

Hands-on tutorials

Getting started using the Amazon Glacier storage classes



Getting started using the Amazon Glacier storage classes: Hands-on tutorials

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Getting started using the Amazon Glacier storage classes	i
Overview	1
What you'll accomplish	2
Implementation	2
Congratulations!	17

Getting started using the Amazon Glacier storage classes

AWS experience	Beginner
Time to complete	20 minutes
Cost to complete	Less than \$1 (Amazon S3 pricing page)
Services used	Amazon S3

Overview

The [Amazon Glacier storage classes](#) are purpose-built for data archiving, providing you with the highest performance, most retrieval flexibility, and the lowest cost archive storage in the cloud. To keep costs low yet suitable for varying retrieval needs, these storage classes support flexible retrieval options from milliseconds to several hours. The purpose of this tutorial is to show you how easy it is to begin storing your archive datasets in the Amazon Glacier storage classes.

You can choose from three archive storage classes optimized for different access patterns and storage duration. For archive data that needs immediate access, choose the [Amazon Glacier Instant Retrieval](#) storage class, an archive storage class that delivers the lowest cost storage with milliseconds retrieval. For archive data that does not require immediate access but needs the flexibility to retrieve large sets of data at no cost, choose Amazon Glacier Flexible Retrieval (formerly Amazon Glacier), with retrieval in minutes or free bulk retrievals in 5-12 hours. To save even more on long-lived archive storage, choose Amazon S3 Glacier Deep Archive, the lowest cost storage in the cloud with data retrieval within twelve hours.

By archiving on AWS you'll have access to very low cost cloud storage, you'll be able to digitally preserve and retain your data for the long term, and you'll be able to leverage comprehensive security and compliance capabilities. The Amazon Glacier storage classes are used by customers for their long-term enterprise archive data, media archives, backup data, and data lake archives.

Use the [S3 console](#) and S3 API to easily archive your data in [Amazon S3](#). The S3 console and S3 API allow you to access all the features and functionality that the Amazon S3 service provides. Follow this tutorial to begin using the S3 console to store your archive datasets in the Amazon Glacier storage classes.

What you'll accomplish

- Create an Amazon S3 bucket
- Upload objects to the Amazon Glacier storage classes
- Restore your objects stored in the Amazon Glacier Flexible Retrieval or S3 Glacier Deep Archive storage classes

Implementation

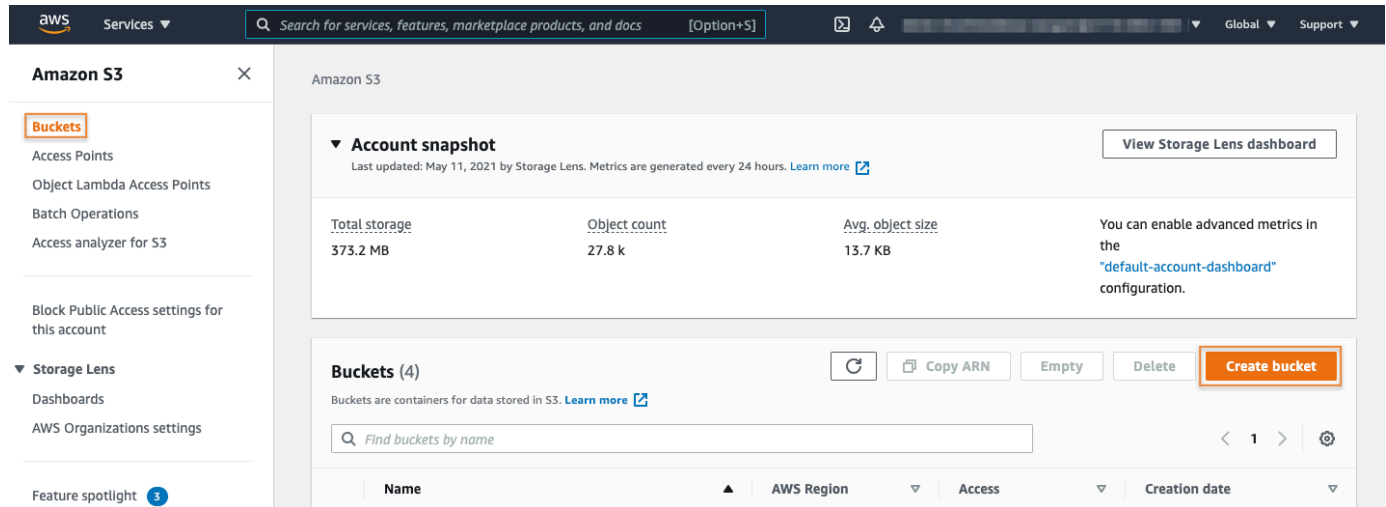
Step 1: Getting started using the Amazon Glacier storage classes

1. Sign into the Amazon S3 console
 - If you have not already done so, create an AWS account. [Access this support page for more information on how to create and activate a new AWS account.](#)
 - Log into the [AWS Management Console](#) using your account information.
 - From the AWS console services search bar, enter '**S3**'. Under the services search results section, select **S3**. You may notice an option for Amazon Glacier. This option is for the Glacier service prior to integration with Amazon S3. We recommend all new Amazon Glacier users use the S3 console.



2. Create an S3 bucket

Choose **Buckets** from the S3 menu on the left rail and then select the **Create bucket** button.



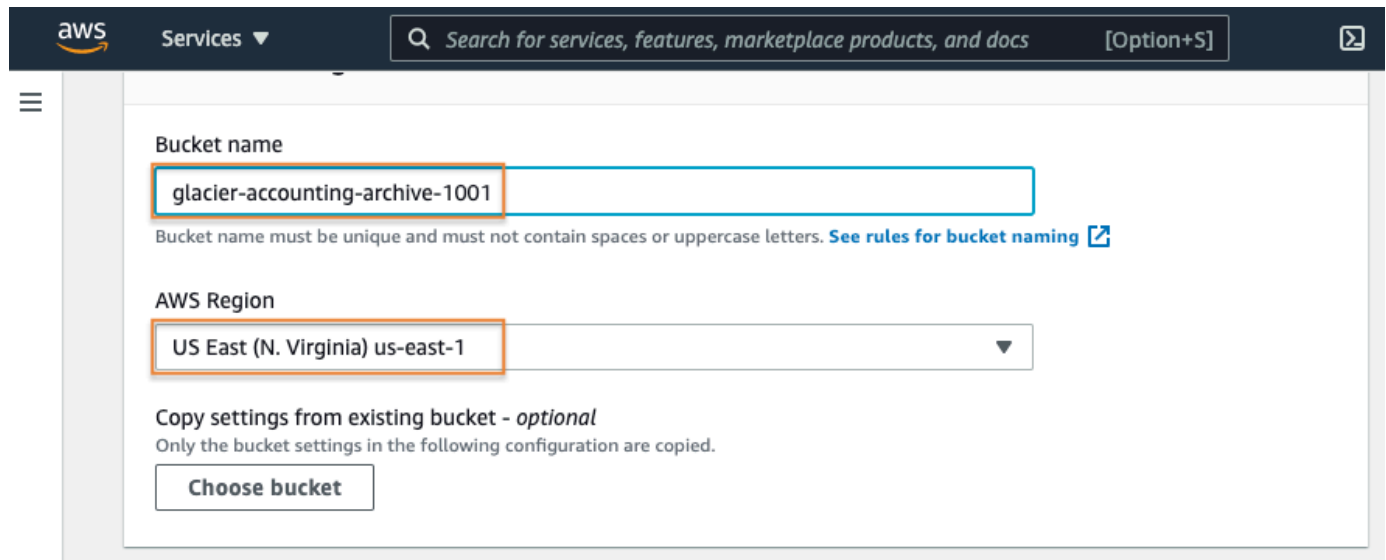
The screenshot shows the Amazon S3 console interface. On the left, the 'Amazon S3' menu is open, with 'Buckets' highlighted. The main content area displays an 'Account snapshot' with the following data:

Total storage	Object count	Avg. object size
373.2 MB	27.8 k	13.7 KB

Below the snapshot, there is a 'Buckets (4)' section with a search bar and a table of buckets. The 'Create bucket' button is highlighted in orange.

3. Configure the bucket

- Enter a descriptive globally unique name for your bucket.
- Select which AWS Region you would like your bucket created in.
- The default Block Public Access setting is appropriate for this workload, so leave this section as is.



The screenshot shows the 'Create bucket' configuration form in the Amazon S3 console. The 'Bucket name' field is filled with 'glacier-accounting-archive-1001' and is highlighted with an orange box. The 'AWS Region' dropdown menu is set to 'US East (N. Virginia) us-east-1' and is also highlighted with an orange box. Below the form, there is a 'Choose bucket' button.

4. Enable versioning

Next, enable bucket versioning to protect your data from accidental or malicious user deletes or overwrites.

[Read more about bucket versioning here](#). Then, add some tags to help track costs associated with our archive data over time.

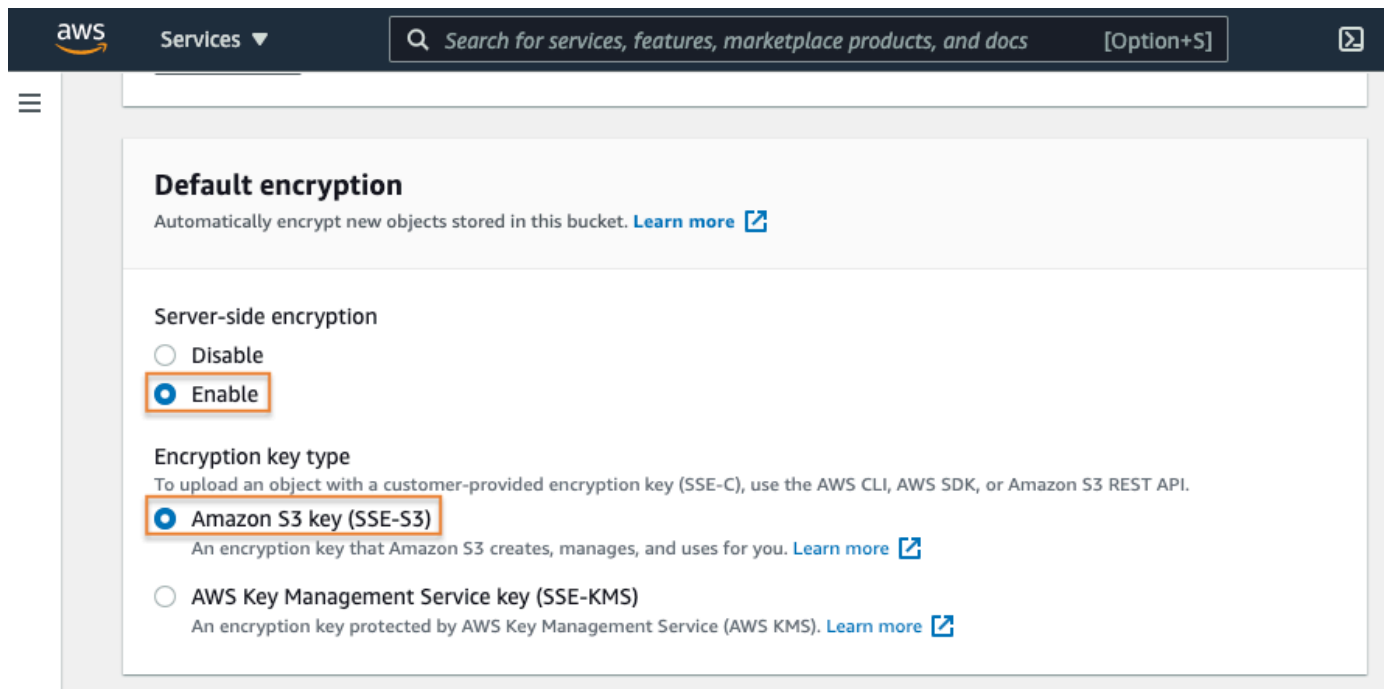
[You can find more information about S3 bucket cost allocation tagging here](#).

The screenshot shows the AWS console interface for configuring an S3 bucket. At the top, there is a navigation bar with the AWS logo, a 'Services' dropdown, a search bar with the placeholder text 'Search for services, features, marketplace products, and docs', and a '[Option+S]' button. Below the navigation bar, there is a hamburger menu icon on the left. The main content area is divided into two sections. The first section is titled 'Bucket Versioning' and contains a description: 'Versioning is a means of keeping multiple variants of an object in the same bucket. You can use versioning to preserve, retrieve, and restore every version of every object stored in your Amazon S3 bucket. With versioning, you can easily recover from both unintended user actions and application failures. [Learn more](#)'. Below the description, there are two radio buttons: 'Disable' and 'Enable'. The 'Enable' radio button is selected and highlighted with an orange border. The second section is titled 'Tags (2) - optional' and contains a description: 'Track storage cost or other criteria by tagging your bucket. [Learn more](#)'. Below the description, there is a table with two columns: 'Key' and 'Value - optional'. The first row has 'dept' in the 'Key' column, 'accounting' in the 'Value - optional' column, and a 'Remove' button. The second row has 'archive' in the 'Key' column, 'true' in the 'Value - optional' column, and a 'Remove' button. At the bottom of the table, there is an 'Add tag' button.

5. Enable default encryption

Next, you have the option of enabling default 'at-rest' encryption for the bucket. The settings here will apply to any objects uploaded to the bucket where you have not defined at-rest encryption details during the upload process.

For this example, enable server-side encryption leveraging S3 service managed keys (SSE-S3). If your workload requirements are not satisfied by SSE-S3, you can also leverage AWS Key Management Service (KMS). [More information about Amazon S3 and AWS KMS can be found here](#).



6. Enable S3 Object Lock

Now you have the option to enable [S3 Object Lock](#) in the **Advanced settings** section. With S3 Object Lock, you can store objects using a write-once-read-many (WORM) model. S3 Object Lock can help prevent objects from being deleted or overwritten for a fixed amount of time, or indefinitely.

S3 Object Lock can be used to help meet regulatory requirements that require WORM storage, or to simply add another layer of protection against object changes and deletion.

For this workload, it is appropriate to enable S3 Object Lock to ensure important archived data is not deleted prematurely by unauthorized users.

- Choose **Enable**.
- Select the check box to acknowledge enabling the S3 Object Lock settings
- Select the **Create bucket** button.

aws Services [Option+S]

An encryption key protected by AWS Key Management Service (AWS KMS). [Learn more](#)

▼ Advanced settings

Object Lock

Store objects using a write-once-read-many (WORM) model to help you prevent objects from being deleted or overwritten for a fixed amount of time or indefinitely. [Learn more](#)

Disable

Enable

Permanently allows objects in this bucket to be locked. Additional Object Lock configuration is required in bucket details after bucket creation to protect objects in this bucket from being deleted or overwritten.

i Object Lock works only in versioned buckets. Enabling Object Lock automatically enables Bucket Versioning.

⚠ **Enabling Object Lock will permanently allow objects in this bucket to be locked**
Enable Object Lock only if you need to prevent objects from being deleted to have data integrity and regulatory compliance. After you enable this feature, anyone with the appropriate permissions can put immutable objects in the bucket. You might be blocked from deleting the objects and the bucket. Additional Object Lock configuration is required in bucket details after bucket creation to protect objects in this bucket from being deleted or overwritten. [Learn more](#)

I acknowledge that enabling Object Lock will permanently allow objects in this bucket to be locked.

i After creating the bucket you can upload files and folders to the bucket, and configure additional bucket settings.

Cancel **Create bucket**

7. Configure S3 Object Lock


Next, the S3 console will present a banner indicating the bucket creation was successful. The S3 console will also present a prompt informing you that additional configuration is needed to enable the S3 Object Lock feature.

Select the **bucket details** link presented in the prompt. Making this selection will open the **Properties** tab for your newly created bucket.

Note

For this exercise, use **Governance** mode for the S3 Object Lock configuration. This will allow you to permanently delete your test object using an admin user after this tutorial has completed.

For more information about S3 Object Lock, read the blog featuring "[Protecting data with Amazon S3 Object Lock.](#)"

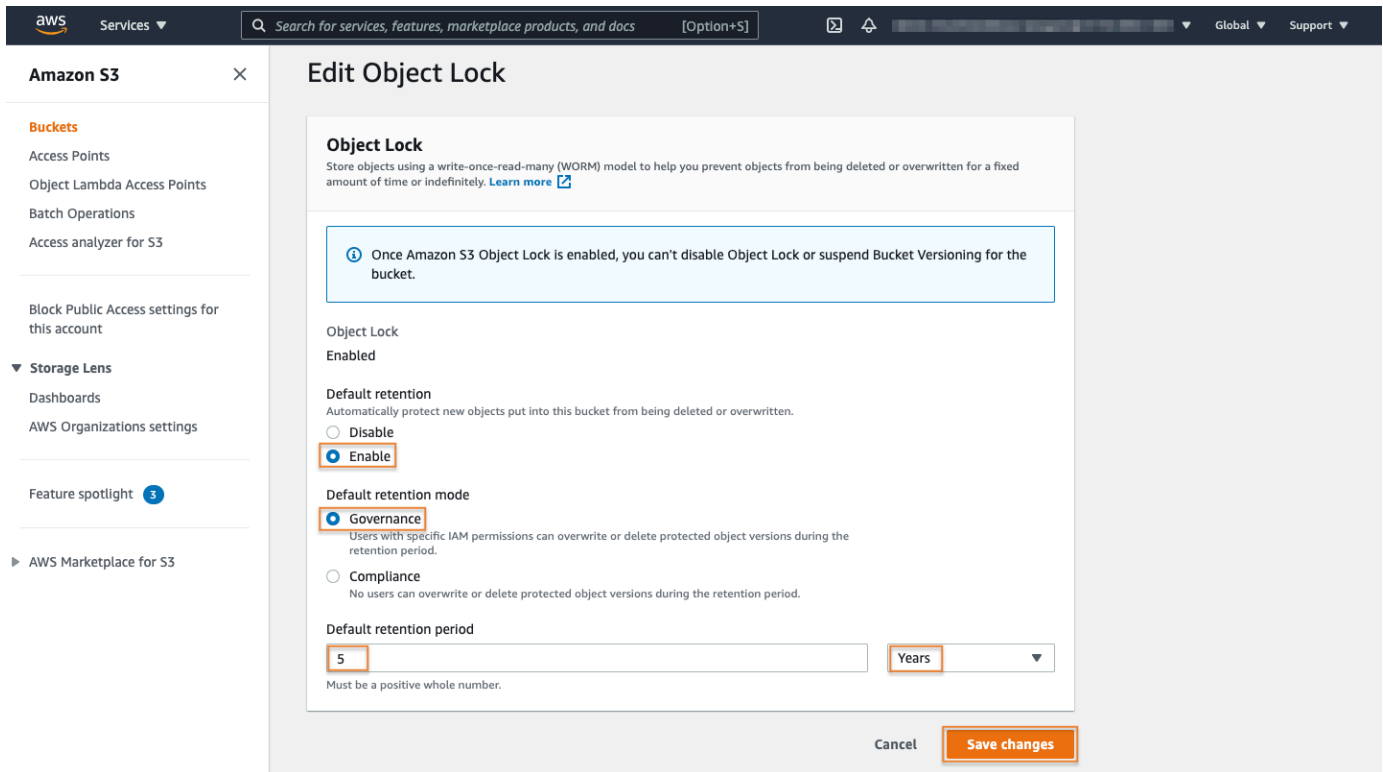


8. Edit the S3 Object Lock

On the bucket **Properties** tab, navigate to the **Object Lock** section and select the **Edit** button. Here you can set your default values for objects uploaded to your bucket.

For this example, you want to enable retention for all objects uploaded to this bucket for 5 years. Select **Enable** for the **Default retention** option, choose governance mode by selecting the **Governance** option under **Default retention mode** and enter **5** as the default retention period.

Lastly, select **Years** for the unit of measure and then select the **Save changes** button.



Step 2: Uploading data to an Amazon S3 bucket

Now that your bucket has been created and configured, you are ready to upload archive data to the Amazon Glacier storage classes.

1. Select the bucket

If you have logged out of your AWS Management Console session, log back in.

Navigate to the [S3 console](#) and select the **Buckets** menu option.

From the list of available buckets, select the bucket name of the bucket you just created.

The screenshot shows the Amazon S3 console interface. On the left, the navigation pane is open to the 'Buckets' section. The main content area displays the 'Account snapshot' and a table of buckets. The bucket 'glacier-accounting-archive-100101' is highlighted with an orange box.

Total storage	Object count	Avg. object size
373.9 MB	28.6 k	13.4 KB

Name	AWS Region
cloudtrail-awslogs-411020521031-jvuc6dah-isengard-do-not-delete	US East (N. Virginia) us-east-1
do-not-delete-gatedgarden-audit-411020521031	US West (Oregon) us-west-2
eab-testdata-prod1a	US East (N. Virginia) us-east-1
glacier-accounting-archive-100101	US East (N. Virginia) us-east-1
prod-datalake-8123	US East (N. Virginia) us-east-1

2. Start the upload

Next, choose the **Objects** tab. Then from within the **Objects** section, select the **Upload** button.

The screenshot shows the Amazon S3 console interface for the bucket 'glacier-accounting-archive-100101'. The 'Objects' tab is selected, and the 'Upload' button is highlighted with an orange box.

Objects (0)

Objects are the fundamental entities stored in Amazon S3. You can use [Amazon S3 inventory](#) to get a list of all objects in your bucket. For others to access your objects, you'll need to explicitly grant them permissions. [Learn more](#)

Copy URL Open Download Delete Actions Create folder **Upload**

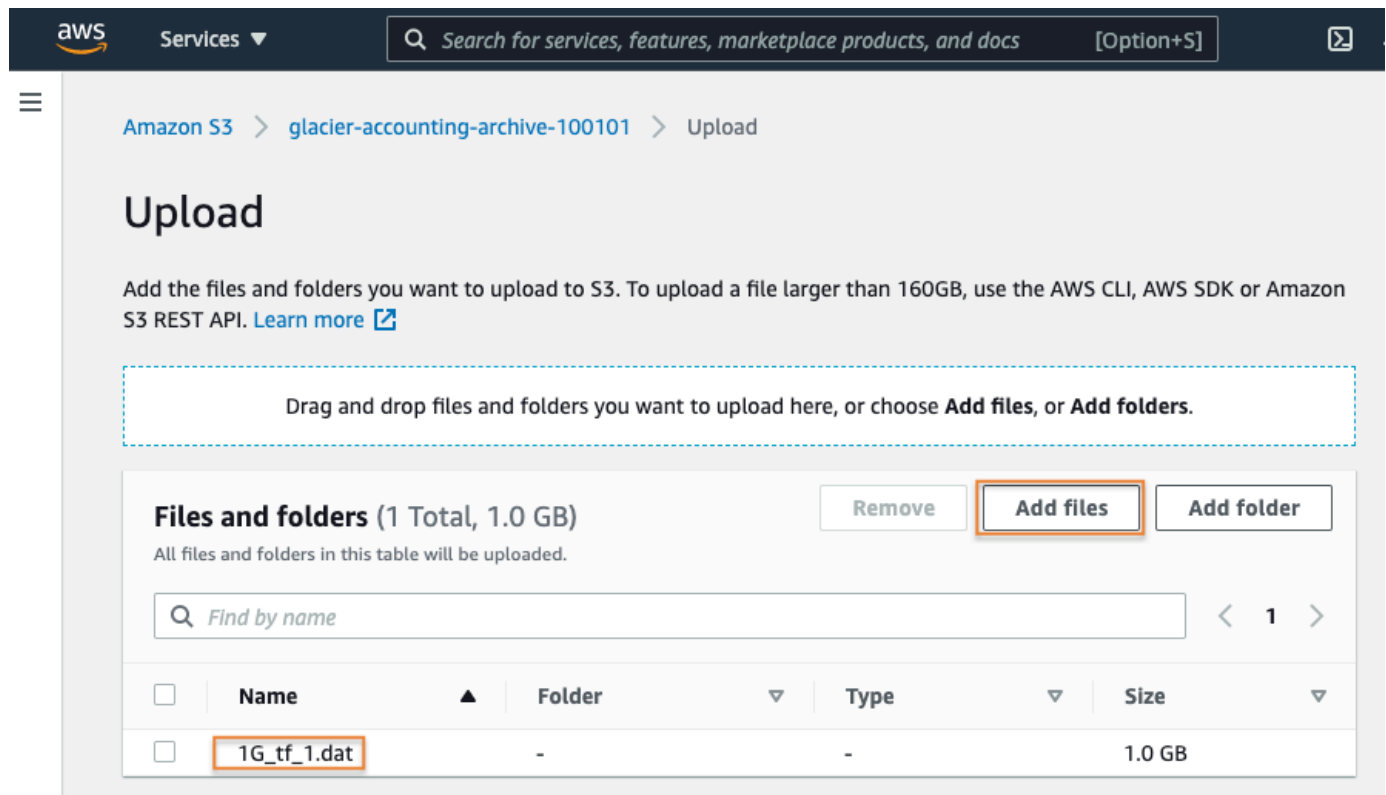
3. Select the file to upload

Then, select the **Add files** button.

Navigate your local file system to locate the archive file you would like to upload.

Select the appropriate file and then select **Open**.

Your file will be listed in the **Files and folders** section.



The screenshot shows the AWS S3 console's 'Upload' page for a bucket named 'glacier-accounting-archive-100101'. The page title is 'Upload'. Below the title, there is a dashed box with the text: 'Drag and drop files and folders you want to upload here, or choose **Add files**, or **Add folders**.' Below this, there is a section titled 'Files and folders (1 Total, 1.0 GB)' with buttons for 'Remove', 'Add files', and 'Add folder'. A search bar labeled 'Find by name' is present. Below the search bar is a table with the following columns: Name, Folder, Type, and Size. The table contains one row with the file name '1G_tf_1.dat', a folder of '-', a type of '-', and a size of '1.0 GB'. The 'Add files' button and the file name in the table are highlighted with orange boxes.

<input type="checkbox"/>	Name	Folder	Type	Size
<input type="checkbox"/>	1G_tf_1.dat	-	-	1.0 GB

4. Select the storage class

In the **Properties** section, select the S3 storage class you would like to upload your archive to.

Select **Glacier Deep Archive**, as the example dataset needs to be retained for 5 years and there is a low probability the data will be accessed often.

Note

If your workload requires milliseconds access and single API call access to your archived data, the Amazon Glacier Instant Retrieval storage class should be selected here instead. More information about the Amazon Glacier storage class options can be viewed [here](#).

Leave the rest of the options on the default settings and select the **Upload** button.

Note

Objects stored in many S3 storage classes have minimum object durations associated with them. In this case, uploading the test file to Glacier Deep Archive will result in 180 days of billing even if it is deleted early. Storing 1 GB in S3 Glacier Deep Archive for 180 days with the retrieval is ~\$0.03. [You can read more about S3 pricing here.](#)


Storage class

Amazon S3 offers a range of storage classes designed for different use cases. [Learn more](#) or see [Amazon S3 pricing](#)

	Storage class	Designed for	Availability Zones	Min storage duration	
<input type="radio"/>	Standard	Frequently accessed data (more than once a month) with milliseconds access	≥ 3	-	-
<input type="radio"/>	Intelligent-Tiering	Data with changing or unknown access patterns	≥ 3	-	-
<input type="radio"/>	Standard-IA	Infrequently accessed data (once a month) with milliseconds access	≥ 3	30 days	1
<input type="radio"/>	One Zone-IA	Recreatable, infrequently accessed data (once a month) stored in a single Availability Zone with milliseconds access	1	30 days	1
<input type="radio"/>	Glacier Instant Retrieval	Long-lived archive data accessed once a quarter with instant retrieval in milliseconds	≥ 3	90 days	1
<input type="radio"/>	Glacier Flexible Retrieval (formerly Glacier)	Long-lived archive data accessed once a year with retrieval of minutes to hours	≥ 3	90 days	-
<input checked="" type="radio"/>	Glacier Deep Archive	Long-lived archive data accessed less than once a year with retrieval of hours	≥ 3	180 days	-
<input type="radio"/>	Reduced redundancy	Noncritical, frequently accessed data with milliseconds access (not recommended as S3 Standard is more cost effective)	≥ 3	-	-

5. Review the status

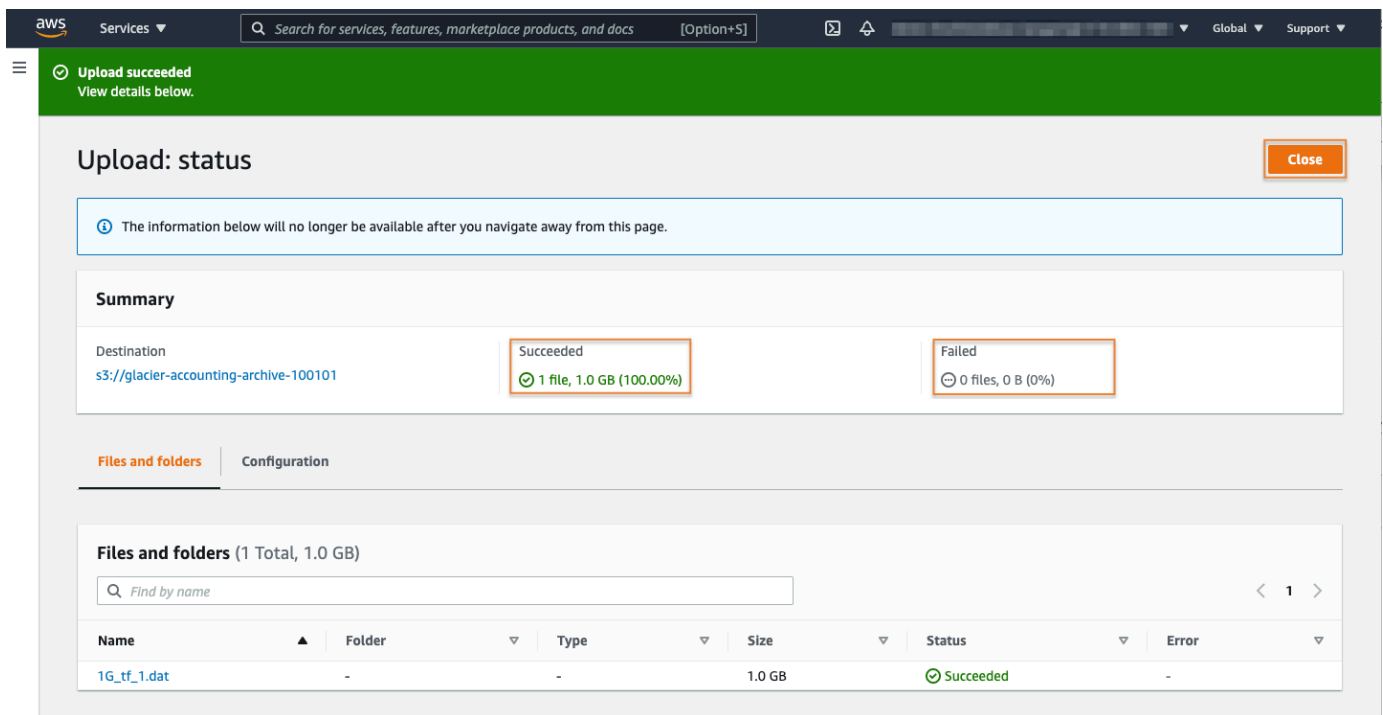
A banner will be displayed providing you with details of the file's upload status.



6. Confirm the upload completed

After your file upload operations have completed, you will be presented with a summary of the operations indicating if it has completed successfully or if it has failed.

In this case, the file has uploaded successfully. Select the **Close** button.



Step 3: Restore your data

Now that you have successfully uploaded your data to S3 Glacier Deep Archive, let's go over the process of restoring your data.

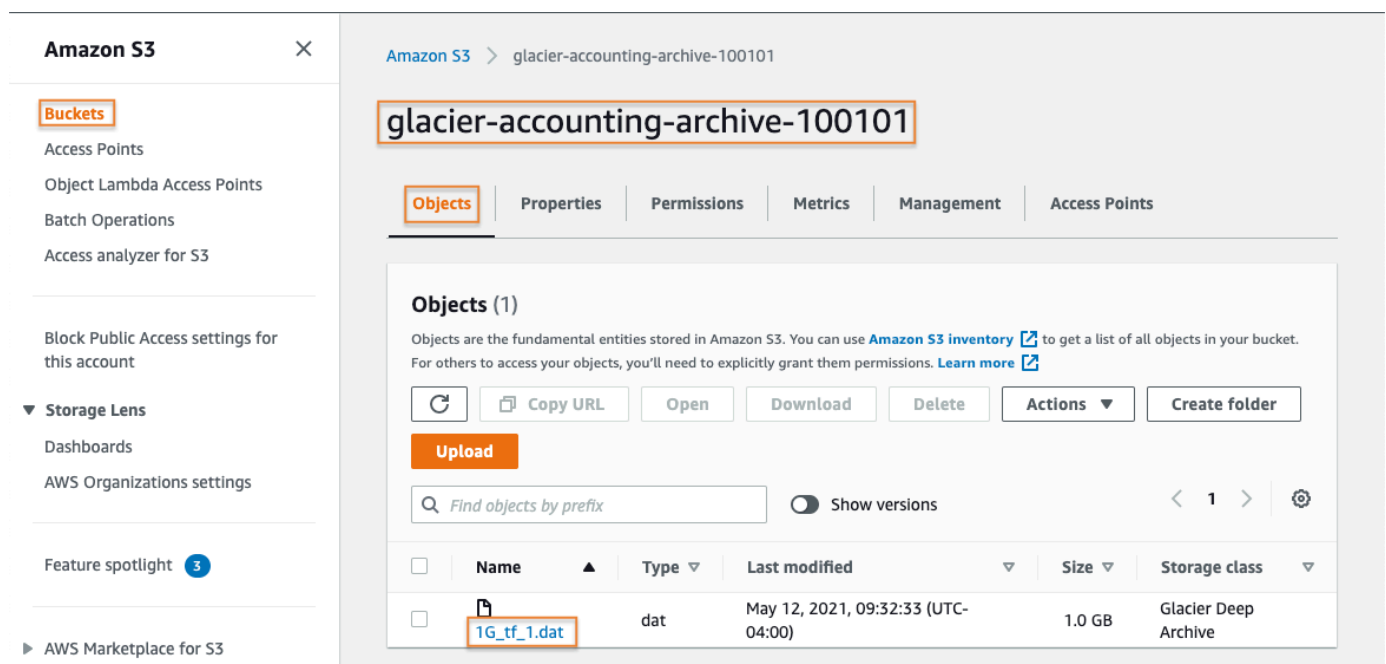
Note

the process of restoring your data before it can be accessed, is required for data that is stored in the Amazon Glacier Flexible Retrieval and S3 Glacier Deep Archive storage classes. Data stored in the Amazon Glacier Instant Retrieval storage class does not require this restore request prior to being accessed. You can learn more about Amazon Glacier Instant Retrieval [here](#).

1. Select the object to restore

If you have logged out of your AWS Management Console session, log back in.

- Navigate to the [S3 console](#) and select the **Buckets** menu option.
- From the list of available buckets, select the bucket name of the bucket you have created for this exercise.
- From the **Objects** menu, select the name of the test file you just uploaded.



The screenshot shows the AWS S3 console interface. On the left, the navigation menu is visible with 'Buckets' selected. The main content area shows the bucket 'glacier-accounting-archive-100101'. The 'Objects' tab is active, displaying a table with one object: '1G_tf_1.dat'. The object's details are as follows:

Name	Type	Last modified	Size	Storage class
1G_tf_1.dat	dat	May 12, 2021, 09:32:33 (UTC-04:00)	1.0 GB	Glacier Deep Archive

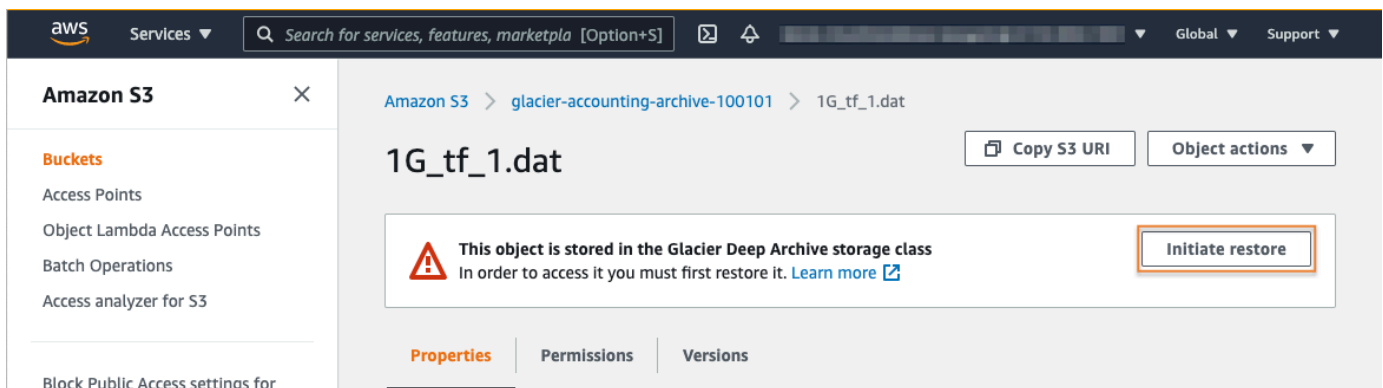
2. Initiate the restore

After selecting your test file's name, you will be presented with a banner indicating that your object is stored in the S3 Glacier Deep Archive storage class and that you need to restore it if you would like to access your data.

You can initiate the restore process by simply selecting the **Initiate restore button** attached to the information banner, or you can choose **Initiate restore** from the **Object actions** menu.

Note

The restore process will create a copy of your archived data and will store that copy in the S3 Standard storage class. During the restore initiation process you will set the number of days that you wish to have your data available. During this time period, you will incur applicable storage charges for your data in both the archive storage class as well as in the active storage class.



3. Configure the restore

From the **Initiate restore** page, you will define the number of days you desire to make your restored copy available.

Next, you will have a choice between standard or bulk retrieval. Data stored in the Amazon Glacier Flexible Retrieval storage class will additionally have an option to select expedited retrieval. [More information about restore options can be found here.](#)

For this exercise, choose the **Standard retrieval** option. Then, select the **Initiate restore** button to continue.

The screenshot shows the 'Initiate restore' page in the Amazon S3 console. The left sidebar contains navigation options for 'Amazon S3', including Buckets, Access Points, Object Lambda Access Points, Batch Operations, Access analyzer for S3, Block Public Access settings, Storage Lens, Dashboards, AWS Organizations settings, Feature spotlight, and AWS Marketplace for S3. The main content area is titled 'Initiate restore' and includes a breadcrumb trail: Amazon S3 > glacier-accounting-archive-100101 > Initiate restore. Below the title is a description: 'To restore objects you must first initiate a restore request, and then wait until the objects are available. Retrieval fees apply. Learn more or see pricing'. The 'Restore objects from Glacier Deep Archive' section explains that temporary copies will be available for a specified number of days. A form field for 'Number of days that the restored copy is available' has the value '5' entered. A message states: 'The restored copy will be available until approximately 05-31-2021.' Under 'Retrieval tier', 'Standard retrieval' (Typically within 12 hours) is selected. The 'Specified objects' section has a search bar and a table with one object:

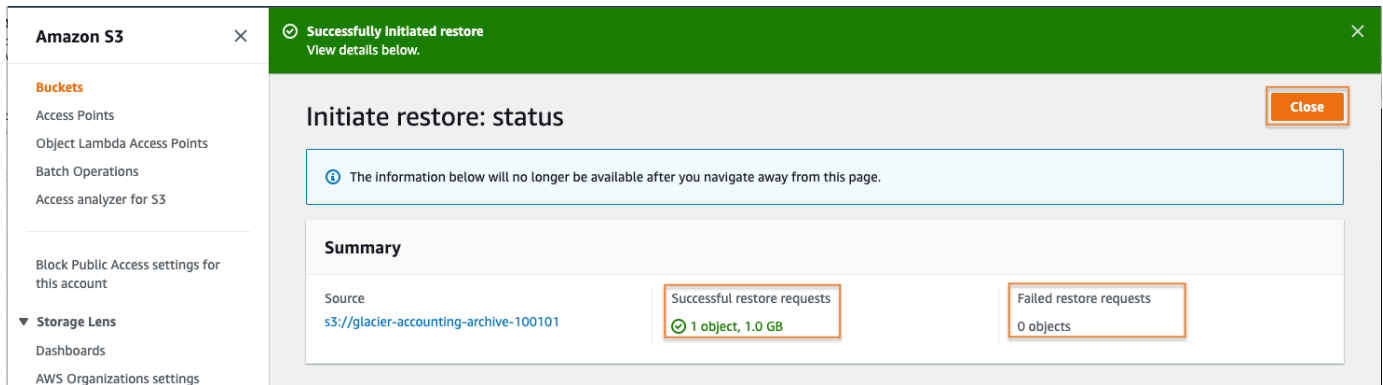
Name	Version ID	Type	Last modified	Size	Storage class	Intelligent-Tiering Access
1G_tf_1.dat	-	dat	May 12, 2021, 09:32:33 (UTC-04:00)	1.0 GB	Glacier Deep Archive	-

At the bottom right, there are 'Cancel' and 'Initiate restore' buttons.

4. Wait for the restore to complete

A summary page will be displayed indicating if the restore request was successful or if any errors occurred. In this case, the restore request was successful. Select the **Close** button to continue.

For this standard restore from S3 Glacier Deep Archive, you will need to wait about 12 hours for the temporary object to be restored to the Amazon S3 Standard-IA storage class. S3 Event notifications support alerting when an object restore event has completed. [More information about S3 Event notifications can be found in the Amazon S3 documentation here.](#)




5. Verify restore has completed

Now you can verify that your object has been restored. After waiting about twelve hours for the restore operation to complete, log back into your [S3 console](#).

Select **Buckets** from the left rail menu, and select the name of your bucket to view its contents. From the **Objects** section, select the file name of the object you have attempted to restore to see its current status.

Here you can see that the object's **Restore status** is listed as **Completed**. The **Restoration expiry date**, which is based on the number of days we defined in the restore process, is listed as well. You have successfully restored your archived object. This object will be available until the time specified in the **Restoration expiry date** section.

You can now perform actions like run S3 select queries against this file, copy the object to another bucket in your account or to another account, or download the data to your local machine.



Clean up resources

In the following steps, you clean up the resources you created in this tutorial. It is a best practice to delete resources that you are no longer using so that you do not incur unintended charges.

1. Delete your test object
 - a. If you have logged out of your AWS Management Console session, log back in.
 - b. Navigate to the [S3 console](#) and select the **Buckets** menu option.
 - c. First you will need to delete the test object from your test bucket. Select the **name** of the bucket you have been working with for this tutorial.
 - d. Put a check mark in the checkbox to the left of your test object name, then select the **Delete** button.
 - e. On the **Delete objects** page, verify that you have selected the proper object to delete and type **permanently delete** into the **Permanently delete objects** confirmation box.
 - f. Then, select the **Delete object** button to continue. Next, you will be presented with a banner indicating if the deletion has been successful.
2. Delete your test bucket
 - a. Finally, you need to delete the test bucket you have created. Return to the list of buckets in your account.
 - b. Select the radio button to the left of the bucket you created for this tutorial, and then select the **Delete** button.
 - c. Review the warning message. If you desire to continue deletion of this bucket, type the bucket name into the **Delete bucket** confirmation box and select **Delete bucket**.

Congratulations!

You have learned how to create an Amazon S3 bucket, upload objects to the Amazon Glacier and S3 Glacier Deep Archive storage classes, and how to restore your objects so that they can be easily retrieved.