iRODS®

S3 Resource Plugin: Glacier Support

Justin James
Applications Engineer
iRODS Consortium

July 5-8, 2022
iRODS User Group Meeting 2022
Leuven, Belgium

Overview of S3 Glacier

Glacier storage classes are the archive tiers for S3.

- Defined by Amazon but other vendors may implement it.
  - We have tested it on both Amazon and FujiFilm.

- Transparent on uploads except for a flag for storage class.

- Asynchronous on downloads except for "instant retrieval" storage classes.

Glacier Support on Object Read

- Prior to reading an object, a **HeadObject** operation must be called to ascertain if the object is currently in archive.

- The x-amz-storage-class header is inspected:
  - If it exists and is either GLACIER or DEEP_ARCHIVE, inspect the x-amz-restore header:
    - If x-amz-restore has ongoing-request=true, then a restore has already been scheduled. Return REPLICA_IS_BEING_STAGED error with message indicating the object is in process of being restored.
    - If x-amz-restore has ongoing-request=false, then the object has already been restored. Proceed as normal.
    - If x-amz-restore does not exist, the object is in archive. Call **RestoreObject** and return REPLICA_IS_BEING_STAGED with message indicating the object is being queued for restoration.
  - If x-amz-storage-class header does not exist or is not GLACIER or DEEP_ARCHIVE, the object can be immediately retrieved. Proceed as normal.

Restoring an Object from Archive

To support the **RestoreObject** operation, two new resource context configuration settings have been introduced.

- S3_RESTORATION_TIER - This is the value sent in the <tier> tag when **RestoreObject** is called. Values are not case sensitive. Valid values are Standard, Bulk, and Expedited. The restoration tier in combination with the storage class defines the length of time needed to complete the restoration. The following are the restoration times for AWS.

	Glacier	**Deep Archive**
Expedited	1-5 minutes	Not Allowed
Standard (default)	3-5 hours	Within 12 hours
Bulk	5-12 hours	Within 48 hours

Note: **RestoreObject** is neither necessary nor allowed for objects stored in Glacier_IR.

- S3_RESTORATION_DAYS - The number of days the object will be restored. The default in the S3 plugin is 7. (According to Amazon, this is overridden if you have the bucket set up with lifecycle configuration.)

Glacier Support on Write or Copy

A new resource context setting has been added to define the storage class used when uploading data objects:

- S3_STORAGE_CLASS - Valid values are (not case-sensitive):
  - STANDARD - default
  - GLACIER
  - DEEP_ARCHIVE
  - GLACIER_IR - Glacier Instant Retrieval

If defined, this setting is sent in the x-amz-storage-class header for **PutObject** and **CopyObject**.

- This header may also have the following values which are either not relevant for Glacier support or have not been implemented:
  - STANDARD_IA - Standard Infrequent Access
  - ONEZONE_IA - One Zone Infrequent Access
  - INTELLIGENT_TIERING
  - OUTPOST

Changes to libs3

The libs3 library did not have support for Glacier and Deep Archive. The following changes were made to our fork of this library:

- Implemented the **RestoreObject** API
- Added the ability to set x-amz-storage-class header on **PutObject** and **CopyObject**
- Added the ability to read x-amz-storage-class and x-amz-restore headers from the **HeadObject** header.
- Will open a pull request to libs3 for these changes.

Summary of New Resource Context Setting Options

Key	Valid Values	Description
S3_STORAGE_CLASS	STANDARD GLACIER DEEP_ARCHIVE GLACIER_IR	The storage class assigned to the data object. Sent on PutObject and CopyObject.
S3_RESTORATION_TIER	STANDARD BULK EXPEDITED	When restoring an object, the restoration tier used.
S3_RESTORATION_DAYS	Any positive integer	The number of days the object will be restored

Example Glacier Setup and File Retrieval

Create a resource that places files in Glacier, performs expedited restorations , and restores for 1 day.

```
$ iadmin mkresc s3resc s3 `hostname`:/justinkylejames-irods1/amazons3resc
"S3_DEFAULT_HOSTNAME=s3.amazonaws.com;S3_AUTH_FILE=/var/lib/irods/amazon.keypair;S3_REGIONNAME=us-east-
1;S3_PROTO=HTTP;HOST_MODE=cacheless_attached;S3_STORAGE_CLASS=Glacier;S3_RESTORATION_TIER=Expedited;S3_RESTORATIO
N_DAYS=1"
Creating resource:
Name:           "s3resc"
Type:           "s3"
Host:           "ce61bbc3beec"
Path:           "/justinkylejames-irods1/amazons3resc"
Context:
"S3_DEFAULT_HOSTNAME=s3.amazonaws.com;S3_AUTH_FILE=/var/lib/irods/amazon.keypair;S3_REGIONNAME=us-east-
1;S3_PROTO=HTTP;HOST_MODE=cacheless_attached;S3_STORAGE_CLASS=Glacier;S3_RESTORATION_TIER=Expedited;S3_RESTORATIO
N_DAYS=1"
```

Create and put a file

```
$ echo test123 > test.txt
$ iput -R s3resc test.txt
```

Example Glacier Setup and File Retrieval

Try to get the object.

```
$ iget test.txt -
remote addresses: 172.17.0.2 ERROR: getUtil: get error for - status = -721000 REPLICA_IS_BEING_STAGED
Level 0: [-]    /github/irods_resource_plugin_s3/s3/s3_transport/src/s3_transport.cpp:208:irods::error
irods::experimental::io::s3_transport::restore_s3_object(const std::string &, libs3_types::bucket_context &,
const unsigned int, const std::string &, const std::string &) :  status [REPLICA_IS_BEING_STAGED]  errno [] --
message [Object is in GLACIER and has been queued for restoration.  Try again later.]
```

Now try to get it again. Note the slightly different error message.

```
$ iget test.txt -
remote addresses: 172.17.0.2 ERROR: getUtil: get error for - status = -721000 REPLICA_IS_BEING_STAGED
Level 0: [-]    /github/irods_resource_plugin_s3/s3/s3_transport/src/s3_transport.cpp:133:irods::error
irods::experimental::io::s3_transport::handle_glacier_status(const std::string &, libs3_types::bucket_context &,
const unsigned int, const std::string &, irods::experimental::io::s3_transport::object_s3_status, const
std::string &) :  status [REPLICA_IS_BEING_STAGED]  errno [] -- message [Object is in GLACIER and is currently
being restored.  Try again later.]
```

Wait a few minutes and retrieve the file again:

```
$ iget test.txt -

test123
```

Release Status

Glacier support has been added to the S3 plugin and has been partially released.

- Restoration from Glacier was added in 4.2.11.0 and included in 4.3.0.0.

- To be released soon in 4.3.0.1:
  - Support for setting the storage class on **PutObject** and **CopyObject**
  - Support for Deep Archive for puts/gets/copies

What's Next:

Release support for Deep Archive and setting the storage class on **PutObject** and **CopyObject**.

Support for the intelligent tiering storage class should be trivial but has not been implemented or tested.

We could write a rule to read metadata on an atomic put and select the storage class dynamically for object level control.