Amazon Simple Storage Service

Console User Guide

Amazon Simple Storage Service: Console User Guide

Copyright © 2020 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

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

Table of Contents

Welcome to the Amazon S3 Console User Guide 1
Changing the console language 2
Creating and configuring a bucket 3
- Creating a bucket 3
  - More info 4
- Deleting a bucket 5
  - More info 5
- Emptying a bucket 6
- Viewing bucket properties 6
- Enabling or disabling versioning 7
- Enabling default encryption 8
  - More info 11
- Enabling server access logging 11
- Enabling object-level logging 14
  - More info 16
- Configuring static website hosting 16
  - Step 1: Configure an Amazon S3 bucket for static website hosting 17
  - Step 2: Editing block public access settings 18
  - Step 3: Adding a bucket policy 20
  - Step 4: Test your website endpoint 21
- Redirecting website requests 21
- Advanced settings 22
  - Setting up a destination for event notifications 22
  - Enabling and configuring event notifications 23
  - Enabling transfer acceleration 28

Access points 30
- Creating an Amazon S3 access point 30
- Managing and using Amazon S3 access points 31

Uploading, downloading, and managing objects 33
- Uploading S3 objects 34
  - Uploading Files and Folders by Using Drag and Drop 35
  - Uploading Files by Pointing and Clicking 39
  - More Info 41
- Copying an object 42
- Moving an object 42
- Downloading S3 objects 43
  - Related topics 46
- Deleting objects 46
  - More info 47
- Undeleting objects 47
  - More info 47
- Restoring archived S3 objects 47
  - Archive Retrieval Options 48
  - Restoring an Archived S3 Object 48
  - Upgrade an In-Progress Restore 51
  - Checking Archive Restore Status and Expiration Date 52
- Locking Amazon S3 objects 53
  - More info 55
- Viewing an overview of an object 55
  - More info 57
- Viewing object versions 58
  - More info 59
- Viewing object properties 59
- Adding encryption to an object 61

More info ... 63
Adding metadata to an object ... 63
Adding system-defined metadata ... 64
Adding user-defined metadata ... 67
Adding tags to an object ... 70
More info ... 72
Using folders ... 72
Creating a folder ... 73
Deleting folders ... 74
Making folders public ... 76
S3 Batch Operations ... 77
Creating an S3 Batch Operations job ... 77
More info ... 77
Managing S3 Batch Operations jobs ... 78
More info ... 78
Storage management ... 79
Creating a lifecycle policy ... 79
Creating replication rules ... 83
Adding a replication rule when the destination bucket is in the same AWS account ... 84
Adding a replication rule when the destination bucket is in a different AWS account ... 91
More info ... 98
Managing replication rules ... 98
More info ... 100
Configuring Storage Class Analysis ... 100
Configuring Amazon S3 Inventory ... 104
Destination Bucket Policy ... 106
Granting Amazon S3 Permission to Use Your AWS KMS CMK for Encryption ... 107
Creating a request metrics filter for a bucket ... 108
Creating a request metrics filter using object tags or prefixes ... 108
Deleting a request metrics filter ... 109
Viewing replication metrics ... 110
Setting permissions ... 112
Blocking public access ... 113
Access status ... 113
More info ... 114
Editing bucket public access settings ... 114
Editing public access settings for an S3 bucket ... 114
Editing public access settings for multiple S3 buckets ... 115
More info ... 116
Editing account public access settings ... 117
More info ... 117
Setting object permissions ... 118
More Info ... 120
Setting ACL bucket permissions ... 121
More info ... 123
Adding a bucket policy ... 124
More info ... 125
Adding cross-domain resource sharing with CORS ... 125
More info ... 126
Setting Object Ownership to bucket owner preferred ... 126
How do I ensure that I take ownership of new objects? ... 127
Using Access Analyzer for S3 ... 127
What information does Access Analyzer for S3 provide? ... 128
Enabling Access Analyzer for S3 ... 128
Blocking all public access ... 129
Reviewing and changing bucket access ... 129
Archiving bucket findings ... 130

Activating an archived bucket finding 131
Viewing finding details 131
Downloading an Access Analyzer for S3 report 131
Document history 133
Earlier updates 133
AWS glossary 136

Welcome to the Amazon S3 Console User Guide

Welcome to the *Amazon Simple Storage Service Console User Guide* for the Amazon Simple Storage Service (Amazon S3) console.

Amazon S3 provides virtually limitless storage on the internet. This guide explains how you can manage buckets, objects, and folders in Amazon S3 by using the AWS Management Console, a browser-based graphical user interface for interacting with AWS services.

For detailed conceptual information about how Amazon S3 works, see What Is Amazon S3? in the *Amazon Simple Storage Service Developer Guide*. The developer guide also has detailed information about Amazon S3 features and code examples to support those features.

Topics

- Creating and configuring an S3 bucket (p. 3)
- Uploading, downloading, and managing objects (p. 33)
- Storage management (p. 79)
- Setting bucket and object access permissions (p. 112)

How do I change the language of the AWS Management Console?

You can change the display language of the AWS Management Console. Several languages are supported.

To change the console language

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. Scroll down until you see the bar at the bottom of the window, and then choose the language on the left side of the bar.

3. Choose the language that you want from the menu. This will change the language for the entire AWS Management Console.

Creating and configuring an S3 bucket

To upload your data (photos, videos, documents etc.) to Amazon S3, you must first create an S3 bucket in one of the AWS Regions. You can then upload your data objects to the bucket.

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

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

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

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

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

Topics

- How do I create an S3 Bucket? (p. 3)
- How do I delete an S3 Bucket? (p. 5)
- How do I empty an S3 Bucket? (p. 6)
- How do I view the properties for an S3 Bucket? (p. 6)
- How do I enable or suspend versioning for an S3 bucket? (p. 7)
- How do I enable default encryption for an Amazon S3 bucket? (p. 8)
- How do I enable server access logging for an S3 bucket? (p. 11)
- How do I enable object-level logging for an S3 bucket with AWS CloudTrail data events? (p. 14)
- How do I configure an S3 bucket for static website hosting? (p. 16)
- How do I redirect requests to an S3 bucket hosted website to another host? (p. 21)
- Advanced settings for S3 bucket properties (p. 22)

How do I create an S3 Bucket?

Before you can upload data to Amazon S3, you must create a bucket in one of the AWS Regions to store your data. After you create a bucket, you can upload an unlimited number of data objects to the bucket.

The AWS account that creates the bucket owns it. By default, you can create up to 100 buckets in each of your AWS accounts. If you need additional buckets, you can increase your account bucket quota to a maximum of 1,000 buckets by submitting a service quota increase. For information about how to increase your bucket quota, see AWS Service Quotas in the *AWS General Reference*.

Buckets have configuration properties, including geographical Region, access settings for the objects in the bucket, and other metadata.

To create a bucket

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. Choose **Create bucket**.
3. In **Bucket name**, enter a DNS-compliant name for your bucket.

   The bucket name must:

   - Be unique across all of Amazon S3.
   - Be between 3 and 63 characters long.
   - Not contain uppercase characters.
   - Start with a lowercase letter or number.

   After you create the bucket, you can't change its name. For information about naming buckets, see Rules for bucket naming in the *Amazon Simple Storage Service Developer Guide*.

   > **Important**
   > Avoid including sensitive information, such as account numbers, in the bucket name. The bucket name is visible in the URLs that point to the objects in the bucket.

4. In **Region**, choose the AWS Region where you want the bucket to reside.

   Choose a Region close to you to minimize latency and costs and address regulatory requirements. Objects stored in a Region never leave that Region unless you explicitly transfer them to another Region. For a list of Amazon S3 AWS Regions, see AWS service endpoints in the *Amazon Web Services General Reference*.
5. In **Bucket settings for Block Public Access**, choose the Block Public Access settings that you want to apply to the bucket.

   We recommend that you leave all settings enabled unless you know you need to turn one or more of them off for your use case, such as to host a public website. Block public access settings that you enable for the bucket will also be enabled for all access points that you create on the bucket. For more information about blocking public access, see Using Amazon S3 Block Public Access in the *Amazon Simple Storage Service Developer Guide*.
6. (Optional) If you want to enable S3 Object Lock:

   a. Choose **Advanced settings**, and read the message that appears.

      > **Important**
      > You can only enable S3 Object Lock for a bucket when you create it. If you enable Object Lock for the bucket, you can't disable it later. Enabling Object Lock also enables versioning for the bucket. After you enable Object Lock for the bucket, you must configure the Object Lock settings before any objects in the bucket will be protected. For more information about configuring protection for objects, see How do I lock an Amazon S3 object? (p. 53).

   b. If you want to enable Object Lock, enter *enable* in the text box and choose **Confirm**.

   For more information about the S3 Object Lock feature, see Locking Objects Using Amazon S3 Object Lock in the *Amazon Simple Storage Service Developer Guide*.
7. Choose **Create bucket**.

More info

- How do I delete an S3 Bucket? (p. 5)

- How do I set ACL bucket permissions? (p. 121)

How do I delete an S3 Bucket?

You can delete an empty bucket, and when you're using the AWS Management Console, you can delete a bucket that contains objects. If you delete a bucket that contains objects, all the objects in the bucket are permanently deleted.

When you delete a bucket with versioning enabled, all versions of all the objects in the bucket are permanently deleted. For more information about versioning, see Managing Objects in a Versioning-Enabled Bucket in the *Amazon Simple Storage Service Developer Guide*.

Before deleting a bucket, consider the following:

- Bucket names are unique. If you delete a bucket, another AWS user can use the name.
- When you delete a bucket that contains objects, all the objects in the bucket are permanently deleted, including objects that transitioned to the `S3 Glacier` storage class.
- If the bucket hosts a static website, and you created and configured an Amazon Route 53 hosted zone as described in Create and Configure Amazon Route 53 Hosted Zone: You must clean up the Route 53 hosted zone settings that are related to the bucket as described in Delete the Route 53 Hosted Zone.
- If the bucket receives log data from Elastic Load Balancing (ELB): We recommend that you stop the delivery of ELB logs to the bucket before deleting it. After you delete the bucket, if another user creates a bucket using the same name, your log data could potentially be delivered to that bucket. For information about ELB access logs, see Access Logs in the *User Guide for Classic Load Balancers* and Access Logs in the *User Guide for Application Load Balancers*.

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

To delete an S3 bucket

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the **Buckets** list, select the option next to the name of the bucket that you want to delete and then choose **Delete** at the top of the page.
3. On the **Delete bucket** page, confirm that you want to delete the bucket by entering the bucket name into the text field, and then choose **Delete bucket**.

   > **Note**
   > If the bucket contains any objects, empty the bucket before deleting it by selecting the *empty bucket configuration* link in the **This bucket is not empty** error alert and following the instructions on the **Empty bucket** page. Then return to the **Delete bucket** page and delete the bucket.

More info

- How do I empty an S3 Bucket? (p. 6)
- How do I delete objects from an S3 bucket? (p. 46)

How do I empty an S3 Bucket?

You can empty a bucket, which deletes all of the objects in the bucket without deleting the bucket. When you empty a bucket with versioning enabled, all versions of all the objects in the bucket are deleted. For more information, see Managing Objects in a Versioning-Enabled Bucket and Deleting/Emptying a Bucket in the *Amazon Simple Storage Service Developer Guide*.

To empty an S3 bucket

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the **Bucket name** list, select the option next to the name of the bucket that you want to empty and then choose **Empty**.
3. On the **Empty bucket** page, confirm that you want to empty the bucket by entering the bucket name into the text field, and then choose **Empty**.
4. (Optional) Monitor the progress of the bucket emptying process on the **Empty bucket: Status** page.

   Warning
   This action deletes all objects in the bucket. Wait for the empty bucket action to finish before adding new objects. New objects might be deleted if they are added while the empty bucket action is in progress.

How do I view the properties for an S3 Bucket?

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

To view the properties for an S3 bucket

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

3. Choose **Properties**.

4. On the **Properties** page, you can configure the following properties for the bucket.

   a. **Versioning** – Versioning enables you to keep multiple versions of an object in one bucket. By default, versioning is disabled for a new bucket. For information about enabling versioning, see How do I enable or suspend versioning for an S3 bucket? (p. 7).

   b. **Server access logging** – Server access logging provides detailed records for the requests that are made to your bucket. By default, Amazon S3 does not collect server access logs. For information about enabling server access logging, see How do I enable server access logging for an S3 bucket? (p. 11).

c. **Static website hosting** – You can host a static website on Amazon S3. To enable static website hosting, choose **Static website hosting** and then specify the settings you want to use. For more information, see How do I configure an S3 bucket for static website hosting? (p. 16).

d. **Object-level logging** – Object-level logging records object-level API activity by using CloudTrail data events. For information about enabling object-level logging, see How do I enable object-level logging for an S3 bucket with AWS CloudTrail data events? (p. 14).

e. **Tags** – With AWS cost allocation, you can use bucket tags to annotate billing for your use of a bucket. A tag is a key-value pair that represents a label that you assign to a bucket. To add tags, choose **Tags**, and then choose **Add tag**. For more information, see Using Cost Allocation Tags for S3 Buckets in the *Amazon Simple Storage Service Developer Guide*.

f. **Transfer acceleration** – Amazon S3 Transfer Acceleration enables fast, easy, and secure transfers of files over long distances between your client and an S3 bucket. For information about enabling transfer acceleration, see How do I enable transfer acceleration for an S3 bucket? (p. 28).

g. **Events** – You can enable certain Amazon S3 bucket events to send a notification message to a destination whenever the events occur. To enable events, choose **Events** and then specify the settings you want to use. For more information, see How do I enable and configure event notifications for an S3 Bucket? (p. 23).

h. **Requester Pays** – You can enable Requester Pays so that the requester (instead of the bucket owner) pays for requests and data transfers. For more information, see Requester Pays Buckets in the *Amazon Simple Storage Service Developer Guide*.

How do I enable or suspend versioning for an S3 bucket?

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

To enable or disable versioning on an S3 bucket

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

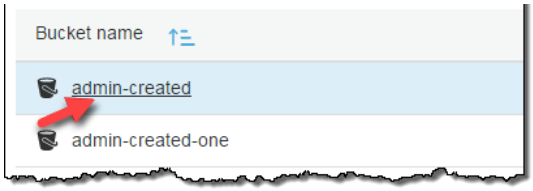

3. Choose **Properties**.

4. Choose **Versioning**.

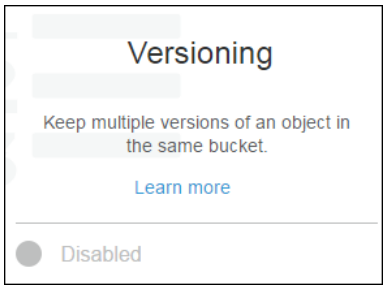

5. Choose **Enable versioning** or **Suspend versioning**, and then choose **Save**.

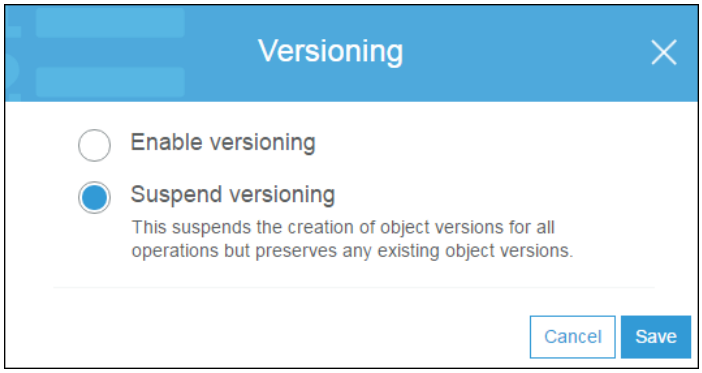

Note
You can use AWS Multi-Factor Authentication (MFA) with versioning. When you use MFA with versioning, you must provide your AWS account's access keys and a valid code from the account's MFA device in order to permanently delete an object version or suspend or reactivate versioning. To use MFA with versioning, you enable `MFA Delete`. However, you cannot enable `MFA Delete` using the AWS Management Console. You must use the AWS CLI or API. For more information, see MFA Delete.

How do I enable default encryption for an Amazon S3 bucket?

Amazon S3 default encryption provides a way to set the default encryption behavior for an Amazon S3 bucket. You can set default encryption on a bucket so that all objects are encrypted when they are stored in the bucket. The objects are encrypted using server-side encryption with either Amazon S3-managed keys (SSE-S3) or AWS Key Management Service (AWS KMS) customer master keys (CMKs).

When you use server-side encryption, Amazon S3 encrypts an object before saving it to disk in its data centers and decrypts it when you download the objects. For more information about protecting data using server-side encryption and encryption key management, see Protecting Data Using Server-Side Encryption in the *Amazon Simple Storage Service Developer Guide*.

Default encryption works with all existing and new Amazon S3 buckets. Without default encryption, to encrypt all objects stored in a bucket, you must include encryption information with every object storage request. You must also set up an Amazon S3 bucket policy to reject storage requests that don't include encryption information.

There are no new charges for using default encryption for S3 buckets. Requests to configure the default encryption feature incur standard Amazon S3 request charges. For information about pricing, see

Amazon S3 Pricing. For SSE-KMS CMK storage, AWS KMS charges apply and are listed at AWS KMS Pricing.

To enable default encryption on an Amazon S3 bucket

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

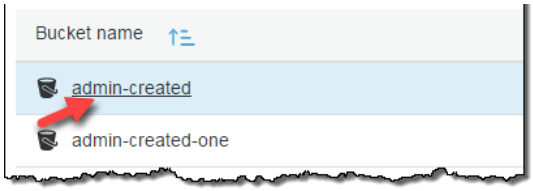

3. Choose **Properties**.

4. Choose **Default encryption**.

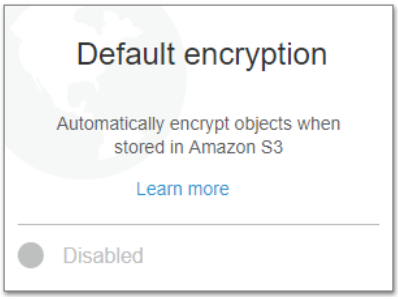

5. If you want to use keys that are managed by Amazon S3 for default encryption, choose **AES-256**, and choose **Save**.

   For more information about using Amazon S3 server-side encryption to encrypt your data, see Protecting Data with Amazon S3-Managed Encryption Keys in the *Amazon Simple Storage Service Developer Guide*.

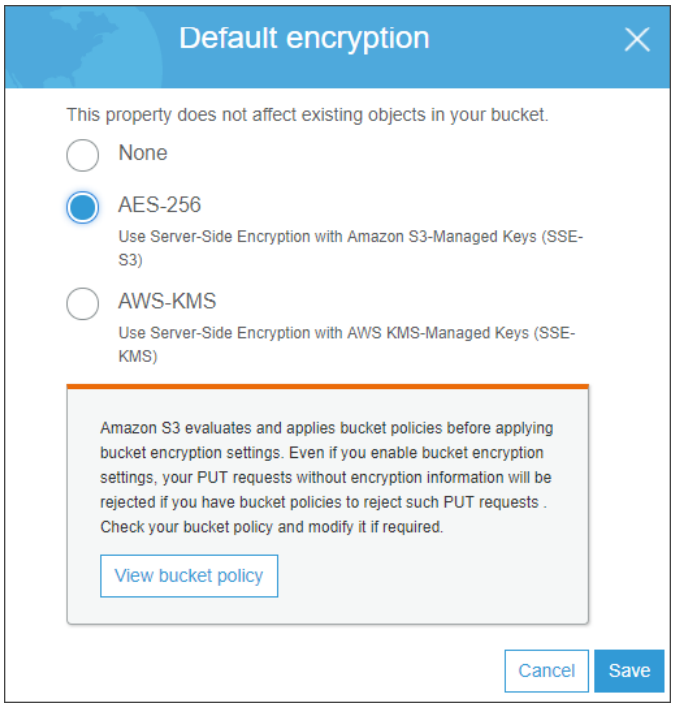

Important
You might need to update your bucket policy when enabling default encryption. For more information, see Moving to Default Encryption from Using Bucket Policies for Encryption Enforcement in the *Amazon Simple Storage Service Developer Guide*.

6. If you want to use CMKs that are stored in AWS KMS for default encryption, follow these steps:

   a. Choose **AWS-KMS**.

   b. To choose a customer-managed AWS KMS CMK that you have created, use one of these methods:

      - In the list that appears, choose the AWS KMS CMK.
      - In the list that appears, choose **Custom KMS ARN**, and then enter the Amazon Resource Name of the AWS KMS CMK.

      Important
      You can use only KMS CMKs that are enabled in the same AWS Region as the bucket. The S3 console only lists 100 KMS CMKs per Region. If you have more than 100 CMKs in the same Region, you can only see the first 100 CMKs in the S3 console. To use a KMS CMK that is not listed in the console, choose **Custom KMS ARN**, and enter the KMS CMK ARN.
      When you use an AWS KMS CMK for server-side encryption in Amazon S3, you must choose a symmetric CMK. Amazon S3 only supports symmetric CMKs and not asymmetric CMKs. For more information, see Using Symmetric and Asymmetric Keys in the *AWS Key Management Service Developer Guide*.

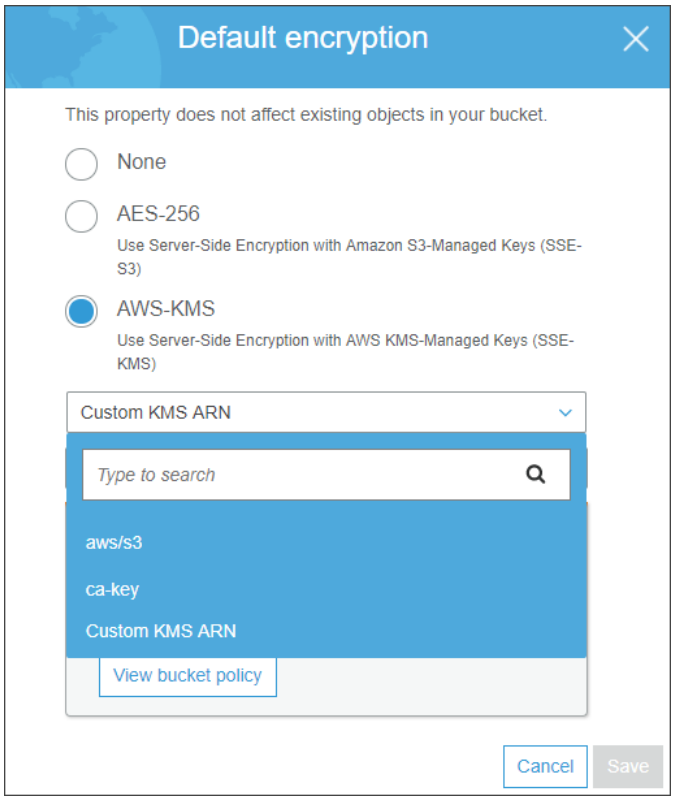

Important
If you use the AWS KMS option for your default encryption configuration, you are subject to the RPS (requests per second) limits of AWS KMS. For more information about AWS KMS limits and how to request a limit increase, see AWS KMS limits.

For more information about creating an AWS KMS CMK, see Creating Keys in the *AWS Key Management Service Developer Guide*. For more information about using AWS KMS with Amazon S3, see Protecting Data with Keys Stored in AWS KMS in the *Amazon Simple Storage Service Developer Guide*.

7. Choose **Save**.

More info

- Amazon S3 Default Encryption for S3 Buckets in the *Amazon Simple Storage Service Developer Guide*
- How do I add encryption to an S3 object? (p. 61)

How do I enable server access logging for an S3 bucket?

This topic describes how to enable server access logging for an Amazon S3 bucket using the AWS Management Console. For information about how to enable logging programmatically and details about how logs are delivered, see Server Access Logging in the *Amazon Simple Storage Service Developer Guide*.

By default, Amazon Simple Storage Service (Amazon S3) doesn't collect server access logs. When you enable logging, Amazon S3 delivers access logs for a source bucket to a target bucket that you choose. The target bucket must be in the same AWS Region as the source bucket and must not have a default retention period configuration.

Server access logging provides detailed records for the requests that are made to an S3 bucket. Server access logs are useful for many applications. For example, access log information can be useful in security and access audits. It can also help you learn about your customer base and understand your Amazon S3 bill.

An access log record contains details about the requests that are made to a bucket. This information can include the request type, the resources that are specified in the request, and the time and date that the request was processed. For more information, see Server Access Log Format in the *Amazon Simple Storage Service Developer Guide*.

> **Important**
> There is no extra charge for enabling server access logging on an Amazon S3 bucket. However, any log files that 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 server access logging for an S3 bucket

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

3. Choose **Properties**.

4. Choose **Server access logging**.

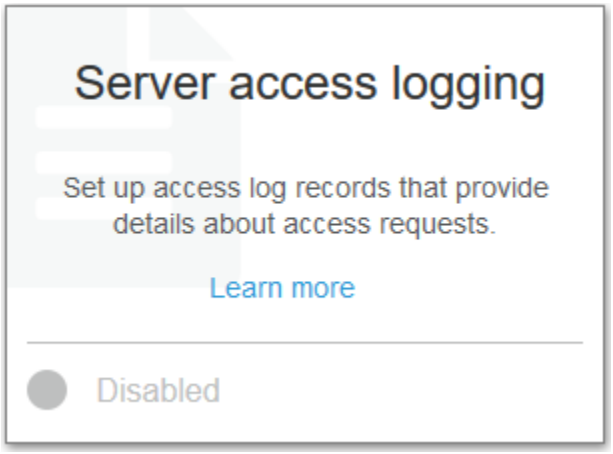

5. Choose **Enable Logging**. For **Target**, choose the name of the bucket that you want to receive the log record objects. The target bucket must be in the same Region as the source bucket and must not have a default retention period configuration.

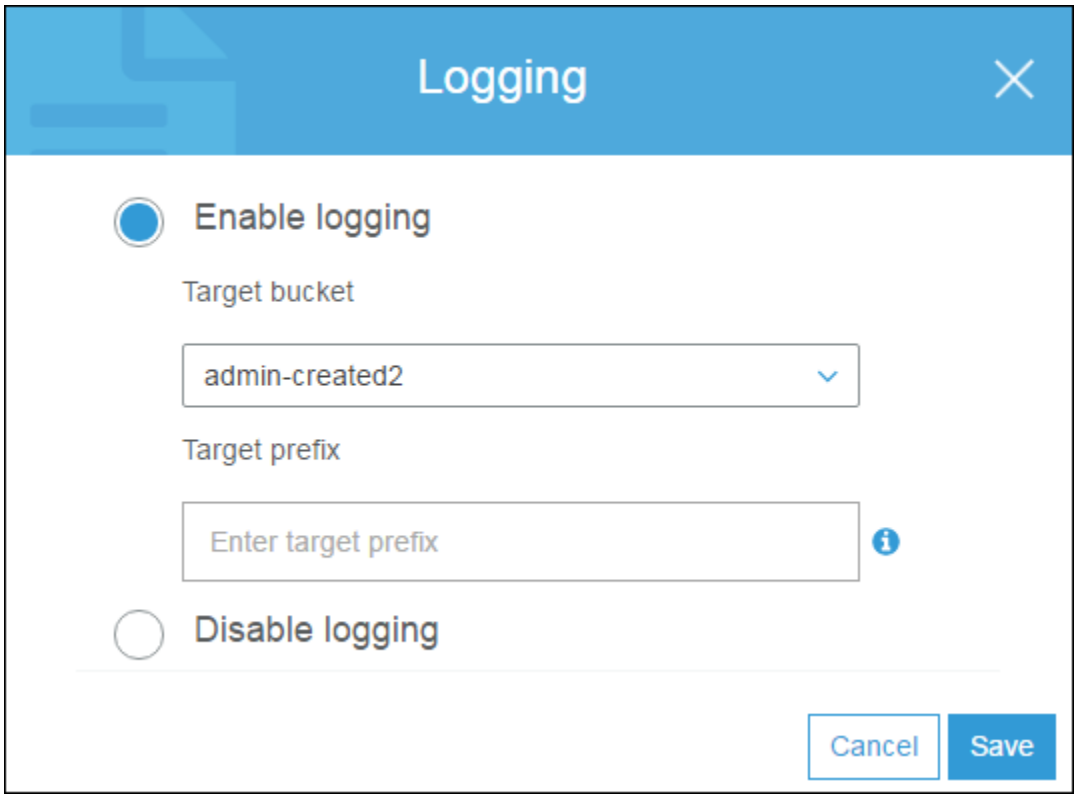

6. (Optional) For **Target prefix**, type a key name prefix for log objects, so that all of the log object names begin with the same string.
7. Choose **Save**.

   You can view the logs in the target bucket. If you specified a prefix, the prefix shows as a folder in the target bucket in the console. After you enable server access logging, it might take a few hours before the logs are delivered to the target bucket. For more information about how and when logs are delivered, see Server Access Logging in the *Amazon Simple Storage Service Developer Guide*.

More Info

How do I view the properties for an S3 Bucket? (p. 6)

How do I enable object-level logging for an S3 bucket with AWS CloudTrail data events?

This section describes how to enable an AWS CloudTrail trail to log data events for objects in an S3 bucket by using the Amazon S3 console. CloudTrail supports logging Amazon S3 object-level API operations such as `GetObject`, `DeleteObject`, and `PutObject`. These events are called data events. By default, CloudTrail trails don't log data events, but you can configure trails to log data events for S3 buckets that you specify, or to log data events for all the Amazon S3 buckets in your AWS account. For more information, see Logging Amazon S3 API Calls Using AWS CloudTrail. CloudTrail does not populate data events in the CloudTrail event history. Additionally, not all bucket-level actions are populated in the CloudTrail event history. For more information, see Using Amazon CloudWatch Logs filter patterns and Amazon Athena to query CloudTrail logs.

To configure a trail to log data events for an S3 bucket, you can use either the AWS CloudTrail console or the Amazon S3 console. If you are configuring a trail to log data events for all the Amazon S3 buckets in your AWS account, it's easier to use the CloudTrail console. For information about using the CloudTrail console to configure a trail to log S3 data events, see Data Events in the *AWS CloudTrail User Guide*.

> **Important**
> Additional charges apply for data events. For more information, see AWS CloudTrail Pricing.

The following procedure shows how to use the Amazon S3 console to enable a CloudTrail trail to log data events for an S3 bucket.

To enable CloudTrail data events logging for objects in an S3 bucket

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

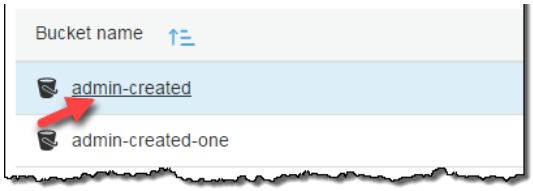

3. Choose **Properties**.

4. Choose **Object-level logging**.

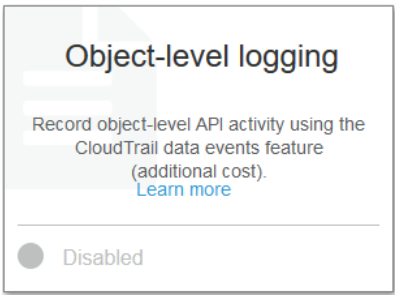

5. Choose an existing CloudTrail trail in the drop-down menu.

   The trail you select must be in the same AWS Region as your bucket, so the drop-down list contains only trails that are in the same Region as the bucket or trails that were created for all Regions.

   If you need to create a trail, choose the **CloudTrail console** link to go to the CloudTrail console. For information about how to create trails in the CloudTrail console, see Creating a Trail with the Console in the *AWS CloudTrail User Guide*.

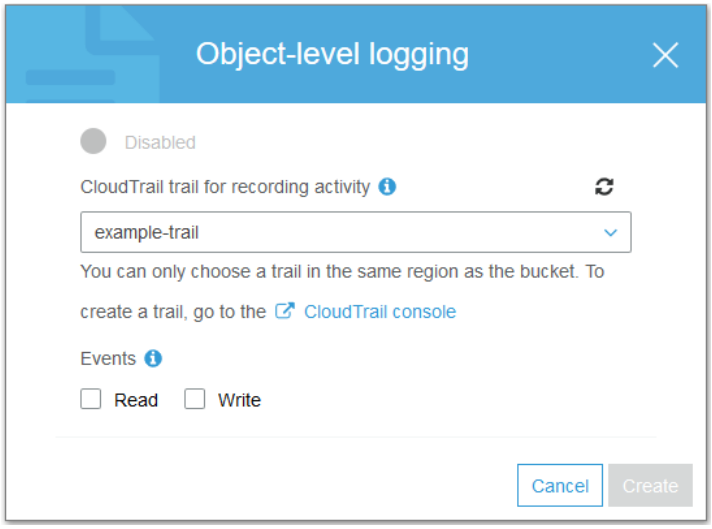

6. Under **Events**, choose one of the following:

   - **Read** to specify that you want CloudTrail to log Amazon S3 read APIs such as `GetObject`.
   - **Write** to log Amazon S3 write APIs such as `PutObject`.
   - **Read** and **Write** to log both read and write object APIs.

   For a list of supported data events that CloudTrail logs for Amazon S3 objects, see Amazon S3 Object-Level Actions Tracked by CloudTrail Logging in the *Amazon Simple Storage Service Developer Guide*.

7. Choose **Create** to enable object-level logging for the bucket.

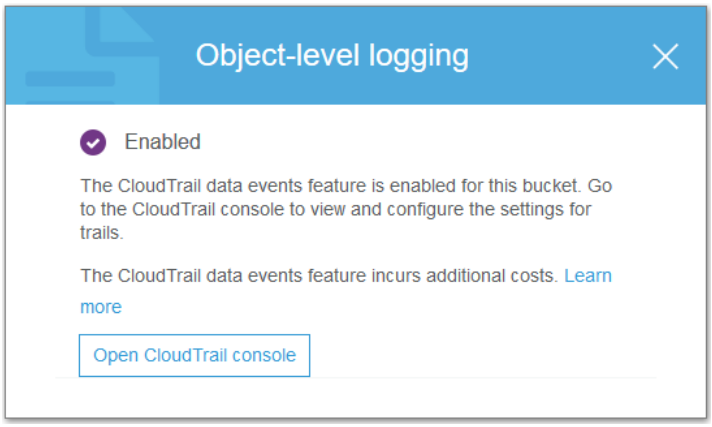

To disable object-level logging for the bucket, you must go to the CloudTrail console and remove the bucket name from the trail's **Data events**.

> **Note**
> If you use the CloudTrail console or the Amazon S3 console to configure a trail to log data events for an S3 bucket, the Amazon S3 console shows that object-level logging is enabled for the bucket.

For information about enabling object-level logging when you create an S3 bucket, see How do I create an S3 Bucket? (p. 3).

More info

- How do I view the properties for an S3 Bucket? (p. 6)
- Logging Amazon S3 API Calls By Using AWS CloudTrail in the *Amazon Simple Storage Service Developer Guide*
- Working with CloudTrail Log Files in the *AWS CloudTrail User Guide*

How do I configure an S3 bucket for static website hosting?

You can host a static website on Amazon S3. On a static website, individual webpages include static content. A static website might also contain client-side scripts. By contrast, a dynamic website relies on server-side processing, including server-side scripts such as PHP, JSP, or ASP.NET. Amazon S3 does not support server-side scripting.

You can use the following quick procedures to configure an S3 bucket for static website hosting in the Amazon S3 console. For more information and detailed walkthroughs, see Hosting a Static Website on Amazon S3 in the *Amazon Simple Storage Service Developer Guide*.

Topics

- Step 1: Configure an Amazon S3 bucket for static website hosting (p. 17)
- Step 2: Editing block public access settings (p. 18)
- Step 3: Adding a bucket policy (p. 20)
- Step 4: Test your website endpoint (p. 21)

Step 1: Configure an Amazon S3 bucket for static website hosting

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

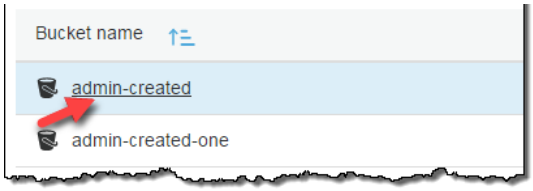

3. Choose **Properties**.

4. Choose **Static website hosting**.

5. Choose **Use this bucket to host a website**.

6. In **Index document**, enter the name of the index document, typically `index.html`.

   When you configure a bucket for website hosting, you must specify an index document. Amazon S3 returns this index document when requests are made to the root domain or any of the subfolders. For more information, see Configuring Index Document Support in the *Amazon Simple Storage Service Developer Guide*.
7. (Optional) If you want to provide your own custom error document for 4XX class errors, in **Error Document**, enter the custom error document filename.

   If you do not specify a custom error document and an error occurs, Amazon S3 returns a default HTML error document. For more information, see Custom Error Document Support in the *Amazon Simple Storage Service Developer Guide*.
8. (Optional) If you want to specify advanced redirection rules, in the **Edit redirection rules** box, enter XML to describe the rules.

For example, you can conditionally route requests according to specific object key names or prefixes in the request. For more information, see Advanced Conditional Redirects in the *Amazon Simple Storage Service Developer Guide*.

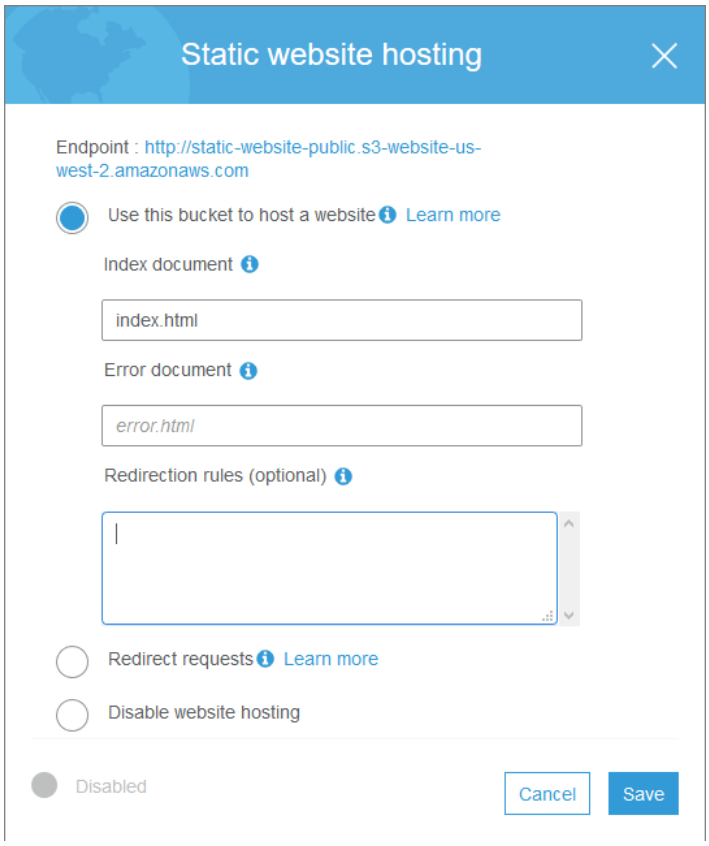

9. Choose **Save**.
10. Upload the index document to your bucket.

    For step-by-step instructions on uploading an object to an S3 bucket, see Uploading Files by Pointing and Clicking (p. 39).
11. Upload other files for your website, including optional custom error documents.

In the next section, you set the permissions required to access your bucket as a static website..

Step 2: Editing block public access settings

By default, Amazon S3 blocks public access to your account and buckets. If you want to use a bucket to host a static website, you can use these steps to edit your block public access settings.

> **Warning**
> Before you complete this step, review Using Amazon S3 Block Public Access to ensure that you understand and accept the risks involved with allowing public access. When you turn off block public access settings to make your bucket public, anyone on the internet can access your bucket. We recommend that you block all public access to your buckets.

1. Open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. Choose the name of the bucket that you have configured as a static website.
3. Choose **Permissions**.

4. Choose **Edit**.
5. Clear **Block *all* public access**, and choose **Save**.

   Warning
   Before you complete this step, review Using Amazon S3 Block Public Access to ensure you understand and accept the risks involved with allowing public access. When you turn off block public access settings to make your bucket public, anyone on the internet can access your bucket. We recommend that you block all public access to your buckets.

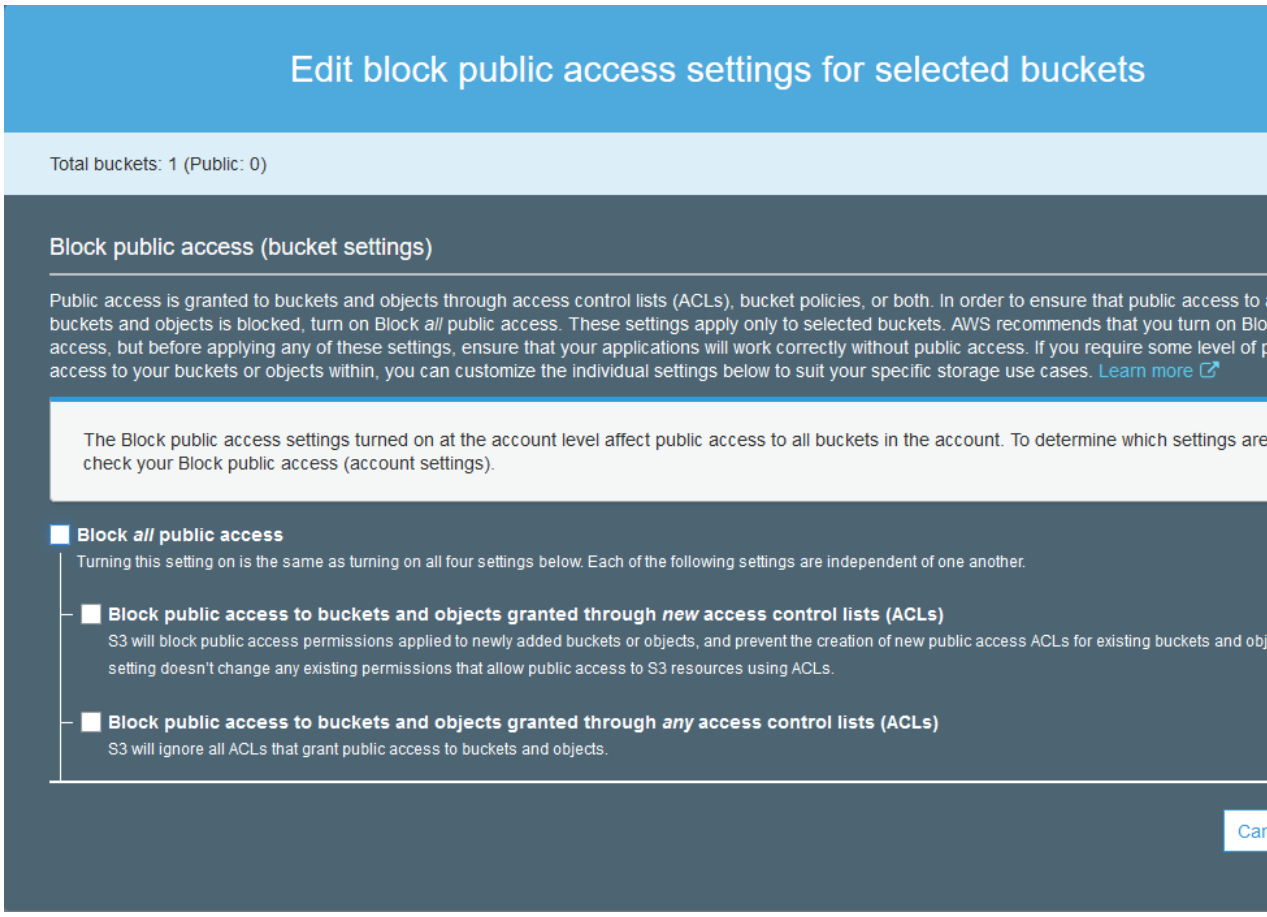

6. In the confirmation box, enter **`confirm`**, and then choose **Confirm**.

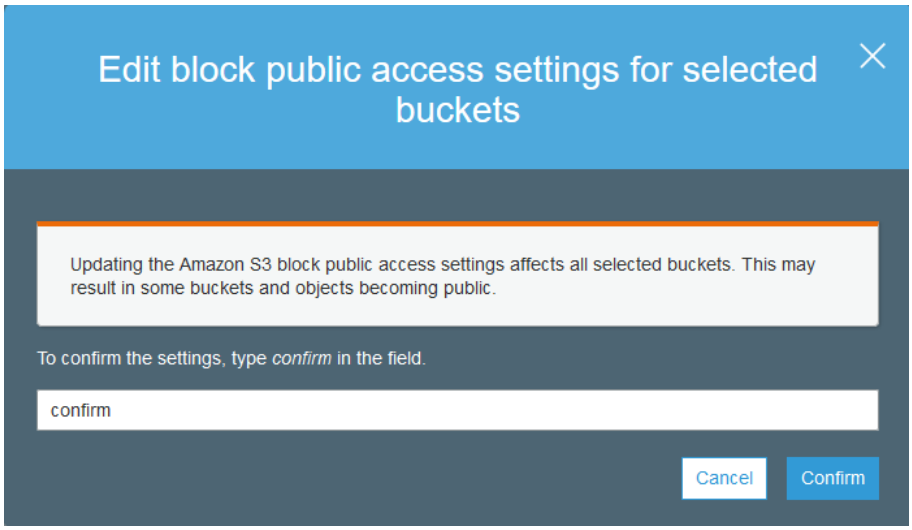

Under **S3 buckets**, the **Access** for your bucket updates to **Objects can be public**. You can now add a bucket policy to make the objects in the bucket publicly readable. If the **Access** still displays as **Bucket and objects not public**, you might have to edit the block public access settings for your account before adding a bucket policy.

Step 3: Adding a bucket policy

After you edit S3 Block Public Access settings, you can add a bucket policy to grant public read access to your bucket. When you grant public read access, anyone on the internet can access your bucket.

> **Important**
> The following policy is an example only and allows full access to the contents of your bucket. Before you proceed with this step, review How can I secure the files in my Amazon S3 bucket? to ensure that you understand the best practices for securing the files in your S3 bucket and risks involved in granting public access.

1. Under **Buckets**, choose the name of your bucket.
2. Choose **Permissions**.
3. Choose **Bucket Policy**.
4. To grant public read access for your website, copy the following bucket policy, and paste it in the **Bucket policy editor**.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "PublicReadGetObject",
            "Effect": "Allow",
            "Principal": "*",
            "Action": [
                "s3:GetObject"
            ],
            "Resource": [
                "arn:aws:s3:::example.com/*"
            ]
        }
    ]
}
```

5. Update the `Resource` to include your bucket name.

   In the preceding example bucket policy, *`example.com`* is the bucket name. To use this bucket policy with your own bucket, you must update this name to match your bucket name.
6. Choose **Save**.

   A warning appears indicating that the bucket has public access. In **Bucket Policy**, a **Public** label appears.

   If you see an error that says `Policy has invalid resource`, confirm that the bucket name in the bucket policy matches your bucket name. For information about adding a bucket policy, see How Do I Add an S3 Bucket Policy?

   If you get an **Error - Access denied** warning and the **Bucket policy editor** does not allow you to save the bucket policy, check your account-level and bucket-level block public access settings to confirm that you allow public access to the bucket.

Step 4: Test your website endpoint

Once you configure your bucket as a static website and set permissions, you can access your website through an Amazon S3 website endpoint. For more information, see Website Endpoints in the *Amazon Simple Storage Service Developer Guide*. For a complete list of Amazon S3 website endpoints, see Amazon S3 Website Endpoints in the *Amazon Web Services General Reference*.

How do I redirect requests to an S3 bucket hosted website to another host?

For more detailed information about configuring a redirect in Amazon S3, see Configuring a webpage redirect in the *Amazon Simple Storage Service Developer Guide*.

You can redirect all requests for a website endpoint for a bucket to another host. If you redirect all requests, any request made to the website endpoint is redirected to the specified host name.

For example, if your root domain is `example.com`, and you want to serve requests for both `http://example.com` and `http://www.example.com`, you can create two buckets named `example.com` and `www.example.com`. Then, maintain the content in the `example.com` bucket, and configure the other `www.example.com` bucket to redirect all requests to the `example.com` bucket. For more information, see Configuring a Static Website Using a Custom Domain Name.

To redirect requests for a bucket website endpoint

1. Open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. Choose the name of the bucket that you have configured as a static website (for example, `example.com`).
3. Choose **Properties**.
4. Choose **Static website hosting**.
5. Choose **Redirect requests**.
6. In the **Target bucket or domain** box, enter the bucket or domain that you want to redirect to.

   For example, if you are redirecting to a root domain address, you would enter **`example.com`**.
7. In the **Protocol** box, enter the protocol for the redirected requests (**`http`** or **`https`**).

If you do not specify a protocol, the protocol of the original request is used.

8. Choose **Save**.

Advanced settings for S3 bucket properties

This section describes how to configure advanced S3 bucket property settings for object replication, event notification, and transfer acceleration.

Topics

- How do I set up a destination to receive event notifications? (p. 22)
- How do I enable and configure event notifications for an S3 Bucket? (p. 23)
- How do I enable transfer acceleration for an S3 bucket? (p. 28)

How do I set up a destination to receive 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

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

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

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

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

An Amazon SQS queue

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

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

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

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

A Lambda function

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

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

Warning
If your notification ends up writing to the bucket that triggers the notification, this could cause an execution loop. For example, if the bucket triggers a Lambda function each time an object is uploaded, and the function uploads an object to the bucket, then the function indirectly triggers itself. To avoid this, use two buckets, or configure the trigger to only apply to a prefix used for incoming objects.
For more information and an example of using Amazon S3 notifications with AWS Lambda, see Using AWS Lambda with Amazon S3 in the *AWS Lambda Developer Guide*.

How do I enable and configure event notifications for an S3 Bucket?

You can enable certain Amazon S3 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 information about using event notifications with the AWS SDKs and the Amazon S3 REST APIs, see Configuring Amazon S3 Event Notifications in the *Amazon Simple Storage Service Developer Guide*.

Topics

- Event notification types (p. 24)
- Event notification destinations (p. 25)
- Enabling and configuring event notifications (p. 25)
- More info (p. 28)

Event notification types

When you configure event notifications for a bucket, you must specify the type of events for which you want to receive notifications. For a complete list of event types, see Supported Event Types section in the *Amazon Simple Storage Service Developer Guide*.

In the Amazon S3 console, you have the following options for configuring event notifications. You can choose a single option or multiple options.

- **Object creation**
  - **All object create events** – Receive a notification when an object is created in your bucket by any of the following object creation actions: **Put**, **Post**, **Copy**, and **Multipart upload completed**.
  - **Put**, **Post**, **Copy**, and **Multipart upload completed** – Receive a notification for one of these specific object creation actions.
- **Object deletion**
  - **All object delete events** – Receive a notification any time an object in your bucket is deleted.
  - **Delete marker created** – Receive a notification 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.
- **Object restoration from S3 Glacier or S3 Glacier Deep Archive storage class**
  - **Restore initiated** – Receive a notification for *Initiation* of object restoration.
  - **Restore completed** – Receive a notification for *Completion* of object restoration.
- **Reduced Redundancy Storage (RRS) object lost events**
  - **Object in RSS Lost** – Receive a notification that an object of the RRS storage class has been lost
- **Objects eligible for replication using Amazon S3 Replication Time Control**
  - **Replication time missed threshold** – Receive a notification that an object exceeded the 15-minute threshold for replication.
  - **Replication time completed after threshold** – Receive a notification that an object replicated after the 15-minute threshold.
  - **Replication time not tracked** – Receive a notification that an object that was eligible for replication is no longer being tracked by replication metrics.
  - **Replication time failed** – Receive a notification that an object failed to replicate.

    Note
    When you delete the last object from a folder, Amazon S3 can generate an object creation event. When there are multiple objects with the same prefix with a trailing slash (/) as part of their names, those objects are shown as being part of a folder in the Amazon S3 console. The name of the folder is formed from the characters preceding the trailing slash (/).
    When you delete all the objects listed under that folder, no actual object is available to represent the empty folder. Under such circumstances, the Amazon S3 console creates a zero-byte object to represent that folder. If you enabled event notification for the creation of objects, the zero-byte object creation action that is taken by the console triggers an object creation event.

The Amazon S3 console displays a folder under the following circumstances:

- When a zero-byte object has a trailing slash (/) in its name. In this case, there is an actual Amazon S3 object of 0 bytes that represents a folder.
- If the object has a slash (/) within its name. In this case, there isn't an actual object representing the folder.

Event notification destinations

When you configure event notifications for a bucket, you also choose the notification destination. Event notification messages can be sent to the following destinations:

- **Amazon Simple Notification Service (Amazon SNS) topic** – Coordinates and manages the delivery or sending of messages to subscribing endpoints or clients. For information about the Amazon SNS topic format, see SNS FAQ.
- **Amazon Simple Queue Service (Amazon SQS) queue** – Offers reliable and scalable hosted queues for storing messages as they travel between computers. For information about Amazon SQS, see What is Amazon Simple Queue Service? in the *Amazon Simple Queue Service Developer Guide*.
- **AWS Lambda** – Invoke a Lambda function and provide the event message as an argument. When you create a Lambda function, you package up and upload your custom code to AWS Lambda. AWS Lambda uses the AWS infrastructure to run the code on your behalf. For information about using Lambda with Amazon S3, see Using AWS Lambda with Amazon S3 in the *AWS Lambda Developer Guide*.

For more information about granting the Amazon S3 service principal the permissions required to publish event notifications to a destination, see Granting Permissions to Publish Event Notification Messages to a Destination in the *Amazon S3 Developer Guide*.

> **Warning**
> If your notification ends up writing to the bucket that triggers the notification, this could cause an execution loop. For example, if the bucket triggers a Lambda function each time an object is uploaded, and the function uploads an object to the bucket, then the function indirectly triggers itself. To avoid this, use two buckets, or configure the trigger to only apply to a prefix used for incoming objects.
> For more information and an example of using Amazon S3 notifications with AWS Lambda, see Using AWS Lambda with Amazon S3 in the *AWS Lambda Developer Guide*.

Enabling and configuring event notifications

Before you can enable event notifications for your bucket, you must set up one of the destination types. For more information, see How do I set up a destination to receive event notifications? (p. 22)

To enable and configure event notifications for an S3 bucket

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

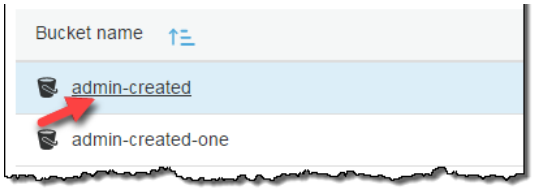

3. Choose **Properties**.

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

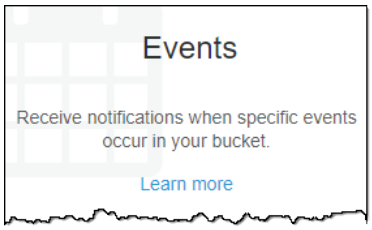

5. Choose **Add notification**.

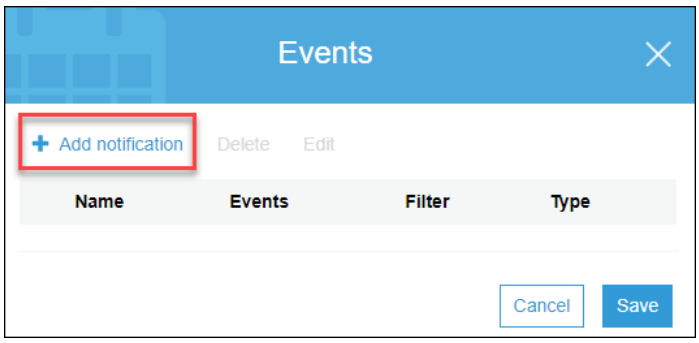

6. In **Name**, enter a descriptive name for your event notification.

   If you don't enter a name, a GUID is generated and used for the name.

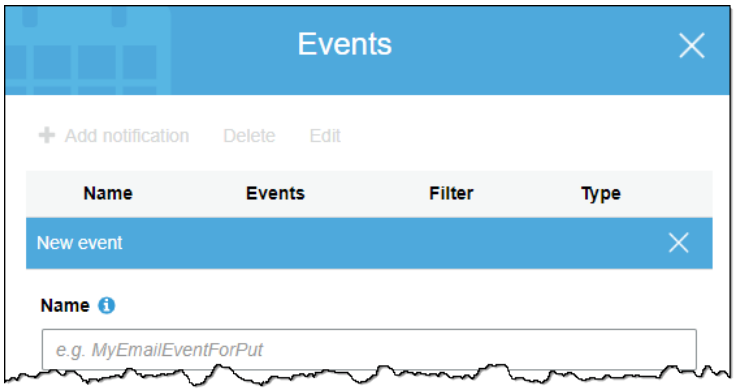

7. Under **Events**, select one or more events.

   For a listing of the event types, see Event notification types (p. 24).

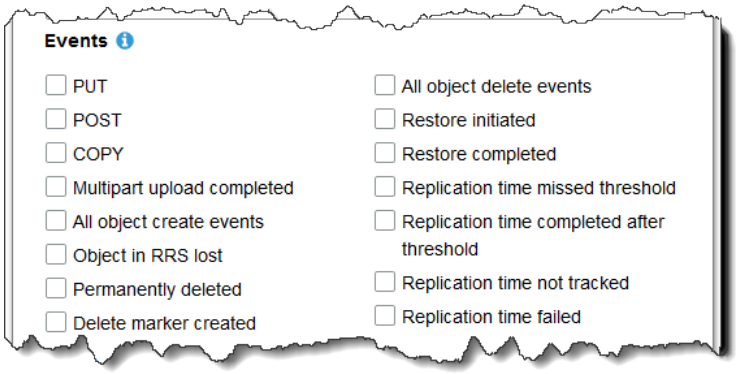

8. To filter event notifications by a prefix or suffix, enter a **Prefix** or a **Suffix**.

   For example, you can set up a prefix filter so that you receive notifications only when files are added to a specific folder (for example, `images/`). For more information, see Configuring Notifications with Object Key Name Filtering.

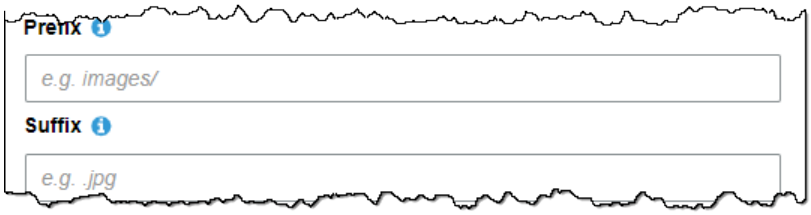

9. Choose the event notification destination: **SNS Topic**, **SQS Queue**, or **Lambda Function**.

   For a description of the destinations, see Event notification destinations (p. 25).

   When you choose your **Send to** destination, a box appears for you to enter your specific SNS, SQS, or Lambda function destination. In the example image below, the **Send To** location is **SNS Topic**, and you can see a **SNS** box for the SNS topic name.

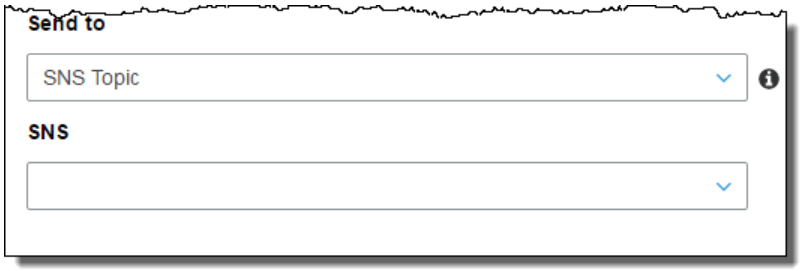

10. In the box that appears, choose or enter the destination SNS, SQS, or Lambda function.

    You can choose or enter the SNS, SQS, or Lambda function name, or you can choose to the destination Amazon Resource Name (ARN). The example screenshot below shows the **Add SNS topic ARN** option.

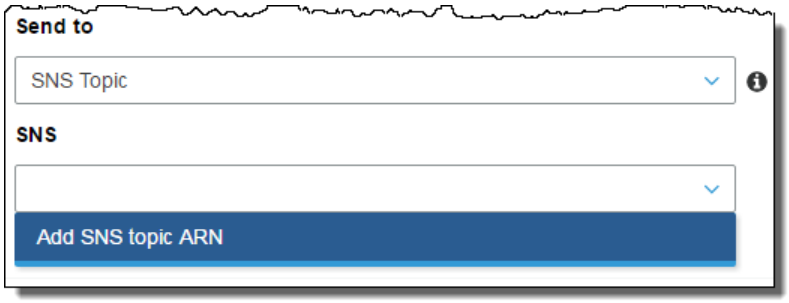

11. If you chose **Add ARN**, enter the SNS topic, SQS queue, or Lambda function ARN.

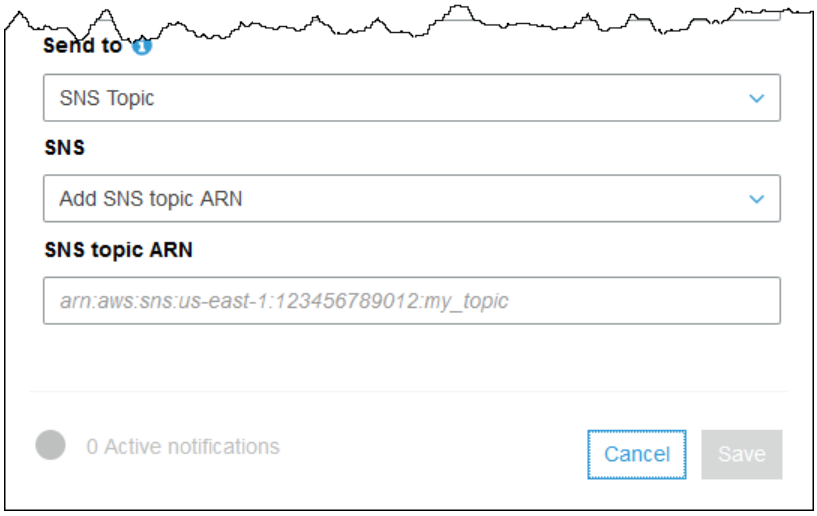

12. Choose **Save**.

    Amazon S3 sends a test message to the event notification destination.

More info

- How do I restore an S3 object that has been archived? (p. 47)

How do I enable transfer acceleration for an S3 bucket?

Amazon Simple Storage Service (Amazon S3) transfer acceleration enables fast, easy, and secure file transfers between your client and an S3 bucket over long distances. This topic describes how to enable Amazon S3 transfer acceleration for a bucket. For more information, see Amazon S3 Transfer Acceleration in the *Amazon Simple Storage Service Developer Guide*.

To enable transfer acceleration for an S3 bucket

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the **Buckets** list, choose the name of the bucket that you want to enable transfer acceleration for.
3. Choose **Properties**.

4. Under **Advanced settings**, choose **Transfer acceleration**.

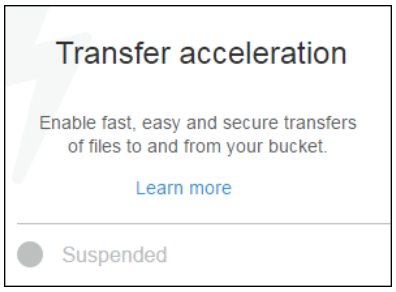

5. (Optional) If you want to compare accelerated and non-accelerated upload speeds, to open the Amazon S3 Transfer Acceleration Speed Comparison tool, choose the **Want to compare your data transfer speed by region?**.

   The Speed Comparison tool uses multipart upload to transfer a file from your browser to various AWS Regions with and without Amazon S3 transfer acceleration. You can compare the upload speed for direct uploads and transfer accelerated uploads by Region.
6. Choose **Enabled**, and choose **Save**.

   Endpoint displays the transfer acceleration endpoint for your bucket. You use this endpoint to access accelerated data transfers to and from your bucket. If you suspend transfer acceleration, the accelerate endpoint no longer works.

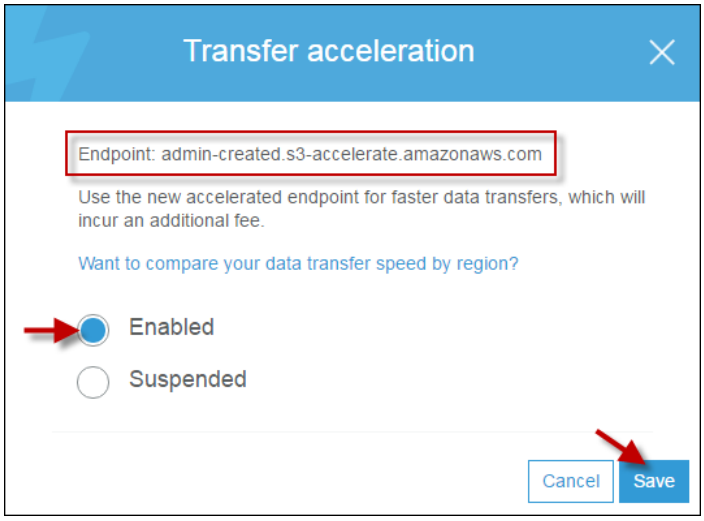

Introduction to Amazon S3 access points

You can use Amazon S3 access points to manage access to your S3 objects. Amazon S3 access points are named network endpoints that are attached to buckets that you can use to perform S3 object operations, such as uploading and retrieving objects. A bucket can have up to 1,000 access points attached, and each access point enforces distinct permissions and network controls to give you fine-grained control over access to your S3 objects.

For more information about Amazon S3 Access Points, see Managing Data Access with Amazon S3 Access Points in the *Amazon Simple Storage Service Developer Guide*.

The following topics explain how to use the S3 Management Console to create, manage, and use Amazon S3 Access Points.

Topics

- Creating an Amazon S3 access point (p. 30)
- Managing and using Amazon S3 access points (p. 31)

Creating an Amazon S3 access point

This section explains how to create an Amazon S3 access point using the AWS Management Console. For information about creating access points using the AWS CLI, AWS SDKs, and the Amazon S3 REST APIs, see Managing Data Access with Amazon S3 Access Points in the *Amazon Simple Storage Service Developer Guide*.

An access point is associated with exactly one Amazon S3 bucket. Before you begin, make sure that you have created a bucket that you want to use with this access point. For more information about creating buckets, see Creating and configuring an S3 bucket (p. 3).

To create an access point

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the **S3 buckets** section, select the bucket that you want to attach this access point to.
3. On the bucket detail page, choose the **Access points** tab.
4. Choose **Create access point**.
5. Enter your desired name for the access point in the **Access point name** field.
6. Choose a **Network access type**. If you choose **Virtual private cloud (VPC)**, enter the **VPC ID** that you want to use with the access point.

   For more information about network access type for access points, see Creating Access Points Restricted to a Virtual Private Cloud in the *Amazon Simple Storage Service Developer Guide*.
7. Select the block public access settings that you want to apply to the access point. All block public access settings are enabled by default for new access points, and we recommend that you leave all settings enabled unless you know you have a specific need to disable any of them. Amazon S3

currently doesn't support changing an access point's block public access settings after the access point has been created.

For more information about using Amazon S3 Block Public Access with access points, see Managing Public Access to Access Points in the *Amazon Simple Storage Service Developer Guide*.

8. (Optional) Specify the access point policy. The console automatically displays the Amazon Resource Name (ARN) for the access point, which you can use in the policy.
9. Choose **Create access point**.

Managing and using Amazon S3 access points

This section explains how to manage and use your Amazon S3 access points using the AWS Management Console. Each access point is associated with a single Amazon S3 bucket. Before you begin, navigate to the list of your access points for a bucket as described in the following procedure.

To find a list of access points for 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 **S3 buckets** section, select the bucket whose access points you want to manage.
3. On the bucket detail page, choose the **Access points** tab.

From the **Access points** tab, you can view an access point's configuration details, edit an access point's policy, use an access point to access your bucket, or delete an access point. The following procedures explain how to perform each of these tasks.

To view access point configuration details

1. Navigate to the **Access points** tab for your bucket.
2. Locate the access point whose configuration you want to view. You can browse for the access point in the access point list, or you can search for a specific access point using the **Search by name** field.
3. Choose the name of the access point whose configuration you want to view.

   Note
   To view an access point's configuration, choose (click on) the name of the access point, not the option button next to the access point name.

To edit an access point policy

1. Navigate to the **Access points** tab for your bucket.
2. Select the option button next to the name of the access point whose policy you want to edit.
3. Choose **Edit access point policy**.
4. Enter the policy in the text field. The console automatically displays the Amazon Resource Name (ARN) for the access point, which you can use in the policy. You can also choose **Policy generator** to use the AWS Policy Generator to help construct the policy.
5. Choose **Save**.

To use an access point to access your bucket

1. Navigate to the **Access points** tab for your bucket.
2. Select the option button next to the name of the access point you want to use.

3. Choose **Use this access point**.
4. The console displays a label above the name of your bucket that shows the access point that you're currently using. While you're using the access point, you can only perform the object operations that are allowed by the access point permissions.

   Note
   The console view always shows all objects in the bucket. Using an access point as described in this procedure restricts the operations you can perform on those objects, but not whether you can see that they exist in the bucket.
5. To exit the access point view of your bucket, choose **Exit access point**.

Note
The S3 Management Console doesn't support using virtual private cloud (VPC) access points to access bucket resources. To access bucket resources from a VPC access point, use the AWS CLI, AWS SDKs, or Amazon S3 REST APIs.

To delete an access point

1. Navigate to the **Access points** tab for your bucket.
2. Select the option button next to the name of the access point that you want to delete.
3. Choose **Delete**.
4. Confirm that you want to delete your access point by entering its name in the text field that appears, and choose **Confirm**.

Uploading, downloading, and managing objects

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

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

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

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

Objects can be any file type: images, backups, data, movies, etc. You can have an unlimited number of objects in a bucket. The maximum size of file you can upload by using the Amazon S3 console is 160 GB. To upload a file larger than 160 GB, use the AWS CLI, AWS SDK, or Amazon S3 REST API. For more information, see Uploading Objects in the *Amazon Simple Storage Service Developer Guide*.

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

> **Note**
> If you rename an object, or change any of the properties; **Storage Class**, **Enryption**, **Metadata**, a new object is created to replace the old one. If S3 Versioning is enabled, a new version of the object is created, and the existing object becomes an older version. The role that changes the property also becomes the owner of the new object or (object version).

Topics

- How do I upload files and folders to an S3 bucket? (p. 34)
- Copying an object (p. 42)
- Moving an object (p. 42)
- How do I download an object from an S3 bucket? (p. 43)
- How do I delete objects from an S3 bucket? (p. 46)
- How do I undelete a deleted S3 object? (p. 47)
- How do I restore an S3 object that has been archived? (p. 47)
- How do I lock an Amazon S3 object? (p. 53)
- How do I see an overview of an object? (p. 55)
- How do I see the versions of an S3 object? (p. 58)
- How do I view the properties of an object? (p. 59)
- How do I add encryption to an S3 object? (p. 61)
- How do I add metadata to an S3 object? (p. 63)

- How do I add tags to an S3 object? (p. 70)
- How do I use folders in an S3 bucket? (p. 72)

How do I upload files and folders to an S3 bucket?

This topic explains how to use the AWS Management Console to upload one or more files or entire folders to an Amazon S3 bucket. Before you can upload files and folders to an Amazon S3 bucket, you need write permissions for the bucket. For more information about access permissions, see Setting bucket and object access permissions (p. 112). For information about uploading files programmatically, see Uploading Objects in the *Amazon Simple Storage Service Developer Guide*.

When you upload a file to Amazon S3, it is stored as an S3 object. Objects consist of the file data and metadata that describes the object. You can have an unlimited number of objects in a bucket.

You can upload any file type—images, backups, data, movies, etc.—into an S3 bucket. The maximum size of a file that you can upload by using the Amazon S3 console is 160 GB. To upload a file larger than 160 GB, use the AWS CLI, AWS SDK, or Amazon S3 REST API. For more information, see Uploading Objects in the *Amazon Simple Storage Service Developer Guide*.

> **Note**
> To upload *folders*, you must drag and drop them. To upload *files*, you can drag and drop or point and click. Drag and drop functionality is supported only for Chrome and Firefox browsers.

For information about which Chrome and Firefox browser versions are supported, see What browsers are supported for use with the AWS Management Console?.

When you upload a folder, Amazon S3 uploads all of the files and subfolders from the specified folder to your bucket. It then assigns an object key name that is a combination of the uploaded file name and the folder name. For example, if you upload a folder called `/images` that contains two files, `sample1.jpg` and `sample2.jpg`, Amazon S3 uploads the files and then assigns the corresponding key names, `images/sample1.jpg` and `images/sample2.jpg`. The key names include the folder name as a prefix. The Amazon S3 console displays only the part of the key name that follows the last "/". For example, within an images folder the `images/sample1.jpg` and `images/sample2.jpg` objects are displayed as `sample1.jpg` and a `sample2.jpg`.

If you upload individual files and you have a folder open in the Amazon S3 console, when Amazon S3 uploads the files, it includes the name of the open folder as the prefix of the key names. For example, if you have a folder named `backup` open in the Amazon S3 console and you upload a file named `sample1.jpg`, the key name is `backup/sample1.jpg`. However, the object is displayed in the console as `sample1.jpg` in the `backup` folder.

If you upload individual files and you do not have a folder open in the Amazon S3 console, when Amazon S3 uploads the files, it assigns only the file name as the key name. For example, if you upload a file named `sample1.jpg`, the key name is `sample1.jpg`. For more information on key names, see Object Key and Metadata in the *Amazon Simple Storage Service Developer Guide*.

If you upload an object with a key name that already exists in a versioning-enabled bucket, Amazon S3 creates another version of the object instead of replacing the existing object. For more information about versioning, see How do I enable or suspend versioning for an S3 bucket? (p. 7).

Topics

- Uploading Files and Folders by Using Drag and Drop (p. 35)
- Uploading Files by Pointing and Clicking (p. 39)
- More Info (p. 41)

Uploading Files and Folders by Using Drag and Drop

If you are using the Chrome or Firefox browsers, you can choose the folders and files to upload, and then drag and drop them into the destination bucket. Dragging and dropping is the *only* way that you can upload folders.

To upload folders and files to an S3 bucket by using drag and drop

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

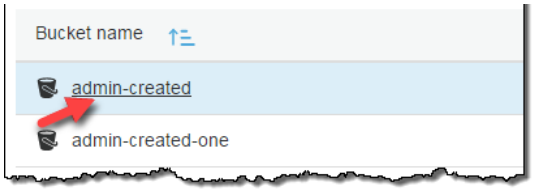

3. In a window other than the console window, select the files and folders that you want to upload. Then drag and drop your selections into the console window that lists the objects in the destination bucket.

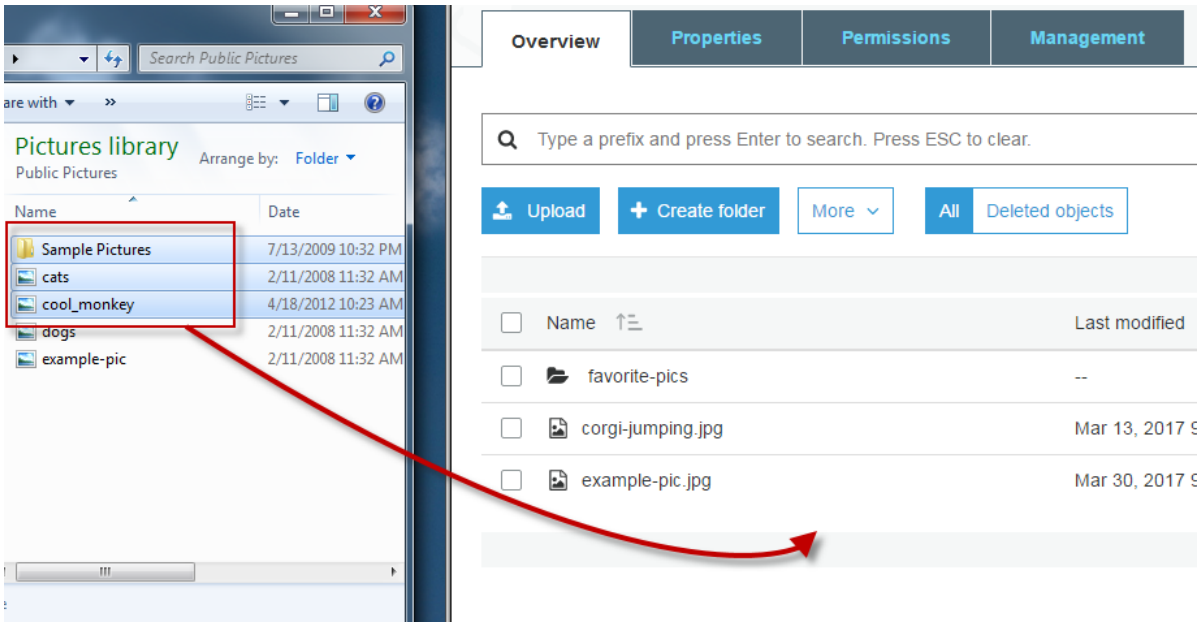

The files you chose are listed in the **Upload** dialog box.

4. In the **Upload** dialog box, do one of the following:

   a. Drag and drop more files and folders to the console window that displays the **Upload** dialog box. To add more files, you can also choose **Add more files**. This option works *only* for files, not folders.

   b. To immediately upload the listed files and folders without granting or removing permissions for specific users or setting public permissions for all of the files that you're uploading, choose **Upload**. For information about object access permissions, see How do I set permissions on an object? (p. 118).

   c. To set permissions or properties for the files that you are uploading, choose **Next**.

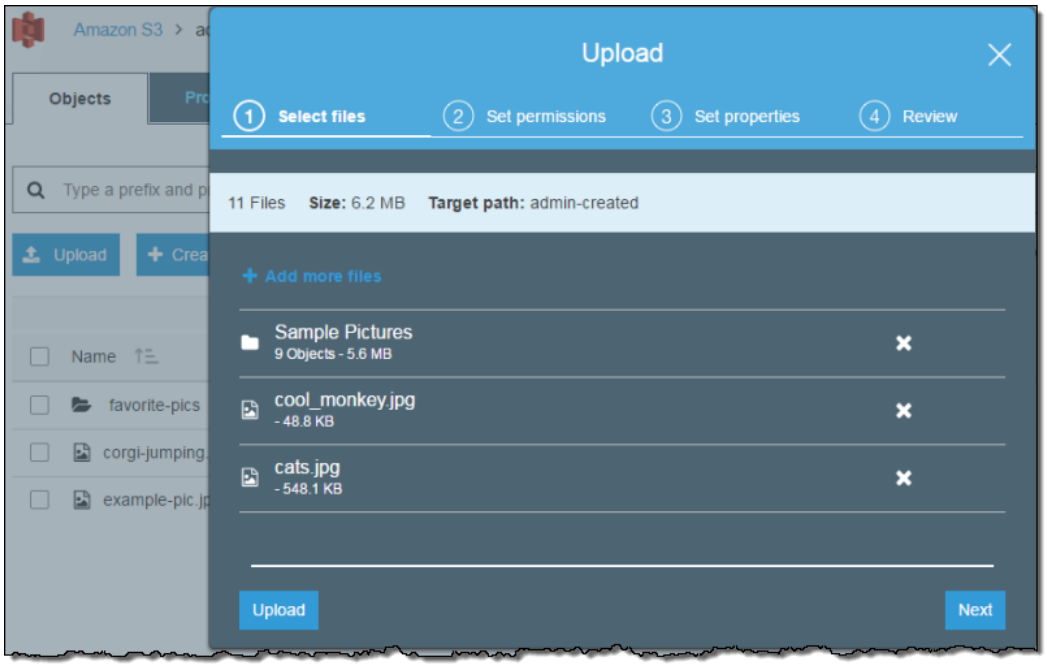

5. On the **Set Permissions** page, under **Manage users** you can change the permissions for the AWS account owner. The *owner* refers to the AWS account root user, and not an AWS Identity and Access Management (IAM) user. For more information about the root user, see The AWS Account Root User.

   Choose **Add account** to grant access to another AWS account. For more information about granting permissions to another AWS account, see How do I set ACL bucket permissions? (p. 121).

   Under **Manage public permissions** you can grant read access to your objects to the general public (everyone in the world), for all of the files that you're uploading. Granting public read access is applicable to a small subset of use cases such as when buckets are used for websites. We recommend that you do not change the default setting of **Do not grant public read access to this object(s)**. You can always make changes to object permissions after you upload the object. For information about object access permissions, see How do I set permissions on an object? (p. 118).

   When you're done configuring permissions, choose **Next**.

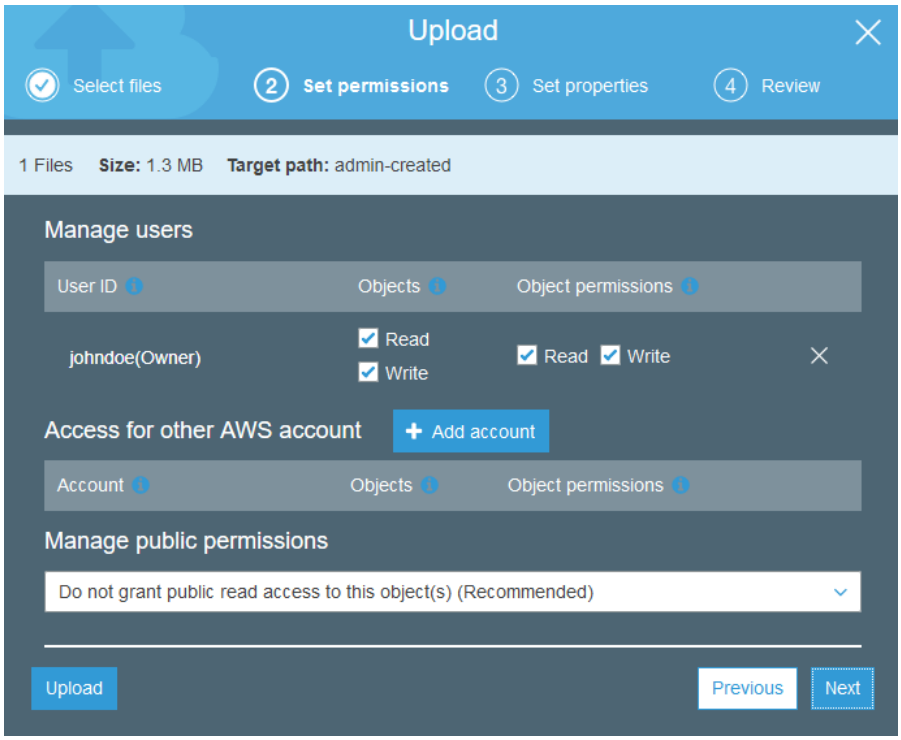

6. On the **Set Properties** page, choose the storage class and encryption method to use for the files that you are uploading. You can also add or modify metadata.

   a. Choose a storage class for the files you're uploading. For more information about storage classes, see Storage Classes in the *Amazon Simple Storage Service Developer Guide*.

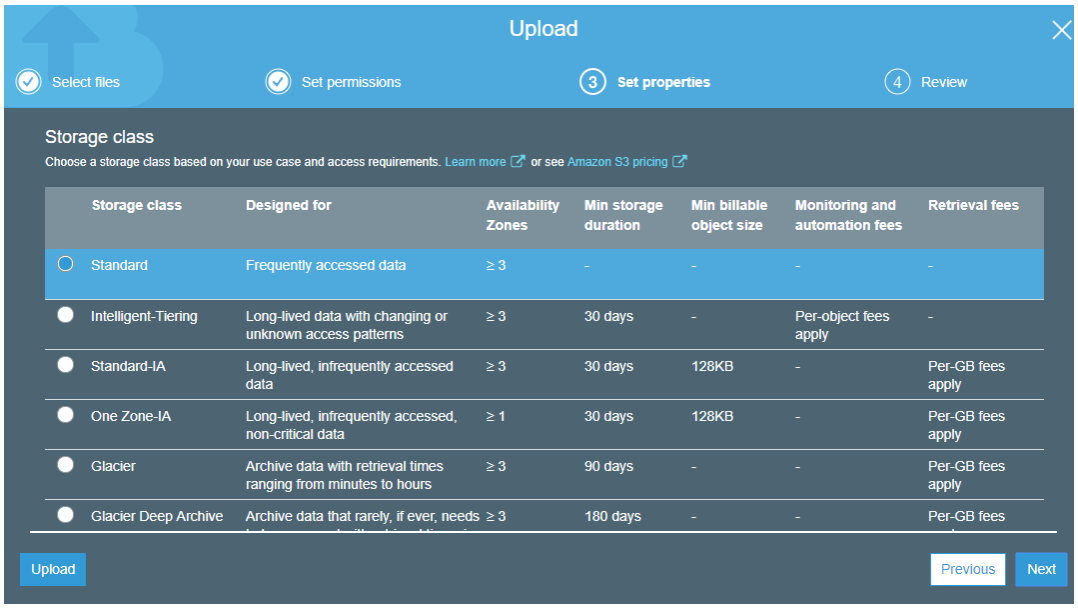

   b. Choose the type of encryption for the files that you're uploading. If you don't want to encrypt them, choose **None**.

      i. To encrypt the uploaded files using keys that are managed by Amazon S3, choose **Amazon S3 master-key**. For more information, see Protecting Data with Amazon S3-Managed Encryption Keys Classes in the *Amazon Simple Storage Service Developer Guide*.

ii. To encrypt the uploaded files using the AWS Key Management Service (AWS KMS), choose **AWS KMS master-key**. Then choose a customer master key (CMK) from the list of AWS KMS CMKs.

> **Note**
> To encrypt objects in a bucket, you can use only CMKs that are available in the same AWS Region as the bucket.

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

For more information about creating an AWS KMS CMK, see Creating Keys in the *AWS Key Management Service Developer Guide*. For more information about protecting data with AWS KMS, see Protecting Data Using Keys Stored in AWS KMS (SSE-KMS) in the *Amazon Simple Storage Service Developer Guide*.

c. Metadata for Amazon S3 objects is represented by a name-value (key-value) pair. There are two kinds of metadata: system-defined metadata and user-defined metadata.

If you want to add Amazon S3 system-defined metadata to all of the objects you are uploading, for **Header**, select a header. You can select common HTTP headers, such as **Content-Type** and **Content-Disposition**. Type a value for the header, and then choose **Save**. For a list of system-defined metadata and information about whether you can add the value, see System-Defined Metadata in the *Amazon Simple Storage Service Developer Guide*.

d. Any metadata starting with prefix `x-amz-meta-` is treated as user-defined metadata. User-defined metadata is stored with the object, and is returned when you download the object.

To add user-defined metadata to all of the objects that you are uploading, type `x-amz-meta-` plus a custom metadata name in the **Header** field. Type a value for the header, and then choose **Save**. Both the keys and their values must conform to US-ASCII standards. User-defined metadata can be as large as 2 KB. For more information about user-defined metadata, see User-Defined Metadata in the *Amazon Simple Storage Service Developer Guide*.

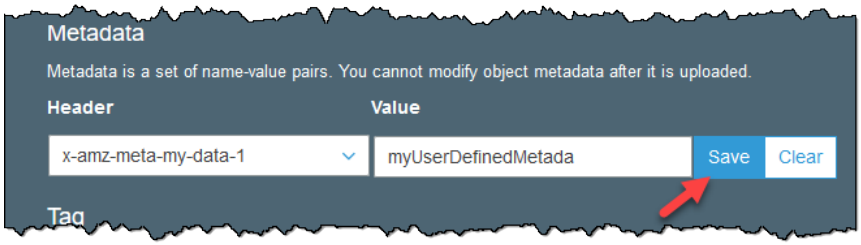

e. Object tagging gives you a way to categorize storage. Each tag is a key-value pair. Key and tag values are case sensitive. You can have up to 10 tags per object.

To add tags to all of the objects that you are uploading, type a tag name in the **Key** field. Type a value for the tag, and then choose **Save**. A tag key can be up to 128 Unicode characters in length and tag values can be up to 255 Unicode characters in length. For more information about object tags, see Object Tagging in the *Amazon Simple Storage Service Developer Guide*.

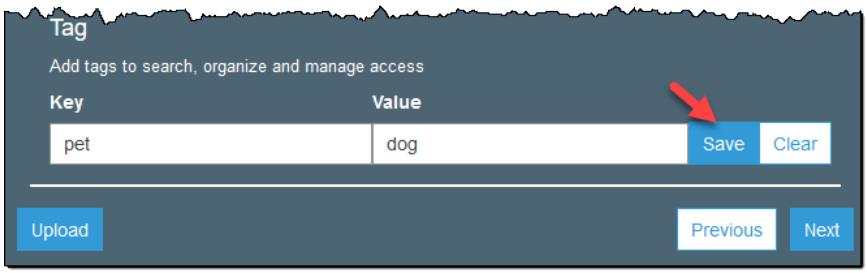

7. Choose **Next**.
8. On the **Upload** review page, verify that your settings are correct, and then choose **Upload**. To make changes, choose **Previous**.
9. To see the progress of the upload, choose **In progress** at the bottom of the browser window.

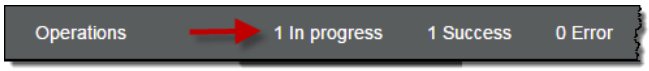

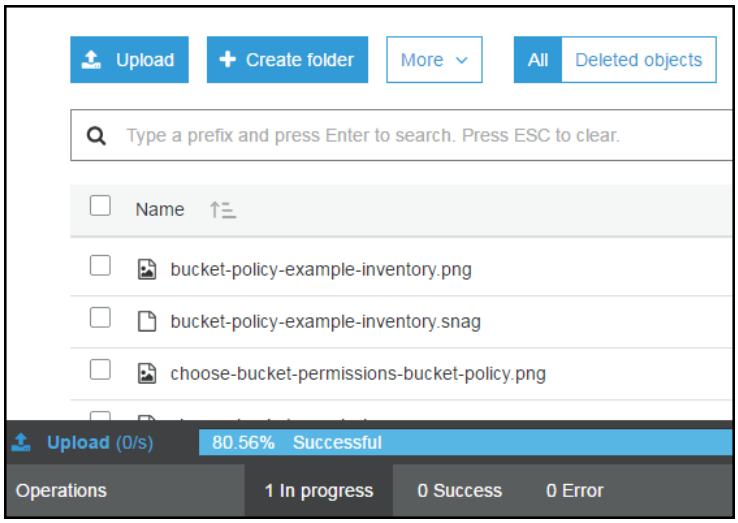

To see a history of your uploads and other operations, choose **Success**.

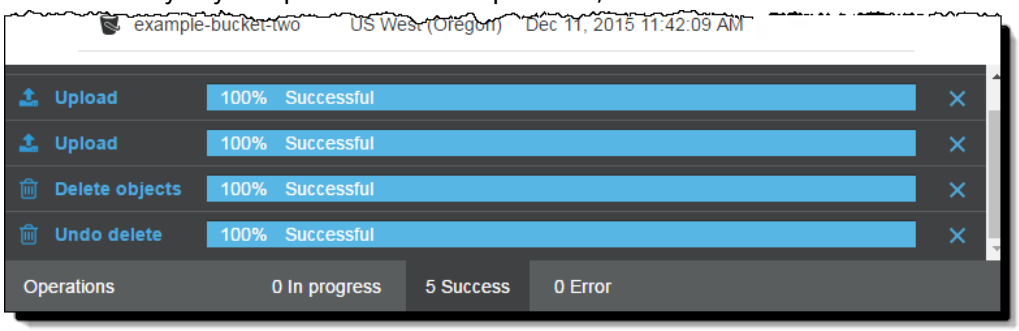

Uploading Files by Pointing and Clicking

This procedure explains how to upload files into an S3 bucket by choosing **Upload**.

To upload files to an S3 bucket by pointing and clicking

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

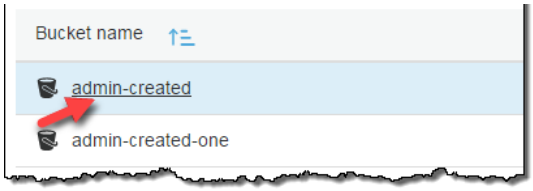

3. Choose **Upload**.

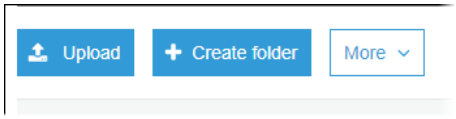

4. In the **Upload** dialog box, choose **Add files**.

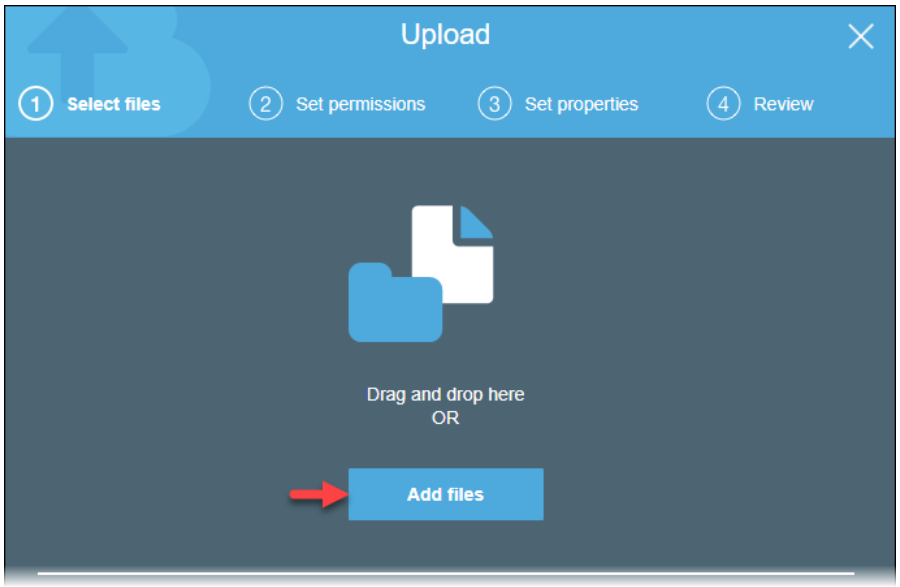

5. Choose one or more files to upload, and then choose **Open.**

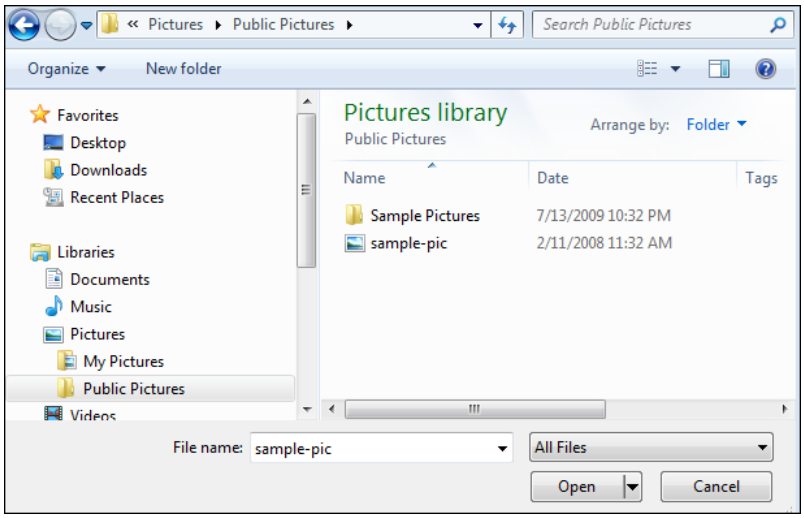

6. After you see the files that you chose listed in the **Upload** dialog box, do one of the following:

   a. To add more files, choose **Add more files**.

   b. To immediately upload the listed files, choose **Upload**.

   c. To set permissions or properties for the files that you are uploading, choose **Next**.

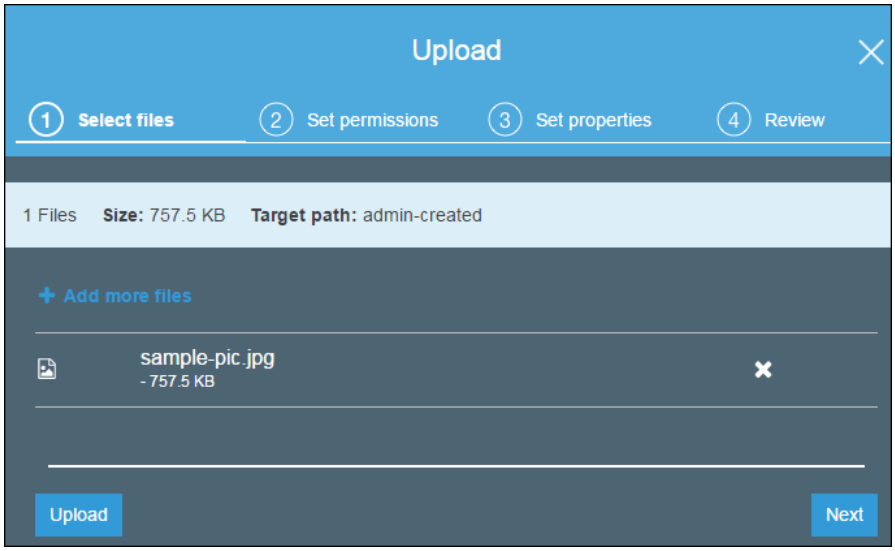

7. To set permissions and properties, start with **Step 5** of Uploading Files and Folders by Using Drag and Drop (p. 35).

More Info

- How do I set permissions on an object? (p. 118).
- How do I download an object from an S3 bucket? (p. 43)

Copying an object

In the Amazon S3 console, you can copy an object to a bucket or a folder. For more information, see Copying objects in the *Amazon Simple Storage Service Developer Guide*.

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. In the **Buckets** list, choose your bucket name.
3. On the **Overview** tab, select the check box beside the object that you want to copy.
4. Choose **Actions**, and choose **Copy**.
5. Choose the copy destination:
   - If you want to copy your object to a bucket, select the bucket.
   - If you want to copy your object to a folder in a bucket, choose the bucket name, and then choose the folder.
6. Under **Review**, confirm the copy details, and choose **Copy**.

   Amazon S3 copies your object to the destination.

   Note
   This action copies all specified objects. When copying folders, wait for the copy action to finish before making additional changes in the folders.

Moving an object

In the Amazon S3 console, you can move an object to a bucket or a folder.

To move 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 **Buckets** list, choose your bucket name.
3. Select the check box beside the object that you want to move.
4. Choose **Actions**, and choose **Move**.
5. Choose the move destination:
   - If you want to move your object to a bucket, select the bucket.
   - If you want to move your object to a folder in a bucket, choose the bucket name, and then choose the folder.
6. Under **Review**, confirm the move details, and choose **Move**.

   Amazon S3 moves your object to the destination.

   Note
   This action moves all specified objects. When moving folders, wait for the move action to finish before making additional changes in the folders.

How do I download an object from an S3 bucket?

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

Data transfer fees apply when you download objects. For information about Amazon S3 features, and pricing, see Amazon S3.

Important

- If an object key name consists of a single period (.), or two periods (..), you can't download the object using the Amazon S3 console. To download an object with a key name of "." or "..", you must use the AWS CLI, AWS SDKs, or REST API. For more information about naming objects, see Object Key Naming Guidelines in the *Amazon Simple Storage Service Developer Guide*.
- You can download a single object per request using the Amazon S3 console. To download multiple objects, use the AWS CLI, AWS SDKs, or REST API.

To download an object from an S3 bucket

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

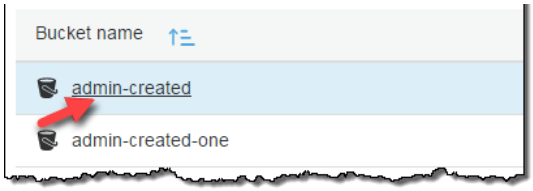

3. You can download an object from an S3 bucket in any of the following ways:
   - In the **Name** list, select the check box next to the object you want to download, and then choose **Download** on the object description page that appears.

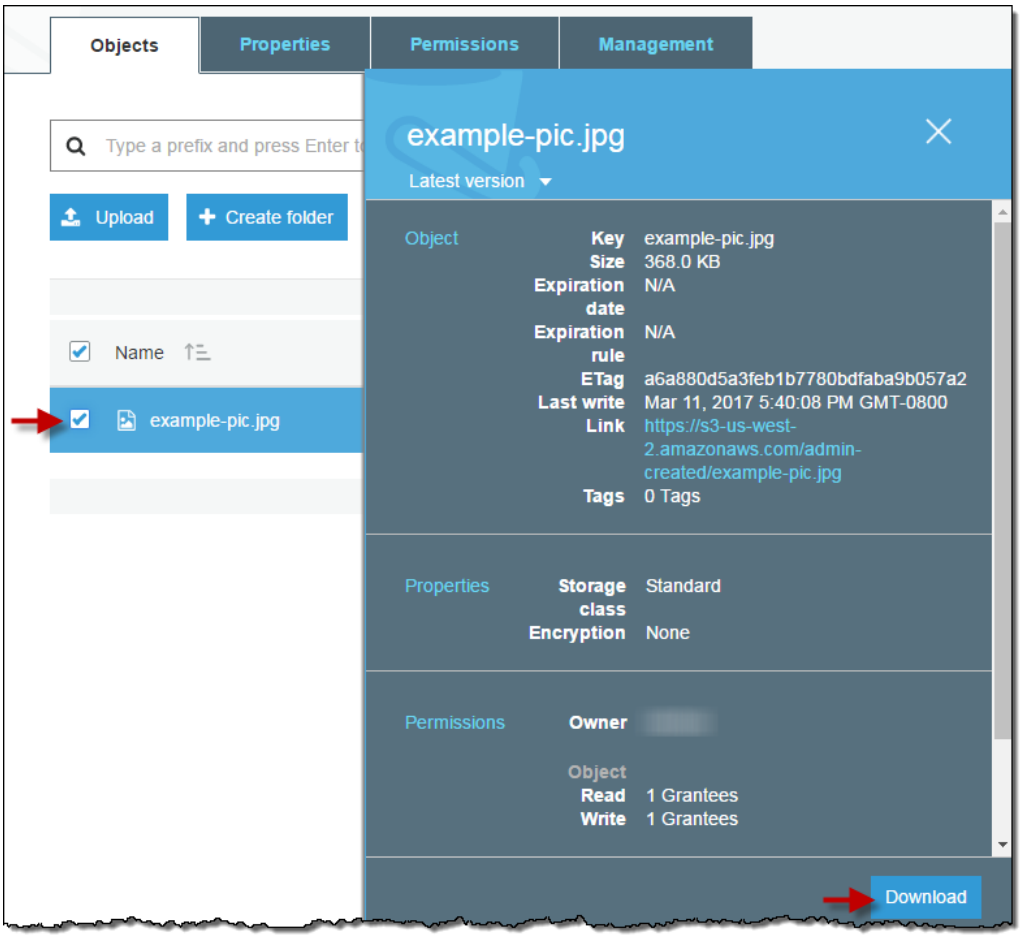

- Choose the name of the object that you want to download.

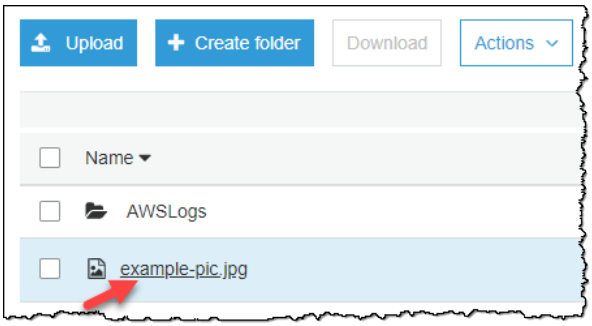

On the **Overview** page, choose **Download**.

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

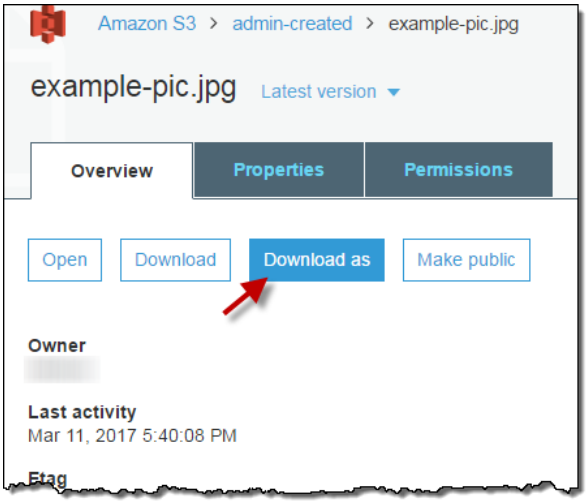

- Choose the name of the object that you want to download. Choose **Latest version** and then choose the download icon.

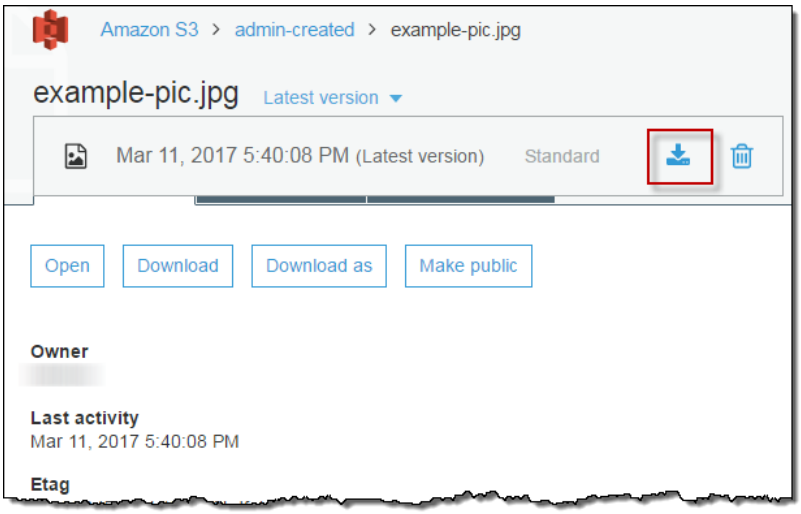

Related topics

- How do I upload files and folders to an S3 bucket? (p. 34)

How do I delete objects from an S3 bucket?

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

For information about Amazon S3 features and pricing, see Amazon S3.

To delete objects from an S3 bucket

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the **Bucket name** list, choose the name of the bucket that you want to delete an object from.
3. You can delete objects from an S3 bucket in any of the following ways:

   - In the **Name** list, select the check box next to the objects and folders that you want to delete, choose **Actions**, and then choose **Delete** from the drop-down menu.

     In the **Delete objects** dialog box, verify that the name(s) of the object(s) and/or folder(s) you selected for deletion are listed and then choose **Delete**.
   - Or, choose the name of the object that you want to delete, choose **Latest version**, and then choose the **trash can** icon.

   Warning
   This action deletes all specified objects. When deleting folders, wait for the delete action to finish before adding new objects to the folder. Otherwise, new objects might be deleted as well.

More info

- How do I undelete a deleted S3 object? (p. 47)
- How do I create a lifecycle policy for an S3 Bucket? (p. 79)

How do I undelete a deleted S3 object?

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

To be able to undelete a deleted object, you must have had versioning enabled on the bucket that contains the object before the object was deleted. For information about enabling versioning, see How do I enable or suspend versioning for an S3 bucket? (p. 7).

When you delete an object in a versioning-enabled bucket, all versions remain in the bucket and Amazon S3 creates a delete marker for the object. To undelete the object, you must delete this delete marker. For more information about versioning and delete markers, see Object Versioning in the *Amazon Simple Storage Service Developer Guide*.

To recover deleted objects from an S3 bucket

The following steps describe how to recover deleted objects that are not folders from your S3 bucket including objects that are within those folders.

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the **Bucket name** list, choose the name of the bucket that you want.
3. To see a list of the **versions** of the objects in the bucket, select **Show**. You'll be able to see the delete markers for deleted objects.
4. To undelete an object, you must delete the delete marker. Select the check box next to the **delete marker** of the object to recover, and then choose **delete** from the **Actions** menu.
5. Then, choose **Hide** and you'll see the undeleted object listed.

   Note
   You can't use the Amazon S3 console to undelete folders. You must use the AWS CLI or SDK. For examples, see How can I retrieve an Amazon S3 object that was deleted in a versioning-enabled bucket?

More info

- How do I see the versions of an S3 object? (p. 58)
- How do I enable or suspend versioning for an S3 bucket? (p. 7)
- Using Versioning in the *Amazon Simple Storage Service Developer Guide*

How do I restore an S3 object that has been archived?

This section explains how to use the Amazon S3 console to restore an object that has been archived to the S3 Glacier or S3 Glacier Deep Archive storage classes. Objects stored in the S3 Glacier or S3

Glacier Deep Archive are not immediately accessible. To access an object in this class, you must restore a temporary copy of it to its S3 bucket for the duration (number of days) that you specify. For information about the S3 Glacier or S3 Glacier Deep Archive storage classes, see Storage Classes in the *Amazon Simple Storage Service Developer Guide*.

When you restore an archive, you pay for both the archive and the restored copy. Because there is a storage cost for the copy, restore objects only for the duration you need them. If you want a permanent copy of the object, create a copy of it in your S3 bucket. For information about Amazon S3 features and pricing, see Amazon S3.

After restoring an object, you can download it from the **Overview** page. For more information, see How do I see an overview of an object? (p. 55).

Topics

- Archive Retrieval Options (p. 48)
- Restoring an Archived S3 Object (p. 48)
- Upgrade an In-Progress Restore (p. 51)
- Checking Archive Restore Status and Expiration Date (p. 52)

Archive Retrieval Options

The following are the available retrieval options when restoring an archived object:

- **`Expedited`** - Expedited retrievals allow you to quickly access your data stored in the S3 Glacier storage class when occasional urgent requests for a subset of archives are required. For all but the largest archived objects (250 MB+), data accessed using Expedited retrievals is typically made available within 1–5 minutes. Provisioned capacity ensures that retrieval capacity for Expedited retrievals is available when you need it. For more information, see Provisioned Capacity. Expedited retrievals and provisioned capacity are not available for objects stored in the S3 Glacier Deep Archive storage class.
- **`Standard`** - Standard retrievals allow you to access any of your archived objects within several hours. This is the default option for the S3 Glacier and S3 Glacier Deep Archive retrieval requests that do not specify the retrieval option. Standard retrievals typically finish within 3–5 hours for objects stored in the S3 Glacier storage class. They typically finish within 12 hours for objects stored in the S3 Glacier Deep Archive storage class.
- **`Bulk`** - Bulk retrievals are the lowest-cost retrieval option in Amazon S3 Glacier, enabling you to retrieve large amounts, even petabytes, of data inexpensively. Bulk retrievals typically finish within 5–12 hours for objects stored in the S3 Glacier storage class. They typically finish within 48 hours for objects stored in the S3 Glacier Deep Archive storage class.

For more information about retrieval options, see Restoring Archived Objects in the *Amazon Simple Storage Service Developer Guide*.

Restoring an Archived S3 Object

This topic explains how to use the Amazon S3 console to restore an object that has been archived to the S3 Glacier or S3 Glacier Deep Archive storage classes. (The console uses the names Glacier and Glacier Deep Archive for these storage classes.)

To restore archived S3 objects

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

2. In the **Bucket name** list, choose the name of the bucket that contains the objects that you want to restore.

3. In the **Name** list, select the object or objects that you want to restore, choose **Actions**, and then choose **Restore**.

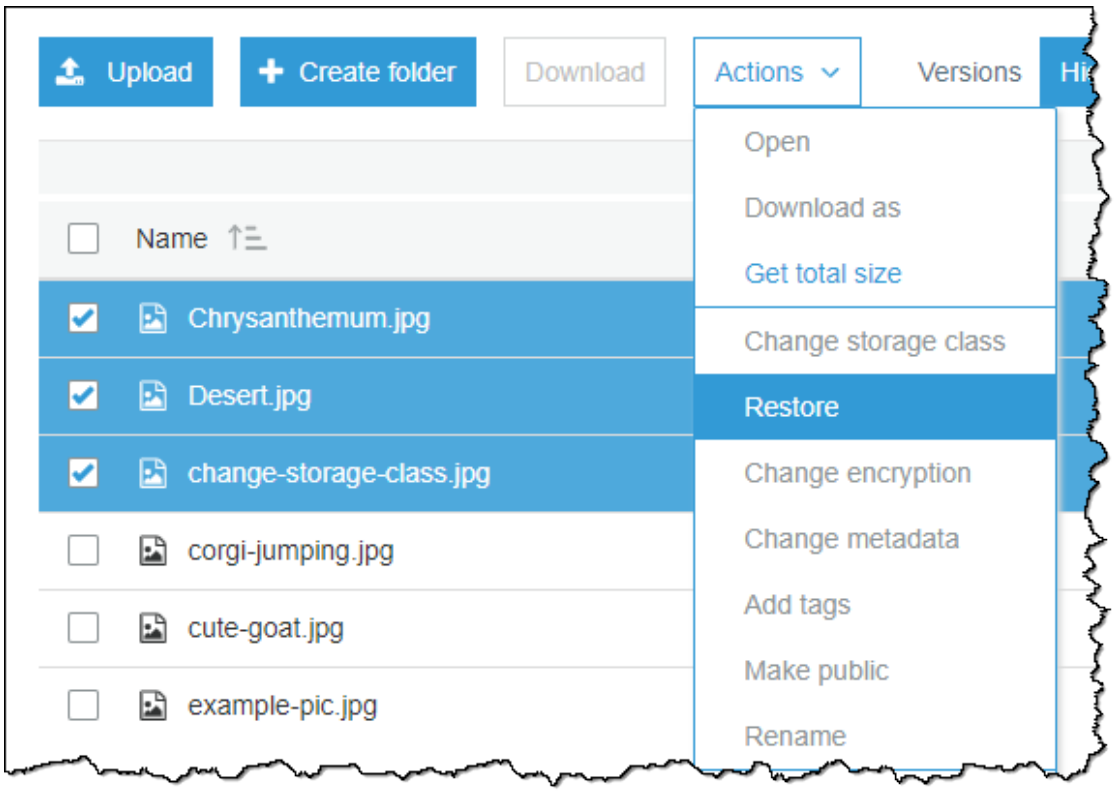

4. In the **Initiate restore** dialog box, type the number of days that you want your archived data to be accessible.
5. Choose one of the following retrieval options from the **Retrieval options** menu.
   - Choose **Bulk retrieval** or **Standard retrieval**, and then choose **Restore**.
   - Choose **Expedited retrieval** (available only for the Glacier storage class).

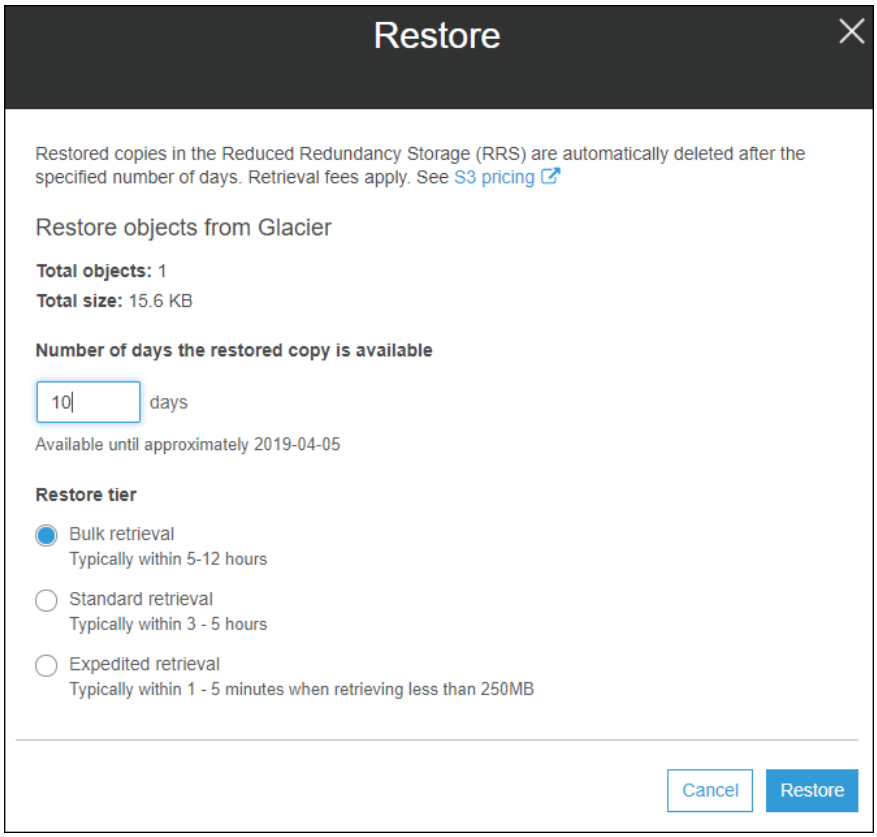

6. Provisioned capacity is only available for the Glacier storage class. If you have provisioned capacity, choose **Restore** to start a provisioned retrieval. If you have provisioned capacity, all of your expedited retrievals are served by your provisioned capacity. For more information about provisioned capacity, see Provisioned Capacity.

   - If you don't have provisioned capacity and you don't want to buy it, choose **Restore**.
   - If you don't have provisioned capacity, but you want to buy it, choose **Add capacity unit**, and then choose **Buy**. When you get the **Purchase succeeded** message, choose **Restore** to start provisioned retrieval.

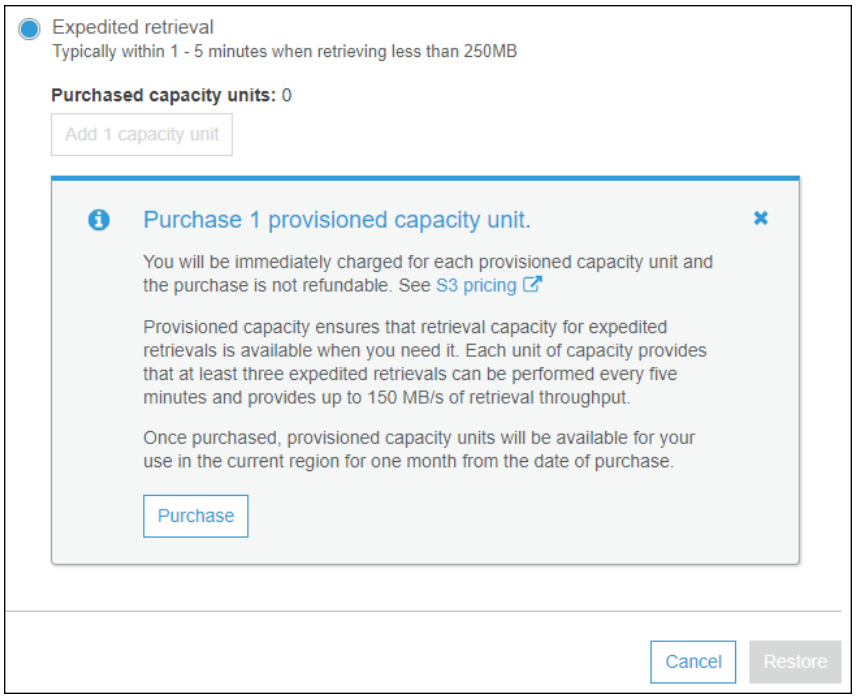

Upgrade an In-Progress Restore

You can upgrade the speed of your restoration while it is in progress.

To upgrade an in-progress restore to a faster tier

1. In the **Name** list, select one or more of the objects that you are restoring, choose **Actions**, and then choose **Restore from Glacier**. For information about checking the restoration status of an object, see Checking Archive Restore Status and Expiration Date (p. 52).

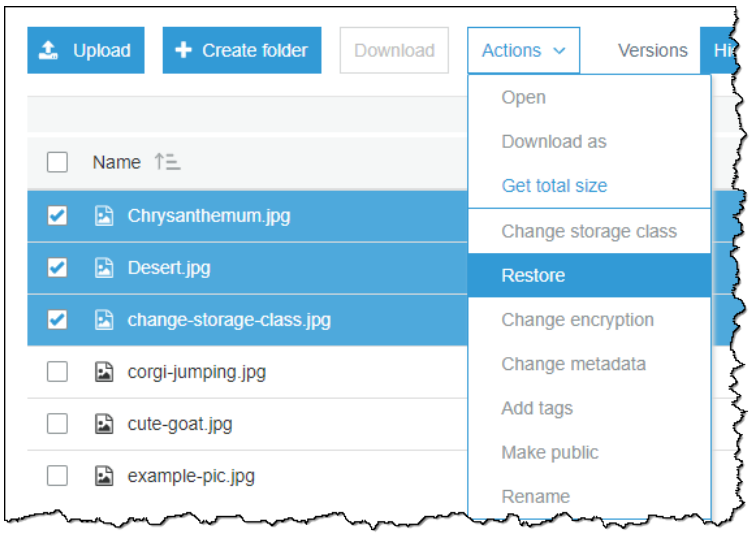

2. Choose the tier that you want to upgrade to and then choose **Restore**. For more information about upgrading to a faster restore tier, see see Restoring Archived Objects in the *Amazon Simple Storage Service Developer Guide*.

Checking Archive Restore Status and Expiration Date

To check the progress of the restoration, see the object overview panel. For information about the overview panel, see How do I see an overview of an object? (p. 55).

The **Overview** section shows that restoration is **In progress**.

When the temporary copy of the object is available, the object's **Overview** section shows the **Restoration expiry date**. This is when Amazon S3 will remove the restored copy of your archive.

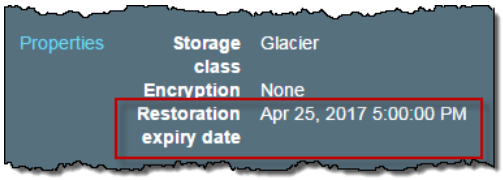

Restored objects are stored only for the number of days that you specify. If you want a permanent copy of the object, create a copy of it in your Amazon S3 bucket.

Amazon S3 calculates the expiry date by adding the number of days that you specify to the time you request to restore the object, and then rounding to the next day at midnight UTC. This calculation applies to the initial restoration of the object and to any extensions to availability that you request. For example, if an object was restored on 10/15/2012 10:30 AM UTC and the number of days that you specified is 3, then the object is available until 10/19/2012 00:00 UTC. If, on 10/16/2012 11:00 AM UTC you change the number of days that you want it to be accessible to 1, then Amazon S3 makes the restored object available until 10/18/2012 00:00 UTC.

After restoring an object, you can download it from the **Overview** page. For more information, see How do I see an overview of an object? (p. 55).

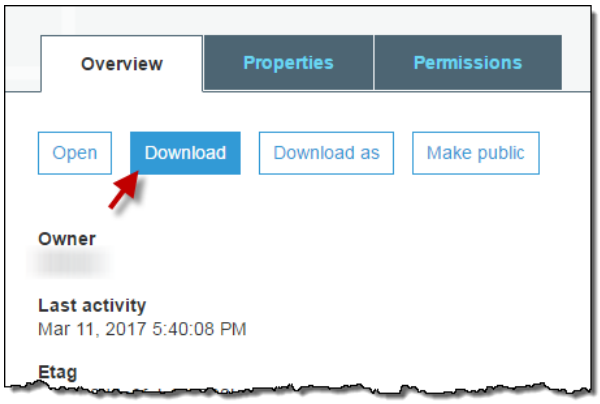

More Info

- Restoring Archived Objects in the Amazon S3 Developer Guide.
- restore-object in the AWS CLI Command Reference.
- Identity and Access Management in Amazon S3 Glacier in the S3 Glacier Developer Guide.
- How do I create a lifecycle policy for an S3 Bucket? (p. 79)
- How do I undelete a deleted S3 object? (p. 47)

How do I lock an Amazon S3 object?

With S3 Object Lock, you can store objects in Amazon S3 using a *write-once-read-many* (WORM) model. You can use S3 Object Lock to prevent an object from being deleted or overwritten for a fixed amount of time or indefinitely. For information about object locking using the AWS CLI, AWS SDKs, and the Amazon S3 REST APIs, see Locking Objects Using Object Lock in the *Amazon Simple Storage Service Developer Guide*.

Before you lock any objects, you have to enable a bucket to use S3 Object Lock. You enable Object Lock when you create a bucket. After you enable Object Lock on a bucket, you can lock objects in that bucket. When you create a bucket with Object Lock enabled, you can't disable Object Lock or suspend versioning for that bucket.

For information about creating a bucket with S3 Object Lock enabled, see How do I create an S3 Bucket? (p. 3).

To lock an Amazon S3 object

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

3. In the **Name** list, choose the name of the object that you want to lock.

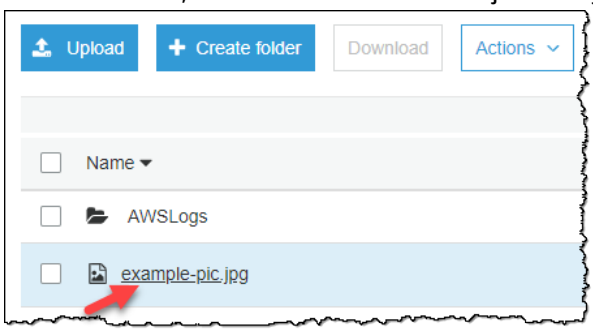

4. Choose **Properties**.

5. Choose **Object Lock**.

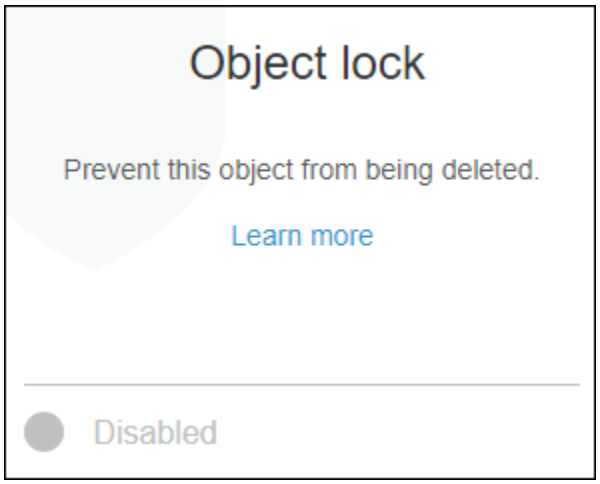

6. Choose a retention mode. You can change the **Retain until date**. You can also choose to enable a legal hold. For more information, see S3 Object Lock Overview in the *Amazon Simple Storage Service Developer Guide*.

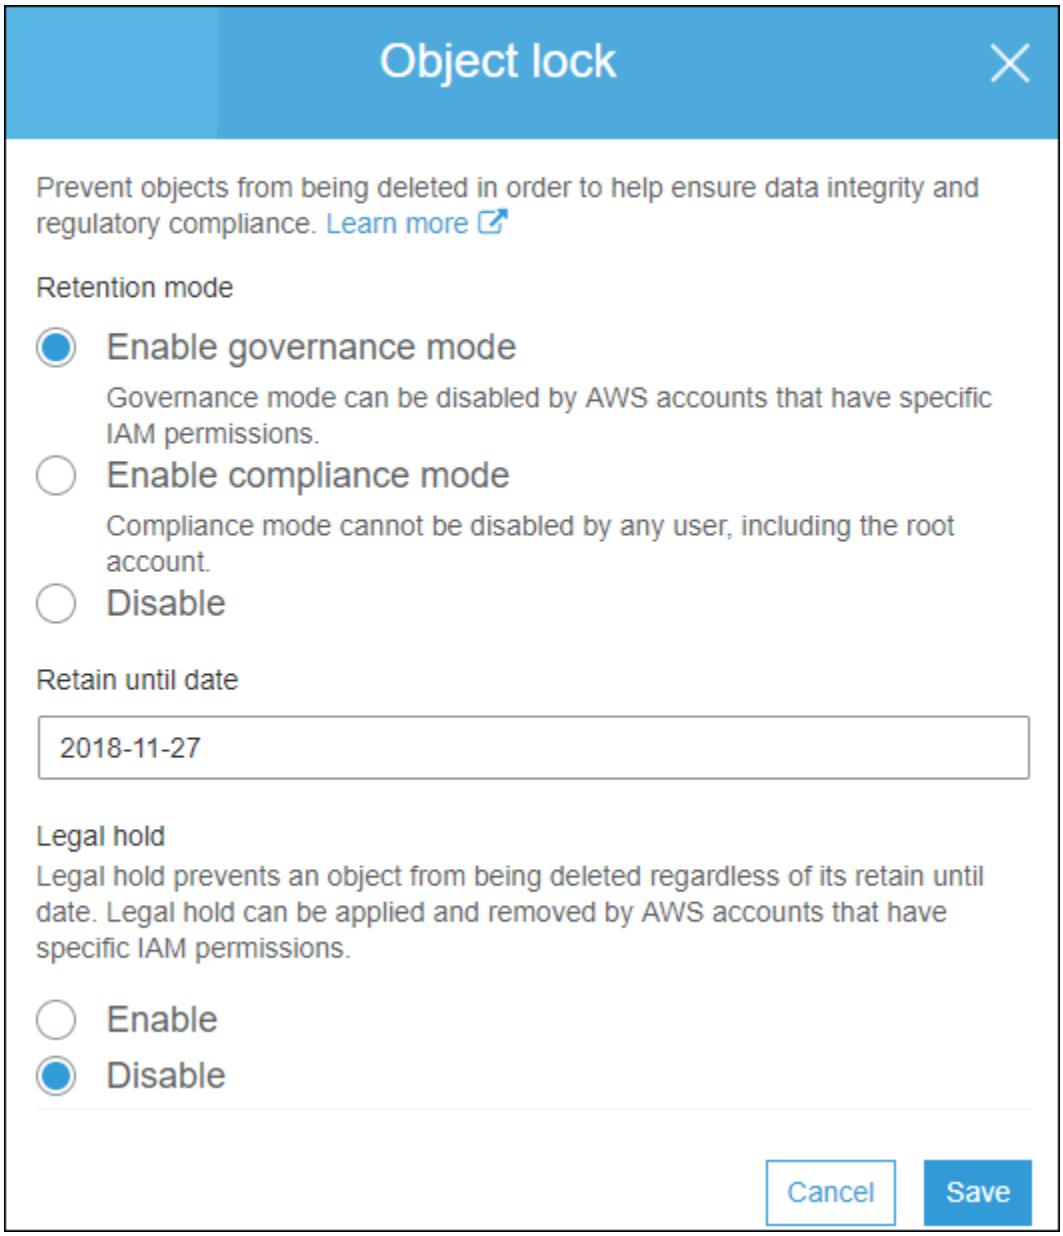

7. Choose **Save**.

More info

- Setting bucket and object access permissions (p. 112)

How do I see an overview of an object?

This section explains how to use the Amazon S3 console to view the object overview panel. This panel provides an overview of all the essential information for an object in one place.

To see the overview panel for an object

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

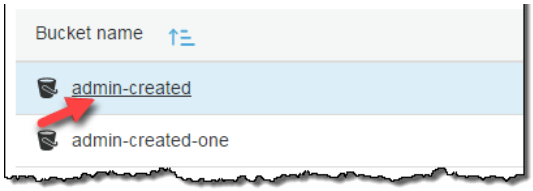

3. In the **Name** list, select the check box next to the name of the object for which you want an overview.

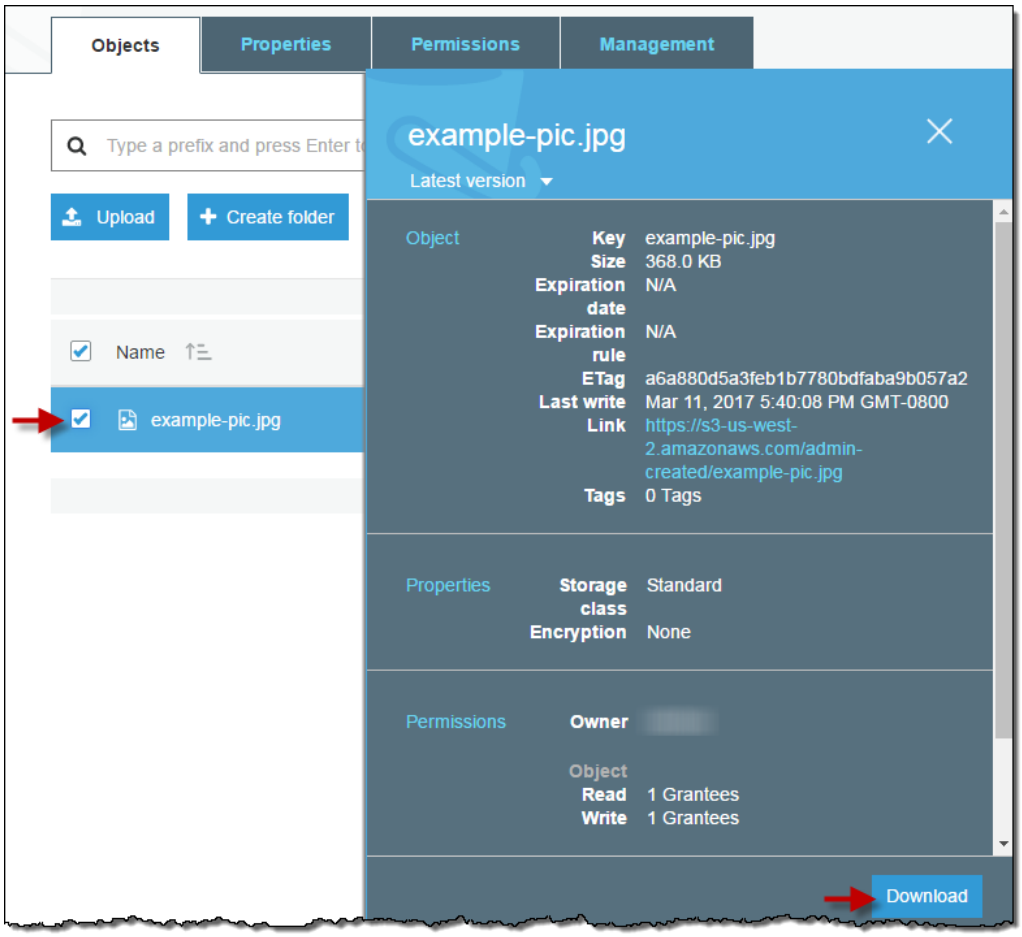

4. To download the object, choose **Download** in the object overview panel. To copy the path of the object to the clipboard, choose **Copy Path**.

5. If versioning is enabled on the bucket, choose **Latest versions** to list the versions of the object. You can then choose the download icon to download an object version, or choose the trash can icon to delete an object version.

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

More info

- How do I see the versions of an S3 object? (p. 58)

How do I see the versions of an S3 object?

This section explains how to use the Amazon S3 console to see the different versions of an object.

A versioning-enabled bucket can have many versions of the same object:, one current (latest) version and zero or more noncurrent (previous) versions. Amazon S3 assigns each object a unique version ID. For information about enabling versioning, see How do I enable or suspend versioning for an S3 bucket? (p. 7).

If a bucket is versioning-enabled, Amazon S3 creates another version of an object under the following conditions:

- If you upload an object that has the same name as an object that already exists in the bucket, Amazon S3 creates another version of the object instead of replacing the existing object.
- If you update any object properties after you upload the object to the bucket, such as changing the storage details or other metadata , Amazon S3 creates a new object version in the bucket.

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

To see multiple versions of an object

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

3. To see a list of the versions of the objects in the bucket, choose **Show**. For each object version, the console shows a unique version ID, the date and time the object version was created, and other properties. (Objects stored in your bucket before you set the versioning state have a version ID of **null**.)

   To list the objects without the versions, choose **Hide**.

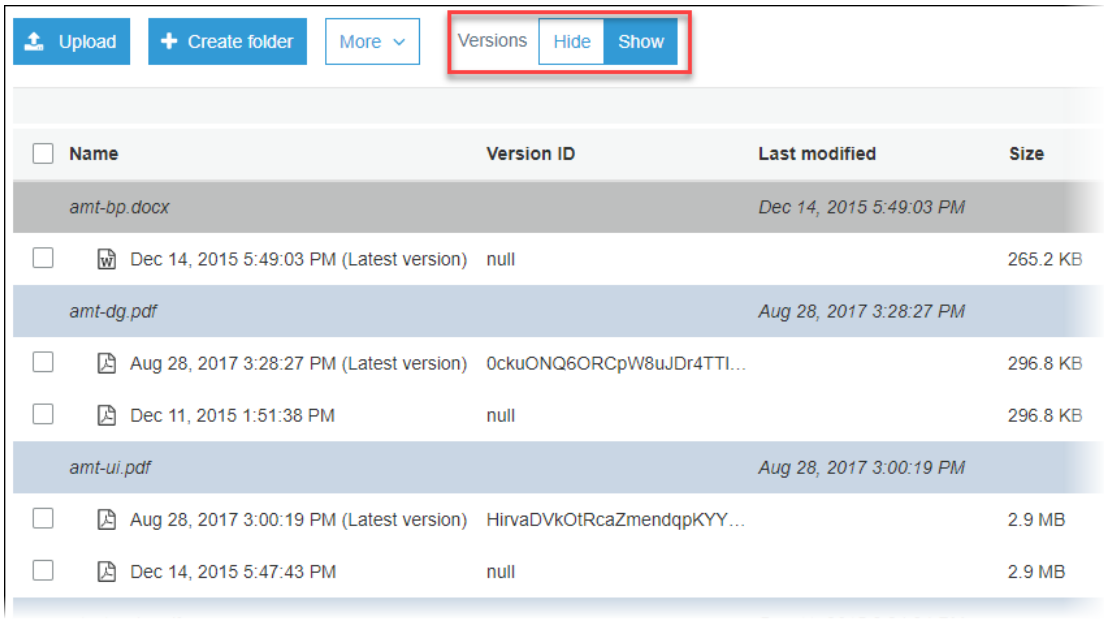

You also can view, download, and delete object versions in the object overview panel. For more information, see How do I see an overview of an object? (p. 55).

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

More info

- How do I enable or suspend versioning for an S3 bucket? (p. 7)
- How do I create a lifecycle policy for an S3 Bucket? (p. 79)

How do I view the properties of an object?

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

To view the properties of an object

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

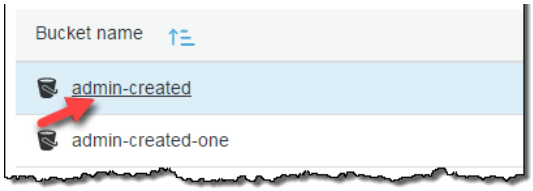

3. In the **Name** list, choose the name of the object you want to view the properties for.

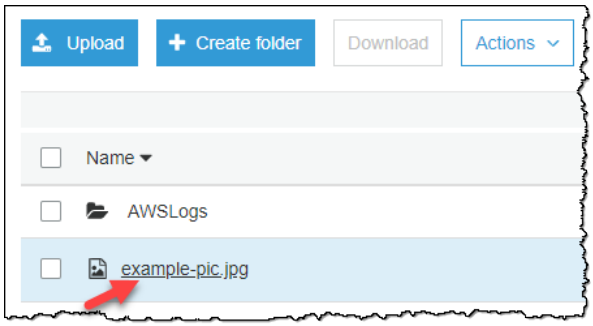

4. Choose **Properties**.

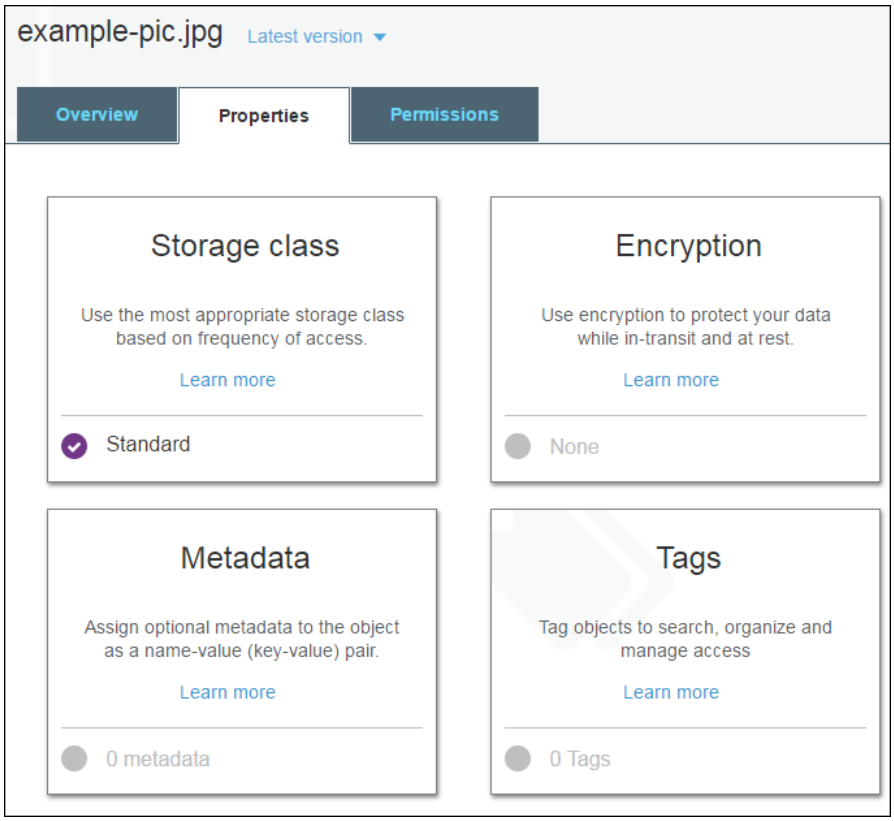

5. On the **Properties** page, you can configure the following properties for the object.

   Note
   If you change the properties; **Storage Class**, **Enryption**, **Metadata**, a new object is created to replace the old one. If S3 Versioning is enabled, a new version of the object is created, and the existing object becomes an older version. The role that changes the property also becomes the owner of the new object or (object version).

   a. **Storage class** – Each object in Amazon S3 has a storage class associated with it. The storage class that you choose to use depends on how frequently you access the object. The default storage class for S3 objects is STANDARD. You choose which storage class to use when you upload an object. For more information about storage classes, see Storage Classes in the *Amazon Simple Storage Service Developer Guide*.

      To change the storage class after you upload an object, choose **Storage class**. Choose the storage class that you want, and then choose **Save**.

b. **Encryption** – You can encrypt your S3 objects. For more information, see How do I add encryption to an S3 object? (p. 61).

c. **Metadata** – Each object in Amazon S3 has a set of name-value pairs that represents its metadata. For information on adding metadata to an S3 object, see How do I add metadata to an S3 object? (p. 63).

d. **Tags** – You can add tags to an S3 object. For more information, see How do I add tags to an S3 object? (p. 70).

How do I add encryption to an S3 object?

This topic describes how to set or change the type of encryption an object using the Amazon S3 console.

> **Note**
> If you change an object's encryption, a new object is created to replace the old one. If S3 Versioning is enabled, a new version of the object is created, and the existing object becomes an older version. The role that changes the property also becomes the owner of the new object or (object version).

To add or change encryption for an object

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the **Bucket name** list, choose the name of the bucket that contains the object.
3. In the **Name** list, choose the name of the object that you want to add or change encryption for.
4. Choose **Properties**, and then choose **Encryption**.

   The **Encryption** dialog box opens, giving you three choices for object encryption:

   - **None**- No object encryption.
   - **AES-256** - Server-side encryption with Amazon S3 managed keys (SSE-S3).
   - **AWS-KMS** - Server-side encryption with AWS Key Management Service (AWS KMS) customer master keys (SSE-KMS).
5. If you want to remove encryption from an object that already has encryption settings, choose **None** and then choose **Save**.

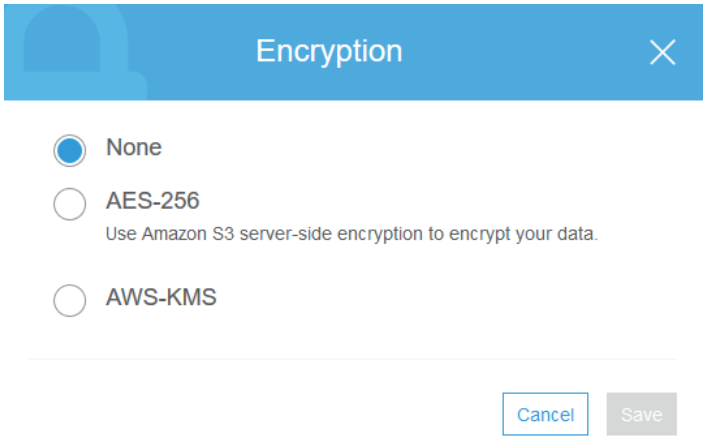

6. If you want to encrypt your object using keys that are managed by Amazon S3, follow these steps:

   a. Choose**AES-256**.

For more information about using Amazon S3 server-side encryption to encrypt your data, see Protecting Data with Amazon S3-Managed Encryption Keys Classes in the *Amazon Simple Storage Service Developer Guide*.

b. Choose **Save**.

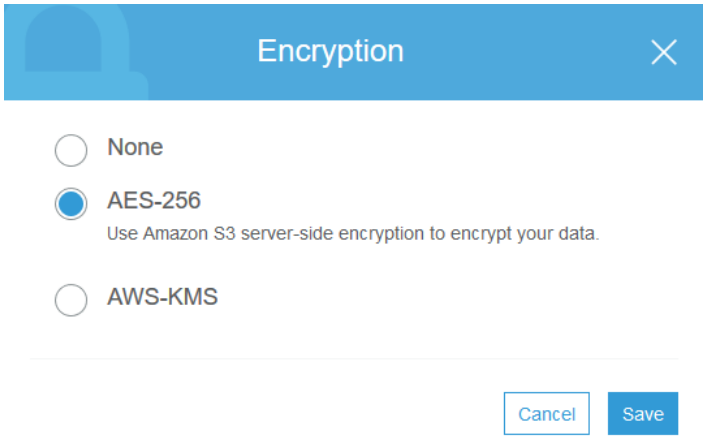

7. If you want to encrypt your object using AWS KMS, follow these steps:

   a. Choose **AWS-KMS**.

   b. Choose an AWS KMS customer master key (CMK).

   The list shows customer managed CMKs that you have created and your AWS managed CMK for Amazon S3. For more information about creating a customer managed AWS KMS CMK, see Creating Keys in the *AWS Key Management Service Developer Guide*.

   Important
   The Amazon S3 console lists only 100 AWS KMS CMKs per AWS Region. If you have more than 100 CMKs in the same Region, you can see only the first 100 CMKs in the S3 console. To use a KMS CMK that is not listed in the console, choose **Custom KMS ARN**, and enter the KMS CMK ARN.

   c. Choose **Save**.

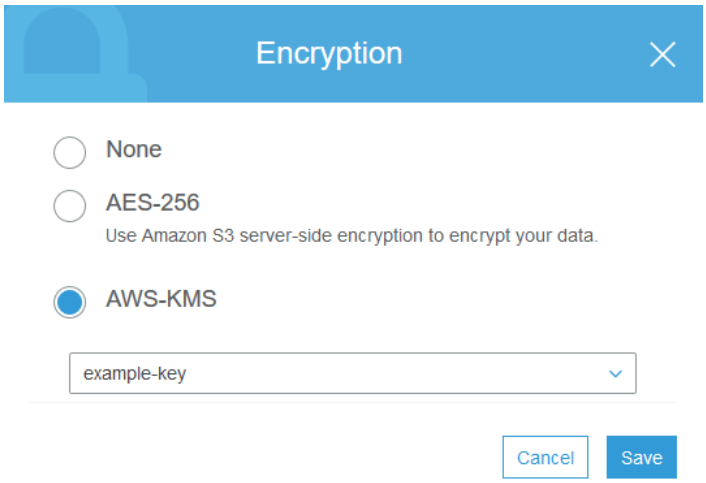

Important
To encrypt objects in the bucket, you can use only CMKs that are enabled in the same AWS Region as the bucket. Amazon S3 only supports symmetric CMKs. Amazon S3 does not

support asymmetric CMKs. For more information, see Using Symmetric and Asymmetric Keys.

8. To give an external account the ability to use an object that is protected by an AWS KMS CMK, follow these steps:

   a. Choose **AWS-KMS**.

   b. Enter the Amazon Resource Name (ARN) for the external account.

   c. Choose **Save**.

   Administrators of an external account that have usage permissions to an object protected by your AWS KMS CMK can further restrict access by creating a resource-level AWS Identity and Access Management (IAM) policy.

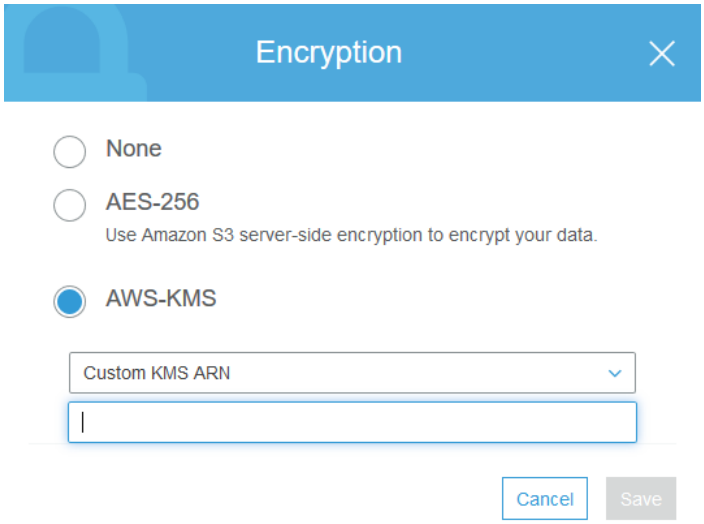

Note
This action applies encryption to all specified objects. When encrypting folders, wait for the save operation to finish before adding new objects to the folder.

More info

- How do I enable default encryption for an Amazon S3 bucket? (p. 8)
- Amazon S3 default encryption for S3 buckets in the *Amazon Simple Storage Service Developer Guide*
- How do I view the properties of an object? (p. 59)
- Uploading, downloading, and managing objects (p. 33)

How do I add metadata to an S3 object?

Each object in Amazon Simple Storage Service (Amazon S3) has a set of name-value pairs that provides metadata about the object. *Metadata* is additional information about the object. Some metadata is set by Amazon S3 when you upload the object, for example,`Date` and `Content-Length`. You can also set some metadata when you upload the object, or you can add it later. This section explains how to use the Amazon S3 console to add metadata to an S3 object.

Object metadata is a set of name-value (key-value) pairs. For example, the metadata for content length, `Content-Length`, is the name (key) and the size of the object in bytes (value). For more information about object metadata, see Object Metadata in the *Amazon Simple Storage Service Developer Guide*.

There are two kinds of metadata for an S3 object, Amazon S3 system metadata and user-defined metadata:

- **System metadata**–There are two categories of system metadata. Metadata such as the `Last-Modified` date is controlled by the system. Only Amazon S3 can modify the value. There is also system metadata that you control, for example, the storage class configured for the object.
- **User-defined metadata**–You can define your own custom metadata, called user-defined metadata. You can assign user-defined metadata to an object when you upload the object or after the object has been uploaded. User-defined metadata is stored with the object and is returned when you download the object. Amazon S3 does not process user-defined metadata.

The following topics describe how to add metadata to an object using the Amazon S3 console.

> **Note**
> If you change an object's metadata, a new object will be created and will replace the old one. If S3 Versioning is enabled, a new version of the object is created, and the existing object becomes an older version. The role that changes the property also becomes the owner of the new object or (object version).

Topics

- Adding system-defined metadata to an S3 object (p. 64)
- Adding user-defined metadata to an S3 object (p. 67)

Adding system-defined metadata to an S3 object

You can configure some system metadata for an S3 object. For a list of system-defined metadata and whether you can modify their values, see System-Defined Metadata in the *Amazon Simple Storage Service Developer Guide*.

To add system metadata to an object

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

3. In the **Name** list, choose the name of the object that you want to add metadata to.

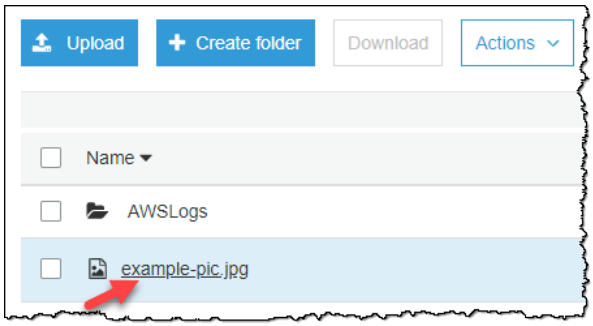

4. Choose **Properties**, and then choose **Metadata**.

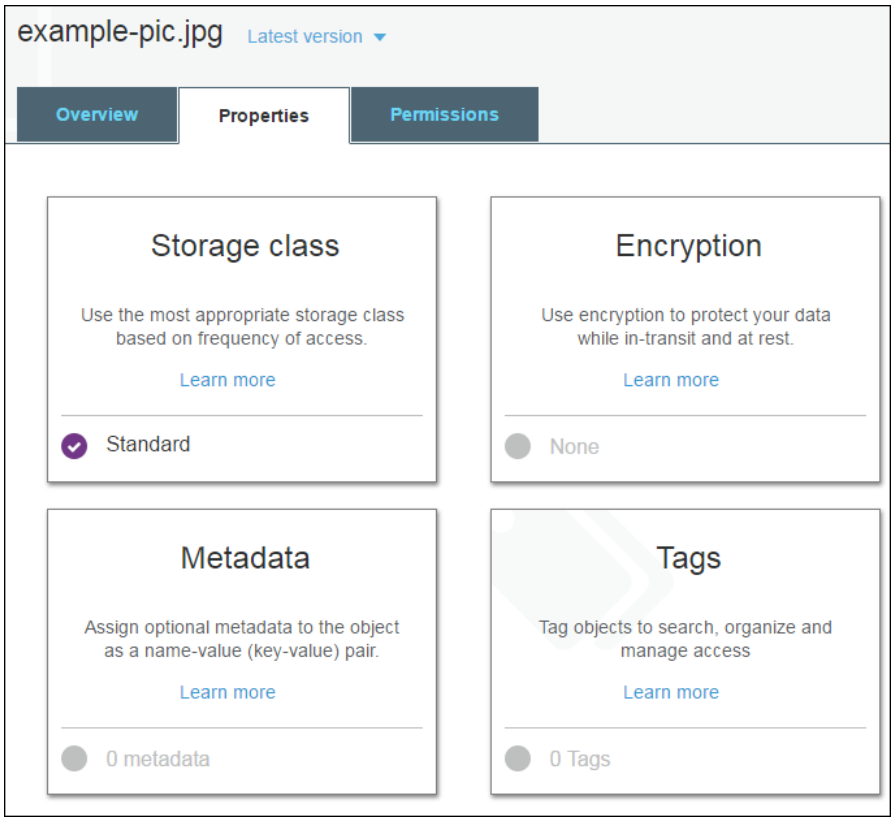

5. Choose **Add Metadata**, and then choose a key from the **Select a key** menu.

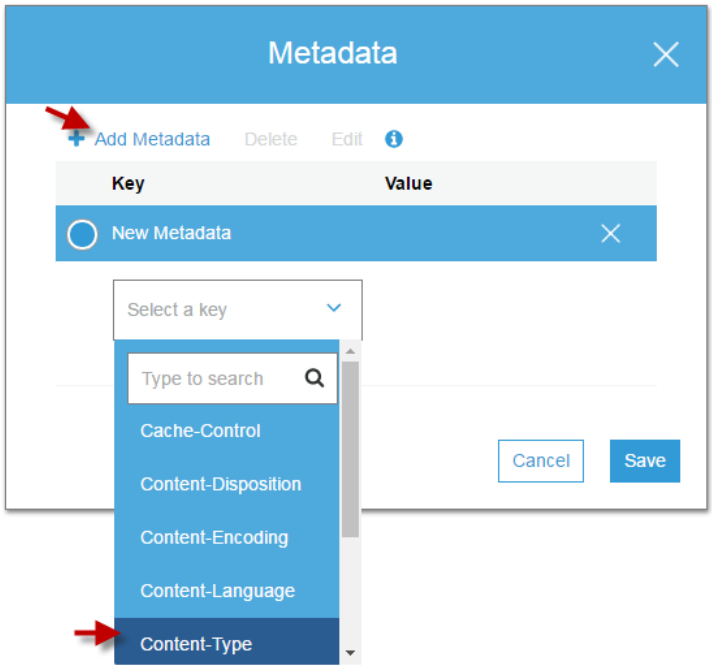

6. Depending on which key you chose, choose a value from the **Select a value** menu or type a value.

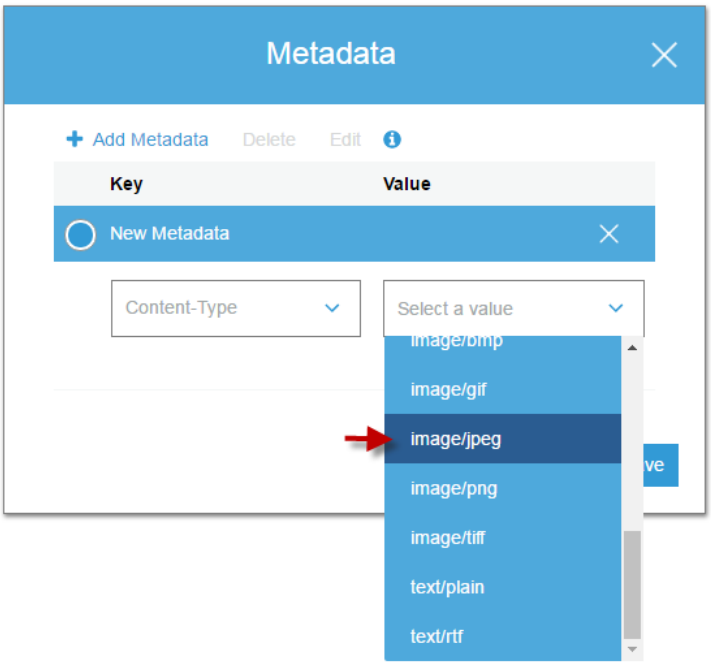

7. Choose **Save**.

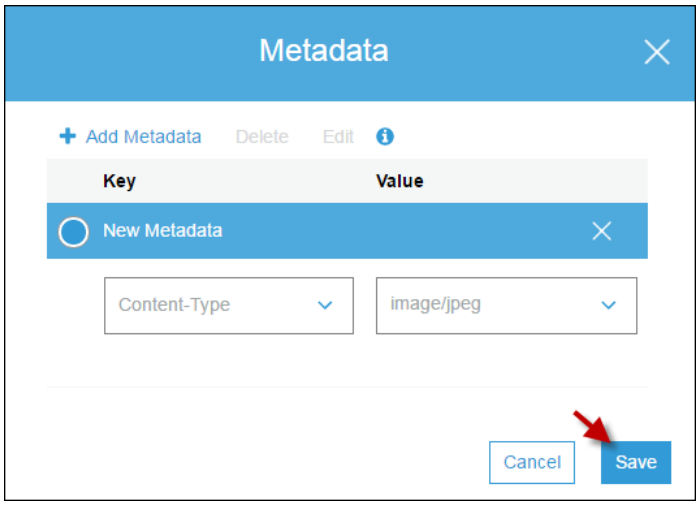

Note
This action adds metadata to all specified objects. When adding metadata to folders, wait for the save operation to finish before adding new objects to the folder.

Adding user-defined metadata to an S3 object

You can assign user-defined metadata to an object. User-defined metadata must begin with the prefix "`x-amz-meta-`", otherwise Amazon S3 will not set the key value pair as you define it. You define custom metadata by adding a name that you choose to the `x-amz-meta-` key. This creates a custom key. For example, if you add the custom name `alt-name`, the metadata key would be `x-amz-meta-alt-name`.

User-defined metadata can be as large as 2 KB. Both keys and their values must conform to US-ASCII standards. For more information, see User-Defined Metadata in the *Amazon Simple Storage Service Developer Guide*.

To add user-defined metadata to an object

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

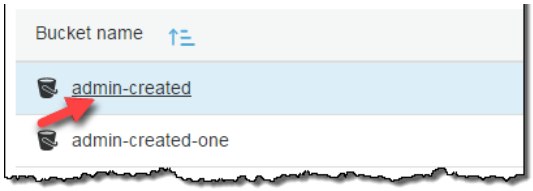

3. In the **Name** list, choose the name of the object that you want to add metadata to.

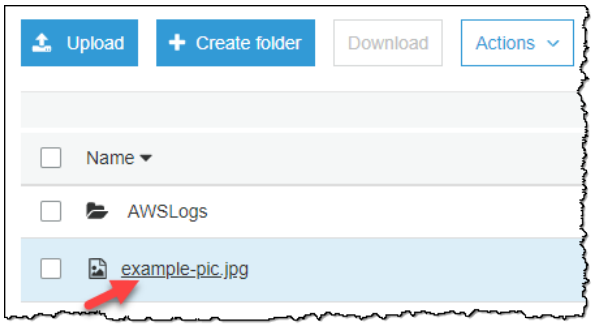

4. Choose **Properties**, and then choose **Metadata**.

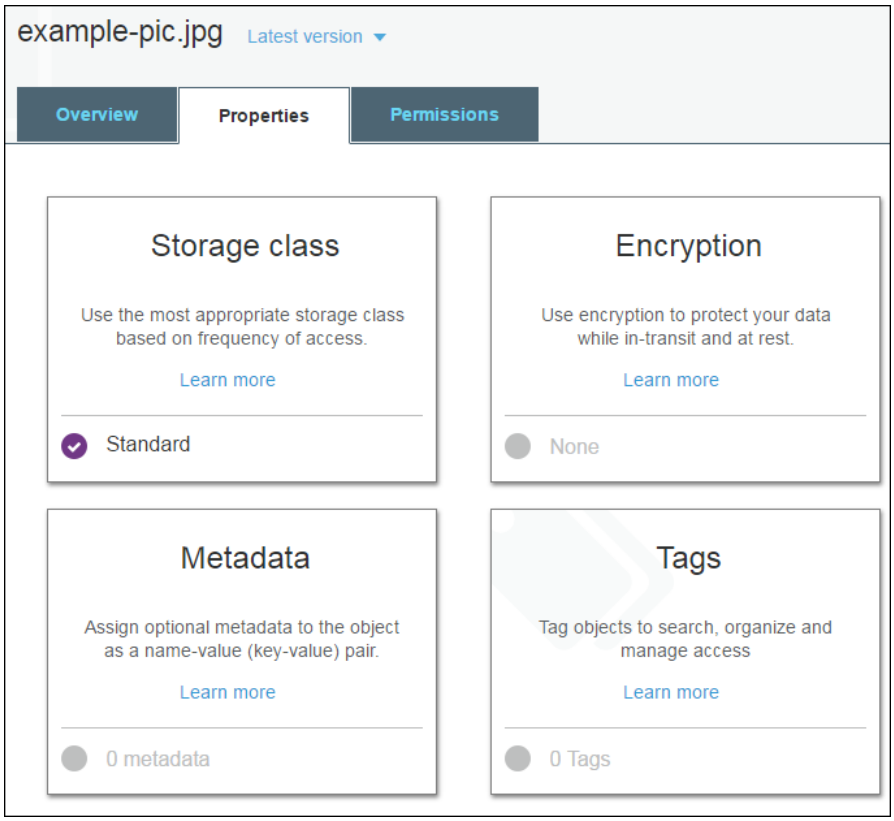

5. Choose **Add Metadata**, and then choose the `x-amz-meta-` key from the **Select a key** menu. Any metadata starting with the prefix `x-amz-meta-` is user-defined metadata.

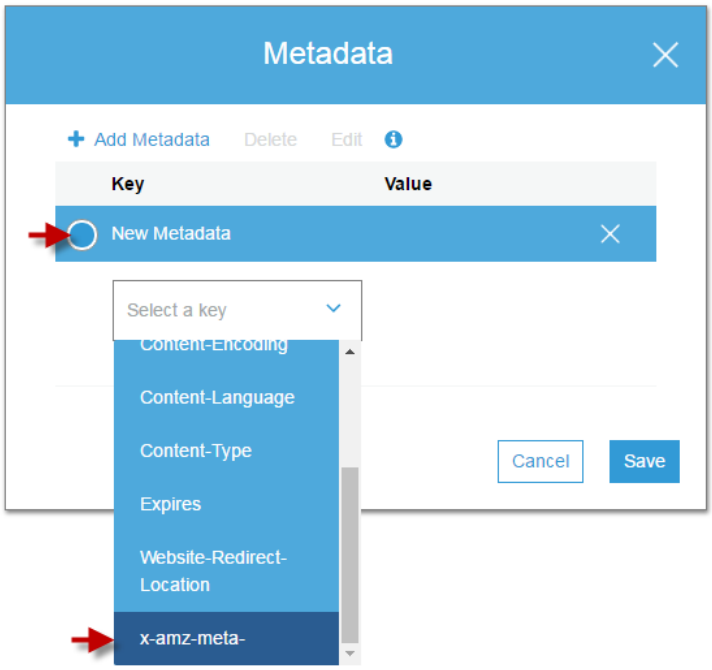

6. Type a custom name following the `x-amz-meta-` key. For example, for the custom name `alt-name`, the metadata key would be `x-amz-meta-alt-name`. Enter a value for the custom key, and then choose **Save**.

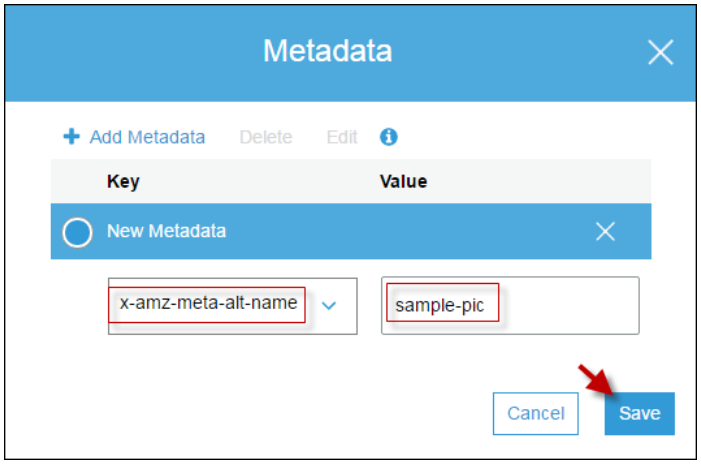

Note
When adding metadata to folders, wait for the save to complete before adding new objects. Otherwise, new objects might be edited as well.

More info

- How do I view the properties of an object? (p. 59)
- Uploading, downloading, and managing objects (p. 33)

How do I add tags to an S3 object?

Object tagging gives you a way to categorize storage. This topic explains how to use the console to add tags to an S3 object after the object has been uploaded. For information about adding tags to an object when the object is being uploaded, see How do I upload files and folders to an S3 bucket? (p. 34).

Each tag is a key-value pair that adheres to the following rules:

- You can associate up to 10 tags with an object. Tags associated with an object must have unique tag keys.
- A tag key can be up to 128 Unicode characters in length and tag values can be up to 255 Unicode characters in length.
- Key and tag values are case sensitive.

For more information about object tags, see Object Tagging in the *Amazon Simple Storage Service Developer Guide*.

To add tags to an object

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

3. In the **Name** list, choose the name of the object you want to add tags to.

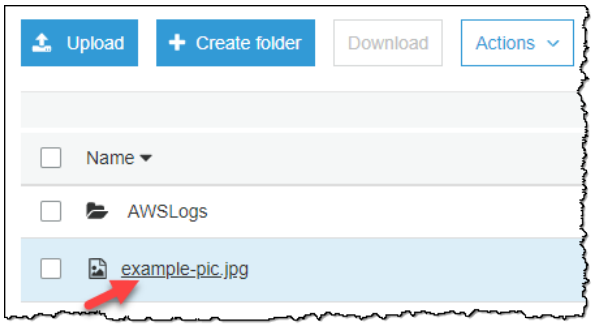

4. Choose **Properties**.

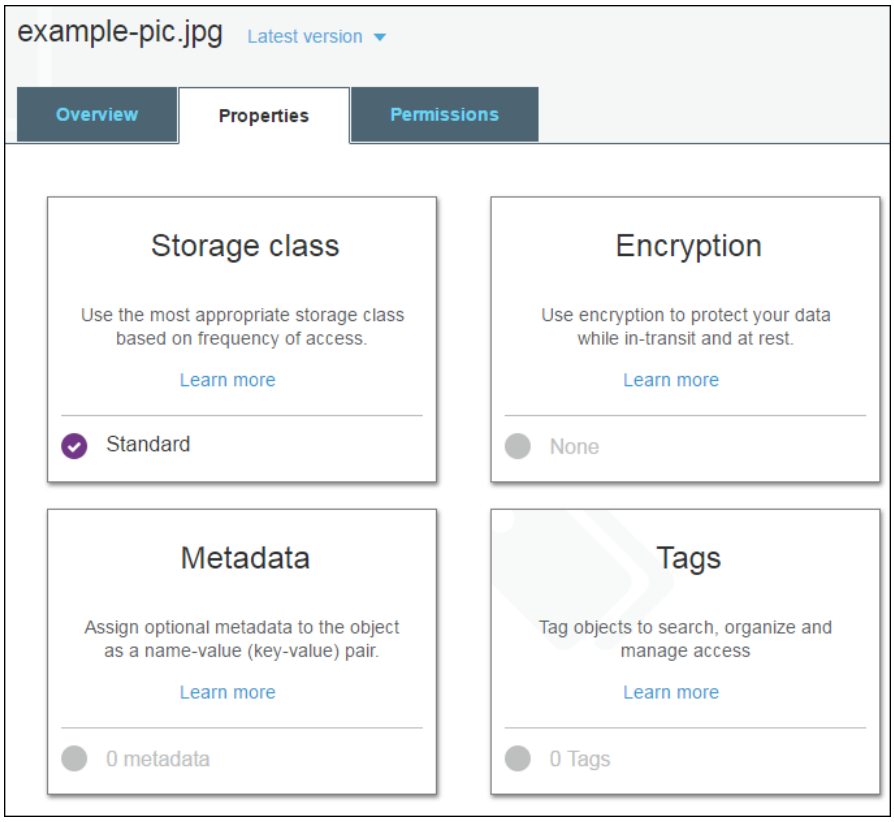

5. Choose **Tags** and then choose **Add Tag**.

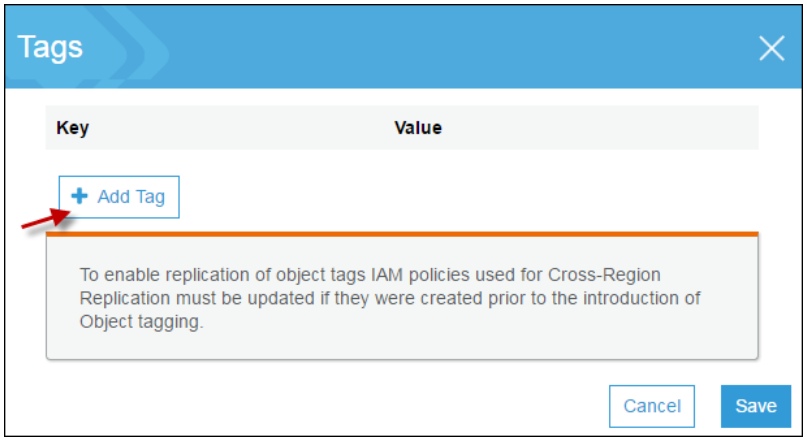

6. Each tag is a key-value pair. Type a **Key** and a **Value**. Then choose **Add Tag** to add another tag or choose **Save**.

   You can enter up to 10 tags for an object.

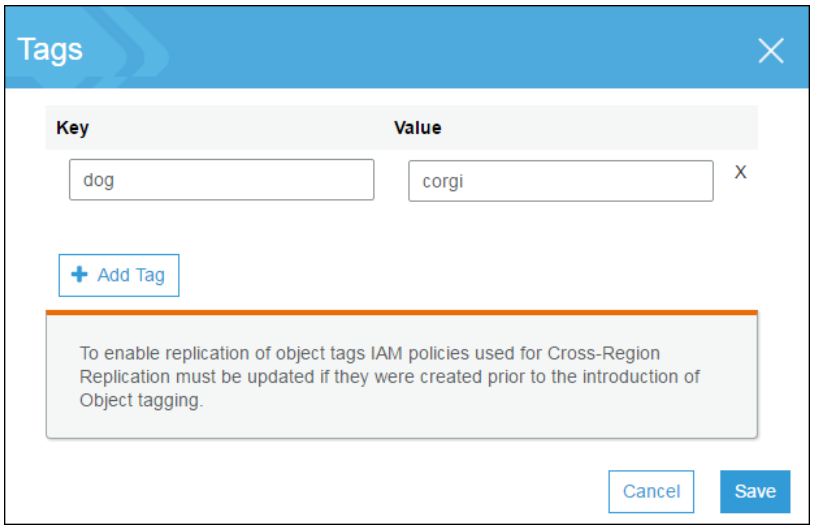

More info

- How do I view the properties of an object? (p. 59)
- Uploading, downloading, and managing objects (p. 33)

How do I use folders in an S3 bucket?

In Amazon S3, buckets and objects are the primary resources, and objects are stored in buckets. Amazon S3 has a flat structure instead of a hierarchy like you would see in a 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 a shared name prefix for objects (that is, objects have names that begin with a common string). Object names are also referred to as *key names*.

For example, you can create a folder on the console named `photos` and store an object named `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/2017/example.jpg`, the console will show you a folder named `photos` containing the folder `2017` and the object `example.jpg`.

Topics

- Creating a folder (p. 73)
- How do I delete folders from an S3 bucket? (p. 74)
- Making folders public (p. 76)

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 copied from one folder to another.

Important

- The Amazon S3 console implements folder object creation by creating a zero-byte objects with the folder *prefix and delimiter* value as the key. These folder objects don't appear in the console, but otherwise behave like any other objects and can be viewed and manipulated through the API, AWS CLI and SDKs.
- The Amazon S3 console treats all objects that have a forward slash "/" character as the last (trailing) character in the key name as a folder, for example `examplekeyname/`. You can't upload an object that has a key name with a trailing "/" character using the Amazon S3 console. However, you can upload objects that are named with a trailing "/" with the Amazon S3 API by using the AWS CLI, AWS SDKs, or REST API.
- An object that is named with a trailing "/" appears as a folder in the Amazon S3 console. The Amazon S3 console does not display the content and metadata for such an object. When you use the console to copy an object named with a trailing "/", a new folder is created in the destination location, but the object's data and metadata are not copied.

Creating a folder

This section describes how to use the Amazon S3 console to create a folder.

To create a folder

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

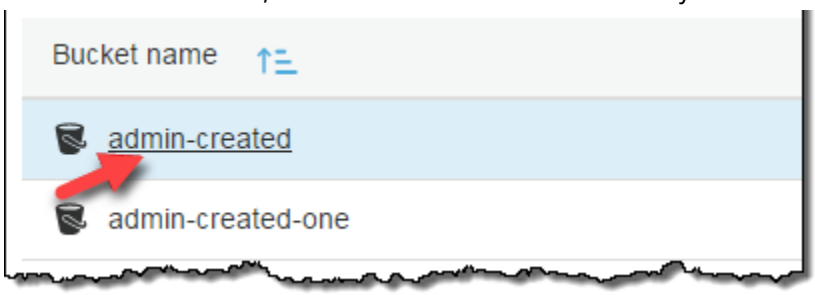

3. Choose **Create folder**.

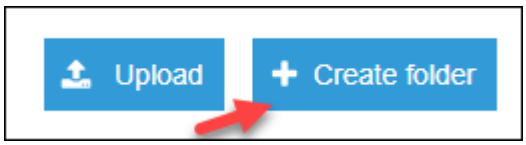

4. Enter a name for the folder (for example, **`favorite-pics`**). Choose the encryption setting for the folder object, and then choose **Save**.

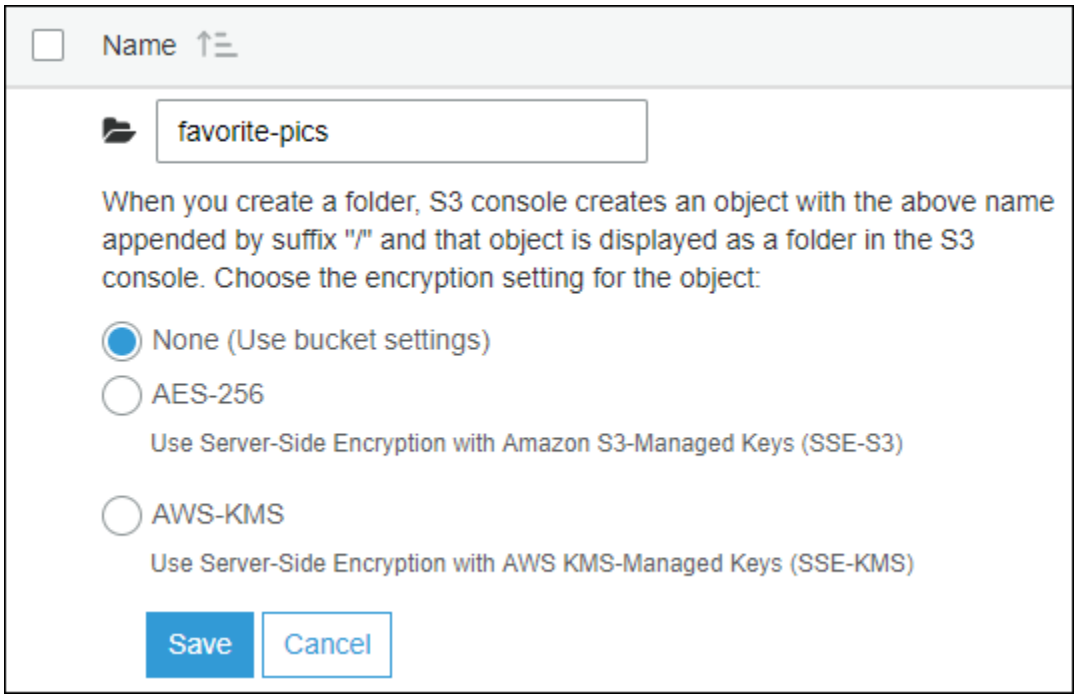

5. If you don't want to use encryption with your folder, choose **None**.
6. If you want to encrypt your object using keys that are managed by Amazon S3, follow these steps:
   - Choose **AES-256**.

     For more information about using Amazon S3 server-side encryption to encrypt your data, see Protecting data with Amazon S3 managed encryption keys classes in the *Amazon Simple Storage Service Developer Guide*.
7. If you want to encrypt your object using AWS KMS, follow these steps:
   a. Choose **AWS-KMS**.
   b. Choose an AWS KMS customer master key (CMK).

      The list shows customer managed CMKs that you have created and your AWS managed CMK for Amazon S3. For more information about creating a customer managed AWS KMS CMK, see Creating Keys in the *AWS Key Management Service Developer Guide*.

      Important
      To encrypt objects in the bucket, you can use only the CMKs that are enabled in the same AWS Region as the bucket. Amazon S3 only supports *symmetric* CMKs and does not support *asymmetric* CMKs. For more information, see Using symmetric and asymmetric keys.
      The Amazon S3 console only lists 100 KMS CMKs per Region. If you have more than 100 CMKs in the same Region, you can see only the first 100 CMKs in the S3 console. To use anAWS KMS KMS CMK that is not listed in the console, choose **Custom KMS ARN**, and enter the KMS CMK ARN.
8. Choose **Save**.

How do I delete folders from an S3 bucket?

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

For information about Amazon S3 features and pricing, see Amazon S3.

To delete folders from an S3 bucket

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

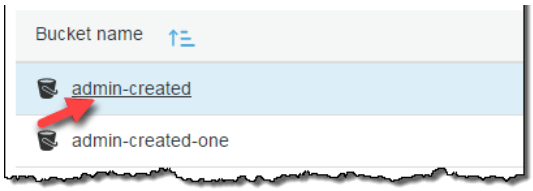

3. In the **Name** list, select the check box next to the folders and objects that you want to delete, choose **More**, and then choose **Delete**.

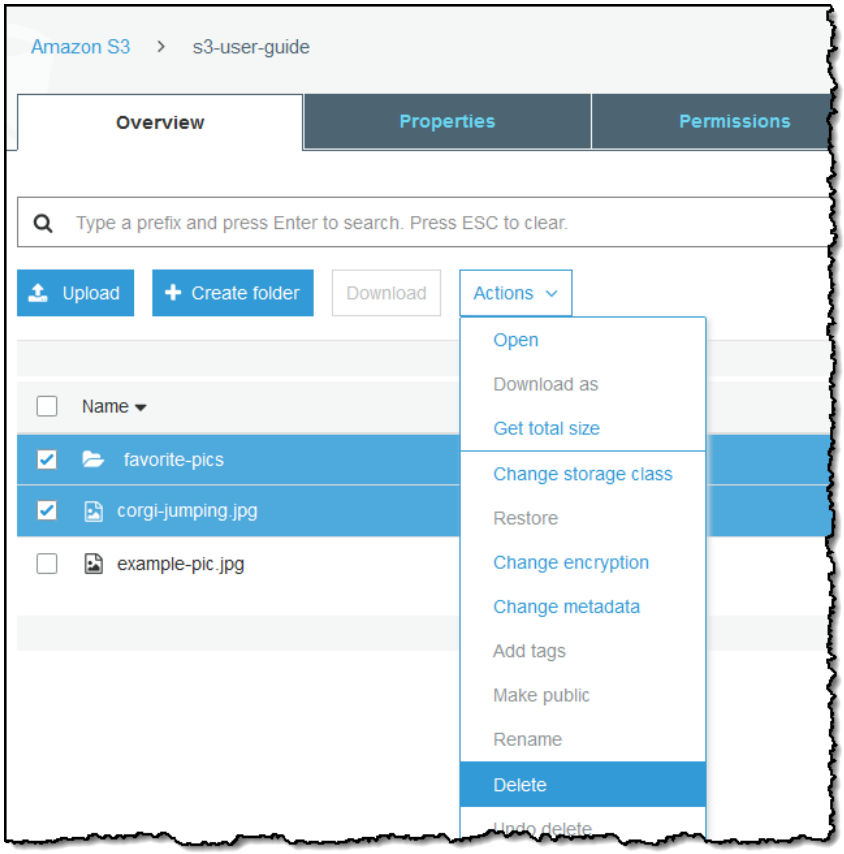

In the **Delete objects** dialog box, verify that the names of the folders you selected for deletion are listed and then choose **Delete**.

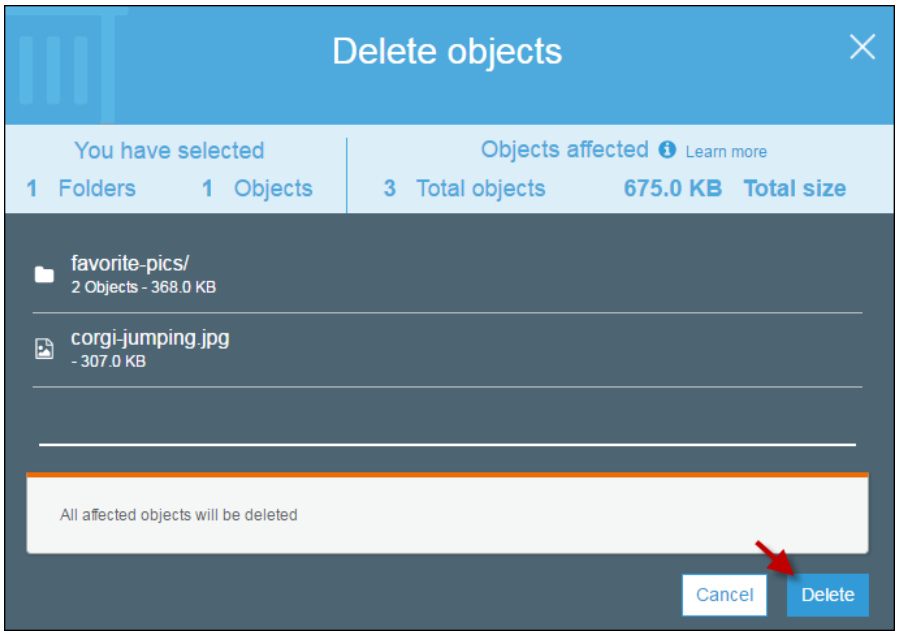

Warning
This action deletes all specified objects. When deleting folders, wait for the delete action to finish before adding new objects to the folder. Otherwise, new objects might be deleted as well.

Related topics

- How do I delete objects from an S3 bucket? (p. 46)

Making folders public

Amazon S3 has a flat structure instead of a hierarchy like you would typically see in a file system. However, for the sake of organizational simplicity, the Amazon S3 console supports a folder concept as a way to group objects. In Amazon S3, the folder is a naming prefix for an object or group of objects. For more information, see How do I use folders in an S3 bucket? (p. 72)

We recommend blocking all public access to your Amazon S3 folders and buckets unless you specifically require a public folder or bucket. When you make a folder public, anyone on the internet can view all the objects that are grouped in that folder. In the Amazon S3 console, you can make a folder public. You can also make a folder public by creating a bucket policy that limits access by prefix. For more information, see Setting bucket and object access permissions (p. 112).

Warning
After you make a folder public in the Amazon S3 console, you can't make it private again. Instead, you must set permissions on each individual object in the public folder so that the objects have no public access. For more information, see How do I set permissions on an object? (p. 118)

More info

- How do I delete folders from an S3 bucket? (p. 74)
- How do I set ACL bucket permissions? (p. 121)
- How do I block public access to S3 buckets? (p. 113)

Introduction to S3 Batch Operations

S3 Batch Operations performs large-scale batch operations on Amazon S3 objects. You can use S3 Batch Operations to copy objects, set object tags or access control lists (ACLs), initiate object restores from Amazon S3 Glacier, or invoke an AWS Lambda function to perform custom actions using your objects. You can perform these operations on a custom list of objects, or you can use an Amazon S3 inventory report to make generating even the largest lists of objects easy. S3 Batch Operations use the same Amazon S3 APIs that you already use, so you'll find the interface familiar. For information about performing S3 Batch Operations using the AWS CLI, AWS SDKs, and the Amazon S3 REST APIs, see Performing S3 Batch Operations in the *Amazon Simple Storage Service Developer Guide*.

The following topics explain how to use the Amazon S3 console to configure and run batch operations.

Topics

- Creating an S3 Batch Operations job (p. 77)
- Managing S3 Batch Operations jobs (p. 78)

Creating an S3 Batch Operations job

This section describes how to create a S3 Batch Operations job. For information about performing Batch Operations using the AWS CLI, AWS SDKs, and the Amazon S3 REST APIs, see Performing S3 Batch Operations in the *Amazon Simple Storage Service Developer Guide*.

To create a batch job

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. Choose **Batch Operations** on the navigation pane of the Amazon S3 console.
3. Choose **Create job**.
4. Choose the **Region** where you want to create your job.
5. Under **Manifest format**, choose the type of manifest object to use.
   - If you choose **S3 inventory report**, enter the path to the manifest.json object that Amazon S3 generated as part of the CSV-formatted Inventory report, and optionally the version ID for the manifest object if you want to use a version other than the most recent.
   - If you choose **CSV**, enter the path to a CSV-formatted manifest object. The manifest object must follow the format described in the console. You can optionally include the version ID for the manifest object if you want to use a version other than the most recent.
6. Choose **Next**
7. Under **Operation**, choose the operation that you want to perform on all objects listed in the manifest. Fill out the information for the operation you chose and then choose **Next**.
8. Fill out the information for **Configure additional options** and then choose **Next**.
9. For **Review**, verify the settings. If you need to make changes, choose **Previous**. Otherwise, choose **Create Job**.

More info

- The Basics: S3 Batch Operations Jobs in the *Amazon Simple Storage Service Developer Guide*

- Creating a S3 Batch Operations Job in the *Amazon Simple Storage Service Developer Guide*
- Operations in the *Amazon Simple Storage Service Developer Guide*

Managing S3 Batch Operations jobs

Amazon S3 provides a set of tools to help you manage your S3 Batch Operations jobs after you create them. For more information about managing S3 Batch Operations, see Managing S3 Batch Operations Jobs in the *Amazon Simple Storage Service Developer Guide*.

More info

- The Basics: S3 Batch Operations Jobs in the *Amazon Simple Storage Service Developer Guide*
- Creating a S3 Batch Operations Job in the *Amazon Simple Storage Service Developer Guide*
- Operations in the *Amazon Simple Storage Service Developer Guide*

Storage management

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

Topics

- How do I create a lifecycle policy for an S3 Bucket? (p. 79)
- How do I add a replication rule to an S3 bucket? (p. 83)
- How do I manage the replication rules for an S3 Bucket? (p. 98)
- How Do I Configure Storage Class Analysis? (p. 100)
- How Do I Configure Amazon S3 Inventory? (p. 104)
- How do I create a request metrics filter for all the objects in my S3 bucket? (p. 108)
- How do I create a request metrics filter that limits scope by object tag or prefix? (p. 108)
- How do I delete a request metrics filter? (p. 109)
- How do I view replication metrics? (p. 110)

How do I create a lifecycle policy for an S3 Bucket?

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

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

Using a lifecycle policy, you can define actions specific to current and noncurrent object versions. For more information, see Object Lifecycle Management and Object Versioning and Using Versioning in the *Amazon Simple Storage Service Developer Guide*.

To create a lifecycle policy

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

4. In the **Lifecycle rule** dialog box, type a name for your rule.

The name must be unique within the bucket.

5. Choose the scope of the lifecycle rule: all objects with a specific prefix or tag or all the objects in the bucket.

   - To apply this lifecycle rule to *all objects with a specific prefix or tag*, choose **Limit the scope to specific prefixes or tags**. In the **Add prefix or tag filter** box, type the prefix or tag name, and press **Enter**. For more information about object name prefixes, see Object Keys in the *Amazon Simple Storage Service Developer Guide*.For more information about object tags, see Object Tagging in the *Amazon Simple Storage Service Developer Guide*.

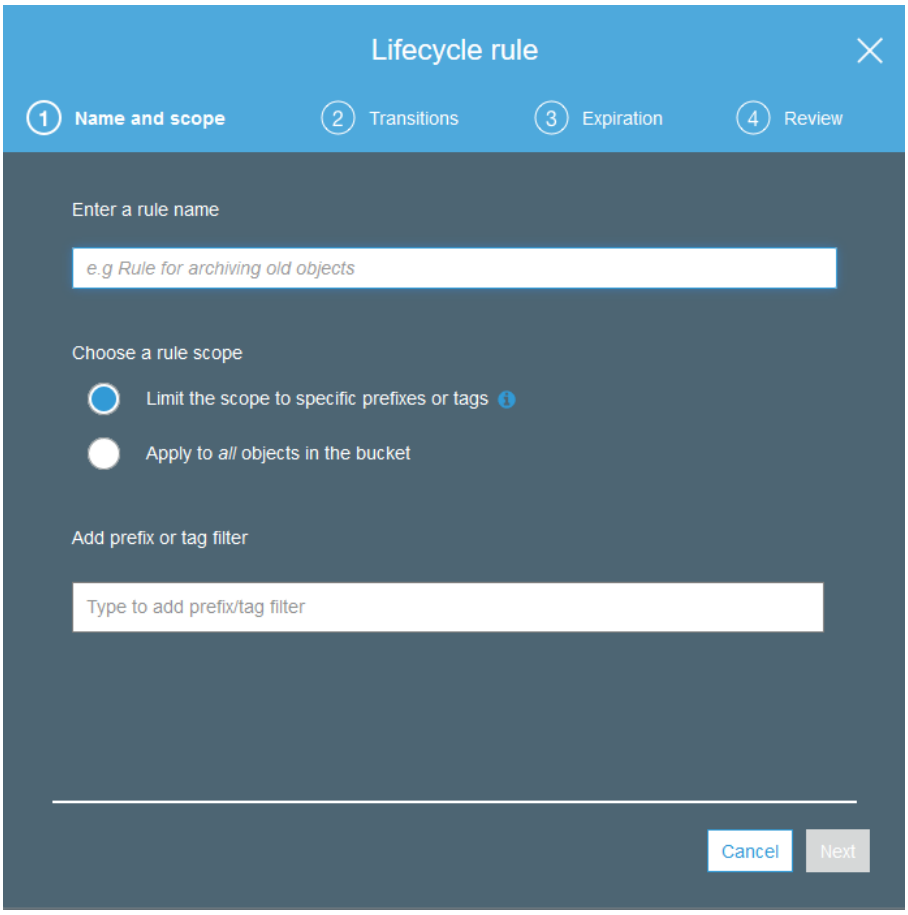

   - To apply this lifecycle rule to *all objects in the bucket*, choose **Apply to *all* objects in the bucket**.

6. Choose **Next**.

   The **Storage class transition** page opens. When you configure your storage class transitions, you define the rules to transition objects to the Standard-IA, One Zone-IA, Glacier, and Deep Archive storage classes. For more information, see Storage Classes in the *Amazon Simple Storage Service Developer Guide*. You can define transitions for current or previous object versions or for both current and previous versions. Versioning enables you to keep multiple versions of an object in one bucket. For more information about versioning, see How do I enable or suspend versioning for an S3 bucket? (p. 7).

7. Choose the versions for which you want to define transitions, current or noncurrent:

   - To define transitions that are applied to the current verion of the object, choose **Current version**.
   - To define transitions that are applied to all previous versions of the object, choose **Previous versions**.

In the screenshot below, **Current version** is chosen.

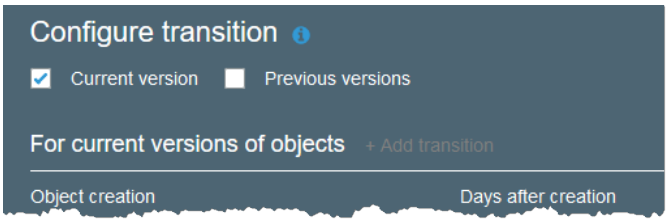

8. To add a transition:

   a. For a current version, under **For current object versions**, choose **Add transition**.

   b. For a non-current version, under **For non-current object versions**, choose **Add transition**.

9. For each transition that you add, choose one of the following:

   - **Transition to Standard-IA after**.
   - **Transition to Intelligent-Tiering after**.
   - **Transition to One Zone-IA after**.
   - **Transition to Glacier after**.
   - **Transition to Glacier Deep Archive after**.

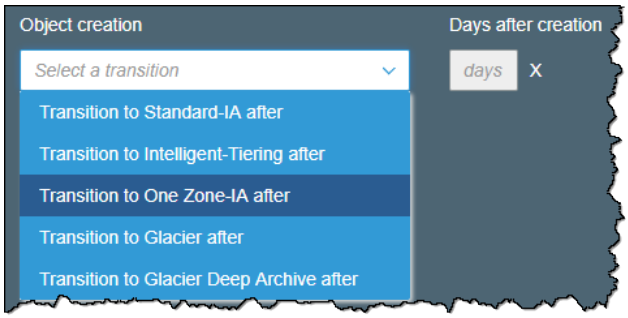

Important
When you choose the Glacier or Glacier Deep Archive storage class, your objects remain in Amazon S3. You cannot access them directly through the separate Amazon S3 Glacier service. For more information, see Transitioning Objects Using Amazon S3 Lifecycle.

10. In the **Days after creation** box, enter the number of days after the creation of the object that you want the transition to be applied (for example, 30 or 100 days).

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

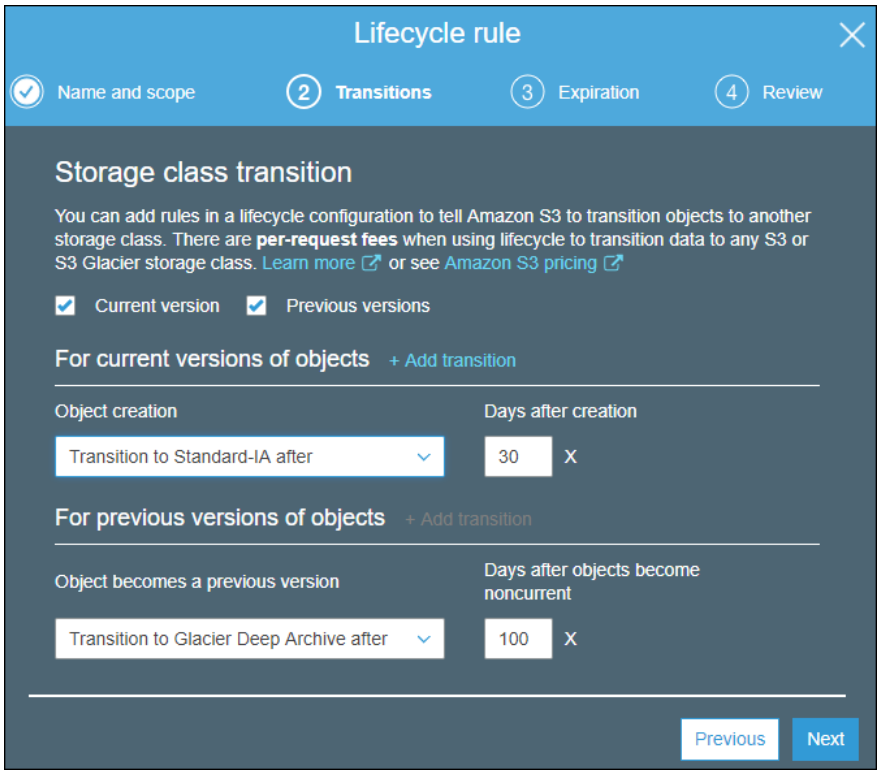

12. Under **Configure expiration**, for this example, choose both **Current version** and **Previous versions**.

    Important
    In a non-versioned bucket the expiration action results in Amazon S3 permanently removing the object. For more information about lifecycle actions, see Elements to describe lifecycle actions in the *Amazon Simple Storage Service Developer Guide*.

13. Choose **Expire current version of object**, and then enter the number of days after object creation to delete the object (for example, 395 days).

    If you choose this expire option, you cannot choose the option to clean up expired delete markers.

14. Choose **Permanently delete previous versions**, and then enter the number of days after an object becomes a previous version to permanently delete the object (for example, 465 days).

15. To clean up incomplete multipart uploads, we recommend that you choose **Clean up incomplete multipart uploads** and enter the number of days after the multipart upload initiation that you want to end and clean up incomplete multipart uploads(for example, 7 days).

    For more information about multipart uploads, see Multipart Upload Overview in the Amazon Simple Storage Service Developer Guide.

16. Choose **Next**.

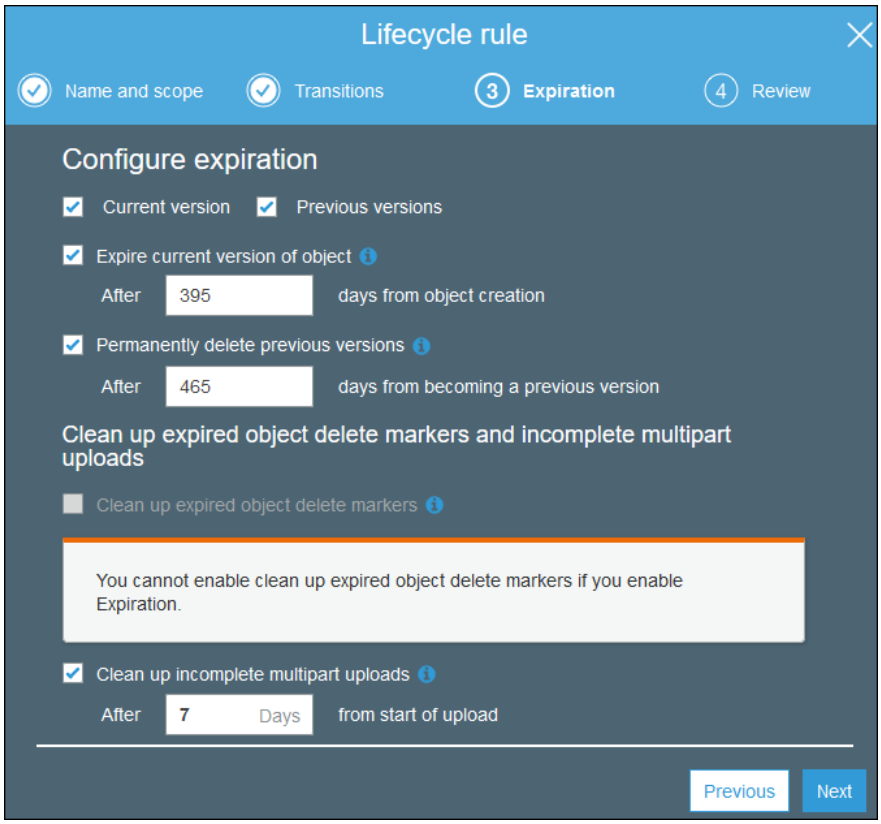

17. For **Review**, verify the settings for your rule. If you need to make changes, choose **Previous**. Otherwise, choose **Save**.

    If the rule does not contain any errors, it is enabled and you can see it on the **Lifecycle** page.

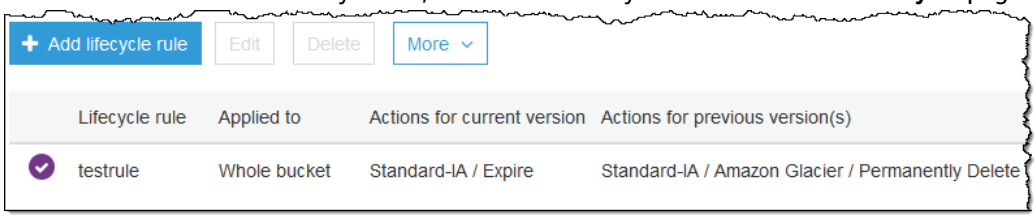

How do I add a replication rule to an S3 bucket?

Replication is the automatic, asynchronous copying of objects across buckets in the same or different AWS Regions. Replication copies newly created objects and object updates from a source bucket to a destination bucket. For more information about replication concepts and how to use replication with the AWS CLI, AWS SDKs, and the Amazon S3 REST APIs, see Replication in the *Amazon Simple Storage Service Developer Guide*.

Replication requires versioning to be enabled on both the source and destination buckets. To review the full list of requirements, see Requirements for Replication in the *Amazon Simple Storage Service Developer Guide*. For more information about versioning, see How do I enable or suspend versioning for an S3 bucket? (p. 7)

The object replicas in the destination bucket are exact replicas of the objects in the source bucket. They have the same key names and the same metadata—for example, creation time, owner, user-defined

metadata, version ID, access control list (ACL), and storage class. Optionally, you can explicitly specify a different storage class for object replicas. And regardless of who owns the source bucket or the source object, you can choose to change replica ownership to the AWS account that owns the destination bucket. For more information, see Changing the replica owner in the *Amazon Simple Storage Service Developer Guide*.

You can use S3 Replication Time Control (S3 RTC) to replicate your data in the same AWS Region or across different AWS Regions in a predictable timeframe. S3 RTC replicates 99.99 percent of new objects stored in Amazon S3 within 15 minutes and most objects within seconds. For more information, see Replicating Objects Using S3 Replication Time Control (S3 RTC) in the *Amazon Simple Storage Service Developer Guide*.

> **Note about replication and lifecycle rules**
> Metadata for an object remains identical between original objects and replica objects. Lifecycle rules abide by the creation time of the original object, and not by when the replicated object becomes available in the destination bucket. However, lifecycle does not act on objects that are pending replication until replication is complete.

You use the Amazon S3 console to add replication rules to the source bucket. Replication rules define which source bucket objects to replicate and the destination bucket where the replicated objects are stored. You can create a rule to replicate all the objects in a bucket or a subset of objects with a specific key name prefix, one or more object tags, or both. A destination bucket can be in the same AWS account as the source bucket, or it can be in a different account.

If you specify an object version ID to delete, Amazon S3 deletes that object version in the source bucket. But it doesn't replicate the deletion in the destination bucket. In other words, it doesn't delete the same object version from the destination bucket. This protects data from malicious deletions.

If the destination bucket is in a different account from the source bucket, you must add a bucket policy to the destination bucket to grant the owner of the source bucket account permission to replicate objects in the destination bucket. The Amazon S3 console builds this required bucket policy for you to copy and add to the destination bucket in the other account.

When you add a replication rule to a bucket, the rule is enabled by default, so it starts working as soon as you save it.

Topics

- Adding a replication rule when the destination bucket is in the same AWS account (p. 84)
- Adding a replication rule when the destination bucket is in a different AWS account (p. 91)
- More info (p. 98)

Adding a replication rule when the destination bucket is in the same AWS account

Follow these steps to configure a replication rule when the destination bucket is in the same AWS account as the source bucket.

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

3. Choose **Management**, choose **Replication**, and then choose **Add rule**.

4. In the **Replication rule** wizard, under **Set source**, you have the following options for setting the replication source:

   - To replicate the whole bucket, choose **Entire bucket** ***`bucket-name`***.
   - To replicate all objects that have the same prefix (for example, all objects that have names that begin with the string `pictures`), choose **Prefix or tags**. Enter a prefix in the box, choose the prefix from the drop-down list, and then press **Enter**. If you enter a prefix that is the name of a folder, you must use **/** (forward slash) as the last character (for example, `pictures/`). For more information about prefixes, see Object Keys in the *Amazon Simple Storage Service Developer Guide*.
   - To replicate all objects with one or more object tags, enter a tag in the box, choose the tag from the drop-down list, and then press **Enter**. Enter a tag value and then press **Enter**. Repeat the procedure to add another tag. You can combine a prefix and tags. For more information about object tags, see Object Tagging in the *Amazon Simple Storage Service Developer Guide*.

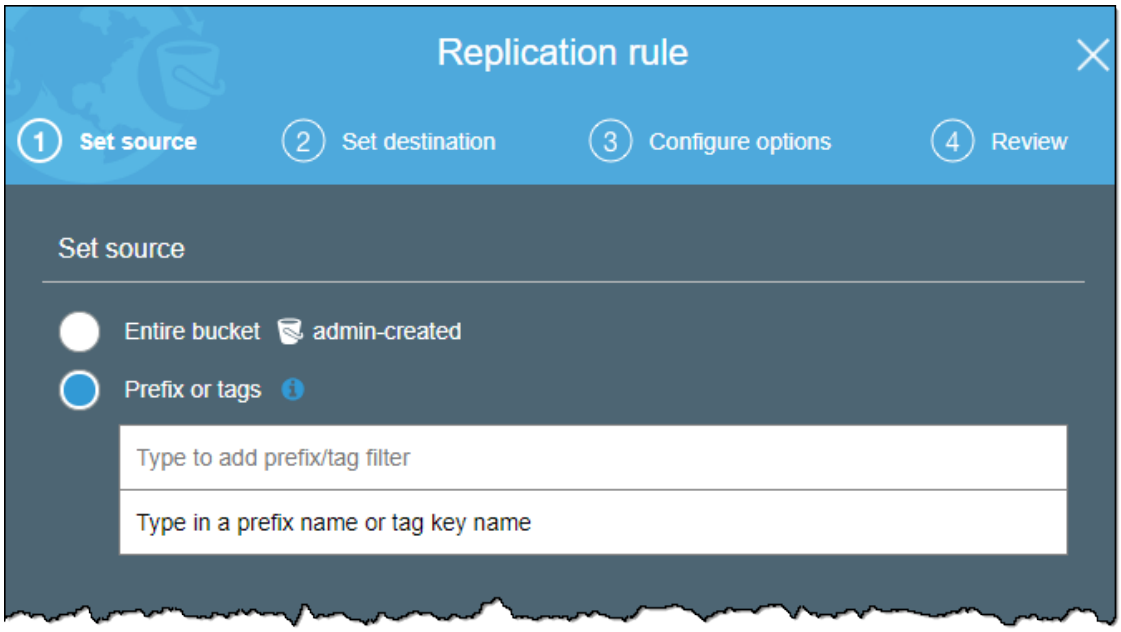

The new schema supports prefix and tag filtering and the prioritization of rules. For more information about the new schema, see Replication Configuration Backward Compatibility in the *Amazon Simple Storage Service Developer Guide*. The developer guide describes the XML used with the Amazon S3 API that works behind the user interface. In the developer guide, the new schema is described as *replication configuration XML V2*.

5. To replicate objects in the source bucket that are encrypted with AWS Key Management Service (AWS KMS), under **Replication criteria**, select **Replicate objects encrypted with AWS KMS**. Under **Choose one or more keys for decrypting source objects** are the source AWS KMS customer master keys (CMKs) that you allow replication to use. All source CMKs are included by default. You can choose to narrow the CMK selection.

   Objects encrypted by AWS KMS CMKs that you do not select are not replicated. A CMK or a group of CMKs is chosen for you, but you can choose the CMKs if you want. For information about using AWS KMS with replication, see Replicating Objects Created with Server-Side Encryption (SSE) Using Encryption Keys Stored in AWS KMS in the *Amazon Simple Storage Service Developer Guide*.

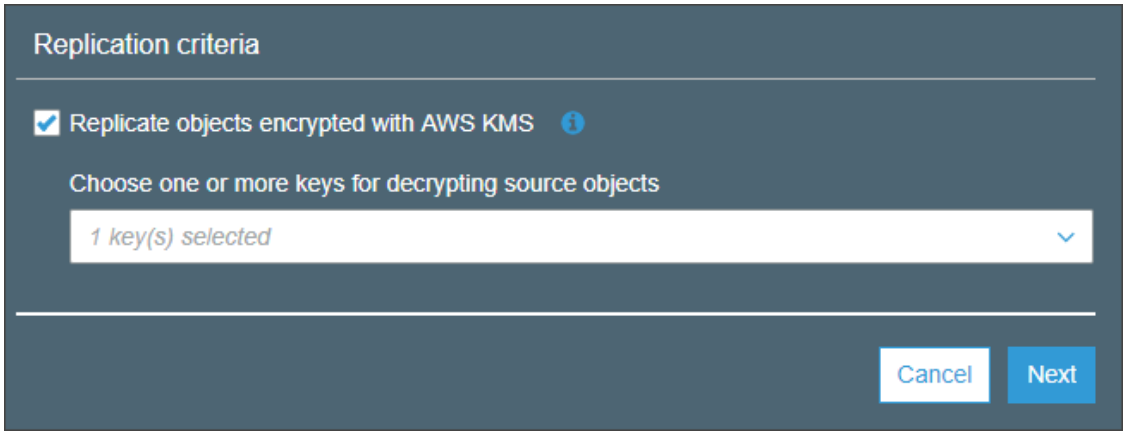

Important

When you replicate objects that are encrypted with AWS KMS, the AWS KMS request rate doubles in the source Region and increases in the destination Region by the same amount. These increased call rates to AWS KMS are due to the way that data is re-encrypted using the customer master key (CMK) that you define for the replication destination Region. AWS KMS has a request rate limit that is per calling account per Region. For information about the limit defaults, see AWS KMS Limits - Requests per Second: Varies in the *AWS Key Management Service Developer Guide*.

If your current Amazon S3 PUT object request rate during replication is more than half the default AWS KMS rate limit for your account, we recommend that you request an increase to your AWS KMS request rate limit. To request an increase, create a case in the AWS Support Center at Contact Us. For example, suppose that your current PUT object request rate is 1,000 requests per second and you use AWS KMS to encrypt your objects. In this case, we recommend that you ask AWS Support to increase your AWS KMS rate limit to 2,500 requests per second, in both your source and destination Regions (if different), to ensure that there is no throttling by AWS KMS.

To see your PUT object request rate in the source bucket, view `PutRequests` in the Amazon CloudWatch request metrics for Amazon S3. For information about viewing CloudWatch metrics, see How do I create a request metrics filter for all the objects in my S3 bucket? (p. 108)

Choose **Next**.

6. To choose a destination bucket from the account that you're currently using, on the **Set destination** page, under **Destination bucket**, choose **Buckets in this account**. Enter the name of the destination bucket for the replication, or choose a name in the drop-down list.

   If you want to choose a destination bucket from a different AWS account, see Adding a replication rule when the destination bucket is in a different AWS account (p. 91).

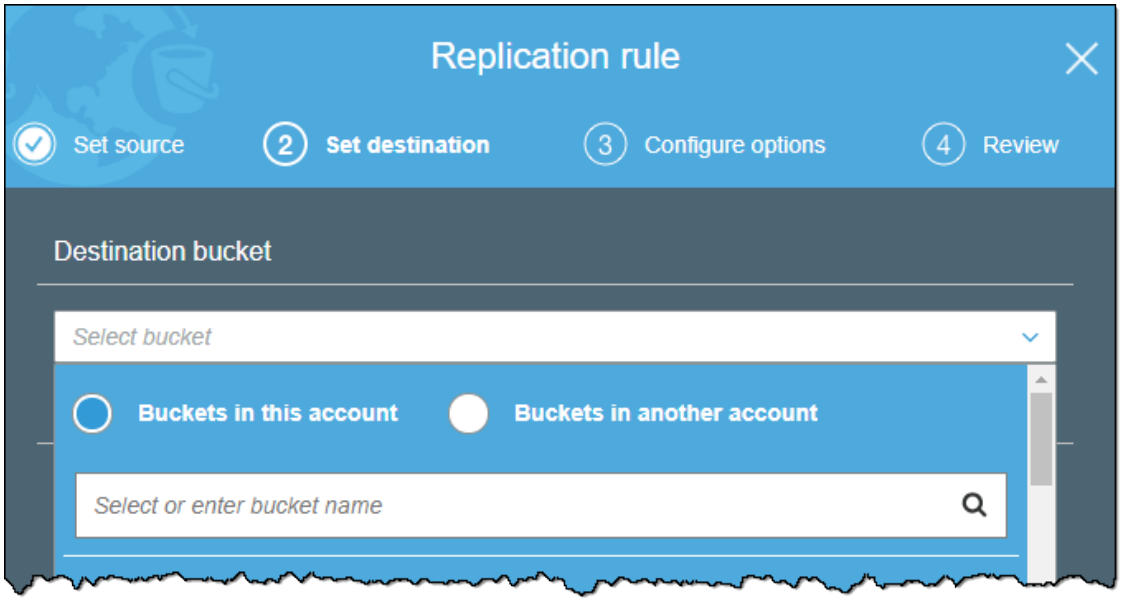

   If versioning is not enabled on the destination bucket, you get a warning message that contains an **Enable versioning** button. Choose this button to enable versioning on the bucket.

7. If you chose to replicate objects encrypted with AWS KMS, under **Destination encryption settings**, enter the Amazon Resource Name (ARN) of the AWS KMS CMK to use to encrypt the replicas in the destination bucket. You can find the ARN for your AWS KMS CMK in the IAM console, under **Encryption keys**. Or, you can choose a CMK name from the drop-down list.

For more information about creating an AWS KMS CMK, see Creating Keys in the *AWS Key Management Service Developer Guide*.

> **Important**
> The Amazon S3 console lists only 100 AWS KMS CMKs per AWS Region. If you have more than 100 CMKs in the same Region, you can see only the first 100 CMKs in the S3 console. To use a KMS CMK that is not listed in the console, choose **Custom KMS ARN**, and enter the KMS CMK ARN.

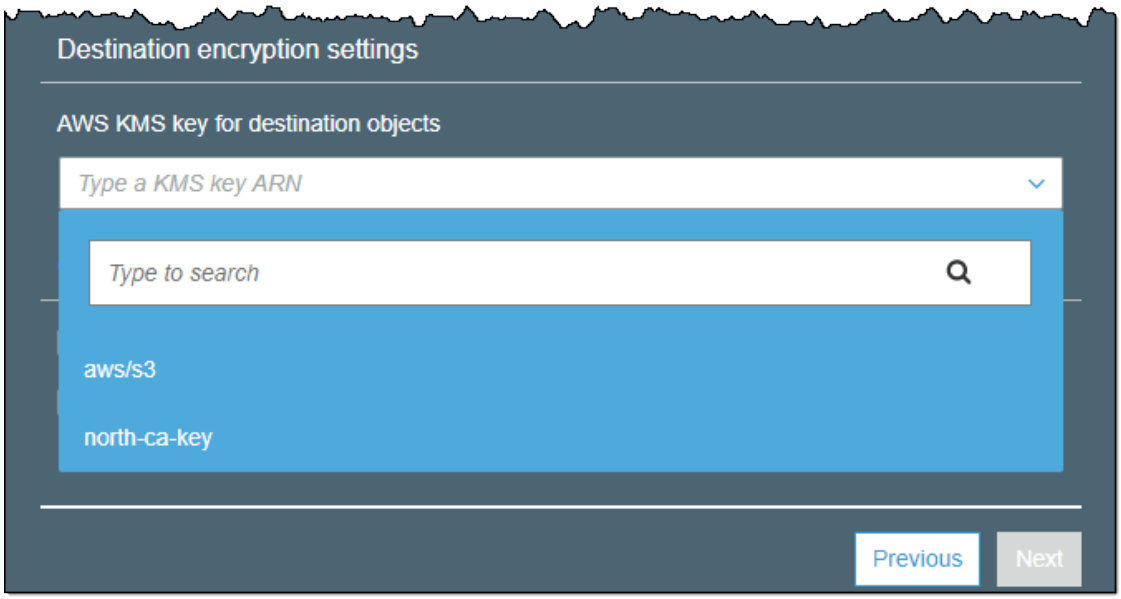

8. If you want to replicate your data into a specific storage class in the destination bucket, on the **Set destination** page, under **Destination Options**, select **Change the storage class for the replicated object(s)**. Then choose the storage class that you want to use for the replicated objects in the destination bucket. If you don't select this option, the storage class for replicated objects is the same class as the original objects.

   Similarly, if you want to change **Object Ownership** in the destination bucket, choose **Change object ownership to the destination bucket owner**. For more information about this option, see Adding a replication rule when the destination bucket is in a different AWS account (p. 91).

   If you want to enable S3 Replication Time Control (S3 RTC) in your replication configuration, select **Replication time control**.

   > **Note**
   > When you use S3 RTC, additional per-GB data transfer fees and CloudWatch metrics fees apply.

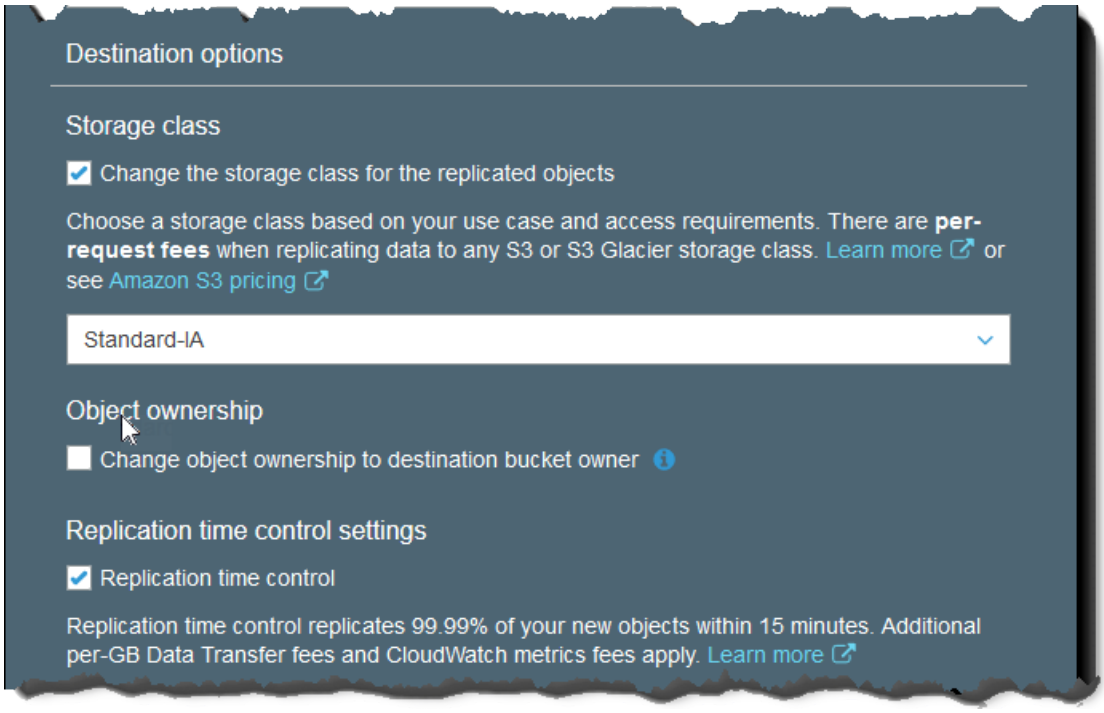

Choose **Next**.

9. Set up an AWS Identity and Access Management (IAM) role that Amazon S3 can assume to replicate objects on your behalf.

   To set up an IAM role, on the **Configure options** page, under **Select role**, do one of the following:

   - We highly recommend that you choose **Create new role** to have Amazon S3 create a new IAM role for you. When you save the rule, a new policy is generated for the IAM role that matches the source and destination buckets that you choose. The name of the generated role is based on the bucket names and uses the following naming convention: **replication_role_for_*source-bucket*_to_*destination-bucket***.
   - You can choose to use an existing IAM role. If you do, you must choose a role that grants Amazon S3 the necessary permissions for replication. Replication fails if this role does not grant Amazon S3 sufficient permissions to follow your replication rule.

     Important
     When you add a replication rule to a bucket, you must have the `iam:PassRole` permission to be able to pass the IAM role that grants Amazon S3 replication permissions. For more information, see Granting a User Permissions to Pass a Role to an AWS Service in the *IAM User Guide*.

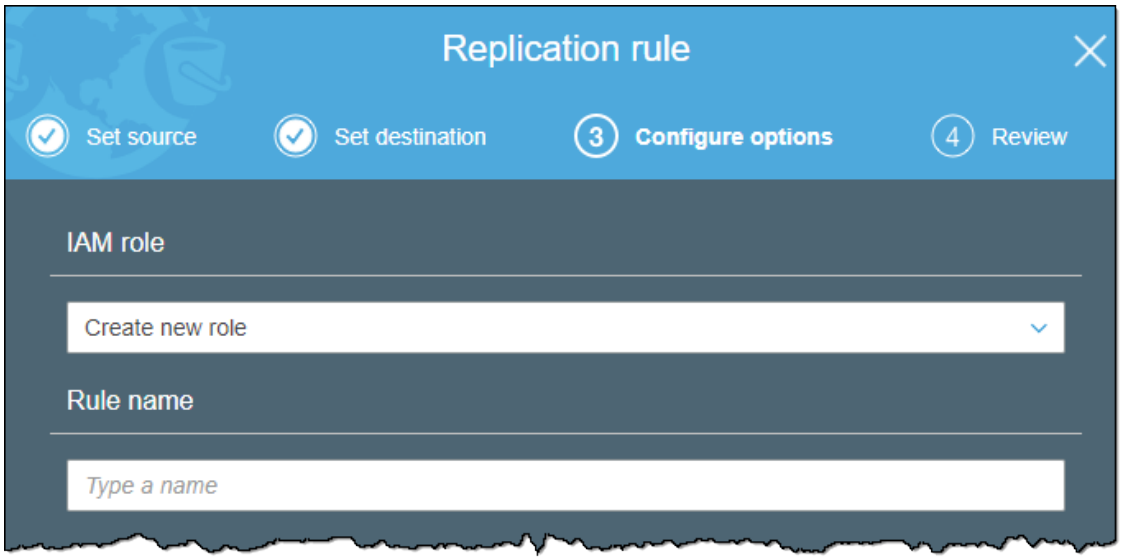

Under **Rule name**, enter a name for your rule to help identify the rule later. The name is required and must be unique within the bucket.

10. If the bucket has existing replication rules, you are instructed to set a priority for the rule. You must set a priority for the rule to avoid conflicts caused by objects that are included in the scope of more than one rule. In the case of overlapping rules, Amazon S3 uses the rule priority to determine which rule to apply. The higher the number, the higher the priority. For more information about rule priority, see Replication configuration overview in the *Amazon Simple Storage Service Developer Guide*.

    Under **Status**, **Enabled** is selected by default. An enabled rule starts to work as soon as you save it. If you want to enable the rule later, select **Disabled**.

    Choose **Next**.

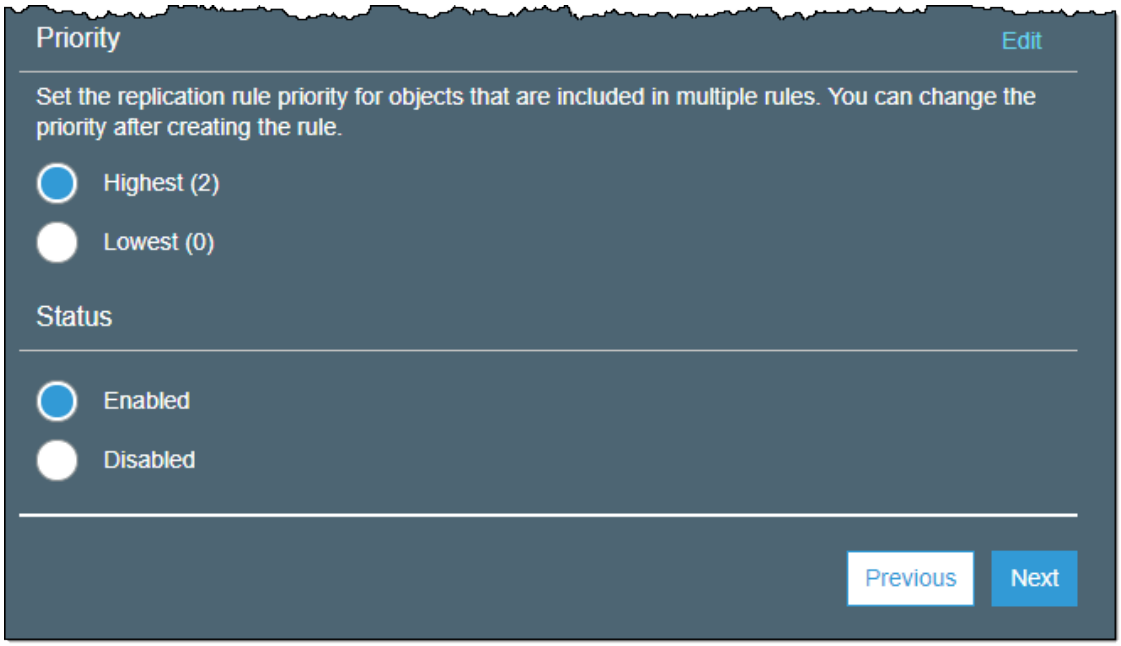

11. On the **Review** page, review your replication rule. If it looks correct, choose **Save**. Otherwise, choose **Previous** to edit the rule before saving it.
12. After you save your rule, you can edit, enable, disable, or delete your rule on the **Replication** page.

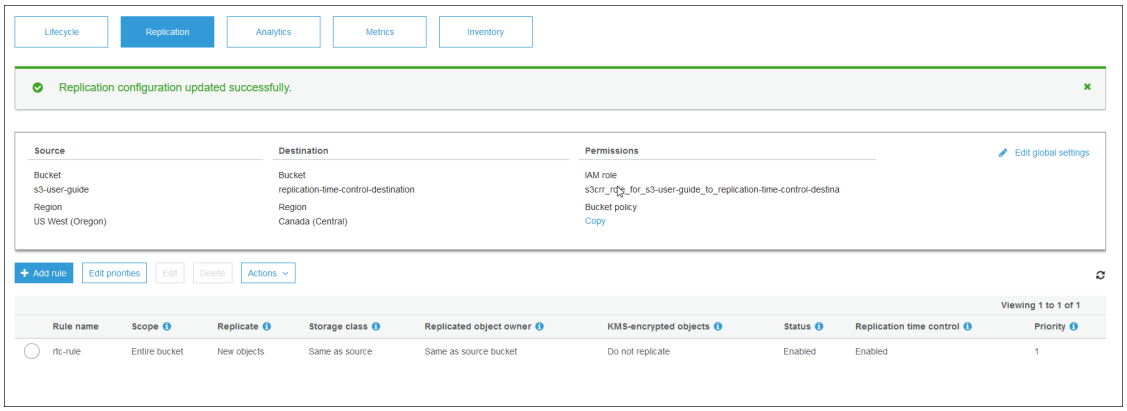

Adding a replication rule when the destination bucket is in a different AWS account

Follow these steps to configure a replication rule when the destination bucket is in a different AWS account than the source bucket.

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

3. Choose **Management**, choose **Replication**, and then choose **Add rule**.

4. If you have never created a replication rule before, start with Adding a replication rule when the destination bucket is in the same AWS account (p. 84).

   On the **Replication rule** wizard **Set destination** page, under **Destination bucket**, choose **Buckets in another account**. Then enter the name of the destination bucket and the account ID from a different AWS account. Choose **Save**.

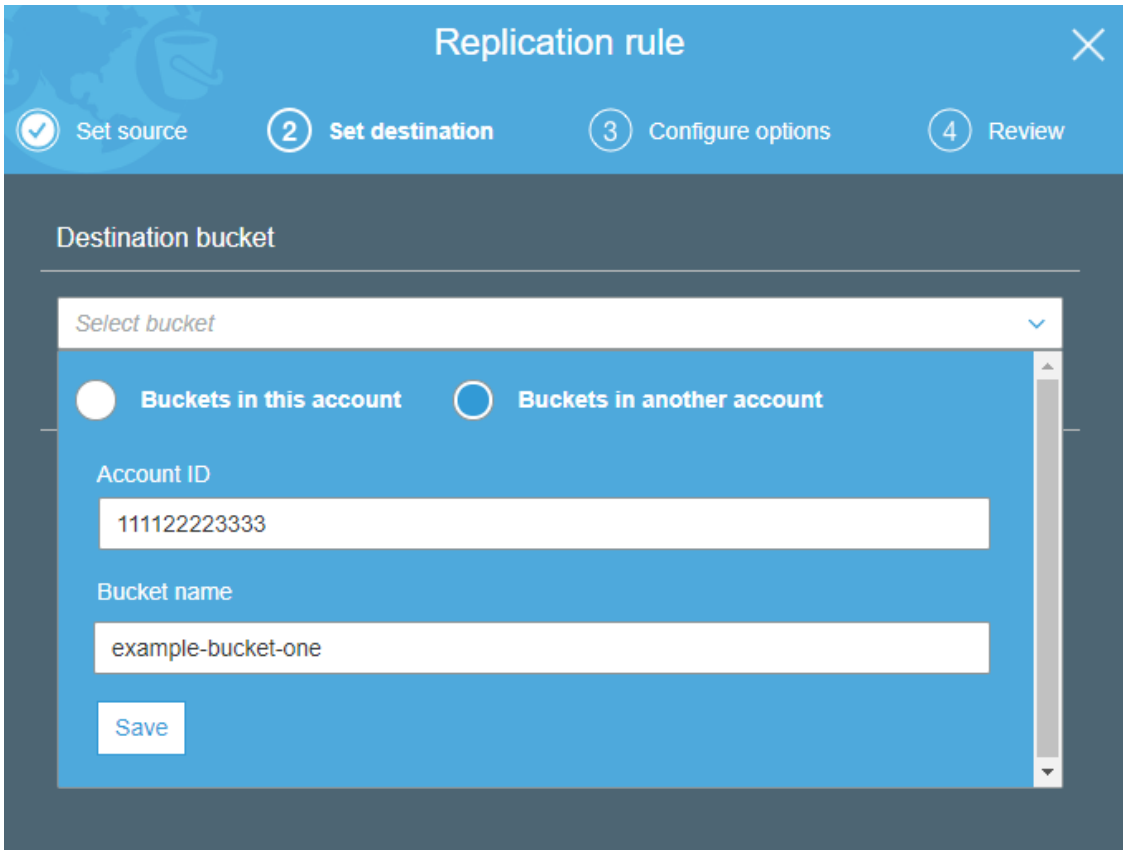

After you save the destination bucket name and account ID, you might get a warning message telling you to add a bucket policy to the destination bucket so that Amazon S3 can verify whether versioning is enabled on the bucket. You are presented with a bucket policy in a few steps that you can copy and add to the destination bucket in the other account. For information about adding a bucket policy to an S3 bucket and versioning, see How do I add an S3 Bucket policy? (p. 124) and How do I enable or suspend versioning for an S3 bucket? (p. 7)

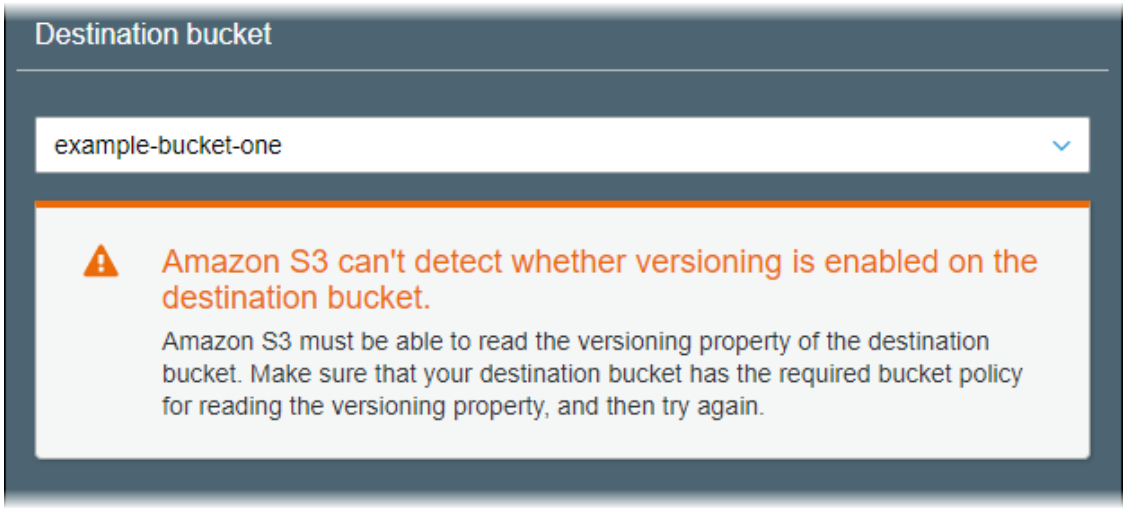

5. If you chose to replicate objects encrypted with AWS KMS, under **Destination encryption settings**, enter the Amazon Resource Name (ARN) AWS KMS CMK to use to encrypt the replicas in the destination bucket.

For more information about creating an AWS KMS CMK, see Creating Keys in the *AWS Key Management Service Developer Guide*.

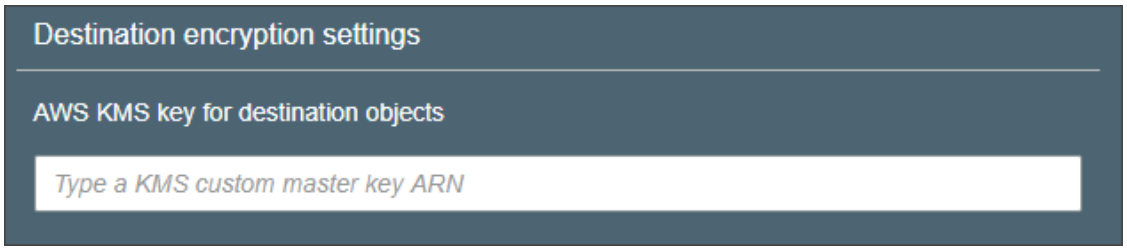

6. On the **Set destination** page, under **Destination Options**:

   - To replicate your data into a specific storage class in the destination bucket, select **Change the storage class for the replicated object(s)**. Then choose the storage class that you want to use for the replicated objects in the destination bucket. If you don't select this option, the storage class for replicated objects is the same class as the original objects.
   - To change the object ownership of the replica objects to the destination bucket owner, select **Change object ownership to destination owner**. This option enables you to separate object ownership of the replicated data from the source. If asked, type the account ID of the destination bucket.

     When you select this option, regardless of who owns the source bucket or the source object, the AWS account that owns the destination bucket is granted full permission to replica objects. For more information, see Changing the Replica Owner in the *Amazon Simple Storage Service Developer Guide*.
   - If you want to add S3 Replication Time Control (S3 RTC) to your replication configuration, select **Replication time control**.

     Note
     When you use S3 RTC, additional per-GB data transfer fees and CloudWatch metrics fees apply.

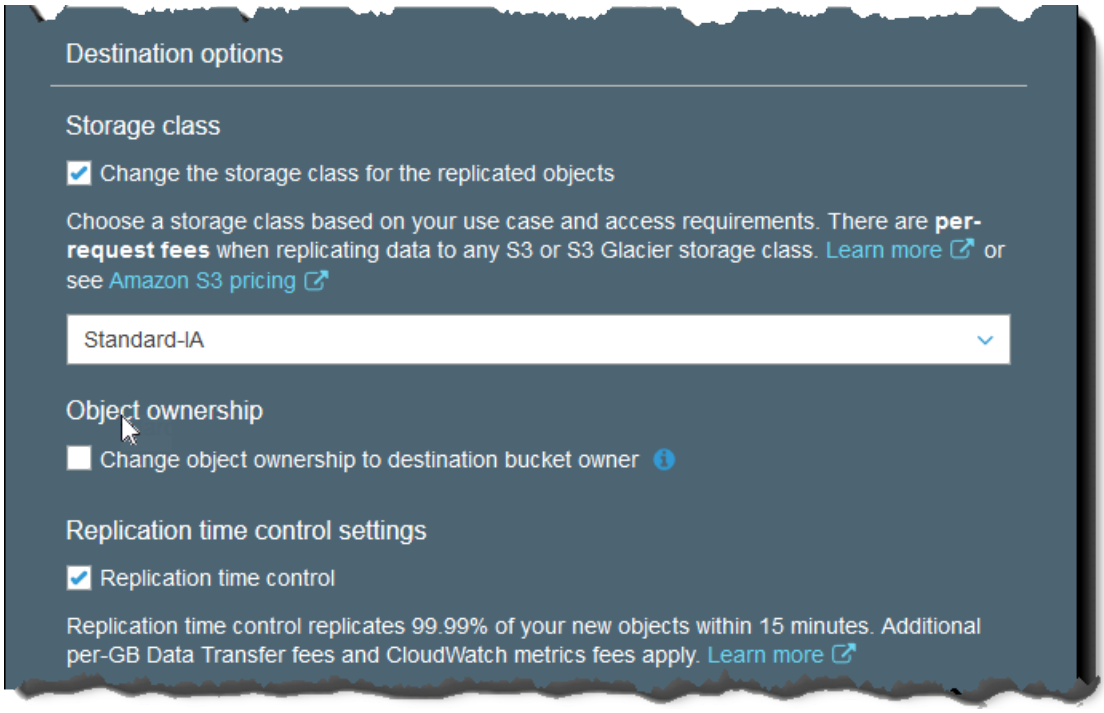

Choose **Next**.

7. Set up an AWS Identity and Access Management (IAM) role that Amazon S3 can assume to perform replication of objects on your behalf.

   To set up an IAM role, on the **Configure options** page, under **Select role**, do one of the following:

   - We highly recommend that you choose **Create new role** to have Amazon S3 create a new IAM role for you. When you save the rule, a new policy is generated for the IAM role that matches the source and destination buckets that you choose. The name of the generated role is based on the bucket names and uses the following naming convention: **replication_role_for_*source-bucket*_to_*destination-bucket***.
   - You can choose to use an existing IAM role. If you do, you must choose a role that allows Amazon S3 to replicate objects from the source bucket to the destination bucket on your behalf.

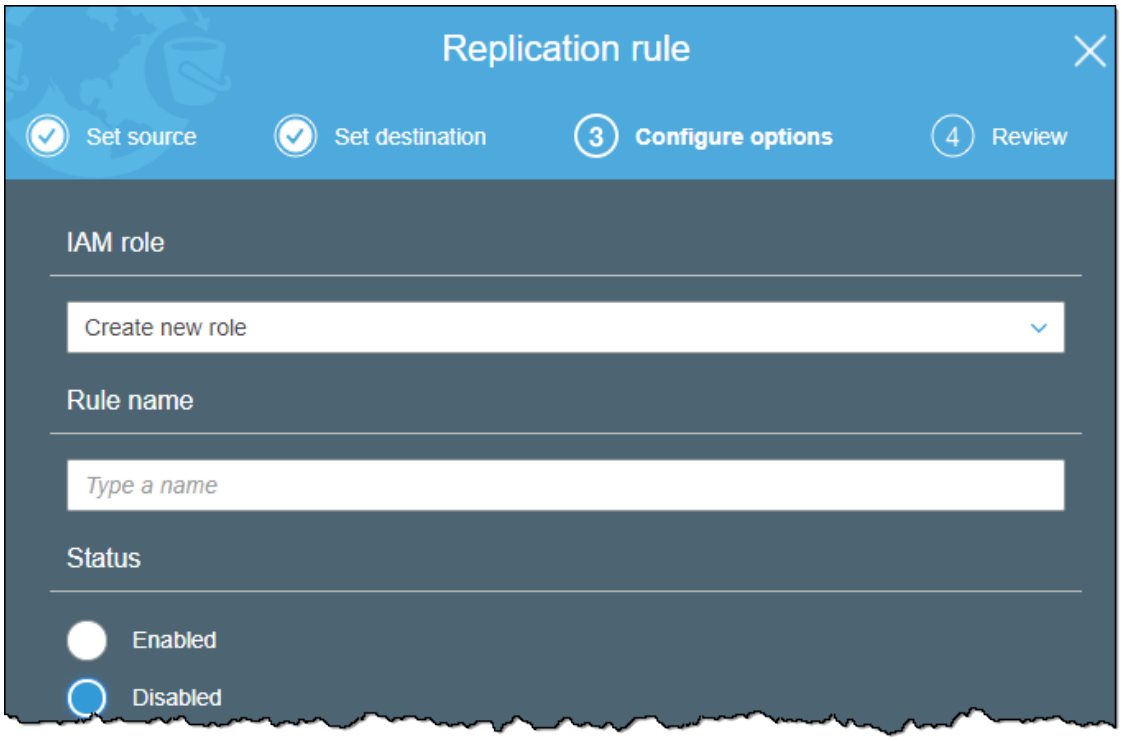

8. A bucket policy is provided on the **Configure options** page that you can copy and add to the destination bucket in the other account. For information about adding a bucket policy to an S3 bucket, see How do I add an S3 Bucket policy? (p. 124)

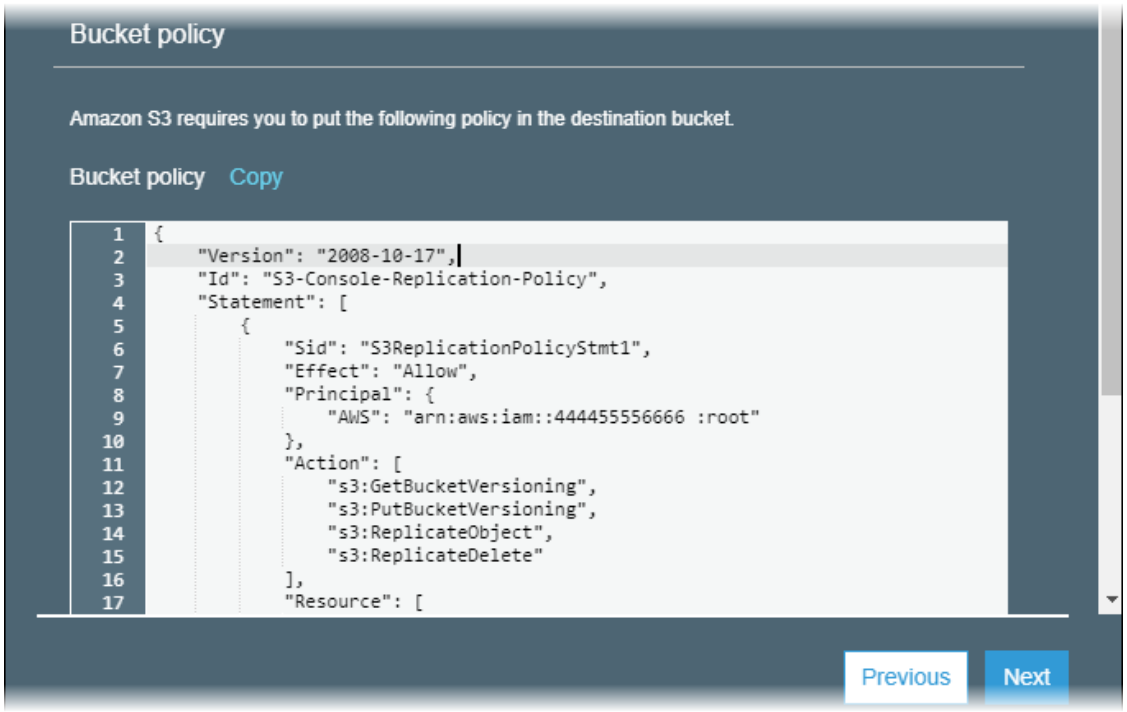

9. If you chose to replicate objects encrypted with AWS KMS, an AWS KMS key policy is provided on the **Configure options** page. You can copy this policy to add to the key policy for the AWS KMS CMK

that you are using. The key policy grants the source bucket owner permission to use the CMK. For information about updating the key policy, see Grant the source bucket owner permission to encrypt using the AWS KMS CMK (p. 97).

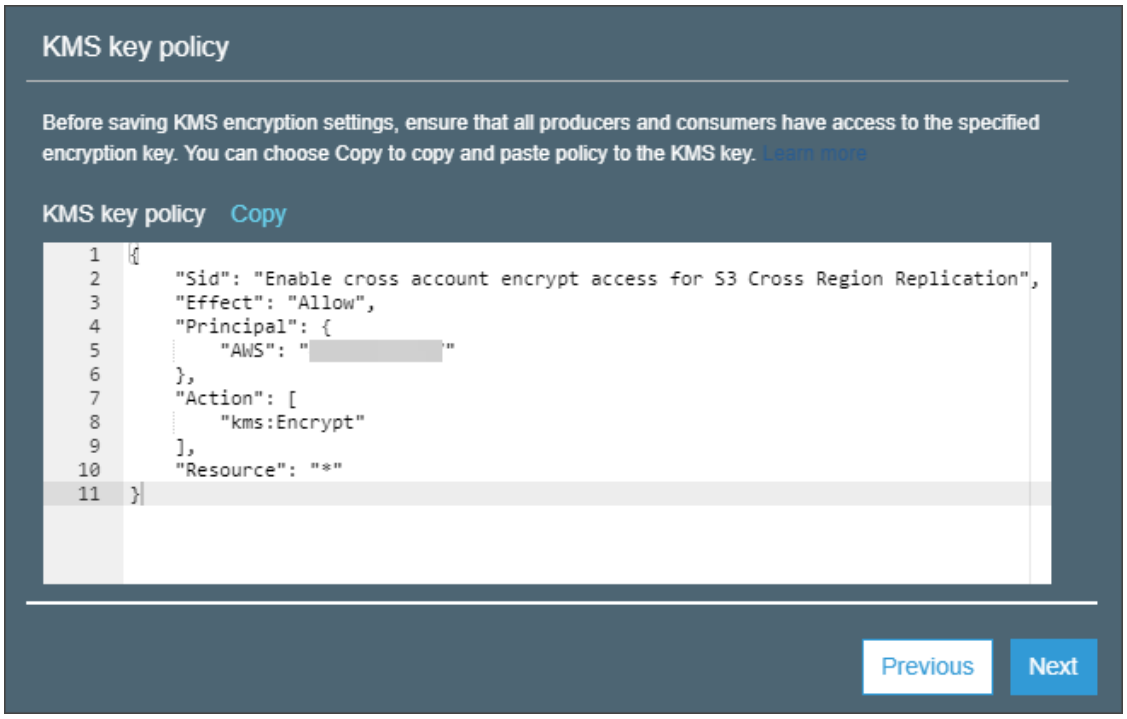

10. On the **Review** page, review your replication rule. If it looks correct, choose **Save**. Otherwise, choose **Previous** to edit the rule before saving it.
11. After you save your rule, you can edit, enable, disable, or delete your rule on the **Replication** page.

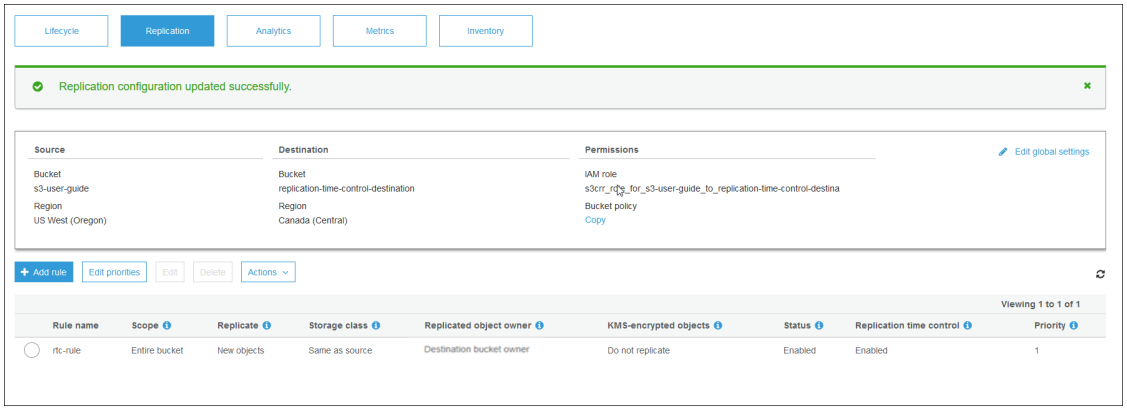

12. Follow the instructions given on the **Replication** page under the warning message, The replication rule is saved, but additional settings are required in the destination account. Sign out of the AWS account that you are currently in, and then sign in to the destination account.

    Important
    Replication fails until you sign in to the destination account and complete the following steps.

13. After you sign in to the destination account, choose the **Management** tab, choose **Replication**, and then choose **Receive objects** on the **Actions** menu.
14. On the **Receive objects** page, you can do the following:

    - Enable versioning on the destination bucket.

- Apply the bucket policy provided by Amazon S3 to the destination bucket.
- Copy the AWS KMS key policy that you need to update the AWS KMS CMK that is being used to encrypt the replica objects in the destination bucket. For information about updating the key policy, see Grant the source bucket owner permission to encrypt using the AWS KMS CMK (p. 97).

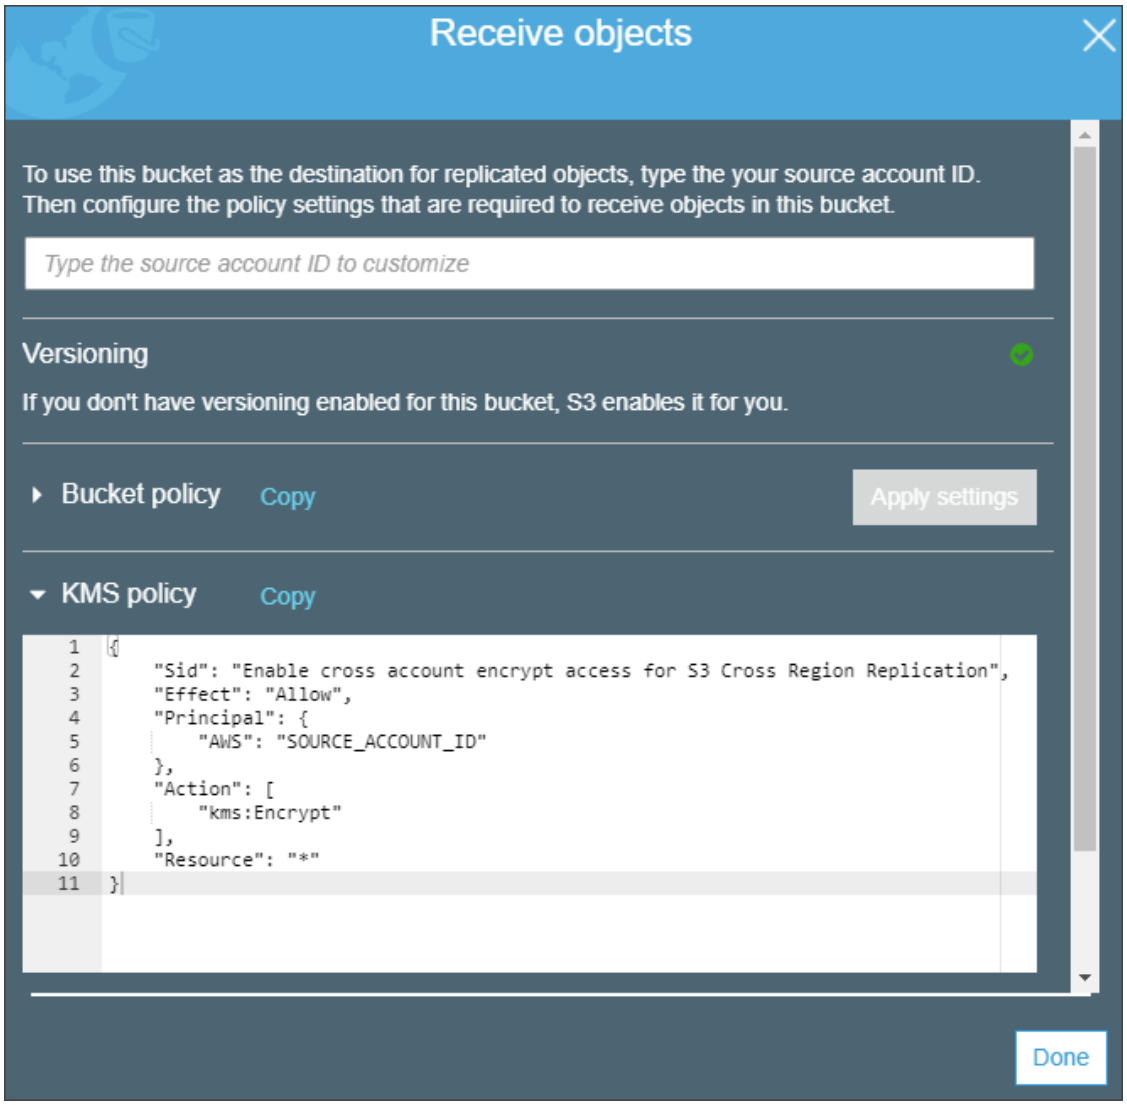

Grant the source bucket owner permission to encrypt using the AWS KMS CMK

You must grant permissions to the account of the source bucket owner to encrypt using your AWS KMS CMK with a key policy. The following procedure describes how to use the AWS Identity and Access Management (IAM) console to modify the key policy for the AWS KMS CMK that is being used to encrypt the replica objects in the destination bucket.

To grant permissions to encrypt using your AWS KMS CMK

1. Sign in to the AWS Management Console using the AWS account that owns the AWS KMS CMK. Open the AWS KMS console at https://console.aws.amazon.com/kms.
2. Choose the alias of the CMK that you want to encrypt with.
3. In the **Key Policy** section of the page, choose **Switch to policy view**.
4. Choose **Edit** to edit **Key Policy**.
5. Using the **Key Policy** editor, insert the key policy provided by Amazon S3 into the existing key policy, and then choose **Save Changes**. You might want to add the policy to the end of the existing policy.

For more information about creating and editing AWS KMS CMKs, see Getting Started in the *AWS Key Management Service Developer Guide*.

More info

- How do I manage the replication rules for an S3 Bucket? (p. 98)
- How do I enable or suspend versioning for an S3 bucket? (p. 7)
- Replication in the *Amazon Simple Storage Service Developer Guide*

How do I manage the replication rules for an S3 Bucket?

Replication is the automatic, asynchronous copying of objects across buckets in the same or different AWS Regions. It replicates newly created objects and object updates from a source bucket to a specified destination bucket.

You use the Amazon S3 console to add replication rules to the source bucket. Replication rules define the source bucket objects to replicate and the destination bucket where the replicated objects are stored. For more information about replication, see Replication in the *Amazon Simple Storage Service Developer Guide*.

You can manage replication rules on the **Replication** page. You can add, view, enable, disable, delete, and change the priority of the replication rules. For information about adding replication rules to a bucket, see How do I add a replication rule to an S3 bucket? (p. 83).

To manage the replication rules for an S3 bucket

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

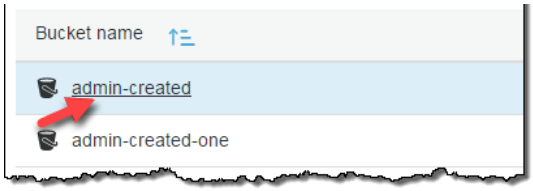

3. Choose **Management**, and then choose **Replication**.

4. You change the replication rules in the following ways.
   - To change settings that affect all the replication rules in the bucket, choose **Edit global settings**.

     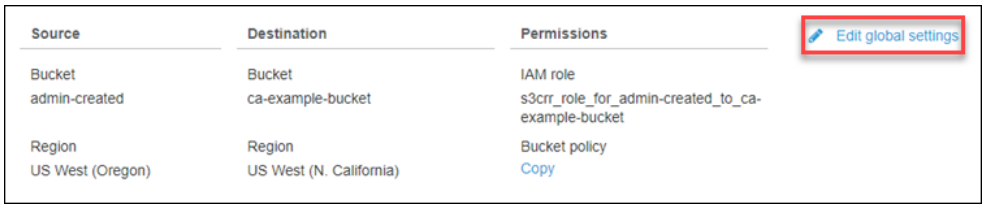

     You can change the destination bucket, and the IAM role. If needed, you can copy the required bucket policy for cross-account destination buckets.

     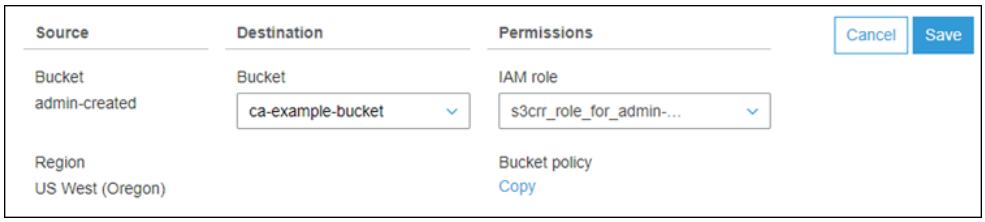

   - To change a replication rule, select the rule and choose **Edit**, which starts the Replication wizard to help you make the change. For information about using the wizard, see How do I add a replication rule to an S3 bucket? (p. 83).

     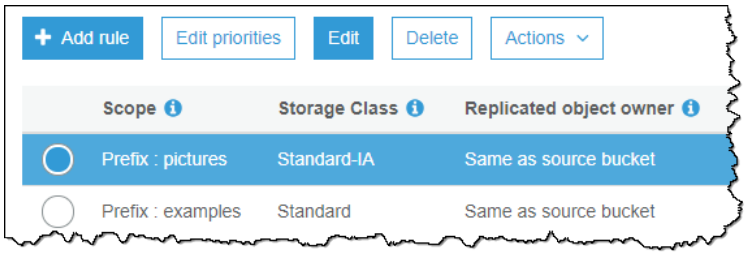

   - To enable or disable a replication rule, select the rule, choose **More**, and in the drop-down list, choose **Enable rule** or **Disable rule**. You can also disable, enable, or delete all the rules in the bucket from the **More** drop-down list.

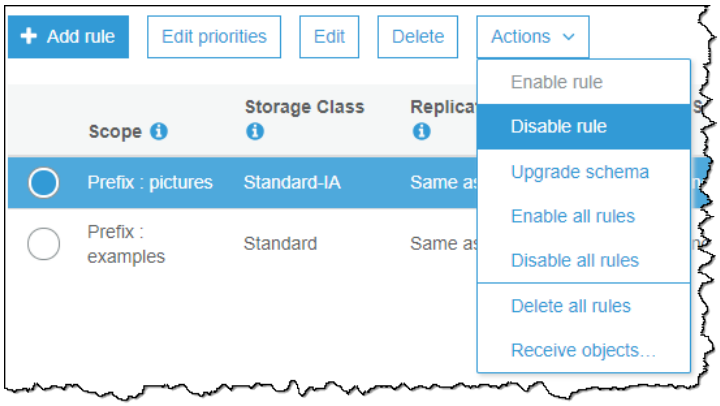

- To change the rule priorities, choose **Edit priorities**. You can then change the priority for each rule under the **Priority** column heading. Choose **Save** to save your changes.

  You set rule priorities to avoid conflicts caused by objects that are included in the scope of more than one rule. In the case of overlapping rules, Amazon S3 uses the rule priority to determine which rule to apply. The higher the number, the higher the priority. For more information about rule priority, see Replication Configuration Overview in the *Amazon Simple Storage Service Developer Guide*.

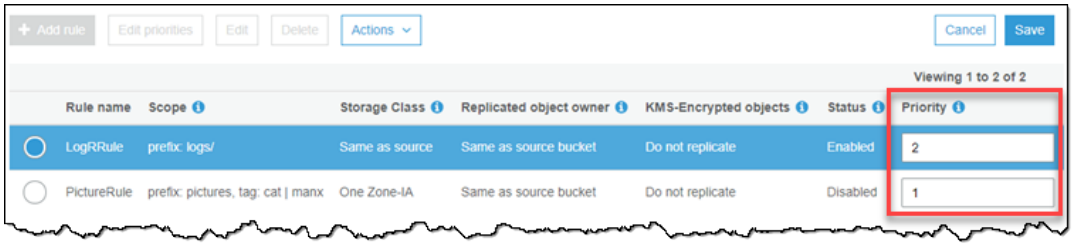

More info

- How do I add a replication rule to an S3 bucket? (p. 83)
- Replication in the *Amazon Simple Storage Service Developer Guide*

How Do I Configure Storage Class Analysis?

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

> **Important**
> Storage class analysis does not give recommendations for transitions to the ONEZONE_IA or S3 Glacier storage classes.

For more information about analytics, see Amazon S3 Analytics – Storage Class Analysis in the *Amazon Simple Storage Service Developer Guide*.

To configure storage class analysis

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

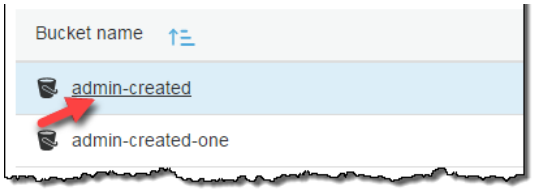

3. Choose the **Management** tab, and then choose **Analytics**.

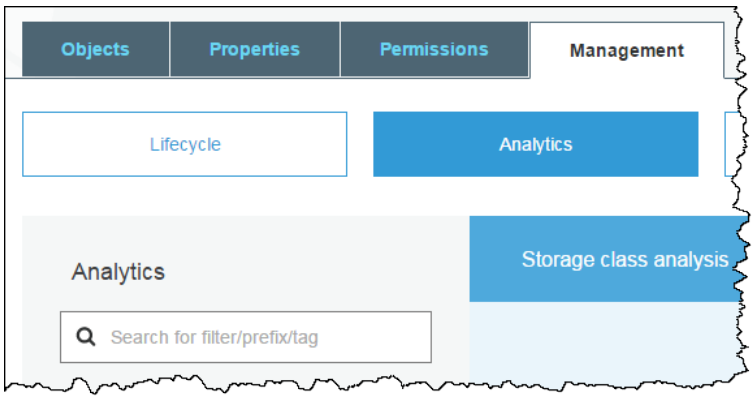

4. Choose **Add**.

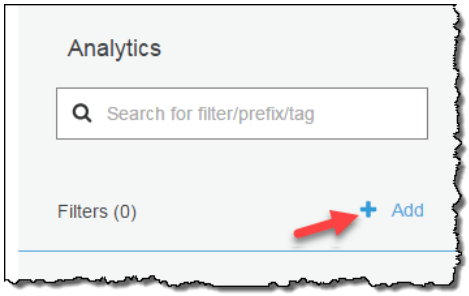

5. Type a name for the filter. If you want to analyze the whole bucket, leave the **Prefix / tags** field empty.

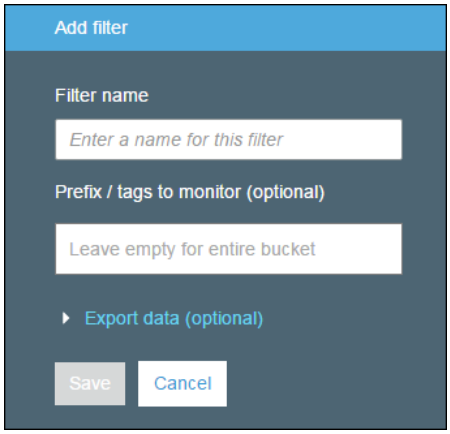

6. In the **Prefix / tags** field, type text for the prefix or tag for the objects that you want to analyze, or choose from the dropdown list that appears when you start typing.

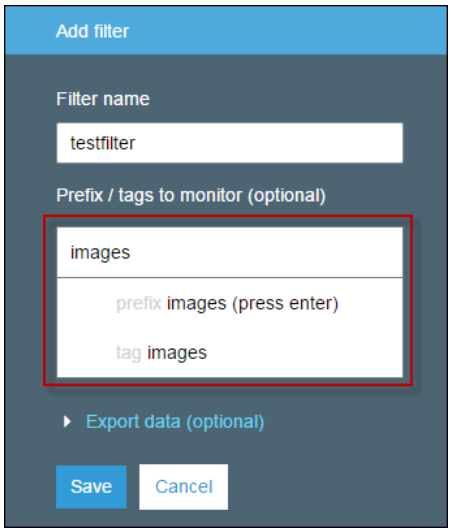

7. If you chose **tag**, enter a value for the tag. You can enter one prefix and multiple tags.

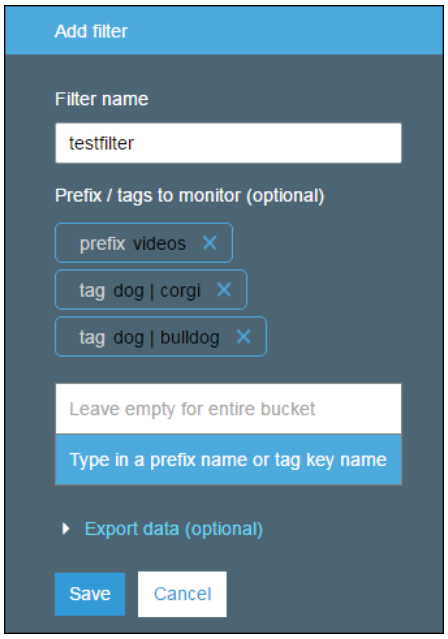

8. Optionally, you can choose **Export data** to export analysis reports to a comma-separated values (.csv) flat file. Choose a destination bucket where the file can be stored. You can type a prefix for the destination bucket. The destination bucket must be in the same AWS Region as the bucket for which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

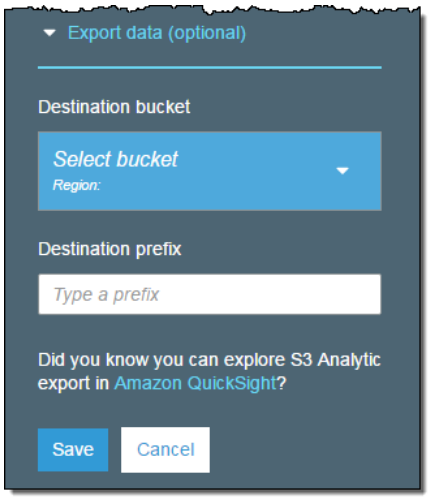

9. Choose **Save**.

   Amazon S3 creates a bucket policy on the destination bucket that grants Amazon S3 write permission. This allows it to write the export data to the bucket.

   Note
   This action configures storage class analysis for all specified buckets.

If an error occurs when you try to create the bucket policy, you'll be given instructions on how to fix it. For example, if you chose a destination bucket in another AWS account and do not have permissions to read and write to the bucket policy, you'll see the following message. You must have the destination bucket owner add the displayed bucket policy to the destination bucket. If the policy is not added to the destination bucket you won't get the export data because Amazon S3 doesn't have permission to write to the destination bucket. If the source bucket is owned by a different account than that of the current user, then the correct account ID of the source bucket must be substituted in the policy.

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

More Info

Storage management (p. 79)

How Do I Configure Amazon S3 Inventory?

Amazon S3 inventory provides a flat file list of your objects and metadata, which is a scheduled alternative to the Amazon S3 synchronous `List` API operation. Amazon S3 inventory provides comma-separated values (CSV) or Apache optimized row columnar (ORC) or Apache Parquet (Parquet) output files that list your objects and their corresponding metadata on a daily or weekly basis for an S3 bucket or for objects that share a prefix (objects that have names that begin with the same string). For more information, see Amazon S3 Inventory in the *Amazon Simple Storage Service Developer Guide*.

To configure inventory

> **Note**
> It may take up to 48 hours to deliver the first report.

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

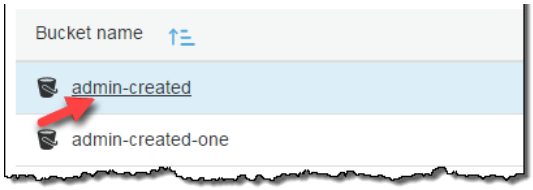

3. Choose the **Management** tab, and then choose **Inventory**.
4. Choose **Add new**.
5. Type a name for the inventory and set it up as follows:
   - Optionally, add a prefix for your filter to inventory only objects whose names begin with the same string.
   - Choose the destination bucket where you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

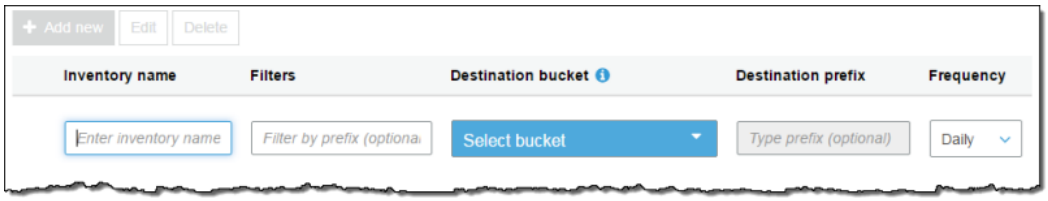

6. Under **Advanced settings**, you can set the following:
   a. Choose either the CSV, ORC, or Parquet output file format for your inventory. For more information about these formats, see Amazon S3 Inventory in the *Amazon Simple Storage Service Developer Guide*.

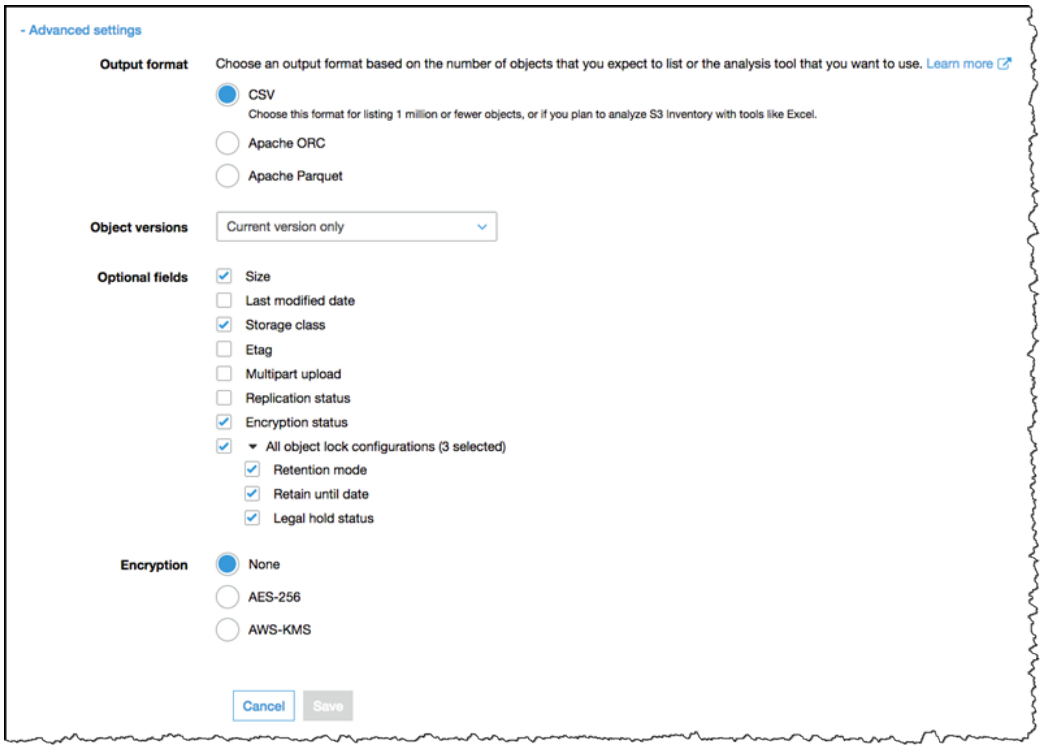

b. To include all versions of the objects in the inventory, choose **Include all versions** in the **Object versions** list. By default, the inventory includes only the current versions of the objects.

c. For **Optional fields**, select one or more of the following to add to the inventory report:

- **Size** – Object size in bytes.
- **Last modified date** – The object creation date or the last modified date, whichever is the latest.
- **Storage class** – The storage class used for storing the object.
- **ETag** – The entity tag is a hash of the object. The ETag reflects changes only to the contents of an object, and not its metadata. The ETag may or may not be an MD5 digest of the object data. Whether it is depends on how the object was created and how it is encrypted.
- **Multipart upload** – Specifies that the object was uploaded as a multipart upload. For more information, see Multipart Upload Overview in the *Amazon Simple Storage Service Developer Guide*.
- **Replication status** – The replication status of the object. For more information, see How do I add a replication rule to an S3 bucket? (p. 83).
- **Encryption status** – The server-side encryption used to encrypt the object. For more information, see Protecting Data Using Server-Side Encryption in the *Amazon Simple Storage Service Developer Guide*.
- **S3 Object Lock configurations** – The Object Lock status of the object, including the following settings:
  - **Retention mode** – The level of protection applied to the object, either *Governance* or *Compliance*.
  - **Retain until date** – The date until which the locked object cannot be deleted.
  - **Legal hold status** – The legal hold status of the locked object.

  For information about S3 Object Lock, see S3 Object Lock Overview in the *Amazon Simple Storage Service Developer Guide*.

For more information about the contents of an inventory report, see What's Included in an Amazon S3 Inventory? in the *Amazon Simple Storage Service Developer Guide*.

d. For **Encryption**, choose a server-side encryption option to encrypt the inventory report, or choose **None**:

- **None** – Do not encrypt the inventory report.
- **AES-256** – Encrypt the inventory report using server-side encryption with Amazon S3-managed keys (SSE-S3). Amazon S3 server-side encryption uses 256-bit Advanced Encryption Standard (AES-256). For more information, see Amazon S3-Managed Encryption Keys (SSE-S3) in the *Amazon Simple Storage Service Developer Guide*.
- **AWS-KMS** – Encrypt the report using server-side encryption with AWS Key Management Service (AWS KMS) customer master keys (CMKs). For more information, see AWS KMS CMKs in the *Amazon Simple Storage Service Developer Guide*.

  Note
  To encrypt the inventory list file with SSE-KMS, you must grant Amazon S3 permission to use the AWS KMS CMK. For instructions, see Grant Amazon S3 Permission to Encrypt Using Your AWS KMS CMK (p. 107) .

7. Choose **Save**.

Destination Bucket Policy

Amazon S3 creates a bucket policy on the destination bucket that grants Amazon S3 write permission. This allows Amazon S3 to write data for the inventory reports to the bucket.

If an error occurs when you try to create the bucket policy, you are given instructions on how to fix it. For example, if you choose a destination bucket in another AWS account and don't have permissions to read and write to the bucket policy, you see the following message.

In this case, the destination bucket owner must add the displayed bucket policy to the destination bucket. If the policy is not added to the destination bucket, you won't get an inventory report because Amazon S3 doesn't have permission to write to the destination bucket. If the source bucket is owned by

a different account than that of the current user, the correct account ID of the source bucket must be substituted in the policy.

For more information, see Amazon S3 Inventory in the *Amazon Simple Storage Service Developer Guide*.

Granting Amazon S3 Permission to Use Your AWS KMS CMK for Encryption

To grant Amazon S3 permission to encrypt using a customer managed AWS Key Management Service (AWS KMS) customer master key (CMK), you must use a key policy. To update your key policy so that you can use an AWS KMS customer managed CMK to encrypt the inventory file, follow the steps below.

To grant permissions to encrypt using your AWS KMS CMK

1. Using the AWS account that owns the customer managed CMK, sign into the AWS Management Console.
2. Open the AWS KMS console at https://console.aws.amazon.com/kms.
3. To change the AWS Region, use the Region selector in the upper-right corner of the page.
4. In the left navigation pane, choose **Customer managed keys**.
5. Under **Customer managed keys**, choose the customer managed CMK that you want to use to encrypt the inventory file.
6. Under **Key policy**, choose **Switch to policy view**.
7. To update the key policy, choose **Edit**.
8. Under **Edit key policy**, add the following key policy to the existing key policy.

```
{
    "Sid": "Allow Amazon S3 use of the CMK",
    "Effect": "Allow",
    "Principal": {
        "Service": "s3.amazonaws.com"
    },
    "Action": [
        "kms:GenerateDataKey"
    ],
    "Resource": "*"
}
```

9. Choose **Save changes**.

   For more information about creating customer managed CMKs AWS KMS and using key policies, see the following links in the *AWS Key Management Service Developer Guide*:

   - Getting Started
   - Using Key Policies in AWS KMS

More Info

Storage management (p. 79)

How do I create a request metrics filter for all the objects in my S3 bucket?

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

For more information about CloudWatch metrics for Amazon S3, see Monitoring metrics with Amazon CloudWatch in the *Amazon Simple Storage Service Developer Guide*.

Configuring request metrics for all the objects in an S3 bucket

To configure request metrics on a bucket

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the **Bucket name** list, choose the name of the bucket that contains the objects you want request metrics for.
3. Choose the **Management** tab, and then choose **Metrics**.
4. Choose **Requests**.
5. Choose the edit icon next to the bucket icon.
6. To create request metrics configuration for all the objects in your bucket, select the **Request metrics** check box.
7. Choose **Save**.

   After about 15 minutes, CloudWatch begins tracking these request metrics. You can see graphs for the metrics on the Amazon S3 or CloudWatch console. Request metrics are billed at the standard CloudWatch rate. For more information, see Amazon CloudWatch pricing.

   You can also create a filter that limits the scope of replication metrics by prefixes or tags. For more information, see How do I create a request metrics filter that limits scope by object tag or prefix? (p. 108)

How do I create a request metrics filter that limits scope by object tag or prefix?

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

For more information about CloudWatch metrics for Amazon S3, see Monitoring metrics with Amazon CloudWatch in the *Amazon Simple Storage Service Developer Guide*.

Creating a request metrics filter that limits scope by object tag or prefix

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

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the **Bucket name** list, choose the name of the bucket that contains the objects you want request metrics for.
3. Choose the **Management** tab, and then choose **Metrics**.
4. Choose **Requests**.
5. In the left pane, beside **Filters**, choose **Add**.
6. In **Filter name**, enter a name for your request metrics filter.
7. In **Prefix/tags that you want to monitor**, enter a prefix or tag, and choose the prefix or tag from the dropdown list that appears.
8. To add additional tags or filters, repeat the previous step.
9. Choose **Save**.

   You have now created a filter that limits the request metrics scope by object tags and prefixes. About 15 minutes after CloudWatch begins tracking these request metrics, you can see charts for the metrics on both the Amazon S3 and CloudWatch consoles. Request metrics are billed at the standard CloudWatch rate. For more information, see Amazon CloudWatch pricing.

   You can also configure request metrics at the bucket level. For information, see How do I create a request metrics filter for all the objects in my S3 bucket? (p. 108)

How do I delete a request metrics filter?

In the Amazon S3 console, you can delete a request metrics filter. When you delete a filter, you are no longer charged for request metrics that use that *specific filter*. However, you will continue to be charged for any other filter configurations that exist. When you delete a filter, you can no longer use the filter for request metrics. Deleting a filter cannot be undone.

For more information about creating a request metrics filter, see How do I create a request metrics filter for all the objects in my S3 bucket? (p. 108) and How do I create a request metrics filter that limits scope by object tag or prefix? (p. 108)

Deleting a request metrics filter

To delete a filter that limits scope with tags or 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, choose the name of the bucket that contains the objects you want request metrics for.
3. Choose the **Management** tab, and then choose **Metrics**.
4. Choose **Requests**.
5. Under **Filters**, choose the edit icon beside your filter.

   Important
   Deleting a filter cannot be undone.

6. Choose **Delete filter**.

   Amazon S3 deletes your filter.

To delete a filter for all the objects in your 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, choose the name of the bucket that contains the objects you want request metrics for.
3. Choose the **Management** tab, and then choose **Metrics**.
4. Choose **Requests**.
5. Choose the edit icon beside the bucket icon.
6. Clear **Request metrics**.

   Clearing **Request metrics** also clears **Data transfer metrics**.
7. Choose **Save**.

How do I view replication metrics?

There are three types of Amazon CloudWatch metrics for Amazon S3: storage metrics, request metrics, and replication metrics. *Replication* metrics are turned on automatically when you enable replication with S3 Replication Time Control (S3 RTC) using the AWS Management Console or the Amazon S3 API. Replication metrics are available 15 minutes after you enable a replication rule with S3 Replication Time Control (S3 RTC) (S3 RTC).

Replication metrics track the rule IDs of the replication configuration. A replication rule ID can be specific to a prefix, a tag, or a combination of both. For more information about S3 Replication Time Control (S3 RTC), see Replicating Objects Using S3 Replication Time Control (S3 RTC) in the *Amazon Simple Storage Service Developer Guide*.

For more information about CloudWatch metrics for Amazon S3, see Monitoring Metrics with Amazon CloudWatch in the *Amazon Simple Storage Service Developer Guide*.

Viewing replication metrics

Prerequisites

Enable a replication rule that has S3 RTC.

To view replication metrics

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the **Bucket name** list, choose the name of the bucket that contains the objects you want replication metrics for.
3. Choose the **Management** tab, and then choose **Metrics**.
4. Choose **Replication**.
5. In the **Rule IDs** list in the left pane, select the rule IDs that you want. If you have several rule IDs to choose from, you can search for the IDs that you want.
6. After choosing the rule IDs that you want, choose **Display graphs** below the **Rule IDs** selection box.

You can then view the replication metrics **Replication Latency (in seconds)**, **Operations pending replication**, and **Bytes pending replication** for the rules that you selected. Amazon CloudWatch begins reporting replication metrics 15 minutes after you enable S3 RTC on the respective replication rule. You can view replication metrics on the Amazon S3 or CloudWatch console. For information, see Replication metrics overview in the *Amazon Simple Storage Service Developer Guide*.

Setting bucket and object access permissions

This section explains how to use the Amazon Simple Storage Service (Amazon S3) console to grant access permissions to your buckets and objects. It also explains how to use Amazon S3 block public access to prevent the application of any settings that allow public access to data within S3 buckets.

Buckets and objects are Amazon S3 resources. You grant access permissions to your buckets and objects by using resource-based access policies. You can associate an access policy with a resource. An access policy describes who has access to resources. The resource owner is the AWS account that creates the resource. For more information about resource ownership and access policies, see Overview of Managing Access in the *Amazon Simple Storage Service Developer Guide*.

Bucket access permissions specify which users are allowed access to the objects in a bucket and which types of access they have. *Object access permissions* specify which users are allowed access to the object and which types of access they have. For example, one user might have only read permission, while another might have read and write permissions.

Bucket and object permissions are independent of each other. An object does not inherit the permissions from its bucket. For example, if you create a bucket and grant write access to a user, you can't access that user's objects unless the user explicitly grants you access. Bucket permissions generally allow a user to list information about a bucket and add and delete objects from a bucket. Object permissions generally allow a user to download, replace or delete objects.

> **Note**
> You do not necessarily need to grant bucket permissions in order to grant object permissions, and vice versa. For example, you can use the AWS console to grant a user update permissions on an object without granting that user permissions to the bucket containing that object. However, if you were to grant permissions only to the object, and not the bucket, the grantee would not be able to use the AWS console to access the object. (They would not be able to view the object in the console because they would not be able to view the bucket containing the object.) The grantee instead would have to access the object programmatically, such as with the AWS CLI.

To grant access to your buckets and objects to other AWS accounts and to the general public, you use resource-based access policies known as *access control lists* (ACLs).

A *bucket policy* is a resource-based AWS Identity and Access Management (IAM) policy that grants other AWS accounts or IAM users access to an S3 bucket. Bucket policies supplement, and in many cases, replace ACL-based access policies. For more information about using IAM with Amazon S3, see Managing Access Permissions to Your Amazon S3 Resources in the *Amazon Simple Storage Service Developer Guide*.

For more in-depth information about managing access permissions, see Introduction to Managing Access Permissions to Your Amazon S3 Resources in the *Amazon Simple Storage Service Developer Guide*.

This section also explains how to use the Amazon S3 console to add a cross-origin resource sharing (CORS) configuration to an S3 bucket. CORS allows client web applications that are loaded in one domain to interact with resources in another domain.

Topics

- How do I block public access to S3 buckets? (p. 113)
- How do I edit public access settings for S3 buckets? (p. 114)

- How do I edit public access settings for all the S3 buckets in an AWS account? (p. 117)
- How do I set permissions on an object? (p. 118)
- How do I set ACL bucket permissions? (p. 121)
- How do I add an S3 Bucket policy? (p. 124)
- How do I add cross-domain resource sharing with CORS? (p. 125)
- Setting S3 Object Ownership to bucket owner preferred in the AWS Management Console (p. 126)
- Using Access Analyzer for S3 (p. 127)

How do I block public access to S3 buckets?

Amazon S3 block public access prevents the application of any settings that allow public access to data within S3 buckets. You can configure block public access settings for an individual S3 bucket or for all the buckets in your account. For information about blocking public access using the AWS CLI, AWS SDKs, and the Amazon S3 REST APIs, see Using Amazon S3 Block Public Access in the *Amazon Simple Storage Service Developer Guide*.

The following topics explain how to use the Amazon S3 console to configure block public access settings:

- How do I edit public access settings for S3 buckets? (p. 114)
- How do I edit public access settings for all the S3 buckets in an AWS account? (p. 117)

The following sections explain viewing bucket access status and searching by access types.

Viewing access status

The list buckets view shows whether your bucket is publicly accessible. Amazon S3 labels the permissions for a bucket as follows:

- **Public** – Everyone has access to one or more of the following: List objects, Write objects, Read and write permissions.
- **Objects can be public** – The bucket is not public, but anyone with the appropriate permissions can grant public access to objects.
- **Buckets and objects not public** – The bucket and objects do not have any public access.
- **Only authorized users of this account** – Access is isolated to IAM users and roles in this account and AWS service principals because there is a policy that grants public access.

The access column shows the access status of the listed buckets.

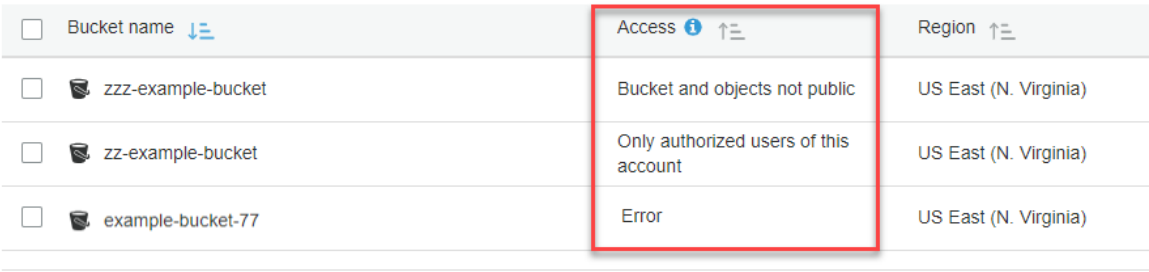

You can also filter bucket searches by access type. Choose an access type from the drop-down list that is next to the **Search for buckets** bar.

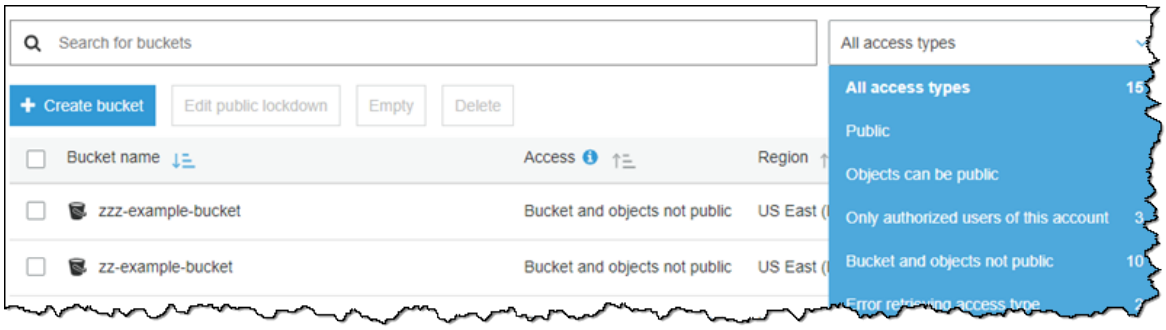

More info

- How do I edit public access settings for S3 buckets? (p. 114)
- How do I edit public access settings for all the S3 buckets in an AWS account? (p. 117)
- Setting bucket and object access permissions (p. 112)
- Restricting Access Using an Origin Access Identity in the *Amazon Simple Storage Service Developer Guide*
- Accessing Private Content in Amazon CloudFront in the *AWS Developer Blog*

How do I edit public access settings for S3 buckets?

Amazon S3 Block Public Access prevents the application of any settings that allow public access to data within S3 buckets. This section describes how to edit Block Public Access settings for one or more S3 buckets. For information about blocking public access using the AWS CLI, AWS SDKs, and the Amazon S3 REST APIs, see Using Amazon S3 Block Public Access in the *Amazon Simple Storage Service Developer Guide*.

Topics

- Editing public access settings for an S3 bucket (p. 114)
- Editing public access settings for multiple S3 buckets (p. 115)
- More info (p. 116)

Editing public access settings for an S3 bucket

Follow these steps if you need to change the public access settings for a single S3 bucket.

To edit the Amazon S3 block public access settings for an S3 bucket

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

3. Choose **Permissions**.
4. Choose **Edit** to change the public access settings for the bucket. For more information about the four Amazon S3 Block Public Access Settings, see Block Public Access Settings in the Amazon Simple Storage Service Developer Guide.

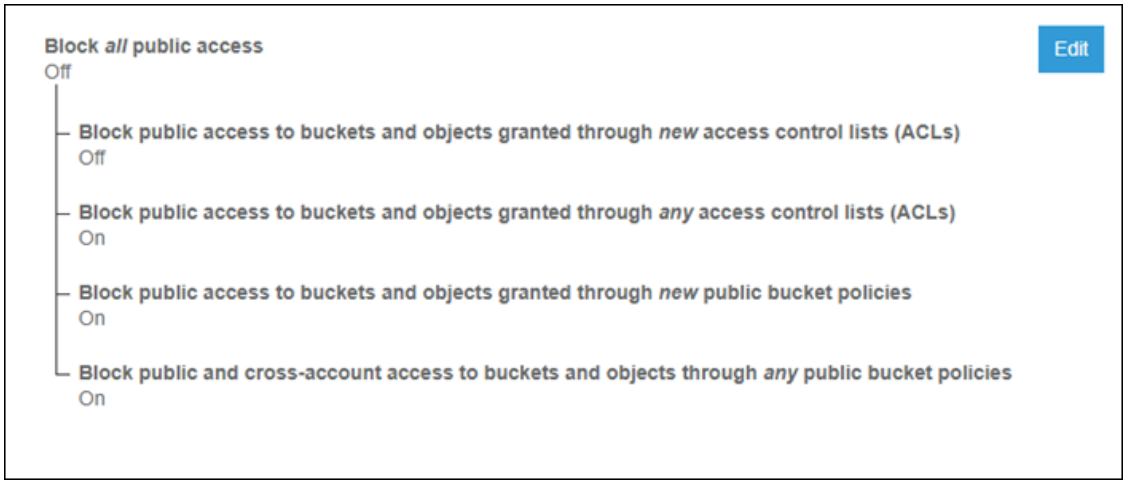

5. Choose the setting that you want to change, and then choose **Save**.
6. When you're asked for confirmation, enter **`confirm`**. Then choose **Confirm** to save your changes.

Editing public access settings for multiple S3 buckets

Follow these steps if you need to change the public access settings for more than one S3 bucket.

To edit the Amazon S3 block public access settings for multiple S3 buckets

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the **Bucket name** list, choose the buckets that you want, and then choose **Edit public access settings**.

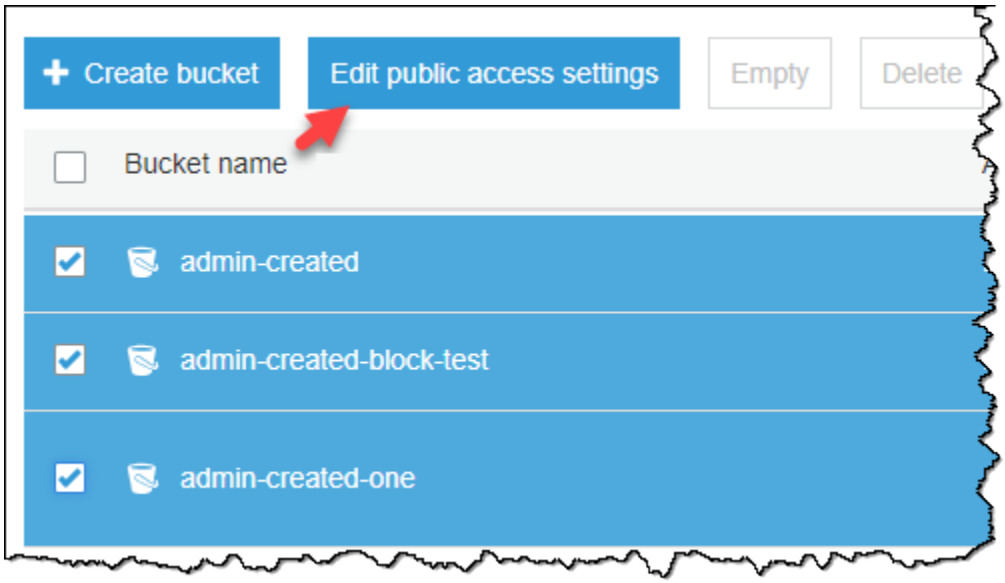

3. Choose the setting that you want to change, and then choose **Save**.

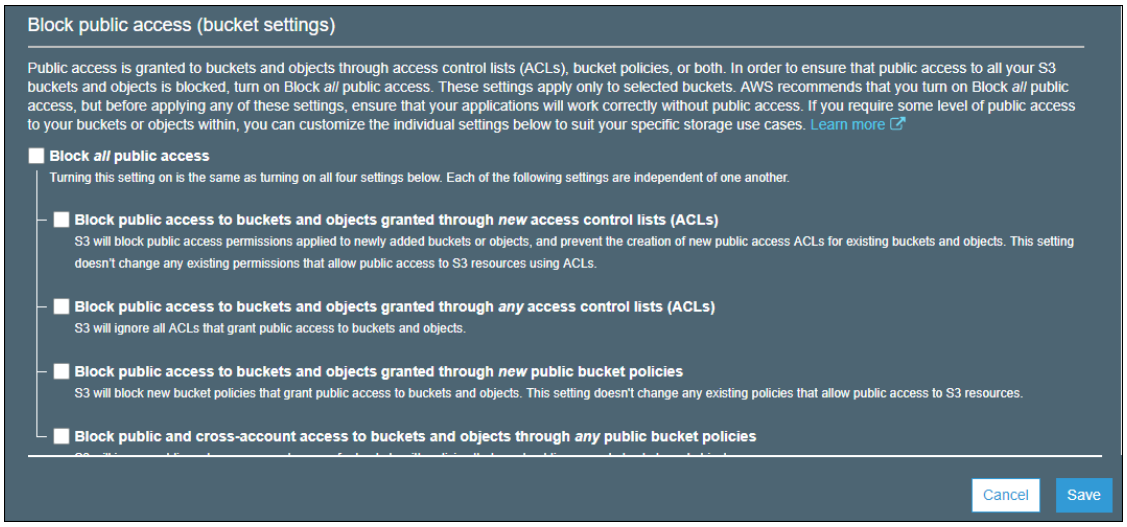

4. When you're asked for confirmation, enter **`confirm`**. Then choose **Confirm** to save your changes.

You can change Amazon S3 Block Public Access settings when you create a bucket. For more information, see How do I create an S3 Bucket? (p. 3).

More info

- How do I block public access to S3 buckets? (p. 113)
- How do I edit public access settings for all the S3 buckets in an AWS account? (p. 117)
- Setting bucket and object access permissions (p. 112)

How do I edit public access settings for all the S3 buckets in an AWS account?

Amazon S3 block public access prevents the application of any settings that allow public access to data within S3 buckets. This section describes how to edit block public access settings for all the S3 buckets in your AWS account. For information about blocking public using the AWS CLI, AWS SDKs, and the Amazon S3 REST APIs, see Using Amazon S3 Block Public Access in the *Amazon Simple Storage Service Developer Guide*.

To edit block public access settings for all the S3 buckets in an AWS account

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. Choose **Block public access (account settings)**.

3. Choose **Edit** to change the block public access settings for all the buckets in your AWS account.

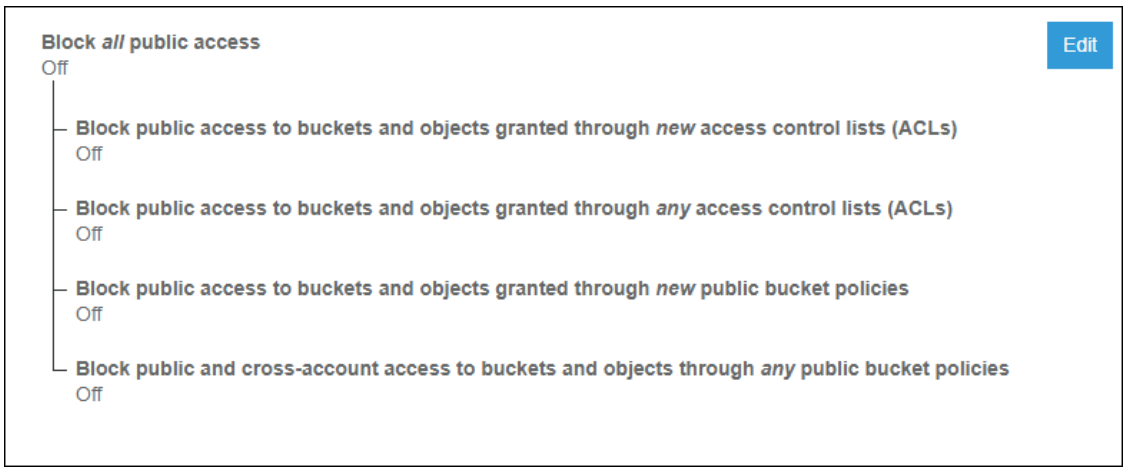

4. Choose the settings that you want to change, and then choose **Save**.
5. When you're asked for confirmation, enter **`confirm`**. Then choose **Confirm** to save your changes.

More info

- How do I block public access to S3 buckets? (p. 113)
- How do I edit public access settings for S3 buckets? (p. 114)
- Setting bucket and object access permissions (p. 112)

How do I set permissions on an object?

This section explains how to use the Amazon Simple Storage Service (Amazon S3) console to manage access permissions for an Amazon S3 object by using access control lists (ACLs). ACLs are resource-based access policies that grant access permissions to buckets and objects. For more information about managing access permissions with resource-based policies, see Overview of Managing Access in the *Amazon Simple Storage Service Developer Guide*.

Bucket and object permissions are independent of each other. An object does not inherit the permissions from its bucket. For example, if you create a bucket and grant write access to a user, you can't access that user's objects unless the user explicitly grants you access.

You can grant permissions to other AWS accounts or predefined groups. The user or group that you grant permissions to is called the *grantee*. By default, the owner, which is the AWS account that created the bucket, has full permissions.

Each permission you grant for a user or a group adds an entry in the ACL that is associated with the object. The ACL lists grants, which identify the grantee and the permission granted. For more information about ACLs, see Managing Access with ACLs in the *Amazon Simple Storage Service Developer Guide*.

To set permissions for an object

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

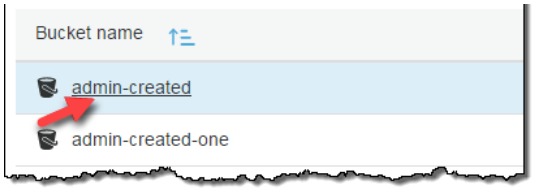

3. In the **Name** list, choose the name of the object for which you want to set permissions.

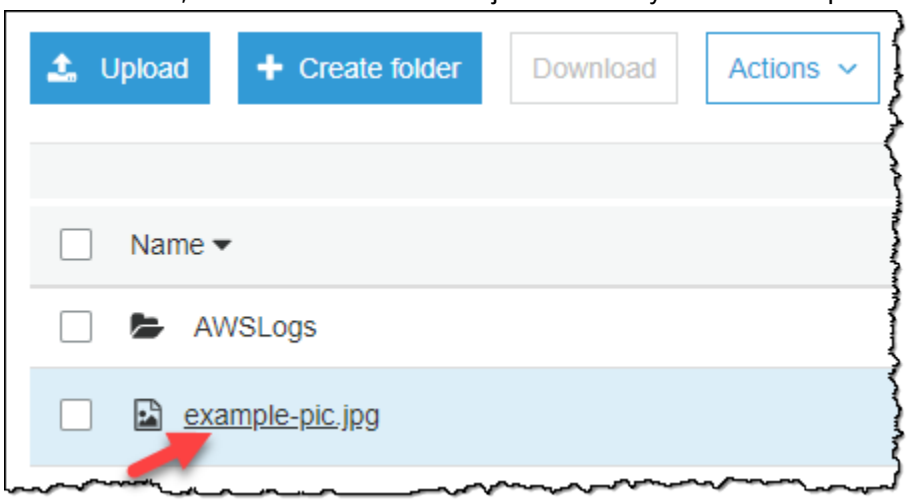

4. Choose **Permissions**.
5. You can manage object access permissions for the following:

a. **Access for object owner**

The *owner* refers to the AWS account root user, and not an AWS Identity and Access Management (IAM) user. For more information about the root user, see The AWS Account Root User in the *IAM User Guide*.

To change the owner's object access permissions, under **Access for object owner**, choose **Your AWS Account (owner)**.

Select the check boxes for the permissions that you want to change, and then choose **Save**.

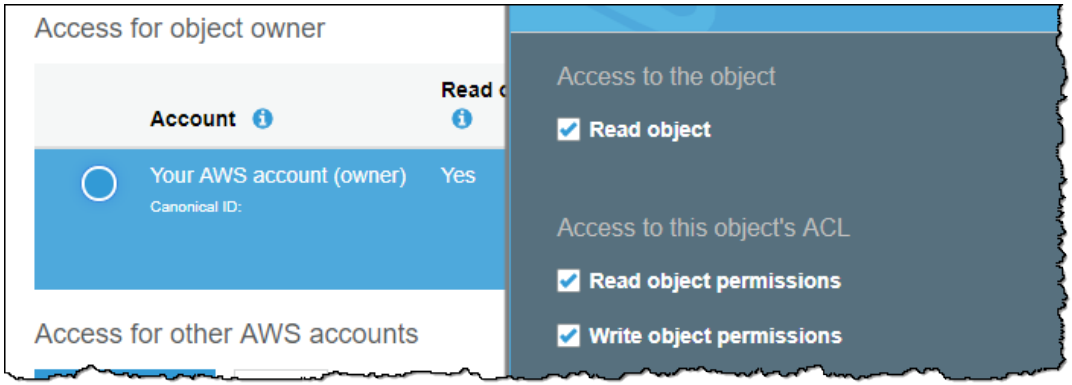

b. **Access for other AWS accounts**

To grant permissions to an AWS user from a different AWS account, under **Access for other AWS accounts**, choose **Add account**. In the **Enter an ID** field, enter the canonical ID of the AWS user that you want to grant object permissions to. For information about finding a canonical ID, see AWS Account Identifiers in the *Amazon Web Services General Reference*. You can add as many as 99 users.

Select the check boxes for the permissions that you want to grant to the user, and then choose **Save**. To display information about the permissions, choose the Help icons.

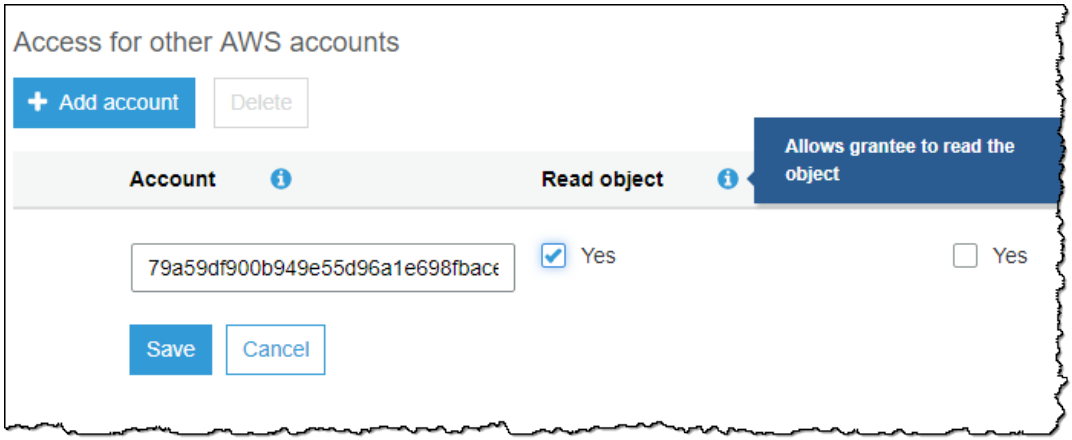

c. **Public access**

To grant access to your object to the general public (everyone in the world), under **Public access**, choose **Everyone**. Granting public access permissions means that anyone in the world can access the object.

Select the check boxes for the permissions that you want to grant, and then choose **Save**.

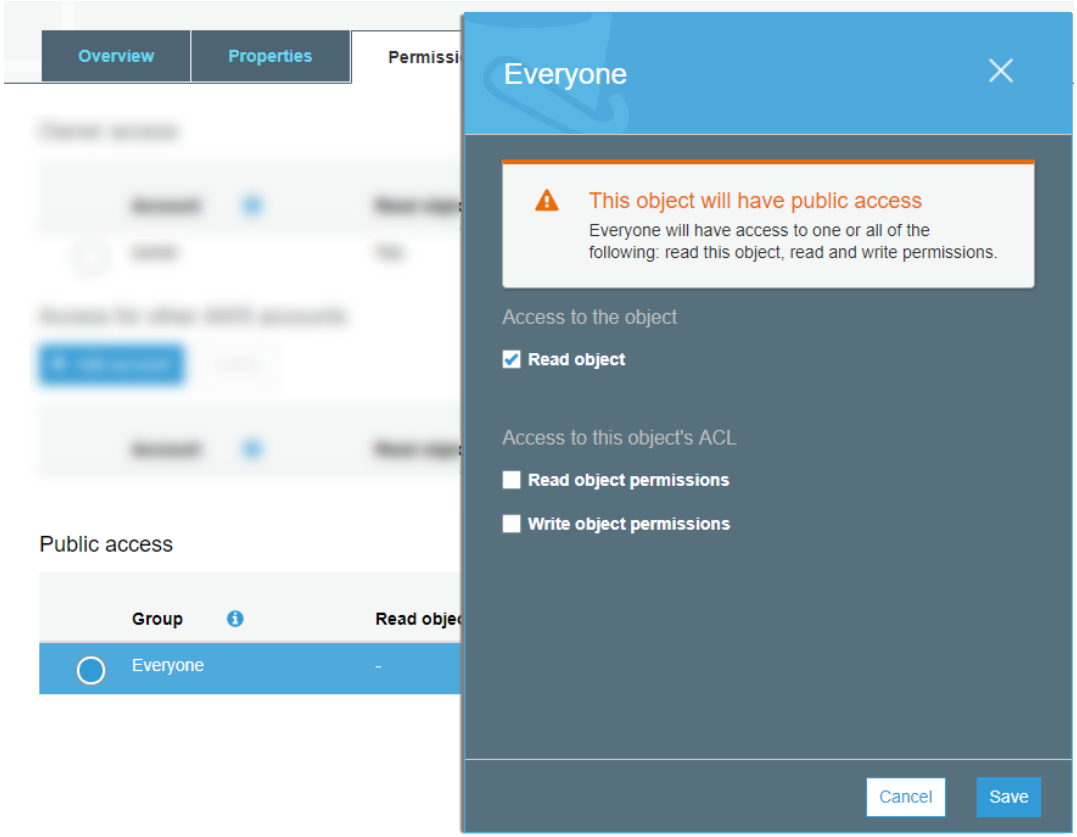

Warning

- Use caution when granting the **Everyone** group anonymous access to your Amazon S3 objects. When you grant access to this group, anyone in the world can access your object. If you need to grant access to everyone, we highly recommend that you only grant permissions to **Read objects**.
- We highly recommend that you *do not* grant the **Everyone** group write object permissions. Doing so allows anyone to overwrite the ACL permissions for the object.

Note
This action applies permissions to all specified objects. When applying permissions to folders, wait for the save operation to finish before adding new objects.

You can also set object permissions when you upload objects. For more information about setting permissions when uploading objects, see How do I upload files and folders to an S3 bucket? (p. 34).

More Info

- Setting bucket and object access permissions (p. 112)
- How do I set ACL bucket permissions? (p. 121)

How do I set ACL bucket permissions?

This section explains how to use the Amazon Simple Storage Service (Amazon S3) console to manage access permissions for S3 buckets by using access control lists (ACLs). ACLs are resource-based access policies that grant access permissions to buckets and objects. For more information about managing access permissions with resource-based policies, see Overview of Managing Access in the *Amazon Simple Storage Service Developer Guide*.

You can grant permissions to other AWS account users or to predefined groups. The user or group that you are granting permissions to is called the *grantee*. By default, the owner, which is the AWS account that created the bucket, has full permissions.

Each permission you grant for a user or group adds an entry in the ACL that is associated with the bucket. The ACL lists grants, which identify the grantee and the permission granted. For more information about ACLs, see Managing Access with ACLs in the *Amazon Simple Storage Service Developer Guide*.

> **Warning**
> We highly recommend that you avoid granting write access to the **Everyone (public access)** or **Authenticated Users group (all AWS authenticated users)** groups. For more information about the effects of granting write access to these groups, see Amazon S3 Predefined Groups in the *Amazon Simple Storage Service Developer Guide*.

To set ACL access permissions for an S3 bucket

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

3. Choose **Permissions**, and then choose **Access Control List**.
4. You can manage bucket access permissions for the following:

   a. **Access for your AWS account root user**

      The *owner* refers to the AWS account root user, and not an AWS Identity and Access Management (IAM) user. For more information about the root user, see The AWS Account Root User in the *IAM User Guide*.

      To change the owner's bucket access permissions, under **Access for your AWS account root user**, choose **Your AWS Account (owner)**.

      Select the check boxes for the permissions that you want to change, and then choose **Save**.

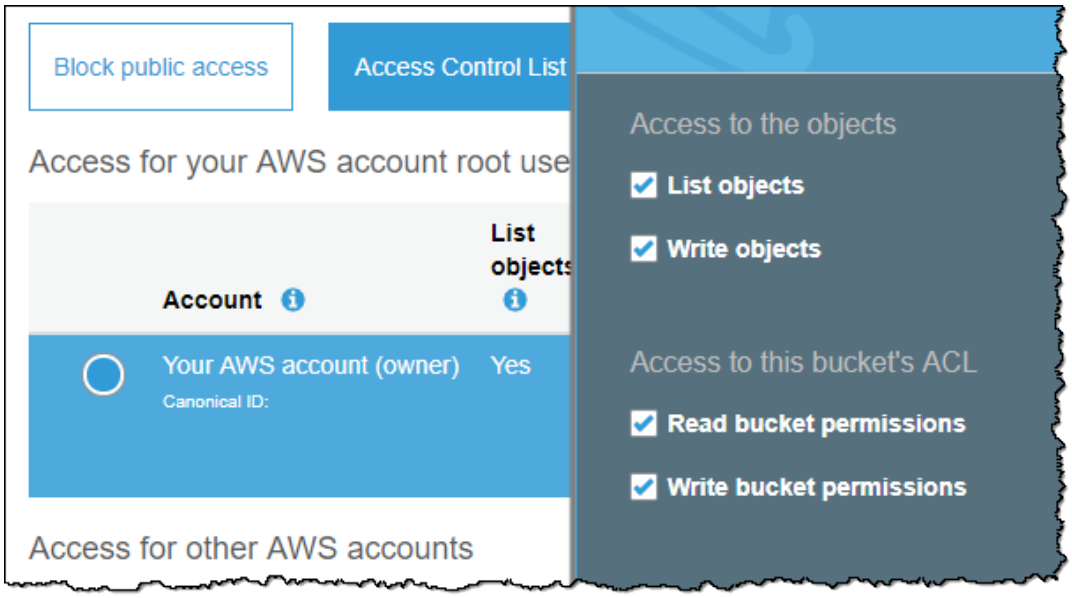

b. **Access for other AWS accounts**

To grant permissions to an AWS user from a different AWS account, under **Access for other AWS accounts**, choose **Add account**. In the **Enter an ID** field, enter the canonical ID of the AWS user that you want to grant bucket permissions to. For information about finding a canonical ID, see AWS Account Identifiers in the *Amazon Web Services General Reference*. You can add as many as 99 users.

Select the check boxes next to the permissions that you want to grant to the user, and then choose **Save**. To display information about the permissions, choose the Help icons.

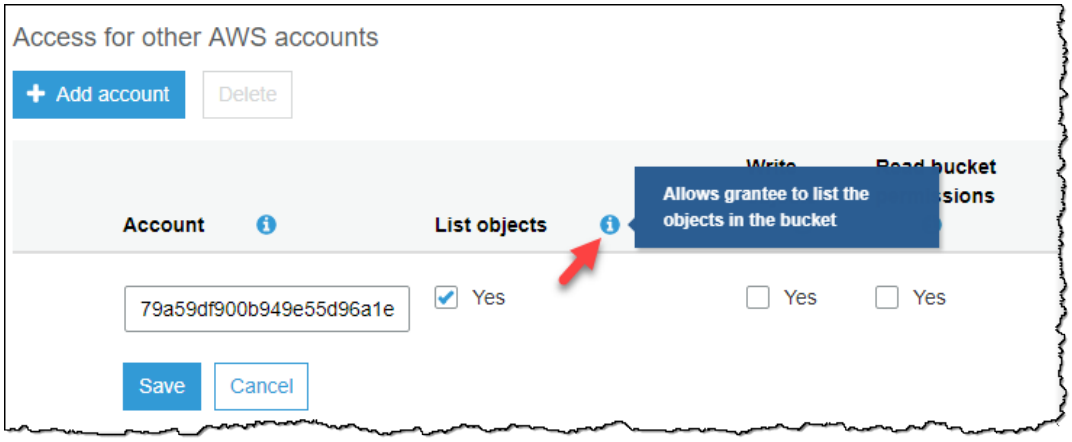

> **Warning**
> When you grant other AWS accounts access to your resources, be aware that the AWS accounts can delegate their permissions to users under their accounts. This is known as *cross-account access*. For information about using cross-account access, see Creating a Role to Delegate Permissions to an IAM User in the *IAM User Guide*.

c. **Public access**

To grant access to your bucket to the general public (everyone in the world), under **Public access**, choose **Everyone**. Granting public access permissions means that anyone in the world can access the bucket. Select the check boxes for the permissions that you want to grant, and then choose **Save**.

To undo public access to your bucket, under **Public access**, choose **Everyone**. Clear all the permissions check boxes, and then choose **Save**.

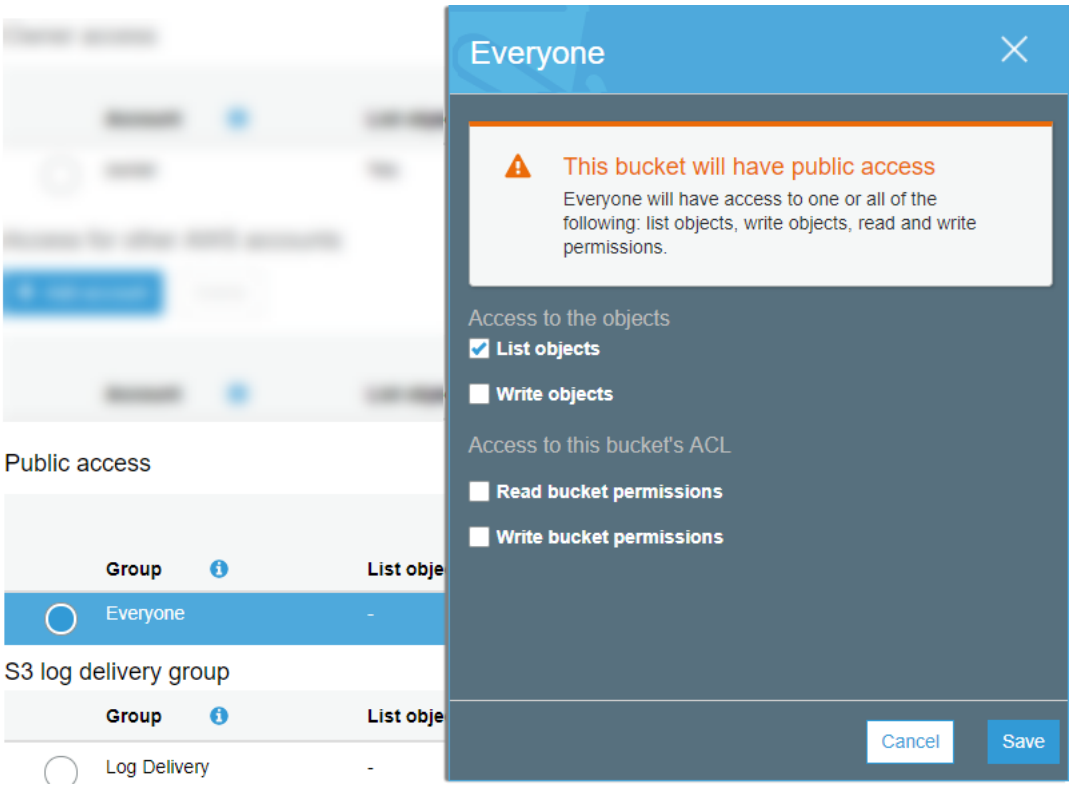

Warning
Use caution when granting the **Everyone** group public access to your S3 bucket. When you grant access to this group, anyone in the world can access your bucket. We highly recommend that you never grant any kind of public write access to your S3 bucket.

d. **S3 log delivery group**

To grant access to Amazon S3 to write server access logs to the bucket, under **S3 log delivery group**, choose **Log Delivery**.

If a bucket is set up as the target bucket to receive access logs, the bucket permissions must allow the **Log Delivery** group write access to the bucket. When you enable server access logging on a bucket, the Amazon S3 console grants write access to the **Log Delivery** group for the target bucket that you choose to receive the logs. For more information about server access logging, see How do I enable server access logging for an S3 bucket? (p. 11).

You can also set bucket permissions when you are creating a bucket. For more information about setting permissions when creating a bucket, see How do I create an S3 Bucket? (p. 3).

More info

- Setting bucket and object access permissions (p. 112)
- How do I set permissions on an object? (p. 118)
- How do I add an S3 Bucket policy? (p. 124)

How do I add an S3 Bucket policy?

This section explains how to use the Amazon Simple Storage Service (Amazon S3) console to add a new bucket policy or edit an existing bucket policy. A bucket policy is a resource-based AWS Identity and Access Management (IAM) policy. You add a bucket policy to a bucket to grant other AWS accounts or IAM users access permissions for the bucket and the objects in it. Object permissions apply only to the objects that the bucket owner creates. For more information about bucket policies, see Overview of Managing Access in the *Amazon Simple Storage Service Developer Guide*.

For examples of Amazon S3 bucket policies, see Bucket Policy Examples in the *Amazon Simple Storage Service Developer Guide*.

To create or edit a bucket policy

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

3. Choose **Permissions**, and then choose **Bucket Policy**.

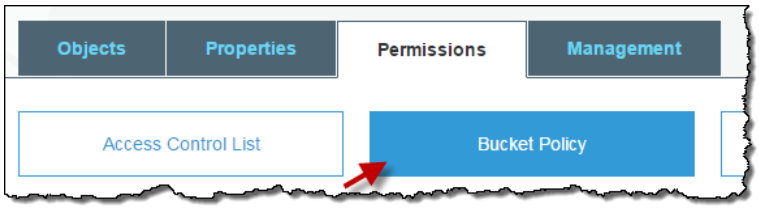

4. In the **Bucket policy editor** text box, type or copy and paste a new bucket policy, or edit an existing policy. The bucket policy is a JSON file. The text you type in the editor must be valid JSON.

5. Choose **Save**.

> **Note**
> Amazon S3 displays the Amazon Resource Name (ARN) for the bucket next to the **Bucket policy editor** title. For more information about ARNs, see Amazon Resource Names (ARNs) and AWS Service Namespaces in the *Amazon Web Services General Reference*.
> Directly below the bucket policy editor text box is a link to the **Policy Generator**, which you can use to create a bucket policy.

More info

- Setting bucket and object access permissions (p. 112)
- How do I set ACL bucket permissions? (p. 121)

How do I add cross-domain resource sharing with CORS?

This section explains how to use the Amazon S3 console to add a cross-origin resource sharing (CORS) configuration to an S3 bucket. CORS allows client web applications that are loaded in one domain to interact with resources in another domain.

To configure your bucket to allow cross-origin requests, you add CORS configuration to the bucket. A CORS configuration is an XML document that defines rules that identify the origins that you will allow to access your bucket, the operations (HTTP methods) supported for each origin, and other operation-specific information. For more information about CORS and examples, see Cross-Origin Resource Sharing (CORS) in the *Amazon Simple Storage Service Developer Guide*.

When you enable CORS on the bucket, the access control lists (ACLs) and other access permission policies continue to apply.

To add a CORS configuration to an S3 bucket

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

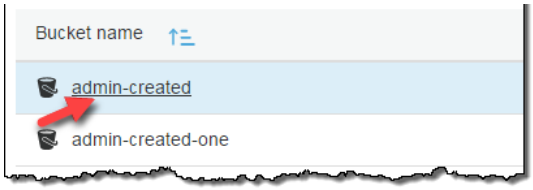

3. Choose **Permissions**, and then choose **CORS configuration**.

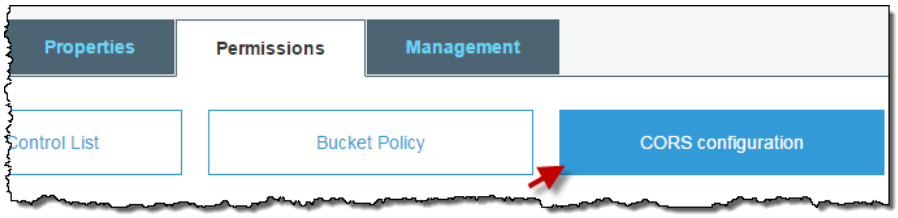

4. In the **CORS configuration editor** text box, type or copy and paste a new CORS configuration, or edit an existing configuration. The CORS configuration is an XML file. The text that you type in the editor must be valid XML. For more information, see How Do I Configure CORS on My Bucket?
5. Choose **Save**.

> **Note**
> Amazon S3 displays the Amazon Resource Name (ARN) for the bucket next to the **CORS configuration editor** title. For more information about ARNs, see Amazon Resource Names (ARNs) and AWS Service Namespaces in the *Amazon Web Services General Reference*.

More info

- Setting bucket and object access permissions (p. 112)
- How do I set ACL bucket permissions? (p. 121)
- How do I add an S3 Bucket policy? (p. 124)

Setting S3 Object Ownership to bucket owner preferred in the AWS Management Console

S3 Object Ownership is currently under preview and can be configured through the AWS Management Console, AWS Command Line Interface, AWS SDKs, or Amazon S3 REST APIs. AWS CloudFormation support is planned for general availability.

S3 Object Ownership enables you to take ownership of new objects that other AWS accounts upload to your bucket with the `bucket-owner-full-control` canned access control list (ACL). This section describes how to set Object Ownership using the console.

Setting Object Ownership to bucket owner preferred on an S3 bucket

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

2. In the **Bucket name** list, choose the name of the bucket that you want to enable S3 Object Ownership for.
3. Choose the **Permissions** tab.
4. Choose **S3 Object Ownership**.
5. Choose **Bucket owner preferred**, and then choose **Save**.

How do I ensure that I take ownership of new objects?

With the above steps Object Ownership will take ownership of any new objects that are written by other accounts with the `bucket-owner-full-control` canned ACL. For more information about enforcing Object Ownership, see How do I ensure that I take ownership of new objects? in the *Amazon Simple Storage Service Developer Guide*.

Using Access Analyzer for S3

Access Analyzer for S3 alerts you to S3 buckets that are configured to allow access to anyone on the internet or other AWS accounts, including AWS accounts outside of your organization. For each public or shared bucket, you receive findings into the source and level of public or shared access. For example, Access Analyzer for S3 might show that a bucket has read or write access provided through a bucket access control list (ACL), a bucket policy, or an access point policy. Armed with this knowledge, you can take immediate and precise corrective action to restore your bucket access to what you intended.

When reviewing an at-risk bucket in Access Analyzer for S3, you can block all public access to the bucket with a single click. We recommend that you block all access to your buckets unless you require public access to support a specific use case. Before you block all public access, ensure that your applications will continue to work correctly without public access. For more information, see Using Amazon S3 Block Public Access in the *Amazon Simple Storage Service Developer Guide*.

You can also drill down into bucket-level permission settings to configure granular levels of access. For specific and verified use cases that require public access, such as static website hosting, public downloads, or cross-account sharing, you can acknowledge and record your intent for the bucket to remain public or shared by archiving the findings for the bucket. You can revisit and modify these bucket configurations at any time. You can also download your findings as a CSV report for auditing purposes.

Access Analyzer for S3 is available at no extra cost on the Amazon S3 console. Access Analyzer for S3 is powered by AWS Identity and Access Management (IAM) Access Analyzer. To use Access Analyzer for S3 in the Amazon S3 console, you must visit the IAM console and enable IAM Access Analyzer on a per-Region basis.

For more information about IAM Access Analyzer, see What is Access Analyzer? in the *IAM User Guide*. For more information about Access Analyzer for S3, review the following sections.

Important

- Access Analyzer for S3 requires an account-level analyzer. To use Access Analyzer for S3, you must visit IAM Access Analyzer and create an analyzer that has an account as the zone of trust. For more information, see Enabling Access Analyzer in *IAM User Guide*.
- When a bucket policy or bucket ACL is added or modified, Access Analyzer generates and updates findings based on the change within 30 minutes. Findings related to account level block public access settings may not be generated or updated for up to 6 hours after you change the settings.

Topics

- What information does Access Analyzer for S3 provide? (p. 128)
- Enabling Access Analyzer for S3 (p. 128)
- Blocking all public access (p. 129)
- Reviewing and changing bucket access (p. 129)
- Archiving bucket findings (p. 130)
- Activating an archived bucket finding (p. 131)
- Viewing finding details (p. 131)
- Downloading an Access Analyzer for S3 report (p. 131)

What information does Access Analyzer for S3 provide?

Access Analyzer for S3 provides findings for buckets that can be accessed outside your AWS account. Buckets that are listed under **Buckets with public access** can be accessed by anyone on the internet. If Access Analyzer for S3 identifies public buckets, you also see a warning at the top of the page that shows you the number of public buckets in your Region. Buckets listed under **Buckets with access from other AWS accounts — including third-party AWS accounts** are shared conditionally with other AWS accounts, including accounts outside of your organization.

For each bucket, Access Analyzer for S3 provides the following information:

- **Bucket name**
- **Discovered by Access analyzer** - When Access Analyzer for S3 discovered the public or shared bucket access.
- **Shared through** - How the bucket is shared—through a bucket policy, a bucket ACL, or an access point policy. A bucket can be shared through both policies and ACLs. If you want to find and review the source for your bucket access, you can use the information in this column as a starting point for taking immediate and precise corrective action.
- **Status** - The status of the bucket finding. Access Analyzer for S3 displays findings for all public and shared buckets.
  - **Active** - Finding has not been reviewed.
  - **Archived** - Finding has been reviewed and confirmed as intended.
  - **All** - All findings for buckets that are public or shared with other AWS accounts, including AWS accounts outside of your organization.
- **Access level** - Access permissions granted for the bucket:
  - **List** - List resources.
  - **Read** - Read but not edit resource contents and attributes.
  - **Write** - Create, delete, or modify resources.
  - **Permissions** - Grant or modify resource permissions.
  - **Tagging** - Update tags associated with the resource.

Enabling Access Analyzer for S3

To use Access Analyzer for S3, you must complete the following prerequisite steps.

1. Grant the required permissions.

For more information, see Permissions Required to use Access Analyzer in the *IAM User Guide*.

2. Visit IAM to create an account-level analyzer for each Region where you want to use Access Analyzer.

   Access Analyzer for S3 requires an account-level analyzer. To use Access Analyzer for S3, you must create an analyzer that has an account as the zone of trust. For more information, see Enabling Access Analyzer in *IAM User Guide*.

Blocking all public access

If you want to block all access to a bucket in a single click, you can use the **Block all public access** button in Access Analyzer for S3. When you block all public access to a bucket, no public access is granted. We recommend that you block all public access to your buckets unless you require public access to support a specific and verified use case. Before you block all public access, ensure that your applications will continue to work correctly without public access.

If you don't want to block all public access to your bucket, you can edit your block public access settings on the Amazon S3 console to configure granular levels of access to your buckets. For more information, see Using Amazon S3 Block Public Access in the *Amazon Simple Storage Service Developer Guide*.

In rare events, Access Analyzer for S3 might report no findings for a bucket that an Amazon S3 block public access evaluation reports as public. This happens because Amazon S3 block public access reviews policies for current actions and any potential actions that might be added in the future, leading to a bucket becoming public. On the other hand, Access Analyzer for S3 only analyzes the current actions specified for the Amazon S3 service in the evaluation of access status.

To block all public access to a bucket using Access Analyzer for S3

1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the navigation pane on the left, under **Dashboards**, choose **Access analyzer for S3**.
3. In Access Analyzer for S3, choose a bucket.
4. Choose **Block all public access**.
5. To confirm your intent to block all public access to the bucket, in **Block all public access (bucket settings)**, enter **`confirm`**.

   Amazon S3 blocks all public access to your bucket. The status of the bucket finding updates to **resolved**, and the bucket disappears from the Access Analyzer for S3 listing. If you want to review resolved buckets, open IAM Access Analyzer on the IAM console.

Reviewing and changing bucket access

If you did not intend to grant access to the public or other AWS accounts, including accounts outside of your organization, you can modify the bucket ACL, bucket policy, or access point policy to remove the access to the bucket. The **Shared through** column shows all sources of bucket access: bucket policy, bucket ACL, and/or access point policy.

To review and change a bucket policy, a bucket ACL, or an access point policy

1. Open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the navigation pane, choose **Access analyzer for S3**.
3. To see whether public access or shared access is granted through a bucket policy, a bucket ACL, or an access point policy, look in the **Shared through** column.

4. Under **Bucket name**, choose the name of the bucket with the bucket policy, bucket ACL, or access point policy that you want to change or review.
5. If you want to change or view a bucket ACL:
   a. Choose **Permissions**.
   b. Choose **Access Control List**.
   c. Review your bucket ACL, and make changes as required.

      For more information, see How do I set ACL bucket permissions? (p. 121)
6. If you want to change or review a bucket policy:
   a. Choose **Permissions**.
   b. Choose **Bucket Policy**.
   c. Review or change your bucket policy as required.

      For more information, see How do I add an S3 Bucket policy? (p. 124)
7. If you want to review or change an access point policy:
   a. Choose **Access points**.
   b. Choose the access point name.
   c. Review or change access as required.

      For more information, see Managing and using Amazon S3 access points (p. 31).

   If you edit or remove a bucket ACL, a bucket policy, or an access point policy to remove public or shared access, the status for the bucket findings updates to resolved. The resolved bucket findings disappear from the Access Analyzer for S3 listing, but you can view them in IAM Access Analyzer.

Archiving bucket findings

If a bucket grants access to the public or other AWS accounts, including accounts outside of your organization, to support a specific use case (for example, a static website, public downloads, or cross-account sharing), you can archive the finding for the bucket. When you archive bucket findings, you acknowledge and record your intent for the bucket to remain public or shared. Archived bucket findings remain in your Access Analyzer for S3 listing so that you always know which buckets are public or shared.

To archive bucket findings in Access Analyzer for S3

1. Open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the navigation pane, choose **Access analyzer for S3**.
3. In Access Analyzer for S3, choose an active bucket.
4. To acknowledge your intent for this bucket to be accessed by the public or other AWS accounts, including accounts outside of your organization, choose **Archive**.
5. Enter `confirm`, and choose **Archive**.

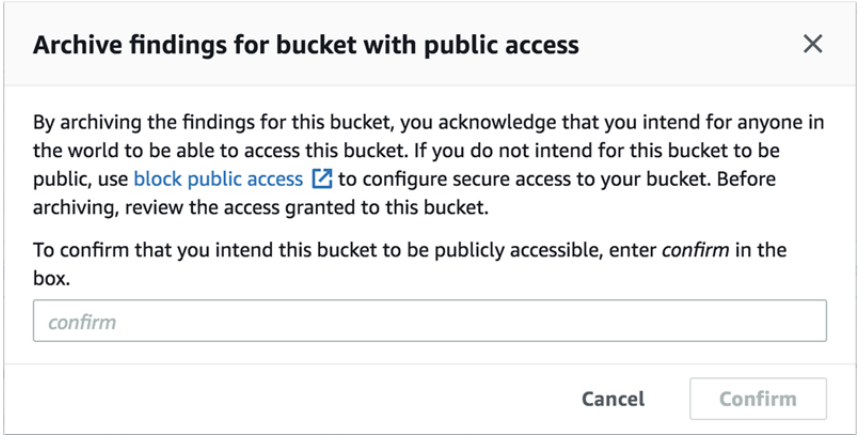

Activating an archived bucket finding

After you archive findings, you can always revisit them and change their status back to active, indicating that the bucket requires another review.

To activate an archived bucket finding in Access Analyzer for S3

1. Open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the navigation pane, choose **Access analyzer for S3**.
3. Choose the archived bucket findings.
4. Choose **Mark as active**.

Viewing finding details

If you need to see more information about a bucket, you can open the bucket finding details in IAM Access Analyzer on the IAM console.

To view finding details in Access Analyzer for S3

1. Open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the navigation pane, choose **Access analyzer for S3**.
3. In Access Analyzer for S3, choose a bucket.
4. Choose **View details**.

   The finding details open in IAM Access Analyzer on the IAM console.

Downloading an Access Analyzer for S3 report

You can download your bucket findings as a CSV report that you can use for auditing purposes. The report includes the same information that you see in Access Analyzer for S3 on the Amazon S3 console.

To download a report

1. Open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the navigation pane on the left, choose **Access analyzer for S3**.
3. In the Region filter, choose the Region.

Access Analyzer for S3 updates to shows buckets for the chosen Region.

4. Choose **Download report**.

   A CSV report is generated and saved to your computer.

Document history

Latest documentation update: March 27, 2019

The following table describes the important changes in each release of the *Amazon Simple Storage Service Console User Guide* from June 19, 2018, onward. For notification about updates to this documentation, you can subscribe to an RSS feed.

update-history-change	update-history-description	update-history-date
New archive storage class (p. 133)	Amazon S3 now offers a new archive storage class, S3 Glacier Deep Archive, for storing rarely accessed objects. For more information, see How Do I Restore an S3 Object That Has Been Archived? and Storage Classes in the *Amazon Simple Storage Service Developer Guide*.	March 27, 2019
Blocking public access to S3 buckets (p. 133)	Amazon S3 block public access prevents the application of any settings that allow public access to data within S3 buckets. For more information, see Blocking Public Access to S3 Buckets.	November 15, 2018
Filtering enhancements in cross-region replication (CRR) rules (p. 133)	In a CRR rule, you can specify an object filter to choose a subset of objects to apply the rule to. Previously, you could filter only on an object key prefix. In this release, you can filter on an object key prefix, one or more object tags, or both. For more information, see How Do I Add a Replication Rule to an S3 Bucket?.	September 19, 2018
Updates now available over RSS (p. 133)	You can now subscribe to an RSS feed to receive notifications about updates to the Amazon Simple Storage Service Console User Guide guide.	June 19, 2018

Earlier updates

The following table describes the important changes in each release of the *Amazon Simple Storage Service Console User Guide* before June 19, 2018.

Change	Description	Date changed
New storage class	Amazon S3 now offers a new storage class, ONEZONE_IA (IA, for infrequent access) for storing objects. For more information, see Storage Classes in the *Amazon Simple Storage Service Developer Guide*.	April 4, 2018
Support for ORC-formatted Amazon S3 inventory files	Amazon S3 now supports the Apache optimized row columnar (ORC) format in addition to comma-separated values (CSV) file format for inventory output files. For more information, see How Do I Configure Amazon S3 Inventory? (p. 104).	November 17, 2017
Bucket permissions check	Bucket permissions check in the Amazon S3 console checks bucket policies and bucket access control lists (ACLs) to identify publicly accessible buckets. Bucket permissions check makes it easier to identify S3 buckets that provide public read and write access. .	November 06, 2017
Default encryption for S3 buckets	Amazon S3 default encryption provides a way to set the default encryption behavior for an S3 bucket. You can set default encryption on a bucket so that all objects are encrypted when they are stored in the bucket. The objects are encrypted using server-side encryption with either Amazon S3-managed keys (SSE-S3) or AWS KMS-managed keys (SSE-KMS). For more information, see How do I enable default encryption for an Amazon S3 bucket? (p. 8).	November 06, 2017
Encryption status in Amazon S3 inventory	Amazon S3 now supports including encryption status in Amazon S3 inventory so you can see how your objects are encrypted at rest for compliance auditing or other purposes. You can also configure to encrypt Amazon S3 inventory with server-side encryption (SSE) or SSE-KMS so that all inventory files are encrypted accordingly. For more information, see How Do I Configure Amazon S3 Inventory? (p. 104).	November 06, 2017
Cross-region replication enhancements	Cross-region replication now supports the following: • By default, Amazon S3 does not replicate objects in your source bucket that are created using server-side encryption using AWS KMS-managed keys. You can now configure a replication rule to replicate these objects. For more information, see How do I add a replication rule to an S3 bucket? (p. 83). • In a cross-account scenario, you can configure a replication rule to change replica ownership to the AWS account that owns the destination bucket. For more information, see Adding a replication rule when the destination bucket is in a different AWS account (p. 91).	November 06, 2017
Added functionality and documentation	The Amazon S3 console now supports enabling object-level logging for an S3 bucket with AWS CloudTrail data events logging. For more information, see How do I enable object-level logging for an S3 bucket with AWS CloudTrail data events? (p. 14).	October 19, 2017
Old Amazon S3 console no longer available	The old version of the Amazon S3 console is no longer available and the old user guide was removed from the Amazon S3 documentation site.	August 31, 2017

Change	Description	Date changed
General availability of New Amazon S3 console	Announced the general availability of the new Amazon S3 console.	May 15, 2017

AWS glossary

For the latest AWS terminology, see the AWS glossary in the *AWS General Reference*.