NetApp

Mirror and backup protection on the local cluster

ONTAP 9

NetApp
January 08, 2026

Table of Contents

Mirror and backup protection on the local cluster

Create a mirror relationship for a new ONTAP S3 bucket on the local cluster

When you create new S3 buckets, you can protect them immediately to a SnapMirror S3 destination on the same cluster. You can mirror data to a bucket in a different storage VM or the same storage VM as the source.

Before you begin

- Requirements for ONTAP versions, licensing, and S3 server configuration have been completed.
- A peering relationship exists between source and destination storage VMs.
- CA Certificates are needed for the source and destination VMs. You can use self-signed CA certificates or certificates signed by an external CA vendor.

System Manager

1. If this is the first SnapMirror S3 relationship for this storage VM, verify that root user keys exist for both source and destination storage VMs and regenerate them if they do not:

   a. Click **Storage > Storage VMs** and then select the storage VM.

   b. In the **Settings** tab, click 🖊 in the S3 tile.

   c. In the **Users** tab, verify that there is an access key for the root user

   d. If there is not, click ⋮ next to **root**, then click **Regenerate Key**.
      Do not regenerate the key if one already exists.

2. Edit the storage VM to add users, and to add users to groups, in both the source and destination storage VMs:
   Click **Storage > storage VMs**, click the storage VM, click **Settings** and then click 🖊 under S3.

   See Add S3 users and groups for more information.

3. Create a SnapMirror S3 policy if you don't have an existing one and you don't want to use the default policy:

   a. Click **Protection > Overview**, and then click **Local Policy Settings**.

   b. Click → next to **Protection Policies**, then click **Add**.

      ▪ Enter the policy name and description.

      ▪ Select the policy scope, cluster or SVM

      ▪ Select **Continuous** for SnapMirror S3 relationships.

      ▪ Enter your **Throttle** and **Recovery Point Objective** values.

4. Create a bucket with SnapMirror protection:

   a. Click **Storage > Buckets** then click **Add**.

   b. Enter a name, select the storage VM, enter a size, then click **More Options**.

   c. Under **Permissions**, click **Add**. Verifying permissions is optional but recommended.

      ▪ **Principal** and **Effect** - select values corresponding to your user group settings, or accept the defaults.

      ▪ **Actions** - make sure the following values are shown:

      ```
      GetObject,PutObject,DeleteObject,ListBucket,GetBucketAcl,GetOb
      jectAcl,ListBucketMultipartUploads,ListMultipartUploadParts
      ```

      ▪ **Resources** - use the defaults (`bucketname, bucketname/*`) or other values you need

         See Manage user access to buckets for more information about these fields.

   d. Under **Protection**, check **Enable SnapMirror (ONTAP or Cloud)**. Then enter the following values:

      ▪ Destination

         ▪ **TARGET**: ONTAP System

         ▪ **CLUSTER**: Select the local cluster.

- **STORAGE VM**: Select a storage VM on the local cluster.
  - **S3 SERVER CA CERTIFICATE**: Copy and paste the contents of the source certificate.
- Source
  - **S3 SERVER CA CERTIFICATE**: Copy and paste the contents of the destination certificate.

5. Check **Use the same certificate on the destination** if you are using a certificate signed by an external CA vendor.

6. If you click **Destination Settings**, you can also enter your own values in place of the defaults for bucket name, capacity, and performance service level.

7. Click **Save**. A new bucket is created in the source storage VM, and it is mirrored to a new bucket that is created the destination storage VM.

Back up locked buckets

Beginning with ONTAP 9.14.1, you can back up locked S3 buckets and restore them as required.

When defining the protection settings for a new or existing bucket, you can enable object locking on destination buckets, provided that the source and destination clusters run ONTAP 9.14.1 or later, and that object locking is enabled on the source bucket. The object locking mode and lock retention tenure of the source bucket become applicable for the replicated objects on the destination bucket. You can also define a different lock retention period for the destination bucket in the **Destination Settings** section. This retention period is also applied to any non-locked objects replicated from the source bucket and S3 interfaces.

For information about how to enable object locking on a bucket, see Create a bucket.

CLI

1. If this is the first SnapMirror S3 relationship for this SVM, verify that root user keys exist for both source and destination SVMs and regenerate them if they do not:

   `vserver object-store-server user show`

   Verify that there is an access key for the root user. If there is not, enter:

   `vserver object-store-server user regenerate-keys -vserver` *svm_name* `-user` *root*

   Do not regenerate the key if one already exists.

2. Create buckets in both the source and destination SVMs:

   ```
   vserver object-store-server bucket create -vserver svm_name -bucket
   bucket_name [-size integer[KB|MB|GB|TB|PB]] [-comment text]
   [additional_options]
   ```

3. Add access rules to the default bucket policies in both the source and destination SVMs:

   ```
   vserver object-store-server bucket policy add-statement -vserver svm_name
   -bucket bucket_name -effect {allow|deny} -action object_store_actions
   -principal user_and_group_names -resource object_store_resources [-sid
   text] [-index integer]
   ```

```
src_cluster::> vserver object-store-server bucket policy add-
statement -bucket test-bucket -effect allow -action
GetObject,PutObject,DeleteObject,ListBucket,GetBucketAcl,GetObjectAc
l,ListBucketMultipartUploads,ListMultipartUploadParts -principal -
-resource test-bucket, test-bucket /*
```

4. Create a SnapMirror S3 policy if you don't have an existing one and you don't want to use the default policy:

```
snapmirror policy create -vserver svm_name -policy policy_name -type
continuous [-rpo integer] [-throttle throttle_type] [-comment text]
[additional_options]
```

Parameters:

   ◦ `continuous` – the only policy type for SnapMirror S3 relationships (required).

   ◦ `-rpo` – specifies the time for recovery point objective, in seconds (optional).

   ◦ `-throttle` – specifies the upper limit on throughput/bandwidth, in kilobytes/seconds (optional).

   Example

   ```
   src_cluster::> snapmirror policy create -vserver vs0 -type
   continuous -rpo 0 -policy test-policy
   ```

5. Install CA server certificates on the admin SVM:

   a. Install the CA certificate that signed the *source* S3 server's certificate on the admin SVM:
   ```
   security certificate install -type server-ca -vserver admin_svm -cert
   -name src_server_certificate
   ```

   b. Install the CA certificate that signed the *destination* S3 server's certificate on the admin SVM:
   ```
   security certificate install -type server-ca -vserver admin_svm -cert
   -name dest_server_certificate
   ```

   If you are using a certificate signed by an external CA vendor, you only need to install this certificate on the admin SVM.

   Learn more about `security certificate install` in the ONTAP command reference.

6. Create a SnapMirror S3 relationship:
   ```
   snapmirror create -source-path src_svm_name:/bucket/bucket_name
   -destination-path dest_peer_svm_name:/bucket/bucket_name, …} [-policy
   policy_name]`
   ```

   You can use a policy you created or accept the default.

Example

```
src_cluster::> snapmirror create -source-path vs0-src:/bucket/test-
bucket -destination-path vs1-dest:/vs1/bucket/test-bucket-mirror
-policy test-policy
```

7. Verify that mirroring is active:
   `snapmirror show -policy-type continuous -fields status`

Related information

- snapmirror create
- snapmirror policy create
- snapmirror show

Create a mirror relationship for an existing ONTAP S3 bucket on the local cluster

You can begin protecting existing S3 buckets on the same cluster at any time; for example, if you upgraded an S3 configuration from a release earlier than ONTAP 9.10.1. You can mirror data to a bucket in a different storage VM or the same storage VM as the source.

Before you begin

- Requirements for ONTAP versions, licensing, and S3 server configuration have been completed.
- A peering relationship exists between source and destination storage VMs.
- CA Certificates are needed for the source and destination VMs. You can use self-signed CA certificates or certificates signed by an external CA vendor.

System Manager

1. If this is the first SnapMirror S3 relationship for this storage VM, verify that root user keys exist for both source and destination storage VMs and regenerate them if they do not:

    a. Click **Storage > Storage VMs** and then select the storage VM.

    b. In the **Settings** tab, click ✏️ in the **S3** tile.

    c. In the **Users** tab, verify that there is an access key for the root user.

    d. If there is not, click ⋮ next to **root**, then click **Regenerate Key**.
       Do not regenerate the key if one already exists

2. Verify that existing users and groups are present and have the correct access in both the source and destination storage VMs:
   Select **Storage > Storage VMs,** then select the storage VM, and then **Settings** tab. Finally, locate the **S3** tile, select ✏️ , and select the **Users** tab and then the **Groups** tab to view user and group access settings.

    See Add S3 users and groups for more information.

3. Create a SnapMirror S3 policy if you don't have an existing one and you don't want to use the default policy:

    a. Click **Protection > Overview**, and then click **Local Policy Setting**.

    b. Click → next to **Protection Policies**, then click **Add**.

       ▪ Enter the policy name and description.

       ▪ Select the policy scope, cluster or SVM

       ▪ Select **Continuous** for SnapMirror S3 relationships.

       ▪ Enter your **Throttle** and **Recovery Point Objective** values.

4. Verify that the bucket access policy of the existing bucket continues to meet your needs:

    a. Click **Storage > Buckets** and then select the bucket you want to protect.

    b. In the **Permissions** tab, click ✏️ **Edit**, then click **Add** under **Permissions**.

       ▪ **Principal** and **Effect** - select values corresponding to your user group settings, or accept the defaults.

       ▪ **Actions** - make sure the following values are shown:

       ```
       GetObject,PutObject,DeleteObject,ListBucket,GetBucketAcl,GetOb
       jectAcl,ListBucketMultipartUploads,ListMultipartUploadParts
       ```

       ▪ **Resources** - use the defaults *(bucketname, bucketname/*)* or other values you need.

         See Manage user access to buckets for more information about these fields.

5. Protect an existing bucket with SnapMirror S3:

    a. Click **Storage** > **Buckets** and then select the bucket you want to protect.

    b. Click **Protect** and enter the following values:

       ▪ Destination

- **TARGET**: ONTAP System
- **CLUSTER**: Select the local cluster.
- **STORAGE VM**: Select the same or a different storage VM.
- **S3 SERVER CA CERTIFICATE**: Copy and paste the contents of the *source* certificate.
        - Source
            - **S3 SERVER CA CERTIFICATE**: Copy and paste the contents of the *destination* certificate.

6. Check **Use the same certificate on the destination** if you are using a certificate signed by an external CA vendor.

7. If you click **Destination Settings**, you can also enter your own values in place of the defaults for bucket name, capacity, and performance service level.

8. Click **Save**. The existing bucket is mirrored to a new bucket in the destination storage VM.

Back up locked buckets

Beginning with ONTAP 9.14.1, you can back up locked S3 buckets and restore them as required.

When defining the protection settings for a new or existing bucket, you can enable object locking on destination buckets, provided that the source and destination clusters run ONTAP 9.14.1 or later, and that object locking is enabled on the source bucket. The object locking mode and lock retention tenure of the source bucket become applicable for the replicated objects on the destination bucket. You can also define a different lock retention period for the destination bucket in the **Destination Settings** section. This retention period is also applied to any non-locked objects replicated from the source bucket and S3 interfaces.

For information about how to enable object locking on a bucket, see Create a bucket.

CLI

1. If this is the first SnapMirror S3 relationship for this SVM, verify that root user keys exist for both source and destination SVMs and regenerate them if they do not:
   ```
   vserver object-store-server user show
   ```

   Verify that there is an access key for the root user. If there is not, enter:
   ```
   vserver object-store-server user regenerate-keys -vserver svm_name -user
   root
   ```

   Do not regenerate the key if one already exists.

2. Create a bucket on the destination SVM to be the mirror target:

   ```
   vserver object-store-server bucket create -vserver svm_name -bucket
   dest_bucket_name [-size integer[KB|MB|GB|TB|PB]] [-comment text]
   [additional_options]
   ```

3. Verify that the access rules to the default bucket policies are correct in both the source and destination SVMs:

   ```
   vserver object-store-server bucket policy add-statement -vserver svm_name
   -bucket bucket_name -effect {allow|deny} -action object_store_actions
   -principal user_and_group_names -resource object_store_resources [-sid
   text] [-index integer]`
   ```

```
clusterA::> vserver object-store-server bucket policy add-statement
-bucket test-bucket -effect allow -action
GetObject,PutObject,DeleteObject,ListBucket,GetBucketAcl,GetObjectAc
l,ListBucketMultipartUploads,ListMultipartUploadParts -principal -
-resource test-bucket, test-bucket /*
```

4. Create a SnapMirror S3 policy if you don't have an existing one and you don't want to use the default policy:

```
snapmirror policy create -vserver svm_name -policy policy_name -type
continuous [-rpo _integer] [-throttle throttle_type] [-comment text]
[additional_options]
```

Parameters:

   ◦ `continuous` – the only policy type for SnapMirror S3 relationships (required).

   ◦ `-rpo` – specifies the time for recovery point objective, in seconds (optional).

   ◦ `-throttle` – specifies the upper limit on throughput/bandwidth, in kilobytes/seconds (optional).

   Example

   ```
   clusterA::> snapmirror policy create -vserver vs0 -type
   continuous -rpo 0 -policy test-policy
   ```

5. Install CA server certificates on the admin SVM:

   a. Install the CA certificate that signed the *source* S3 server's certificate on the admin SVM:
   `security certificate install -type server-ca -vserver admin_svm -cert -name src_server_certificate`

   b. Install the CA certificate that signed the *destination* S3 server's certificate on the admin SVM:
   `security certificate install -type server-ca -vserver admin_svm -cert -name dest_server_certificate`

   If you are using a certificate signed by an external CA vendor, you only need to install this certificate on the admin SVM.

   Learn more about `security certificate install` in the ONTAP command reference.

6. Create a SnapMirror S3 relationship:
   `snapmirror create -source-path src_svm_name:/bucket/bucket_name -destination-path dest_peer_svm_name:/bucket/bucket_name, …} [-policy policy_name]`

   You can use a policy you created or accept the default.

Example

```
src_cluster::> snapmirror create -source-path vs0-src:/bucket/test-
bucket -destination-path vs1-dest:/bucket/test-bucket-mirror -policy
test-policy
```

7. Verify that mirroring is active:
   `snapmirror show -policy-type continuous -fields status`

Related information

- [snapmirror create](#)
- [snapmirror policy create](#)
- [snapmirror show](#)

Take over from the destination ONTAP S3 bucket on the local cluster

If the data in a source bucket becomes unavailable, you can break the SnapMirror relationship to make the destination bucket writable and begin serving data.

About this task

When a takeover operation is performed, source bucket is converted to read-only and original destination bucket is converted to read-write, thereby reversing the SnapMirror S3 relationship.

When the disabled source bucket is available again, SnapMirror S3 automatically resynchronizes the contents of the two buckets. You don't need to explicitly resynchronize the relationship, as is required for standard volume SnapMirror deployments.

If the destination bucket is on a remote cluster, the takeover operation must be initiated from the remote cluster.

System Manager

Failover from the unavailable bucket and begin serving data:

1. Click **Protection > Relationships**, then select **SnapMirror S3**.

2. Click ⋮, select **Failover**, then click **Failover**.

CLI

1. Initiate a failover operation for the destination bucket:
   ```
   snapmirror failover start -destination-path svm_name:/bucket/bucket_name
   ```

2. Verify the status of the failover operation:
   ```
   snapmirror show -fields status
   ```

Example

```
clusterA::> snapmirror failover start -destination-path vs1:/bucket/test-
bucket-mirror
```

Related information

- snapmirror failover start

- snapmirror show

Restore an ONTAP S3 bucket from the destination SVM on the local cluster

When data in a source bucket is lost or corrupted, you can repopulate your data by restoring objects from a destination bucket.

About this task

You can restore the destination bucket to an existing bucket or a new bucket. The target bucket for the restore operation must be larger than the destination bucket's logical used space.

If you use an existing bucket, it must be empty when starting a restore operation. Restore does not "roll back" a bucket in time; rather, it populates an empty bucket with its previous contents.

The restore operation must be initiated from the local cluster.

System Manager

Restore the back-up data:

1. Click **Protection > Relationships**, then select the bucket.

2. Click ⋮ and then select **Restore**.

3. Under **Source**, select **Existing Bucket** (the default) or **New Bucket**.

    ◦ To restore to an **Existing Bucket** (the default), complete these actions:

        ▪ Select the cluster and storage VM to search for the existing bucket.

        ▪ Select the existing bucket.

4. Copy and paste the contents of the destination S3 server CA certificate.

    ◦ To restore to a **New Bucket**, enter the following values:

        ▪ The cluster and storage VM to host the new bucket.

        ▪ The new bucket's name, capacity, and performance service level.
          See Storage service levels for more information.

        ▪ The contents of the destination S3 server CA certificate.

5. Under **Destination**, copy and paste the contents of the source S3 server CA certificate.

6. Click **Protection** > Relationships to monitor the restore progress.

Restore locked buckets

Beginning with ONTAP 9.14.1, you can back up locked buckets and restore them as needed.

You can restore an object-locked bucket to a new or existing bucket. You can select an object-locked bucket as the destination in the following scenarios:

• **Restore to a new bucket**: When object locking is enabled, a bucket can be restored by creating a bucket that also has object locking enabled. When you restore a locked bucket, the object locking mode and retention period of the original bucket are replicated. You can also define a different lock retention period for the new bucket. This retention period is applied to non-locked objects from other sources.

• **Restore to an existing bucket**: An object-locked bucket can be restored to an existing bucket, as long as versioning and a similar object-locking mode are enabled on the existing bucket. The retention tenure of the original bucket is maintained.

• **Restore non-locked bucket**: Even if object locking is not enabled on a bucket, you can restore it to a bucket that has object locking enabled and is on the source cluster. When you restore the bucket, all the non-locked objects become locked, and the retention mode and tenure of the destination bucket become applicable to them.

CLI

1. If you are restoring objects to a new bucket, create the new bucket. For more information, see Create a cloud backup relationship for a new ONTAP S3 bucket.

2. Initiate a restore operation for the destination bucket:
   ```
   snapmirror restore -source-path svm_name:/bucket/bucket_name -destination
   -path svm_name:/bucket/bucket_name
   ```

Example

```
clusterA::> snapmirror restore -source-path vs0:/bucket/test-bucket
-destination-path vs1:/bucket/test-bucket-mirror
```

Learn more about `snapmirror restore` in the ONTAP command reference.