
Amazon Simple Storage Service

Console User Guide

Amazon Simple Storage Service: Console User Guide

Copyright © 2017 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

Welcome to the New Amazon S3 Console User Guide	1
Creating and Configuring a Bucket	2
Creating a Bucket	3
Deleting a Bucket	7
Emptying a Bucket	8
Viewing Bucket Properties	9
Enabling or Disabling Versioning	11
Configuring Static Website Hosting	12
Redirecting Website Requests	15
Advanced Settings	16
Enabling Cross-Region Replication	17
Disabling Cross-Region Replication	20
Setting Up a Destination for Event Notifications	21
Enabling and Configuring Event Notifications	23
Uploading, Downloading, and Managing Objects	28
Uploading Objects	29
More Info	35
Downloading Objects	35
Related Topics	39
Deleting Objects	39
Related Topics	41
Undeleting Objects	41
More Info	42
Deleting Folders	43
Related Topics	44
Viewing an Overview of an Object	45
More Info	46
Viewing Object Versions	46
More Info	48
Viewing Object Properties	48
Adding Tags to an Object	50
Storage Management	53
Creating a Lifecycle Policy	53
Configuring Storage Class Analysis	59
Configuring Storage Inventory	63
Configuring Request Metrics	66
Configuring a Request Metrics Filter	68
Setting Permissions	72
Setting Bucket Permissions	72
Setting Object Permissions	73

Welcome to the New Amazon S3 Console User Guide

*If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.*



Announcement: Object Tagging and new Storage Management features available in new console

[Opt In](#) to try object tagging and storage management.

This is the *Amazon Simple Storage Service Console User Guide* for the new Amazon S3 console.

The Amazon S3 console is one of the interfaces that you can use to work with Amazon S3. The console enables you to perform Amazon S3 tasks without writing any code.

Topics

- [Creating and Configuring an S3 Bucket \(p. 2\)](#)
- [Uploading, Downloading, and Managing Objects \(p. 28\)](#)
- [Storage Management \(p. 53\)](#)
- [Setting Permissions \(p. 72\)](#)

Creating and Configuring an S3 Bucket

*If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.*



Announcement: Object Tagging and new Storage Management features available in new console

[Opt In](#) to try object tagging and storage management.

Amazon S3 is cloud storage for the Internet. To upload your data (photos, videos, documents etc.), you first create a bucket in one of the AWS Regions. You can then upload your data objects to the bucket.

Every object you store in Amazon S3 resides in a bucket. You can use buckets to group related objects in the same way that you use a directory to group files in a file system.

Amazon S3 creates buckets in the AWS Region that you specify. You can choose any AWS Region that is geographically close to you to optimize latency, minimize costs, or address regulatory requirements. For example, if you reside in Europe, you might find it advantageous to create buckets in the EU (Ireland) or EU (Frankfurt) regions. For a list of Amazon S3 AWS Regions, see [Regions and Endpoints](#) in the *Amazon Web Services General Reference*.

You are not charged for creating a bucket. You are only charged for storing objects in the bucket and for transferring objects out of the bucket. For more information about pricing, see [Amazon Simple Storage Service \(S3\) FAQs](#).

Amazon S3 bucket names are globally unique, regardless of the AWS Region in which you create the bucket. You specify the name at the time you create the bucket. For bucket naming guidelines, see [Bucket Restrictions and Limitations](#) in the *Amazon Simple Storage Service Developer Guide*.

The following topics explain how to use the Amazon S3 console to create, delete, and manage buckets.

Topics

- [How Do I Create an S3 Bucket? \(p. 3\)](#)
- [How Do I Delete an S3 Bucket? \(p. 7\)](#)
- [How Do I Empty an S3 Bucket? \(p. 8\)](#)
- [How Do I View the Properties for an S3 Bucket? \(p. 9\)](#)
- [How Do I Enable or Suspend Versioning for an S3 Bucket? \(p. 11\)](#)
- [How Do I Configure an S3 Bucket for Static Website Hosting? \(p. 12\)](#)
- [How Do I Redirect Requests to an S3 Bucket Hosted Website to Another Host? \(p. 15\)](#)
- [Advanced Settings for S3 Bucket Properties \(p. 16\)](#)

How Do I Create an S3 Bucket?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.




Before you can upload data to Amazon Simple Storage Service, you must create a bucket in one of the AWS Regions to store your data in. After you create a bucket, you can upload an unlimited number of data objects to the bucket.

Buckets have configuration properties, including their geographical region, who has access to the objects in the bucket, and other metadata.

To create an S3 bucket

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Choose **Create bucket**.



3. On the **Name and region** page, type a name for your bucket and choose the AWS Region where you want the bucket to reside. Complete the fields on this page as follows:
 - a. For **Bucket name**, type a unique DNS-compliant name for your new bucket. Follow these naming guidelines:
 - The name must be unique across all existing bucket names in Amazon S3.
 - After you create the bucket you cannot change the name, so choose wisely.

- Choose a bucket name that reflects the objects in the bucket because the bucket name is visible in the URL that points to the objects that you're going to put in your bucket.

For information about naming buckets, see [Rules for Bucket Naming](#) in the *Amazon Simple Storage Service Developer Guide*.

- For **Region**, choose the AWS Region where you want the bucket to reside. Choose a Region close to you to minimize latency and costs, or to address regulatory requirements. Objects stored in a Region never leave that Region unless you explicitly transfer them to another Region. For a list of Amazon S3 AWS Regions, see [Regions and Endpoints](#) in the *Amazon Web Services General Reference*.
- (Optional) If you have already set up a bucket that has the same settings that you want to use for the new bucket that you want to create, you can set it up quickly by choosing **Copy settings from an existing bucket**, and then choosing the bucket whose settings you want to copy.
- Do one of the following:
 - If you copied settings from another bucket, choose **Create**. You're done, so skip the following steps.
 - If not, choose **Next**.

The screenshot shows the 'Create bucket' wizard in the AWS Management Console. The title bar is blue with a close button (X) on the right. Below the title bar is a progress bar with four steps: 1. Name and region (highlighted), 2. Set properties, 3. Set permissions, and 4. Review. The main content area is dark blue. Under the heading 'Name and region', there is a 'Bucket name' section with a text input field containing the placeholder 'Enter DNS-compliant bucket name'. Below that is a 'Region' section with a dropdown menu showing 'Select a region'. Further down is a 'Copy settings from an existing bucket' section with a dropdown menu showing 'Select bucket (optional)' and '33 Buckets'. At the bottom of the form are three buttons: 'Create' (white), 'Cancel' (blue), and 'Next' (grey).

- On the **Set properties** page, you can configure the following properties for the bucket. Or, you can configure these properties later, after you create the bucket.
 - Versioning** – You can store every version of every object in your bucket. Versioning is enabled for a new bucket by default. To disable versioning for the bucket, choose **Versioning**, choose **Disable versioning**, and then choose **Save**. For more information, see [How Do I Enable or Suspend Versioning for an S3 Bucket?](#) (p. 11).
 - Logging** – Server access logging provides detailed records for the requests made to your bucket. By default, Amazon S3 does not collect server access logs. To enable logging for the bucket, choose **Logging**. To disable logging, choose **Disable logging**. Choose **Save** to save your

settings. For more information, see [Server Access Logging](#) in the *Amazon Simple Storage Service Developer Guide*

- c. **Tags** – With AWS cost allocation, you can use tags to annotate billing for your use of a bucket. A tag is a key-value pair that represents a label that you assign to a bucket. To add tags, choose **Tags**, and then choose **Add tag**.

The screenshot shows the 'Create bucket' wizard in the AWS Management Console, specifically the 'Set properties' step. The progress bar at the top indicates four steps: 1. Name and region (completed), 2. Set properties (current), 3. Set permissions, and 4. Review. The main content area features three configuration cards: 'Versioning' (description: 'Keep multiple versions of an object in the same bucket'), 'Logging' (description: 'Set up access log records that provide details about access requests'), and 'Tags' (description: 'Use tags to track your cost against projects or other criteria'). Each card has a 'Learn more' link and a radio button set to 'Disabled'. At the bottom right, there are 'Previous' and 'Next' buttons.

5. Choose **Next**.
6. On the **Set permissions** page, you manage permissions. You can make changes to permissions after you create the bucket. When you're done configuring permissions on the bucket, choose **Next**.

Create bucket

✓ Name and region ✓ Set properties **3 Set permissions** 4 Review

▼ Manage users

User ID	Objects	Object permissions
	☒ Read ☒ Write	☒ Read ☒ Write

▼ Manage public permissions

Group	Objects	Object permissions
Any authenticated AWS user	☐ Read ☐ Write	☐ Read ☐ Write
Everyone	☐ Read ☐ Write	☐ Read ☐ Write

7. On the **Review** page, verify the settings. If you see something you want to change, choose **Edit**. If your current settings are correct, choose **Create bucket**.

Create bucket

✓ Name and region ✓ Set properties ✓ Set permissions **4 Review**

Name and region [Edit](#)

Bucket name admin-created-one **Region** US West (Oregon)

Properties [Edit](#)

Versioning	Disabled
Logging	Disabled
Tagging	0 Tags

Permissions [Edit](#)

Users	1
Public permissions	Disabled

[Previous](#) [Create bucket](#)

How Do I Delete an S3 Bucket?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.



Announcement: Object Tagging and new Storage Management features available in new console

[Opt In](#) to try object tagging and storage management.


You can delete a bucket and all of the objects contained in the bucket. You can also delete an empty bucket. When you delete a bucket with versioning enabled, all versions of all the objects in the bucket are deleted. For more information, see [Managing Objects in a Versioning-Enabled Bucket](#) and [Deleting/Emptying a Bucket](#) in the *Amazon Simple Storage Service Developer Guide*.

Important


If you want to continue to use the same bucket name, don't delete the bucket. We recommend that you empty the bucket and keep it. After a bucket is deleted, the name becomes available to reuse, but the name might not be available for you to reuse for various reasons. For example, it might take some time before the name can be reused and some other account could create a bucket with that name before you do.

To delete an S3 bucket

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the bucket icon next to the name of bucket that you want to delete and then choose **Delete bucket**.



3. In the **Delete bucket** dialog box, type the name of the bucket that you want to delete for confirmation and then choose **Confirm**.



How Do I Empty an S3 Bucket?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.



Announcement: Object Tagging and new Storage Management features available in new console


[Opt In](#) to try object tagging and storage management.

You can empty a bucket, which deletes all of the objects in the bucket without deleting the bucket. When you empty a bucket with versioning enabled, all versions of all the objects in the bucket are deleted. For more information, see [Managing Objects in a Versioning-Enabled Bucket](#).


The Amazon S3 console supports emptying your bucket provided that the bucket contains less than 100,000 objects. The Amazon S3 console returns an error if you attempt to empty a bucket that contains more than 100,000 objects. If your bucket contains more than 100,000 objects, you must use other options to empty the bucket, such as the AWS CLI, Amazon S3 bucket lifecycle policies, or programmatically using the AWS SDKs. For more information, see [Deleting/Emptying a Bucket](#) in the *Amazon Simple Storage Service Developer Guide*.

To empty an S3 bucket

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the bucket icon next to the name of bucket that you want to delete and then choose **Empty bucket**.



3. In the **Empty bucket** dialog box, type the name of the bucket you want to empty for confirmation and then choose **Confirm**.



How Do I View the Properties for an S3 Bucket?


If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.

Announcement: Object Tagging and new Storage Management features available in new console
[Opt In](#) to try object tagging and storage management.

This topic explains how to view the properties for an S3 bucket.

To view the properties for an S3 bucket

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that you want to view the properties for.



3. Choose **Properties**.



4. On the **Properties** page, you can configure the following properties for the bucket.
 - a. **Versioning** – You can store every version of every object in your bucket. Versioning is enabled for a new bucket by default. To disable versioning for the bucket, choose **Versioning**, choose **Disable versioning**, and then choose **Save**. For more information, see [How Do I Enable or Suspend Versioning for an S3 Bucket? \(p. 11\)](#).
 - b. **Static website hosting** – You can host a static website on Amazon S3. To enable static website hosting, choose **Static website hosting** and then specify the settings you want to use. For more information, see [How Do I Configure an S3 Bucket for Static Website Hosting? \(p. 12\)](#).
 - c. **Logging** – Server access logging provides detailed records for the requests made to your bucket. By default, Amazon S3 does not collect server access logs. To enable logging for the bucket, choose **Logging**. To disable logging, choose **Disable logging**. Choose **Save** to save your settings.
 - d. **Tags** – With AWS cost allocation, you can use tags to annotate billing for your use of a bucket. A tag is a key-value pair that represents a label that you assign to a bucket. To add tags, choose **Tags** and then choose **Add tag**.
 - e. **Cross-region replication** – Enables automatic, asynchronous copying of objects across buckets in different AWS Regions. To enable cross-region replication, choose **Cross-region replication** and then specify the settings you want to use. For more information, see [How Do I Enable and Configure Cross-Region Replication for an S3 Bucket? \(p. 17\)](#).
 - f. **Events** – You can enable certain Amazon S3 bucket events to send a notification message to a destination whenever the events occur. To enable events, choose **Events** and then specify the settings you want to use. For more information, see [How Do I Enable and Configure Event Notifications for an S3 Bucket? \(p. 23\)](#).

How Do I Enable or Suspend Versioning for an S3 Bucket?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.




Announcement: Object Tagging and new Storage Management features available in new console

Opt In to try object tagging and storage management.

Versioning enables you to keep multiple versions of an object in one bucket. This section describes how to enable object versioning on a bucket. For more information about versioning support in Amazon S3, see [Object Versioning](#) and [Using Versioning](#) in the *Amazon Simple Storage Service Developer Guide*.

To enable or disable versioning on an S3 bucket


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that you want to enable versioning for.




3. Choose **Properties**.



4. Choose **Versioning**.



5. Choose **Enable versioning** or **Suspend versioning**, and then choose **Save**.



How Do I Configure an S3 Bucket for Static Website Hosting?


If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.



You can host a static website on Amazon S3. On a static website, individual web pages include static content and they might also contain client-side scripts. By contrast, a dynamic website relies on server-side processing, including server-side scripts such as PHP, JSP, or ASP.NET. Amazon S3 does not support server-side scripting. For more information, see [Hosting a Static Website on Amazon S3](#) in the *Amazon Simple Storage Service Developer Guide*.

To configure an S3 bucket for static website hosting

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that you want to enable static website hosting for.




3. Choose **Properties**.



4. Choose **Static website hosting**.



After you enable your bucket for static website hosting, web browsers can access all of your content through the Amazon S3 website endpoint for your bucket.



5. Choose **Use this bucket to host**.
 - a. For **Index Document**, type the name of the index document, which is typically named index.html. When you configure a bucket for website hosting, you must specify an index document. Amazon S3 returns this index document when requests are made to the root domain or any of the subfolders. For more information, see [Configure a Bucket for Website Hosting](#) in the *Amazon Simple Storage Service Developer Guide*.

- b. (Optional) For **Error Document**, type the name of a custom error document. If an error occurs, Amazon S3 returns an HTML error document. For 4XX class errors, you can optionally provide your own custom error document, in which you can provide additional guidance to your users. For more information, see [Custom Error Document Support](#) in the *Amazon Simple Storage Service Developer Guide*.
- c. (Optional) For **Edit redirection rules**, describe the rules using XML in the text area if you want to specify advanced redirection rules. For example, you can conditionally route requests according to specific object key names or prefixes in the request. For more information, see [Configure a Bucket for Website Hosting](#) in the *Amazon Simple Storage Service Developer Guide*.

Static website hosting

Endpoint : admin-created3.s3-website-us-west-2.amazonaws.com

☒ Use this bucket to host

Index document

Folder1/index.html

Error document

error.html

Edit redirection rules

☐ Redirect requests

☐ Disable website hosting

Cancel Save

6. Choose **Save**.
7. Add a bucket policy to the website bucket to grant everyone access to the objects in the bucket. When you configure a bucket as a website, you must make the objects that you want to serve publicly readable. To do so, you write a bucket policy that grants everyone `s3:GetObject` permission. The following example bucket policy grants everyone access to the objects in the example-bucket bucket.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PublicReadGetObject",
      "Effect": "Allow",
      "Principal": "*",
      "Action": [
        "s3:GetObject"
      ],
      "Resource": [
```

```
    "arn:aws:s3:::example-bucket/*"  
  }  
}  
}
```

For information about adding a bucket policy, see [How Do I Set Bucket Permissions? \(p. 72\)](#).
For more information, see [Permissions Required for Website](#) in the *Amazon Simple Storage Service Developer Guide*.

Note


If you choose **Disable website hosting**, Amazon S3 removes any existing website configuration from the bucket, and the bucket is not accessible from the website endpoint. However, the bucket is still available at the REST endpoint. For a list of Amazon S3 endpoints, see [Amazon S3 Regions and Endpoints](#) in the *Amazon Web Services General Reference*.

How Do I Redirect Requests to an S3 Bucket Hosted Website to Another Host?

You can redirect all requests to your S3 bucket hosted static website to another host.

To redirect all requests to an S3 bucket's website endpoint to another host

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that you want to redirect all requests from.



3. Choose **Properties**.



4. Choose **Static website hosting**.



5. Choose **Redirect requests**.

A configuration dialog for "Static website hosting". It shows the endpoint "admin-created3.s3-website-us-west-2.amazonaws.com". There are three radio buttons: "Use this bucket to host", "Redirect requests" (which is selected), and "Disable website hosting". Below "Redirect requests" is a text field for "Target Bucket or Domain" containing "Bucketname1 or www.exampledomain.com". Below that is a text field for "Protocol" containing "https or http". At the bottom right are "Cancel" and "Save" buttons.

- a. For **Target bucket or domain**, type the name of the bucket or the domain name where you want requests to be redirected. To redirect requests to another bucket, type the name of the target bucket. For example, if you are redirecting to a root domain address, you would type **www.example.com**. For more information, see [Configure a Bucket for Website Hosting](#) in the *Amazon Simple Storage Service Developer Guide*.
- b. For **Protocol**, type the protocol (http, https) for the redirected requests. If no protocol is specified, the protocol of the original request is used. If you redirect all requests, any request made to the bucket's website endpoint will be redirected to the specified host name.
6. Choose **Save**.

Advanced Settings for S3 Bucket Properties

*If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.*



Announcement: Object Tagging and new Storage Management features available in new console

[Opt In](#) to try object tagging and storage management.

This section describes how to configure advanced S3 bucket property settings for tags, cross-region replication, and event notification.

Topics

- [How Do I Enable and Configure Cross-Region Replication for an S3 Bucket?](#) (p. 17)
- [How Do I Disable Cross-Region Replication for an S3 Bucket?](#) (p. 20)
- [How Do I Set Up a Destination to Receive Event Notifications?](#) (p. 21)
- [How Do I Enable and Configure Event Notifications for an S3 Bucket?](#) (p. 23)

How Do I Enable and Configure Cross-Region Replication for an S3 Bucket?

*If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.*



Announcement: Object Tagging and new Storage Management features available in new console


[Opt In](#) to try object tagging and storage management.

Cross-region replication is the automatic, asynchronous copying of objects across buckets in different AWS Regions. When you enable cross-region replication, Amazon S3 replicates newly created objects, object updates, and object deletions from a source bucket to a destination bucket in a different region. Cross-region replication has specific requirements that define what can and cannot be replicated across regions based on how the object is created and how it is encrypted. For more information, see [Cross-Region Replication](#) in the *Amazon Simple Storage Service Developer Guide*.

Cross-region replication requires that versioning must be enabled on both your source bucket and your destination bucket that is in a different region. For more information, see [How Do I Enable or Suspend Versioning for an S3 Bucket?](#) (p. 11).

To enable cross-region replication of an S3 bucket to another bucket


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that you want to enable cross-region replication for.




3. Choose **Properties**.



4. Under **Advanced settings**, choose **Cross-region replication**.



5. Choose **Enable cross-region replication**, and then configure your settings as follows:
- For **Destination**, choose the region of the destination bucket and then choose the destination bucket. If you do not see your desired destination bucket in the list, confirm that the bucket exists in the region you selected and that you have enabled versioning on that bucket.
 - For **Source**, choose **Whole bucket** to replicate the whole bucket or choose **Prefix in this bucket** to replicate all objects with the same prefix (for example, all objects in a specific folder).
 - If you choose **Prefix in this bucket**, choose the arrow next to **Current prefixes**, choose **+Add**, type a prefix to use, and then choose a destination storage class. You can add more than one prefix.



- For **Destination storage class**, choose the storage class you want to use for the replicated objects.

- d. To perform cross-region replication of objects on your behalf, you need to set up an AWS Identity and Access Management (IAM) role that Amazon S3 can use. For **Select role**, do one of the following:
- If you want Amazon S3 to create a new IAM role for you, choose **Create new role** and then choose **Save**. Amazon S3 will generate a policy for the IAM role that matches the source and destination buckets you choose. The generated role is named based on the bucket names using the following naming convention: **replication_role_for_*source-bucket*_to_*destination-bucket***
 - If you want to use an existing IAM role, choose an IAM role that allows Amazon S3 to replicate objects from the source bucket to the destination bucket on your behalf and then choose **Save**.

Cross-region replication ✕

☒ Enable cross-region replication

Source

Region: US West (Oregon) (us-west-2)

Whole bucket

Destination

US West (N. California)

ca-example-bucket

Destination storage class

Standard - IA

Select role

replication_role_for_admin-created2_to_ca-example-bucket

☐ Disable cross-region replication

Cancel Save

You have now enabled cross-region replication of one bucket to another. The time it takes for Amazon S3 to replicate an object depends on the object size. It can take up to several hours to replicate a large-sized object.

Note

Metadata for an object remains identical between original objects and replica objects. Lifecycle rules abide by the creation time of the original object, and not by when the replicated object becomes available in the destination bucket. However, lifecycle actions on objects pending replication do not resolve until the replication has completed.

How Do I Disable Cross-Region Replication for an S3 Bucket?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.


Announcement: Object Tagging and new Storage Management features available in new console
[Opt In](#) to try object tagging and storage management.

Cross-region replication is the automatic, asynchronous copying of objects across buckets in different AWS Regions. For more information, see [Cross-Region Replication](#) in the Amazon Simple Storage Service Developer Guide.

Cross-region replication requires that versioning must be enabled on both your source bucket and your destination bucket that is in a different region. For more information, see [How Do I Enable or Suspend Versioning for an S3 Bucket?](#) (p. 11).

To disable cross-region replication of an S3 bucket to another bucket


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that you want to enable versioning for.




3. Choose **Properties**.



4. Under **Advanced settings**, choose **Cross-region replication**.




5. Choose **Disable cross-region replication**.



6. Choose **Save**.

How Do I Set Up a Destination to Receive Event Notifications?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.



Before you can enable event notifications for your bucket you must set up one of the following destination types:

An Amazon SNS topic

Amazon Simple Notification Service (Amazon SNS) is a web service that coordinates and manages the delivery or sending of messages to subscribing endpoints or clients. You can use the Amazon SNS console to create an Amazon SNS topic that your notifications can be sent to. The Amazon SNS topic must be in the same region as your Amazon S3 bucket. For information about creating an Amazon SNS topic, see [Getting Started](#) in the *Amazon Simple Notification Service Developer Guide*.

Before you can use the Amazon SNS topic that you create as an event notification destination, you need the following:

- The Amazon Resource Name (ARN) for the Amazon SNS topic
- A valid Amazon SNS topic subscription (the topic subscribers are notified when a message is published to your Amazon SNS topic)
- A permissions policy that you set up in the Amazon SNS console (as shown in the following example)

```
{
  "Version": "2012-10-17",
  "Id": "__example_policy_ID",
  "Statement": [
    {
      "Sid": "example-statement-ID",
      "Effect": "Allow",
      "Principal": "*",
      "Action": "SNS:Publish",
      "Resource": "arn:aws:sns:region:account-number:topic-name",
      "Condition": {
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:s3::bucket-name"
        }
      }
    }
  ]
}
```

An Amazon SQS queue

You can use the Amazon SQS console to create an Amazon SQS queue that your notifications can be sent to. The Amazon SQS queue must be in the same region as your Amazon S3 bucket. For information about creating an Amazon SQS queue, see [Getting Started with Amazon SQS](#) in the *Amazon Simple Queue Service Developer Guide*.

Before you can use the Amazon SQS queue as an event notification destination, you need the following:

- The Amazon Resource Name (ARN) for the Amazon SQS topic
- A permissions policy that you set up in the Amazon SQS console (as shown in the following example)

```
{
  "Version": "2012-10-17",
  "Id": "__example_policy_ID",
  "Statement": [
    {
      "Sid": "example-statement-ID",
      "Effect": "Allow",
      "Principal": "*",
      "Action": "SQS:*",
      "Resource": "arn:aws:sqs:region:account-number:queue-name",
      "Condition": {
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:s3::bucket-name"
        }
      }
    }
  ]
}
```

A Lambda function

You can use the AWS Lambda console to create a Lambda function. The Lambda function must be in the same region as your S3 bucket. For information about creating a Lambda function, see the [AWS Lambda Developer Guide](#).

Before you can use the Lambda function as an event notification destination, you must have the name or the ARN of a Lambda function to set up the Lambda function as a event notification destination.

For information about using Lambda with Amazon S3, see [Using AWS Lambda: with Amazon S3](#) in the *AWS Lambda Developer Guide*.

How Do I Enable and Configure Event Notifications for an S3 Bucket?

*If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.*



Announcement: Object Tagging and new Storage Management features available in new console

[Opt In](#) to try object tagging and storage management.

You can enable certain Amazon S3 bucket events to send a notification message to a destination whenever the events occur. This section explains how to use the Amazon S3 console to enable event notifications. For more information about using event notifications, see [Configuring Notifications for Amazon S3 Events](#) in the *Amazon Simple Storage Service Developer Guide*.

Amazon S3 can send notifications for the following events:

- **An object created event** – You choose **ObjectCreated (All)** when configuring your events in the console to enable notifications for anytime an object is created in your bucket. Or, you can select one or more of the specific object-creation actions to trigger event notifications. These actions are **Put**, **Post**, **Copy**, and **CompleteMultiPartUpload**.
- **An object delete event** – You select **ObjectDelete (All)** when configuring your events in the console to enable notification for anytime an object is deleted. Or, you can select **Delete** to trigger event notifications when an unversioned object is deleted or a versioned object is permanently deleted. You select **Delete Marker Created** to trigger event notifications when a delete marker is created for a versioned object.
- **A Reduced Redundancy Storage (RRS) object lost event** – You select **RRSObjectLost** to be notified when Amazon S3 detects that an object of the RRS storage class has been lost.


Event notification messages can be sent to the following types of destinations:

- **An Amazon Simple Notification Service (Amazon SNS) topic** – A web service that coordinates and manages the delivery or sending of messages to subscribing endpoints or clients.
- **An Amazon Simple Queue Service (Amazon SQS) queue** – Offers reliable and scalable hosted queues for storing messages as they travel between computer.
- **A Lambda function** – AWS Lambda is a compute service where you can upload your code and the service can run the code on your behalf using the AWS infrastructure. You package up and upload your custom code to AWS Lambda when you create a Lambda function

Before you can enable event notifications for your bucket you must set up one of these destination types. For more information, see [How Do I Set Up a Destination to Receive Event Notifications?](#) (p. 21).

To enable and configure event notifications for an S3 bucket


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that you want to enable events for.




3. Choose **Properties**.




4. Under Advanced settings, choose **Events**.



5. Choose **Add notification**.



6. In **Name**, type a descriptive name for your event configuration. If you do not enter a name, a GUID is autogenerated and used for the name.




7. Under **Events**, select one or more of the type of event occurrences that you want to receive notifications for. When the event occurs a notification is sent to a destination that you choose. For example, you could do any of the following:
 - Select **ObjectCreate (All)** to enable event notifications for anytime an object is created in the bucket.
 - Select **Put** and **Complete MultipartUpload** to trigger event notifications anytime a new object is put into a bucket and anytime a multipart upload completes.
 - Select **ObjectDelete (All)** to enable event notifications for anytime an object is deleted in the bucket.
 - Select **Delete** or **Delete Marker Created** to trigger notifications for specific types of object deletes.

For information about deleting versioned objects, see [Deleting Object Versions](#). For information about object versioning, see [Object Versioning](#) and [Using Versioning](#).

Note

When you delete the last object from a folder Amazon S3 can generate an object creation event. The Amazon S3 console displays a folder under the following circumstances: 1) when a zero byte object has a trailing slash (/) in its name (in this case there is an actual Amazon S3 object of 0 bytes that represents a folder), and 2) if the object has a slash (/) within its name (in this case there isn't an actual object representing the folder). When there are multiple objects with the same prefix with a trailing slash (/) as part of their names, those objects are shown as being part of a folder. The name of the folder is formed from the characters preceding the trailing slash (/). When you delete all the objects listed under that folder, there is no actual object available to represent the empty folder. Under such circumstance the Amazon S3 console creates a zero byte object to represent that folder. If you enabled event notification for creation of objects, the zero byte object creation action that is taken by the console will trigger an object creation event.




8. Type an object name **Prefix** and/or a **Suffix** to filter the event notifications by the prefix and/or suffix. For example, you can set up a filter so that you are sent a notification only when files are added

to an image folder (for example, objects with the name prefix `images/`). For more information, see [Configuring Notifications with Object Key Name Filtering](#).




The screenshot shows two input fields. The first field is labeled "Prefix" and contains the text "e.g. images/". Below it is an information icon (i). The second field is labeled "Suffix" and contains the text "e.g. .jpg".

9. Select the type of destination to have the event notifications sent to.




The screenshot shows a dropdown menu labeled "Send to". The menu is open, showing three options: "SNS Topic", "SQS Queue", and "Lambda Function". The "SNS Topic" option is highlighted in blue. There is an information icon (i) to the right of the dropdown.

- a. If you select the **SNS Topic** destination type.
 - i. In the **SNS topic** box, type the name or select from the menu, the Amazon SNS topic that will receive notifications from Amazon S3. For information about the Amazon SNS topic format, see [SNS FAQ](#).



The screenshot shows the "Send to" dropdown menu set to "SNS Topic". Below it is a text input field labeled "SNS" with a dropdown arrow on the right. There is an information icon (i) to the right of the "Send to" dropdown.

- ii. (Optional) You can also select **Add SNS topic ARN** from the menu and type the **ARN** of the SNS topic in **SNS topic ARN**.



The screenshot shows the "Send to" dropdown menu set to "SNS Topic". Below it is a text input field labeled "SNS" with a dropdown arrow on the right. The "Add SNS topic ARN" option is highlighted in blue. There is an information icon (i) to the right of the "Send to" dropdown.

- b. If you select the **SQS queue** destination type, do the following:

- i. In **SQS queue**, type or choose a name from the menu of the Amazon SQS queue that you want to receive notifications from Amazon S3. For information about Amazon SQS, see [What is Amazon Simple Queue Service?](#) in the *Amazon Simple Queue Service Developer Guide*.
- ii. (Optional) You can also select **Add SQS topic ARN** from the menu and type the ARN of the SQS queue in **SQS queue ARN**.
- c. If you select the **Lambda Function** destination type, do the following:
 - i. In **Lambda Function**, type or choose the name of the Lambda function that you want to receive notifications from Amazon S3.
 - ii. If you don't have any Lambda functions in the region that contains your bucket, you'll be prompted to enter a Lambda function ARN. In **Lambda Function ARN**, type the ARN of the Lambda function that you want to receive notifications from Amazon S3.
 - iii. (Optional) You can also choose **Add Lambda function ARN** from the menu and type the ARN of the Lambda function in **Lambda function ARN**.

For information about using Lambda with Amazon S3, see [Using AWS Lambda: with Amazon S3](#) in the *AWS Lambda Developer Guide*.

10. Choose **Save**. Amazon S3 will send a test message to the event notification destination.

Uploading, Downloading, and Managing Objects

*If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.*



Announcement: Object Tagging and new Storage Management features available in new console

[Opt In](#) to try object tagging and storage management.

Amazon S3 is cloud storage for the Internet. To upload your data (photos, videos, documents etc.), you first create a bucket in one of the AWS Regions. You can then upload an unlimited number of data objects to the bucket.

The data that you store in Amazon S3 consists of objects. Every object resides within a bucket that you create in a specific AWS Region. Every object that you store in Amazon S3 resides in a bucket.

Objects stored in a region never leave the region unless you explicitly transfer them to another region. For example, objects stored in the EU (Ireland) region never leave it. The objects stored in an AWS region physically remain in that region. Amazon S3 does not keep copies of objects or move them to any other region. However, you can access the objects from anywhere, as long as you have necessary permissions to do so.

Before you can upload an object into Amazon S3, you must have write permissions to a bucket.

Objects can be any file type: images, backups, data, movies, etc. The maximum size of file you can upload by using the Amazon S3 console is 78GB. You can have an unlimited number of objects in a bucket.

The following topics explain how to use the Amazon S3 console to upload, delete, and manage objects.

Topics

- [How Do I Upload an Object to an S3 Bucket? \(p. 29\)](#)
- [How Do I Download an Object from an S3 Bucket? \(p. 35\)](#)
- [How Do I Delete Objects from an S3 Bucket? \(p. 39\)](#)
- [How Do I Undelete a Deleted S3 Object? \(p. 41\)](#)
- [How Do I Delete Folders from an S3 Bucket? \(p. 43\)](#)
- [How Do I See an Overview of an Object? \(p. 45\)](#)
- [How Do I See the Versions of an S3 Object? \(p. 46\)](#)

- [How Do I View the Properties of an Object? \(p. 48\)](#)
- [How Do I Add Tags to an Object? \(p. 50\)](#)

How Do I Upload an Object to an S3 Bucket?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.



Announcement: Object Tagging and new Storage Management features available in new console

[Opt In](#) to try object tagging and storage management.

This section explains how to use the AWS Management Console to upload one or more files or entire folders to an Amazon S3 bucket.

Before you can upload an object to an Amazon S3 bucket, you must have write permissions for the bucket.

Objects can be any file type: images, backups, data, movies, etc. The maximum size of a file that you can upload by using the Amazon S3 console is 78 GB. You can have an unlimited number of objects in a bucket.


When you upload a folder, Amazon S3 uploads all of the files and subfolders from the specified folder to your bucket. It then assigns a key name that is a combination of the uploaded file name and the folder name. For example, if you upload a folder called /images that contains two files, sample1.jpg and sample2.jpg, Amazon S3 uploads the files and then assigns the corresponding object key names, images/sample1.jpg and images/sample2.jpg. The key names include the folder name as a prefix.

If you upload files that are not in a folder, when Amazon S3 uploads the files, it assigns only the file names as the key name for the objects created.

If you upload an object with a key name that already exists in a versioning-enabled bucket, Amazon S3 creates another version of the object instead of replacing the existing object.

To upload an object to an S3 bucket


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that you want to upload your objects to.




3. Choose **Upload**.




- If the bucket is empty, choose **Get started** or **Upload an object**.




4. In the **Upload** dialog box, choose **Add files** to select the files to upload.




5. In the dialog box, use one of the following methods to add the files that you want to upload:
- a. Choose one or more files and folders to upload, and then choose **Open**.




- b. If you are using the Google Chrome browser, you can choose one or more files to upload, and then drag and drop your selection into the **Upload** dialog box.



6. The files you chose are listed in the **Upload** dialog box.
- a. To add more files, choose **Add more files**.
 - b. To immediately upload the files, choose **Upload**.
 - c. To continue on to setting permissions or properties for the files that you are uploading, choose **Next**.



7. On the **Set Permissions** page, you can grant or remove permissions for specific users and set public permissions for the files you are uploading. Make the changes, and then choose **Next**.



8. On the **Set Properties** page, choose the storage class and encryption method to use for the objects you are uploading. You can also add or modify metadata.
- Choose a storage class for the objects you are uploading. For more information about storage classes, see [Storage Classes](#) in the *Amazon Simple Storage Service Developer Guide*.

- b. Choose the type of encryption for the objects you're uploading, or, if you don't want to encrypt the objects you're uploading, choose **None**.

The screenshot shows the 'Upload' console in the Amazon S3 console. The progress bar at the top indicates four steps: 'Select files' (completed), 'Set permissions' (completed), 'Set properties' (current step, indicated by a circled 3), and 'Review' (indicated by a circled 4). The 'Set properties' section is divided into three parts: 'Storage class', 'Encryption', and 'Metadata'. Under 'Storage class', the instruction is 'Choose one depending on your use case scenario and performance access requirements.' There are three radio buttons: 'Standard' (selected), 'Standard-IA', and 'Reduced redundancy'. Under 'Encryption', the instruction is 'Protect data in transit by using Amazon S3 master-key or by using AWS KMS master-key.' There are three radio buttons: 'None' (selected), 'Amazon S3 master-key', and 'AWS KMS master-key'. The 'Metadata' section is partially visible at the bottom.

- i. To encrypt your uploaded objects using keys that are managed by Amazon S3, choose **Amazon S3 master-key**. For more information, see [Protecting Data with Amazon S3-Managed Encryption Keys Classes](#) in the *Amazon Simple Storage Service Developer Guide*.
- ii. To encrypt your uploaded objects using the AWS Key Management Service (AWS KMS), choose **AWS KMS master-key** and then choose a master key from the list of the AWS KMS master keys that you have previously created.

Note

You can use only keys in the same AWS Region as this bucket to encrypt objects in this bucket.

- You can give an external account the ability to use an object that is protected by an AWS KMS key. To do this, select **Custom KMS ARN** from the list and enter the Amazon Resource Name (ARN) for the external account. Administrators of an external account that have usage permissions to an object protected by your AWS KMS key can further restrict access by creating a resource-level IAM policy.

For more information about creating an AWS KMS key, see [Creating Keys](#) in the *AWS Key Management Service Developer Guide*. For more information, see [Protecting Data with AWS KMS-Managed Key](#) in the *Amazon Simple Storage Service Developer Guide*.

- c. If you want to add Amazon S3 metadata to all of the objects you are uploading, for **Header**, select a header. You can select common HTTP headers, such as **Content-Type** and **Content-Disposition**. Type a value for the header, and then choose **Save**.
- d. To add user-defined custom metadata to all of the objects that you are uploading, type `x-amz-meta-` plus the metadata name in the **Header** field. Type a value for the header, and then choose **Save**.
- Amazon S3 object metadata is represented by a key-value pair. User metadata is stored with the object and returned when you download the object. Amazon S3 does not process custom metadata. Custom user-defined metadata can be as large as 2 KB, and both the keys and their values must conform to US-ASCII standards. Any metadata starting with prefix `x-amz-meta-` is treated as user-defined metadata.


The screenshot shows the 'Metadata' configuration page. At the top, it says 'Metadata is a set of name-value pairs. You cannot modify object metadata after it is uploaded.' Below this is a table with two columns: 'Header' and 'Value'. The first row has a dropdown menu with 'Select a header ...' and a text input field with 'Header value'. The second row has a text input field with 'x-amz-meta-' and a text input field with 'Header value'. Each row has a 'Save' button to its right. At the bottom of the form are three buttons: 'Upload', 'Previous', and 'Next'.

9. Choose **Next**.
10. On the **Upload** review page, verify that your settings are correct, and then choose **Upload**. To make changes, choose **Previous**.


The screenshot shows the 'Upload' review page. At the top is a blue header with the word 'Upload' and a close button. Below the header is a progress bar with four steps: 'Select files', 'Set permissions', 'Set properties', and 'Review' (which is the current step, indicated by a circled 4). Below the progress bar are three sections: 'Files', 'Permissions', and 'Properties'. The 'Files' section shows '9 Files' and 'Size: 5.6 MB'. The 'Permissions' section shows '1 grantees'. The 'Properties' section shows 'Encryption: No' and 'Storage class: Standard'. At the bottom right are two buttons: 'Previous' and 'Upload'.

11. To see the progress of your upload, choose **In progress** at the bottom of the browser window.

The screenshot shows a progress bar with four items: 'Operations', '1 In progress', '1 Success', and '0 Error'. A red arrow points to the '1 In progress' item.



To see a history of your uploads and other operations, choose **Success**.



More Info

- [How Do I Download an Object from an S3 Bucket? \(p. 35\)](#)

How Do I Download an Object from an S3 Bucket?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.


Announcement: Object Tagging and new Storage Management features available in new console
[Opt In](#) to try object tagging and storage management.

This section explains how to use the Amazon S3 console to download objects from an S3 bucket.


Data transfer fees apply when you download objects. For information about Amazon S3 features, and pricing, see [Amazon S3](#).

To download an object from an S3 bucket


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that you want to download an object from.




3. You can download an object from an S3 bucket in any of the following ways:
 - In the **Name** list, select the check box next to the object you want to download, and then choose **Download** on the object description page that appears.




- Choose the name of the object that you want to download.




On the **Overview** page, choose **Download**.



- Choose the name of the object that you want to download and then choose **Download as** on the **Overview** page.



- Choose the name of the object that you want to download. Choose **Latest version** and then choose the download icon.



Related Topics

- [How Do I Upload an Object to an S3 Bucket?](#) (p. 29)

How Do I Delete Objects from an S3 Bucket?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.


 **Announcement: Object Tagging and new Storage Management features available in new console**
[Opt In](#) to try object tagging and storage management.

This section explains how to use the Amazon S3 console to delete objects. Because all objects in your S3 bucket incur storage costs, you should delete objects that you no longer need. If you are collecting log files, for example, it's a good idea to delete them when they're no longer needed. You can set up a lifecycle rule to automatically delete objects such as log files.


For information about Amazon S3 features and pricing, see [Amazon S3](#).

To delete objects from an S3 bucket


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that you want to delete an object from.
3. You can delete objects from an S3 bucket in any of the following ways:
 - In the **Name** list, select the check box next to the objects and folders that you want to delete, choose **More**, and then choose **Delete**.



In the **Delete objects** dialog box, verify that the names of the objects and folders you selected for deletion are listed and then choose **Delete**.



- Choose the name of the object that you want to delete, choose **Latest version**, and then choose the trash can icon.



Related Topics

- [How Do I Undelete a Deleted S3 Object? \(p. 41\)](#)
- [How Do I Create a Lifecycle Policy for an S3 Bucket? \(p. 53\)](#)

How Do I Undelete a Deleted S3 Object?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.

 **Announcement: Object Tagging and new Storage Management features available in new console**
[Opt In](#) to try object tagging and storage management.

This section explains how to use the Amazon S3 console to recover (undelete) deleted objects.

To be able to undelete a deleted object, you must have had versioning enabled on the bucket that contains the object before the object was deleted. For information about enabling versioning, see [How Do I Enable or Suspend Versioning for an S3 Bucket? \(p. 11\)](#).


You can undelete only an object that was deleted as the latest (current) version. You can't undelete a previous version of an object that has been deleted. For more information, see [Object Versioning](#) and [Using Versioning](#) in the *Amazon Simple Storage Service Developer Guide*.

To recover deleted objects from an S3 bucket


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that you want to recover a deleted object from.
3. Choose **Deleted objects** to see the objects that have been deleted from the bucket.

4. Select the check box next to the object or objects that you want to recover, and then choose **Undo delete** from the **More** menu.

If the object you want to recover is in a folder. Choose the name of the folder, select the object or objects that you want to recover, and then choose **Undo delete** from the **More** menu.



5. On the **Undo delete** review page, verify that the objects that are listed are correct, and then choose **Undo delete**. Otherwise, choose **Cancel** and return to step 3.



More Info

- [How Do I Enable or Suspend Versioning for an S3 Bucket?](#) (p. 11)
- [How Do I Delete Objects from an S3 Bucket?](#) (p. 39)

How Do I Delete Folders from an S3 Bucket?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.



Announcement: Object Tagging and new Storage Management features available in new console


[Opt In](#) to try object tagging and storage management.

This section explains how to use the Amazon S3 console to delete folders from an S3 bucket.


For information about Amazon S3 features and pricing, see [Amazon S3](#).

To delete folders from an S3 bucket


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that you want to delete folders from.



3. In the **Name** list, select the check box next to the folders and objects that you want to delete, choose **More**, and then choose **Delete**.



In the **Delete objects** dialog box, verify that the names of the folders you selected for deletion are listed and then choose **Delete**.



Related Topics

- [How Do I Delete Objects from an S3 Bucket? \(p. 39\)](#)

How Do I See an Overview of an Object?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.



Announcement: Object Tagging and new Storage Management features available in new console


[Opt In](#) to try object tagging and storage management.

This section explains how to use the console to view the object overview panel.


The overview panel provides an overview of the essential information about an object in one place.

To see the overview panel for an object

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that contains the object.
3. In the **Name** list, select the check box next to the name of the object for which you want an overview.



4. If the bucket is versioning-enabled, you can choose **Latest versions** to see all of the versions of the object.



5. You can then choose the download icon to download an object version, or choose the trash can icon to delete an object version.

Important

You can undelete only an object that was deleted as the latest (current) version. You can't undelete a previous version of an object that has been deleted. For more information, see [Object Versioning](#) and [Using Versioning](#) in the *Amazon Simple Storage Service Developer Guide*.

More Info

- [How Do I See the Versions of an S3 Object?](#) (p. 46)

How Do I See the Versions of an S3 Object?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.



Announcement: Object Tagging and new Storage Management features available in new console

[Opt In](#) to try object tagging and storage management.

This section explains how to use the console to see the multiple versions of an object.

A versioning-enabled bucket can have many versions of the same object, one current (latest) version and zero or more noncurrent (previous) versions. For information about enabling versioning, see [How Do I Enable or Suspend Versioning for an S3 Bucket?](#) (p. 11).


If a bucket is versioning-enabled, Amazon Simple Storage Service creates another version of an object under the following conditions:

- When you upload an object with a name that already exists in the bucket, Amazon S3 creates another version of the object instead of replacing the existing object.
- If you update any object properties after the object is first uploaded, such as changing the storage details or any other metadata changes, then Amazon S3 creates a new object version in the bucket.


For more information about versioning support in Amazon S3, see [Object Versioning](#) and [Using Versioning](#) in the *Amazon Simple Storage Service Developer Guide*.

To see multiple versions of an object

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that contains the object.
3. In the **Name** list, choose the name of the object for which you want to see the versions.



4. Choose **Latest versions** to see a list of the versions of the object.



5. You can then choose the download icon to download an object version, or choose the trash can icon to delete an object version.

Important

You can undelete only an object that was deleted as the latest (current) version. You cannot undelete a previous version of an object that has been deleted. For more information, see [Object Versioning](#) and [Using Versioning](#) in the *Amazon Simple Storage Service Developer Guide*.

You also can view object versions in the object overview panel, see [How Do I See an Overview of an Object?](#) (p. 45).

More Info

- [How Do I Enable or Suspend Versioning for an S3 Bucket?](#) (p. 11)
- [How Do I Create a Lifecycle Policy for an S3 Bucket?](#) (p. 53)

How Do I View the Properties of an Object?


*If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.*

 **Announcement: Object Tagging and new Storage Management features available in new console**
[Opt In](#) to try object tagging and storage management.


This section explains how to use the console to view the properties of an object.

To view the properties of an object

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that contains the object.
3. In the **Name** list, choose the name of the object you want to view the properties for.



4. Choose **Properties**.



How Do I Add Tags to an Object?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.




Announcement: Object Tagging and new Storage Management features available in new console

[Opt In](#) to try object tagging and storage management.


This section explains how to use the console to add tags to an object. For information about object tags, see [Object Tagging](#) in the *Amazon Simple Storage Service Developer Guide*.

To add tags to an object


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that contains the object.
3. In the **Name** list, choose the name of the object you want to add tags to.



4. Choose **Properties**.



5. Choose **Tags** and then choose **Add Tag**.



6. Choose **Tags** and then choose **Add Tag**.

Tags

Key	Value
+ Add Tag	

To enable replication of object tags IAM policies used for Cross-Region Replication must be updated if they were created prior to the introduction of Object tagging.

[Cancel](#) [Save](#)

7. Each tag is a key-value pair. Type a **Key** and a **Value**. Then choose **Add Tag** to add another tag or choose **Save**.

You can enter up to 10 tags for an object.

Tags

Key	Value
dog	corgi X

[+ Add Tag](#)

To enable replication of object tags IAM policies used for Cross-Region Replication must be updated if they were created prior to the introduction of Object tagging.

[Cancel](#) [Save](#)

Storage Management

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.



Announcement: Object Tagging and new Storage Management features available in new console

[Opt In](#) to try object tagging and storage management.

This section explains how to configure Amazon S3 storage management tools.

Topics

- [How Do I Create a Lifecycle Policy for an S3 Bucket? \(p. 53\)](#)
- [How Do I Configure Storage Class Analysis? \(p. 59\)](#)
- [How Do I Configure Storage Inventory? \(p. 63\)](#)
- [How Do I Configure Request Metrics for an S3 Bucket? \(p. 66\)](#)
- [How Do I Configure a Request Metrics Filter? \(p. 68\)](#)

How Do I Create a Lifecycle Policy for an S3 Bucket?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.



Announcement: Object Tagging and new Storage Management features available in new console

[Opt In](#) to try object tagging and storage management.


You can use lifecycle policies to define actions you want Amazon S3 to take during an object's lifetime (for example, transition objects to another storage class, archive them, or delete them after a specified period of time).

You can define a lifecycle policy for all objects or a subset of objects in the bucket by using a shared prefix (that is, objects that have names that begin with a common string).


A versioning-enabled bucket can have many versions of the same object, one current version and zero or more noncurrent (previous) versions. Using a lifecycle policy, you can define actions specific to current and noncurrent object versions. For more information, see [Object Lifecycle Management](#) and [Object Versioning](#) and [Using Versioning](#) in the *Amazon Simple Storage Service Developer Guide*.

To create a lifecycle policy

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that you want to create a lifecycle policy for.



3. Choose the **Management** tab, and then choose **Add lifecycle rule**.



- If the bucket does not have a lifecycle policy, you can choose **Get started**.

There is no lifecycle policy applied to this bucket.
Here is how to get started.



Use lifecycle policies to manage your objects

You can manage an object's lifecycle by using a lifecycle policy, which defines how Amazon S3 manages objects during their lifetime.

[Learn more](#)



Automate transition to tiered storage

Lifecycle policies enable you to automatically transition objects to the Standard - IA and/or to the Amazon Glacier storage class.

[Learn more](#)



Expire your objects

Using a lifecycle policy, you can automatically expire objects based on your retention needs or clean up incomplete multipart uploads.

[Learn more](#)

[Get started](#)

4. In the **Lifecycle rule** dialog box, type a name for your rule to help identify the rule later. The name must be unique within the bucket. Configure the rule as follows:
- To apply this lifecycle rule to all objects with a specified name prefix (i.e., objects whose name begins with a common string), type in a prefix. You can also limit the lifecycle rule scope to one or more object tags. You can combine a prefix and tags. For more information about object name prefixes, see [Object Keys](#) in the *Amazon Simple Storage Service Developer Guide*. For more information about object tags, see [Object Tagging](#) in the *Amazon Simple Storage Service Developer Guide*
 - To apply this lifecycle rule to all objects in the bucket, choose **Next**.

Lifecycle rule

1 Name and scope

2 Transitions

3 Expiration

4 Review

Enter a rule name

testrule

Add filter to limit scope to prefix/tags ⓘ

Type to add prefix/tag filter

Type in a prefix name or tag key name

Cancel


Next

5. You configure lifecycle rules by defining rules to transition objects to the Standard-IA and Amazon Glacier storage classes. For more information, see [Storage Classes](#) in the *Amazon Simple Storage Service Developer Guide*.

You can define transitions for current or previous object versions, or for both current and previous versions.

- a. Select **Current version** to define transitions that are applied when an object is created that is within the scope of the rule.

Select **Previous version** to define transitions that are applied when an object is created that is within the scope of the rule.



The screenshot shows a dark-themed dialog box titled "Configure transition". It has two sections. The first section, "Object transitions and/or expiration", contains two radio buttons: "Current version" (which is selected) and "Previous versions". The second section, "For current version of objects", contains a table with two columns: "Object creation" and "Days after object creation". At the bottom left of the dialog is a "+ Add transition" link.

- b. Choose **Add transitions** and specify one of the following transitions:
- Choose **Transition to Standard-IA after**, and then type the number of days after the creation of an object that you want the transition to be applied (for example, 30 days).
 - Choose **Transition to Amazon Glacier after**, and then type the number of days after the creation of an object that you want the transition to be applied (for example, 100 days).



The screenshot shows a close-up of the transition selection process. It features a "+ Add transition" link at the top left. Below it is a dropdown menu with the placeholder text "Select a transition". The dropdown is open, showing two options: "Transition to Standard-IA after" and "Transition to Amazon Glacier after". To the right of the dropdown is a text input field labeled "Days" with a clear button (X) to its right.

6. When you are done configuring transitions, choose **Next**.

The screenshot shows the 'Lifecycle rule' configuration window in the Amazon S3 console. The progress bar at the top indicates four steps: 1. Name and scope (checked), 2. Transitions (active), 3. Expiration (highlighted), and 4. Review. The main content area is divided into two sections: 'For current version of objects' and 'For previous versions of objects'. Each section has a table with columns for 'Object creation' (or 'Object becomes a previous version'), 'Days after object creation', and a checkbox. In the 'For current version of objects' section, there is one rule: 'Transition to Standard-IA after' 30 Days, with the checkbox checked. In the 'For previous versions of objects' section, there are two rules: 'Transition to Standard-IA after' 30 Days (checkbox checked) and 'Transition to Amazon Glacier after' 100 Days (checkbox checked). At the bottom right, there are 'Previous' and 'Next' buttons.

Object creation	Days after object creation	
+ Add transition		
Transition to Standard-IA after	30 Days	X

Object becomes a previous version	Days after object creation	
+ Add transition		
Transition to Standard-IA after	30 Days	X
Transition to Amazon Glacier after	100 Days	X

7. Select **Expiration** and then enter the number of days after object creation to delete the object (for example, 455 days).
8. Select **Permanently delete previous versions** and then enter the number of days after an object becomes a previous version to permanently delete the object (for example, 455 days).
9. It is a recommended best practice to always select **Clean up incomplete multipart uploads**. For example, type 7 for the number of days after the multipart upload initiation date that you want to end and clean up any multipart uploads that have not completed. For more information about multipart uploads, see [Multipart Upload Overview](#) in the Amazon Simple Storage Service Developer Guide.
10. Choose **Next**.

The screenshot shows the 'Lifecycle rule' console with the 'Expiration' step selected. The progress bar at the top indicates the following steps: 1. Name and scope (checked), 2. Transitions (checked), 3. Expiration (selected), and 4. Review (circled). The 'Expiration' section contains three checked options:

- Expiration**: After 455 Days from object creation
- Permanently delete previous versions**: After 455 Days from becoming a previous version
- Clean up incomplete multipart uploads**: After 7 Days from start of upload

At the bottom right, there are 'Previous' and 'Next' buttons.


11. For **Review**, verify the settings for your rule. If you need to make changes, choose **Previous**. Otherwise, choose **Save**.

The screenshot shows the 'Lifecycle rule' console with the 'Review' step selected. The progress bar at the top indicates the following steps: 1. Name and scope (checked), 2. Transitions (checked), 3. Expiration (checked), and 4. Review (selected). The 'Review' section displays the following information:

- Name and scope**: Name: TestRule, Scope: videos/ (with an 'Edit' link)
- Transitions** (with an 'Edit' link):
 - For current version of objects:
 - 1st transition to Standard-IA after 30 days
 - 2nd transition to Amazon Glacier after 100 days
 - For previous versions of objects:
 - 1st transition to Standard-IA after 30 days
 - 2nd transition to Amazon Glacier after 100 days
- Expiration** (with an 'Edit' link):
 - Expire after after 455 days
 - Permanently delete after after 455 days
 - Clean up incomplete multipart uploads after 7 days

At the bottom right, there are 'Previous' and 'Save' buttons.

12. If the rule does not contain any errors, it is listed on the **Lifecycle** page and is enabled.



How Do I Configure Storage Class Analysis?


If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.




By using the Amazon S3 analytics storage class analysis tool you can analyze storage access patterns to help you decide when to transition the right data to the right storage class. Storage class analysis observes data access patterns to help you determine when to transition less frequently accessed STANDARD storage to the STANDARD_IA (IA, for infrequent access) storage class. For more information, see [Amazon S3 Analytics – Storage Class Analysis](#) in the *Amazon Simple Storage Service Developer Guide*.

To configure storage class analysis


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket for which you want to configure storage class analysis.



3. Choose the **Management** tab, and then choose **Analytics**.




4. Choose **Add**.



5. Type a name for the filter. If you want to analyze the whole bucket, leave the **Prefix / tags** field empty.

A screenshot of the 'Add filter' dialog box in the Amazon S3 console. The dialog has a blue header with the text 'Add filter'. Below the header, there are two text input fields. The first field is labeled 'Filter name' and has a placeholder text 'Enter a name for this filter'. The second field is labeled 'Prefix / tags to monitor (optional)' and has a placeholder text 'Leave empty for entire bucket'. Below these fields, there is a link that says 'Export data (optional)'. At the bottom of the dialog, there are two buttons: 'Save' and 'Cancel'.

6. In the **Prefix / tags** field, type text for the prefix or tag for the objects that you want to analyze, or choose from the dropdown list that appears when you start typing.



Add filter

Filter name

testfilter

Prefix / tags to monitor (optional)

images


prefix images (press enter)

tag images

► Export data (optional)

Save Cancel

7. If you chose **tag**, enter a value for the tag. You can enter one prefix and multiple tags.



Add filter

Filter name

testfilter

Prefix / tags to monitor (optional)

prefix videos X

tag dog | corgi X

tag dog | bulldog X


Leave empty for entire bucket

Type in a prefix name or tag key name


► Export data (optional)

Save Cancel

8. Optionally, you can choose **Export data** to export analysis reports to a comma-separated values (.csv) flat file. Choose a destination bucket where the file can be stored. You can type a prefix for the destination bucket. The destination bucket must be in the same AWS Region as the bucket for which you are setting up the analysis.




9. Choose **Save**.
10. If you chose the **Export data** option, you must create a bucket policy on the destination bucket to grant permissions to Amazon S3 to write objects to the bucket.
 - a. Choose the name of the destination bucket from the **Bucket name** list. (If the bucket is in a different account, sign into that account.)
 - b. Choose **Permissions**, and then choose **Bucket Policy**.



- c. Copy the bucket policy from [Granting Permissions for Amazon S3 Inventory and Amazon S3 Analytics](#) in the *Amazon Simple Storage Service Developer Guide* and paste the policy into the **Bucket policy editor** text box.

In the **Bucket policy editor**, make the following changes to the example policy:

- In Resource, replace destination-bucket with the name of the destination bucket that you chose in Step 8. This is the bucket that you're adding the policy to.
- In Condition - ArnLike, replace source-bucket with the name of the source bucket for the analysis.
- In Condition - StringEquals, replace 1234567890 with the account number of the source bucket.



d. Choose **Save**.

For information about the exported data and how the filter works, see [Amazon S3 Analytics – Storage Class Analysis](#) in the *Amazon Simple Storage Service Developer Guide*.

How Do I Configure Storage Inventory?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.




Announcement: Object Tagging and new Storage Management features available in new console

[Opt In](#) to try object tagging and storage management.


Amazon S3 inventory is one of the tools Amazon S3 provides to help you manage your storage. You can simplify and speed up business workflows and big data jobs using the Amazon S3 inventory, which provides a scheduled alternative to the Amazon S3 synchronous List API operation. Amazon S3 inventory provides a comma-separated values (.csv) flat-file output of your objects and their corresponding metadata on a daily or weekly basis for an S3 bucket or for objects that share a prefix (objects that have names that begin with the same string). For more information, see [Amazon S3 Storage Inventory](#) in the *Amazon Simple Storage Service Developer Guide*.

To configure inventory


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket for which you want to configure storage inventory.



3. Choose the **Management** tab, and then choose **Inventory**.



4. Choose **Add new**.




5. Type a name for the inventory and set it up as follows:
- Optionally, add a prefix for your filter to inventory only objects whose names begin with the same string.
 - Choose the destination bucket where you want reports to be saved. The destination bucket must be in the same AWS Region as the bucket for which you are setting up the inventory.
 - Optionally, choose a prefix for the destination bucket.
 - Choose how frequently to generate the inventory.

The screenshot shows the 'Add new' form in the Amazon S3 console. It has five main sections: 'Inventory name' with a text input field 'Enter inventory name'; 'Filters' with a text input field 'Filter by prefix (optional)'; 'Destination bucket' with a dropdown menu 'Select bucket'; 'Destination prefix' with a text input field 'Type prefix (optional)'; and 'Frequency' with a dropdown menu 'Daily'. Below these fields is a '+ Advanced settings' link. At the bottom are 'Cancel' and 'Save' buttons.

6. Choose **Save**.


7. Create a bucket policy on the destination bucket to grant Amazon S3 permission to write objects to the bucket.
 - a. From the **Bucket name** list, choose the name of the destination bucket. If the bucket is in a different account, sign into that account.
 - b. Choose **Permissions**, and then choose **Bucket Policy**.



- c. Copy the bucket policy from [Granting Permissions for Amazon S3 Inventory and Amazon S3 Analytics](#) in the *Amazon Simple Storage Service Developer Guide* and paste it into the **Bucket policy editor** text box.

In the **Bucket policy editor**, make the following changes to the example policy:

- In Resource, replace destination-bucket with the name of the destination bucket that you chose in Step 5. This is the bucket that you're adding the policy to.
- In Condition - ArnLike, replace source-bucket with the name of the source bucket for the inventory.
- In Condition - StringEquals, replace 1234567890 with the account number of the source bucket.



- d. Choose **Save**.

For more information about storage inventories, see [Amazon S3 Storage Inventory](#) in the *Amazon Simple Storage Service Developer Guide*.

How Do I Configure Request Metrics for an S3 Bucket?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.



Announcement: Object Tagging and new Storage Management features available in new console


[Opt In](#) to try object tagging and storage management.

There are two types of CloudWatch metrics for Amazon S3: storage metrics and request metrics. Storage metrics are reported once per day and are provided to all customers at no additional cost. Request metrics are available at 1-minute intervals after some latency to process, and metrics are billed at the standard CloudWatch rate. To get request metrics, you must opt into them by configuring them in the console or with the Amazon S3 API.


For more conceptual information about CloudWatch metrics for Amazon S3, see [Monitoring Metrics with Amazon CloudWatch](#) in the *Amazon Simple Storage Service Developer Guide*.

To configure request metrics on a bucket


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that has the objects you want to get request metrics for.




3. Choose the **Management** tab, and then choose **Metrics**.




4. Choose **Requests**.



5. From the name of your bucket in the left-side pane, choose the edit icon.



6. Choose the **Request metrics** check box. This also enables Data Transfer metrics.



7. Choose **Save**.

You have now created a metrics configuration for all the objects in an Amazon S3 bucket. About 15 minutes after CloudWatch begins tracking these request metrics, you can see graphs for the metrics in both the Amazon S3 or CloudWatch consoles. You can also define a filter so the metrics are only collected and reported on a subset of objects in the bucket. For more information, see [How Do I Configure a Request Metrics Filter?](#) (p. 68).

How Do I Configure a Request Metrics Filter?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.



Announcement: Object Tagging and new Storage Management features available in new console


[Opt In](#) to try object tagging and storage management.

There are two types of CloudWatch metrics for Amazon S3: storage metrics and request metrics. Storage metrics are reported once per day and are provided to all customers at no additional cost. Request metrics are available at 1 minute intervals after some latency to process, and metrics are billed at the standard CloudWatch rate. To get request metrics, you must opt into them by configuring them in the console or with the Amazon S3 API.


For more conceptual information about CloudWatch metrics for Amazon S3, see [Monitoring Metrics with Amazon CloudWatch](#) in the *Amazon Simple Storage Service Developer Guide*.

To filter request metrics on a subset of objects in a bucket


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that has the objects you want to get request metrics for.




3. Choose the **Management** tab, and then choose **Metrics**.




4. Choose **Requests**.



5. From **Filters** in the left-side pane, choose **Add**.




6. Provide a name for this metrics configuration.



The screenshot shows the 'Metrics' section of the Amazon S3 console. At the top, there is a search bar labeled 'Search for filter/prefix/tag'. Below it, there is a list of filters, currently empty, with a '+ Add' button. The 'Add filter' dialog is open, showing the 'Filter name' field with the text 'Monthly Release'. A red arrow points to this field. Below the 'Filter name' field is the 'Prefix / tags that you want to monitor' section, which contains a text input field with the placeholder 'Type to add prefix/tag filter'. At the bottom of the dialog are 'Save' and 'Cancel' buttons.

7. Provide one or more prefixes or tags, separated by commas, in **Prefix /tags that you want to monitor**. From the drop down, select whether the value you provided is a tag or a prefix.



The screenshot shows the 'Metrics' section of the Amazon S3 console. At the top, there is a search bar labeled 'Search for filter/prefix/tag'. Below it, there is a list of filters, currently empty, with a '+ Add' button. The 'Add filter' dialog is open, showing the 'Filter name' field with the text 'Monthly Release'. Below the 'Filter name' field is the 'Prefix / tags that you want to monitor' section. This section contains a dropdown menu with the selected value 'prefix music' and a text input field with the text 'music'. A red arrow points to the 'music' text. At the bottom of the dialog are 'Save' and 'Cancel' buttons.

8. Choose **Save**.

You have now created a metrics configuration for request metrics on a subset of the objects in an Amazon S3 bucket. About 15 minutes after CloudWatch begins tracking these request metrics, you can see graphs for the metrics in both the Amazon S3 or CloudWatch consoles. You can also request metrics at the bucket level. For information, see [How Do I Configure Request Metrics for an S3 Bucket? \(p. 66\)](#)

Setting Permissions

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.



Announcement: Object Tagging and new Storage Management features available in new console

[Opt In](#) to try object tagging and storage management.

The topics in this section explain how to use the Amazon S3 console to set bucket and object permissions.

Bucket permissions specify which users are allowed access to the objects in a bucket and what permissions you have granted them. For example, one user might have only read permission while another might have read and write permissions.

Bucket and object permissions are completely independent. An object does not inherit the permissions from its bucket. For example, if you create a bucket and grant write access to another user, you will not be able to access that user's objects unless the user explicitly grants you access.

Topics

- [How Do I Set Bucket Permissions? \(p. 72\)](#)
- [How Do I Set Permissions on an Object? \(p. 73\)](#)

How Do I Set Bucket Permissions?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.




Announcement: Object Tagging and new Storage Management features available in new console

[Opt In](#) to try object tagging and storage management.


Bucket permissions specify which users are allowed access to the objects in an S3 bucket and which permissions they have. For example, one user might have only read permission, while another might have read and write permissions.

To set permissions for an S3 bucket

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that you want to set permissions for.



3. Choose **Permissions**.



Then choose one of following:

- To grant permissions to a person or group, choose **Access Control Lists** (ACLs). This is the default. For more information, see [Access Control List \(ACL\) Overview](#) in the *Amazon Simple Storage Service Developer Guide*.
- To create or edit a bucket policy, choose **Bucket Policy**. In the **Bucket policy editor**, paste your bucket policy into the text box. For help generating a policy, use the **Policy Generator**. For examples of Amazon S3 bucket policies, see [Bucket Policy Examples](#) in the *Amazon Simple Storage Service Developer Guide*.
- To create a cross-origin resource sharing (CORS) configuration, choose **CORS configuration**. For more information, see [Cross-Origin Resource Sharing \(CORS\)](#) in the *Amazon Simple Storage Service Developer Guide*.

How Do I Set Permissions on an Object?

If you are in the old Amazon S3 console, to use the new console, choose **Opt In** in the following box, which appears on the old Amazon S3 console home page. Note that **Opt In** is not available in all Regions.




Announcement: Object Tagging and new Storage Management features available in new console

[Opt In](#) to try object tagging and storage management.

This section explains how to use the console to edit AWS account permissions for an object. In this topic, each permission you grant adds an entry in the access control list (ACL) associated with the object. You also can grant permissions to other AWS accounts or built-in groups. By default, the owner has full permissions. For more information, see [Access Control List \(ACL\) Overview](#) in the *Amazon Simple Storage Service Developer Guide*.

To set permissions for an object

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Bucket name** list, choose the name of the bucket that contains the object.
3. In the **Name** list, choose the name of the object for which you want to set permissions.



4. Choose **Permissions**.

