NetApp

Protect buckets with SnapMirror S3

ONTAP 9

NetApp
December 21, 2024

Table of Contents

Protect buckets with SnapMirror S3

SnapMirror S3 overview

Beginning with ONTAP 9.10.1, you can protect buckets in ONTAP S3 object stores using SnapMirror mirroring and backup functionality. Unlike standard SnapMirror, SnapMirror S3 enables mirroring and backups to non-NetApp destinations like AWS S3.

SnapMirror S3 supports active mirrors and backup tiers from ONTAP S3 buckets to the following destinations:

Target	Supports active mirrors and takeover?	Supports backup and restore?
ONTAP S3 • buckets in the same SVM • buckets in different SVMs on the same cluster • buckets in SVMs on different clusters	Yes	Yes
StorageGRID	No	Yes
AWS S3	No	Yes
Cloud Volumes ONTAP for Azure	Yes	Yes
Cloud Volumes ONTAP for AWS	Yes	Yes
Cloud Volumes ONTAP for Google Cloud	Yes	Yes

You can protect existing buckets on ONTAP S3 servers or you can create new buckets with data protection enabled immediately.

SnapMirror S3 requirements

- ONTAP version

  ONTAP 9.10.1 or later must be running on source and destination clusters.

- Licensing

  The following licenses are available in the ONTAP One software suite are required on ONTAP source and destination systems to provide access for:

  ◦ ONTAP S3 protocol and storage

  ◦ SnapMirror S3 to target other NetApp object store targets (ONTAP S3, StorageGRID, and Cloud Volumes ONTAP)

  ◦ SnapMirror S3 to target third-party object stores, including AWS S3 (available in the ONTAP One Compatibility bundle)

- ONTAP S3

  ◦ ONTAP S3 servers must be running source and destination SVMs.

◦ It is recommended but not required that CA certificates for TLS access are installed on systems that host S3 servers.

   ▪ The CA certificates used to sign the S3 servers' certificates must be installed on the admin storage VM of the clusters that host S3 servers.

   ▪ You can use a self-signed CA certificate or a certificate signed by an external CA vendor.

   ▪ If the source or destination storage VMs are not listening on HTTPS, it is not necessary to install CA certificates.

• Peering (for ONTAP S3 targets)

   ◦ Intercluster LIFs must be configured (for remote ONTAP targets), and the intercluster LIFs of the source and destination cluster can connect to the source and destination S3 server data LIFs.

   ◦ Source and destination clusters are peered (for remote ONTAP targets).

   ◦ Source and destination storage VMs are peered (for all ONTAP targets).

• SnapMirror policy

   ◦ An S3-specific SnapMirror policy is required for all SnapMirror S3 relationships, but you can use the same policy for multiple relationships.

   ◦ You can create your own policy or accept the default **Continuous** policy, which includes the following values:

      ▪ Throttle (upper limit on throughput/bandwidth) - unlimited.

      ▪ Time for recovery point objective: 1 hour (3600 seconds).

> ⓘ You should be aware that when two S3 buckets are in a SnapMirror relationship, if there are lifecycle policies configured so that the current version of an object expires (is deleted), the same action is replicated to the partner bucket. This is true even if the partner bucket is read-only or passive.

• Root user keys
  Storage VM root user access keys are required for SnapMirror S3 relationships; ONTAP does not assign them by default. The first time you create an SnapMirror S3 relationship, you must verify that the keys exist on both source and destination storage VMs and regenerate them if they do not. If you need to regenerate them, you must ensure that all clients and all SnapMirror object-store configurations using the access and secret key pair are updated with the new keys.

For information about S3 server configuration, see the following topics:

• Enable an S3 server on a storage VM
• About the ONTAP S3 configuration process

For information about cluster and storage VM peering, see the following topic:

• Prepare for mirroring and vaulting (System Manager, steps 1-6)
• Cluster and SVM peering (CLI)

Supported SnapMirror relationships

SnapMirror S3 supports fan-out and cascade relationships. For an overview, see Fan-out and cascade data protection deployments.

SnapMirror S3 does not support fan-in deployments (data protection relationships between multiple source buckets and a single destination bucket). SnapMirror S3 can support multiple bucket mirrors from multiple clusters to a single secondary cluster, but each source bucket must have its own destination bucket on the secondary cluster.

Control access to S3 buckets

When you create new buckets, you can control access by creating users and groups. For more information, see the following topics:

- Add S3 users and groups (System Manager)
- Create an S3 user (CLI)
- Create or modify S3 groups (CLI)

Mirror and backup protection on a remote cluster

Create a mirror relationship for a new bucket (remote cluster)

When you create new S3 buckets, you can protect them immediately to an SnapMirror S3 destination on a remote cluster.

About this task

You will need to perform tasks on both source and destination systems.

Before you begin

- Requirements for ONTAP versions, licensing, and S3 server configuration have been completed.
- A peering relationship exists between source and destination clusters, and 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. On the source cluster, create an 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**. Verifying permissions is optional but recommended.

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

   c. Under **Permissions**, click **Add**.

      - **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 remote cluster.
        - **STORAGE VM**: Select a storage VM on the remote 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]
   ```

Example

```
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. On the source SVM, create an 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:

- type `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 SVMs of the source and destination clusters:

a. On the source cluster, install the CA certificate that signed the *destination* S3 server certificate:
```
security certificate install -type server-ca -vserver src_admin_svm
-cert-name dest_server_certificate
```

b. On the destination cluster, install the CA certificate that signed the *source* S3 server certificate:
```
security certificate install -type server-ca -vserver dest_admin_svm
-cert-name src_server_certificate
```

If you are using a certificate signed by an external CA vendor, install the same certificate on the source and destination admin SVM.

See the `security certificate install` man page for details.

6. On the source SVM, create an 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`

Create a mirror relationship for an existing bucket (remote cluster)

You can begin protecting existing S3 buckets at any time; for example, if you upgraded an S3 configuration from a release earlier than ONTAP 9.10.1.

About this task

You need to perform tasks on both the source and destination clusters.

Before you begin

- Requirements for ONTAP versions, licensing, and S3 server configuration have been completed.

- A peering relationship exists between source and destination clusters, and 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.

Steps

You can create a mirror relationship using System Manager or the ONTAP CLI.

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. Select **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, 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. On the source cluster, create an SnapMirror S3 policy if you don't have an existing one and you don't want to use the default policy:

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

   b. Select ➜ next to **Protection Policies**, then click **Add**.

   c. Enter the policy name and description.

   d. Select the policy scope, either cluster or SVM.

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

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

4. Verify that the bucket access policy of the existing bucket still meets 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 protection:

   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 remote cluster.
  - **STORAGE VM**: Select a storage VM on the remote 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.

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 of 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]
   ```

```
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. On the source SVM, create an 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 certificates on the admin SVMs of source and destination clusters:

   a. On the source cluster, install the CA certificate that signed the *destination* S3 server certificate:
   ```
   security certificate install -type server-ca -vserver src_admin_svm
   -cert-name dest_server_certificate
   ```

   b. On the destination cluster, install the CA certificate that signed the *source* S3 server certificate:
   ```
   security certificate install -type server-ca -vserver dest_admin_svm
   -cert-name src_server_certificate
   ```

   If you are using a certificate signed by an external CA vendor, install the same certificate on the source and destination admin SVM.

   See the `security certificate install` man page for details.

6. On the source SVM, create an 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.

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

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

Takeover and serve data from the destination bucket (remote 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. It is not necessary to explicitly resynchronize the relationship, as is required for volume SnapMirror deployments.

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

   ```
   dest_cluster::> snapmirror failover start -destination-path
   dest_svm1:/bucket/test-bucket-mirror
   ```

Restore a bucket from the destination storage VM (remote cluster)

If 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 remote cluster.

System Manager

Restore the backed up data:

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

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.

      ▪ 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.

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

5. 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. Create the new destination bucket for restore. For more information, see Create a backup relationship for a new bucket (cloud target).

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
   ```

```
dest_cluster::> snapmirror restore -source-path src_vs1:/bucket/test-
bucket -destination-path dest_vs1:/bucket/test-bucket-mirror
```

Mirror and backup protection on the local cluster

Create a mirror relationship for a new bucket (local cluster)

When you create new S3 buckets, you can protect them immediately to an 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 an 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]
   ```

Example

```
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 an 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.

   See the `security certificate install` man page for details.

6. Create an 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.

```
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`

Create a mirror relationship for an existing bucket (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 an 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 an 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.

      See the `security certificate install` man page for details.

6. Create an 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
   ```

Takeover and serve data from the destination bucket (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
```

Restore a bucket from the destination storage VM (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 backup relationship for a new bucket (cloud target).

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
   ```

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

Backup protection with cloud targets

Requirements for cloud target relationships

Make sure that your source and target environments meet the requirements for SnapMirror S3 backup protection to cloud targets.

You must have valid account credentials with the object store provider to access the data bucket.

Intercluster LIFs and an IPspace should be configured on the cluster before the cluster can connect to a cloud object store. You should create intercluster LIFs on each node to seamlessly transfer data from the local storage to the cloud object store.

For StorageGRID targets, you need to know the following information:

- server name, expressed as a fully-qualified domain name (FQDN) or IP address
- bucket name; the bucket must already exist
- access key
- secret key

In addition, the CA certificate used to sign the StorageGRID server certificate needs to be installed on the ONTAP S3 cluster's admin storage VM using the `security certificate install command`. For more information, see Installing a CA certificate if you use StorageGRID.

For AWS S3 targets, you need to know the following information:

- server name, expressed as a fully-qualified domain name (FQDN) or IP address
- bucket name; the bucket must already exist
- access key
- secret key

The DNS server for the ONTAP cluster's admin storage VM must be able to resolve FQDNs (if used) to IP addresses.

Create a backup relationship for a new bucket (cloud target)

When you create new S3 buckets, you can back them up immediately to an SnapMirror S3 target bucket on an object store provider, which can be a StorageGRID system or an Amazon S3 deployment.

Before you begin

- You have valid account credentials and configuration information for the object store provider.

- Intercluster network interfaces and an IPspace have been configured on the source system.
- • The DNS configuration for the source storage VM must be able to resolve the target's FQDN.

System Manager

1. Edit the storage VM to add users, and to add users to groups:

   a. Click **Storage > storage VMs**, click the storage VM, click **Settings** and then click ✏️ under **S3**.

      See Add S3 users and groups for more information.

2. Add a Cloud Object Store on the source system:

   a. Click **Protection > Overview**, then select **Cloud Object Stores**.

   b. Click **Add**, then select **Amazon S3** or **StorageGRID**.

   c. Enter the following values:

      - Cloud object store name

      - URL style (path or virtual-hosted)

      - storage VM (enabled for S3)

      - Object store server name (FQDN)

      - Object store certificate

      - Access key

      - Secret key

      - Container (bucket) name

3. Create an 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,GetO
        bjectAcl,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)**, select **Cloud Storage**, then select the **Cloud Object Store**.

When you click **Save**, a new bucket is created in the source storage VM, and it is backed up to the cloud object store.

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
   ```

   Confirm 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 in the source SVM:

   ```
   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 policy:

   ```
   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]
   ```

   Example

   ```
   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 an 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:
   * `type 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. If the target is a StorageGRID system, install the StorageGRID CA server certificate on the admin SVM of the source cluster:

`security certificate install -type server-ca -vserver` *src_admin_svm* `-cert -name` *storage_grid_server_certificate*

See the `security certificate install` man page for details.

6. Define the SnapMirror S3 destination object store:

`snapmirror object-store config create -vserver` *svm_name* `-object-store-name` *target_store_name* `-usage data -provider-type {AWS_S3|SGWS} -server` *target_FQDN* `-container-name` *remote_bucket_name* `-is-ssl-enabled true -port` *port_number* `-access-key target_access_key -secret-password` *target_secret_key*

Parameters:
* `-object-store-name` – the name of the object store target on the local ONTAP system.
* `-usage` – use `data` for this workflow.
* `-provider-type` – `AWS_S3` and `SGWS` (StorageGRID) targets are supported.
* `-server` – the target server's FQDN or IP address.
* `-is-ssl-enabled` –enabling SSL is optional but recommended.

See the `snapmirror object-store config create` man page for details.

Example

```
src_cluster::> snapmirror object-store config create -vserver vs0
-object-store-name sgws-store -usage data -provider-type SGWS
-server sgws.example.com -container-name target-test-bucket -is-ssl
-enabled true -port 443 -access-key abc123 -secret-password xyz890
```

7. Create an SnapMirror S3 relationship:

`snapmirror create -source-path` *svm_name*`:/bucket/`*bucket_name* `-destination -path` *object_store_name*`:/objstore -policy` *policy_name*

Parameters:
* `-destination-path` - the object store name you created in the previous step and the fixed value `objstore`.

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

Example

```
src_cluster::> snapmirror create -source-path vs0:/bucket/test-
bucket -destination-path sgws-store:/objstore -policy test-policy
```

8. Verify that mirroring is active:

`snapmirror show -policy-type continuous -fields status`

Create a backup relationship for an existing bucket (cloud target)

You can begin backing up existing S3 buckets at any time; for example, if you upgraded an S3 configuration from a release earlier than ONTAP 9.10.1.

Before you begin

- You have valid account credentials and configuration information for the object store provider.
- Intercluster network interfaces and an IPspace have been configured on the source system.
- The DNS configuration for the source storage VM must be able to resolve the target's FQDN.

System Manager

1. Verify that the users and groups are correctly defined:
   Click **Storage > storage VMs**, click the storage VM, click **Settings** and then click ✏ under S3.

   See Add S3 users and groups for more information.

2. Create an 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**.

   c. Enter the policy name and description.

   d. Select the policy scope, cluster or SVM

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

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

3. Add a Cloud Object Store on the source system:

   a. Click **Protection > Overview**, then select **Cloud Object Store**.

   b. Click **Add**, then select **Amazon S3** or **Others** for StorageGRID Webscale.

   c. Enter the following values:

      ▪ Cloud object store name

      ▪ URL style (path or virtual-hosted)

      ▪ storage VM (enabled for S3)

      ▪ Object store server name (FQDN)

      ▪ Object store certificate

      ▪ Access key

      ▪ Secret key

      ▪ Container (bucket) name

4. Verify that the bucket access policy of the existing bucket still meets 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,GetObjectAcl`
        `,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. Back up the bucket using SnapMirror S3:

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

   b. Click **Protect**, select **Cloud Storage** under **Target**, then select the **Cloud Object Store**.

When you click **Save**, the existing bucket is backed up to the cloud object store.

CLI

1. Verify that the access rules in the default bucket policy are correct:
   ```
   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]
   ```

   Example

   ```
   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 /*
   ```

2. Create an 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:
   * `type 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
   ```

3. If the target is a StorageGRID system, install the StorageGRID CA certificate on the admin SVM of the source cluster:
   ```
   security certificate install -type server-ca -vserver src_admin_svm -cert
   -name storage_grid_server_certificate
   ```

   See the `security certificate install` man page for details.

4. Define the SnapMirror S3 destination object store:
   ```
   snapmirror object-store config create -vserver svm_name -object-store-name
   target_store_name -usage data -provider-type {AWS_S3|SGWS} -server
   target_FQDN -container-name remote_bucket_name -is-ssl-enabled true -port
   port_number -access-key target_access_key -secret-password
   target_secret_key
   ```

   Parameters:
   * `-object-store-name` – the name of the object store target on the local ONTAP system.
   * `-usage` – use `data` for this workflow.

* `-provider-type` – `AWS_S3` and `SGWS` (StorageGRID) targets are supported.
* `-server` – the target server's FQDN or IP address.
* `-is-ssl-enabled` –enabling SSL is optional but recommended.

See the `snapmirror object-store config create` man page for details.

Example

```
src_cluster::> snapmirror object-store config create -vserver vs0
-object-store-name sgws-store -usage data -provider-type SGWS
-server sgws.example.com -container-name target-test-bucket -is-ssl
-enabled true -port 443 -access-key abc123 -secret-password xyz890
```

5. Create an SnapMirror S3 relationship:
`snapmirror create -source-path` *svm_name*`:/bucket/`*bucket_name* `-destination`
`-path` *object_store_name*`:/objstore -policy` *policy_name*

Parameters:
* `-destination-path` - the object store name you created in the previous step and the fixed value
`objstore`.

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

Example

```
src_cluster::> snapmirror create -source-path vs0:/bucket/buck-evp
-destination-path sgws-store:/objstore -policy test-policy
```

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

Restore a bucket from a cloud target

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

System Manager

Restore the back-up data:

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

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.

       ▪ 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.

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

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

CLI procedure

1. Create the new destination bucket for restore. For more information, see Create a backup relationship for a bucket (cloud target).

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

    Example

    The following example restores a destination bucket to an existing bucket.
    ```
    clusterA::> snapmirror restore -source-path sgws.store:/objstore
    -destination-path vs0:/bucket/test-bucket
    ```

Modify a mirror policy

You might want to modify an S3 mirror policy; for example, if you want to adjust the RPO and throttle values.

System Manager

If you want to adjust these values, you can edit an existing protection policy.

1. Click **Protection > Relationships**, and then select the protection policy for the relationship you want to modify.
2. Click ⋮ next to the policy name, then click **Edit**.

CLI

Modify an SnapMirror S3 policy:

```
snapmirror policy modify -vserver svm_name -policy policy_name [-rpo integer]
[-throttle throttle_type] [-comment text]
```

Parameters:

* `-rpo` – specifies the time for recovery point objective, in seconds.
* `-throttle` – specifies the upper limit on throughput/bandwidth, in kilobytes/seconds.

Example

```
clusterA::> snapmirror policy modify -vserver vs0 -policy test-policy
-rpo 60
```