tenant accounts](#) on page 7
2. [Configuring client connections](#) on page 8
3. [Testing your connection in the Swift API configuration](#) on page 10

## Creating and configuring Swift tenant accounts

A Swift tenant account is required before Swift API clients can store and retrieve objects on StorageGRID. Each tenant account has its own account ID, groups and users, and containers and objects.

Swift tenant accounts are created by a StorageGRID grid administrator using the Grid Manager or the Grid Management API.

When creating a Swift tenant account, the grid administrator specifies the following information:

- Display name for the tenant (the tenant's account ID is assigned automatically and cannot be changed)
- Optionally, a storage quota for the tenant account—the maximum number of gigabytes, terabytes, or petabytes available for the tenant's objects. A tenant's storage quota represents a logical amount (object size), not a physical amount (size on disk).
- If single sign-on (SSO) is not in use for the StorageGRID system, whether the tenant account will use its own identity source or share the grid's identity source, and the initial password for the tenant's local root user.
- If SSO is enabled, which federated group has Root Access permission to configure the tenant account.

After a Swift tenant account is created, users with the Root Access permission can access the Tenant Manager to perform tasks such as the following:

- Setting up identity federation (unless the identity source is shared with the grid), and creating local groups and users
- Monitoring storage usage

**Attention:** Swift users must have the Root Access permission to access the Tenant Manager. However, the Root Access permission does not allow users to authenticate into the Swift REST API to create containers and ingest objects. Users must have the Administrator permission to authenticate into the Swift REST API.

## Related references

[Supported Swift API endpoints](#) on page 12

## Related information

[Administering StorageGRID](#)  
[Using tenant accounts](#)

## Configuring client connections

A grid administrator makes configuration choices that affect how Swift clients connect to StorageGRID to store and retrieve data. The specific information you need to make a connection depends upon the configuration that was chosen.

Client applications can store or retrieve objects by connecting to any of the following:

- The Load Balancer service on Admin Nodes or Gateway Nodes, or optionally, the virtual IP address of a high availability (HA) group of Admin Nodes or Gateway Nodes
- The legacy CLB service on Gateway Nodes, or optionally, the virtual IP address of a high availability group of Gateway Nodes
- Storage Nodes, with or without an external load balancer

The Load Balancer service is new in StorageGRID 11.3. Existing clients can continue to use the legacy CLB service on Gateway Nodes to connect to StorageGRID. New client applications that depend on StorageGRID to provide load balancing should connect using the Load Balancer service.

When configuring StorageGRID, a grid administrator can use the Grid Manager or the Grid Management API to perform the following steps, all of which are optional:

1. **Configure endpoints for the Load Balancer service.**  
You must configure endpoints to use the Load Balancer service. The Load Balancer service on Admin Nodes or Gateway Nodes distributes incoming network connections from client applications to Storage Nodes. When creating a load balancer endpoint, the grid administrator specifies a port number, whether the endpoint accepts HTTP or HTTPS connections, the type of client (S3 or Swift) that will use the endpoint, and the certificate to be used for HTTPS connections (if applicable).
2. **Configure Untrusted Client Networks.**  
If a grid administrator configures a node's Client Network to be untrusted, the node only accepts inbound connections on the Client Network on ports that are explicitly configured as load balancer endpoints.
3. **Configure high availability groups.**  
If an administrator creates an HA group, the network interfaces of multiple Admin Nodes or Gateway Nodes are placed into an active-backup configuration. Client connections are made using the virtual IP address of the HA group.

For more information about each option, see the instructions for administering StorageGRID.

### Information required to make client connections

The table summarizes the information that you require to make a Swift client connection for each configuration.

Where connection is made	Service that client connects to	IP address	Port
HA group	Load Balancer	Virtual IP of HA group	Load balancer endpoint port
Admin Node	Load Balancer	Admin Node IP address	Load balancer endpoint port
Gateway Node	Load Balancer	Gateway Node IP address	Load balancer endpoint port



Where connection is made	Service that client connects to	IP address	Port
Gateway Node	Legacy CLB	Gateway Node IP address	<ul style="list-style-type: none"> <li>• HTTPS:8083</li> <li>• HTTP:8085</li> </ul>
Storage Node	LDR	Storage Node IP address	<ul style="list-style-type: none"> <li>• HTTPS: 18083</li> <li>• HTTP:18085</li> </ul>

It is possible to configure a DNS name for the IP address that clients use to connect to StorageGRID. Contact your local network administrator.

## Identifying IP addresses for client connections

Client applications can use the IP address of a Gateway Node, Admin Node, or Storage Node to connect to StorageGRID. If high availability groups are configured, client applications can connect using the virtual IP address of the HA group.

### About this task

You can use the Grid Manager to look up the IP address or virtual IP address that you require to make a client connection to StorageGRID.

### Steps

1. Sign in to the Grid Manager using a supported browser.
2. To find the IP address of a grid node:
  - a. Select **Nodes**.
  - b. Select the Admin Node, Gateway Node, or Storage Node to which you want to connect.
  - c. Select the **Overview** tab.
  - d. In the Node Information section, note the IP addresses for the node.
  - e. Click **Show more** to view IPv6 addresses and interface mappings.

You can establish connections from client applications to any of the IP addresses in the list:

- **eth0:** Grid Network
- **eth1:** Admin Network (optional)
- **eth2:** Client Network (optional)

**Note:** If you are viewing an Admin Node or a Gateway Node and it is the active node in a high availability group, the virtual IP address of the HA group is shown on eth2.

3. To find the virtual IP address of a high availability group:
  - a. Select **Configuration > High Availability Groups**.
  - b. In the table, note the virtual IP address of the HA group.

### Related information

[Administering StorageGRID](#)

## Deciding to use HTTPS or HTTP connections

When client connections are made using a Load Balancer endpoint, connections must be made using the protocol (HTTP or HTTPS) that was specified for that endpoint. To use HTTP for client connections to Storage Nodes or to the legacy CLB service on Gateway Nodes, you must enable its use.

By default, when connecting to Storage Nodes or the legacy CLB service on Gateway Nodes, client applications must use encrypted HTTPS for all connections. Optionally, you can enable less-secure HTTP connections by enabling the **HTTP** option from the Grid Manager. For example, a client application might use HTTP when testing the connection to a Storage Node in a non-production environment.

**Attention:** Be careful when enabling HTTP for a production grid since requests will be sent unencrypted.

To learn how to enable this option, see information about administering StorageGRID.

If the **HTTP** option has been enabled for the grid, for connections to Storage Nodes or the legacy CLB service on Gateway Nodes you use different ports for HTTP than you use for HTTPS.

### Related information

[Administering StorageGRID](#)

## Testing your connection in the Swift API configuration

You can use the Swift CLI to test your connection to the StorageGRID system and to verify that you can read and write objects to the system.

### Before you begin

- You must have downloaded and installed *python-swiftclient*, the Swift command-line client, at <https://swiftstack.com/docs/integration/python-swiftclient.html>.
- You must have a Swift tenant account in the StorageGRID system.

### About this task

If you have not configured security, you must add the `--insecure` flag to each of these commands.

### Steps

1. Query the info URL for your StorageGRID Swift deployment:

```
swift
-U <Tenant_Account_ID:User_Name>
-K <User_Password>
-A https://<FQDN | IP>:<Port>/info
capabilities
```

This is sufficient to test that your Swift deployment is functional. To further test account configuration by storing an object, continue with the additional steps.

2. Put an object in the container:

```
touch test_object
swift
-U <Tenant_Account_ID:Account_User_Name>
```

```
-K <User_Password>
-A https://<FQDN | IP>:<Port>/auth/v1.0
upload test_container test_object
--object-name test_object
```

3. Get the container to verify the object:

```
swift
-U <Tenant_Account_ID:Account_User_Name>
-K <User_Password>
-A https://<FQDN | IP>:<Port>/auth/v1.0
list test_container
```

4. Delete the object:

```
swift
-U <Tenant_Account_ID:Account_User_Name>
-K <User_Password>
-A https://<FQDN | IP>:<Port>/auth/v1.0
delete test_container test_object
```

5. Delete the container:

```
swift
-U <Tenant_Account_ID:Account_User_Name>
-K <User_Password>
-A https://<FQDN | IP>:<Port>/auth/v1.0
delete test_container
```

#### Related concepts

[Creating and configuring Swift tenant accounts](#) on page 7

[Configuring security for the REST API](#) on page 29

## Swift REST API supported operations

---

The StorageGRID system supports most operations in the OpenStack Swift API. Before integrating Swift REST API clients with StorageGRID, review the implementation details for account, container, and object operations.

### Operations supported in StorageGRID

The following Swift API operations are supported:

- [Account operations](#) on page 14
- [Container operations](#) on page 16
- [Object operations](#) on page 19

### Common response headers for all operations

The StorageGRID system implements all common headers for supported operations as defined by the OpenStack Swift Object Storage API v1.

### Related information

[OpenStack: Object Storage API](#)

## Supported Swift API endpoints

StorageGRID supports the following Swift API endpoints: the info URL, the auth URL, and the storage URL.

### info URL

You can determine the capabilities and limitations of the StorageGRID Swift implementation by issuing a GET request to the Swift base URL with the `/info` path.

```
https://FQDN | Node_IP:Swift_Port/info/
```

In the request:

- *FQDN* is the fully qualified domain name.
- *Node\_IP* is the IP address for the Storage Node or the Gateway Node on the StorageGRID network.
- *Swift\_Port* is the port number used for Swift API connections on the Storage Node or Gateway Node.

For example, the following info URL would request information from a Storage Node with the IP address of 10.99.106.103 and using port 18083.

```
https://10.99.106.103:18083/info/
```

The response includes the capabilities of the Swift implementation as a JSON dictionary. A client tool can parse the JSON response to determine the capabilities of the implementation and use them as constraints for subsequent storage operations.

The StorageGRID implementation of Swift allows unauthenticated access to the info URL.

### auth URL

A client can use the Swift auth URL to authenticate as a tenant account user.

```
https://FQDN | Node_IP:Swift_Port/auth/v1.0/
```

You must provide the tenant account ID, user name, and password as parameters in the X-Auth-User and X-Auth-Key request headers, as follows:

```
X-Auth-User: Tenant_Account_ID:Username
```

```
X-Auth-Key: Password
```

In the request headers:

- *Tenant\_Account\_ID* is the account ID assigned by StorageGRID when the Swift tenant was created. This is the same tenant account ID used on the Tenant Manager sign-in page.
- *Username* is the name of a tenant user that has been created in the Tenant Manager. This user must belong to a group that has the Administrator permission. The tenant's root user cannot be configured to use the Swift REST API.

If Identity Federation is enabled for the tenant account, provide the username and password of the federated user from the LDAP server. Alternatively, provide the LDAP user's domain name. For example:

```
X-Auth-User: Tenant_Account_ID:Username@Domain_Name
```

- *Password* is the password for the tenant user. User passwords are created and managed in the Tenant Manager.

The response to a successful authentication request returns a storage URL and an auth token, as follows:

```
X-Storage-Url: https://FQDN | Node_IP:Swift_Port/v1/Tenant_Account_ID
X-Auth-Token: token
X-Storage-Token: token
```

By default, the token is valid for 24 hours from generation time.

Tokens are generated for a specific tenant account. A valid token for one account does not authorize a user to access another account.

### storage URL

A client application can issue Swift REST API calls to perform supported account, container, and object operations against a Gateway Node or Storage Node. Storage requests are addressed to the storage URL returned in the authentication response. The request must also include the X-Auth-Token header and value returned from the auth request.

```
https://FQDN | IP:Swift_Port/v1/Tenant_Account_ID
[/container][/object]
X-Auth-Token: token
```

Some storage response headers that contain usage statistics might not reflect accurate numbers for recently modified objects. It might take a few minutes for accurate numbers to appear in these headers.

The following response headers for account and container operations are examples of those that contain usage statistics:

- X-Account-Bytes-Used
- X-Account-Object-Count
- X-Container-Bytes-Used
- X-Container-Object-Count

#### Related concepts

[Configuring client connections](#) on page 8

[Creating and configuring Swift tenant accounts](#) on page 7

#### Related tasks

[Identifying IP addresses for client connections](#) on page 9

#### Related references

[Account operations](#) on page 14

[Container operations](#) on page 16

[Object operations](#) on page 19

## Account operations

The following Swift API operations are performed on accounts.

### GET account

This operation retrieves the container list associated with the account and account usage statistics.

The following request parameter is required:

- Account

The following request header is required:

- X-Auth-Token

The following supported request query parameters are optional:

- Delimiter
- End\_marker
- Format
- Limit
- Marker
- Prefix

A successful execution returns the following headers with an “HTTP/1.1 204 No Content” response if the account is found and has no containers or the container list is empty; or an “HTTP/1.1 200 OK” response if the account is found and the container list is not empty:

- Accept-Ranges
- Content-Length
- Content-Type
- Date
- X-Account-Bytes-Used
- X-Account-Container-Count
- X-Account-Object-Count
- X-Timestamp
- X-Trans-Id

### **HEAD account**

This operation retrieves account information and statistics from a Swift account.

The following request parameter is required:

- Account

The following request header is required:

- X-Auth-Token

A successful execution returns the following headers with an “HTTP/1.1 204 No Content” response:

- Accept-Ranges
- Content-Length
- Date
- X-Account-Bytes-Used
- X-Account-Container-Count
- X-Account-Object-Count
- X-Timestamp
- X-Trans-Id

### **Related references**

*Swift operations tracked in the audit logs* on page 34

## Container operations

StorageGRID supports a maximum of 1,000 containers per Swift account. The following Swift API operations are performed on containers.

### DELETE container

This operation removes an empty container from a Swift account in a StorageGRID system.

The following request parameters are required:

- Account
- Container

The following request header is required:

- X-Auth-Token

A successful execution returns the following headers with an "HTTP/1.1 204 No Content" response:

- Content-Length
- Content-Type
- Date
- X-Trans-Id

### GET container

This operation retrieves the object list associated with the container along with container statistics and metadata in a StorageGRID system.

The following request parameters are required:

- Account
- Container

The following request header is required:

- X-Auth-Token

The following supported request query parameters are optional:

- Delimiter
- End\_marker
- Format
- Limit
- Marker
- Path
- Prefix

A successful execution returns the following headers with an "HTTP/1.1 200 Success" or a "HTTP/1.1 204 No Content" response:

- Accept-Ranges
- Content-Length



- Content-Type
- Date
- X-Container-Bytes-Used
- X-Container-Object-Count
- X-Timestamp
- X-Trans-Id

When this operation is set in Keystone enabled configurations, the following headers are returned to admin users:

- X-Container-Read
- X-Container-Write

**Note:** Keystone is disabled by default. To enable Keystone, contact your NetApp representative.

### HEAD container

This operation retrieves container statistics and metadata from a StorageGRID system.

The following request parameters are required:

- Account
- Container

The following request header is required:

- X-Auth-Token

A successful execution returns the following headers with an "HTTP/1.1 204 No Content" response:

- Accept-Ranges
- Content-Length
- Date
- X-Container-Bytes-Used
- X-Container-Object-Count
- X-Timestamp
- X-Trans-Id

When this operation is set in Keystone enabled configurations, the following headers are returned to admin users:

- X-Container-Read
- X-Container-Write

**Note:** Keystone is disabled by default. To enable Keystone, contact your NetApp representative.

### POST container

This operation creates, modifies, or deletes the ACL metadata for an existing container by an admin user in a Keystone configured StorageGRID system.

**Note:** This operation is supported only for Swift Keystone accounts. Keystone is disabled by default. To enable Keystone, contact your NetApp representative.

When Keystone is disabled (default), a status of Not Implemented is returned for this operation.

The following request parameters are required:

- Account
- Container

The following request header is required:

- X-Auth-Token

The following request headers are optional:

- X-Container-Read

**Note:** You can specify a referrer in the X-Container-Read metadata, but StorageGRID ignores this value.

- X-Container-Write
- X-Remove-Container-Read
- X-Remove-Container-Write

**Note:** No other metadata operations are supported and will result in the operation being ignored.

A successful execution returns the following headers with an "HTTP/1.1 204 No Content" response:

- Content-Length
- Date
- X-Timestamp
- X-Trans-Id

## PUT container

This operation creates a container for an account in a StorageGRID system.

The following request parameters are required:

- Account
- Container

The following request header is required:

- X-Auth-Token

A successful execution returns the following headers with an "HTTP/1.1 201 Created" or "HTTP/1.1 202 Accepted" (if the container already exists under this account) response:

- Content-Length
- Date
- X-Timestamp
- X-Trans-Id

A container name must be unique in the StorageGRID namespace. If the container exists under another account, the following header is returned: "HTTP/1.1 409 Conflict."

The following optional headers are supported only for admin users in Keystone enabled configurations:

- X-Container-Read

**Note:** You can specify a referrer in the X-Container-Read metadata, but StorageGRID ignores this value.

- X-Container-Write
- X-Remove-Container-Read
- X-Remove-Container-Write

**Note:** These optional headers are supported only for Swift Keystone accounts. Keystone is disabled by default. To enable Keystone, contact your NetApp representative.

#### Related references

[Swift operations tracked in the audit logs](#) on page 34

## Object operations

The following Swift API operations are performed on objects.

### DELETE object

This operation deletes an object's content and metadata from the StorageGRID system.

The following request parameters are required:

- Account
- Container
- Object

The following request header is required:

- X-Auth-Token

A successful execution returns the following response headers with an "HTTP/1.1 204 No Content" response:

- Content-Length
- Content-Type
- Date
- X-Trans-Id

When processing a DELETE Object request, StorageGRID attempts to immediately remove all copies of the object from all stored locations. If successful, StorageGRID returns a response to the client immediately. If all copies cannot be removed within 30 seconds (for example, because a location is temporarily unavailable), StorageGRID queues the copies for removal and then indicates success to the client.

In previous releases, StorageGRID always queued objects for removal. With the introduction of immediate deletion, clients might sometimes receive a slower response, even though object copies are generally being removed more quickly.

For more information on how objects are deleted, see the instructions for administering StorageGRID.

### **GET object**

This operation retrieves the object content and gets the object metadata from a StorageGRID system.

The following request parameters are required:

- Account
- Container
- Object

The following request header is required:

- X-Auth-Token

The following request headers are optional:

- Accept-Encoding
- If-Match
- If-Modified-Since
- If-None-Match
- If-Unmodified-Since
- Range

A successful execution returns the following headers with an "HTTP/1.1 200 OK" response:

- Accept-Ranges
- Content-Disposition, returned only if Content-Disposition metadata was set
- Content-Encoding, returned only if Content-Encoding metadata was set
- Content-Length
- Content-Type
- Date
- ETag
- Last-Modified
- X-Timestamp
- X-Trans-Id

### **HEAD object**

This operation retrieves metadata and properties of an ingested object from a StorageGRID system.

The following request parameters are required:

- Account

- Container
- Object

The following request header is required:

- X-Auth-Token

A successful execution returns the following headers with an "HTTP/1.1 200 OK" response:

- Accept-Ranges
- Content-Disposition, returned only if Content-Disposition metadata was set
- Content-Encoding, returned only if Content-Encoding metadata was set
- Content-Length
- Content-Type
- Date
- ETag
- Last-Modified
- X-Timestamp
- X-Trans-Id

### PUT object

This operation creates a new object with data and metadata, or replaces an existing object with data and metadata in a StorageGRID system.

**Note:** StorageGRID supports objects up to 5 TB in size.

**Attention:** Conflicting client requests, such as a two clients writing to the same key, are resolved on a “latest-wins” basis. The timing for the “latest-wins” evaluation is based on when the StorageGRID system completes a given request, and not on when Swift clients begin an operation.

The following request parameters are required:

- Account
- Container
- Object

The following request header is required:

- X-Auth-Token

The following request headers are optional:

- Content-Disposition
- Content-Encoding

Do not use chunked Content-Encoding if the ILM rule that applies to an object filters objects based on size and uses synchronous placement on ingest (the Balanced or Strict options for Ingest Behavior).

- Transfer-Encoding

Do not use compressed or chunked `Transfer-Encoding` if the ILM rule that applies to an object filters objects based on size and uses synchronous placement on ingest (the `Balanced` or `Strict` options for `Ingest Behavior`).

- `Content-Length`

If an ILM rule filters objects by size and uses synchronous placement on ingest, you must specify `Content-Length`.

**Note:** If you do not follow these guidelines for `Content-Encoding`, `Transfer-Encoding`, and `Content-Length`, StorageGRID must save the object before it can determine object size and apply the ILM rule. In other words, StorageGRID must default to creating interim copies of an object on ingest. That is, StorageGRID must use the `Dual Commit` option for `Ingest Behavior`.

For more information about synchronous placement and ILM rules, see `Information Lifecycle Management` in the instructions for administering StorageGRID.

- `Content-Type`
- `ETag`
- `X-Object-Meta-<name>` (object-related metadata)

If you want to use the **User Defined Creation Time** option as the `Reference Time` for an ILM rule, you must store the value in a user-defined header named `X-Object-Meta-Creation-Time`. For example:

```
X-Object-Meta-Creation-Time: 1443399726
```

This field is evaluated as seconds since January 1, 1970.

- `X-Storage-Class: reduced_redundancy`

This header affects how many object copies StorageGRID creates if the ILM rule that matches an ingested object specifies an `Ingest Behavior` of `Dual Commit` or `Balanced`.

- **Dual commit:** If the ILM rule specifies the `Dual commit` option for `Ingest Behavior`, StorageGRID creates a single interim copy as the object is ingested (single commit).
- **Balanced:** If the ILM rule specifies the `Balanced` option, StorageGRID makes a single interim copy only if the system cannot immediately make all copies specified in the rule. If StorageGRID can perform synchronous placement, this header has no effect.

The `reduced_redundancy` header is best used when the ILM rule that matches the object creates a single replicated copy. In this case using `reduced_redundancy` eliminates the unnecessary creation and deletion of an extra object copy for every ingest operation.

Using the `reduced_redundancy` header is not recommended in other circumstances because it increases the risk the loss of object data during ingest. For example, you might lose data if the single copy is initially stored on a Storage Node that fails before ILM evaluation can occur.

Note that specifying `reduced_redundancy` only affects how many copies are created when an object is first ingested. It does not affect how many copies of the object are made when the object is evaluated by the active ILM policy and does not result in data being stored at lower levels of redundancy in the StorageGRID system.

A successful execution returns the following headers with an "HTTP/1.1 201 Created" response:

- `Content-Length`
- `Content-Type`
- `Date`

- ETag
- Last-Modified
- X-Trans-Id

**Related references**

[Swift operations tracked in the audit logs](#) on page 34

**Related information**

[Administering StorageGRID](#)

## OPTIONS request

The OPTIONS request checks the availability of an individual Swift service. The OPTIONS request is processed by the Storage Node or Gateway Node specified in the URL.

For example, client applications can issue an OPTIONS request to the Swift port on a Storage Node, without providing Swift authentication credentials, to determine whether the Storage Node is available. You can use this request for monitoring or to allow external load balancers to identify when a Storage Node is down.

**OPTIONS method**

When used with the info URL or the storage URL, the OPTIONS method returns a list of supported verbs for the given URL (for example, HEAD, GET, OPTIONS, and PUT). The OPTIONS method cannot be used with the auth URL.

The following request parameter is required:

- Account

The following request parameters are optional:

- Container
- Object

A successful execution returns the following headers with an “HTTP/1.1 204 No Content” response. The OPTIONS request to the storage URL does not require that the target exists.

- Allow (a list of supported verbs for the given URL, for example, HEAD, GET, OPTIONS, and PUT)
- Content-Length
- Content-Type
- Date
- X-Trans-Id

**Related references**

[Supported Swift API endpoints](#) on page 12

## Error responses to Swift API operations

Understanding the possible error responses can help you troubleshoot operations.

The following HTTP status codes might be returned when errors occur during an operation:

Swift error name	HTTP status
AccountNameTooLong, ContainerNameTooLong, HeaderTooBig, InvalidContainerName, InvalidRequest, InvalidURI, MetadataNameTooLong, MetadataValueTooBig, MissingSecurityHeader, ObjectNameTooLong, TooManyContainers, TooManyMetadataItems, TotalMetadataTooLarge	400 Bad Request
AccessDenied	403 Forbidden
ContainerNotEmpty, ContainerAlreadyExists	409 Conflict
InternalError	500 Internal Server Error
InvalidRange	416 Requested Range Not Satisfiable
MethodNotAllowed	405 Method Not Allowed
MissingContentLength	411 Length Required
NotFound	404 Not Found
NotImplemented	501 Not Implemented
PreconditionFailed	412 Precondition Failed
ResourceNotFound	404 Not Found
Unauthorized	401 Unauthorized
UnprocessableEntity	422 Unprocessable Entity



## StorageGRID Swift REST API operations

---

There are operations added on to the Swift REST API that are specific to StorageGRID system.

### GET container consistency request

Consistency level makes a trade-off between the availability of the objects and the consistency of those objects across different Storage Nodes and sites. The GET container consistency request allows you to determine the consistency level being applied to a particular container.

#### Request

Request HTTP Header	Description
X-Auth-Token	Specifies the Swift authentication token for the account to use for the request.
x-ntap-sg-consistency	Specifies the type of request, where <code>true</code> = GET container consistency, and <code>false</code> = GET container.
Host	The hostname to which the request is directed.

#### Request example

```
GET /v1/28544923908243208806/Swift container
X-Auth-Token: SGRD_3a877009a2d24cb1801587bfa9050f29
x-ntap-sg-consistency: true
Host: test.com
```

#### Response

Response HTTP Header	Description
Date	The date and time of the response.
Connection	Whether the connection to the server is open or closed.
X-Trans-Id	The unique transaction identifier for the request.
Content-Length	The length of the response body.

Response HTTP Header	Description
x-ntap-sg-consistency	<p>The consistency control level being applied to the container. The following values are supported:</p> <ul style="list-style-type: none"> <li>• <b>all</b>: All nodes receive the data immediately or the request will fail.</li> <li>• <b>strong-global</b>: Guarantees read-after-write consistency for all client requests across all sites.</li> <li>• <b>strong-site</b>: Guarantees read-after-write consistency for all client requests within a site.</li> <li>• <b>read-after-new-write</b>: Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees.</li> </ul> <p><b>Note:</b> If your application uses HEAD requests on objects that do not exist, you might receive a high number of 500 Internal Server errors if one or more Storage Nodes are unavailable. To prevent these errors, use the “available” level.</p> <ul style="list-style-type: none"> <li>• <b>available</b> (eventual consistency for HEAD operations): Behaves the same as the “read-after-new-write” consistency level, but only provides eventual consistency for HEAD operations. Offers higher availability for HEAD operations than “read-after-new-write” if Storage Nodes are unavailable.</li> <li>• <b>weak</b>: Provides eventual consistency and high availability, with minimal data protection guarantees, especially if a Storage Node fails or is unavailable. Suitable only for write-heavy workloads that require high availability, do not require read-after-write consistency, and can tolerate the potential loss of data if a node fails.</li> </ul>

### Response example

```

HTTP/1.1 204 No Content
Date: Sat, 29 Nov 2015 01:02:18 GMT
Connection: CLOSE
X-Trans-Id: 1936575373
Content-Length: 0
x-ntap-sg-consistency: strong-site

```

### Related information

[Using tenant accounts](#)

## PUT container consistency request

The PUT container consistency request allows you to specify the consistency level to apply to operations performed on a container. By default, new containers are created using the “read-after-new-write” consistency level.

### Request

Request HTTP Header	Description
X-Auth-Token	The Swift authentication token for the account to use for the request.
x-ntap-sg-consistency	<p>The consistency control level to apply to operations on the container. The following values are supported:</p> <ul style="list-style-type: none"> <li>• <b>all</b>: All nodes receive the data immediately or the request will fail.</li> <li>• <b>strong-global</b>: Guarantees read-after-write consistency for all client requests across all sites.</li> <li>• <b>strong-site</b>: Guarantees read-after-write consistency for all client requests within a site.</li> <li>• <b>read-after-new-write</b>: Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees. <ul style="list-style-type: none"> <li><b>Note</b>: If your application uses HEAD requests on objects that do not exist, you might receive a high number of 500 Internal Server errors if one or more Storage Nodes are unavailable. To prevent these errors, use the “available” level.</li> </ul> </li> <li>• <b>available</b> (eventual consistency for HEAD operations): Behaves the same as the “read-after-new-write” consistency level, but only provides eventual consistency for HEAD operations. Offers higher availability for HEAD operations than “read-after-new-write” if Storage Nodes are unavailable.</li> <li>• <b>weak</b>: Provides eventual consistency and high availability, with minimal data protection guarantees, especially if a Storage Node fails or is unavailable. Suitable only for write-heavy workloads that require high availability, do not require read-after-write consistency, and can tolerate the potential loss of data if a node fails.</li> </ul>
Host	The hostname to which the request is directed.

### Request example

```
PUT /v1/28544923908243208806/Swift container
X-Auth-Token: SGRD_3a877009a2d24cb1801587bfa9050f29
```

```
x-ntap-sg-consistency: strong-site  
Host: test.com
```

### Response

Response HTTP Header	Description
Date	The date and time of the response.
Connection	Whether the connection to the server is open or closed.
X-Trans-Id	The unique transaction identifier for the request.
Content-Length	The length of the response body.

### Response example

```
HTTP/1.1 204 No Content  
Date: Sat, 29 Nov 2015 01:02:18 GMT  
Connection: CLOSE  
X-Trans-Id: 1936575373  
Content-Length: 0
```

### Related information

[Using tenant accounts](#)

## Configuring security for the REST API

---

You should review the security measures implemented for the REST API and understand how to secure your system.

### How StorageGRID provides security for the REST API

You should understand how the StorageGRID system implements security, authentication, and authorization for the REST API.

StorageGRID uses the following security measures.

- Client communications with the Load Balancer service use HTTPS if HTTPS is configured for the load balancer endpoint.  
When you configure a load balancer endpoint, HTTP can optionally be enabled. For example, you might want to use HTTP for testing or other non-production purposes. See the instructions for administering StorageGRID for more information.
- By default, StorageGRID uses HTTPS for client communications with Storage Nodes and the legacy CLB service on Gateway Nodes.  
HTTP can optionally be enabled for these connections. For example, you might want to use HTTP for testing or other non-production purposes. See the instructions for administering StorageGRID for more information.
- Communications between StorageGRID and the client are encrypted using TLS.
- Communications between the Load Balancer service and Storage Nodes within the grid are encrypted whether the load balancer endpoint is configured to accept HTTP or HTTPS connections.
- Clients must supply HTTP authentication headers to StorageGRID to perform REST API operations.

#### Security certificates and client applications

Clients can connect to the Load Balancer service on Gateway Nodes or Admin Nodes, directly to Storage Nodes, or to the legacy CLB service on Gateway Nodes.

In all cases, client applications can make TLS connections using either a custom server certificate uploaded by the grid administrator or a certificate generated by the StorageGRID system:

- When client applications connect to the Load Balancer service, they do so using the certificate that was configured for the specific load balancer endpoint used to make the connection. Each endpoint has its own certificate, which is either a custom server certificate uploaded by the grid administrator or a certificate that the grid administrator generated in StorageGRID when configuring the endpoint.
- When client applications connect directly to a Storage Node or to the legacy CLB service on Gateway Nodes, they use either the system-generated server certificates that were generated for Storage Nodes when the StorageGRID system was installed (which are signed by the system certificate authority), or a single custom server certificate that is supplied for the grid by a grid administrator.

Clients should be configured to trust the certificate authority that signed whichever certificate they use to establish TLS connections.

See the instructions for administering StorageGRID for information on configuring load balancer endpoints, and for instructions on adding a single custom server certificate for TLS connections directly to Storage Nodes or to the legacy CLB service on Gateway Nodes.

## Summary

The following table shows how security issues are implemented in the S3 and Swift REST APIs:

Security issue	Implementation for REST API
Connection security	TLS
Server authentication	X.509 server certificate signed by system CA or custom server certificate supplied by administrator
Client authentication	<ul style="list-style-type: none"> <li>S3: S3 account (access key ID and secret access key)</li> <li>Swift: Swift account (user name and password)</li> </ul> <p><b>Note:</b> By request, you can enable OpenStack's Keystone identity service for use with the Swift REST API. If Keystone is enabled, you must use an additional token for validation. To enable Keystone support, contact your NetApp representative.</p>
Client authorization	<ul style="list-style-type: none"> <li>S3: Bucket ownership and all applicable access control policies</li> <li>Swift: Administrator role access</li> </ul>

## Related information

[Administering StorageGRID](#)

## Supported hashing and encryption algorithms for TLS libraries

The StorageGRID system supports a limited set of cipher suites that client applications can use when establishing a Transport Layer Security (TLS) session.

### Supported versions of TLS

StorageGRID supports TLS 1.2.

**Attention:** SSLv3 and TLS v1.1 (or earlier versions) are no longer supported.

### Supported cipher suites

IANA Name
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256

**Deprecated cipher suites**

The following cipher suites are deprecated. Support for these ciphers will be removed in a future release.

<b>IANA Name</b>
TLS_RSA_WITH_AES_128_GCM_SHA256
TLS_RSA_WITH_AES_256_GCM_SHA384

## Monitoring and auditing operations

---

You can monitor workloads and efficiencies for client operations by viewing transaction trends for the entire grid, or for specific nodes. You can use audit messages to monitor client operations and transactions.

### Steps

1. [Monitoring object ingest and retrieval rates](#) on page 32
2. [Accessing and reviewing audit logs](#) on page 34

## Monitoring object ingest and retrieval rates

You can monitor object ingest and retrieval rates as well as metrics for object counts, queries, and verification. You can view the number of successful and failed attempts by client applications to read, write, and modify objects in the StorageGRID system.

### Steps

1. Sign in to the Grid Manager using a supported browser.
2. On the Dashboard, locate the Protocol Operations section.  

This section summarizes the number of client operations performed by your StorageGRID system. Protocol rates are averaged over the last two minutes.
3. Select **Nodes**.
4. From the Nodes home page (deployment level), click the **Load Balancer** tab.  

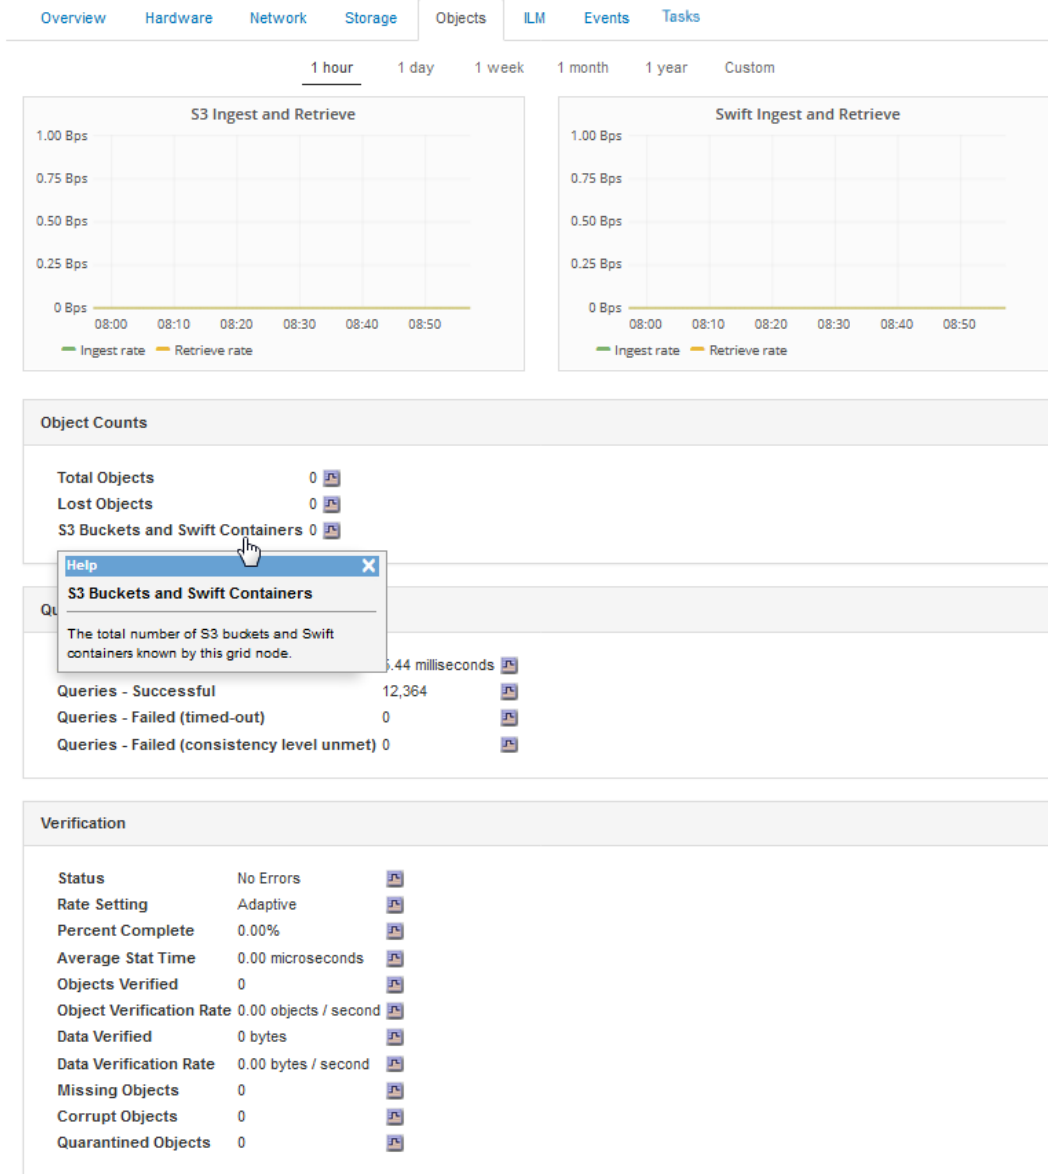
The charts show trends for all client traffic directed to load balancer endpoints within the grid. You can select a time interval in hours, days, weeks, months, or years, or you can apply a custom interval.
5. From the Nodes home page (deployment level), click the **Objects** tab.  

The chart shows ingest and retrieve rates for your entire StorageGRID system in bytes per second and total bytes. You can select a time interval in hours, days, weeks, months, or years, or you can apply a custom interval.
6. To see information for a particular Storage Node, select the node from the list on the left, and click the **Objects** tab.  

The chart shows the object ingest and retrieval rates for this Storage Node. The tab also includes metrics for object counts, queries, and verification. You can click the labels to see the definitions of these metrics.



DC1-S2 (Storage Node)



7. If you want even more detail:

- a. Select **Support > Grid Topology**.
- b. Select *site* > **Overview > Main**.

The API Operations section displays summary information for the entire grid.

- c. Select *Storage Node* > **LDR > client application > Overview > Main**

The Operations section displays summary information for the selected Storage Node.

## Accessing and reviewing audit logs

Audit messages are generated by StorageGRID services and stored in text log files. API-specific audit messages in the audit logs provide critical security, operation, and performance monitoring data that can help you evaluate the health of your system.

### Before you begin

- You must have specific access permissions.
- You must have the `Passwords.txt` file.
- You must know the IP address of an Admin Node.

### About this task

The active audit log file is named `audit.log`, and it is stored on Admin Nodes.

Once a day, the active `audit.log` file is saved, and a new `audit.log` file is started. The name of the saved file indicates when it was saved, in the format `yyyy-mm-dd.txt`.

After a day, the saved file is compressed and renamed, in the format `yyyy-mm-dd.txt.gz`, which preserves the original date.

This example shows the active `audit.log` file, the previous day's file (`2018-04-15.txt`), and the compressed file for the prior day (`2018-04-14.txt.gz`).

```
audit.log
2018-04-15.txt
2018-04-14.txt.gz
```

### Steps

1. Log in to an Admin Node:
  - a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
  - b. Enter the password listed in the `Passwords.txt` file.
2. Go to the directory containing the audit log files:
 

```
cd /var/local/audit/export
```
3. View the current or a saved audit log file, as required.

### Related information

[Understanding audit messages](#)

## Swift operations tracked in the audit logs

All successful storage DELETE, GET, HEAD, POST, and PUT operations are tracked in the StorageGRID audit log. Failures are not logged, nor are info, auth, or OPTIONS requests.

See [Understanding audit messages](#) for details about the information tracked for the following Swift operations.

### Account operations

- GET account

- HEAD account

### **Container operations**

- DELETE container
- GET container
- HEAD container
- POST container
- PUT container

### **Object operations**

- DELETE object
- GET object
- HEAD object
- PUT object

### **Related references**

[\*Account operations\*](#) on page 14

[\*Container operations\*](#) on page 16

[\*Object operations\*](#) on page 19

### **Related information**

[\*Understanding audit messages\*](#)

## Copyright

---

Copyright © 2019 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

## 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., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277