

Hands-on tutorials

Replicate Data within and between AWS Regions



Replicate Data within and between AWS Regions: Hands-on tutorials

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

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

Table of Contents

Replicate Data within and between AWS Regions using Amazon S3 Replication	i
Overview	1
What you will accomplish	2
Prerequisites	2
Implementation	2
Conclusion	20

Replicate Data within and between AWS Regions using Amazon S3 Replication

AWS experience	Beginner
Time to complete	20 minutes
Cost to complete	Less than \$1 (Amazon S3 pricing page)
Services used	Amazon S3
Last updated	October 17, 2022

Overview

[Amazon S3 Replication](#) is an elastic, fully managed, low-cost feature that replicates objects between Amazon S3 buckets. S3 Replication gives you the ability to replicate data from one source bucket to multiple destination buckets in the same, or different, AWS Regions. Whether you want to maintain a secondary copy of your data for data protection, or have data in multiple geographies to provide users with the lowest latency, S3 Replication gives you the controls you need to meet your business needs. This Amazon S3 getting started guide shows you how to follow S3 Replication best practices with S3 Same-Region Replication (SRR), S3 Cross-Region Replication (CRR), S3 Replication Time Control (S3 RTC), and S3 Batch Replication.

With [S3 Same-Region Replication \(SRR\)](#), you can automatically replicate data between buckets within the same AWS Region to help aggregate logs into a single bucket, replicate between developer and test accounts, and abide by data sovereignty laws. With [S3 Cross-Region Replication \(CRR\)](#), you can replicate objects (and their respective metadata and object tags) into other AWS Regions for reduced latency, compliance, security, disaster recovery, and regional efficiency. You can also enable [S3 Replication Time Control \(S3 RTC\)](#) to help you meet compliance or business requirements for data replication. S3 RTC replicates most objects that you upload to Amazon S3 in seconds, and 99.99 percent of those objects within 15 minutes. To replicate existing objects, you can use [S3 Batch Replication](#) to backfill a newly created bucket with existing objects, retry objects that were previously unable to replicate, migrate data across accounts, or add new buckets to your data lake. For more information on S3 Replication, visit the [Replicating Objects](#) section in

the **Amazon S3 User Guide**. By the end of this tutorial, you will be able to replicate data within and between AWS Regions using Amazon S3 Replication.

What you will accomplish

In this tutorial, you will:

- Create an S3 bucket
- Create an S3 Replication rule on your S3 bucket
- Choose destination S3 bucket
- Choose or create IAM roles for replication
- Specify encryption type (*optional*)
- Choose destination S3 storage class
- Enable additional replication options (*optional*)

Prerequisites

To complete this tutorial, you need an AWS account. [Access this support page for more information on how to create and activate a new AWS account.](#)

Implementation


Step 1: Create the source and destination buckets

1. Sign in to the Amazon S3 console

If you have not already done so, create an [AWS account](#).


Log in to the [AWS Management Console](#) using your account information.

From the AWS console services search bar, enter **S3**. Under the services search results section, select **S3**.




2. Create your first bucket

Choose **Buckets** from the Amazon S3 menu in the left navigation pane and then choose the **Create bucket** button.



3. Configure bucket properties

Enter a descriptive, globally unique name for your bucket. Select which **AWS Region** you would like your bucket created in. For this example, the **EU (Frankfurt) eu-central-1 Region** is selected.



Amazon S3 > Buckets > Create bucket

Create bucket Info

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

General configuration

Bucket name

Bucket name must be globally unique and must not contain spaces or uppercase letters. [See rules for bucket naming](#)


AWS Region

Copy settings from existing bucket - *optional*
Only the bucket settings in the following configuration are copied.

4. Enable bucket versioning

S3 Replication requires **Bucket Versioning** to be enabled for both source and destination S3 buckets. For more information about versioning, see [Using versioning in S3 buckets](#).

You can leave the remaining options as defaults. Navigate to the bottom of the page and choose **Create bucket**.



Bucket Versioning

Versioning is a means of keeping multiple variants of an object in the same bucket. You can use versioning to preserve, retrieve, and restore every version of every object stored in your Amazon S3 bucket. With versioning, you can easily recover from both unintended user actions and application failures. [Learn more](#)


Bucket Versioning

Disable

Enable

5. Create your destination bucket

Repeat the above steps to create another S3 bucket to serve as the destination bucket for replicating objects. Make sure to enable **Bucket Versioning** for the destination S3 bucket as well.



Amazon S3 > Buckets > Create bucket

Create bucket [Info](#)

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

General configuration

Bucket name

aws-s3-replication-tutorial-destination-bucket

Bucket name must be globally unique and must not contain spaces or uppercase letters. [See rules for bucket naming](#)

AWS Region

US West (N. California) us-west-1


Copy settings from existing bucket - optional
Only the bucket settings in the following configuration are copied.

[Choose bucket](#)

Step 2: Create an S3 Replication on your S3 bucket

1. Choose the source bucket

From your list of S3 buckets, choose the S3 bucket that you want to configure as your source for replication.



Amazon S3 > Buckets

Account snapshot
Last updated: Aug 31, 2022 by Storage Lens. Metrics are generated every 24 hours. [Learn more](#)

[View Storage Lens dashboard](#)

Buckets (26) [Info](#)

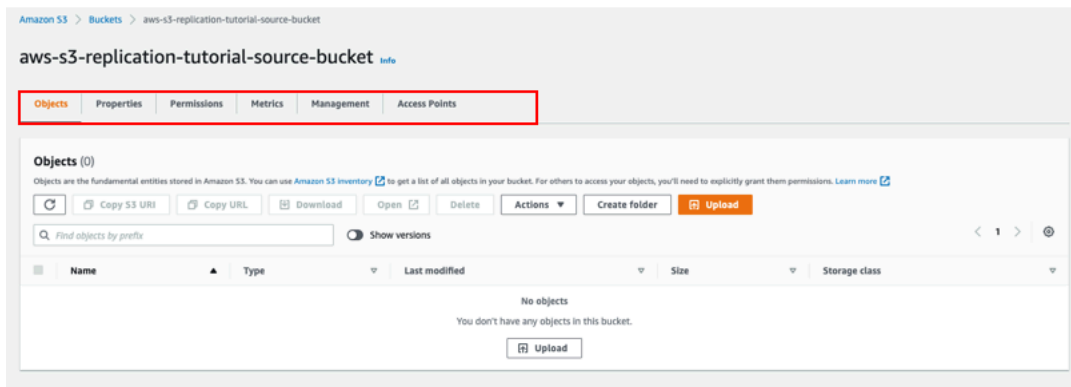
[Refresh](#) [Copy ARN](#) [Empty](#) [Delete](#) [Create bucket](#)

Find buckets by name

Name	AWS Region	Access	Creation date
<input type="radio"/> ack-test-3-us-east-2	US West (N. California) us-west-1	Bucket and objects not public	August 24, 2022, 18:02:19 (UTC-05:00)
<input type="radio"/> ack-test-bucket-us-east-1	US East (N. Virginia) us-east-1	Bucket and objects not public	August 24, 2022, 18:01:16 (UTC-05:00)
<input type="radio"/> adi-test-bucket-us-west-1	US West (N. California) us-west-1	Bucket and objects not public	August 24, 2022, 18:01:45 (UTC-05:00)
<input type="radio"/> adi-test-bucket0	EU (Frankfurt) eu-central-1	Bucket and objects not public	August 24, 2022, 18:00:31 (UTC-05:00)
<input type="radio"/> aws-s3-replication-tutorial-destination-bucket	US West (N. California) us-west-1	Bucket and objects not public	September 1, 2022, 16:54:47 (UTC-05:00)
<input type="radio"/> aws-s3-replication-tutorial-source-bucket	EU (Frankfurt) eu-central-1	Bucket and objects not public	September 1, 2022, 16:47:22 (UTC-05:00)

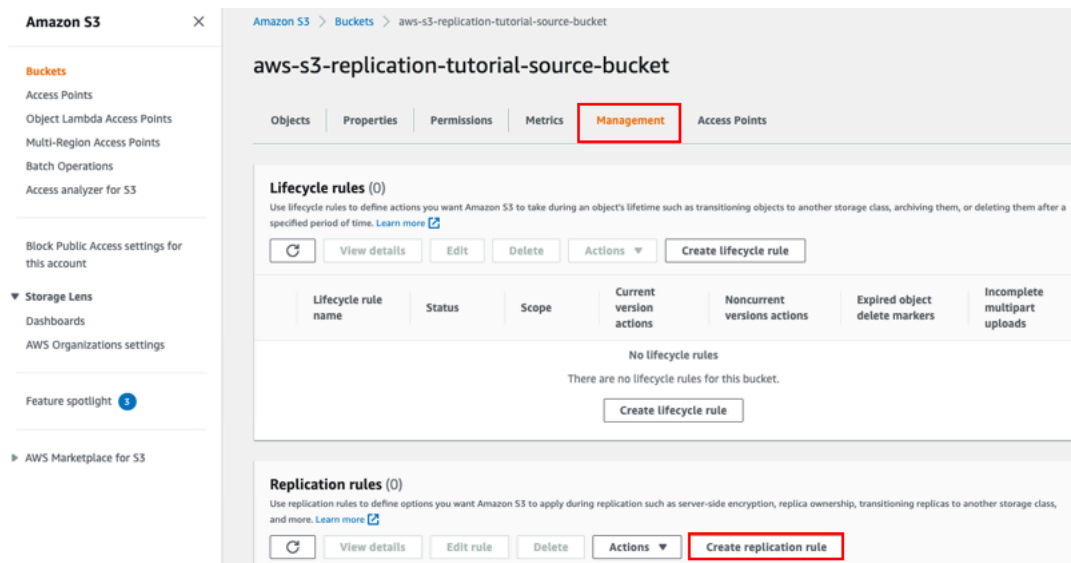
2. Review bucket details

Once you select the source S3 bucket, the console takes you to the S3 bucket landing page, as shown in the following screenshot. Here, you can review the **Objects**, **Properties**, **Permissions**, **Metrics**, **Management**, and **Access Points** for the selected S3 bucket.



3. Create a replication rule

Choose the **Management** tab of the replication source bucket. Under **Management**, you will see **Replication rules**. Select **Create replication rule**.



Step 3: Configure a replication rule

1. Enable replication

Provide a **Replication rule name** and enable the replication rule by selecting **Enabled** under the **Status** section. If the replication rule is disabled, it will not run.

Priority indicates which rule has precedence whenever two or more replication rules conflict. You will have the option to edit the priority of each replication rule on the replication configuration page. Amazon S3 attempts to replicate objects according to all replication rules. However, if there are two or more rules with the same destination bucket, then objects are

replicated according to the rule with the highest priority. A rule with priority 1 is executed before a rule with priority 2. The lower the number, the higher the priority. For example, say you have a replication rule to replicate all objects with tag **foo1** and another replication rule to replicate all objects with tag **foo2**. If you have one object tagged with **foo1** and **foo2**, it will only get replicated with the replication rule with the higher priority. When you have only one replication rule going to one destination bucket, priority is not considered.

Replication rule configuration

Replication rule name

Up to 255 characters. In order to be able to use CloudWatch metrics to monitor the progress of your replication rule, the replication rule name must only contain English characters.

Status

Choose whether the rule will be enabled or disabled when created.

 Enabled Disabled

Priority

The priority value resolves conflicts that occur when an object is eligible for replication under multiple rules to the same destination. The rule is added to the configuration at the highest priority and the priority can be changed on the replication rules table.

2. Set the replication scope

Narrow the scope of replication by defining a **Filter type (Prefix or Tags)**, or choose to replicate the entire bucket. For example, if you want to only replicate objects that are in the **Prefix** "Finance", specify that scope. For more information on filtering objects for replication, visit the documentation on [specifying a filter in the S3 User Guide](#).

Source bucket

Source bucket name

Source Region

Choose a rule scope

 Limit the scope of this rule using one or more filters Apply to all objects in the bucket

Filter type

You can filter objects by prefix, object tags, or a combination of both.

Prefix

Add a filter to limit the scope of this rule to a single prefix.

Don't include the bucket name in the prefix. Using certain characters in key names can cause problems with some applications and protocols.

Tags

You can limit the scope of this rule to the key value pairs added below.

3. Choose the destination bucket

Choose the destination bucket by selecting the **Browse S3** button. You can replicate to a destination bucket in the same or different AWS Region, and in the same or different AWS account. Note that you will need two different S3 buckets to configure replication, and both buckets (source and destination) must have S3 Versioning enabled. The S3 console does not provide you with a way to create a new S3 bucket in the replication setup process. In this example, we chose the destination bucket to be “aws-s3-replication-tutorial-destination-bucket.”

Destination

Destination

You can replicate objects across buckets in different AWS Regions (Cross-Region Replication) or you can replicate objects across buckets in the same AWS Region (Same-Region Replication). You can also specify a different bucket for each rule in the configuration. [Learn more](#) or see [Amazon S3 pricing](#)

- Choose a bucket in this account
 Specify a bucket in another account

Bucket name

Choose the bucket that will receive replicated objects.

aws-s3-replication-tutorial-destination-bucket

Browse S3

Destination Region

US West (N. California) us-west-1

4. Set up IAM permissions

When creating new replication rules from the same source bucket, make sure that the **IAM role** associated with this configuration has sufficient permissions to write new objects in the new destination bucket. You can choose to create a new IAM role or select an existing IAM role with the right set of permissions. For more information, see the documentation on [setting up permissions for S3 Replication](#).

Destination Region

US West (N. California) us-west-1

IAM role

- Choose from existing IAM roles
 Enter IAM role ARN

IAM role

Create new role



View

5. (Optional) Configure encryption

Skip this step if your objects are not encrypted.

If your objects are encrypted with [Amazon S3-managed encryption keys \(SSE-S3\)](#) or [AWS Key Management Service \(AWS KMS\)](#), you will need to specify the encryption options while setting up replication. S3 Replication supports SSE-S3 (default encryption) and AWS KMS server-side encryption. If you choose AWS KMS encryption, you will need to provide the AWS KMS keys to decrypt in source and re-encrypt in destination. To save on [AWS KMS costs](#), you also have the option to enable [Amazon S3 Bucket Keys](#).

Encryption

Replicate objects encrypted with AWS KMS
You can replicate objects that are encrypted with AWS Key Management Service (AWS KMS) keys.

⚠ Replication might increase the number of KMS requests you will make in the source and destination AWS Regions. [Learn more about KMS operation quotas](#)

AWS KMS keys for decrypting source objects (1)

Search by alias or key ID

< 1 >

Alias Key ID

aws/s3

Customer ARNs - optional

Add custom ARNs using the key ARN or the alias ARN

Add key

AWS KMS key for encrypting destination objects

AWS managed key (aws/s3)

Choose from your AWS KMS keys

Enter AWS KMS key ARN

6. Select the storage class

Next, you will have the option to choose a different S3 storage class for your replicated objects at the destination bucket. Consider choosing lower cost storage classes as appropriate for your workloads. For example, you can choose the **Amazon Glacier Instant Retrieval** storage class if your replicated objects will be infrequently accessed but need to be retrieved in milliseconds, **S3 Glacier Deep Archive** to archive data that rarely needs to be accessed, and **S3 Intelligent-Tiering** to optimize storage costs for data with unpredictable or changing access patterns. For more information, see the documentation on [using Amazon S3 storage classes](#).

Destination storage class

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

Change the storage class for the replicated objects

Storage class

Storage class	Designed for	Availability Zones	Min storage duration
<input checked="" type="radio"/> Standard	Frequently accessed data (more than once a month) with milliseconds access	≥ 3	-
<input type="radio"/> Intelligent-Tiering	Data with changing or unknown access patterns	≥ 3	-
<input type="radio"/> Standard-IA	Infrequently accessed data (once a month) with milliseconds access	≥ 3	30 days
<input type="radio"/> One Zone-IA	Recreateable, infrequently accessed data (once a month) stored in a single Availability Zone with milliseconds access	1	30 days
<input type="radio"/> Glacier Instant	Long-lived archive data accessed once a quarter with instant retrieval in	≥ 3	90 days

7. Choose additional options

Choose **Additional replication options**, such as enabling [S3 Replication Time Control \(S3 RTC\)](#), [Replication metrics and notifications](#), [Delete marker replication](#), and [Replica modifications sync](#). **S3 RTC** helps you meet compliance and business requirements as it provides an SLA of 15 minutes to replicate 99.99% of your objects. RTC can be enabled along with S3 Cross-Region Replication (S3 CRR) and S3 Same-Region Replication (S3 SRR) and has replication metrics and notifications enabled, by default. For non-RTC rules, you have the option to select **Replication metrics and notifications**, which provide detailed metrics to track minute-by-minute progress of bytes pending, operations pending, and replication latency for the replication rule. Selecting **Delete marker replication** means deletes on the source bucket will be replicated to the destination bucket. This should be enabled if you want to keep the source and destination buckets in sync, but not if the goal is to protect against accidental or malicious deletes. To establish two-way replication between two S3 buckets, create bidirectional replication rules (A to B, and B to A) and enable **Replica modification sync** for both of the replication rules in the source and destination S3 buckets. This will help you to keep object metadata such as tags, ACLs, and Object Lock settings in sync between replicas and source objects.

Review the replication configuration, and choose **Save**.

Additional replication options

- Replication Time Control (RTC)**
Replication Time Control replicates 99.99% of new objects within 15 minutes and provides replication metrics and notifications. Additional fees will apply. [Learn more](#)
- Replication metrics and notifications**
Monitor the progress of your replication rule through Cloudwatch Metrics. Cloudwatch metrics fees apply. [Learn more](#) or see [Amazon Cloudwatch pricing](#)
- Delete marker replication**
Delete markers created by S3 delete operations will be replicated. Delete markers created by lifecycle rules are not replicated. [Learn more](#)
- Replica modification sync**
Replicate metadata changes made to replicas in this bucket to the destination bucket. [Learn more](#)

Cancel **Save**

Step 4: Create another replication rule

Create another S3 Replication rule on the same source S3 bucket to another destination S3 bucket.

1. Name and enable your rule

Repeat the previous steps to create another S3 Replication rule from the same source S3 bucket to another destination S3 bucket. Provide a **Replication rule name** and enable the replication rule by selecting **Enabled** under the **Status** section. Choose what to replicate by choosing the scope of the replication rule.

Amazon S3 > Buckets > aws-s3-replication-tutorial-source-bucket > Replication rules > Create replication rule

Create replication rule

Replication rule configuration

Replication rule name
CRR-second-destination-noRTC
Up to 255 characters. In order to be able to use CloudWatch metrics to monitor the progress of your replication rule, the replication rule name must only contain English characters.

Status
Choose whether the rule will be enabled or disabled when created.
 Enabled
 Disabled

Priority
The priority value resolves conflicts that occur when an object is eligible for replication under multiple rules to the same destination. The rule is added to the configuration at the highest priority and the priority can be changed on the replication rules table.
2

Source bucket

Source bucket name
aws-s3-replication-tutorial-source-bucket

Source Region
EU (Frankfurt) eu-central-1

Choose a rule scope
 Limit the scope of this rule using one or more filters
 Apply to all objects in the bucket

Destination

Destination
You can replicate objects across buckets in different AWS Regions (Cross-Region Replication) or you can replicate objects across buckets in the same AWS Region (Same-Region Replication). You can also specify a different bucket for each rule in the configuration. [Learn more](#) or see [Amazon S3 pricing](#)

Choose a bucket in this account
 Specify a bucket in another account

2. Choose objects to replicate

Choose the destination bucket by selecting the **Browse S3** button. In this example, we chose the destination bucket to be "ack-test-bucket-us-east-1".

Destination

Destination
You can replicate objects across buckets in different AWS Regions (Cross-Region Replication) or you can replicate objects across buckets in the same AWS Region (Same-Region Replication). You can also specify a different bucket for each rule in the configuration. [Learn more](#) or see [Amazon S3 pricing](#)

Choose a bucket in this account
 Specify a bucket in another account

Bucket name
Choose the bucket that will receive replicated objects.
ack-test-bucket-us-east-1

Destination Region
US East (N. Virginia) us-east-1

3. Select the storage class

Select the **Destination S3 storage class**. In this example, we chose to replicate to the **S3 Standard** storage class in the destination bucket.

Destination storage class

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

Change the storage class for the replicated objects

Storage class

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

4. Choose additional options

Choose **Additional replication options** for the replication rule. In this example, we chose to enable **S3 Replication Time Control (RTC)**.

Additional replication options

Replication Time Control (RTC)

Replication Time Control replicates 99.99% of new objects within 15 minutes and provides replication metrics and notifications. Additional fees will apply. [Learn more](#)

Replication metrics and notifications

Monitor the progress of your replication rule through Cloudwatch Metrics. Cloudwatch metrics fees apply. [Learn more](#) or see [Amazon Cloudwatch pricing](#)

Delete marker replication

Delete markers created by S3 delete operations will be replicated. Delete markers created by lifecycle rules are not replicated. [Learn more](#)

Replica modification sync

Replicate metadata changes made to replicas in this bucket to the destination bucket. [Learn more](#)


Cancel

Save

Step 5: Review replication configuration

- After you save the replication rule, you are back on the S3 Replication landing page, as shown in the following screenshot. Here, you can review the replication configuration with all the different replication rules, and the rule priorities and the additional options, such as encryption and RTC.

Next, upload a new object to the replication source bucket to test the newly added replication configuration. Confirm that you see that object replicated to the new destination bucket. Replication metrics can take a few minutes to show up in the S3 console.



The screenshot shows the AWS S3 console interface for managing replication rules. At the top, there is a breadcrumb trail: Amazon S3 > Buckets > aws-s3-replication-tutorial-source-bucket > Replication rules. Below this, the page title is "Replication rules" with an "Info" link. A brief description states: "Replication enables automatic and asynchronous copying of objects across buckets in the same or different AWS Regions. A replication configuration is a set of rules that define what options should be applied to a group of objects during replication." There are two buttons: "Create replication job" and "Edit".

The "Replication configuration settings" section includes:

- Source bucket: aws-s3-replication-tutorial-source-bucket
- Source Region: EU (Frankfurt) eu-central-1
- IAM role: s3crr_role_for_aws-s3-replication-tutorial-source-bucket

 Buttons for "Create replication job" and "Edit" are present.

The "Replication rules (2)" section includes a description: "Use replication rules to define options you want Amazon S3 to apply during replication such as server-side encryption, replica ownership, transitioning replicas to another storage class, and more. [Learn more](#)". There are buttons for "View details", "Edit rule", "Delete", "Actions", and "Create replication rule". A pagination control shows "1" of 2 rules.

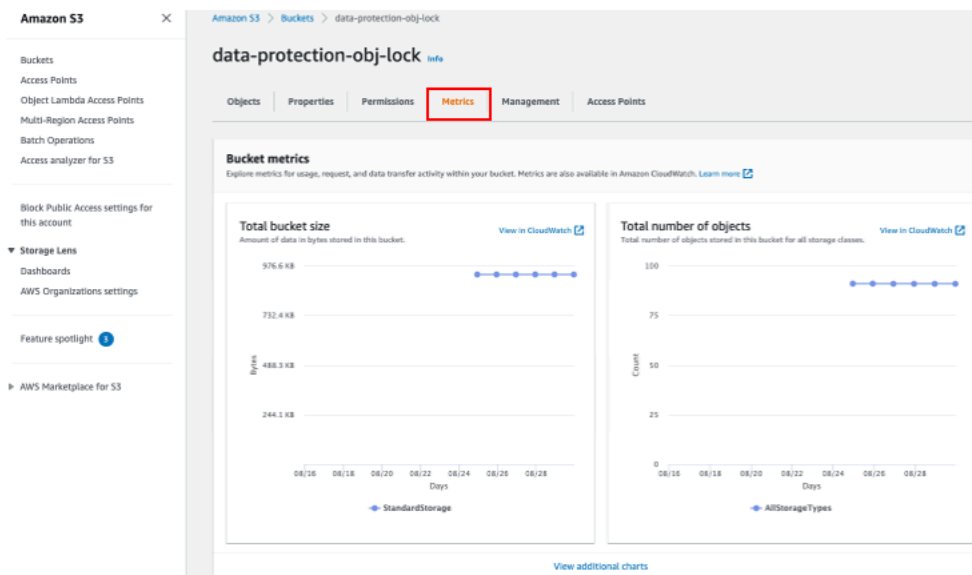
Replication rule name	Status	Destination bucket	Destination Region	Priority	Scope	Storage class	Replica owner	Replication Time Control	KMS-encrypted objects	Replica modification sync
CRR-second-destination-noRTC	Enabled	s3://ack-test-bucket-us-east-1	US East (N. Virginia) us-east-1	1	Entire bucket	Transition to Standard	Same as source	Enabled	Do not replicate	Disabled
testing 1:N replication	Enabled	s3://aws-s3-replication-tutorial-destination-bucket	US West (N. California) us-west-1	0	Prefix: Finance	Same as source	Same as source	Disabled	Do not replicate	Disabled

Step 6: Monitor replication progress

Now that you've configured replication for this bucket, you can track per-destination metrics and notifications.

- View the Metrics tab

Open the **Metrics** tab for the source bucket.

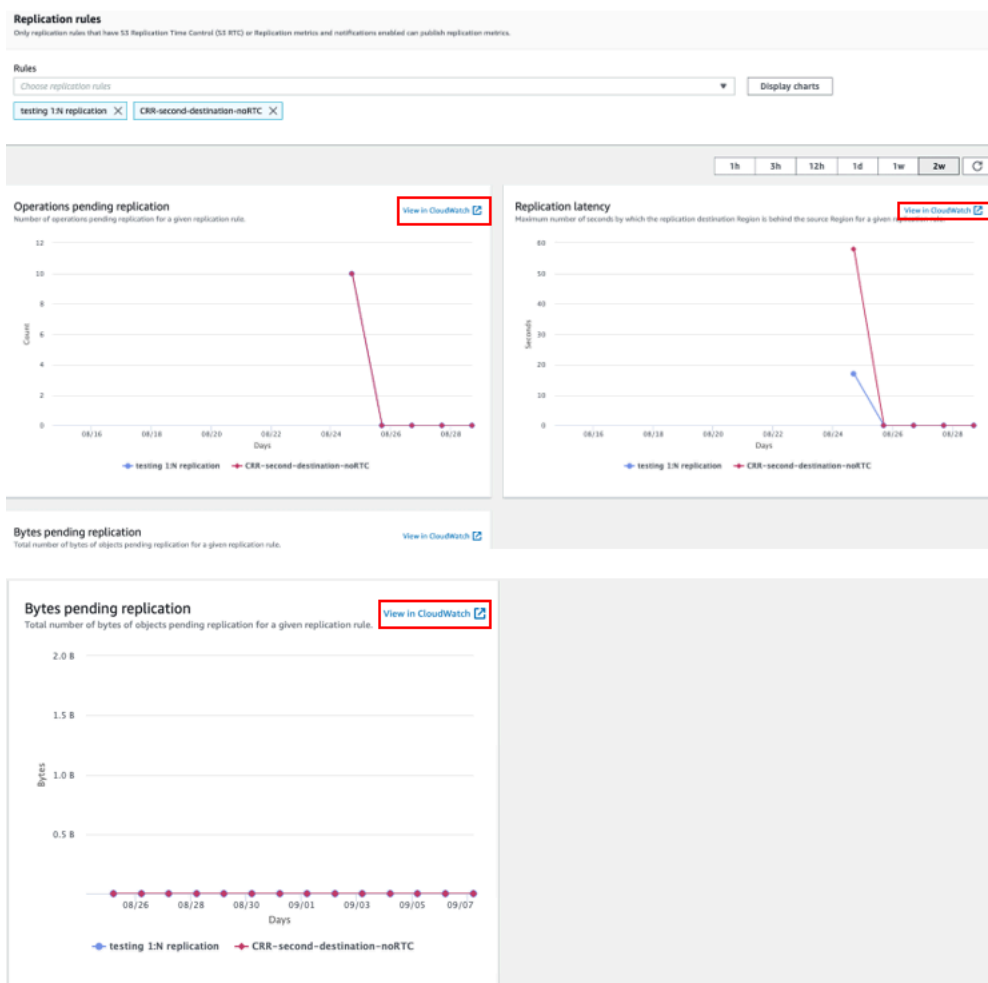


2. Analyze replication metrics

Navigate down to view **Replication metrics** and select one or more **Replication rules** to monitor. Choose **Display charts** to see **Operations pending replication**, **Bytes pending replication**, and **Replication latency** for all replication rules.

3. (Optional) Monitor with CloudWatch

Additionally, you can use the **View in CloudWatch** link to view the **Replication metrics** on [Amazon CloudWatch](#). Here you can get a comprehensive view of the replication metrics for each replication rule, source bucket, and destination bucket in one place. Additionally, you can gather actionable insights and set alarms to monitor the metrics. For more information, see [Using Amazon CloudWatch alarms](#).




Step 7: Monitor replication status for individual objects

You can use [Amazon S3 Inventory](#) to audit and report on the replication status of your objects for business, compliance, and regulatory needs. Replication status can help you determine the current state of an object being replicated. The replication status of a source object will return either **Pending**, **Completed**, or **Failed**. The replication status of a replica will return **Replica**. For more details on replication status, see [Getting replication status information](#). For more information on configuring the Amazon S3 Inventory report, see the documentation on [managing and analyzing your data at scale using Amazon S3 Inventory and Amazon Athena](#).

Clean up resources

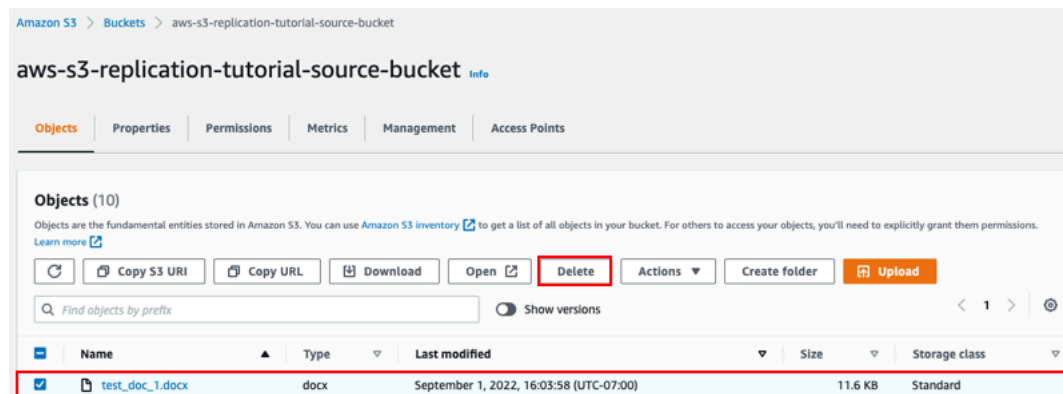
1. Open the S3 console

If you have logged out of your AWS Management Console session, log back in. Navigate to the [S3 console](#) and select the **Buckets** menu option. First, you will need to delete the test object from your test bucket. Select the name of the bucket you have been working with for this tutorial.



2. Select the object

Put a check mark in the check box to the left of your test object name, then choose the **Delete** button.



3. Delete the object

On the **Delete objects** page, verify that you have selected the proper object to delete and enter **delete** into the **Permanently delete objects** confirmation box. Then, choose the **Delete object** button to continue.

Amazon S3 > Buckets > aws-s3-replication-tutorial-source-bucket > Delete objects

Delete objects [Info](#)

⚠ If a folder is selected for deletion, all objects in the folder will be deleted, and any new objects added while the delete action is in progress might also be deleted. If an object is selected for deletion, any new objects with the same name that are uploaded before the delete action is completed will also be deleted. [Learn more](#)

i Deleting the specified objects adds delete markers to them. If you need to undo the delete action, you can delete the delete markers. [Learn more](#)

Specified objects

Find objects by name < 1 >

Name	Type	Last modified	Size
test_doc_1.docx	docx	September 1, 2022, 16:03:58 (UTC-07:00)	11.6 KB

Delete objects?

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

Cancel **Delete objects**

4. Confirm deletion

Next, you will be presented with a banner indicating if the deletion has been successful.

Successfully deleted objects
View details below.

Delete objects: status Close

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

Summary

Source	Successfully deleted	Failed to delete
s3://aws-s3-replication-tutorial-source-bucket	1 object, 11.6 KB	0 objects

Failed to delete Configuration


Failed to delete (0)

Find objects by name < 1 >

Name	Folder	Type	Last modified	Size	Error
No objects failed to delete.					

5. Delete the source bucket

Finally, you need to delete the test buckets you have created. Return to the list of buckets in your account. Select the radio button to the left of the source bucket you created for this tutorial, and then choose the **Delete** button.




The screenshot shows the Amazon S3 Buckets console. On the left is a navigation menu with options like Buckets, Access Points, and Storage Lens. The main area displays an 'Account snapshot' and a list of 26 buckets. The bucket 'aws-s3-replication-tutorial-source-bucket' is selected, indicated by a blue radio button and a red border around its row. The table columns are Name, AWS Region, Access, and Creation date.

Name	AWS Region	Access	Creation date
ack-test-3-us-east-2	US West (N. California) us-west-1	Bucket and objects not public	August 24, 2022, 16:02:19 (UTC-07:00)
ack-test-bucket-us-east-1	US East (N. Virginia) us-east-1	Bucket and objects not public	August 24, 2022, 16:01:16 (UTC-07:00)
ack-test-bucket-us-west-1	US West (N. California) us-west-1	Bucket and objects not public	August 24, 2022, 16:01:45 (UTC-07:00)
ack-test-bucket0	EU (Frankfurt) eu-central-1	Bucket and objects not public	August 24, 2022, 16:00:31 (UTC-07:00)
aws-s3-replication-tutorial-destination-bucket	US West (N. California) us-west-1	Bucket and objects not public	September 1, 2022, 14:54:47 (UTC-07:00)
aws-s3-replication-tutorial-source-bucket	EU (Frankfurt) eu-central-1	Bucket and objects not public	September 1, 2022, 14:47:22 (UTC-07:00)

6. Confirm deletion

Review the warning message. If you desire to continue deletion of this bucket, enter the bucket name into the **Delete bucket** confirmation box, and choose **Delete bucket**.



The screenshot shows the 'Delete bucket' confirmation dialog. It includes a warning message: 'Deleting a bucket cannot be undone. Bucket names are unique. If you delete a bucket, another AWS user can use the name.' Below this, the question 'Delete bucket "aws-s3-replication-tutorial-source-bucket"?' is displayed. A text input field contains the bucket name 'aws-s3-replication-tutorial-source-bucket', which is highlighted with a red box. At the bottom, there are 'Cancel' and 'Delete bucket' buttons.

7. Delete the destination bucket

Repeat the previous steps to delete the destination bucket created as part of this tutorial as well. Return to the list of buckets in your account. Select the radio button to the left of the source bucket you created for this tutorial, and then choose the **Delete** button.

The screenshot shows the Amazon S3 console interface. On the left is a navigation menu with options like Buckets, Access Points, and Storage Lens. The main area displays an 'Account snapshot' and a list of 26 buckets. The bucket 'aws-s3-replication-tutorial-destination-bucket' is selected and highlighted with a red border. The table below shows the details of the buckets.

Name	AWS Region	Access	Creation date
ack-test-3-us-east-2	US West (N. California) us-west-1	Bucket and objects not public	August 24, 2022, 16:02:19 (UTC-07:00)
ack-test-bucket-us-east-1	US East (N. Virginia) us-east-1	Bucket and objects not public	August 24, 2022, 16:01:16 (UTC-07:00)
aci-test-bucket-us-west-1	US West (N. California) us-west-1	Bucket and objects not public	August 24, 2022, 16:01:45 (UTC-07:00)
aci-test-bucket0	EU (Frankfurt) eu-central-1	Bucket and objects not public	August 24, 2022, 16:00:31 (UTC-07:00)
aws-s3-replication-tutorial-destination-bucket	US West (N. California) us-west-1	Bucket and objects not public	September 1, 2022, 14:54:47 (UTC-07:00)

8. Confirm deletion

Review the warning message. If you desire to continue deletion of this bucket, enter the bucket name into the **Delete bucket** confirmation box, and choose **Delete bucket**.

The screenshot shows the 'Delete bucket' confirmation dialog. It includes a warning message: 'Deleting a bucket cannot be undone. Bucket names are unique. If you delete a bucket, another AWS user can use the name.' Below the warning is a question: 'Delete bucket "aws-s3-replication-tutorial-destination-bucket"?'. A text input field contains the bucket name 'aws-s3-replication-tutorial-destination-bucket', which is highlighted with a red box.

Conclusion

Congratulations! You have learned how to use S3 Replication to replicate objects from source to destination S3 buckets across one or many AWS Regions to meet compliance requirements, minimize latency, and increase operational efficiency.

S3 Replication is a fully managed, low cost, policy-based storage management feature designed to require little to no manual intervention. We recommend you enable metrics and notifications for each replication rule, turn on [Amazon S3 Event Notifications](#) on your source bucket, and enable appropriate [Amazon CloudWatch](#) metrics and alerts. Once enabled, you will be able to track the progress of S3 Replication to one or more S3 buckets.