Hands-on tutorials

Replicate Data within and between AWS Regions

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

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**.

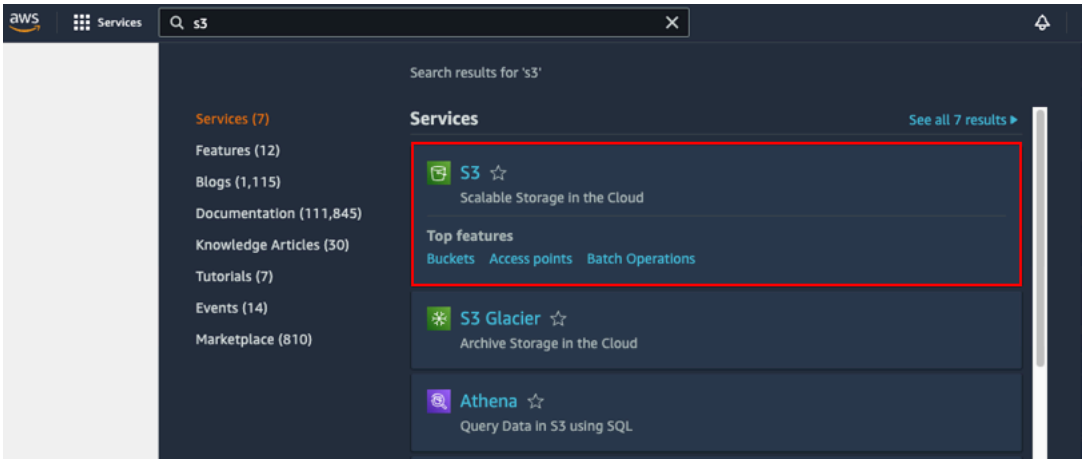

2. Create your first bucket

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

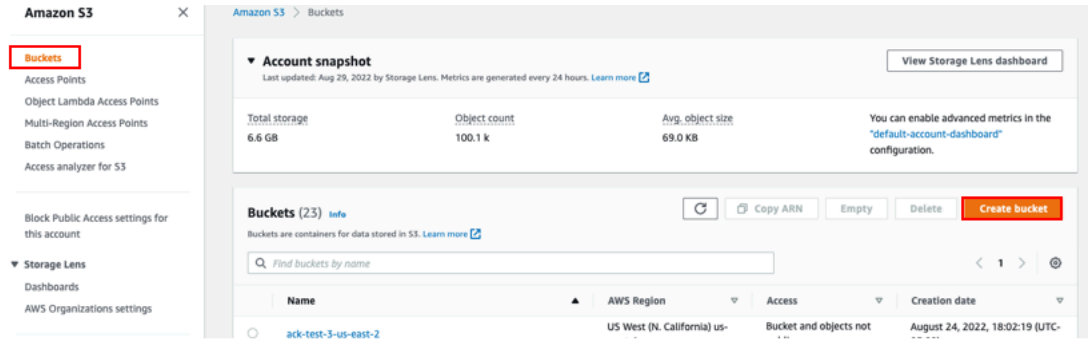

3. Configure bucket properties

   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.

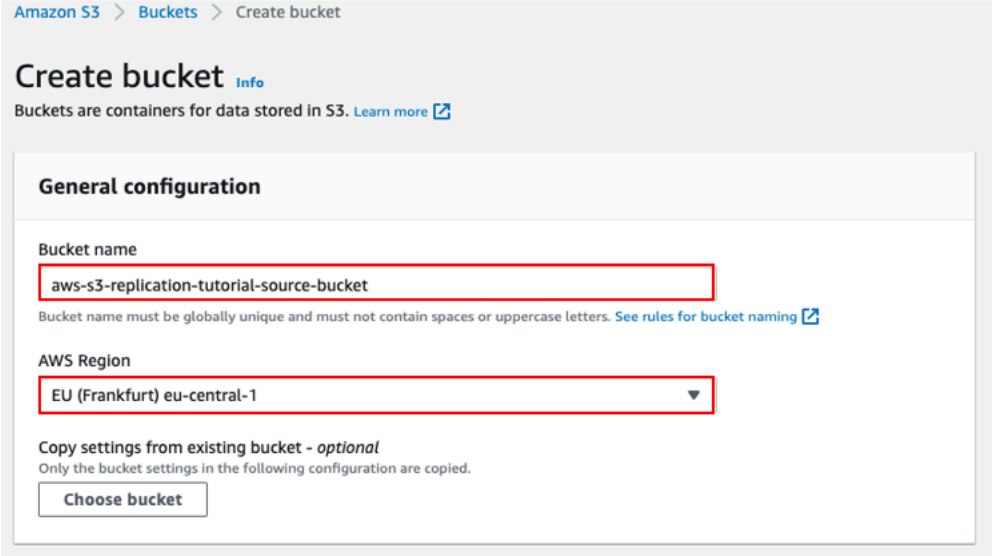

4. Enable bucket versioning

   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.**

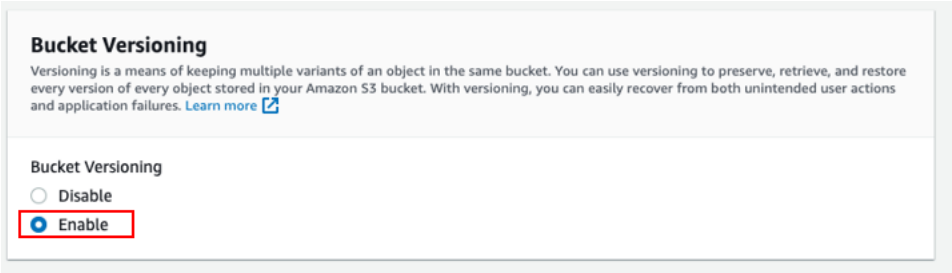

5. Create your destination bucket

   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.

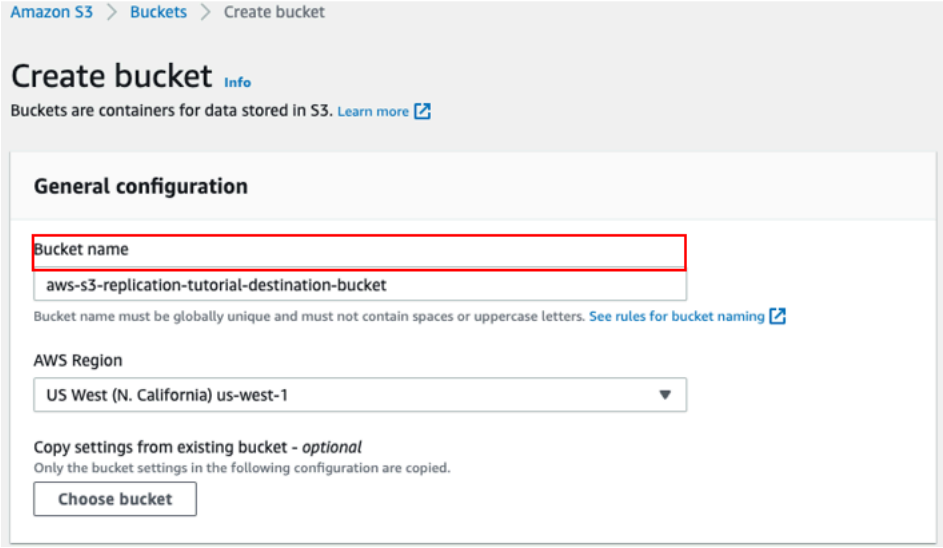

Step 2: Create an S3 Replication on your S3 bucket

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.

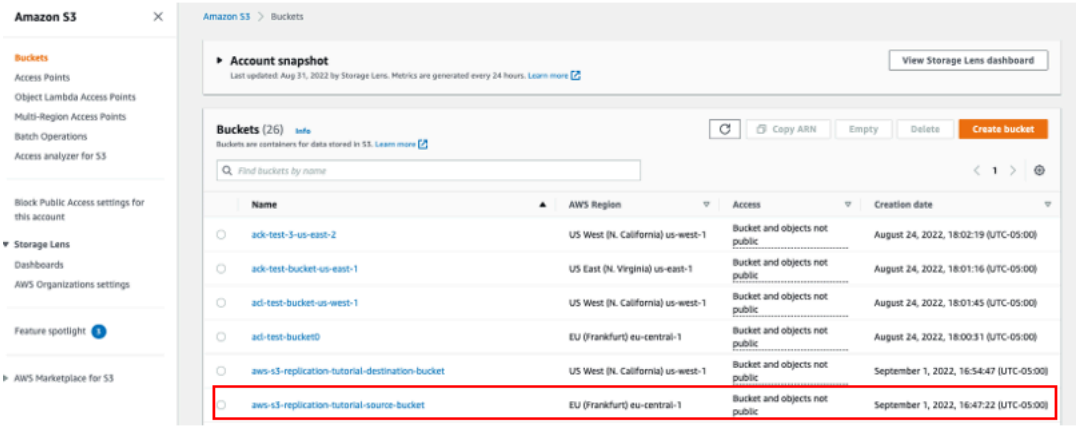

2. Review bucket details

   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.

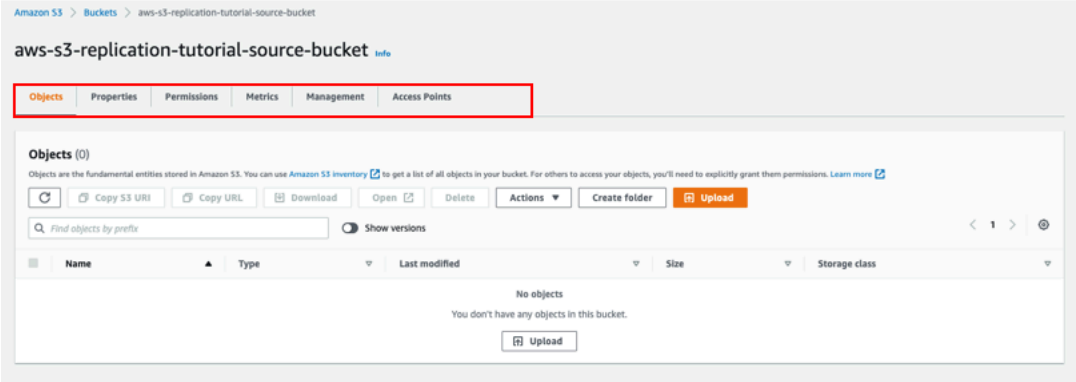

3. Create a replication rule

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

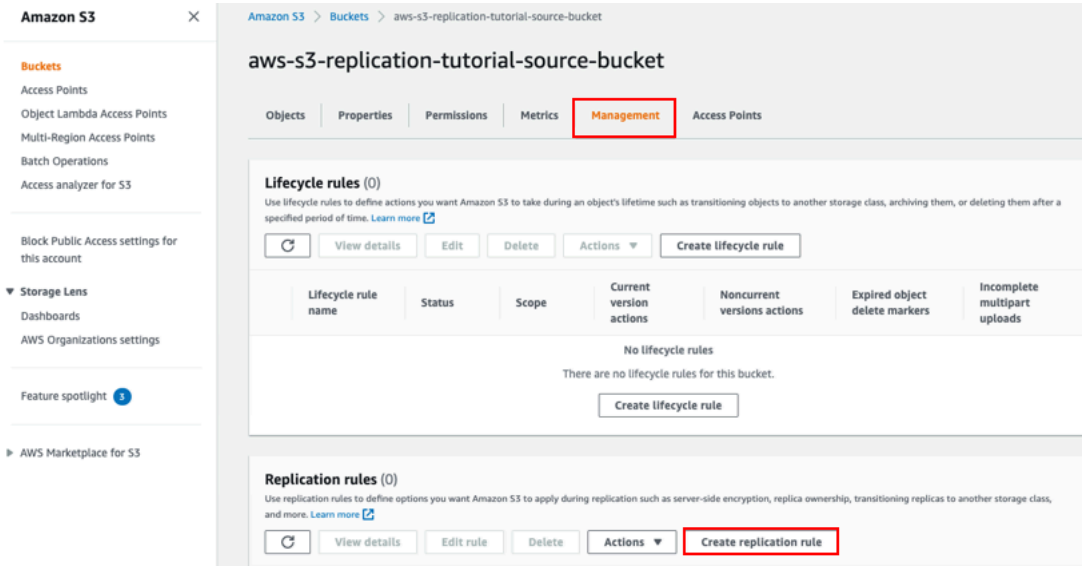

Step 3: Configure a replication rule

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.

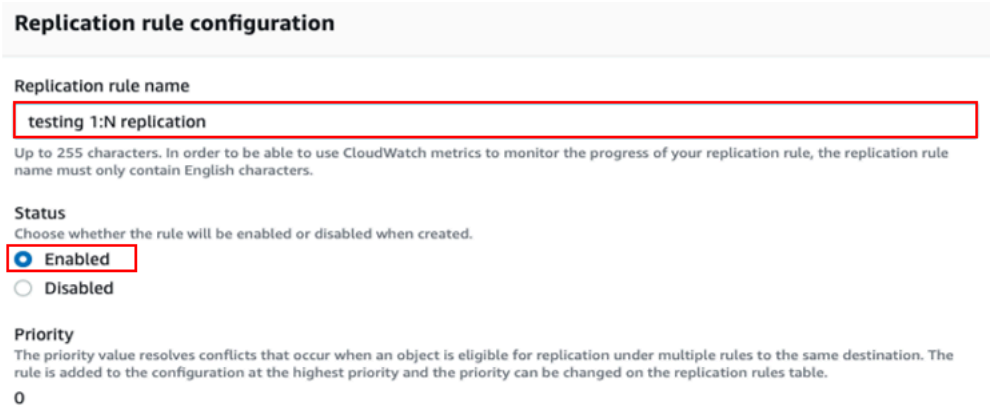

2. Set the replication scope

   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.

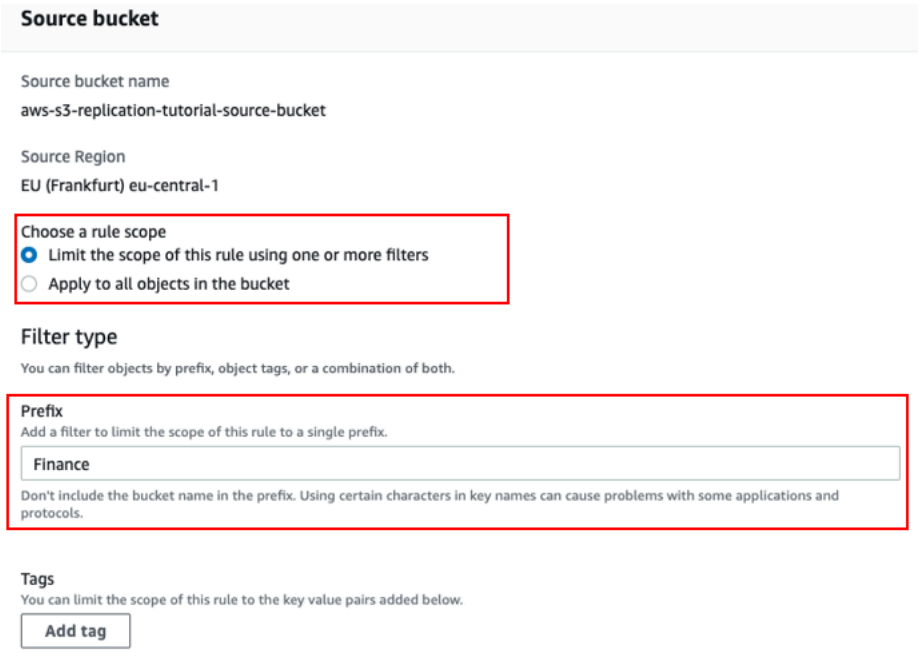

3. Choose the destination bucket

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."

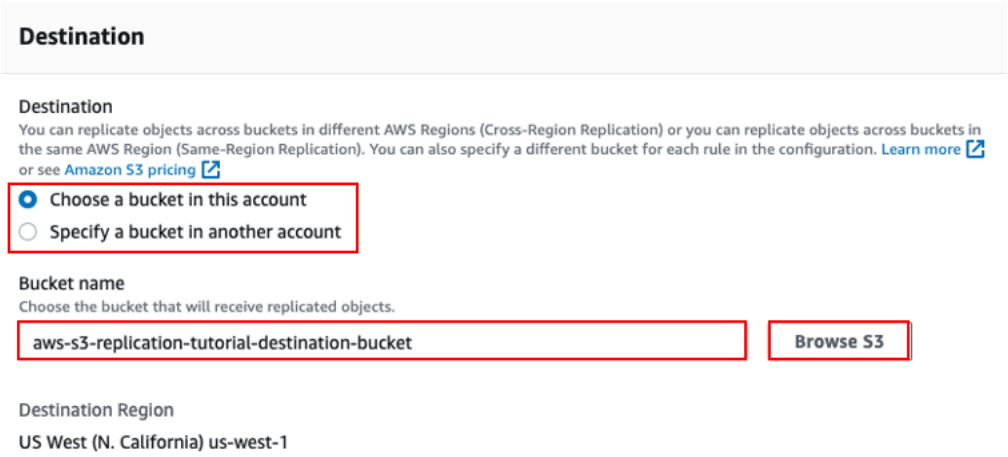

4. Set up IAM permissions

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.

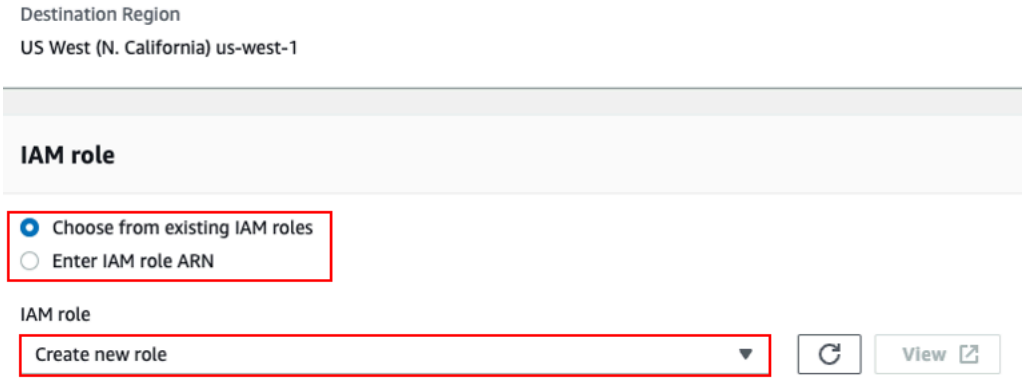

5. (Optional) Configure encryption

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.

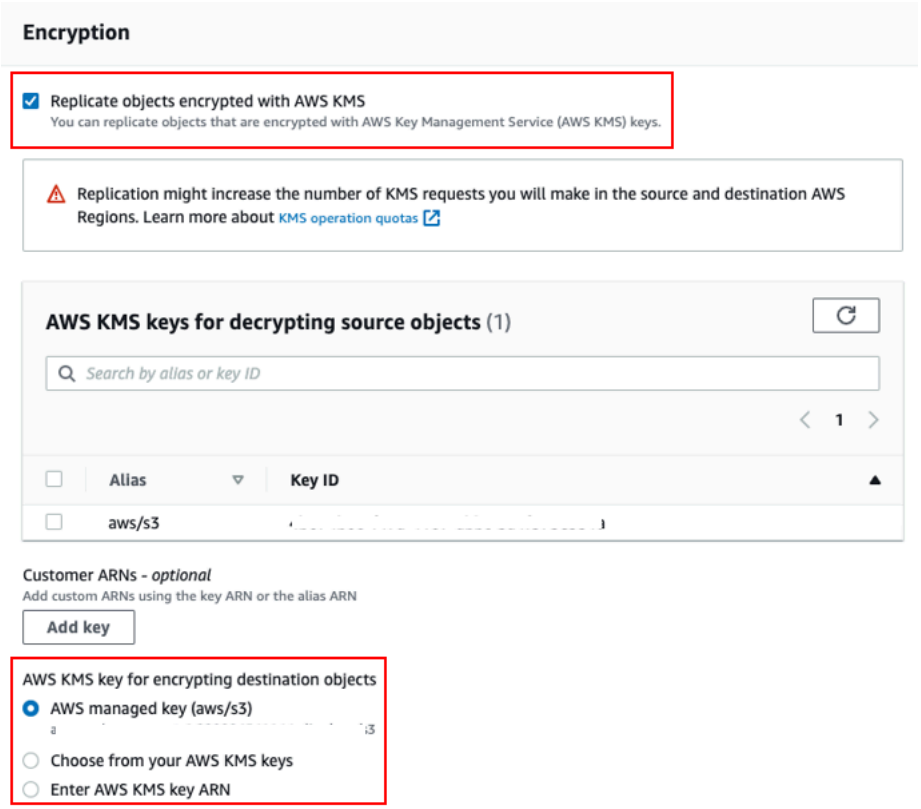

6. Select the storage class

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.

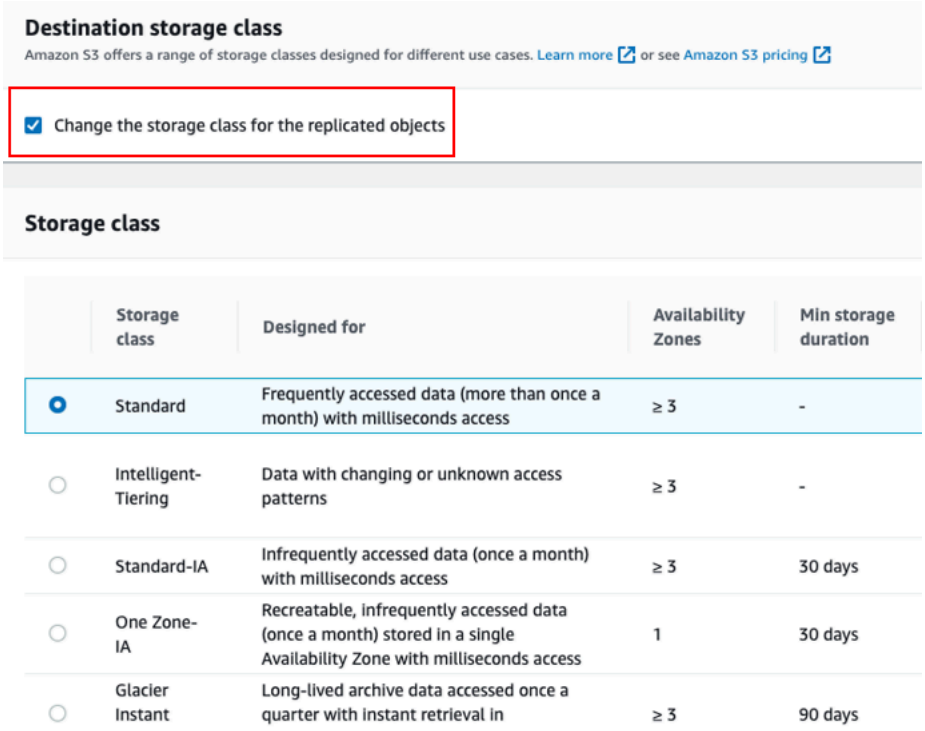

7. Choose additional options

   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**.

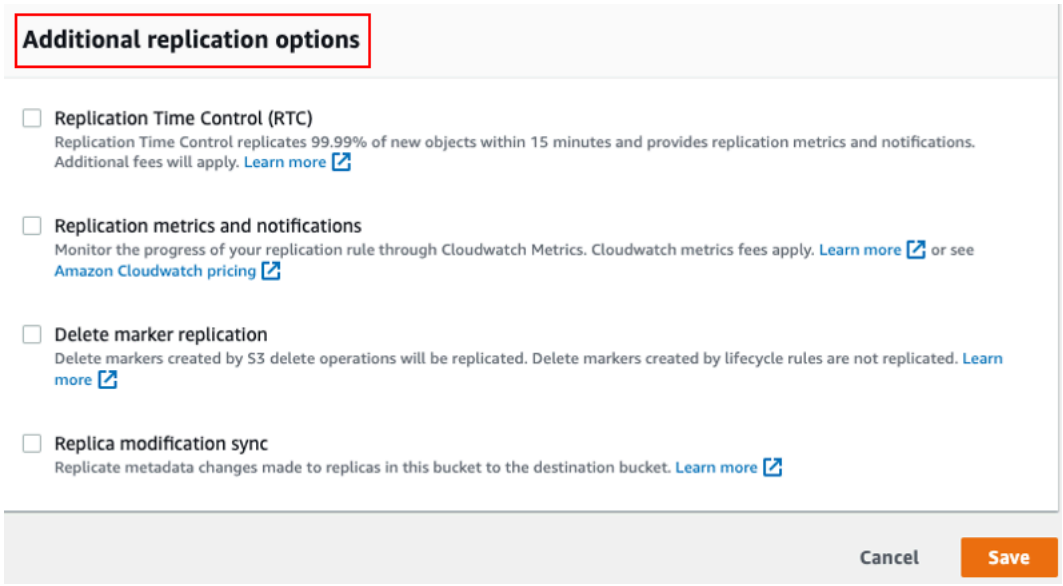

Step 4: Create another replication rule

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.

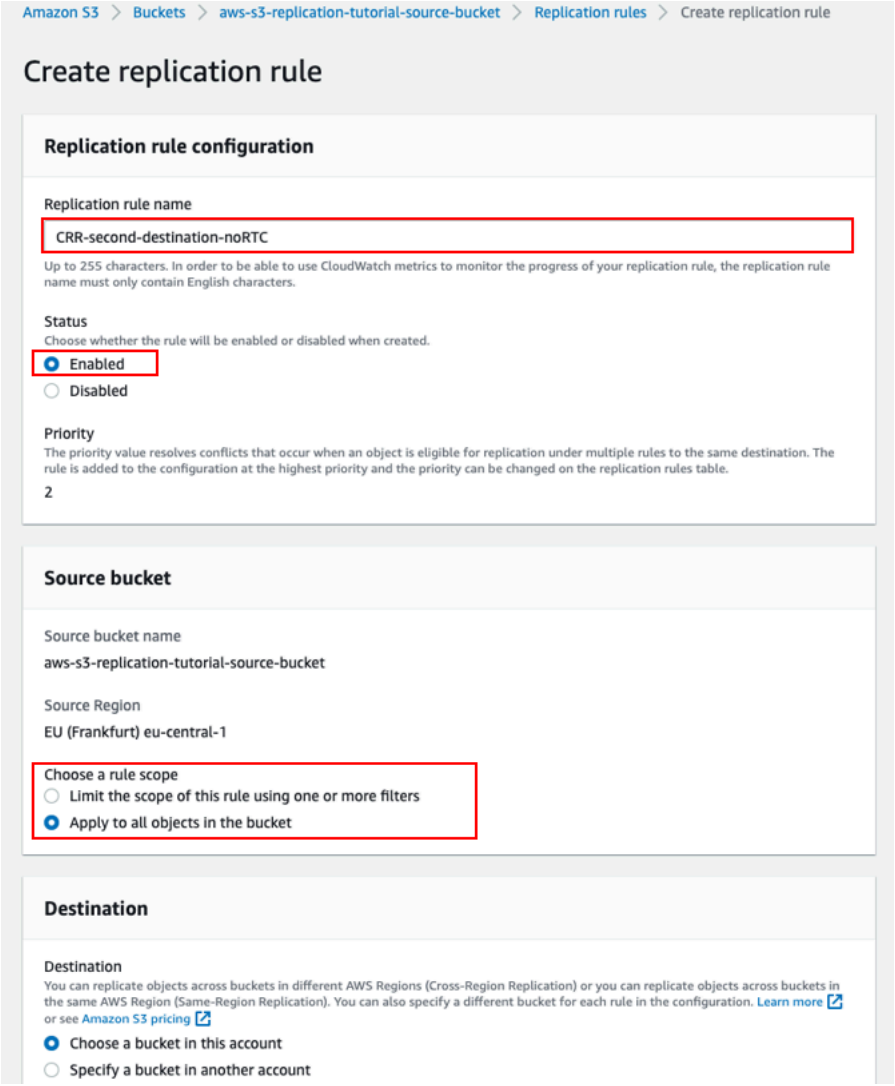

2. Choose objects to replicate

   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”.

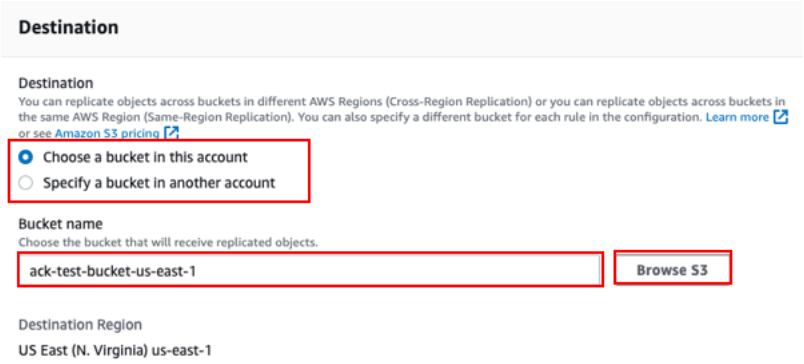

3. Select the storage class

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

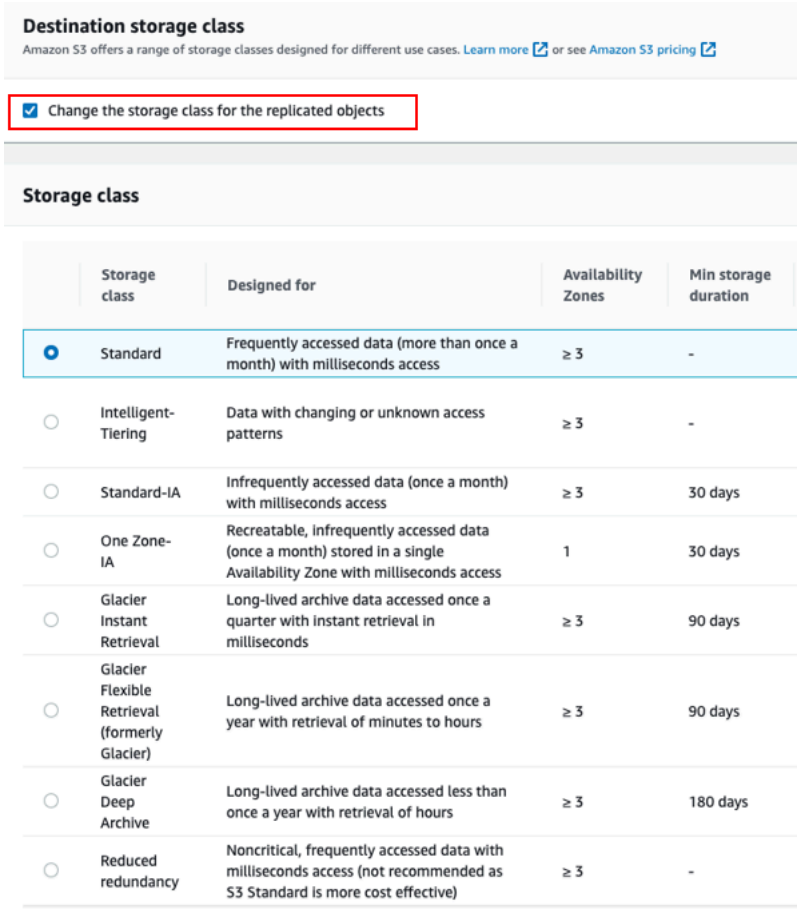

4. Choose additional options

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

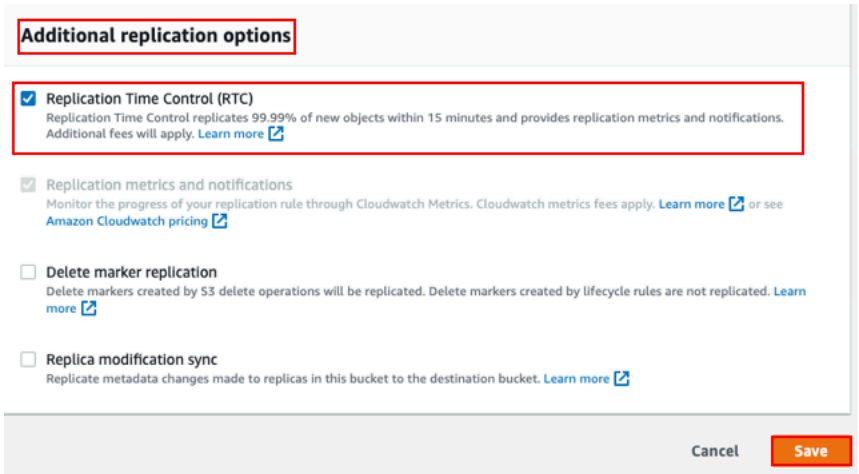

Step 5: Review replication configuration

- 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.

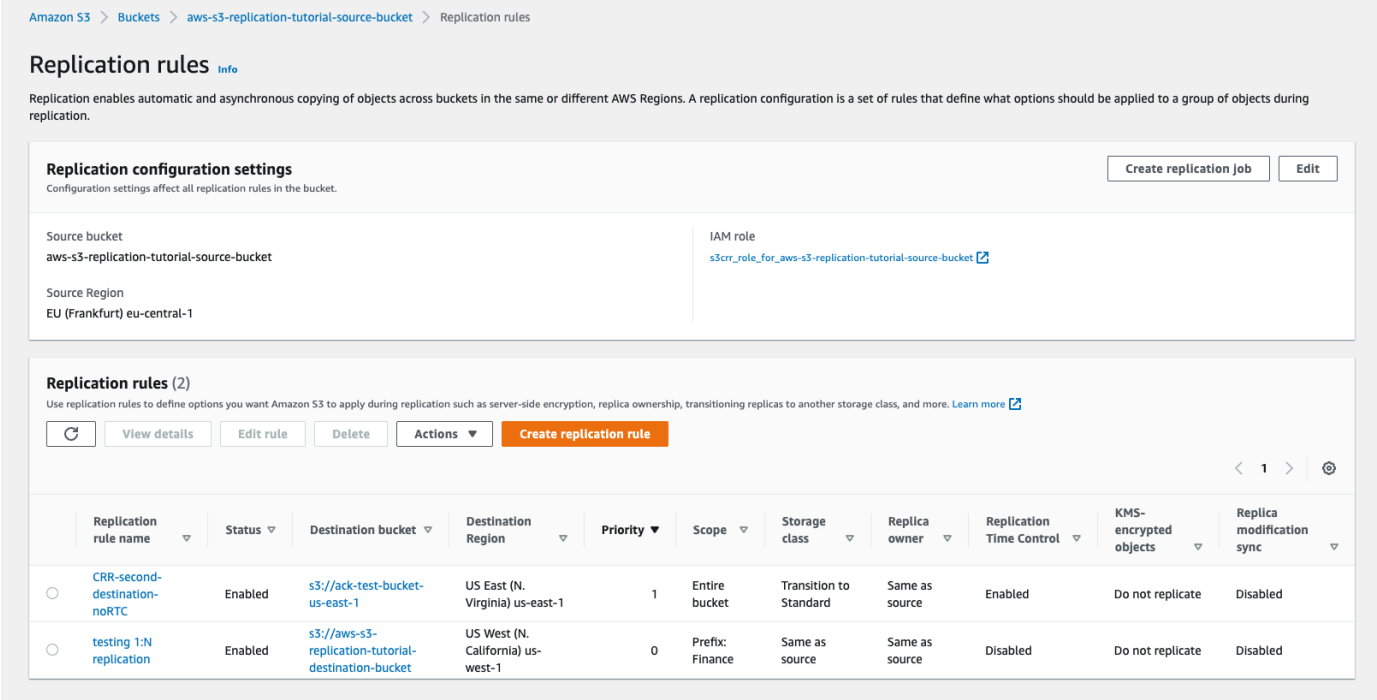

Step 6: Monitor replication progress

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

1. View the Metrics tab

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

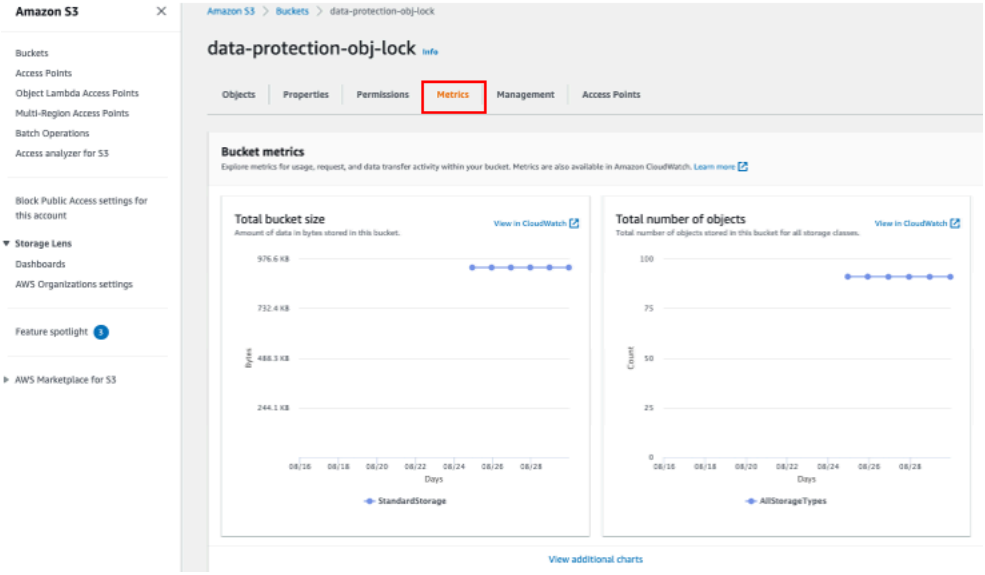

2. Analyze replication metrics

   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.

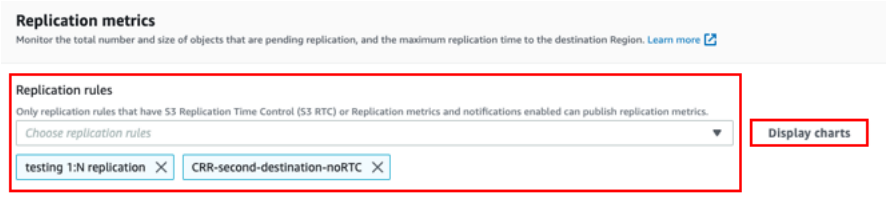

3. (Optional) Monitor with CloudWatch

   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.

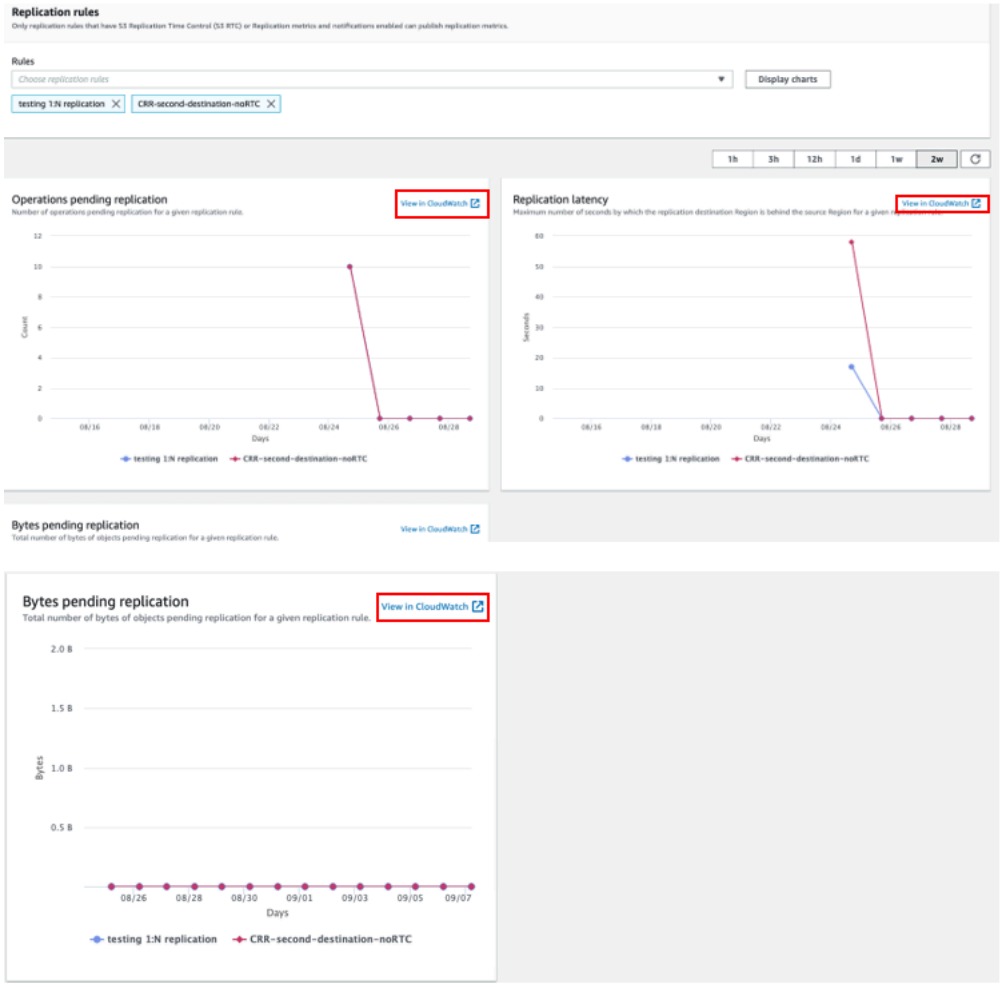

Step 7: Monitor replication status for individual objects

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.

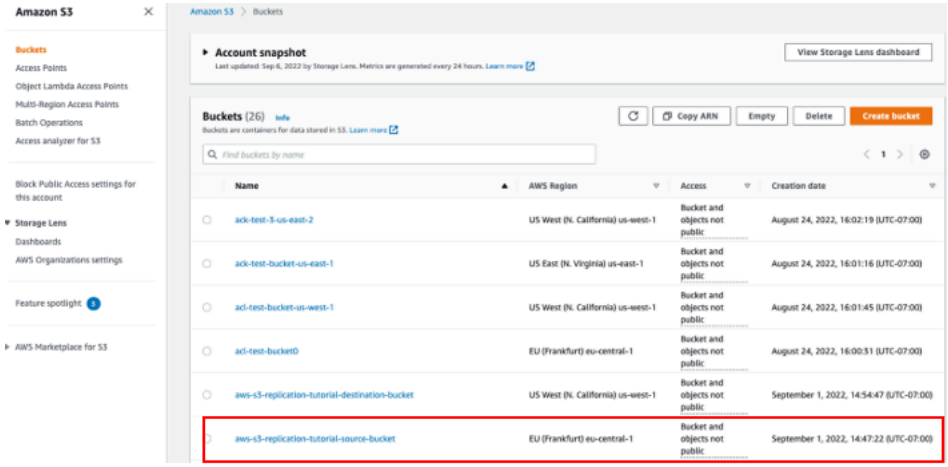

2. Select the object

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

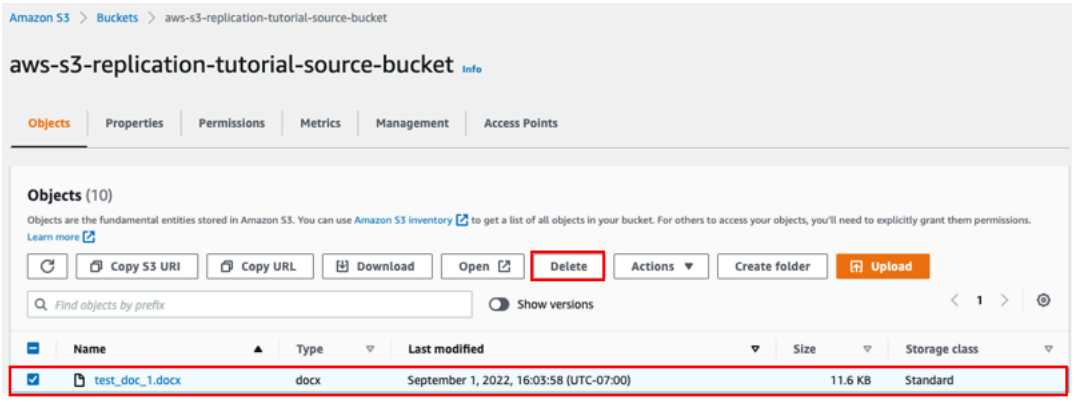

3. Delete the object

   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.

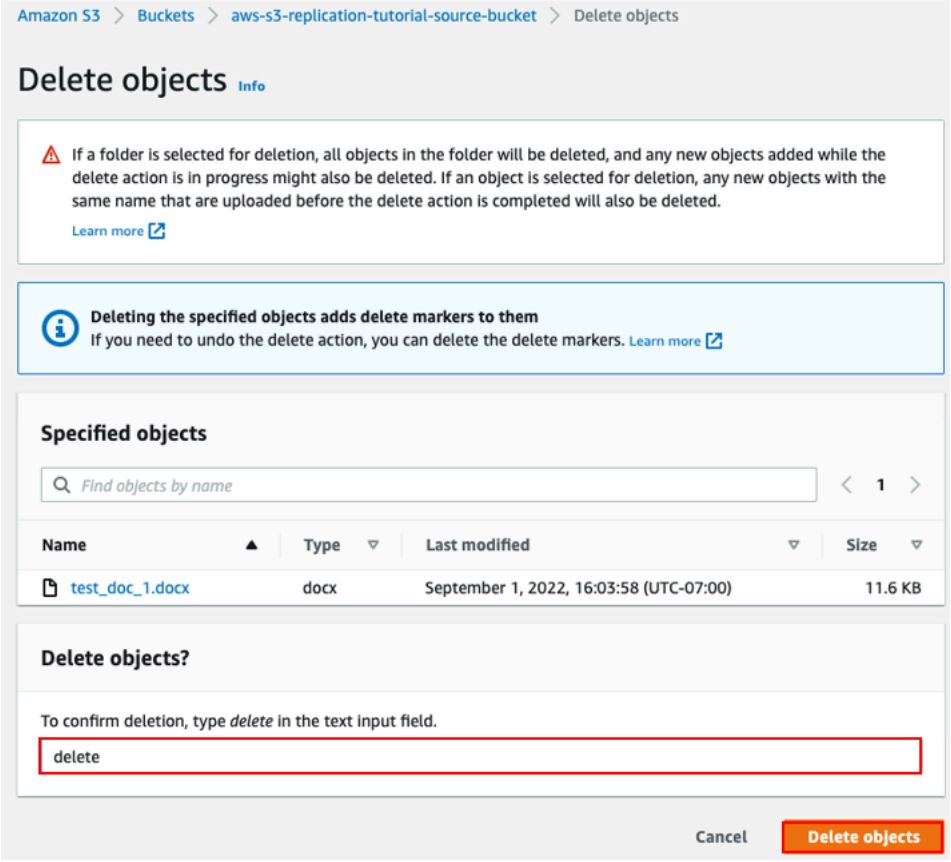

4. Confirm deletion

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

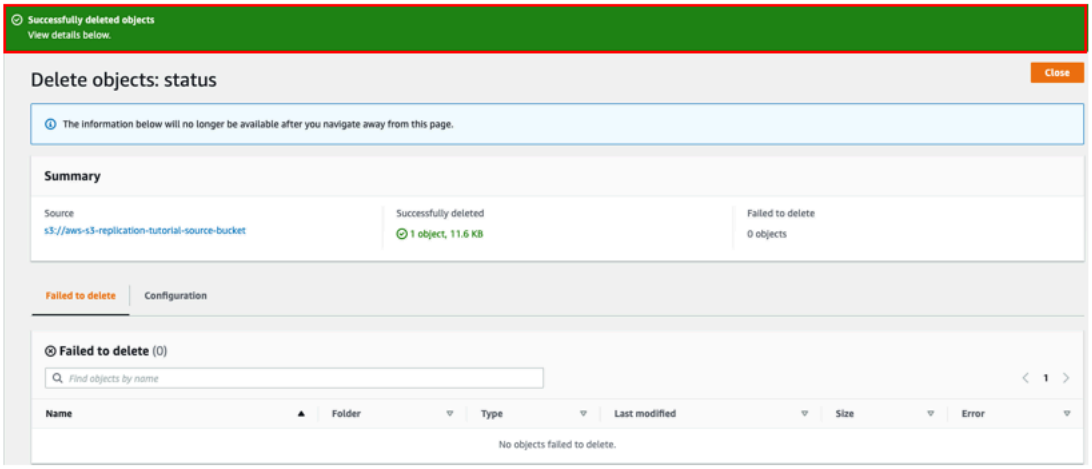

5. Delete the source bucket

   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.

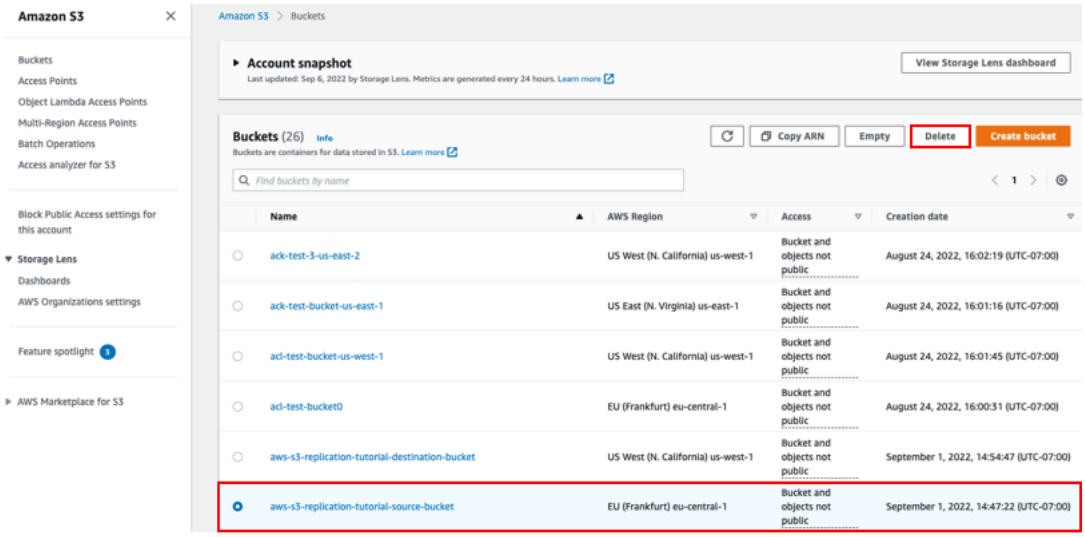

6. Confirm deletion

   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**.

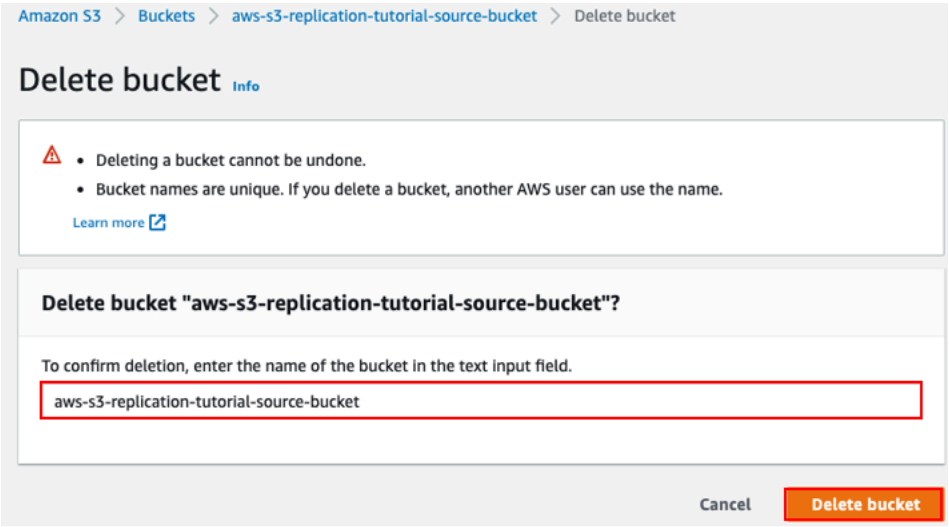

7. Delete the destination bucket

   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.

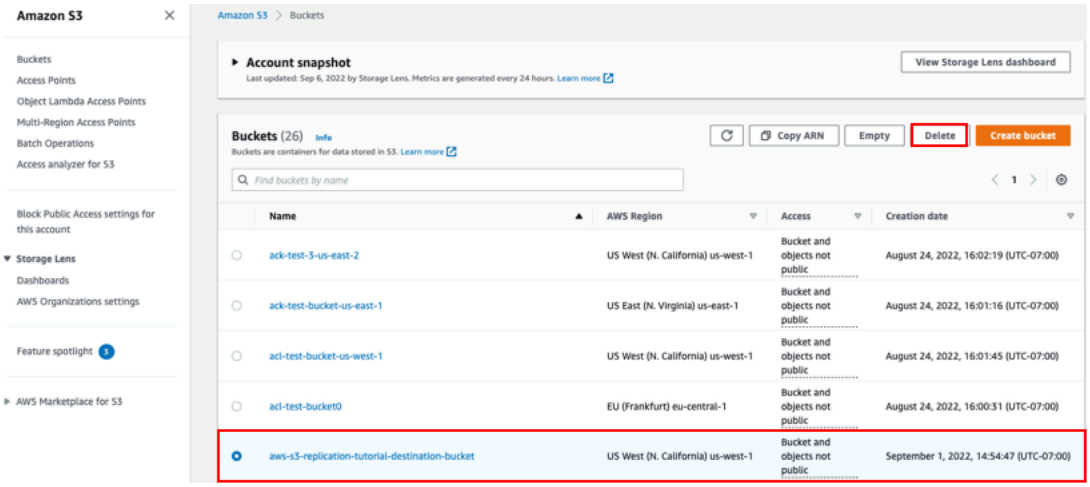

8. Confirm deletion

   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.**

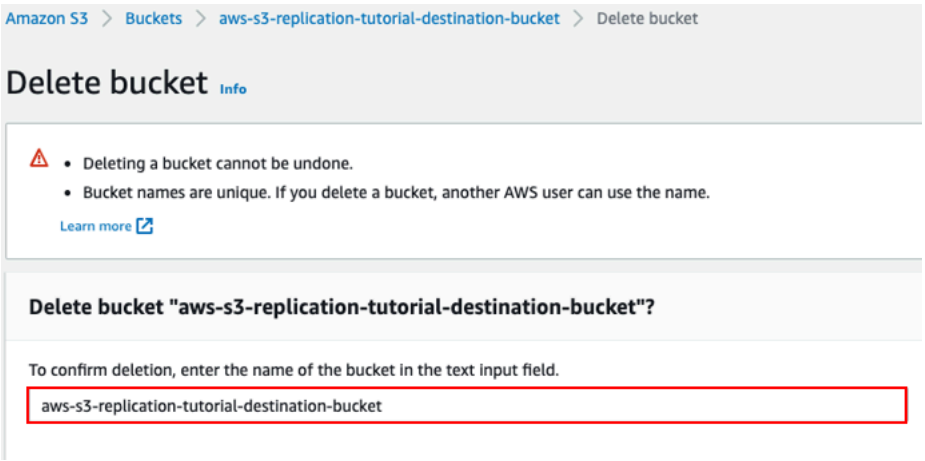

Conclusion

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.