
Amazon Simple Storage Service

Console User Guide

API Version 2006-03-01



Amazon Simple Storage Service: Console User Guide

Copyright © 2015 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

The following are trademarks of Amazon Web Services, Inc.: Amazon, Amazon Web Services Design, AWS, Amazon CloudFront, AWS CloudTrail, AWS CodeDeploy, Amazon Cognito, Amazon DevPay, DynamoDB, ElastiCache, Amazon EC2, Amazon Elastic Compute Cloud, Amazon Glacier, Amazon Kinesis, Kindle, Kindle Fire, AWS Marketplace Design, Mechanical Turk, Amazon Redshift, Amazon Route 53, Amazon S3, Amazon VPC, and Amazon WorkDocs. In addition, Amazon.com graphics, logos, page headers, button icons, scripts, and service names are trademarks, or trade dress of Amazon in the U.S. and/or other countries. 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 Amazon S3	1
Resources and Operations	1
Resource Owner	2
Resource Operations	2
About the Console	2
Support for Viewing Data	3
Support for Properties	3
Support for Folders	4
Support for Moving Data	5
Intuitive UI	6
Easy to Switch to Other AWS Consoles	7
About the Amazon S3 Documentation	7
Working with Buckets	9
Creating a Bucket	9
Deleting or Emptying a Bucket	12
Browsing the Objects in Your Bucket	14
Editing Bucket Permissions	16
Configuring a Bucket for Website Hosting	19
Managing Bucket Logging	20
Enabling Events	21
Set Up a Destination to Receive the Event Notifications	22
Enable Event Notifications	23
Editing and Deleting Event Notifications Configurations	26
Enabling Bucket Versioning	27
Managing Lifecycle Configuration	28
Lifecycle Configuration for a Bucket without Versioning	29
Lifecycle Configuration for a Bucket with Versioning	32
Maintaining Lifecycle Configuration Rules	38
Managing Cost Allocation Tagging	41
Managing Cross-Region Replication	42
Enable Cross-Region Replication	42
Disable or delete Cross-Region Replication	44
Working with Objects	49
Uploading Objects	49
Using the Enhanced Uploader	55
Editing Object Properties	57
Editing Object Details	57
Permissions	60
Metadata	63
Opening an Object	64
Downloading an Object	65
Copying an Object	66
Renaming an Object	67
Deleting an Object	68
Deleting Objects by using Lifecycle Configuration Management	69
Restoring an Object	69
Managing Objects in a Versioning-Enabled Bucket	72
Uploading an Object	73
Updating Object Properties	73
Deleting Objects from a Versioning-Enabled Bucket	73
Working with Folders	75
Public Folders	75
Creating a Folder	76
Deleting a Folder	76
Resources	78

Document History 80
AWS Glossary 84

Welcome to Amazon S3

This is the *Amazon Simple Storage Service Console User Guide*.

The Amazon S3 console is one of the interfaces available to help you work with Amazon S3. The console enables you to perform Amazon S3 tasks without writing any code. This section first introduces Amazon S3 resources and operations and then explains how the console is logically organized to support these operations. The section also introduces console-specific concepts such as folders, properties, and other features that help you easily upload files and folders, move objects around, and manage objects by creating folders. We recommend that you read the following sections:

- [About Amazon S3 Resources and Operations \(p. 1\)](#)
- [About the Amazon S3 Console \(p. 2\)](#)
- [About the Amazon S3 Documentation \(p. 7\)](#)

For information on Amazon S3 features, pricing, and to see the FAQ, go to the [Amazon S3 product page](#).

About Amazon S3 Resources and Operations

Amazon S3 is storage for the Internet. You can think of Amazon S3 as a collection of resources and operations. Buckets and objects are the primary resources. Amazon S3 provides APIs for you to create buckets and upload objects. In addition, there are other resources, many of which store bucket and object specific configuration information. These are referred to as subresources. For example, the following are some of the bucket subresources:

- *lifecycle* – You can define lifecycle configuration rules for objects that have a well-defined lifecycle. For example, archive objects one year after creation, or delete an object 10 years after creation. The *lifecycle* subresource stores the lifecycle configuration rules. For more information, go to [Object Lifecycle Management](#).
- *website* – You can host a static website on Amazon S3. To host your static website, you configure your bucket for website hosting. The *website* subresource stores the website configuration information. For more information, go to [Hosting a Static Website on Amazon S3](#).
- *versioning* – Versioning provides protection from accidental overwrites and deletes. We recommend versioning as a best practice to prevent objects from being deleted or overwritten by mistake. The *versioning* subresource stores versioning configuration information. For more information, go to [Using Versioning](#).

- *policy* and *ACL* (access control list) – These subresources store access permission information. By default, all your resources are private. You as the resource owner must grant permissions for others to access these resources. For more information, see [Resource Owner \(p. 2\)](#).

There are also subresources associated with objects. For example, Amazon S3 provides an *ACL* subresource that helps you manage object-level permissions.

Resource Owner

By default, all Amazon S3 resources are private. Only a resource owner can access the resource. The resource owner refers to the AWS account that creates the resource. The resource owner can optionally grant others permission to access the resources. These can be other AWS accounts, IAM users in an AWS account, or applications that get permissions via the IAM roles. For information about AWS accounts and IAM users, go to [What is IAM?](#) in *IAM User Guide*. For more information about permissions, go to [Managing Access Permissions to Your Amazon S3 Resources](#) in the *Amazon Simple Storage Service Developer Guide*.

Resource Operations

To help you work with buckets, objects, and related subresources, Amazon S3 provides a set of operations. You have the following options to work with Amazon S3:

- Use the Amazon S3 console to perform operations without writing any code.
- Use the AWS SDKs that provide wrapper libraries for Java, .NET, Python, PHP, and other languages. For more information about the available SDKs, go to [Sample Code and Libraries](#).
- Use the AWS Command Line Interface (CLI) to manage Amazon S3 objects by using a command line user interface. For more information about the AWS CLI, go to [AWS Command Line Interface](#).
- Both the console and the AWS SDK libraries internally make the Amazon S3 REST API call described in the API reference. If you need to, you can also write code to make the REST API calls directly from your application.

For a list of Amazon S3 operations go to, [Operations on Buckets](#) and [Operations on Objects](#) in the *Amazon Simple Storage Service API Reference*.

About the Amazon S3 Console

Using the Amazon S3 console, you can create and manage the resources discussed in the preceding section. The console supports additional features that are not natively supported by Amazon S3 (for example, the concept of folders). These additional features are designed to help you manage your resources. Some of the console highlights discussed in this section are:

- Support for viewing data
- Support for properties
- Support for folders

Note

The Amazon S3 data model does not natively support the concept of folders, nor does it provide any APIs for folder-level operations. But the Amazon S3 console supports folders to help you organize your data.

- Support for moving data around
- Visibility into object properties
- Ability to act on groups of data

- Intuitive UI that abstracts the underlying API calls
- Easy to switch to other consoles that are part of the AWS Management Console

Note

You might want to sign into the Amazon S3 console at <https://console.aws.amazon.com/s3> as you read the remainder of this section. Your Session Credentials will keep you logged into the AWS Management Console for approximately twelve hours.

Support for Viewing Data

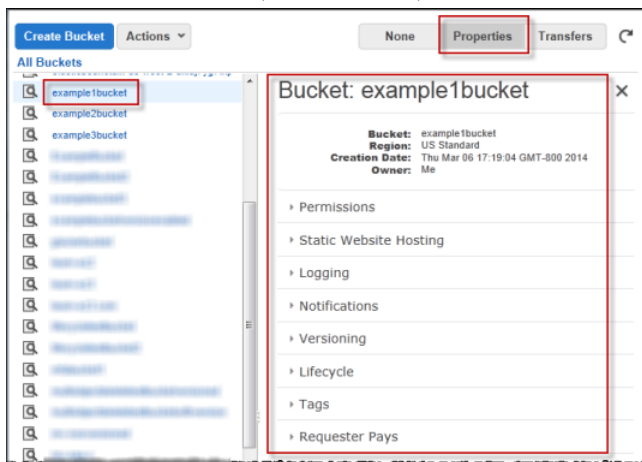
The Amazon S3 console provides a view of your Amazon S3 data. It lists your buckets and the objects in each bucket. When you create a bucket you specify an AWS region where you want the bucket to reside. Amazon S3 bucket names are globally unique and the console lists all buckets, regardless of the region in which the bucket is stored. So the Amazon S3 console does not require any region selection to list buckets and objects.




Support for Properties

The console supports the concept of properties. Using the properties abstraction, the Amazon S3 console shows the metadata and subresources associated with the primary resource (bucket or object).

If you click on a bucket name and then click **Properties**, you will get a list of bucket properties. These properties include bucket subresources, described in the preceding section, and metadata information such as resource name, creation date, and owner.



If you click on an object name and then click **Properties**, the console displays a list of object properties.

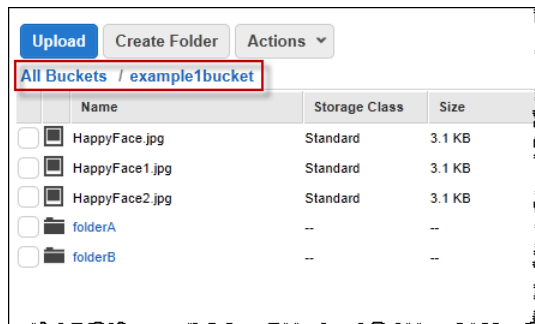


The **Link** property shows the object URL, a valid resource address. But the URL does not contain authentication information. If you click the link Amazon S3 will deny access to the object unless you make the object public (by default all objects are private). For information about downloading, see [Downloading an Object](#) (p. 65).

Support for Folders

The concept of folders is unique to the console. Amazon S3 uses buckets and objects, but the service does not natively support folders, nor does it provide any API to work with folders.

To help you organize your data, however, the Amazon S3 console supports the concept of folders. You can create folders to group your objects. The following screenshot shows a bucket (`example1bucket`) that contains two folders, `folderA` and `folderB`.




Important


In Amazon S3, you create buckets and store objects. The service does not support any hierarchy that you see in a typical file system.

The console uses the object key names to derive the folder hierarchy. It uses the "/" character in the key name to infer hierarchy, as the following examples show:

- If you have three objects—`logs/date1.txt`, `logs/date2.txt`, and `logs/date3.txt`—the console shows a folder named `logs`. If you open the folder, you see three objects: `date1.txt`, `date2.txt`, and `date3.txt`.




- You can nest folders in the console. For example, if you have an object named `photos/2013/example.jpg`, the console shows you a folder named `photos` containing the folder `2013`, and the folder `2013` contains the object `example.jpg`.



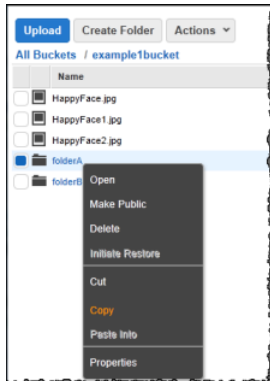
- If you upload an object with key name `myPhoto.jpg`, there is no "/" delimiter in the key name, and the console shows the object at the root level of the bucket.

The console also supports following folder-level actions. For example, for the existing objects in a folder you can request Amazon S3 to store them encrypted using server-side encryption, or change the storage class for those objects. These actions apply only once to the existing objects in the folder. Amazon S3 console does not save this configuration and will not apply to any new objects you add to the bucket.




Support for Moving Data


Using the Amazon S3 console, you can easily move data around. For example, to copy objects between buckets and folders right-click on an object inside the source bucket or folder and then click **Copy**.




Then, right-click on the target bucket or folder and click **Paste Into** to make a copy.



The console also enables you to act on group of data. For example, you can select and copy multiple objects or folders.



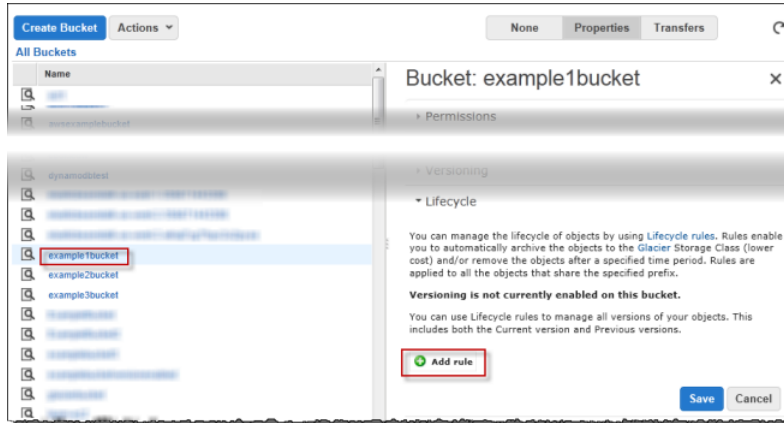
When uploading, you can upload an individual object or a folder. To upload click **Actions** and then click **Upload**. Then you can click **Add Files** or you can drag and drop files and folders to the **Drag and Drop files and folders to upload here.** area of the **Upload** dialog as shown in the following screenshot. Drag and drop does not work a with all Internet browsers.



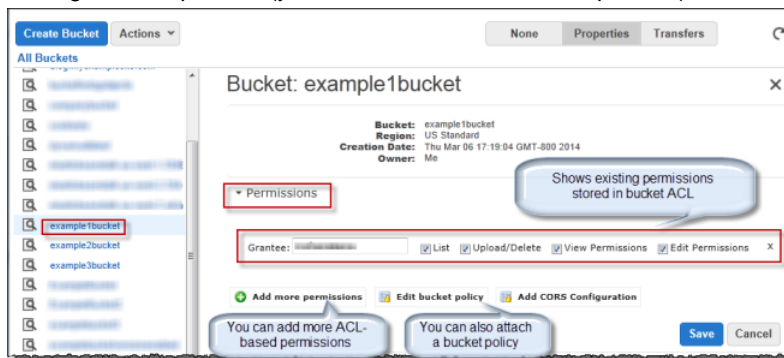
Intuitive UI

The Amazon S3 console provides an intuitive UI for some of the API calls. For example:

- You can set lifecycle policies by adding rules using the console UI.




- Manage bucket policies (you can add or delete bucket policies) and other (ACL-based) permissions.



- You can also configure your bucket as a website.

Easy to Switch to Other AWS Consoles

From the Amazon S3 console, you can switch to other AWS consoles to manage your other AWS resources, such as the IAM console to manage users in your account.



About the Amazon S3 Documentation

Amazon S3 is documented in the following guides.

Amazon S3 Guide	Description
Developer Guide	This is the primary Amazon S3 guide. It provides conceptual information for all Amazon S3 features and provides working examples using some of the AWS SDKs.

Amazon Simple Storage Service Console User Guide
About the Amazon S3 Documentation

Amazon S3 Guide	Description
API Reference	This guide documents the REST API operations that Amazon S3 supports. When sending requests to Amazon S3 using the REST API, you will need to sign the requests. This guide explains signing and authentication.
Getting Started Guide	This guide provides Amazon S3 console–based introductory experience of working with Amazon S3.
Console User Guide (this guide)	This guide provides detailed procedures for console-based operations. The help links in the console link to procedural topics in this guide.

Also, the Amazon S3 product detail page provides pricing and additional product information. You can also engage with the Amazon S3 community in the discussion forum.

Information	Relevant Sections
General product overview and pricing	Amazon Simple Storage Service (Amazon S3)
Discussion forum	Amazon S3 Forum

Working with Buckets

Topics

- [Creating a Bucket \(p. 9\)](#)
- [Deleting or Emptying an Amazon S3 Bucket \(p. 12\)](#)
- [Browsing the Objects in Your Bucket \(p. 14\)](#)
- [Editing Bucket Permissions \(p. 16\)](#)
- [Configuring a Bucket for Website Hosting \(p. 19\)](#)
- [Managing Bucket Logging \(p. 20\)](#)
- [Enabling Event Notifications \(p. 21\)](#)
- [Enabling Bucket Versioning \(p. 27\)](#)
- [Managing Lifecycle Configuration \(p. 28\)](#)
- [Managing Cost Allocation Tagging \(p. 41\)](#)
- [Managing Cross-Region Replication \(p. 42\)](#)

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. Buckets have properties, such as access permissions and versioning status, and you can specify the region where you want them to reside.

This section explains how to use the Amazon S3 console to create, delete, and manage buckets.

As you create buckets, upload objects, and perform various other operations, usage reports are available that you might find useful. For more information, go to [Billing and Cost Management Console](#).

Creating a Bucket

Before you can upload data into Amazon S3, you must create a bucket to store the data in. Buckets have configuration properties, including their geographical region, who has access to the objects in the bucket, and other metadata, such as the storage class of the objects in the bucket.

The console enables you to use folders, which you can store objects in. Folders, like objects, must reside in a bucket. For more information about using folders, see [Working With Folders \(p. 75\)](#).

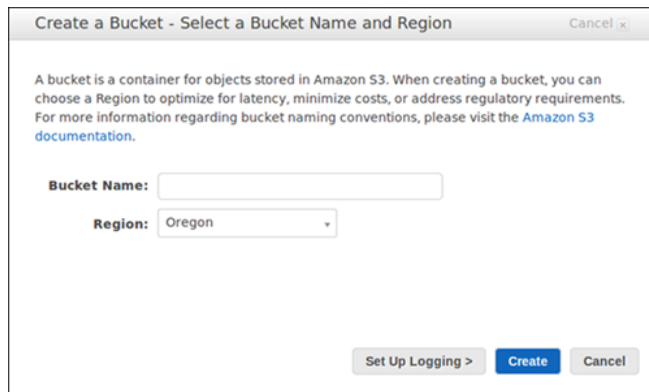
Use the following procedure to create a bucket.

Note

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.

To create a bucket

1. Sign into the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3>.
2. Click **Create Bucket**.
3. In the Create Bucket dialog box, in the Bucket Name box, type a name for your bucket.



The name that you choose must be unique across all existing bucket names in Amazon S3. One way to help ensure uniqueness is to prefix your bucket names with the name of your organization.

The bucket name is visible in the URL that points to the objects that you're going to put in your bucket. For that reason, choose a bucket name that reflects the objects in the bucket.

To ensure a single, consistent naming approach for Amazon S3 buckets across regions and to ensure bucket names conform to DNS naming conventions, bucket names must comply with the following requirements.

- Can contain lowercase letters, numbers, periods (.), and hyphens (-).
- Must start with a number or letter.
- Must be between 3 and 63 characters long.
- Must not be formatted as an IP address (e.g., 192.168.5.4).
- Must not contain underscores (_).
- Must not end with a hyphen.
- Cannot contain two, adjacent periods.
- Cannot contain dashes next to periods (e.g., my-.bucket.com and my.-bucket are invalid).

Note

If you want to use your S3 bucket as an origin for an Amazon CloudFront distribution, the requirements for naming S3 buckets are more restrictive. For more information, see the `DNSName` element in the "S3Origin Child Elements" table in the [DistributionConfig Complex Type](#) section of the *Amazon CloudFront API Reference*.

To take advantage of Amazon S3's CNAME support, you should name your bucket the same as your website's base address (e.g. `www.mysite.com`). For more information about CNAME, go to [Virtual Hosting](#) in the Amazon Simple Storage Service Developer Guide.

Note

Once you create a bucket, you cannot change the name of it. Make sure the bucket name you choose is appropriate.

4. In the **Region** box, click the region where you want the bucket to reside.

You should choose a region close to you to optimize latency, minimize costs, or to address regulatory requirements. Objects stored in a region never leave that region unless you explicitly transfer them to another region. For more information about regions, go to [Regions and Endpoints](#) in the *Amazon Web Services General Reference*.

In the next step, you have the opportunity to set up logging. Server access logging provides detailed records for the requests made against your bucket. An access log record contains details about the request, such as the request type, the resources specified in the request worked, and the time and date the request was processed. Server access logs are useful for many applications because they give bucket owners insight into the nature of requests made by clients not under their control. Amazon S3 delivers access logs to your bucket. By default, Amazon S3 does not collect server access logs.

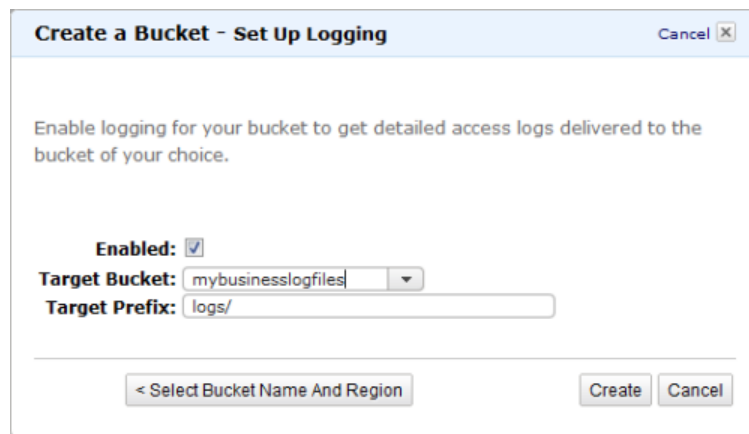
5. Do one of the following.

To...	Do this...
Create a bucket without setting up logging	Click Create
Set up server access logging for the bucket you're creating	Click Set Up Logging

Note

There is no extra charge for enabling server access logging on an Amazon S3 bucket. However, any log files the system delivers to you will accrue the usual charges for storage. (You can delete log files at any time.) We do not assess data transfer charges for delivering log files to your bucket, but we do charge the normal data transfer rate for accessing the log files. For more information, go to [Amazon S3 Pricing](#).

6. If you clicked **Set Up Logging** in the **Create a Bucket - Set Up Logging** dialog box, do the following:

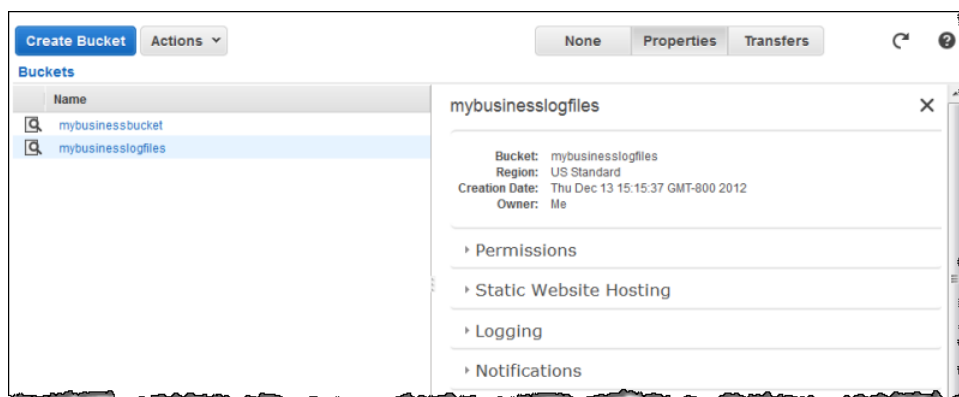


- a. Select the **Enabled** check box.
- b. In the **Target Bucket** box, select the bucket where you want the log files stored.
- c. (Optional) In **Target Prefix** box, specify a prefix for the name of the log files.

Amazon S3 adds the prefix to the log file names when storing them in your bucket. For example, if you specify the prefix "logs/," all logs stored in the target bucket are prefixed with `logs/`, so, all the logs will be stored in the `logs` folder.

7. Click **Create**.

If Amazon S3 successfully creates your bucket, the console displays your empty bucket.



Deleting or Emptying an Amazon S3 Bucket

This section explains how to use the console to delete or empty an Amazon S3 bucket.

You can delete a bucket and all the objects contained in the bucket. For information on the limitations for deleting a bucket, see [Deleting/Emptying a Bucket](#) in the *Amazon Simple Storage Service Developer Guide*.


Note

When you delete a bucket with versioning enabled, all versions of all the objects in the bucket are deleted. For more information about managing objects when versioning is enabled, see [Managing Objects in a Versioning-Enabled Bucket \(p. 72\)](#).

When you delete a bucket, there may be a delay of up to one hour before the bucket name is available for reuse in a new region or by a new bucket owner. If you re-create the bucket in the same region or with the same bucket owner, there is no delay.

To delete a bucket

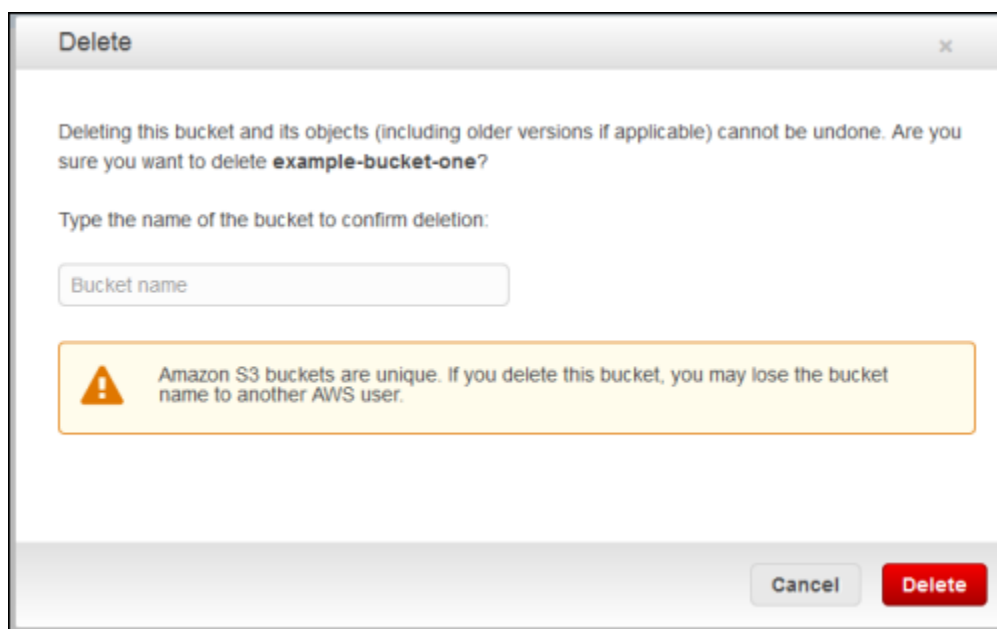
1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Right-click the bucket that you want to delete, and then click **Delete Bucket**.



Tip

Optionally, to get this menu you can click the bucket and then click **Actions**, which is near the top of the console window next to **Create Bucket**.

3. When a confirmation message appears, enter the bucket name and then click **Delete**.



You can empty a bucket, which deletes all the objects in the bucket without deleting the bucket. For information on the limitations for emptying a bucket, see [Deleting/Emptying a Bucket](#) in the *Amazon Simple Storage Service Developer Guide*.

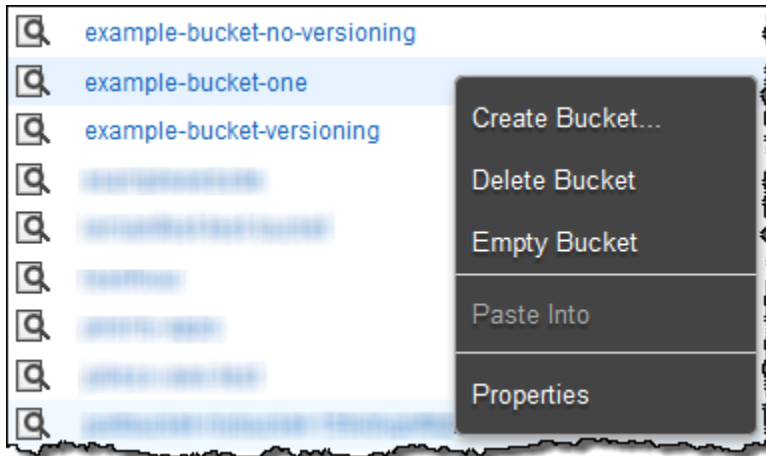
Note

When you empty a bucket with versioning enabled, all versions of all the objects in the bucket are deleted. For more information about managing objects when versioning is enabled, see [Managing Objects in a Versioning-Enabled Bucket](#) (p. 72).

To empty a bucket

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

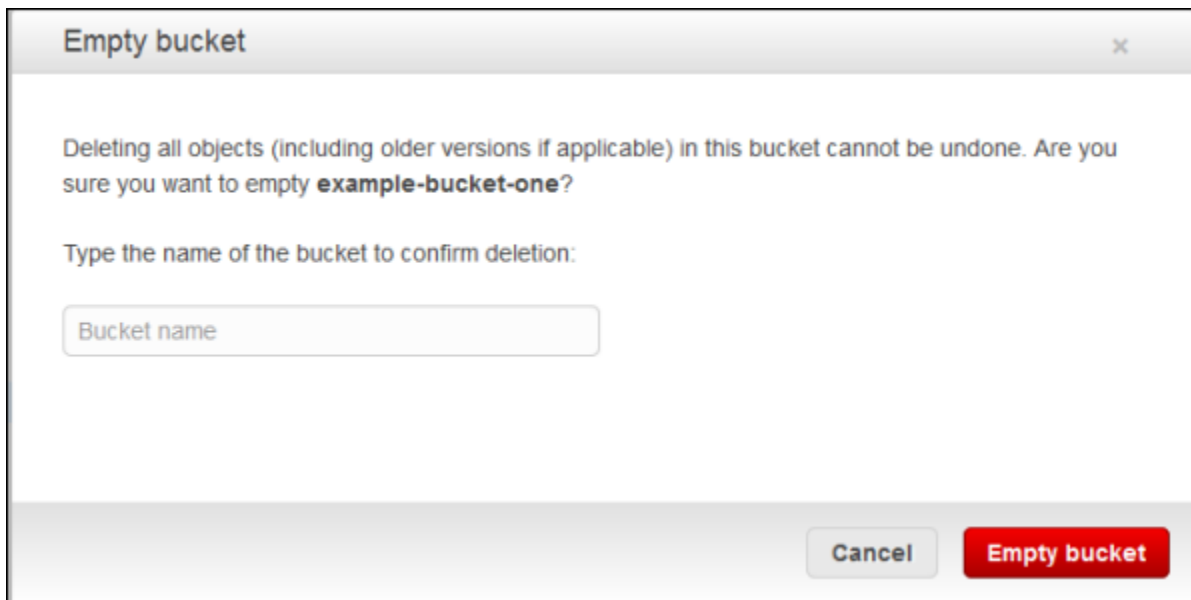
2. Right-click the bucket that you want to empty, and then click **Empty Bucket**.



Tip

Optionally, to get this menu you can click the bucket and then click **Actions**, which is near the top of the console window next to **Create Bucket**.

3. When a confirmation message appears, enter the bucket name and then click **Empty bucket**.



Browsing the Objects in Your Bucket

This section describes how to use the console to browse and display the objects and folders in your bucket.

To list the objects in a bucket

- Click the bucket whose objects you want to display.

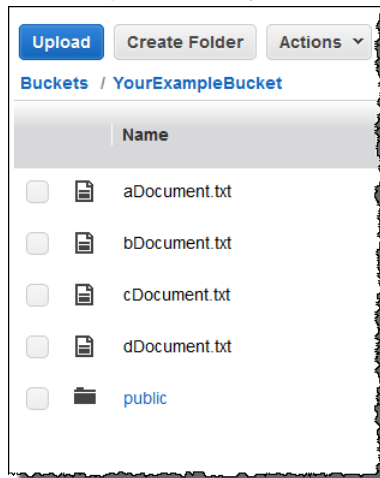
The Objects and Folders list displays the objects and folders in the selected bucket.

Note

If you have a large number of objects in a bucket, you can scroll down to the bottom of the Objects and Folders panel. When the scroll bar reaches the bottom of the list, the AWS Management Console automatically retrieves the next set of keys in your bucket, refreshes the view, and shows them in the console view.

When you click a bucket name, the console lists all the objects in the bucket in alphanumeric order. However, if your bucket contains large number of objects, scrolling down the long list to search for an object can be cumbersome. The jump feature enables you to type a string, and the console skips ahead to the specific object in the object list. If there are no objects whose key name match the specified string, the console jumps to the next object in the list in alphanumeric order.

For example, assume you have a bucket (ExampleBucket) with the following objects.

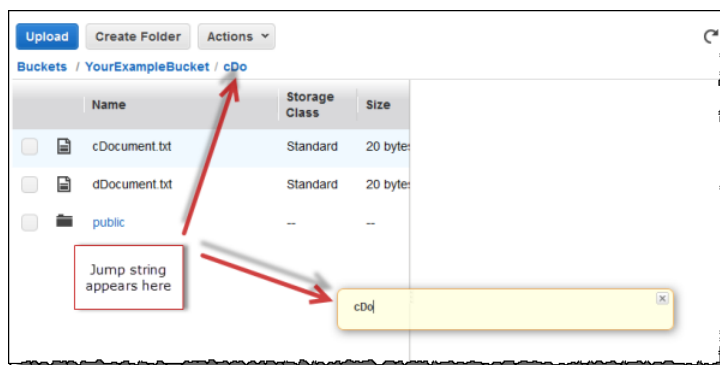


To jump to an object in your list

1. Click the bucket name to display its objects.
2. Begin typing an object key name.

As you begin typing characters, for example, a letter **c**, the console performs the following actions:

- Opens a *jump* dialog box showing the character you typed.
- Skips ahead to the first object whose key name starts with the string you typed.
- Appends the jump string to the existing navigation breadcrumb.



3. While the jump dialog box is visible, do one of the following:
 - **Press Enter** – This closes the jump dialog box. The jump results (such as the **cDo** shown in the preceding example screenshot) remain.
 - **Press Esc** – This cancels the jump operation and the *jump* dialog box closes.

Tip

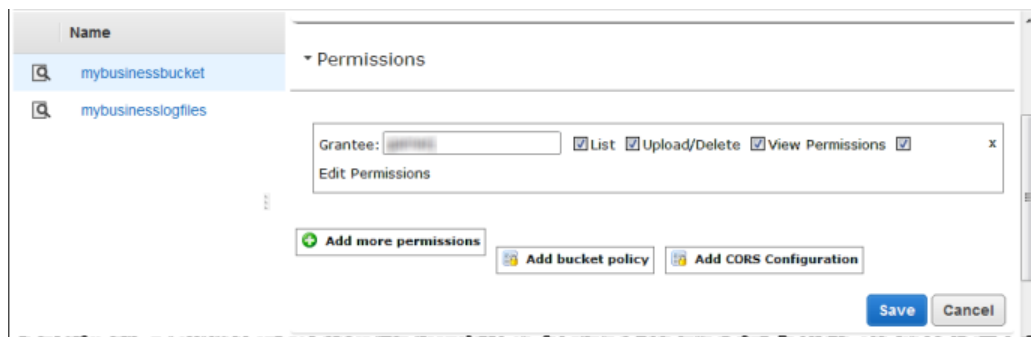
To return to the top of the list, press the **Backspace** key.

Editing Bucket Permissions

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

To edit bucket permissions



1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the Buckets list, click the bucket whose properties you want to view.




3. Click **Permissions**, and then do any of the following:

Amazon Simple Storage Service Console User Guide

Editing Bucket Permissions

To...	Do this...
Add permissions for a person or group	<p>a. Click Add more permissions.</p> <p>b. In the Grantee box of the new line that appears, add the name of the person or group for which you want to set permissions. The name can be the email address associated with an AWS account, a canonical ID, or one of the predefined Amazon S3 groups. For a list of predefined Amazon S3 Groups, go to Who is a Grantee in the <i>Amazon Simple Storage Service Developer Guide</i>. You can add as many as 100 grantees.</p> <p>c. Select the check boxes next to the permissions you want to grant.</p>  <p>The screenshot shows the 'Permissions' section of the console. It includes a 'Grantee' input field, a list of permissions with checkboxes (List, Upload/Delete, View Permissions), and an 'x' icon to remove the grantee. Below this, there are three buttons: 'Add more permissions' (highlighted with a red box), 'Add bucket policy', and 'Add CORS Configuration'. 'Save' and 'Cancel' buttons are at the bottom right.</p>
Remove a person or group from the permission list	Click the "x" on the line of the grantee you want to remove.
Add a bucket policy	<p>a. Click Add bucket policy.</p> <p>b. In the Bucket Policy Editor, paste your bucket policy into the box provided.</p> <p>For help in generating a policy, you can use the AWS Policy Generator . For examples of Amazon S3 bucket policies, see Example Cases for Amazon S3 Bucket Policies in the <i>Amazon Simple Storage Service Developer Guide</i>.</p> <p>c. Click Save.</p>  <p>The screenshot shows the same 'Permissions' section as above. The 'Add bucket policy' button is highlighted with a red box. The 'Add more permissions' button is also visible but not highlighted.</p>

To...	Do this...
Add a Cross-Origin Resource Sharing (CORS) configuration	<p>a. Click Add CORS Configuration. In the CORS Configuration Editor, paste your CORS configuration into the field provided, and then click Save. For information about CORS configuration, see Enabling Cross-Origin Resource Sharing in the <i>Amazon Simple Storage Service Developer Guide</i>.</p> 

There are built-in groups that you can choose from the **Grantee** box:

- **Everyone**—Use this group to grant anonymous access.
- **Authenticated Users**—This group consists of any user that has an Amazon AWS Account. When you grant the Authenticated User group permission, any valid signed request can perform the appropriate action. The request can be signed by either an AWS Account or IAM User.
- **Log Delivery**—This group grants write access to your bucket when the bucket is used to store server access logs. For more information, see [Managing Bucket Logging](#).
- **Me**—This group refers to your AWS root account, and not an IAM user.

You can grant permission to an AWS account by entering the accounts canonical user ID or email address in the **Grantee** field. The email address must be the same one they used when signing up for an AWS account. You can grant a grantee any of the following permissions:

- **Open/Download**—Enables the account to access the object when they are logged in
- **View Permissions**—Can view the permissions associated with the object
- **Edit Permissions**—Can edit the permissions associated with the object

For more information about predefined Amazon S3 Groups, go to [Who is a Grantee](#) in the *Amazon Simple Storage Service Developer Guide*.

You can grant access to an account by using the email address that the user entered when signing up for an AWS account. You can grant an account any of the following permissions:

- **List** – Allows the grantee to view a list of the objects in the bucket.
- **Upload/Delete** – Allows the grantee to access the object when they logged in.
- **View Permissions** – Allows the grantee to view the permissions associated with the object.
- **Edit Permissions** – Allows the grantee to edit the permissions associated with the object.

Caution

We highly recommend against granting the Everyone group **Upload/Delete** permission. Doing so will allow anyone to store objects in your bucket, for which you will be billed, and allows others to delete objects that you may want to keep.

4. Click **Save**.

Configuring a Bucket for Website Hosting

You can host static websites on Amazon S3. For conceptual information, go to [Hosting Websites on Amazon S3](#) in the *Amazon Simple Storage Service Developer Guide*. This section explains how to use the Amazon S3 console to configure a bucket as a website.

To manage a bucket's website configuration

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the Buckets pane, click the bucket that you want to configure.
3. In the result pane, click **Static Website Hosting**.

Bucket: businessbucketlogfiles
Region: Oregon
Creation Date: Wed Jan 28 09:29:18 GMT-800 2015
Owner: Me

› Permissions

▼ Static Website Hosting

You can host your static website entirely on Amazon S3. Once you enable your bucket for static website hosting, all your content is accessible to web browsers via the Amazon S3 website endpoint for your bucket.

Endpoint: businessbucketlogfiles.s3-website-us-west-2.amazonaws.com

Each bucket serves a website namespace (e.g. "www.example.com"). Requests for your host name (e.g. "example.com" or "www.example.com") can be routed to the contents in your bucket. You can also redirect requests to another host name (e.g. redirect "example.com" to "www.example.com"). See our [walkthrough](#) for how to set up an Amazon S3 static website with your host name.

Do not enable website hosting

Enable website hosting

Redirect all requests to another host name

Save Cancel

4. Do one of the following:
 - To configure a bucket for website hosting, click **Enable website hosting**. In the **Index Document** box, type the name of the index document. Optionally, in the **Error Document** box, you can also provide the name of a custom error document and specify custom rules to redirect requests. For more information, go to [Configure a Bucket for Website Hosting](#) in the *Amazon Simple Storage Service Developer Guide*.
 - To redirect all requests to a different web page, click **Redirect all requests to another host name**. In the **Redirect all requests to** box, type the name of the location where you want requests to be redirected, for example, example.com or http://example.com. If you don't specify the protocol (http, https), the protocol of the original request is used. If you redirect all requests, then any request made to the bucket's website endpoint will be redirected to the specified host name.
5. When the settings are as you want them, click **Save**.
6. Add the following policy to the bucket to grant everyone access to the objects in the bucket. For step-by-step instructions, see [Editing Bucket Permissions](#) (p. 16).

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 sample 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 more information, go to [Permissions Required for Website Access](#) in the *Amazon Simple Storage Service Developer Guide*.

Note

If you click **Do not enable 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.

Managing Bucket Logging

Logging provides a way to get detailed access logs delivered to a bucket you choose. An access log record contains details about the request, such as the request type, the resources specified in the request worked, and the time and date the request was processed. For more information about the contents of a log, see [Server Access Log Format](#) in the *Amazon Simple Storage Service Developer Guide*.

Server access logs are useful for many applications because they give bucket owners insight into the nature of requests made by clients not under their control. By default, Amazon S3 doesn't collect service access logs, but when you enable logging Amazon S3 delivers access logs to your bucket on an hourly basis.

This section describes how to use the console to enable and disable logging for a bucket. You can store logs in the same bucket you enable logging for, or you can store the logs in a different bucket. For more information about bucket logging, see [Accessing Server Logs](#) in the *Amazon Simple Storage Service Developer Guide*.


Note

There is no extra charge for enabling server access logging on an Amazon S3 bucket. However, any log files the system delivers to you will accrue the usual charges for storage. (You can delete the log files at any time.) We do not assess data transfer charges for log file delivery, but we do charge the normal data transfer rate for accessing the log files.

To enable logging 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. Under **All Buckets**, click the bucket for which access requests will be logged.
3. In the Details pane, click **Properties**

4. Under **Logging**, do the following:



- Select the **Enabled** check box
 - In the **Target Bucket** box, click the name of the bucket that will receive the log objects.
 - (optional) To specify a key prefix for log objects, in the **Target Prefix** box, type the prefix that you want.
5. Click **Save**.

To disable logging 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. Under **All Buckets**, click the bucket for which access requests will be logged.
3. In the Details pane, click **Properties** Under **Logging**, clear the **Enabled** check box.
4. Click **Save**.

Enabling Event Notifications

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 and how to use the Amazon S3 API to enable 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 select **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 removed event

You select **ObjectRemoved(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. Select **DeleteMarkerCreated** to trigger event notifications when a delete marker is created for a versioned object. For information about deleting versioned objects, see [Deleting Object Versions](#). For information about object versioning, see [Object Versioning](#) and [Using Versioning](#).

A Reduced Redundancy Storage (RRS) object lost event

Amazon S3 sends a notification message when it detects that an object of 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
- An Amazon Simple Queue Service (Amazon SQS) queue
- An AWS Lambda function

Topics

- [Set Up a Destination to Receive the Event Notifications \(p. 22\)](#)
- [Enable Event Notifications \(p. 23\)](#)
- [Editing and Deleting Event Notifications Configurations \(p. 26\)](#)

Set Up a Destination to Receive the Event Notifications

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

An Amazon SNS topic

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 must have the Amazon Resource Name (ARN) for the Amazon SNS topic.
- You must have a valid Amazon SNS topic subscription. The topic subscribers are notified when a message is published to your Amazon SNS topic.
- You must set up a permissions policy through 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 topic, go to [Working 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 must have the Amazon Resource Name (ARN) for the Amazon SQS topic.
- You must set up a permissions policy through 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

For information about creating a Lambda function, see the [AWS Lambda Developer Guide](#).

Before you can set up the Lambda function as a event notification destination:

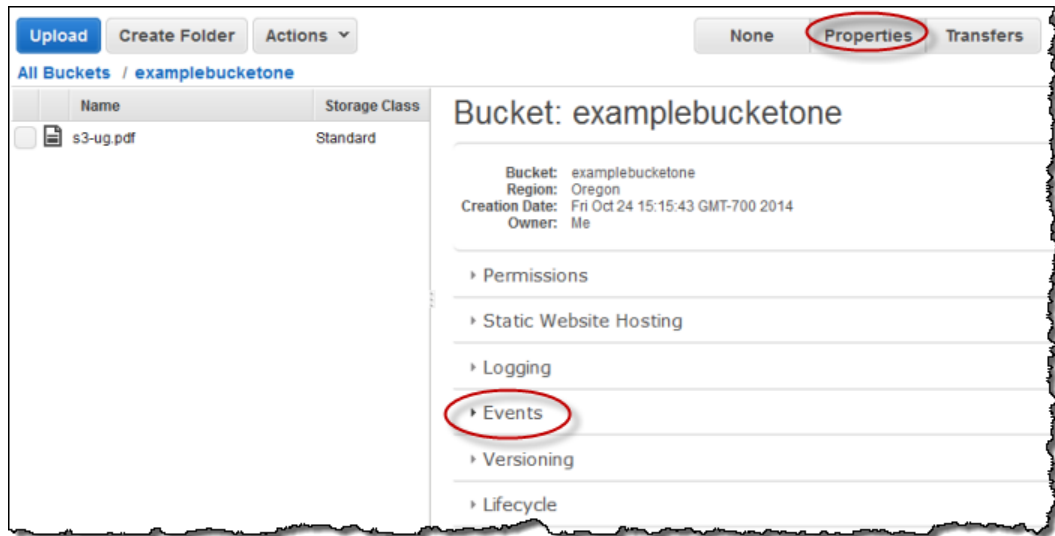
- You must have the ARN of the Lambda function.
- You must have the ARN of the *invocation role* for your Lambda function. For information about creating *invocation role* for Lambda see [Component: Invocation Role](#) in the *AWS Lambda Developer Guide*.

Enable Event Notifications

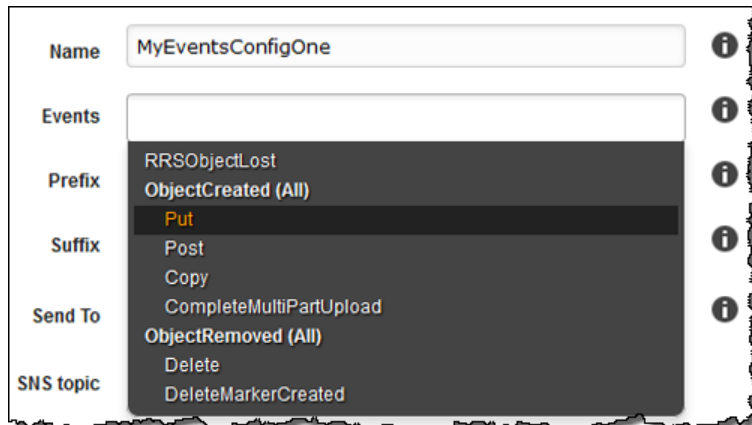
The following procedure shows you how to enable event notifications for a bucket.

To enable bucket event notifications

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Buckets** list, click the bucket whose events you want to configure, click **Properties** and then click **Events**.



3. In the **Name** box, type a descriptive name for your event configuration. If you do not enter a name, a GUID is autogenerated and used for the name.
4. Click in the **Events** box and select the type or types of events that you want to send notifications to a destination when an event occurs.



5. Select **ObjectCreated(All)** to enable event notifications for anytime an object is created in the bucket. Or, you can select specific object creation actions to trigger notifications. For example, you could select **Put** and **CompleteMultiPartUpload** to trigger event notifications anytime a new object is put into a bucket and anytime a multipart upload completes. (Optionally, you could select **ObjectRemoved(All)** to enable event notifications for anytime an object is deleted in the bucket. Or, you could select **Delete** or **DeleteMarkeCreated** to trigger notifications for specific types of object deletes.)

You can configure notifications to be filtered by the prefix and/or suffix of the name of objects. For example, you can set up a configuration so that you are sent a notification only when files are added to an image folder (objects with the name prefix `images/`). For more information on filtering, see [Configuring Notifications with Object Key Name Filtering](#).

Name: MyEventsConfigOne

Events: Put CompleteMultiPartUpload

Prefix: images/

Suffix: e.g. jpg

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

Send To: SNS topic SQS queue Lambda function

- 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, go to <http://aws.amazon.com/sns/faqs/#10>.

Send To: SNS topic SQS queue Lambda function

SNS topic: Select/Enter SNS topic

S3 must have perm...

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

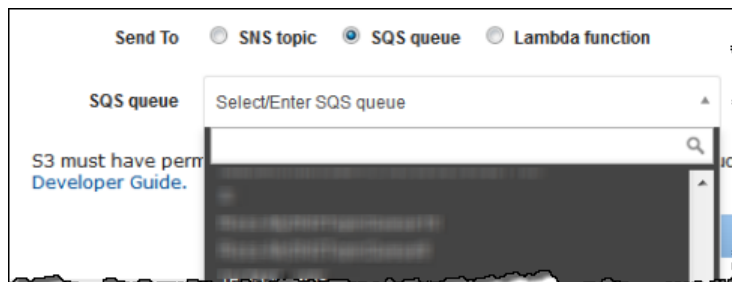
Send To: SNS topic SQS queue Lambda function

SNS topic: Add SNS topic ARN

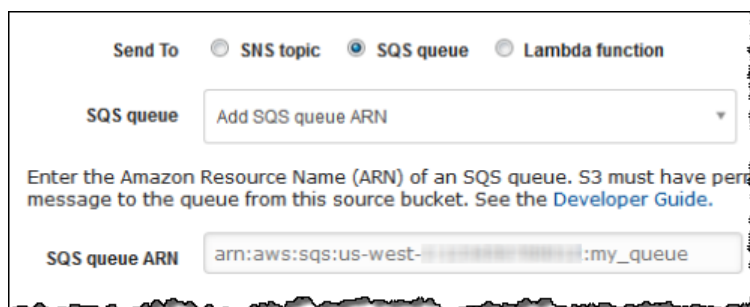
Enter the Amazon Resource Name (ARN) of an SNS topic. S3 must have permis... See the [Developer Guide](#).

SNS topic ARN: arn:aws:sns:us-west-...my_topic

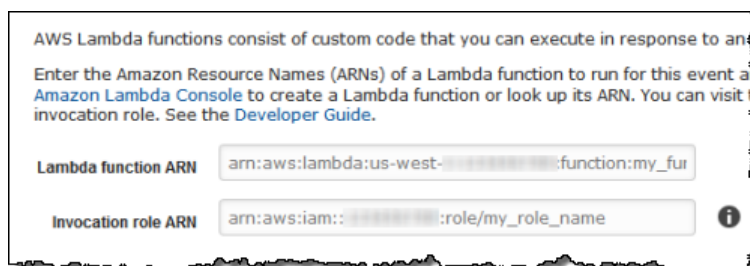
- b. If you select the **SQS queue** destination type.
- i. In the **SQS queue** box, type the name or select from the menu, the name of the Amazon SQS queue that will receive notifications from Amazon S3. For information about Amazon SQS, to [Amazon Simple Queue Service Developer Guide](#).



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



- c. If you select the **Lambda Function** destination type.
 - i. In the **Lambda Function ARN** box, type the ARN of the Lambda function that will receive notifications from Amazon S3.



- ii. Enter the ARN of the IAM Lambda *invocation role* in the **Invocation role ARN** box. For information about Lambda functions and invocation role ARNs, see [AWS Lambda: How it Works](#) in the *AWS Lambda Developer Guide*.


7. Click **Save**. Amazon S3 will send a test message to the event notification destination.

Editing and Deleting Event Notifications Configurations

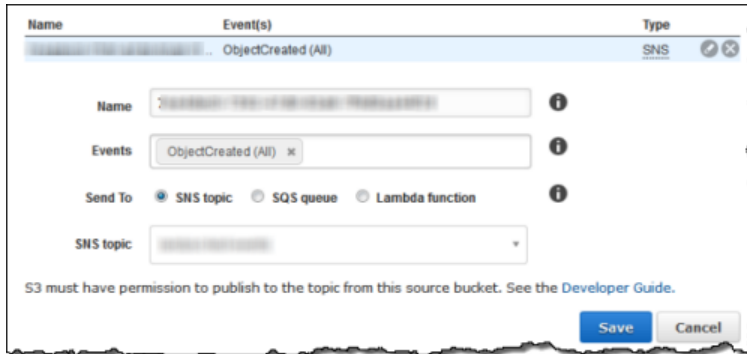
After you have saved an event notifications configuration, you can edit or delete the configuration.

To edit an event notifications configuration

1. In the Event Notifications list, click the pencil icon.



2. Make your changes and then click **Save**.



To delete an event notifications configuration

- In the Event Notifications list, click the x icon that appears on the right side of the screen for the event notification that you want to delete and then click **Save**.




Enabling Bucket Versioning

This section describes how to enable versioning on a bucket. For more information about versioning support in Amazon S3, see [Using Versioning](#) in the *Amazon Simple Storage Service Developer Guide*. For more information about managing objects when versioning is enabled, see [Managing Objects in a Versioning-Enabled Bucket](#) (p. 72).


To enable versioning 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 **Buckets** list, click the details icon on the left of the bucket name and then click **Properties** to display bucket properties.
3. In the **Properties** pane, click **Versioning** and then click **Enable Versioning**.



4. The console displays a confirmation dialog. Click **OK** to enable versioning on the bucket.

Amazon S3 enables versioning on the bucket. Accordingly, the console UI replaces the **Enable Versioning** button with the **Suspend Versioning** button.



After you enable versioning on a bucket, it can be in only the enabled or suspended state; you cannot disable versioning on a bucket. If you suspend versioning, Amazon S3 suspends the creation of object versions for all operations, but preserves any existing object versions. For more information, see [Working with Versioning-Suspended Buckets](#) in the *Amazon Simple Storage Service Developer Guide*.

Managing Lifecycle Configuration

This section explains how to manage lifecycle configuration rules for a bucket: adding, viewing, deleting, and disabling rules. You can use lifecycle configuration rules to archive or delete objects after a specified period of time. A transition action archives an object, and an expiration action deletes the object. For more information about lifecycle configuration transition and expiration actions, go to [Object Lifecycle Management](#) in the *Amazon Simple Storage Service Developer Guide*.

Archiving Objects

You can use a lifecycle configuration rule to archive objects to Amazon Glacier. An archived object is not directly accessible unless you restore a temporary copy. Additionally, you cannot use a lifecycle configuration rule to change the storage class of the archived object from Glacier to Standard or RRS.

Amazon S3 objects that have been archived to the Glacier storage class are visible and available only through the Amazon S3 console or the API, not through the Amazon Glacier console or the API.

Deleting Objects

You can also use a lifecycle configuration rule to delete objects. You might have objects in Amazon S3 or archived to Amazon Glacier that you want to delete using a lifecycle configuration rule. For more information about archiving objects and scheduling object deletions, see [Object Lifecycle Management](#) in the *Amazon Simple Storage Service Developer Guide*.

You can add lifecycle rules to buckets that have object versioning enabled or suspended as well as to buckets that do not. For information on how to enable versioning on a bucket, see [Enabling Bucket Versioning](#) (p. 27).

Topics

- [Lifecycle Configuration for a Bucket without Versioning](#) (p. 29)
- [Lifecycle Configuration for a Bucket with Versioning](#) (p. 32)
- [Maintaining Lifecycle Configuration Rules](#) (p. 38)

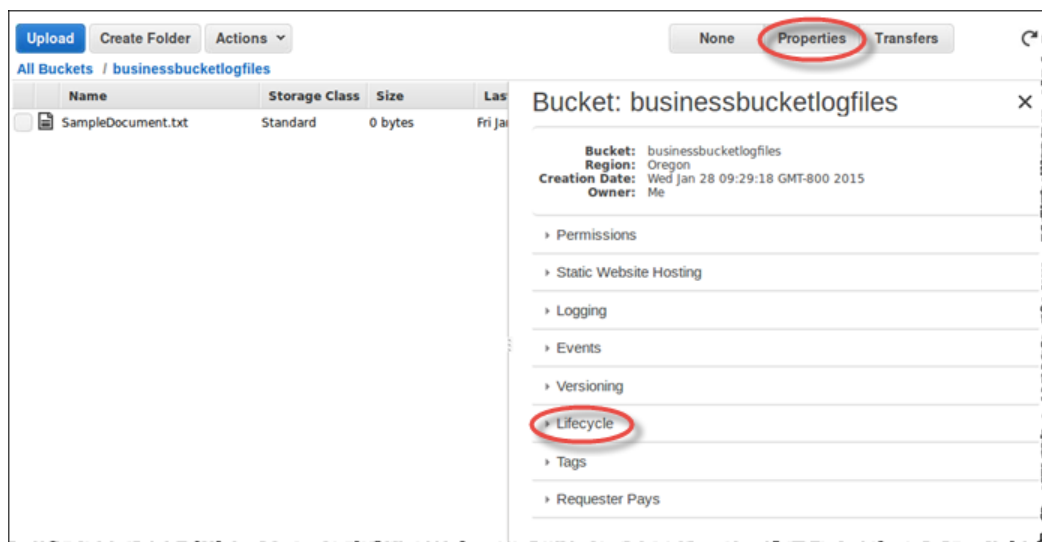
Lifecycle Configuration for a Bucket without Versioning

You can use lifecycle configuration rules to archive or delete objects after a specified period of time. For more information about lifecycle configuration rules, see [Object Lifecycle Management](#) in the *Amazon Simple Storage Service Developer Guide*.

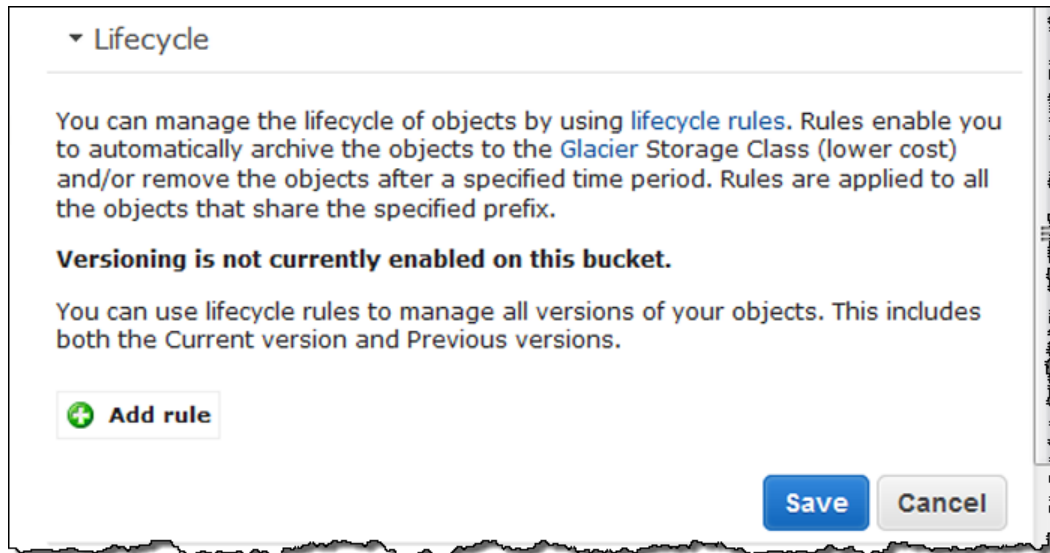
The following example walkthrough creates a lifecycle configuration rule for a bucket that archives your log files after one week and then permanently deletes them after a year.

To add a lifecycle configuration rule to a bucket without versioning

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Buckets** list, choose the bucket whose lifecycle configuration you want to configure, click **Properties** and then choose **Lifecycle**.

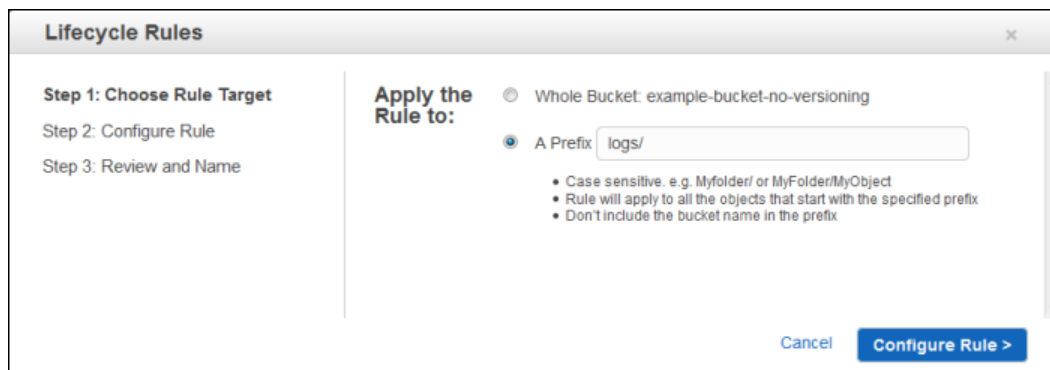


3. Choose **Add rule**.



4. In **Step 1: Choose Rule Target** choose **A Prefix** and type **logs/** as the prefix to specify the subset of objects to which the rule applies and then choose **Configure Rule**. (In our example, by entering **logs/** you will apply the rule to all objects in the bucket's logs folder.)

If you chose **Whole Bucket** instead, the rule would apply to all objects in the bucket.



5. In **Step 2: Configure Rule**, configure the rule as follows:
 - a. Select **Archive to the Glacier Storage Class** and enter the number of days after an object's creation date that you want to archive the object to the Glacier storage class.
 - b. Select **Permanently Delete** and enter the number of days after an object's creation date that you want the object to be permanently deleted. You cannot recover permanently deleted objects.
 - c. Verify that the illustration in the **EXAMPLE** section matches how you want your rule to work.
 - d. Choose **Review**.

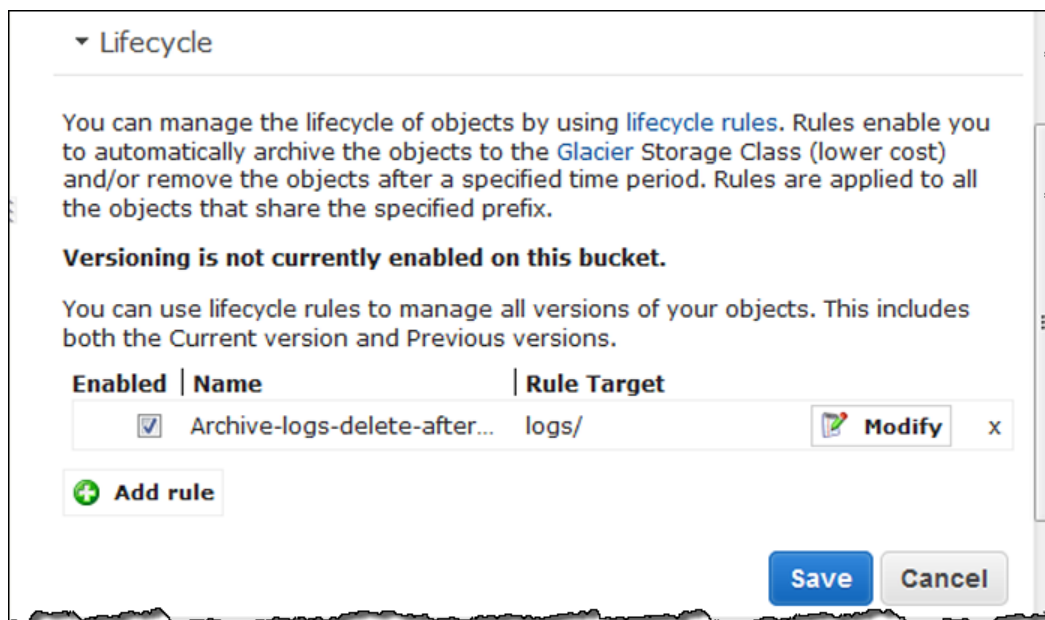
Amazon Simple Storage Service Console User Guide Lifecycle Configuration for a Bucket without Versioning

The screenshot shows the 'Lifecycle Rules' console window. On the left, a sidebar indicates the current step is 'Step 2: Configure Rule'. The main area contains instructions and configuration options. Under 'Action on Objects', two options are checked: 'Archive to the Glacier Storage Class' with a value of 7 days, and 'Permanently Delete' with a value of 365 days. An 'EXAMPLE:' diagram shows a timeline: an object is uploaded on August 3, 2015 (Day 0). On August 10, 2015 (Day 7), the rule 'Archive to Glacier' is applied. On August 2, 2016 (Day 365), the rule 'Rule: Expire' is applied, resulting in the 'Object Deleted'.

6. In **Step 3: Review and Name**, do the following:
 - a. (Optional) You can give your rule a name to identify the rule, if you want. The name must be unique within the bucket. By default, Amazon S3 will generate a unique identifier for the rule.
 - b. Choose **Edit** next to **Rule Target** or **Rule Configuration** if you want to make changes.
 - c. Choose **Create and Activate Rule** when all of the settings are as you want them.

The screenshot shows the 'Lifecycle Rules' console window at the 'Step 3: Review and Name' stage. The 'Rule Name' field contains 'Archive-logs-delete-after-a-year'. The 'Rule Target' is 'logs/' in the 'example-bucket-no-versioning' bucket. The 'Rule Configuration' section shows the same actions as in the previous screenshot: 'Archive to the Glacier Storage Class 7 days after the object's creation date' and 'Permanently Delete 365 days after the object's creation date'. The 'Create and Activate Rule' button is highlighted.

7. If the rule does not contain any errors, it is displayed in the **Lifecycle** pane.



Note

If there is an issue with a rule, an error message appears, providing information about the issue that you need to resolve. For example, if you have multiple rules, Amazon S3 determines if the rule that you're adding will conflict with an existing rule. In that case, the rule cannot be saved.

For more information about modifying, disabling, or deleting an existing lifecycle configuration rule, see [Maintaining Lifecycle Configuration Rules](#) (p. 38).

Lifecycle Configuration for a Bucket with Versioning

You can add lifecycle rules to buckets that have object versioning enabled or suspended. You use object versioning to keep multiple versions of an object in an Amazon S3 bucket. A versioning-enabled bucket can have many versions of the same object, one current version and zero or more previous versions. For more information about versioning, see [Using Versioning](#) and [Object Versioning](#) in the *Amazon Simple Storage Service Developer Guide*.

Note

The *Amazon Simple Storage Service Developer Guide* uses the term *noncurrent* version instead of *previous* version. Both terms mean the same thing.

This topic walks you through creating a lifecycle configuration rule for a bucket that has versioning enabled. You can also add lifecycle configuration rules to a bucket with versioning suspended. For information about how the rules work with a bucket in the versioning-suspended state, see [Object Lifecycle Management](#) in the *Amazon Simple Storage Service Developer Guide*.

The combined functionality of these two Amazon S3 features acts like a recycling bin and provides the following benefits:

- You can recover previous versions for a specified time to protect against unintended overwrites or deletions of your content.

- You can set specific windows of time for retaining the previous versions of your objects in Amazon S3, archiving in Amazon Glacier, or scheduling automatic deletion to help you control storage costs.

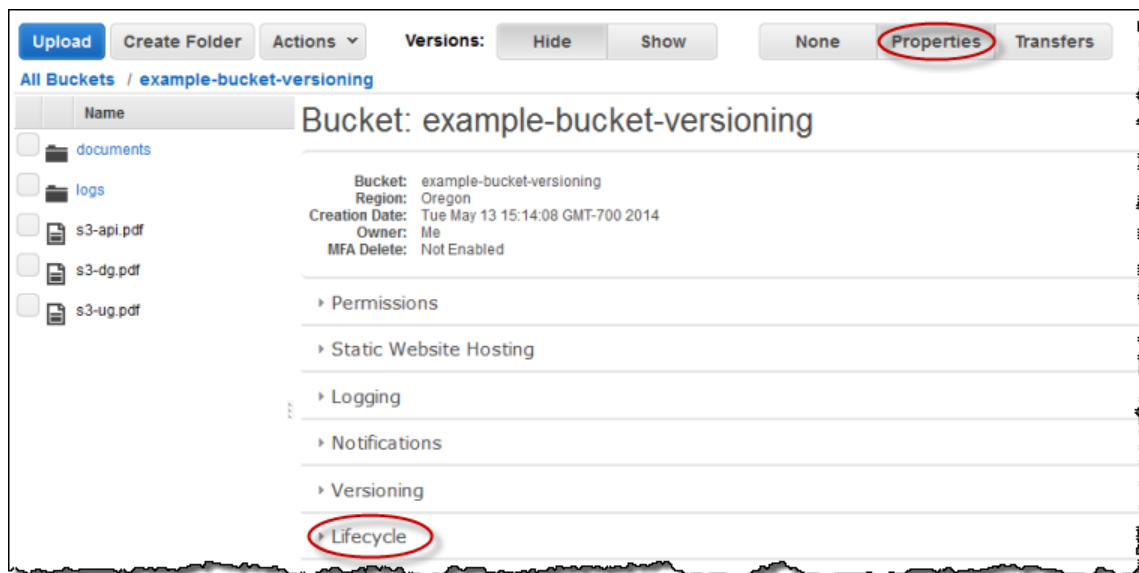
Archiving and Deleting Objects

You can use a lifecycle configuration rule to archive current and previous versions of your objects to Amazon Glacier. You can also use a lifecycle configuration rule to delete current and previous versions of your objects. For more information about archiving and scheduling object deletions, see [Object Lifecycle Management](#) in the *Amazon Simple Storage Service Developer Guide*.

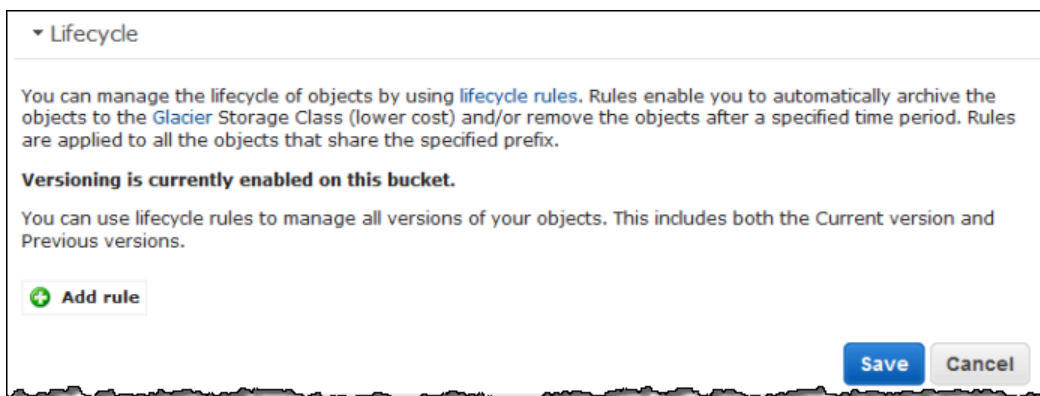
The following example walkthrough adds a lifecycle configuration rule to a bucket with versioning enabled. The configuration rule archives the current version of the files that are in the documents folder after 180 days, expires the files after one year, and then permanently deletes the files after they have been stored for two years. For more information about expiring and permanently deleting objects, see [Deleting Object Versions](#) in the *Amazon Simple Storage Service Developer Guide*.

To add a lifecycle configuration rule to a bucket with versioning enabled

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Buckets** list, choose the bucket whose lifecycle configuration you want to configure, choose **Properties**, and then choose **Lifecycle**.

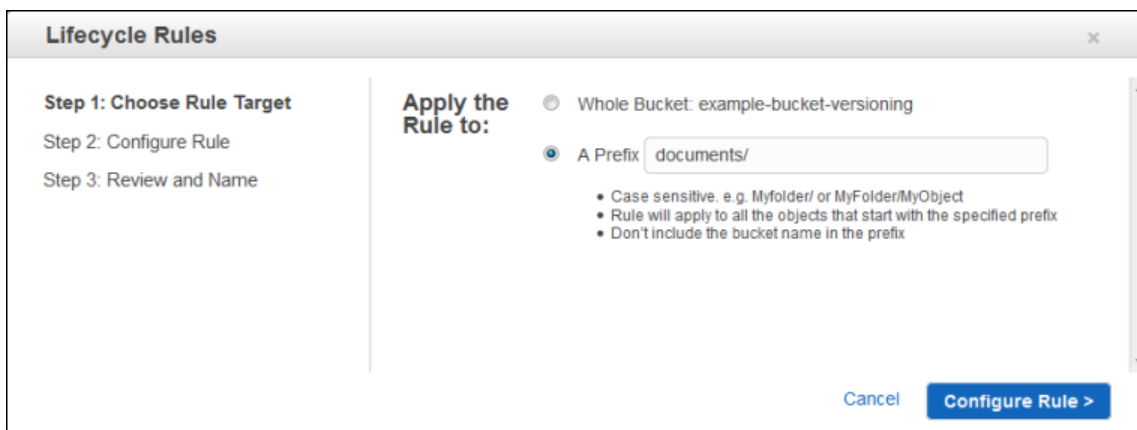


3. Choose **Add rule**.

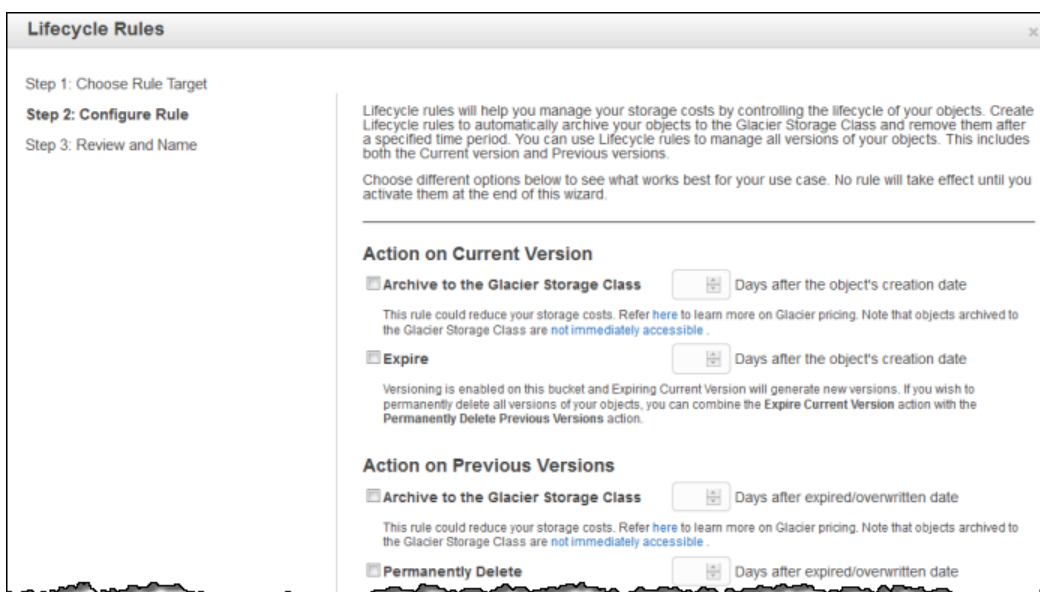


4. In **Step 1: Choose Rule Target**, choose **A Prefix** and type **documents/** as the prefix to specify the subset of objects to which the rule applies, and then choose **Configure Rule**. (In our example, by entering **documents/** you will apply the rule to all objects in the bucket's documents folder.)

If you chose **Whole Bucket** instead, the rule would apply to all objects in the bucket.



5. In **Step 2: Configure Rule**, configure the rule as follows:



- a. Under **Action on Current Version**, do the following:
 - i. Select **Archive to the Glacier Storage Class** and enter the number of days after the object's creation date that you want to archive current version of the object to the Glacier storage class.
 - ii. Select **Expire** and enter the number of days after the object's creation date that you want the current version of the object to expire.
 - iii. Verify that the illustration in the **EXAMPLE** section matches how you want your rule to work.

Action on Current Version

Archive to the Glacier Storage Class 180 Days after the object's creation date

This rule could reduce your storage costs. Refer [here](#) to learn more on Glacier pricing. Note that objects archived to the Glacier Storage Class are **not immediately accessible**.

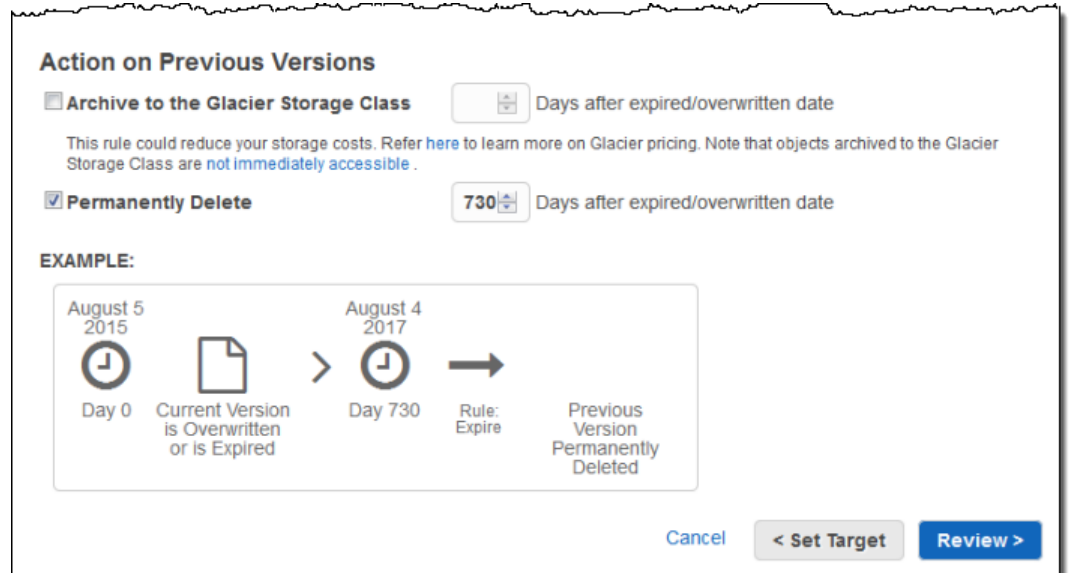
Expire 365 Days after the object's creation date

Versioning is enabled on this bucket and Expiring Current Version will generate new versions. If you wish to permanently delete all versions of your objects, you can combine the **Expire Current Version** action with the **Permanently Delete Previous Versions** action.

EXAMPLE:

August 5 2015 Day 0 Object Uploaded (Current Version) > February 1 2016 Day 180 Rule: Archive to Glacier -> Current Version Storage Class: Glacier > August 4 2016 Day 365 Rule: Expire -> Delete Marker Overwrites Current Version

- b. Under **Action on Previous Versions**, do the following:
 - i. Select **Permanently Delete** and enter the number of days after the object's expired/overwritten date that you want the previous versions to be permanently deleted. You cannot recover permanently deleted objects.
 - ii. Verify that the illustration in the **EXAMPLE** section matches how you want your rule to work.
 - iii. Choose **Review**.



6. In **Step 3: Review and Name**, do the following:
 - a. (Optional) You can give your rule a name to identify the rule, if you want. The name must be unique within the bucket. By default, Amazon S3 will generate a unique identifier for the rule.
 - b. Choose **Edit** next to **Rule Target** or **Rule Configuration** if you want to make changes.
 - c. Choose **Create and Activate Rule** when all of the settings are as you want them.

Lifecycle Rules

Step 1: Choose Rule Target
Step 2: Configure Rule
Step 3: Review and Name

Rule Name
Choose a descriptive name for your rule so you can easily identify it in the future. If you do not want to enter a name now, we will generate one for you.
Rule Name: (Optional)

Rule Target [Edit](#)
This rule will apply to Objects with the prefix: **documents/** in the **example-bucket-versioning** bucket

Rule Configuration [Edit](#)

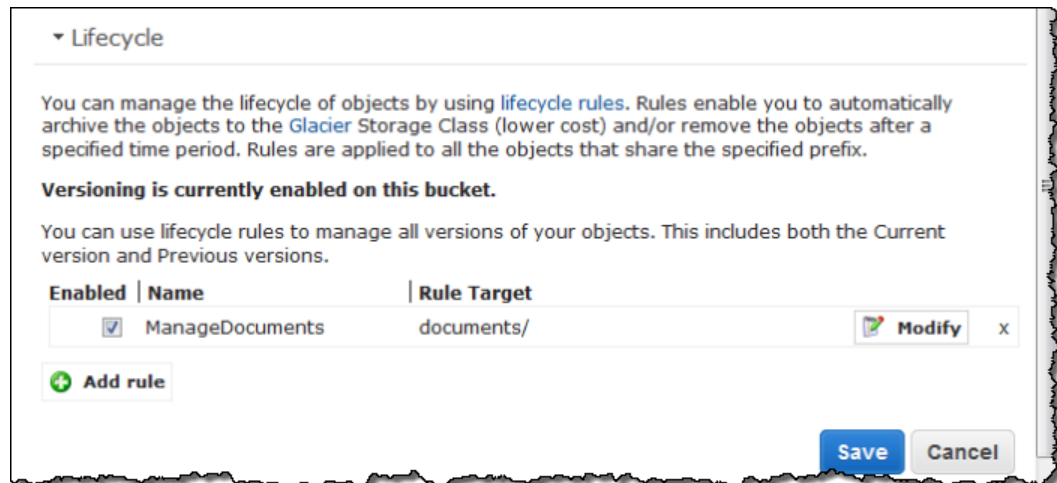
Action on Current Version
Archive to the Glacier Storage Class **180** days after the object's creation date.
This rule could reduce your storage costs. Refer [here](#) to learn more on Glacier pricing. Note that objects archived to the Glacier Storage Class are **not immediately accessible**.

Expire **365** days after the object's creation date
Versioning is enabled on this bucket and Expiring Current Version will generate new versions. If you wish to permanently delete all versions of your objects, you can combine the **Expire Current Version** action with the **Permanently Delete Previous Versions** action.

Action on Previous Versions
Permanently Delete **730** days after overwrite/expiration date.
You cannot recover Permanently Deleted objects.

[Cancel](#) [< Configure Rule](#) [Create and Activate Rule](#)

7. If the rule does not contain any errors, it is displayed in the **Lifecycle** pane.



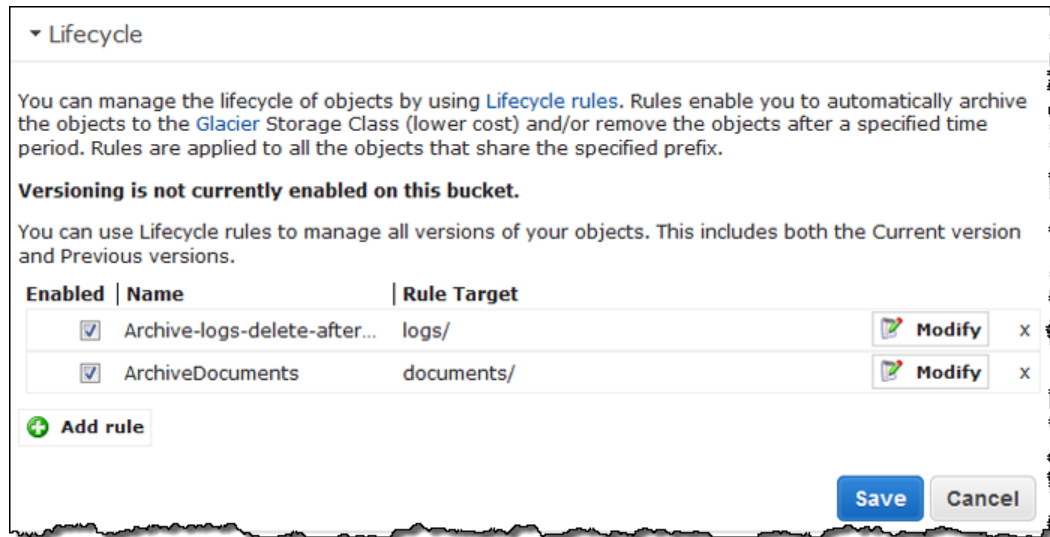
Note

If there is an issue with a rule, an error message appears, providing information about the issue that you need to resolve. For example, if you have multiple rules, Amazon S3 determines if the rule that you're adding will conflict with an existing rule. In that case, the rule cannot be saved.

For more information about modifying, disabling, or deleting an existing lifecycle configuration rule, see [Maintaining Lifecycle Configuration Rules \(p. 38\)](#).

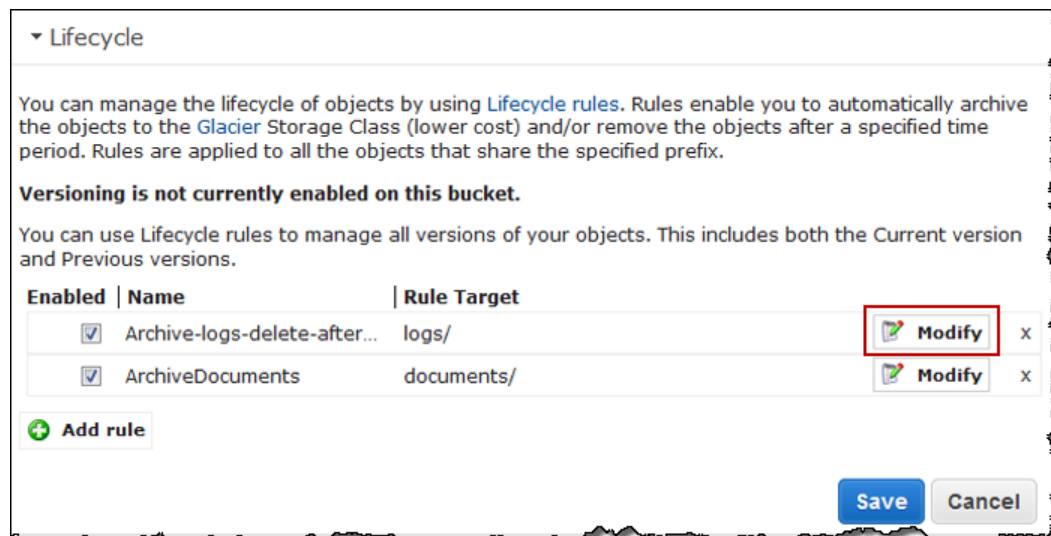
Maintaining Lifecycle Configuration Rules

Lifecycle configuration rules for a bucket are displayed in the **Lifecycle** pane.

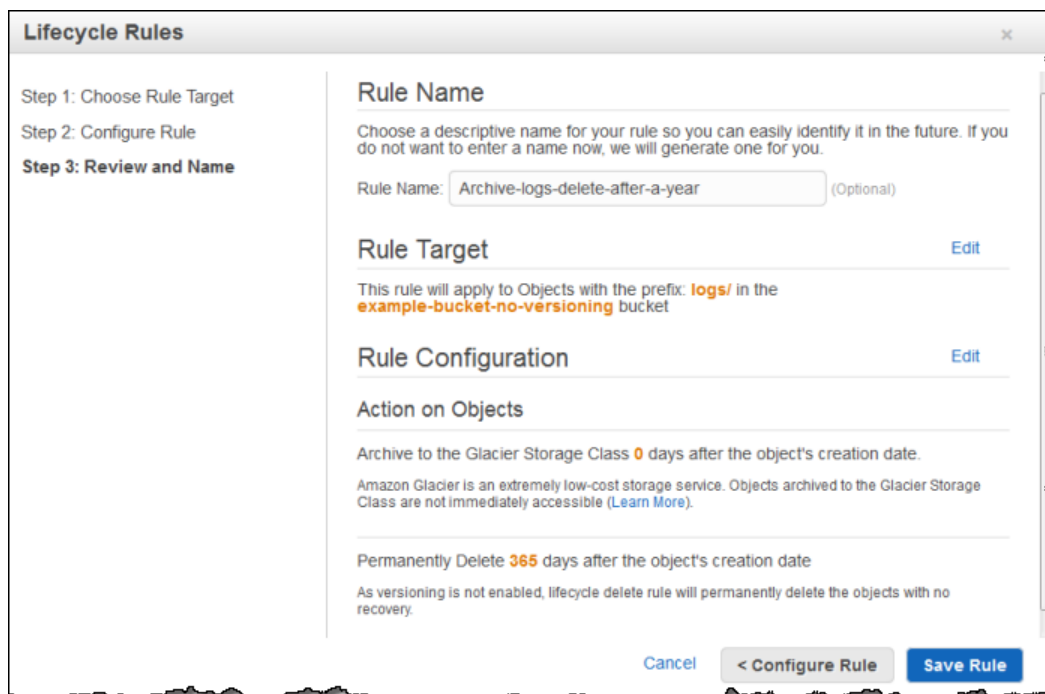


To modify a lifecycle configuration rule

1. In the **Buckets** list, click the name of the bucket that contains the rule, and then click **Lifecycle**.
2. Click **Modify** at the end of the row that describes the rule that you want to delete.



3. Modify your rule.



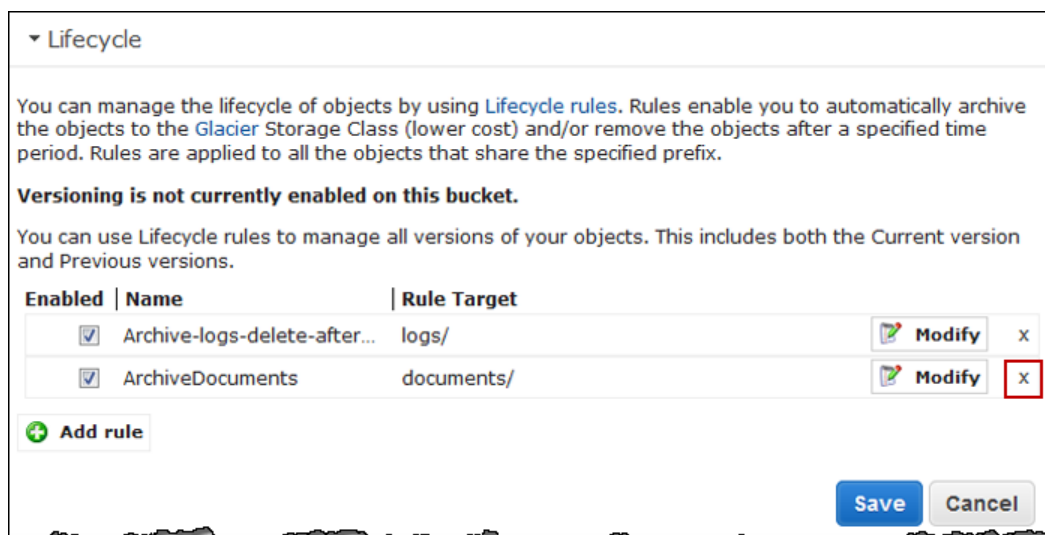
4. Click **Save Rule** when you are finished modifying your rule.

Note

You cannot modify legacy lifecycle configuration rules that use a specific date. The legacy rules will continue to work, but you cannot change them. However, you can disable or delete the date-based rules.

To delete a lifecycle configuration rule

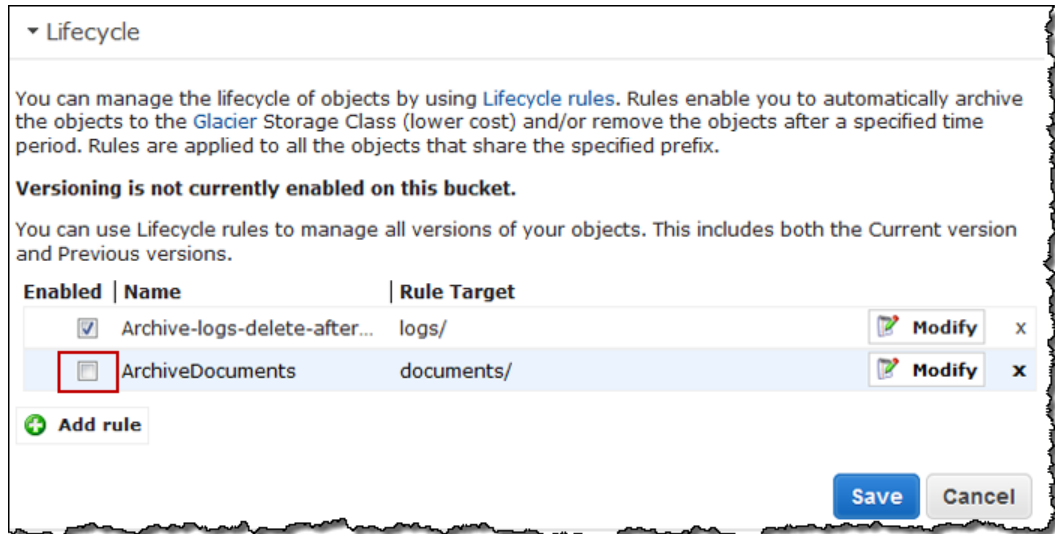
1. In the **Buckets** list, click the name of the bucket that contains the rule, and then click **Lifecycle**.
2. Click the **x** at the end of the row that describes the rule that you want to delete.



3. Click **Save**.

To disable a lifecycle configuration rule

1. In the **Buckets** list, click the name of the bucket that contains the rule, and then click **Lifecycle**.
2. Clear the **Enabled** check box for the rule.



3. Click **Save**.

The rule is not deleted; you can enable it again later if you want.

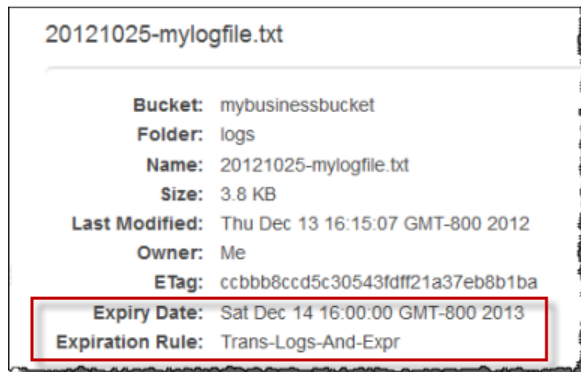
Rules that apply to an object are displayed with the object properties.

To view an object's expiration rule

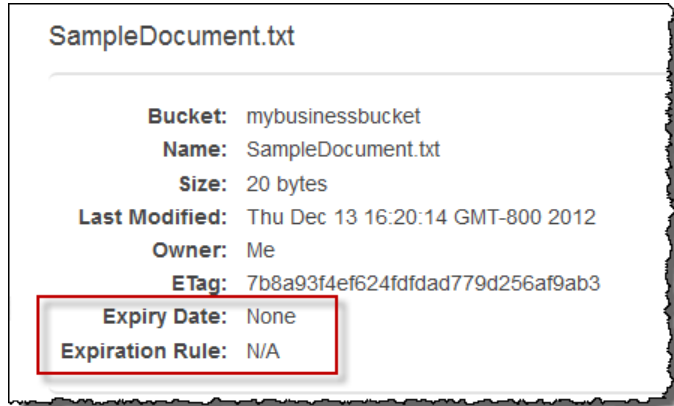
- In the **Object and Folders** list, click the object whose properties you want to view.

Among the object properties, the **Expiry Date** and **Lifecycle Rule** indicate which object expiration rule applies to the object. If no object expiration rule applies to the object, the **Expiry Date** field displays **None**, and the **Lifecycle Rule** field displays **N/A**.

The following example shows the properties for an object in which an rule named "Trans-Logs-And-Expr" applies to the object.



The following examples shows the properties for an object in which no expiration rule applies to the object.



Managing Cost Allocation Tagging

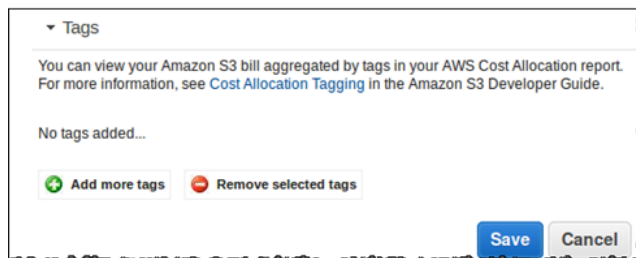
With AWS cost allocation, you can use tags to annotate billing for your use of a bucket. A tag is a key-value pair that represents a label that you assign to a bucket. In your AWS bill, costs are organized by tags that you define.

As a billing resource, a bucket can have as many as ten tags. In the following example, we'll create a tag that associates the bucket with a particular project. For information about cost allocation tagging, go to [Cost Allocation](#) in the *Amazon Simple Storage Service Developer Guide*.

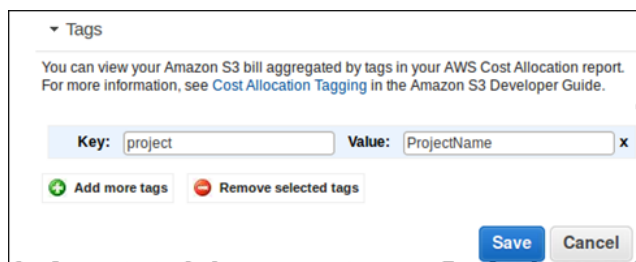
This section explains how to add and remove cost allocation tags for a bucket.

To add a cost allocation tag

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the Buckets list, click the bucket name, and then click **Tags**.



3. Click **Add more tags**.
4. In the **Key** and **Value** boxes, type a key name and a value.

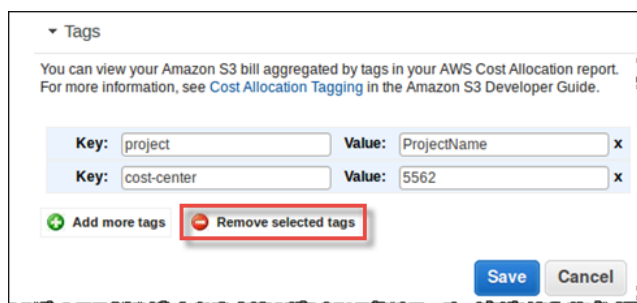


5. Click **Save**.

If there is an issue with a tag, an error message is displayed with information about the issue. For example, if the key-value pair is already in use or a key is missing its associated value, an error message is displayed, and the tag will not be saved.

To delete a cost allocation tag

1. In the Buckets list, click the bucket name, and then click **Tags**.
2. Select one or more tags to delete and click **Remove selected tags**. To select multiple tags, select one tag, and then either press the **Shift** key and drag to select multiple tags or hold down the **Ctrl** key while you click additional tags. The following example shows two tags selected.



You can also click the **x** to the right of a tag's **Value** field to delete just that tag.

3. Click **Save**.

Managing Cross-Region Replication

Cross-region replication is the automatic, asynchronous copying of objects across buckets in different AWS regions. By activating cross-region replication, Amazon S3 will replicate newly created objects, object updates, and object deletions from a source bucket into 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](#) the *Amazon Simple Storage Service Developer Guide*.

Topics

- [Enable Cross-Region Replication](#) (p. 42)
- [Disable or delete Cross-Region Replication](#) (p. 44)

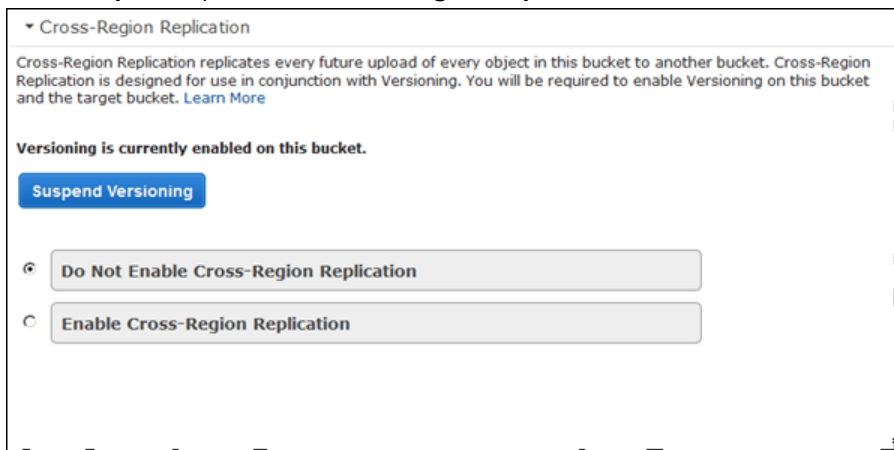
Enable Cross-Region Replication

In this section, you'll learn how to enable cross-region replication in the Amazon S3 console.

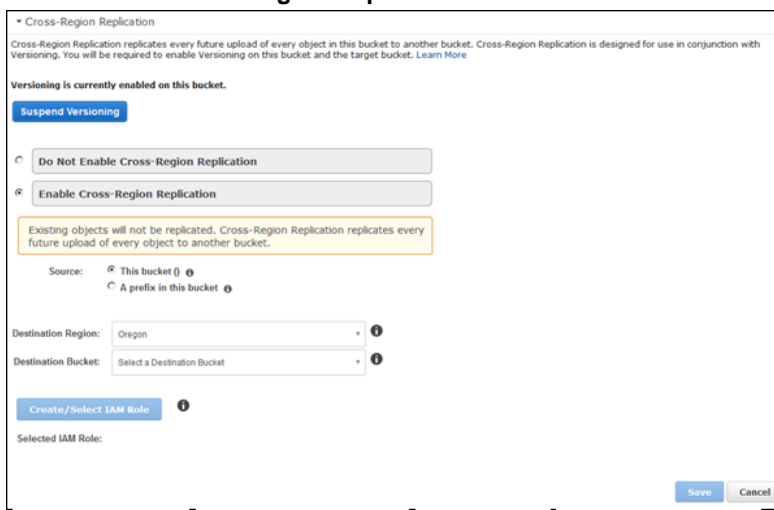
To enable cross-region replication between buckets

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Cross-region replication requires versioning enabled on both your source bucket and your destination bucket in a different region. For more information, see [Enabling Bucket Versioning](#) (p. 27).
3. In the **Buckets** list, select your source bucket and click on the Properties icon on the left of the bucket name.

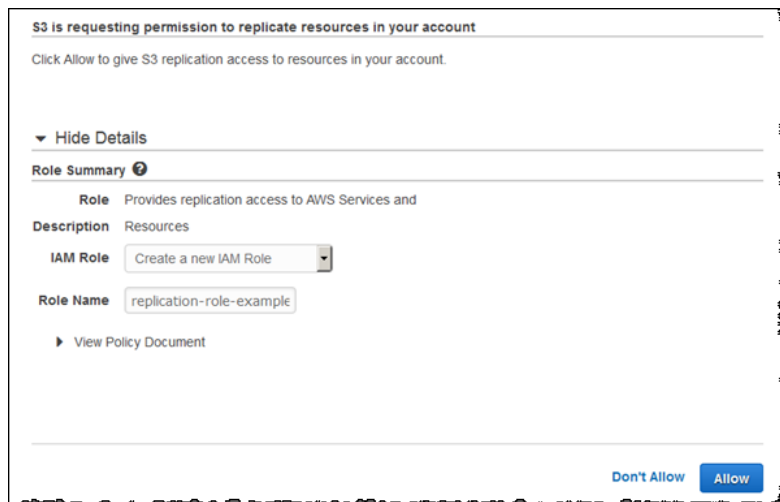
- In the **Properties** pane, click **Cross-Region Replication**.



- Select **Enable Cross-Region Replication**.



- Choose the **Source**—either the entire bucket or a prefix within the bucket.
- Choose the **Destination Region** from the drop-down list.
- 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 above, and that you have enabled versioning on that bucket. If no buckets exist in that region and you click **Create a new bucket** from the list, you'll be prompted to create a new bucket with versioning enabled in that destination region.
- In order to perform cross-region replication of objects on your behalf, Amazon S3 will need to use an IAM role that you have created. Click **Create/Select IAM Role** and a new browser tab will open up within the AWS Identity and Access Management (IAM) console.



On this page, you'll select an existing IAM role or create a new one that will allow Amazon S3 to replicate objects from the source bucket to the destination bucket on your behalf. By default, Amazon S3 will generate a policy document for the IAM role that matches the source and destination buckets you've chosen. To continue, click **Allow** to return to the Amazon S3 console. For more information about IAM roles, see [IAM Roles \(Delegation and Federation\)](#) in *Using IAM*.

10. Click **Save**.

You have now enabled cross-region replication between two buckets. The time it takes for Amazon S3 to replicate an object depends on the object size. For large objects, it can take up to several hours.

Note

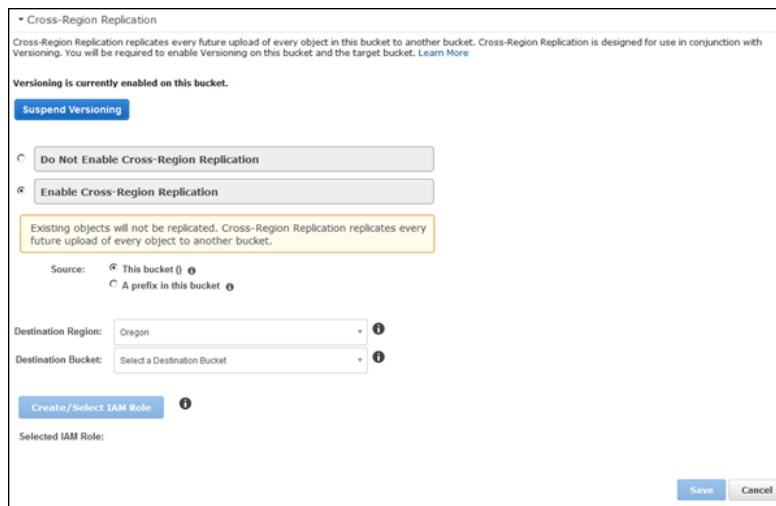
Metadata for an object remains identical between original objects and replica objects. Lifecycle rules abide by the creation time of the original object, and not by when the replicated object became available in the destination bucket. However, lifecycle actions on objects pending replication will not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

Disable or delete Cross-Region Replication

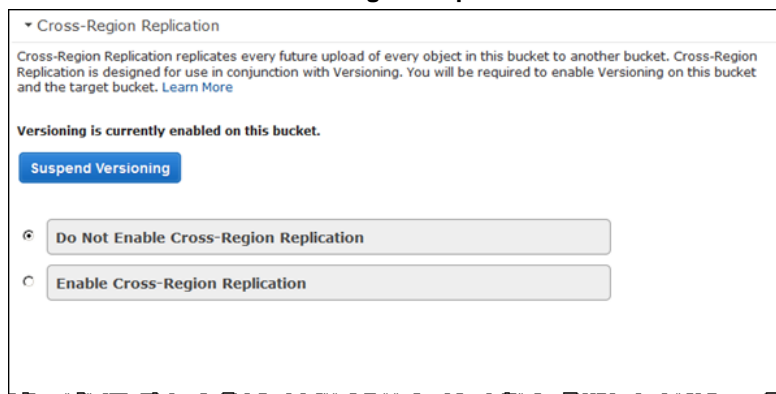
In this section, you'll learn how to disable cross-region replication in the Amazon S3 console. The configuration for cross-region replication can be partially deleted, in the case of removing prefixes, or fully disabled.

To fully disable cross-region replication between two buckets in the Amazon S3 console

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Buckets** list, select your source bucket and click on the Properties icon on the left of the bucket name.
3. In the **Properties** pane, click **Cross-Region Replication**.



4. Select **Do Not Enable Cross-Region Replication**.



5. Click **Save**.

This action fully disables cross-region replication between two buckets. The previous cross-region replication configuration is not deleted, but disabled, and you can re-enable that configuration at any time by clicking **Enable Cross-Region Replication** and then clicking **Save**.

To partially delete the cross-region replication configuration between two buckets by removing prefixes

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. In the **Buckets** list, select your source bucket and click on the Properties icon on the left of the bucket name.
3. In the **Properties** pane, click **Cross-Region Replication**.

▼ Cross-Region Replication

Cross-Region Replication replicates every future upload of every object in this bucket to another bucket. Cross-Region Replication is designed for use in conjunction with Versioning. You will be required to enable Versioning on this bucket and the target bucket. [Learn More](#)

Versioning is currently enabled on this bucket.

Suspend Versioning

Do Not Enable Cross-Region Replication

Enable Cross-Region Replication

Existing objects will not be replicated. Cross-Region Replication replicates every future upload of every object to another bucket.

Source: This bucket () ⓘ
 A prefix in this bucket ⓘ

Prefix1: ×

Prefix2: ×

Destination Region: ⓘ

Destination Bucket: ⓘ

ⓘ

Selected IAM Role: myfinances-examplebucketnameone-s3-repl-role

4. Click the delete icon next to the prefix that you want to remove from this cross-region replication configuration.

▼ Cross-Region Replication

Cross-Region Replication replicates every future upload of every object in this bucket to another bucket. Cross-Region Replication is designed for use in conjunction with Versioning. You will be required to enable Versioning on this bucket and the target bucket. [Learn More](#)

Versioning is currently enabled on this bucket.

Suspend Versioning

Do Not Enable Cross-Region Replication

Enable Cross-Region Replication

Existing objects will not be replicated. Cross-Region Replication replicates every future upload of every object to another bucket.

Source: This bucket () ⓘ A prefix in this bucket ⓘ

Prefix1:

Prefix2:

Destination Region: ⓘ

Destination Bucket: ⓘ

Create/Select IAM Role ⓘ

Selected IAM Role: myfinances-examplebucketnameone-s3-repl-role

5. Click **Save**.

This action deletes the prefix from the configuration for cross-region replication for these two buckets. This means that all objects with that prefix will no longer be replicated across regions.

Note

If you delete all the prefixes from your cross-region replication configuration in the Amazon S3 console, then the Amazon S3 console assumes that you want to enable cross-region replication on every object in the source bucket. That means that every newly created object, object update, and object deletion in the source bucket will be replicated into the destination bucket, regardless of prefix.

You can't suspend versioning on your buckets until the replication configuration is deleted. The Amazon S3 console allows you to disable replication, but it will not delete it. You can delete your cross-region replication configuration using the following AWS Command Line Interface (CLI) command.

```
aws s3api delete-bucket-replication --bucket BUCKETNAME
```

For information about using the CLI, go to [Getting Set Up with the AWS Command Line Interface](#) in the *AWS Command Line Interface User Guide*.

Working with Objects

Topics

- [Uploading Objects into Amazon S3 \(p. 49\)](#)
- [Editing Object Properties \(p. 57\)](#)
- [Opening an Object \(p. 64\)](#)
- [Downloading an Object \(p. 65\)](#)
- [Copying an Object \(p. 66\)](#)
- [Renaming an Object \(p. 67\)](#)
- [Deleting an Object \(p. 68\)](#)
- [Restoring an Object \(p. 69\)](#)
- [Managing Objects in a Versioning-Enabled Bucket \(p. 72\)](#)

Objects are the data that you store in Amazon S3. Every object resides within a bucket you create in specific AWS region.

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 Amazon S3 region physically remain in that region. Amazon S3 does not keep copies or move it to any other region. However, you can access the objects from anywhere, as long as you have necessary permissions.

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

Objects can be any file type: images_backup, data, movies, etc. An object can be as large as 5 TB. You can have an unlimited number of objects in a bucket.

This section explains how to use the console to create, manage, and delete objects.

Uploading Objects into Amazon S3

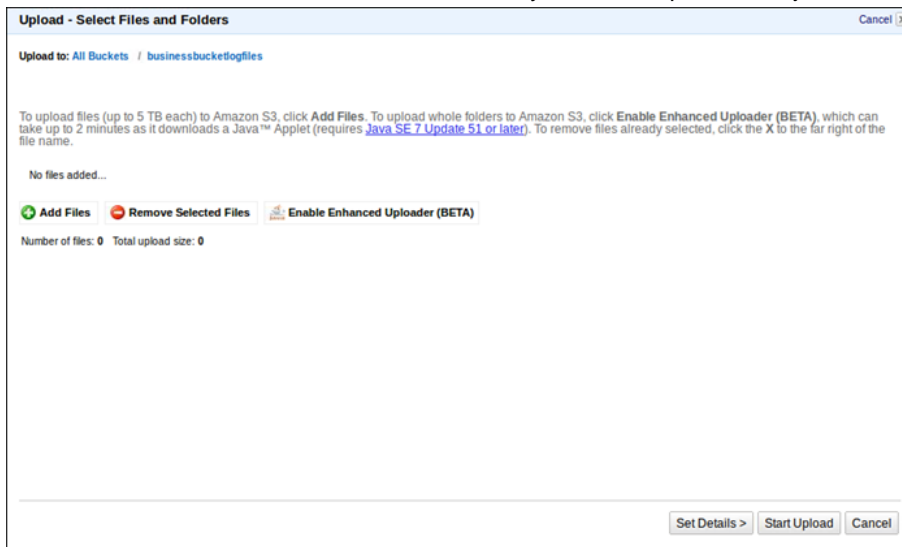
When you upload a folder, Amazon S3 uploads all the files and subfolders from the specified folder to your bucket. It then assigns a key value that is a combination of the uploaded file name and the folder name. For example, if you upload a folder `/images` containing two files, `sample1.jpg` and `sample2.jpg`,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

If you upload one or more files that are not in a folder, Amazon S3 uploads the files and assigns the file names as the key values for the objects created.

This section explains how to use the AWS Management Console to upload one or more files or entire folders into Amazon S3. Amazon S3 stores all files in the specified bucket.

To upload files and folders into Amazon S3

1. Sign into the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3>.
2. In the buckets list, click the name of bucket where you want to upload an object and then click **Upload**.



3. (Optional) In the **Upload - Select Files** wizard, if you want to upload an entire folder, click **Enable Enhanced Uploader** to install the necessary Java applet. After you choose the Enhanced Uploader, if the uploader is not ready to use after two minutes, you might need to change your platform (Windows or Mac) or browser configuration to get the Java applet to work. For instructions on changing your platform and browser configuration, see [Using the Enhanced Uploader \(p. 55\)](#).


You only need to do this step once per console session. After you click **Enable Enhanced Uploader** and then don't want to use it, you can either refresh the browser, or close and reopen the browser to reset the uploader to the default.

The Enhanced Uploader uses a Java applet. .

Note

If you are behind a firewall, you will need to install your organization's supported proxy client for the Java applet to work.

4. Click **Add Files**.



5. In the dialog box that appears, click the file or files that you want to upload, and then click **Open**.
 - If you enabled the advanced uploader in step 2, you see a Java dialog box titled **Select files and folders to upload**, as shown.
 - If not, you see the **File Upload** dialog box associated with your operating system.
6. If you are ready to upload the object immediately, without providing further details about the object, click **Start Upload**. Otherwise, click **Set Details**.

The **Set Details** dialog box gives you the options to **Use Reduced Redundancy Storage** or **Use Server Side Encryption**.

- **Use Reduced Redundancy Storage**— In order to reduce storage costs, you can use reduced redundancy storage for noncritical, reproducible data at lower levels of redundancy than Amazon S3 provides with standard storage. For more information, see [Using Reduced Redundancy Storage](#) in the *Amazon Simple Storage Service Developer Guide*.
- **Use Server Side Encryption**— With server-side encryption (SSE), Amazon S3 encrypts your data at the object level as it writes it to disks in its data centers and decrypts it for you when you access it. For more information about using SSE in Amazon S3 go to [Protecting Data Using Server-Side Encryption](#) in the *Amazon Simple Storage Service Developer Guide*.

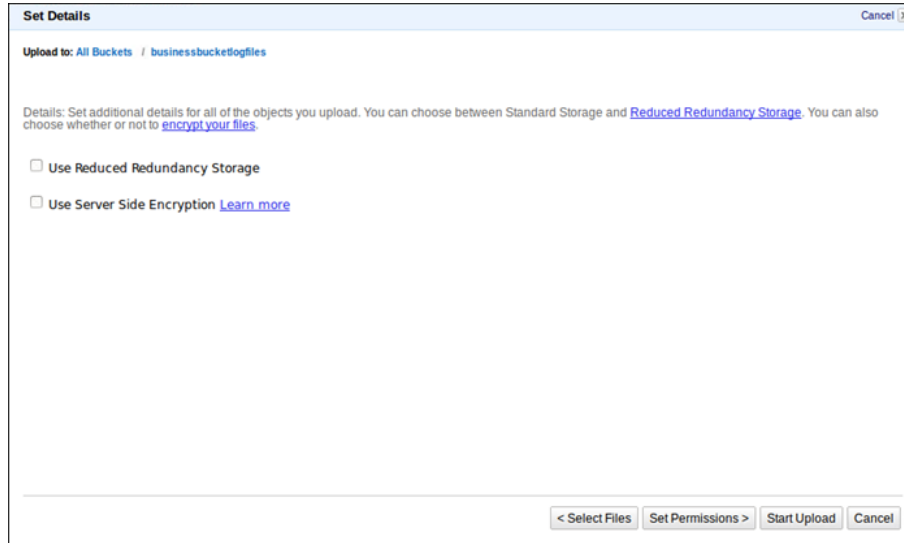
Here in the **Set Details** dialog box, you have two SSE options; **Use the Amazon S3 service master key** or **Use an AWS Key Management Service master key**. Selecting the AWS Key Management Service option enables you to select the **Master Key** from a dropdown list with the following options:

- **aws/s3 (default)**— This is the default AWS KMS master key.
- **Enter a key ARN**— You can give external accounts the ability to use this object protected by a AWS KMS key. To do this, you'll need to provide the Amazon Resource Name (ARN) for the external account in the **ARN / ID** field. 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. The other options in this dropdown list are all AWS KMS master keys that you have previously created. For more information about creating a AWS KMS key, go to [Creating Keys](#) in the *AWS Key Management Service Developer Guide*.

Note

Only keys in the same region as this bucket are available for encrypting objects in this bucket.

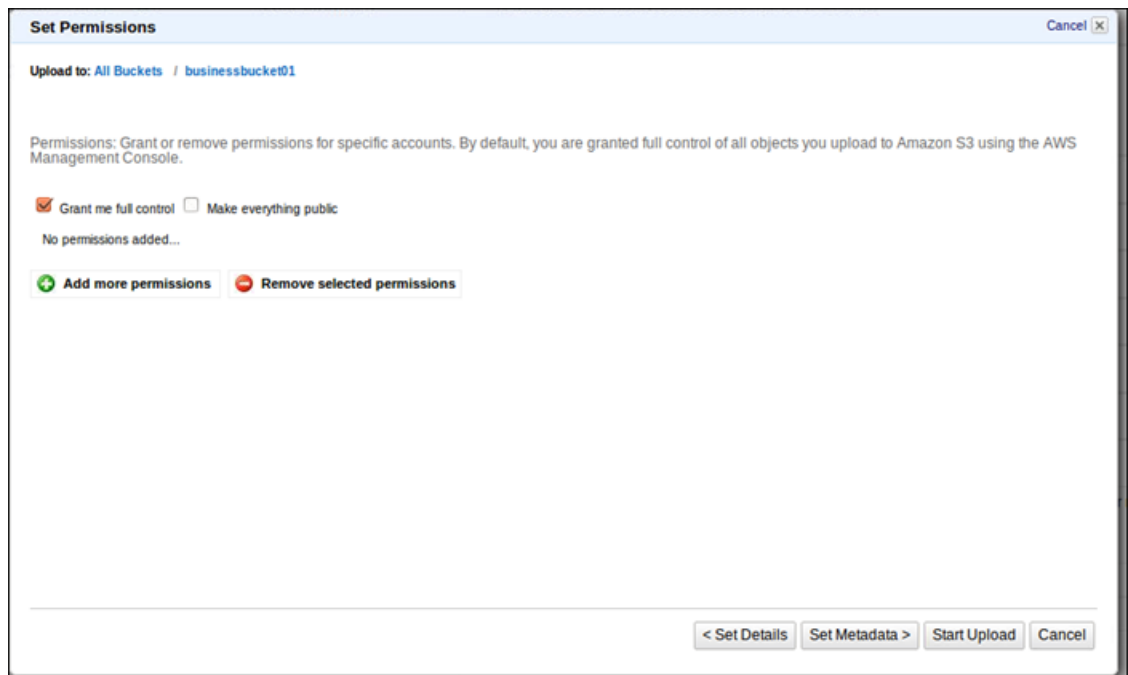
When you've finished setting the object details, click **Set Permissions**.



7. In the **Set Permissions** dialog box, do the following:
- Select (the default) or clear the **Grant me full control** check box.
 - To grant read access to anonymous requests, select the **Make everything public** check box on the **Upload - Set Permissions** panel. By default, the check box is cleared, so no access is granted.

Note

By default, the owner of the upload has full control over all uploaded objects.



8. To grant access to other users and groups for the objects you are uploading, click **Add more permissions**.

In the grantee row that appears:

- For each permission you grant, an entry is made in the object's Access Control List (ACL). For more information, see [Using ACLs](#) in the *Amazon Simple Storage Service Developer Guide*.
- If you click **Add more permissions**, a new **Grantee** row appears. Each **Grantee** row maps to a grant in the Access Control List (For more information, see [Using ACLs](#)) associated with the object. You can grant permission to a user or one of the predefined Amazon S3 groups.

9. There are built-in groups that you can choose from the **Grantee** box:

- **Everyone**—Use this group to grant anonymous access.
- **Authenticated Users**—This group consists of any user that has an Amazon AWS Account. When you grant the Authenticated User group permission, any valid signed request can perform the appropriate action. The request can be signed by either an AWS Account or IAM User.
- **Log Delivery**—This group grants write access to your bucket when the bucket is used to store server access logs. For more information, see [Managing Bucket Logging](#).
- **Me**—This group refers to your AWS root account, and not an IAM user.

You can grant permission to an AWS account by entering the accounts canonical user ID or email address in the **Grantee** field. The email address must be the same one they used when signing up for an AWS account. You can grant a grantee any of the following permissions:

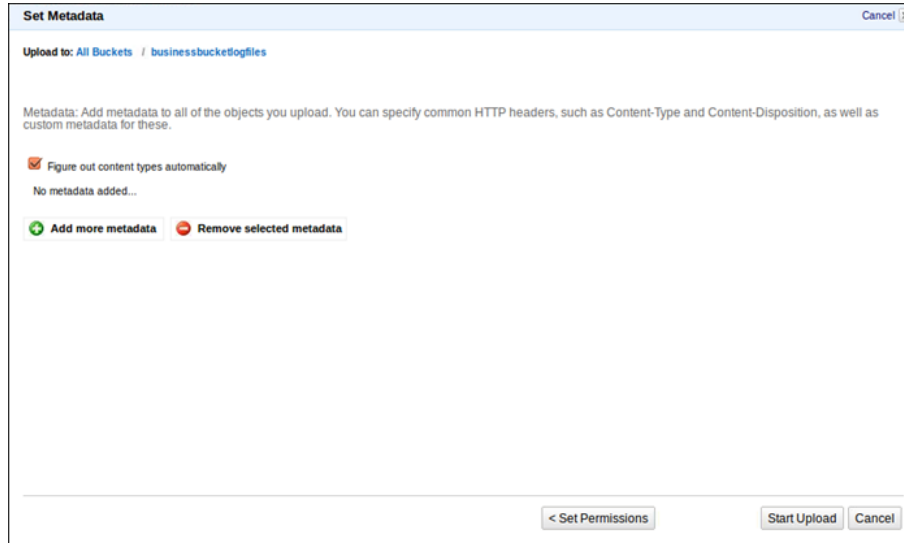
- **Open/Download**—Enables the account to access the object when they are logged in
- **View Permissions**—Can view the permissions associated with the object
- **Edit Permissions**—Can edit the permissions associated with the object

10. To set metadata, click **Set Metadata**.

In the **Upload - Set Metadata** do the following:

- a. If you want the Amazon S3 to infer the content type of the uploaded objects, select the **Figure out content types automatically** check box (default).
- b. To add custom metadata, click **Add more metadata** and enter the key-value pairs that you want.

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 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. When you add user-defined metadata, select `x-amz-meta-` from the **Key** box and then append the metadata name to it.



11. Click **Start Upload**.

You can watch the progress of the upload from within the **Transfers** panel.


Tip

To hide the **Transfer** panel, click **None**. To open it again, click **Transfers**.

When objects upload successfully to Amazon S3, they appear in the Objects and Folders list.

To view file content and properties

- Do either or both of the following:
 - To view the file content, in the Objects and Folders list, double-click the object name.
 - To view object properties, in the Objects and Folders list, click the object.



Note

By default your Amazon S3 resources are private. Only the object owner can click the object link and view the object. If you share this link with others, for example add this link to your web pages, Amazon S3 will deny access. The clickable links on your webpage will work only if you make the object public (see [Editing Object Permissions \(p. 60\)](#)) or you use a pre-signed URL for the clickable link. For more information about pre-signed URL, go to [Share an Object with Others](#) in the *Amazon Simple Storage Service Developer Guide*.

Using the Enhanced Uploader

Topics

- [Using the Enhanced Uploader in Microsoft Windows \(p. 55\)](#)
- [Using the Enhanced Uploader on the Mac \(p. 56\)](#)

The Enhanced Uploader uses a Java applet. After you choose the Enhanced Uploader, if the uploader is not ready to use after two minutes, you might need to change your platform (Windows or Mac) or browser configuration to get the Java applet to work. The instructions in this section describe how to make these changes depending on which platform and browser you are using.

Using the Enhanced Uploader in Microsoft Windows

You need to enable Java in a browser before you can use the Enhanced Uploader on a computer running Microsoft Windows. After you enable Java, you can use the Enhanced Uploader with Internet Explorer or Mozilla Firefox on Windows.

Enable Java for Windows Browsers

Follow the instructions provided at [Launching Java Control Panel on Windows](#) to launch the Java Control Panel. Click the **Security** tab in the Java Control Panel, select **Enable Java content in the browser**, and then click **Apply**. Restart the browser and follow the browser specific steps in the following sections.

Using the Enhanced Uploader with Internet Explorer

This section describes how to use the Enhanced Uploader in Internet Explorer.

To use the Enhanced Uploader in Internet Explorer

1. Open Internet Explorer and sign into the AWS Management Console at <https://console.aws.amazon.com/s3>.
2. Click **Allow the Java(TM) plugin to run on the S3 console**, if your browser displays this message.
3. In the buckets list, click the name of bucket where you want to upload data and then click **Upload**.
4. Click **Enable Enhanced Uploader (BETA)**.
5. In the Security Warning window that asks **Do you want to run this application?**, select **I accept the risk and want to run this application** and then click **Run**.
6. Click **Add Items**.
7. If your browser displays the warning **Allow Access to the following application from this web site**, click **Allow**.
8. In the **Select files and folders to upload** window, select the files and folders that you want to upload and then click **Open**.
9. (Optional) Click **Set Details** to choose a storage class, configure encryption, set permissions, and set metadata.
10. Click **Start Upload**.

Using the Enhanced Uploader with Mozilla Firefox on Windows

This section describes how to use the Enhanced Uploader in Firefox.

To use the Enhanced Uploader in Mozilla Firefox

1. Open Firefox and sign into the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3>.

2. In the buckets list, click the name of bucket where you want to upload data and then click **Upload**.
3. Click **Enable Enhanced Uploader (BETA)**.
4. In the Security Warning window that asks **Do you want to run this application?**, select **I accept the risk and want to run this application** and then click **Run**.
5. Click **Add Items**.
6. In the **Select files and folders to upload** window, select the files and folders that you want to upload and then click **Open**.
7. (Optional) Click **Set Details** to choose a storage class, configure encryption, set permissions, and set metadata.
8. Click **Start Upload**.

Using the Enhanced Uploader on the Mac

You can use the Enhanced Uploader with Safari or Mozilla Firefox on the Mac.

Using the Enhanced Uploader with Safari on the MAC

This section describes how to use the Enhanced Uploader with Safari. You may need to operate in Safari's unsafe mode for the Enhanced Uploader to run, which is described in the following procedure.

To use the Enhanced Uploader in Safari's unsafe mode

1. Open Safari, choose **Safari > Preferences** and then click **Security**.
2. Click **Website Setting** that is next to Internet plug-ins.
3. In the plug-ins windows that is displayed, click **Java** in the left pane.
4. In the **Configured Websites** pane, expand the drop-down next to the URL for the Amazon S3 Management Console website.
5. Click **Run in Unsafe Mode** and then click **Trust** in the warning message that appears.
6. Click **Done**.
7. Sign into the AWS Management Console at <https://console.aws.amazon.com/s3>.
8. In the buckets list, click the name of bucket where you want to upload data and then click **Upload**.
9. Click **Enable Enhanced Uploader (BETA)**.
10. In the Security Warning window that asks **Do you want to run this application?**, select **I accept the risk and want to run this application** and then click **Run**.
11. Click **Add Items**.
12. In the **Select files and folders to upload** window, select the files and folders that you want to upload and then click **Open**.
13. (Optional) Click **Set Details** to choose a storage class, configure encryption, set permissions, and set metadata.
14. Click **Start Upload**.

Using the Enhanced Uploader with Mozilla Firefox on the Mac

This section describes how to use the Enhanced Uploader in Firefox on the Mac.

To use the Enhanced Uploader in Firefox on the Mac

1. Open Firefox and click the menu icon (three horizontal bars on the top right of the window).
2. Click **Preferences** and then click **Content**.
3. If **Block pop-up windows** is selected, clear the check box to disable it.

4. Sign into the AWS Management Console at <https://console.aws.amazon.com/s3>.
5. In the buckets list, click the name of bucket where you want to upload data and then click **Upload**.
6. Click **Enable Enhanced Uploader (BETA)**.
7. Click the plugin icon in the address bar and a message panel opens.
8. In the message panel, click **Allow and Remember**.
9. In the Security Warning window that asks **Do you want to run this application?**, select **I accept the risk and want to run this application** and then click **Run**.
10. Click **Add Items**.
11. In the **Select files and folders to upload** window, select the files and folders that you want to upload and then click **Open**.
12. (Optional) Click **Set Details** to choose a storage class, configure encryption, set permissions, and set metadata.
13. Click **Start Upload**.

Editing Object Properties

Topics

- [Editing Object Details \(p. 57\)](#)
- [Editing Object Permissions \(p. 60\)](#)
- [Editing Object Metadata \(p. 63\)](#)

The properties of an object include the object details, permissions, and metadata that you set when you uploaded the object. You can edit these properties at any time.

This section explains the properties of an object that you can change and includes the object's details, permissions, and metadata.

To access the properties of an object

1. In the Objects and Folders list, click the object.
2. Do any or all of the following:
 - To edit the object details, click **Details**, and then edit the details as explained in [Editing Object Details \(p. 57\)](#).
 - To edit object permissions, click **Permissions**, and then edit the permissions as explained in [Editing Object Permissions \(p. 60\)](#).
 - To edit object metadata, click **Metadata**, and then edit the permissions as explained in [Editing Object Metadata \(p. 63\)](#).

When you select a single object in a bucket you can change all of its properties. When you select multiple objects, you can change only the object details.

Editing Object Details

This section explains how to use the console to edit the details of one or more selected objects. The property details of an object that you see and can change depends on the storage class of the object:

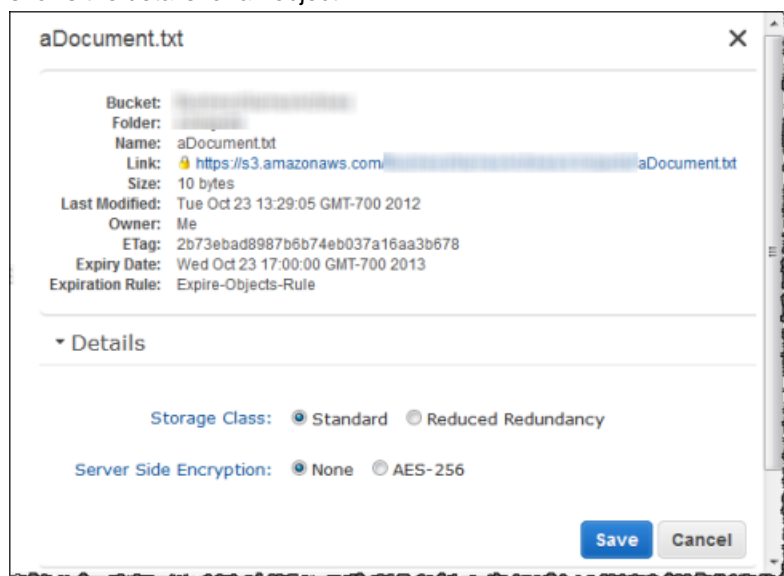
- **Standard and Reduced Redundancy Storage (RRS) Class** – When an object is in the Standard or RRS storage class, the properties of an object you can see and change include the object's storage redundancy and the state of server-side encryption. In general, you use Amazon S3 RRS to reduce

costs by storing noncritical, reproducible data at lower levels of redundancy than Amazon S3 standard storage. For more information, see [Using Reduced Redundancy Storage](#) in the *Amazon Simple Storage Service Developer Guide*. You can use server-side encryption to encrypt objects at rest. For more information, see [Using Encryption](#) in the *Amazon Simple Storage Service Developer Guide*.

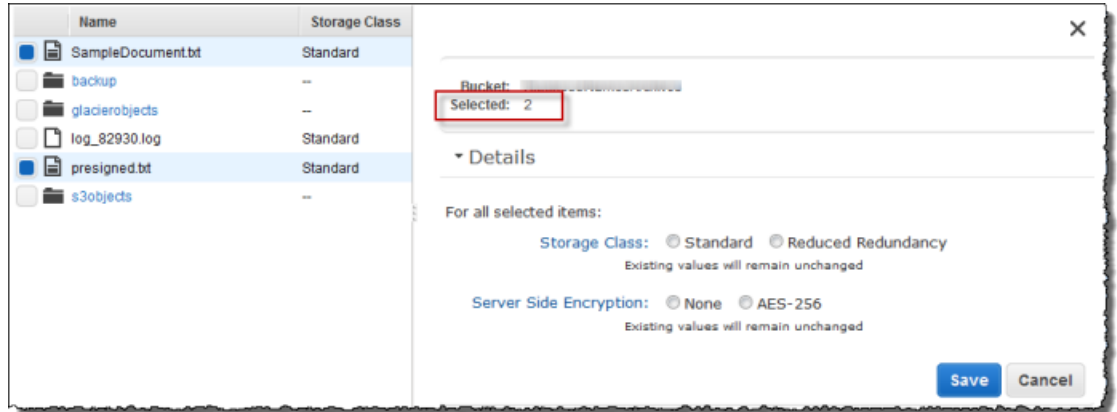
- **Use Server Side Encryption** – With server-side encryption (SSE), Amazon S3 encrypts your data at the object level as it writes it to disks in its data centers and decrypts it for you when you access it. For more information about using SSE in Amazon S3 go to [Protecting Data Using Server-Side Encryption](#) in the *Amazon Simple Storage Service Developer Guide*.

Standard and Reduced Redundancy Storage Class

When you select an object stored in the Standard or Reduced Redundancy Storage (RRS) class and click **Details**, the details become visible. You can change the **Storage Class** property or **Server Side Encryption** property of the object and click **Save** to save change to the properties. The following example shows the details for an object.

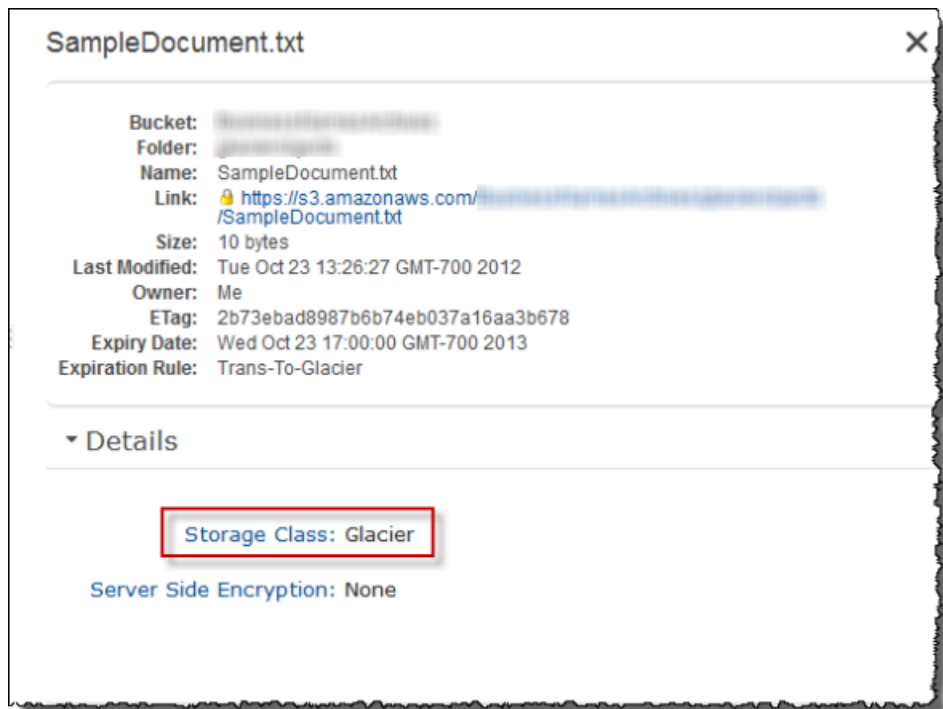


When you select two or more objects in a bucket and click **Details**, no selections for **Storage** or **Server Side Encryption** are shown, regardless of the settings of these properties for the files that are part of the selection. In this multiple object select case, the **Details** panel enables you to change one of the two properties for all of the selected objects. For example, if you select **AES-256** for **Server Side Encryption** and click **Save**, then all of the selected objects will be encrypted. The following example shows the details for two selected items.

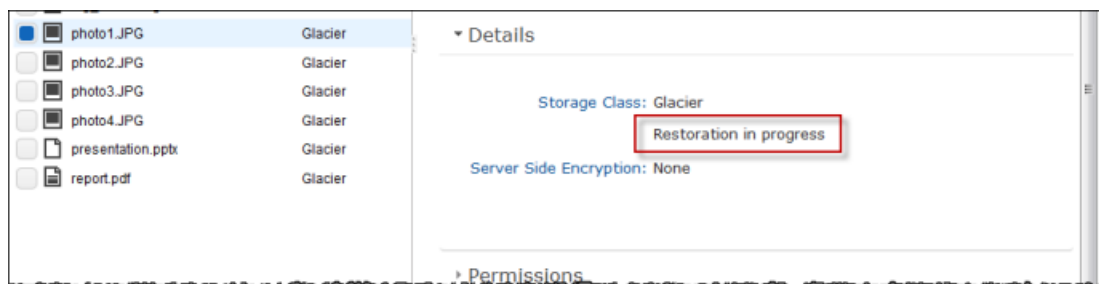


Amazon Glacier Storage Class

When you select an object stored in the Amazon Glacier Storage class and click **Details**, the details appear. If the object has not been restored, the properties of the object are view-only. The following example shows the details properties for an object stored in the Amazon Glacier storage class that has not been restored.



If the object is in the process of being restored, the **Details** tab indicates this. The following example shows the properties for an object stored in the Amazon Glacier storage class that is in the process of being restored. For more information about restoring, see [Restoring an Object \(p. 69\)](#).



If the object is restored, the date until which the object is restored is displayed under **Details**. The following example shows properties of a restored object. You can use the **Modify** button to change the length of time until which the object is restored.



When you select two or more Amazon Glacier Storage Class class objects in a bucket and view the **Properties** of the selected objects, the **Properties** pane shows only the bucket name and the number of objects selected.

Editing Object Permissions

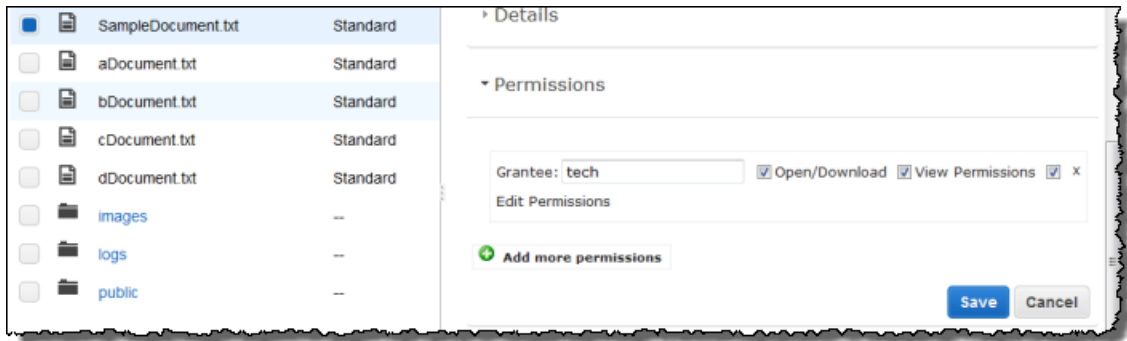
This section explains how to use the console to edit AWS account permissions for an object. In this topic, each permission you grant adds an entry in the Access Control List (ACL) associated with the object. You can grant permission to other AWS accounts or built-in groups. By default, the owner has full 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. This also applies if you grant anonymous write access to a bucket. Only the user `anonymous` can access objects the user created unless permission is explicitly granted to the bucket owner.

To change the 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. Click the object whose permissions you want to change, and then click **Permissions**.



3. Do one of the following:

To...	Do this...
To add permissions for a person or group	<ol style="list-style-type: none"> Click Add more permissions. In the Grantee box of the new line that appears, add the name of the person or group for which you want to set permissions. The name can be the email address associated with an AWS account, a canonical ID, or one of the predefined Amazon S3 groups. For a list of predefined Amazon S3 Groups, go to Who is a Grantee in the <i>Amazon Simple Storage Service Developer Guide</i>. You can add as many as 100 grantees. Select or clear the check boxes, as appropriate, next to the permissions you want to grant or deny.
To remove a person or group from the permission list	Click the "x" on the line of the grantee that you want to remove.

There are built-in groups that you can choose from the **Grantee** box:

- **Everyone**—Use this group to grant anonymous access.
- **Authenticated Users**—This group consists of any user that has an Amazon AWS Account. When you grant the Authenticated User group permission, any valid signed request can perform the appropriate action. The request can be signed by either an AWS Account or IAM User.
- **Log Delivery**—This group grants write access to your bucket when the bucket is used to store server access logs. For more information, see [Managing Bucket Logging](#).
- **Me**—This group refers to your AWS root account, and not an IAM user.

You can grant permission to an AWS account by entering the accounts canonical user ID or email address in the **Grantee** field. The email address must be the same one they used when signing up for an AWS account. You can grant a grantee any of the following permissions:

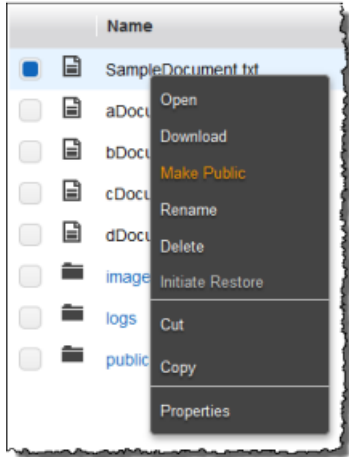
- **Open/Download**—Enables the account to access the object when they are logged in
- **View Permissions**—Can view the permissions associated with the object
- **Edit Permissions**—Can edit the permissions associated with the object

4. Click **Save**.


The console provides a shortcut for making objects accessible to everyone, meaning that everyone can both view and download the object.

To make an object accessible by everyone

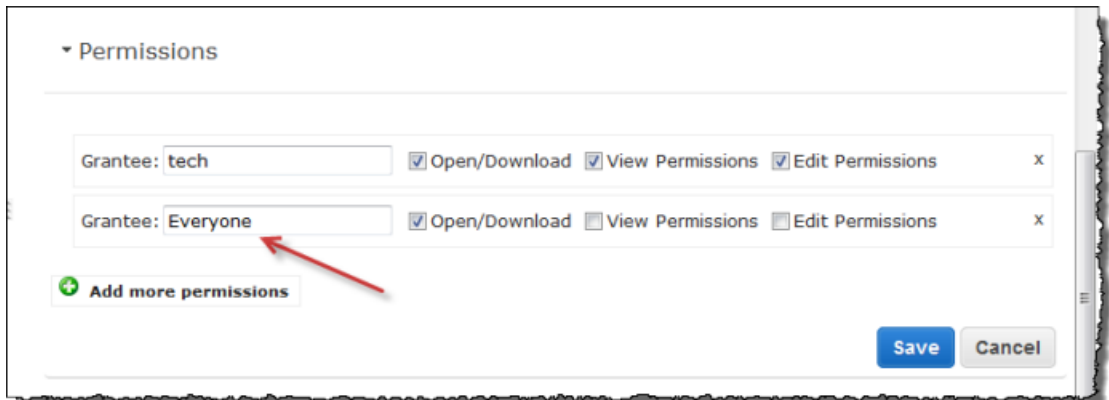
1. Right-click the object that you want to make accessible, and then click **Make Public**.



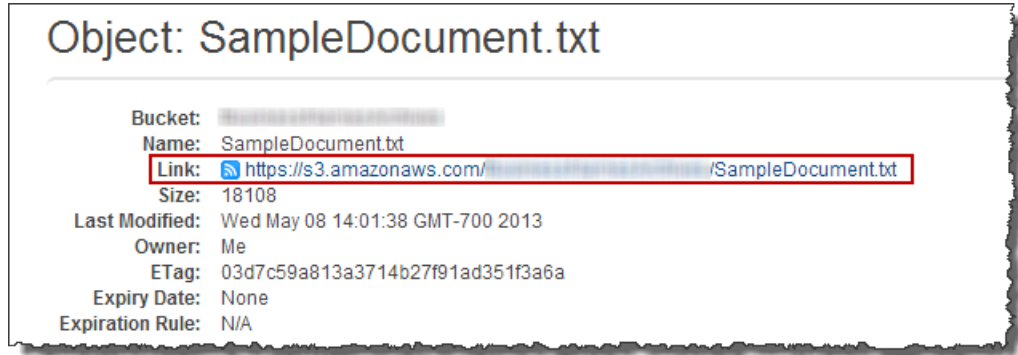
2. The console prompts you to confirm this change. Click **OK**. When the change is complete, click the Close button in the **Transfers** panel.



3. Click **Permissions**. The newly added grantee appears in the display.



4. Get the link for the object to share in the object properties pane as shown in the example below.



Editing Object Metadata

Each object in Amazon S3 has a set of key-value pairs that represents its metadata. There are two types of metadata:

- **System metadata** – Sometimes processed by Amazon S3, e.g., *Content-Type*, and *Content-Length*.
- **User metadata** – Never processed by Amazon S3.

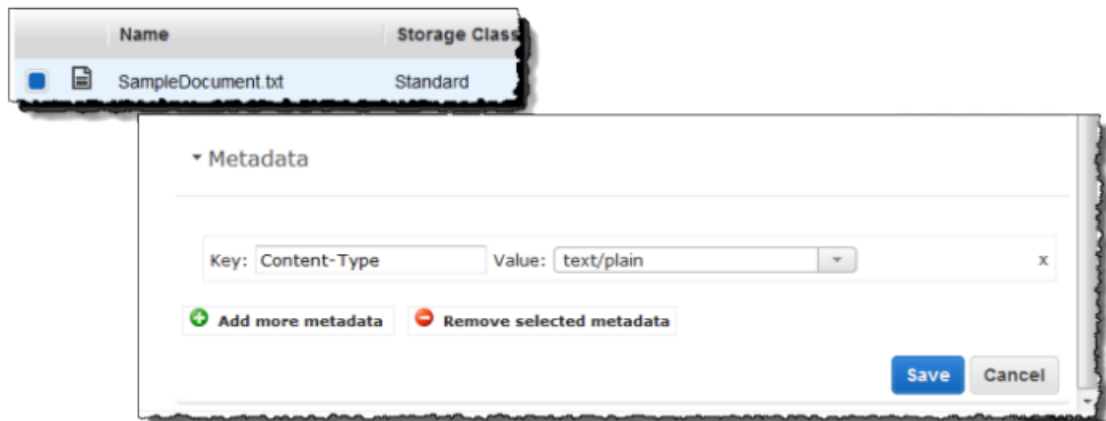
User metadata is stored with the object and returned with it.

The maximum size for user metadata is 2 KB, and both the keys and their values must conform to US-ASCII standards.

This section explains how to use the console to add and remove the metadata associated with an object.

To edit the metadata 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. Click the object whose metadata you want to edit, and then click **Metadata**.



3. Do one of the following:

To...	Do This...
Add metadata	<ol style="list-style-type: none">Click Add more metadata.In the Key box, click one of the available keys, or type a new one, starting with <code>x-amz-meta-</code> (for example, <code>x-amz-meta-<name></code>).In the corresponding Value box, click an entry in the list, if available, or type a value.
Delete metadata	<ol style="list-style-type: none">Click the key-value pair that you want to remove.Click Remove selected metadata, or click the "x" on the line of the key-value pair that you want to remove.

Note

User-defined metadata names must begin with "x-amz-meta-", otherwise Amazon S3 will not set the key value pair as you define it.

- Click **Save**.

Opening an Object

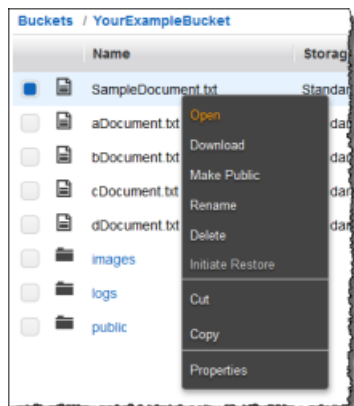
You can open an object to view it in a browser. This section explains how to use the console to open an object.

To open an object

- Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
- Right-click the object that you want to open, and then click **Open**.

Tip

You can use the **SHIFT** and **CTRL** keys to select multiple objects and perform the same action on all of them simultaneously.



Downloading an Object

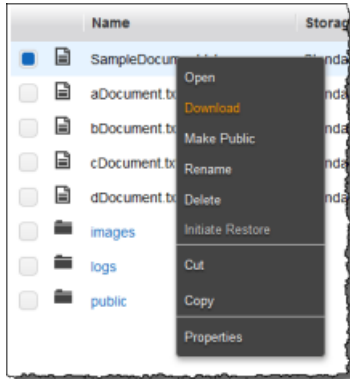
This section explains how to use the Amazon S3 console to download an object from Amazon S3 to your computer.

Note

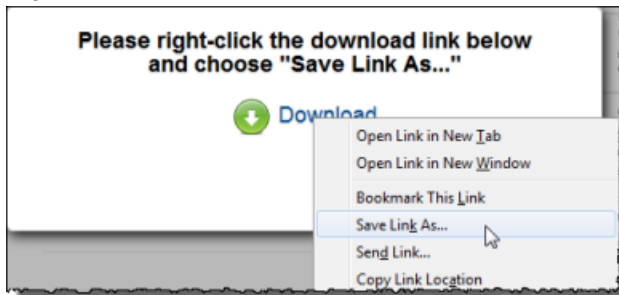
Data transfer fees apply when you download objects.

To download an object


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Right-click the object you want to download, and then click **Download**.



3. Right-click the word **Download**, and then click **Save Link As...**



4. Navigate to the folder on your system where you want to download the object, and then click **Save**.



When the download is complete, click **OK** to return to the console.



Copying an Object

You can also copy or move an object from one place to another by copying or cutting it from one place and pasting it in the new location.

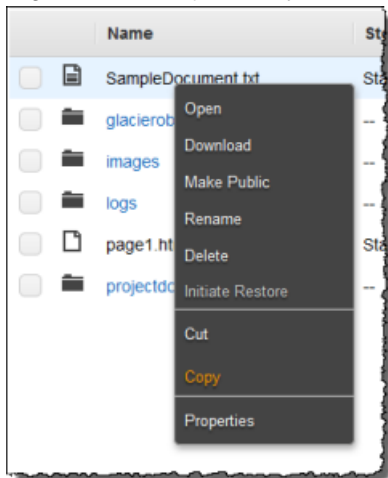
This section explains how to use the Amazon S3 console to copy an object.

Important

Copying and pasting objects protected by AWS Key Management Service (KMS) encryption keys into a new region is not supported in the Amazon S3 console. If you use the following procedure to transfer an AWS KMS protected object out of its home region, the transfer will fail. For more information on using AWS KMS encryption in Amazon S3, see [Protecting Data Using Server-Side Encryption with AWS KMS-Managed Keys \(SSE-KMS\)](#).

To copy an object


1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Right-click the object that you want to copy, and then click **Copy**.



Note

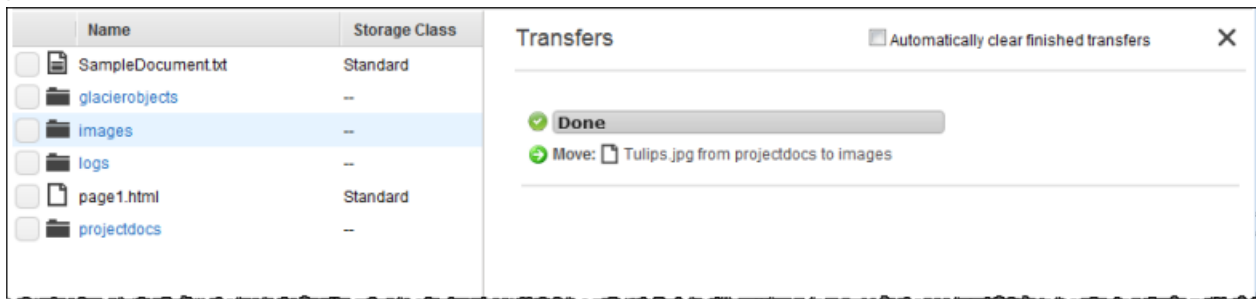
If you click **Cut** instead of **Copy**, you will move your file from its current location to another.

3. Navigate to the bucket and folder where you want to copy the object, right-click the target location, and then click **Paste Into**.



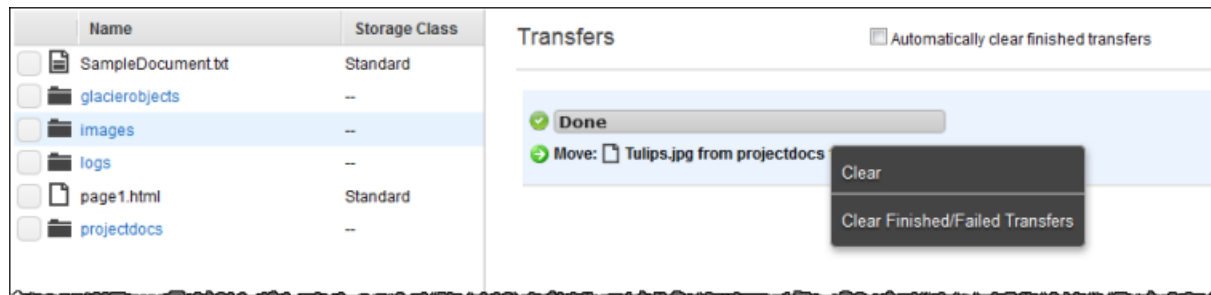
After you initiate the copy process, you must keep the browser open while the copy is in progress.

You can monitor the progress of the copy on the **Transfers** panel. To hide or show the **Transfers** panel, click the **Transfers** button on the console.



Note

To clear individual line items in the **Transfers** panel, right-click the items, and then click **Clear**. To remove all finished or failed transfers, click **Clear Finished/Failed Transfers**.

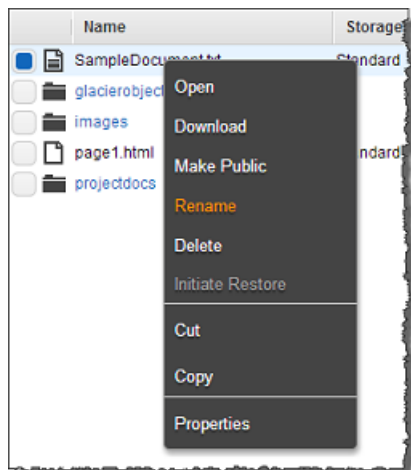


Renaming an Object


This section explains how to use the Amazon S3 console to rename an object. To rename multiple objects, rename each object separately.

To rename an object

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Right-click the object that you want to rename, and then click **Rename**.



3. In the box for the name, type a new name, and then click the checkmark icon to the right of the box to submit the name change.



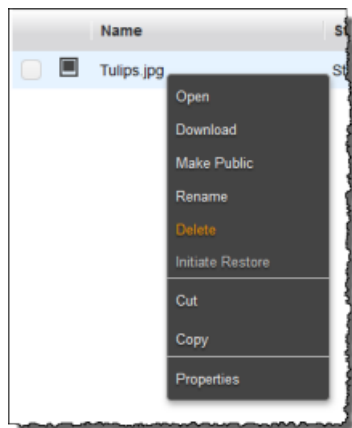
Deleting an Object

Because all objects in your Amazon 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 valuable.

This section explains how to use the Amazon S3 console to delete an object.

To delete 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 **Objects and Folders** list, right-click the object that you want to delete, and then click **Delete**.



3. When a confirmation message appears, click **OK**.

Deleting Objects by using Lifecycle Configuration Management

You can use Amazon S3 lifecycle configuration rules to schedule automatic deletions. For example, you might want to retain log files for 30 days, after which you want to delete them.

Amazon S3 manages object lifetimes with a lifecycle configuration, which is assigned to a bucket and defines rules for individual objects. You can, for example, apply a lifecycle configuration rule to all objects that begin with the prefix `log` to specify that Amazon S3 will delete such objects after 30 days. For more information, see [Managing Lifecycle Configuration \(p. 28\)](#).

Restoring an Object

Objects in the Amazon Glacier storage class are not immediately accessible: you must first restore a temporary copy of the object to its bucket before it is available. For information about when to use the `GLACIER` storage class for objects, go to [Object Lifecycle Management](#) in the *Amazon Simple Storage Service Developer Guide*. Restored objects are stored only for the number of days that you specify. You can modify the number of days an object is retained after it is restored. If you want a permanent copy of the object, create a copy of it within your Amazon S3 bucket.

This section explains how to use the Amazon S3 console to restore an object that is associated with the storage class `GLACIER`. It also provides procedures for both restoring and modifying the number of days.

Note

Amazon S3 calculates the restored date of an object by adding the number of days that you specify to the current time when you are restoring the object and then rounding the resulting time to the next day at midnight UTC. This calculation applies to the initial restoration of the object and to any time you modify the restored object's number of days. For example, if an object was restored on 10/15/2012 10:30 a.m. UTC and the number of days was specified as 3, then the object is restored until 10/19/2012 00:00 UTC. If, on 10/16/2012 11:00 a.m. UTC you change the number of days to 1, then the object is restored until 10/18/2012 00:00 UTC.


To restore an object

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

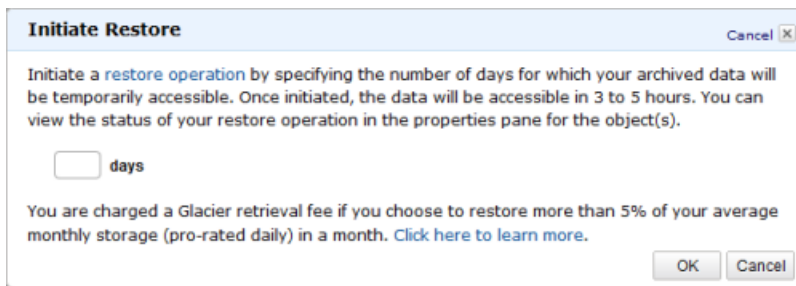
2. Right-click an object in the `GLACIER` storage class that you want to restore, and then click **Initiate Restore**.

Note

The menu shown in the following screenshot is slightly different if you have versioning enabled and you have the **Version: Hide/Show** button set to **Show**.



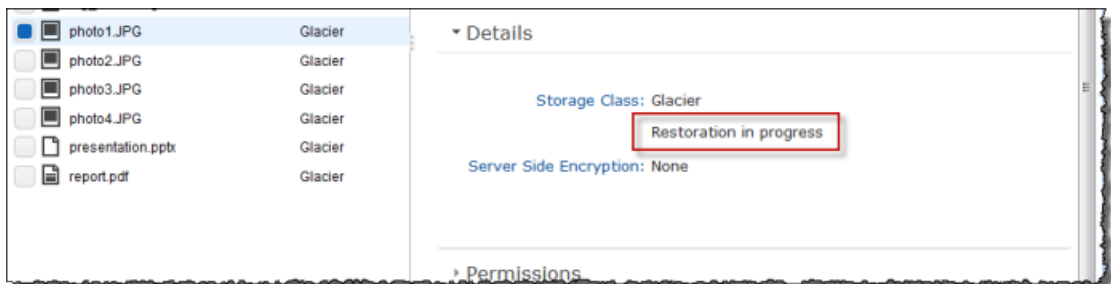
3. In the **Initiate Restore** dialog box, type the number of days until the restored object is deleted.



4. In the confirmation notice that appears, click **OK**.


Use the object **Details** pane to determine the status of the restoration. For more information, see [Editing Object Details \(p. 57\)](#).

The following example indicates that an object is in the process of being restored.



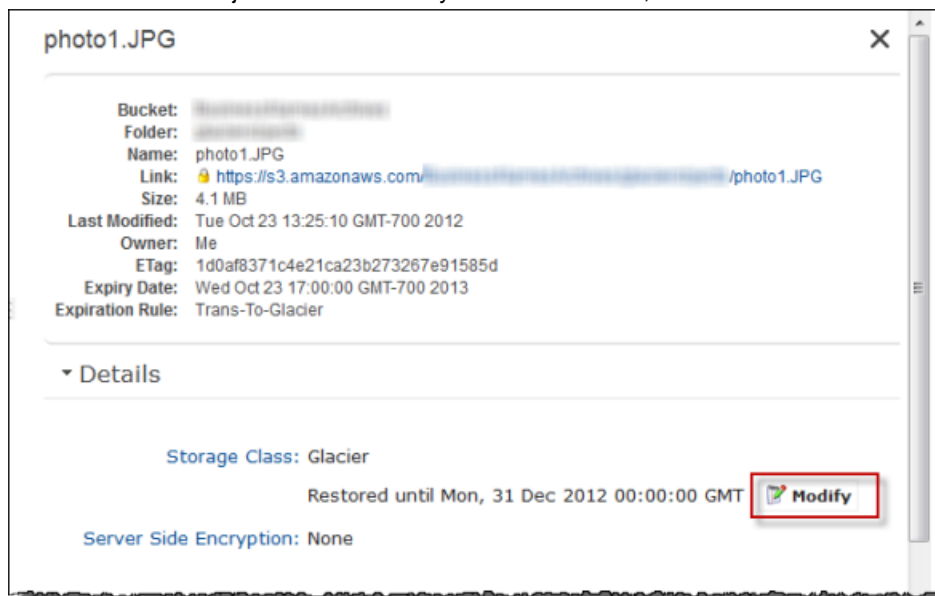
When the object is restored, the **Details** pane shows the date when the copy of object will be deleted.

The following example shows that an object is restored.

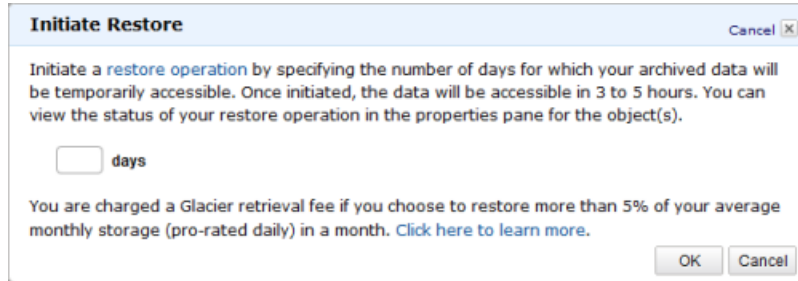


To extend the length of time of a restored object

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Click the restored object whose lifetime you want to extend, and then click **Details**.



3. Click **Modify**.
4. In the **Initiate Restore** dialog box, in the **days** box, type the number of days until the restored object is deleted.




5. In the confirmation message that appears, click **OK**. The **Restored until** date is changed.



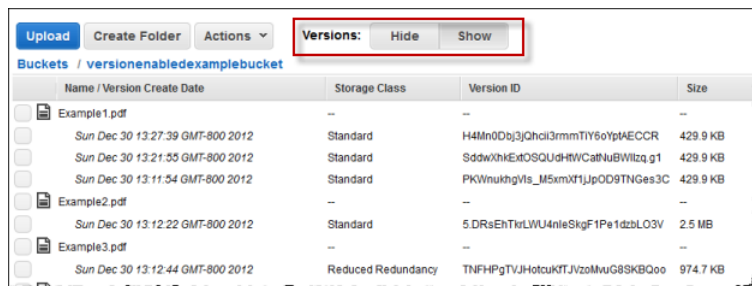
Managing Objects in a Versioning-Enabled Bucket

A versioning-enabled bucket can have multiple versions of objects in the bucket. Amazon S3 assigns each object a unique version ID. For more information about versioning support in Amazon S3, see [Using Versioning](#) in the *Amazon Simple Storage Service Developer Guide*.

When a bucket is versioning-enabled, you can show or hide all the object versions. The following example shows the list of objects in the `versionenabledexamplebucket` bucket. Version information is hidden, so these objects represent the latest version.



If you click **Show**, the console lists all the versions, as shown in the following example:



For each object version, the console shows a unique version ID, the date and time the object version was created, and other properties.

Uploading an Object

If you upload an object with a key name that already exists in the bucket, Amazon S3 creates another version of the object instead of replacing the existing object. For more information about uploading an object, see [Uploading Objects into Amazon S3](#) (p. 49).

Updating Object Properties

If you update any object properties after the initial object upload, such as changing the storage details or any other metadata changes, then Amazon S3 creates a new object version in the bucket. If you rename the object, Amazon S3 creates a new object version.

For example, if you update an object's storage class or change how the object is stored at rest by updating its server-side encryption property, Amazon S3 creates an object version for each property update you save.


When versions are hidden, you can update all the object properties; when versions are shown, you can update only the permissions for the specific object version.

For more information about updating object properties, see [Editing Object Properties](#) (p. 57).

Deleting Objects from a Versioning-Enabled Bucket

In a versioning-enabled bucket, you can either delete an object from the object list (version information hidden) or delete a specific version of the object.


With version information hidden, the console shows the object list as shown in the following example:



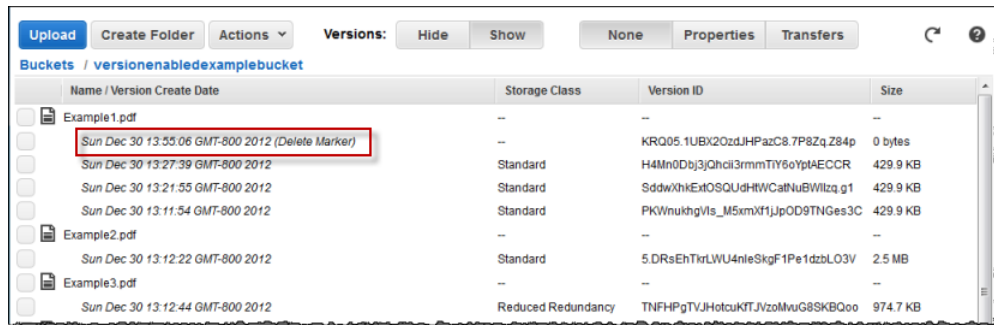
If you select and delete the Example1.pdf object, Amazon S3 adds a delete marker for the object and the object no longer appears in the object list:

Amazon Simple Storage Service Console User Guide

Deleting Objects from a Versioning-Enabled Bucket



However, if you click **Show** to list object versions, the Example1.pdf object appears in the list with all versions and a delete marker at the top.



To delete an object permanently, you must delete all the versions of the object, including the delete marker (if present). If you delete only a specific object version, Amazon S3 permanently deletes only that specific version. If you delete the delete marker, the object reappears in the object list. For more information, see [Deleting an Object](#) (p. 68).

Working with Folders

Topics

- [Public Folders \(p. 75\)](#)
- [Creating a Folder \(p. 76\)](#)
- [Deleting a Folder \(p. 76\)](#)

In Amazon S3, buckets and objects are the primary resources, where objects are stored in buckets. Amazon S3 has a flat structure with no hierarchy like you would see in a typical file system. However, for the sake of organizational simplicity, the Amazon S3 console supports the folder concept as a means of grouping objects. Amazon S3 does this by using key name prefixes for objects.

For example, you can create a folder in the console called `photos`, and store an object called `myphoto.jpg` in it. The object is then stored with the key name `photos/myphoto.jpg`, where `photos/` is the prefix.

Here are two more examples:

- If you have three objects in your bucket—`logs/date1.txt`, `logs/date2.txt`, and `logs/date3.txt`—the console will show a folder named `logs`. If you open the folder in the console, you will see three objects: `date1.txt`, `date2.txt`, and `date3.txt`.
- If you have an object named `photos/2013/example.jpg`, the console will show you a folder named `photos` containing the folder `2013` and the object `example.jpg`.

You can have folders within folders, but not buckets within buckets. You can upload and copy objects directly into a folder. Folders can be created, deleted, and made public, but they cannot be renamed. Objects can be moved from one folder to another. For more information about moving objects, see [Support for Moving Data \(p. 5\)](#).

Public Folders

You can make folders public, which means that all of the objects that appear within a public folder in the console are available for viewing or downloading to anyone on the Internet. However, as mentioned previously, the folder concept is only supported in the console. If you use a web browser to view a folder that you made public, you will get an access denied error because the folder is just a naming prefix, for an object or group of objects.

Note


It is easy to make a folder public, but you cannot make a folder private after you make it public. To make the objects in a public folder private, you have to go through each object in the public folder that you want to make private and set the permissions individually. For more information about how to set an object's permissions, see [Editing Object Permissions \(p. 60\)](#).

Creating a Folder

This section describes how to use the console to create a folder.

To create a folder

1. Click the bucket in the **All Buckets** list in which you want to create a folder.
2. Click **Create Folder**.



3. Under **Name**, in the box that appears, type a name for the folder, and then click the check mark.

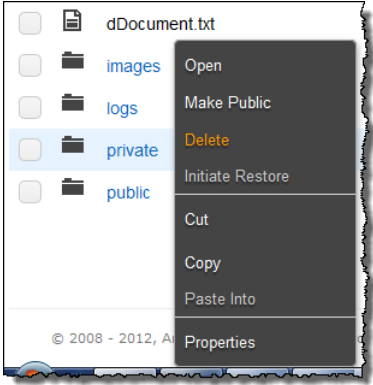
Deleting a Folder

This section describes how to use the console to delete a folder.

Caution

When you delete a folder, any objects or folders contained in the folder will be automatically deleted. If you want to retain those objects, you must move them elsewhere before you delete the folder. For information about moving objects, see [Copying an Object](#).

1. In the **Objects and Folders** list, right-click the folder that you want to delete, and then click **Delete**.



2. When a confirmation message appears, click **OK**.

Amazon S3 Resources

Following is a table that lists related resources that you'll find useful as you work with this service.

Resource	Description
Amazon Simple Storage Service Getting Started Guide	The <i>Amazon Simple Storage Service Getting Started Guide</i> provides a quick tutorial of the service using the AWS Management Console to accomplish basic Amazon S3 tasks.
Amazon Simple Storage Service API Reference	The <i>Amazon Simple Storage Service API Reference</i> describes Amazon S3 operations in detail.
Amazon Simple Storage Service Developer Guide	The <i>Amazon Simple Storage Service Developer Guide</i> describes how to use Amazon S3 operations.
Amazon S3 Technical FAQ	The FAQ covers the top 20 questions developers have asked about this product.
Amazon S3 Release Notes	The Release Notes give a high-level overview of the current release. They specifically note any new features, corrections, and known issues.
AWS Home Page	A central starting point to find documentation, code samples, release notes, and other information to help you build innovative applications with AWS.
AWS Management Console	The console allows you to perform Amazon S3 functions using a simple and intuitive web user interface.
Discussion Forums	A community-based forum for developers to discuss technical questions related to AWS.
AWS Support Center	The home page for AWS Technical Support, including access to our Developer Forums, Technical FAQs, Service Status page, and Premium Support.
AWS Premium Support	The primary web page for information about AWS Premium Support, a one-on-one, fast-response support channel to help you build and run applications on AWS Infrastructure Services.
Amazon S3 product information	The primary web page for information about Amazon S3.

Resource	Description
Amazon S3 pricing information	The primary web page for information about Amazon S3 pricing.
Contact Us	A central contact point for inquiries concerning AWS billing, account, events, abuse etc.
Conditions of Use	Detailed information about the copyright and trademark usage at Amazon.com and other topics.

Document History

The following table describes the important changes to the documentation since the last release of Amazon S3.

- **API version:** 2006-03-01
- **Latest documentation update:** August 11, 2015

Change	Description	Date Changed
Lifecycle configuration update	The Amazon S3 lifecycle configuration rules user interface has been updated to improve usability. For more information, see Lifecycle Configuration for a Bucket without Versioning (p. 29) and Lifecycle Configuration for a Bucket with Versioning (p. 32) .	In this release
Event notifications	Amazon S3 event notifications have been updated to add notifications when objects are deleted and to add filtering on object names with prefix and suffix matching. For more information, see Enabling Event Notifications (p. 21) .	July 28, 2015
Support for deleting and emptying non-empty buckets	Amazon S3 now supports deleting and emptying non-empty buckets. For more information, see Deleting or Emptying an Amazon S3 Bucket (p. 12) .	July 16, 2015
Event notifications	Event notifications have been updated in the Amazon S3 console to support the switch to resource-based permissions for AWS Lambda functions. For more information, see Enabling Event Notifications (p. 21) .	April 9, 2015
Cross-region replication	The Amazon S3 console now supports cross-region replication. Cross-region replication is the automatic, asynchronous copying of objects across buckets in different AWS regions. For more information, see Managing Cross-Region Replication (p. 42) .	March 24, 2015

Change	Description	Date Changed
Event notifications	Amazon S3 now supports new event types and destinations in a bucket notification configuration. Prior to this release, Amazon S3 supported only the <i>s3:ReducedRedundancyLostObject</i> event type and an Amazon SNS topic as the destination. For more information about the new event types, go to Setting Up Notification of Bucket Events .	November, 13, 2014
Amazon S3 now supports lifecycle rules for versioning	The Amazon S3 console now supports lifecycle configuration rules for buckets with versioning. For more information see, Managing Lifecycle Configuration (p. 28) .	May 20, 2014
Console support for enabling bucket versioning	The Amazon S3 console now supports bucket versioning and managing objects in a versioning-enabled bucket. For more information see, Enabling Bucket Versioning (p. 27) , and Managing Objects in a Versioning-Enabled Bucket (p. 72) .	December 31, 2012
Support for static website hosting at the root domain	<p>Amazon S3 now supports hosting static websites at the root domain. Visitors to your website can access your site from their browser without specifying "www" in the web address (e.g., "example.com"). Many customers already host static websites on Amazon S3 that are accessible via a "www" subdomain (e.g., "www.example.com"). Previously, to support root domain access, you needed to run your own web server to proxy root domain requests from browsers to your website on Amazon S3. Running a web server to proxy requests introduces additional costs, operational burden, and another potential point of failure. Now, you can take advantage of the high availability and durability of Amazon S3 for both "www" and root domain addresses.</p> <p>For an example walkthrough, go to Example: Setting Up a Static Website Using a Custom Domain. For conceptual information, go to Hosting Static Websites on Amazon S3 in the <i>Amazon Simple Storage Service Developer Guide</i>.</p>	December 27, 2012
Console revision	Amazon S3 console has been updated. The documentation topics that refer to the console have been revised accordingly.	December 14, 2012
Support for Archiving Data to Amazon Glacier	<p>Amazon S3 now support a storage option that enables you to utilize Amazon Glacier's low-cost storage service for data archival. To archive objects, you define archival rules identifying objects and timeline when you want Amazon S3 to archive these objects to Amazon Glacier. You can easily set the rules on a bucket using the Amazon S3 console or programmatically using the Amazon S3 API or AWS SDKs.</p> <p>In addition to setting object expiration, you can now use lifecycle management to archive data in Amazon S3. For more information, see Managing Lifecycle Configuration (p. 28).</p> <p>For conceptual information, go to Object Lifecycle Management in the <i>Amazon Simple Storage Service Developer Guide</i>.</p>	November 13, 2012

Change	Description	Date Changed
Cross-Origin Resource Sharing (CORS) support	Amazon S3 now supports Cross-Origin Resource Sharing (CORS). CORS defines a way in which client web applications that are loaded in one domain can interact with or access resources in a different domain. With CORS support in Amazon S3, you can build rich client-side web applications on top of Amazon S3 and selectively allow cross-domain access to your Amazon S3 resources. For more information, see Enabling Cross-Origin Resource Sharing in the <i>Amazon Simple Storage Service Developer Guide</i> .	August 31, 2012
AWS Cost Allocation Tagging support	You can use AWS Cost Allocation to control how storage resources are organized on your bill. You do this by defining one or more tags for a bucket. For more information, go to Cost Allocation Tagging in the <i>Amazon Simple Storage Service Developer Guide</i> .	August 21, 2012
Object Expiration support	You can use Object Expiration to schedule automatic removal of data after a configured time period. You set object expiration by adding lifecycle configuration to a bucket. For more information, go to Object Expiration .	December 27, 2011
New region supported	Amazon S3 now supports the South America (Sao Paulo) region. For more information, go to Regions and Endoints in Amazon Web Services General Reference.	December 14, 2011
New region supported	Amazon S3 now supports the US West (Oregon) region. For more information, go to Regions and Endoints in Amazon Web Services General Reference.	November 8, 2011
Documentation Update	This release includes enhancements to the object properties related sections. Information about what the Details properties tab show when you select one or more objects. For more information, see Editing Object Properties (p. 57) .	October 17, 2011
Support for server-side encryption in Amazon S3	This release includes support for server-side encryption in the Amazon S3 console. You can now specify that data stored in Amazon S3 is encrypted at rest. When you upload objects to Amazon S3 using the console, you can choose server-side encryption for your data. For more information, see Uploading Objects into Amazon S3 (p. 49) . For more information about server-side encryption for data stored in Amazon S3, see Using Server-Side Encryption in the <i>Amazon S3 Developer Guide</i> .	October 5, 2011
AWS Management Console enhancements	This release includes the following AWS Management Console enhancements: <ul style="list-style-type: none"> • Folder upload—You can now use AWS Management Console to upload folders into Amazon S3. Amazon S3 uploads all the files, and subfolders from the specified folder to your bucket. For more information, see Uploading Objects into Amazon S3 (p. 49) • Jump feature—Instead of scrolling through a long list to find an object or folder, you can now simply start typing the first few characters of an object or folder name into the browser when looking at a listing. The console will jump to objects that match or follow what you type. For more information, see Browsing the Objects in Your Bucket (p. 14) 	June 6, 2011

Change	Description	Date Changed
Support for hosting static websites in Amazon S3	Amazon S3 introduces enhanced support for hosting static websites. This includes support for index documents and custom error documents. When using these features, requests to the root of your bucket or a subfolder (e.g., http://mywebsite.com/subfolder) returns your index document instead of the list of objects in your bucket. If an error is encountered, Amazon S3 returns your custom error message instead of an Amazon S3 error message. For information on managing website configuration using the AWS Management Console, see Configuring a Bucket for Website Hosting (p. 19) . For more information about Amazon S3's website configuration feature, go to Hosting Websites on Amazon S3 in the <i>Amazon Simple Storage Service Developer Guide</i> .	February 17, 2011
Large object support	Now, you can use AWS Management Console to upload large objects, up to 5 TB each, to an Amazon S3 bucket.	December 9, 2010
Bucket notifications in the console	Now, you can configure bucket properties to enable notifications. These notifications are posted to Amazon SNS (SNS) topic in the event a Reduced Redundancy Storage (RRS) object is lost from the bucket.	September 8, 2010
Bucket policies in the console	Now, you can add and edit Amazon S3 bucket policies using the AWS Management Console. You can access bucket policies in the AWS Management Console by viewing the properties of the specific bucket. Using bucket policies, you can define security rules that apply to all objects or a subset of objects within a bucket. This makes updating and managing permissions easier.	August 13, 2010
New Guide	This is the first release of the <i>Amazon Simple Storage Service Console User Guide</i> . It describes how to use Amazon S3 in the AWS Management Console.	June 8, 2010

AWS Glossary

For the latest AWS terminology, see the [AWS Glossary](#) in the *AWS General Reference*.