

Hands-on tutorials

Store and Retrieve a File with Amazon S3



Store and Retrieve a File with Amazon S3: 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

Store and Retrieve a File with Amazon S3	i
Overview	1
Implementation	1
Congratulations!	14

Store and Retrieve a File with Amazon S3

Cost to complete	Free Tier AWS Free Tier includes 5GB storage, 20,000 Get Requests, and 2,000 Put Requests with Amazon S3. View AWS Free Tier Details »
Services used	Amazon S3
Requires	Storing Your Files with AWS Requires an Account Create a free account in minutes
Last updated	June 1, 2022

Overview

This step-by-step how-to guide will help you store your files in the cloud using Amazon Simple Storage Service (Amazon S3). Amazon S3 is a service that enables you to store your data (referred to as **objects**) at massive scale. In this guide, you will create an Amazon S3 bucket (a container for data stored in Amazon S3), upload a file, retrieve the file, and delete the file.

The resources you create in this guide are [AWS Free Tier](#) eligible.

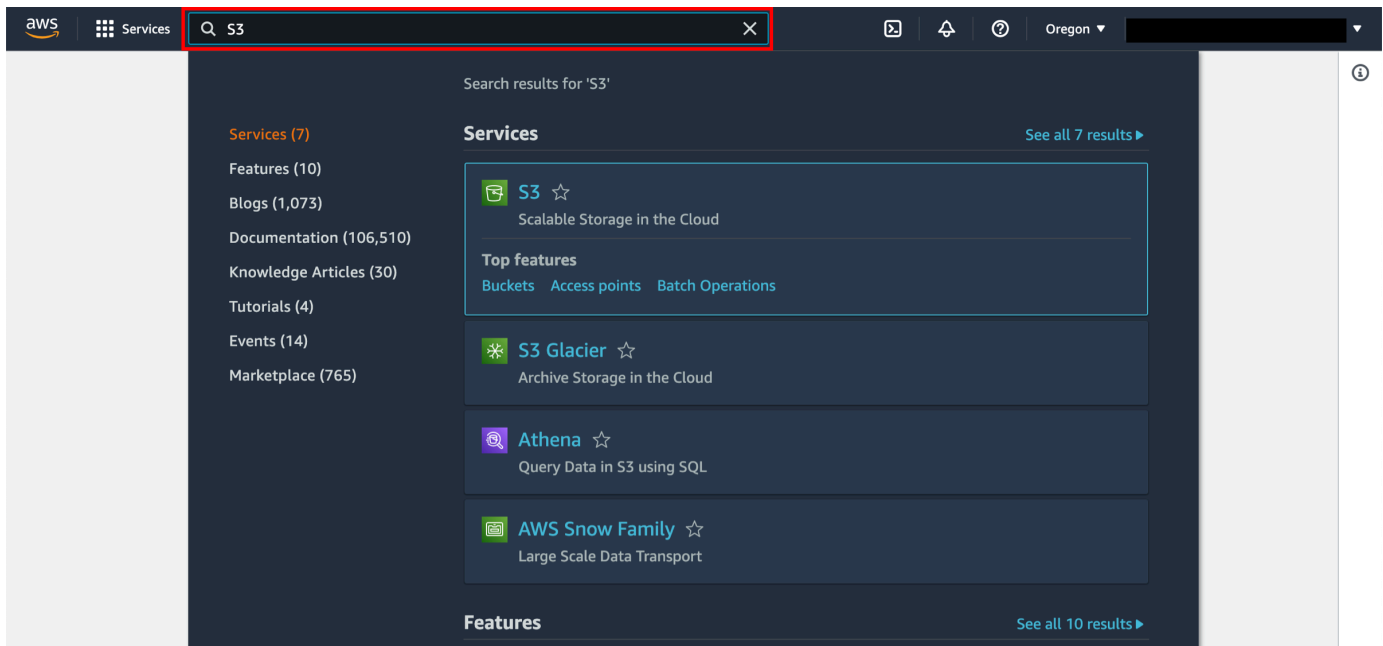
Implementation

Step 1: Upload a file

In this step, you will upload a file to your new Amazon S3 bucket.

1. Open the Amazon S3 console

[Click on](#) the AWS Management Console home to open the console in a new browser window, so you can keep this step-by-step guide open. When the screen loads, enter your user name and password to get started. Then type **S3** in the search bar and select S3 to open the console.

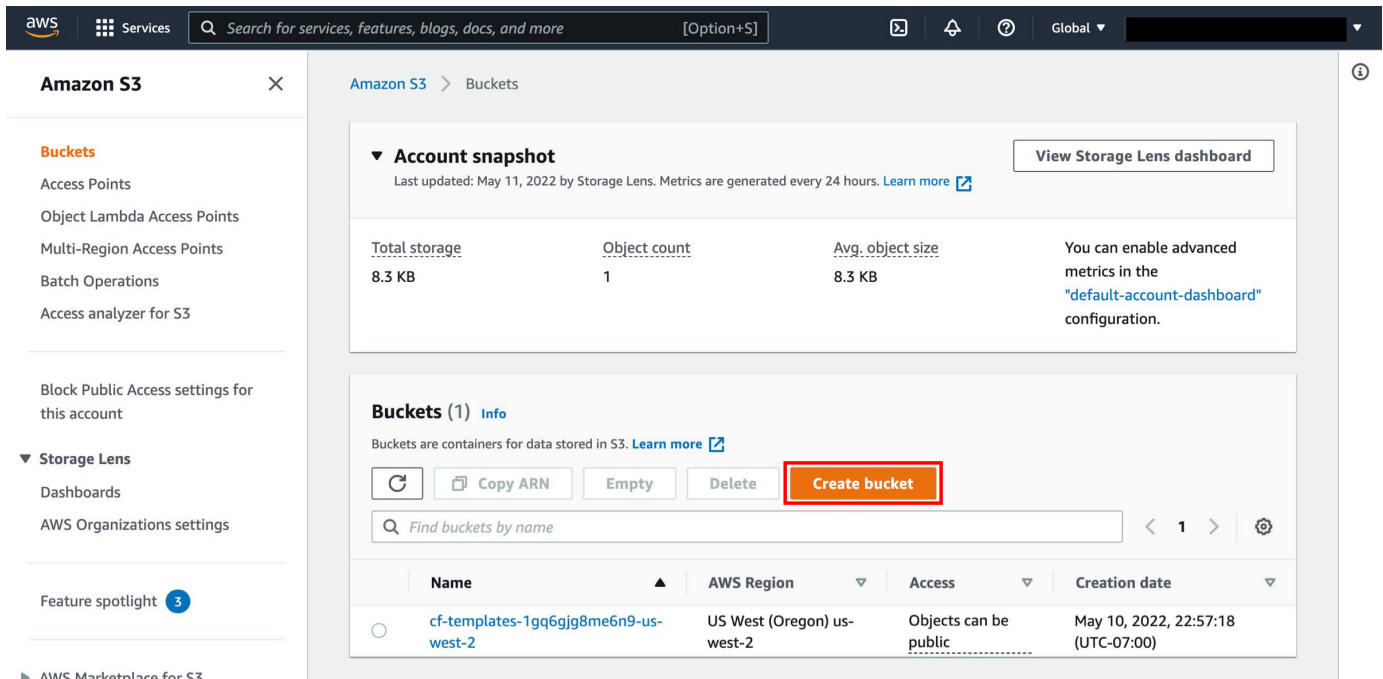


2. Create S3 bucket

In the S3 dashboard, click **Create Bucket**.

If this is the first time you have created a bucket, you will see a screen that looks like the image pictured here.

If you have already created S3 buckets, your S3 dashboard will list all the buckets you have created.

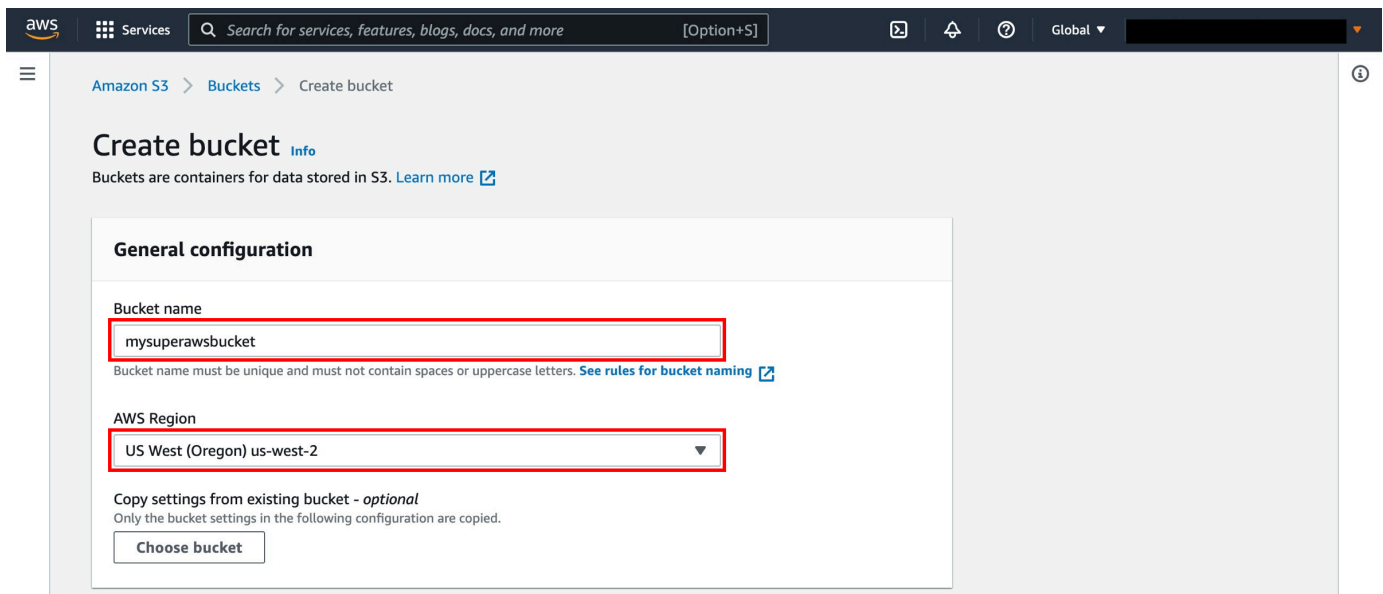


The screenshot shows the Amazon S3 Buckets console. On the left is a navigation sidebar with options like Buckets, Access Points, and Storage Lens. The main content area displays an 'Account snapshot' with metrics: Total storage (8.3 KB), Object count (1), and Avg. object size (8.3 KB). Below this is a 'Buckets (1)' section with a 'Create bucket' button highlighted in red. A table lists one bucket: 'cf-templates-1gq6g98me6n9-us-west-2' in the 'US West (Oregon) us-west-2' region with public access.

Name	AWS Region	Access	Creation date
cf-templates-1gq6g98me6n9-us-west-2	US West (Oregon) us-west-2	Objects can be public	May 10, 2022, 22:57:18 (UTC-07:00)

3. Enter bucket name

Enter a bucket name. Bucket names must be unique across all existing bucket names in Amazon S3. For this guide, we have used **mysuperawsbucket**, but you should choose a name that is relevant and unique to you. There are a number of other [restrictions on S3 bucket names](#) as well. Once you've selected a name, select a Region to create your bucket in.



The screenshot shows the 'Create bucket' configuration page. The 'General configuration' section has a 'Bucket name' field containing 'mysuperawsbucket' and an 'AWS Region' dropdown menu set to 'US West (Oregon) us-west-2'. A 'Choose bucket' button is visible at the bottom.

4. Configure permissions

You have the ability to set permission settings for your S3 bucket. Leave the default values and select **Next**.

Object Ownership Info

Control ownership of objects written to this bucket from other AWS accounts and the use of access control lists (ACLs). Object ownership determines who can specify access to objects.

ACLs disabled (recommended)
All objects in this bucket are owned by this account. Access to this bucket and its objects is specified using only policies.

ACLs enabled
Objects in this bucket can be owned by other AWS accounts. Access to this bucket and its objects can be specified using ACLs.

Object Ownership
Bucket owner enforced

Block Public Access settings for this bucket

Public access is granted to buckets and objects through access control lists (ACLs), bucket policies, access point policies, or all. In order to ensure that public access to this bucket and its objects is blocked, turn on Block all public access. These settings apply only to this bucket and its access points. AWS recommends that you turn on Block all public access, but before applying any of these settings, ensure that your applications will work correctly without public access. If you require some level of public access to this bucket or objects within, you can customize the individual settings below to suit your specific storage use cases. [Learn more](#)

Block all public access
Turning this setting on is the same as turning on all four settings below. Each of the following settings are independent of one another.

- Block public access to buckets and objects granted through *new* access control lists (ACLs)**
S3 will block public access permissions applied to newly added buckets or objects, and prevent the creation of new public access ACLs for existing buckets and objects. This setting doesn't change any existing permissions that allow public access to S3 resources using ACLs.
- Block public access to buckets and objects granted through *any* access control lists (ACLs)**
S3 will ignore all ACLs that grant public access to buckets and objects.
- Block public access to buckets and objects granted through *new* public bucket or access point policies**
S3 will block new bucket and access point policies that grant public access to buckets and objects. This setting doesn't change any existing policies that allow public access to S3 resources.
- Block public and cross-account access to buckets and objects through *any* public bucket or access point policies**
S3 will ignore public and cross-account access for buckets or access points with policies that grant public access to buckets and objects.

5. Review and create

You have many useful options for your S3 bucket including [Versioning](#), [Tags](#), [Default Encryption](#), and [Object Lock](#). We won't enable them for this tutorial.

Select **Create bucket**.

Bucket Versioning
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](#)

Bucket Versioning
 Disable
 Enable

Tags (0) - optional
Track storage cost or other criteria by tagging your bucket. [Learn more](#)

No tags associated with this bucket.

[Add tag](#)

Default encryption
Automatically encrypt new objects stored in this bucket. [Learn more](#)

Server-side encryption
 Disable
 Enable

▼ 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.

[?](#) Object Lock works only in versioned buckets. Enabling Object Lock automatically enables Bucket Versioning.

[?](#) After creating the bucket you can upload files and folders to the bucket, and configure additional bucket settings.

Cancel [Create bucket](#)

Step 2: Create an S3 bucket

In this step, you will create an Amazon S3 bucket. A bucket is the container you store your files in.

1. Open your bucket

You will see your new bucket in the S3 console. Click on your bucket's name to navigate to the bucket.

The screenshot shows the AWS Management Console interface. At the top, there is a navigation bar with the AWS logo, 'Services', a search bar, and a global region selector. A green notification banner at the top right states 'Successfully created bucket "mysuperawsbucket"'. Below this, the 'Buckets' page is displayed, showing an 'Account snapshot' and a table of buckets. The 'mysuperawsbucket' entry is highlighted with a red box.

Name	AWS Region	Access	Creation date
cf-templates-1gq6gjj8me6n9-us-west-2	US West (Oregon) us-west-2	Objects can be public	May 10, 2022, 22:57:18 (UTC-07:00)
mysuperawsbucket	US West (Oregon) us-west-2	Bucket and objects not public	May 12, 2022, 17:34:04 (UTC-07:00)

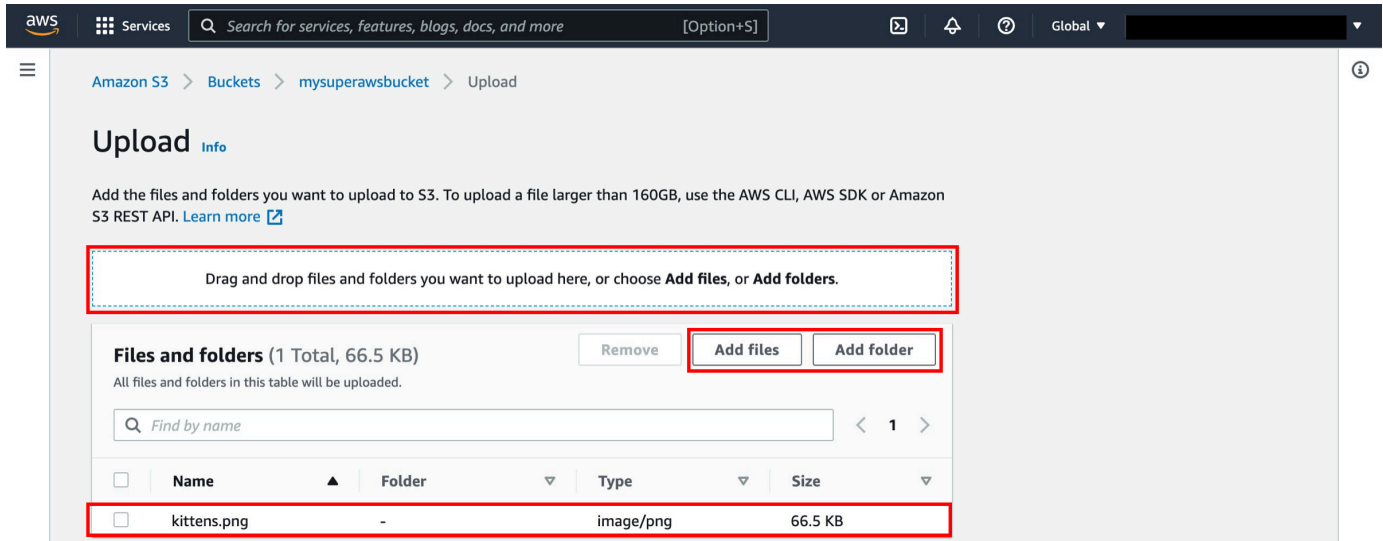
2. Choose Upload

You are in your bucket's home page. Select **Upload**.

The screenshot shows the AWS Management Console interface for the 'mysuperawsbucket' home page. The 'Objects' tab is selected, and the 'Upload' button is highlighted with a red box. Below the table, a message states 'No objects' and another 'Upload' button is also highlighted with a red box.

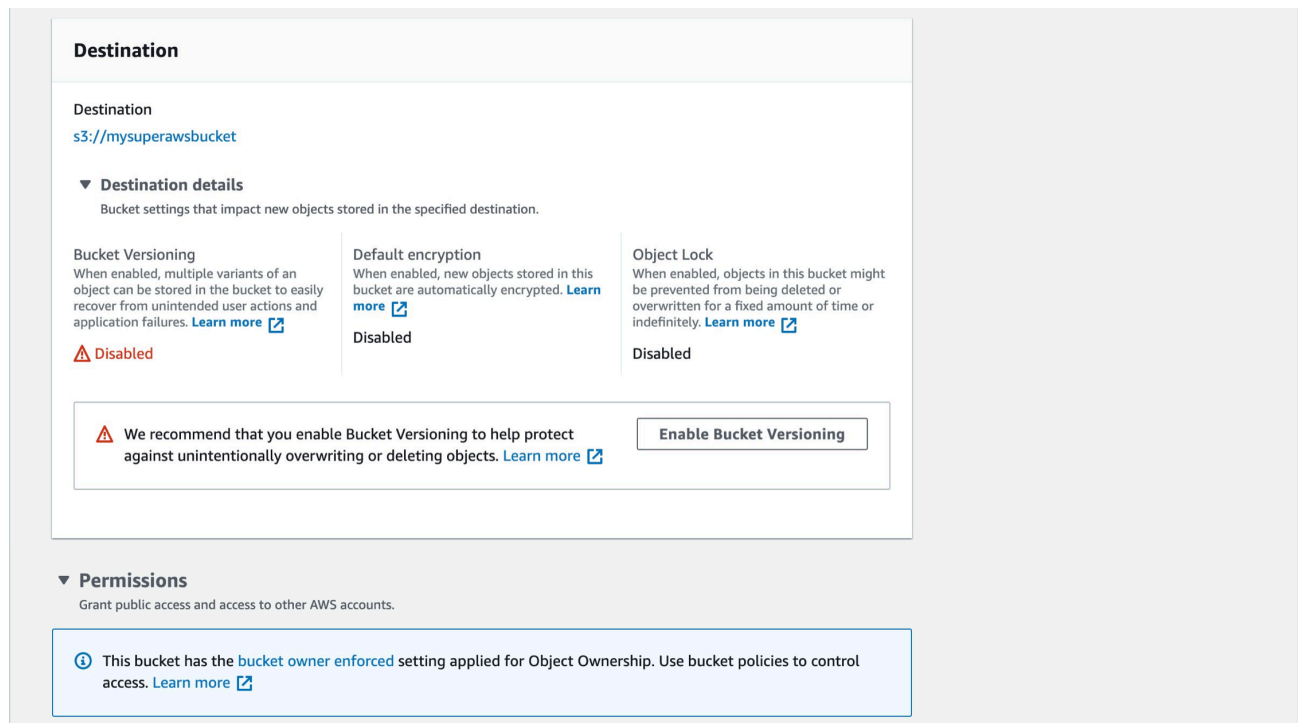
3. Add files

To select a file to upload, either click **Add files** or **Add folder** and select sample file(s) that you would like to store or **Drag and Drop** a file on the upload box. Your file(s) will be displayed after you have selected file(s) to upload.



4. Set permissions

You have the ability to review destination details and permissions. For this tutorial, leave the default values.



5. Configure properties

You have the ability to set property settings like storage class, server-side encryption, additional checksums, tags, and metadata with your object. Leave the default values and select **Upload**.

▼ **Properties**

Specify storage class, encryption settings, tags, and more.

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 checked="" 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 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	-	-

Server-side encryption settings

Server-side encryption protects data at rest. [Learn more](#)

Server-side encryption

- Do not specify an encryption key
- Specify an encryption key

⚠ If your bucket policy requires encrypted uploads, you must specify an encryption key or your upload will fail.

i Since default encryption is disabled for this bucket, no encryption settings will be applied to the objects when storing them in Amazon S3.

Additional checksums

Checksum functions are used for additional data integrity verification of new objects. [Learn more](#)

Additional checksums

- Off**
Amazon S3 will use a combination of MD5 checksums and ETags to verify data integrity.
- On**
Specify a checksum function for additional data integrity validation.

Tags - optional

Track storage cost or other criteria by tagging your objects. [Learn more](#)

No tags associated with this resource.

Metadata - optional

Metadata is optional information provided as a name-value (key-value) pair. [Learn more](#)

6. Confirm upload

You will see your object in your bucket's home screen.

Upload succeeded
View details below.

Upload: status

The information below will no longer be available after you navigate away from this page.

Summary

Destination s3://mysuperawsbucket	Succeeded ✔ 1 file, 66.5 KB (100.00%)	Failed ✘ 0 files, 0 B (0%)
--------------------------------------	--	-------------------------------

Files and folders | Configuration

Files and folders (1 Total, 66.5 KB)

Find by name

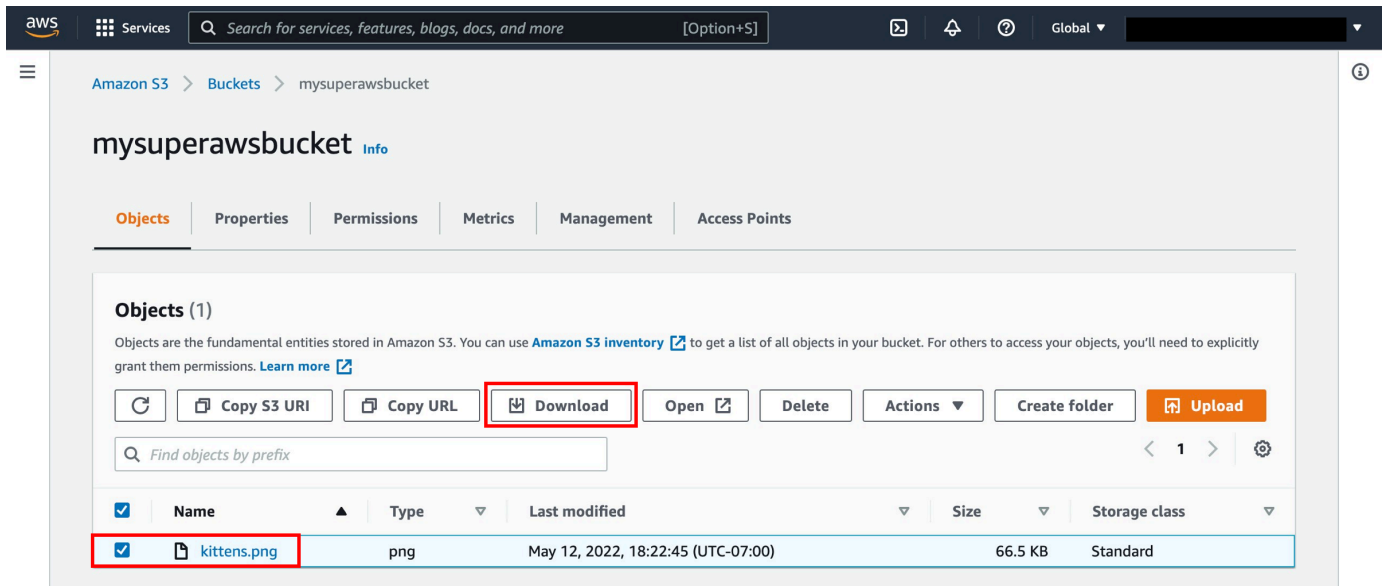
Name	Folder	Type	Size	Status	Error
kittens.png	-	image/png	66.5 KB	✔ Succeeded	-

Step 3: Retrieve the object

In this step, you will download the file from your Amazon S3 bucket.

- Download the object

Select the checkbox next to the file you would like to download, then select **Download**.



Amazon S3 > Buckets > mysuperawsbucket

mysuperawsbucket [Info](#)

Objects | Properties | Permissions | Metrics | Management | Access Points

Objects (1)

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](#)

< 1 >

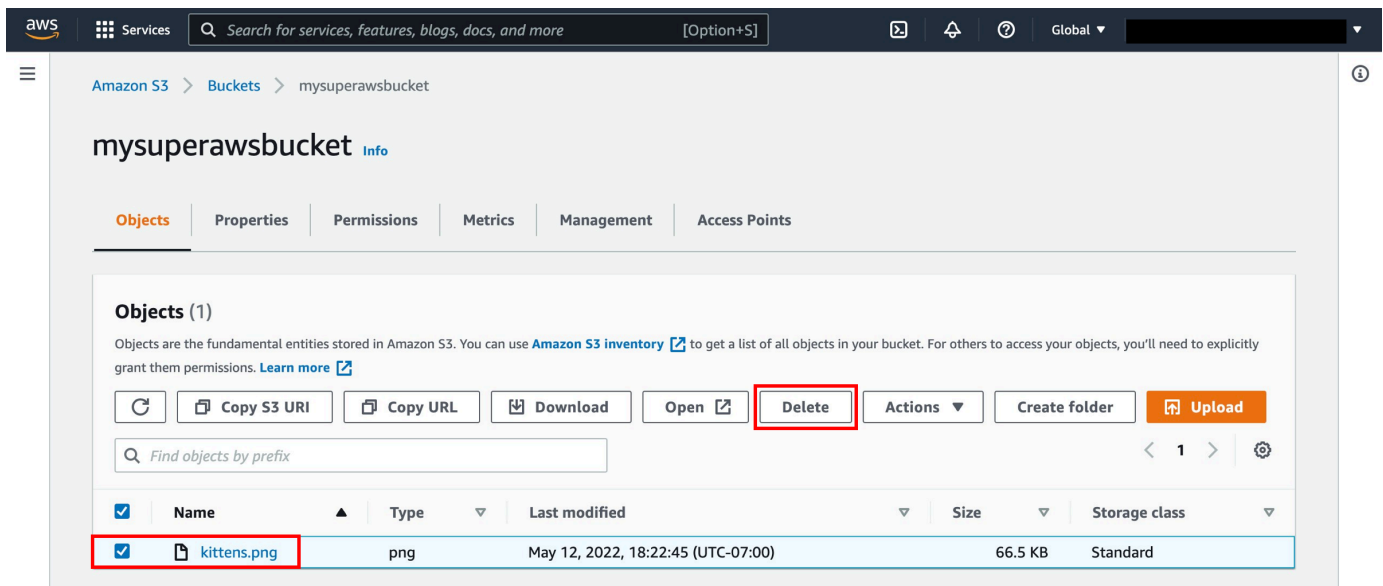
<input checked="" type="checkbox"/>	Name	Type	Last modified	Size	Storage class
<input checked="" type="checkbox"/>	kittens.png	png	May 12, 2022, 18:22:45 (UTC-07:00)	66.5 KB	Standard

(Optional) Clean up resources

You can easily delete your object and bucket from the Amazon S3 console. In fact, it is a best practice to delete resources you are no longer using so you don't keep getting charged for them.

1. Delete the object

You will first delete your object. Select the checkbox next to the file you want to delete and select **Delete**.



Amazon S3 > Buckets > mysuperawsbucket

mysuperawsbucket [Info](#)

Objects | Properties | Permissions | Metrics | Management | Access Points

Objects (1)

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](#)

< 1 >

<input checked="" type="checkbox"/>	Name	Type	Last modified	Size	Storage class
<input checked="" type="checkbox"/>	kittens.png	png	May 12, 2022, 18:22:45 (UTC-07:00)	66.5 KB	Standard

2. Confirm deletion

Review and enter **permanently delete** in the text input field to confirm deletion. Click **Delete objects**.

Amazon S3 > Buckets > mysuperawsbucket > Delete objects

Delete objects Info

⚠ Deleting the specified objects can't be undone. [Learn more](#)

Specified objects

Find objects by name < 1 >

Name	Type	Last modified	Size
kittens.png	png	May 12, 2022, 18:22:45 (UTC-07:00)	66.5 KB

Permanently delete objects?

To confirm deletion, type *permanently delete* in the text input field.

Cancel **Delete objects**

3. Navigate to your bucket

Click on Amazon S3 > Buckets to view all your buckets in the region.

Amazon S3 > **Buckets**

Account snapshot View Storage Lens dashboard

Last updated: May 11, 2022 by Storage Lens. Metrics are generated every 24 hours. [Learn more](#)

Total storage	Object count	Avg. object size	You can enable advanced metrics in the "default-account-dashboard" configuration.
8.3 KB	1	8.3 KB	

Buckets (2) Info

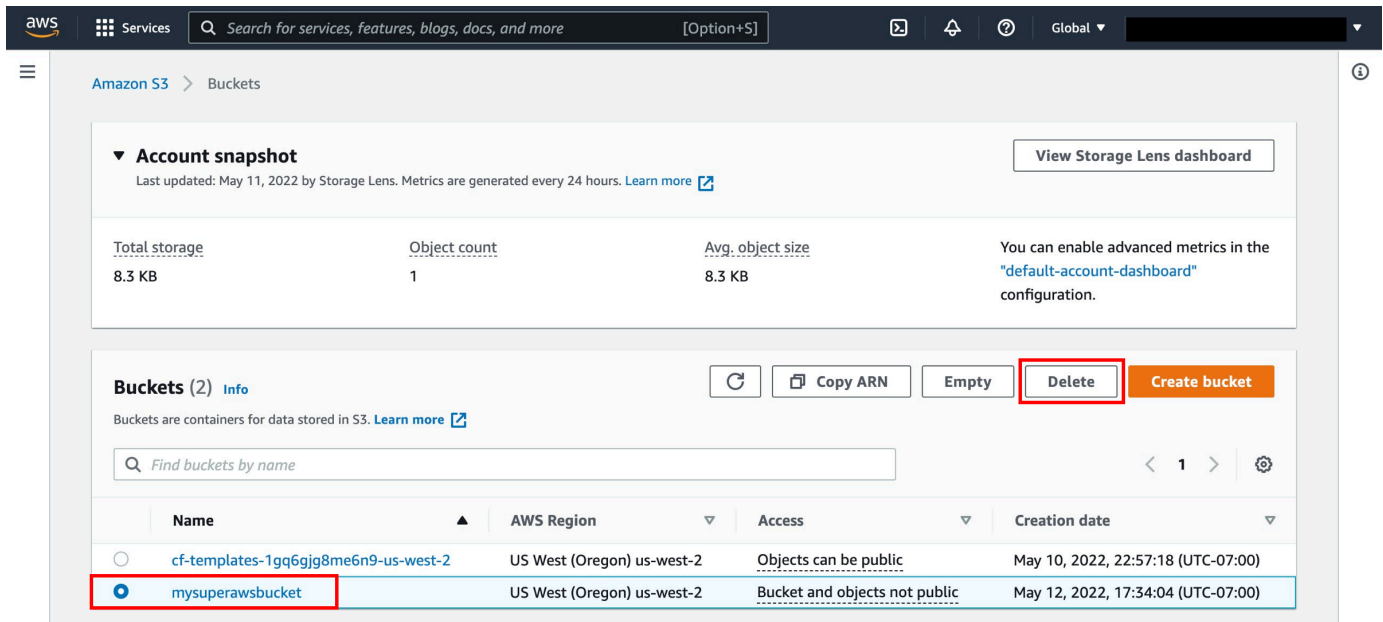
Buckets are containers for data stored in S3. [Learn more](#)

Find buckets by name < 1 >

Name	AWS Region	Access	Creation date
<input type="radio"/> cf-templates-1gq6g9g8me6n9-us-west-2	US West (Oregon) us-west-2	Objects can be public	May 10, 2022, 22:57:18 (UTC-07:00)
<input type="radio"/> mysuperawsbucket	US West (Oregon) us-west-2	Bucket and objects not public	May 12, 2022, 17:34:04 (UTC-07:00)

4. Delete the bucket

Select the radio button to the left of the bucket you created, then choose **Delete**.

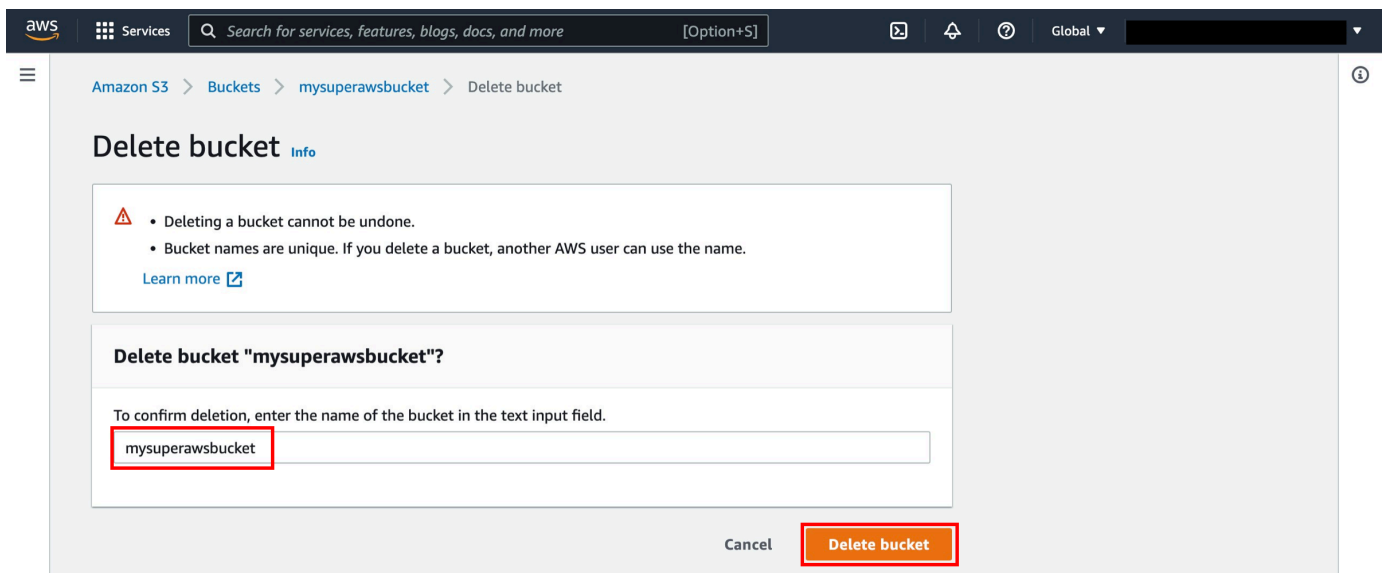


The screenshot shows the AWS Management Console interface for Amazon S3 Buckets. At the top, there's a navigation bar with the AWS logo, 'Services' menu, a search bar, and a global region dropdown. Below the navigation bar, the breadcrumb path is 'Amazon S3 > Buckets'. The main content area is divided into sections. The first section is 'Account snapshot', which includes a 'View Storage Lens dashboard' button and a table with metrics: Total storage (8.3 KB), Object count (1), and Avg. object size (8.3 KB). The second section is 'Buckets (2)', which includes a search bar and a table of buckets. The table has columns for Name, AWS Region, Access, and Creation date. The bucket 'mysuperawsbucket' is selected, and its 'Delete' button is highlighted with a red box.

Name	AWS Region	Access	Creation date
<input type="radio"/> cf-templates-1gq6jg8me6n9-us-west-2	US West (Oregon) us-west-2	Objects can be public	May 10, 2022, 22:57:18 (UTC-07:00)
<input checked="" type="radio"/> mysuperawsbucket	US West (Oregon) us-west-2	Bucket and objects not public	May 12, 2022, 17:34:04 (UTC-07:00)

5. Confirm deletion

To confirm deletion, enter the name of the bucket in the text input field and choose **Delete bucket**.



The screenshot shows the 'Delete bucket' confirmation dialog in the AWS Management Console. The breadcrumb path is 'Amazon S3 > Buckets > mysuperawsbucket > Delete bucket'. The main heading is 'Delete bucket'. Below the heading, there's a warning message: 'Deleting a bucket cannot be undone. Bucket names are unique. If you delete a bucket, another AWS user can use the name.' Below the warning, there's a section titled 'Delete bucket "mysuperawsbucket"?' which contains a text input field with the bucket name 'mysuperawsbucket' entered. At the bottom right, there are two buttons: 'Cancel' and 'Delete bucket', with the latter highlighted by a red box.

Congratulations!

You have backed up your first file to the cloud by creating an Amazon S3 bucket and uploading your file as an S3 object. Amazon S3 is designed for 99.999999999% durability to help ensure that your data is always available when you want it. You've also learned how to retrieve your backed up file and how to delete the file and bucket.