



Protect your data

Element Software

NetApp

November 21, 2024

This PDF was generated from https://docs.netapp.com/us-en/element-software/storage/task_data_protection_create_a_volume_snapshot.html on November 21, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Protect your data 1
 - For more information 1
 - Use volume snapshots for data protection 1
 - Perform remote replication between clusters running NetApp Element software 14
 - Use SnapMirror replication between Element and ONTAP clusters (Element UI) 28
 - Replication between NetApp Element software and ONTAP (ONTAP CLI) 40
 - Back up and restore volumes 59
 - Configure custom Protection Domains 63

Protect your data

NetApp Element software enables you to protect your data in a variety of ways with capabilities such as snapshots for individual volumes or groups of volumes, replication between clusters and volumes running on Element, and replication to ONTAP systems.

- **Snapshots**

Snapshot-only data protection replicates changed data at specific points of time to a remote cluster. Only those snapshots that are created on the source cluster are replicated. Active writes from the source volume are not.

[Use volume snapshots for data protection](#)

- **Remote replication between clusters and volumes running on Element**

You can replicate volume data synchronously or asynchronously from either cluster in a cluster pair both running on running on Element for failover and failback scenarios.

[Perform remote replication between clusters running NetApp Element software](#)

- **Replication between Element and ONTAP clusters using SnapMirror technology**

With NetApp SnapMirror technology, you can replicate snapshots that were taken using Element to ONTAP for disaster recovery purposes. In a SnapMirror relationship, Element is one endpoint and ONTAP is the other.

[Use SnapMirror replication between Element and ONTAP clusters](#)

- **Back up to and restore volumes from SolidFire, S3 or Swift object stores**

You can back up and restore volumes to other SolidFire storage, as well as secondary object stores that are compatible with Amazon S3 or OpenStack Swift.

[Back up and restore volumes to SolidFire, S3, or Swift object stores](#)

For more information

- [SolidFire and Element Software Documentation](#)
- [NetApp Element Plug-in for vCenter Server](#)

Use volume snapshots for data protection

A volume snapshot is a point-in-time copy of a volume. You can take a snapshot of a volume and use the snapshot later if you need to roll a volume back to the state it was in at the time the snapshot was created.

Snapshots are similar to volume clones. However, snapshots are simply replicas of volume metadata, so you cannot mount or write to them. Creating a volume snapshot also takes only a small amount of system resources and space, which makes snapshot creation faster than cloning.

You can take a snapshot of an individual volume or a set of volumes.

Optionally, replicate snapshots to a remote cluster and use them as a backup copy of the volume. This enables you to roll back a volume to a specific point in time by using the replicated snapshot. Alternatively, you can create a clone of a volume from a replicated snapshot.

Find more information

- [Use individual volume snapshots for data protection](#)
- [Using group snapshots for data protection task](#)
- [Scheduling a snapshot](#)

Use individual volume snapshots for data protection

A volume snapshot is a point-in-time copy of a volume. You can use an individual volume rather than a group of volumes for the snapshot.

Find more information

- [Create a volume snapshot](#)
- [Edit snapshot retention](#)
- [Deleting a snapshot](#)
- [Cloning a volume from a snapshot](#)
- [Rolling back a volume to a snapshot](#)
- [Backing up a volume snapshot to an Amazon S3 object store](#)
- [Backing up a volume snapshot to an OpenStack Swift object store](#)
- [Backing up a volume snapshot to a SolidFire cluster](#)

Create a volume snapshot

You can create a snapshot of an active volume to preserve the volume image at any point in time. You can create up to 32 snapshots for a single volume.

1. Click **Management** > **Volumes**.
2. Click the **Actions** icon for the volume you want to use for the snapshot.
3. In the resulting menu, select **Snapshot**.
4. In the **Create Snapshot of Volume** dialog box, enter the new snapshot name.
5. **Optional:** Select the **Include Snapshot in Replication When Paired** check box to ensure that the snapshot is captured in replication when the parent volume is paired.
6. To set the retention for the snapshot, select from one of the following options:
 - Click **Keep Forever** to retain the snapshot on the system indefinitely.
 - Click **Set Retention Period** and use the date spin boxes to choose a length of time for the system to retain the snapshot.
7. To take a single, immediate snapshot, perform the following steps:
 - a. Click **Take Snapshot Now**.

- b. Click **Create Snapshot**.
8. To schedule the snapshot to run at a future time, perform the following steps:
 - a. Click **Create Snapshot Schedule**.
 - b. Enter a **New Schedule Name**.
 - c. Choose a **Schedule Type** from the list.
 - d. **Optional:** Select the **Recurring Schedule** check box to repeat the scheduled snapshot periodically.
 - e. Click **Create Schedule**.

Find more information

[Schedule a snapshot](#)

Edit snapshot retention

You can change the retention period for a snapshot to control when or if the system deletes snapshots. The retention period you specify begins when you enter the new interval. When you set a retention period, you can select a period that begins at the current time (retention is not calculated from the snapshot creation time). You can specify intervals in minutes, hours, and days.

Steps

1. Click **Data Protection > Snapshots**.
2. Click the **Actions** icon for the snapshot you want to edit.
3. In the resulting menu, click **Edit**.
4. **Optional:** Select the **Include Snapshot in Replication When Paired** check box to ensure that the snapshot is captured in replication when the parent volume is paired.
5. **Optional:** Select a retention option for the snapshot:
 - Click **Keep Forever** to retain the snapshot on the system indefinitely.
 - Click **Set Retention Period** and use the date spin boxes to select a length of time for the system to retain the snapshot.
6. Click **Save Changes**.

Delete a snapshot

You can delete a volume snapshot from a storage cluster running Element software. When you delete a snapshot, the system immediately removes it.

You can delete snapshots that are being replicated from the source cluster. If a snapshot is syncing to the target cluster when you delete it, the sync replication completes and the snapshot is deleted from the source cluster. The snapshot is not deleted from the target cluster.

You can also delete snapshots that have been replicated to the target from the target cluster. The deleted snapshot is kept in a list of deleted snapshots on the target until the system detects that you have deleted the snapshot on the source cluster. When the target detects that you have deleted the source snapshot, the target stops replication of the snapshot.

When you delete a snapshot from the source cluster, the target cluster snapshot is not affected (the reverse is

also true).

1. Click **Data Protection > Snapshots**.
2. Click the **Actions** icon for the snapshot you want to delete.
3. In the resulting menu, select **Delete**.
4. Confirm the action.

Clone a volume from a snapshot

You can create a new volume from a snapshot of a volume. When you do this, the system uses the snapshot information to clone a new volume using the data contained on the volume at the time the snapshot was created. This process stores information about other snapshots of the volume in the newly created volume.

1. Click **Data Protection > Snapshots**.
2. Click the **Actions** icon for the snapshot you want to use for the volume clone.
3. In the resulting menu, click **Clone Volume From Snapshot**.
4. Enter a **Volume Name** in the **Clone Volume From Snapshot** dialog box.
5. Select a **Total Size** and size units for the new volume.
6. Select an **Access** type for the volume.
7. Select an **Account** from the list to associate with the new volume.
8. Click **Start Cloning**.

Roll back a volume to a snapshot

You can roll back a volume to a previous snapshot at any time. This reverts any changes made to the volume since the snapshot was created.

Steps

1. Click **Data Protection > Snapshots**.
2. Click the **Actions** icon for the snapshot you want to use for the volume rollback.
3. In the resulting menu, select **Rollback Volume To Snapshot**.
4. **Optional:** To save the current state of the volume before rolling back to the snapshot:
 - a. In the **Rollback To Snapshot** dialog box, select **Save volume's current state as a snapshot**.
 - b. Enter a name for the new snapshot.
5. Click **Rollback Snapshot**.

Back up a volume snapshot

You can use the integrated backup feature to back up a volume snapshot. You can back up snapshots from a SolidFire cluster to an external object store, or to another SolidFire cluster. When you back up a snapshot to an external object store, you must have a connection to the object store that allows read/write operations.

- [Back up a volume snapshot to an Amazon S3 object store](#)
- [Back up a volume snapshot to an OpenStack Swift object store](#)
- [Back up a volume snapshot to a SolidFire cluster](#)

Back up a volume snapshot to an Amazon S3 object store

You can back up SolidFire snapshots to external object stores that are compatible with Amazon S3.

1. Click **Data Protection > Snapshots**.
2. Click the **Actions** icon for the snapshot you want to back up.
3. In the resulting menu, click **Backup to**.
4. In the **Integrated Backup** dialog box under **Backup to**, select **S3**.
5. Select an option under **Data Format**:
 - **Native**: A compressed format readable only by SolidFire storage systems.
 - **Uncompressed**: An uncompressed format compatible with other systems.
6. Enter a hostname to use to access the object store in the **Hostname** field.
7. Enter an access key ID for the account in the **Access Key ID** field.
8. Enter the secret access key for the account in the **Secret Access Key** field.
9. Enter the S3 bucket in which to store the backup in the **S3 Bucket** field.
10. **Optional**: Enter a nametag to append to the prefix in the **Nametag** field.
11. Click **Start Read**.

Back up a volume snapshot to an OpenStack Swift object store

You can back up SolidFire snapshots to secondary object stores that are compatible with OpenStack Swift.

1. Click **Data Protection > Snapshots**.
2. Click the **Actions** icon for the snapshot you want to back up.
3. In the resulting menu, click **Backup to**.
4. In the **Integrated Backup** dialog box, under **Backup to**, select **Swift**.
5. Select an option under **Data Format**:
 - **Native**: A compressed format readable only by SolidFire storage systems.
 - **Uncompressed**: An uncompressed format compatible with other systems.
6. Enter a **URL** to use to access the object store.
7. Enter a **Username** for the account.
8. Enter the **Authentication Key** for the account.
9. Enter the **Container** in which to store the backup.
10. **Optional**: Enter a **Nametag**.
11. Click **Start Read**.

Back up a volume snapshot to a SolidFire cluster

You can back up volume snapshots residing on a SolidFire cluster to a remote SolidFire cluster.

Ensure that the source and target clusters are paired.

When backing up or restoring from one cluster to another, the system generates a key to be used as authentication between the clusters. This bulk volume write key allows the source cluster to authenticate with the destination cluster, providing a level of security when writing to the destination volume. As part of the backup or restore process, you need to generate a bulk volume write key from the destination volume before starting the operation.

1. On the destination cluster, click **Management > Volumes**.
2. Click the **Actions** icon for the destination volume.
3. In the resulting menu, click **Restore from**.
4. In the **Integrated Restore** dialog box under **Restore from**, select **SolidFire**.
5. Select a data format under **Data Format**:
 - **Native**: A compressed format readable only by SolidFire storage systems.
 - **Uncompressed**: An uncompressed format compatible with other systems.
6. Click **Generate Key**.
7. Copy the key from the **Bulk Volume Write Key** box to your clipboard.
8. On the source cluster, click **Data Protection > Snapshots**.
9. Click the Actions icon for the snapshot you want to use for the backup.
10. In the resulting menu, click **Backup to**.
11. In the **Integrated Backup** dialog box under **Backup to**, select **SolidFire**.
12. Select the same data format you selected earlier in the **Data Format** field.
13. Enter the management virtual IP address of the destination volume's cluster in the **Remote Cluster MVIP** field.
14. Enter the remote cluster user name in the **Remote Cluster Username** field.
15. Enter the remote cluster password in the **Remote Cluster Password** field.
16. In the **Bulk Volume Write Key** field, paste the key you generated on the destination cluster earlier.
17. Click **Start Read**.

Using group snapshots for data protection task

You can create a group snapshot of a related set of volumes to preserve a point-in-time copy of the metadata for each volume. You can use the group snapshot in the future as a backup or rollback to restore the state of the group of volumes to a previous state.

Find more information

- [Create a group snapshot](#)
- [Edit group snapshots](#)

- [Edit members of group snapshot](#)
- [Delete a group snapshot](#)
- [Roll back volumes to a group snapshot](#)
- [Clone multiple volumes](#)
- [Clone multiple volumes from a group snapshot](#)

Group snapshot details

The Group Snapshots page on the Data Protection tab provides information about the group snapshots.

- **ID**

The system-generated ID for the group snapshot.

- **UUID**

The unique ID of the group snapshot.

- **Name**

User-defined name for the group snapshot.

- **Create Time**

The time at which the group snapshot was created.

- **Status**

The current status of the snapshot. Possible values:

- **Preparing:** The snapshot is being prepared for use and is not yet writable.
- **Done:** This snapshot has finished preparation and is now usable.
- **Active:** The snapshot is the active branch.

- **# Volumes**

The number of volumes in the group.

- **Retain Until**

The day and time the snapshot will be deleted.

- **Remote Replication**

Indication of whether or not the snapshot is enabled for replication to a remote SolidFire cluster. Possible values:

- **Enabled:** The snapshot is enabled for remote replication.
- **Disabled:** The snapshot is not enabled for remote replication.

Creating a group snapshot

You can create a snapshot of a group of volumes, and you can also create a group snapshot schedule to automate group snapshots. A single group snapshot can consistently snapshot up to 32 volumes at one time.

Steps

1. Click **Management > Volumes**.
2. Use the check boxes to select multiple volumes for a group of volumes.
3. Click **Bulk Actions**.
4. Click **Group Snapshot**.
5. Enter a new group snapshot name in the Create Group Snapshot of Volumes dialog box.
6. **Optional:** Select the **Include Each Group Snapshot Member in Replication When Paired** check box to ensure that each snapshot is captured in replication when the parent volume is paired.
7. Select a retention option for the group snapshot:
 - Click **Keep Forever** to retain the snapshot on the system indefinitely.
 - Click **Set Retention Period** and use the date spin boxes to choose a length of time for the system to retain the snapshot.
8. To take a single, immediate snapshot, perform the following steps:
 - a. Click **Take Group Snapshot Now**.
 - b. Click **Create Group Snapshot**.
9. To schedule the snapshot to run at a future time, perform the following steps:
 - a. Click **Create Group Snapshot Schedule**.
 - b. Enter a **New Schedule Name**.
 - c. Select a **Schedule Type** from the list.
 - d. **Optional:** Select the **Recurring Schedule** check box to repeat the scheduled snapshot periodically.
 - e. Click **Create Schedule**.

Editing group snapshots

You can edit the replication and retention settings for existing group snapshots.

1. Click **Data Protection > Group Snapshots**.
2. Click the Actions icon for the group snapshot you want to edit.
3. In the resulting menu, select **Edit**.
4. **Optional:** To change the replication setting for the group snapshot:
 - a. Click **Edit** next to **Current Replication**.
 - b. Select the **Include Each Group Snapshot Member in Replication When Paired** check box to ensure that each snapshot is captured in replication when the parent volume is paired.
5. **Optional:** To change the retention setting for the group snapshot, select from the following options:
 - a. Click **Edit** next to **Current Retention**.
 - b. Select a retention option for the group snapshot:

- Click **Keep Forever** to retain the snapshot on the system indefinitely.
- Click **Set Retention Period** and use the date spin boxes to choose a length of time for the system to retain the snapshot.

6. Click **Save Changes**.

Deleting a group snapshot

You can delete a group snapshot from the system. When you delete the group snapshot, you can choose whether all snapshots associated with the group are deleted or retained as individual snapshots.

If you delete a volume or snapshot that is a member of a group snapshot, you can no longer roll back to the group snapshot. However, you can roll back each volume individually.

1. Click **Data Protection > Group Snapshots**.
2. Click the Actions icon for the snapshot you want to delete.
3. In the resulting menu, click **Delete**.
4. Select from one of the following options in the confirmation dialog box:
 - Click **Delete group snapshot AND all group snapshot members** to delete the group snapshot and all member snapshots.
 - Click **Retain group snapshot members as individual snapshots** to delete the group snapshot but keep all member snapshots.
5. Confirm the action.

Roll back volumes to a group snapshot

You can roll back a group of volumes at any time to a group snapshot.

When you roll back a group of volumes, all volumes in the group are restored to the state they were in at the time the group snapshot was created. Rolling back also restores volume sizes to the size recorded in the original snapshot. If the system has purged a volume, all snapshots of that volume were also deleted at the time of the purge; the system does not restore any deleted volume snapshots.

1. Click **Data Protection > Group Snapshots**.
2. Click the Actions icon for the group snapshot you want to use for the volume rollback.
3. In the resulting menu, select **Rollback Volumes To Group Snapshot**.
4. **Optional:** To save the current state of the volumes before rolling back to the snapshot:
 - a. In the **Rollback To Snapshot** dialog box, select **Save volumes' current state as a group snapshot**.
 - b. Enter a name for the new snapshot.
5. Click **Rollback Group Snapshot**.

Editing members of group snapshot

You can edit the retention settings for members of an existing group snapshot.

1. Click **Data Protection > Snapshots**.

2. Click the **Members** tab.
3. Click the Actions icon for the group snapshot member you want to edit.
4. In the resulting menu, select **Edit**.
5. To change the replication setting for the snapshot, select from the following options:
 - Click **Keep Forever** to retain the snapshot on the system indefinitely.
 - Click **Set Retention Period** and use the date spin boxes to choose a length of time for the system to retain the snapshot.
6. Click **Save Changes**.

Clone multiple volumes

You can create multiple volume clones in a single operation to create a point-in-time copy of the data on a group of volumes.

When you clone a volume, the system creates a snapshot of the volume and then creates a new volume from the data in the snapshot. You can mount and write to the new volume clone. Cloning multiple volumes is an asynchronous process and takes a variable amount of time depending on the size and number of the volumes being cloned.

Volume size and current cluster load affect the time needed to complete a cloning operation.

Steps

1. Click **Management > Volumes**.
2. Click the **Active** tab.
3. Use the check boxes to select multiple volumes, creating a group of volumes.
4. Click **Bulk Actions**.
5. Click **Clone** in the resulting menu.
6. Enter a **New Volume Name Prefix** in the **Clone Multiple Volumes** dialog box.

The prefix is applied to all volumes in the group.

7. **Optional:** Select a different account to which the clone will belong.

If you do not select an account, the system assigns the new volumes to the current volume account.

8. **Optional:** Select a different access method for the volumes in the clone.

If you do not select an access method, the system uses the current volume access.

9. Click **Start Cloning**.

Cloning multiple volumes from a group snapshot

You can clone a group of volumes from a point-in-time group snapshot. This operation requires that a group snapshot of the volumes already exist, because the group snapshot is used as the basis to create the volumes. After you create the volumes, you can use them like any other volume in the system.

Volume size and current cluster load affect the time needed to complete a cloning operation.

1. Click **Data Protection > Group Snapshots**.
2. Click the Actions icon for the group snapshot you want to use for the volume clones.
3. In the resulting menu, select **Clone Volumes From Group Snapshot**.
4. Enter a **New Volume Name Prefix** in the **Clone Volumes From Group Snapshot** dialog box.

The prefix is applied to all volumes created from the group snapshot.

5. **Optional:** Select a different account to which the clone will belong.

If you do not select an account, the system assigns the new volumes to the current volume account.

6. **Optional:** Select a different access method for the volumes in the clone.

If you do not select an access method, the system uses the current volume access.

7. Click **Start Cloning**.

Schedule a snapshot

You can protect data on a volume or a group of volumes by scheduling volume snapshots to occur at specified intervals. You can schedule either single volume snapshots or group snapshots to run automatically.

When you configure a snapshot schedule, you can choose from time intervals based on days of the week or days of the month. You can also specify the days, hours, and minutes before the next snapshot occurs. You can store the resulting snapshots on a remote storage system if the volume is being replicated.

Find more information

- [Create a snapshot schedule](#)
- [Edit a snapshot schedule](#)
- [Delete a snapshot schedule](#)
- [Copy a snapshot schedule](#)

Snapshot schedule details

On the **Data Protection > Schedules** page, you can view the following information in the list of snapshot schedules.

- **ID**

The system-generated ID for the snapshot.

- **Type**

The type of schedule. Snapshot is currently the only type supported.

- **Name**

The name given to the schedule when it was created. Snapshot schedule names can be up to 223 characters in length and contain a-z, 0-9, and dash (-) characters.

- **Frequency**

The frequency at which the schedule is run. The frequency can be set in hours and minutes, weeks, or months.

- **Recurring**

Indication of whether the schedule is to run only once or at regular intervals.

- **Manually Paused**

Indication of whether or not the schedule has been manually paused.

- **Volume IDs**

The ID of the volume the schedule will use when the schedule is run.

- **Last Run**

The last time the schedule was run.

- **Last Run Status**

The outcome of the last schedule execution. Possible values:

- Success
- Failure

Create a snapshot schedule

You can schedule a snapshot of a volume or volumes to automatically occur at specified intervals.

When you configure a snapshot schedule, you can choose from time intervals based on days of the week or days of the month. You can also create a recurring schedule and specify the days, hours, and minutes before the next snapshot occurs.

If you schedule a snapshot to run at a time period that is not divisible by 5 minutes, the snapshot will run at the next time period that is divisible by 5 minutes. For example, if you schedule a snapshot to run at 12:42:00 UTC, it will run at 12:45:00 UTC. You cannot schedule a snapshot to run at intervals of less than 5 minutes.

Beginning with Element 12.5, you can enable serial creation and select to retain the snapshots on a First-In-First-Out (FIFO) basis from the UI.

- The **Enable Serial Creation** option specifies that only one snapshot is replicated at a time. The creation of a new snapshot fails when a previous snapshot replication is still in progress. If the checkbox is not selected, a snapshot creation is allowed when another snapshot replication is still in progress.
- The **FIFO** option adds the capability to retain a consistent number of the latest snapshots. When the checkbox is selected, snapshots are retained on a FIFO basis. After the queue of FIFO snapshots reaches its maximum depth, the oldest FIFO snapshot is discarded when a new FIFO snapshot is inserted.

Steps

1. Select **Data Protection > Schedules**.
2. Select **Create Schedule**.
3. In the **Volume IDs CSV** field, enter a single volume ID or a comma-separated list of volume IDs to include in the snapshot operation.
4. Enter a new schedule name.
5. Select a schedule type and set the schedule from the options provided.
6. **Optional:** Select **Recurring Schedule** to repeat the snapshot schedule indefinitely.
7. **Optional:** Enter a name for the new snapshot in the **New Snapshot Name** field.

If you leave the field blank, the system uses the time and date of the snapshot's creation as the name.

8. **Optional:** Select the **Include Snapshots in Replication When Paired** check box to ensure that the snapshots are captured in replication when the parent volume is paired.
9. **Optional:** Select the **Enable Serial Creation** check box to ensure that only one snapshot is replicated at a time.
10. To set the retention for the snapshot, select from the following options:
 - **Optional:** Select the **FIFO (First In First out)** check box to retain a consistent number of the latest snapshots.
 - Select **Keep Forever** to retain the snapshot on the system indefinitely.
 - Select **Set Retention Period** and use the date spin boxes to choose a length of time for the system to retain the snapshot.
11. Select **Create Schedule**.

Edit a snapshot schedule

You can modify existing snapshot schedules. After modification, the next time the schedule runs it uses the updated attributes. Any snapshots created by the original schedule remain on the storage system.

Steps

1. Click **Data Protection > Schedules**.
2. Click the **Actions** icon for the schedule you want to change.
3. In the resulting menu, click **Edit**.
4. In the **Volume IDs CSV** field, modify the single volume ID or comma-separated list of volume IDs currently included in the snapshot operation.
5. To pause or resume the schedule, select from the following options:
 - To pause an active schedule, select **Yes** from the **Manually Pause Schedule** list.
 - To resume a paused schedule, select **No** from the **Manually Pause Schedule** list.
6. Enter a different name for the schedule in the **New Schedule Name** field if desired.
7. To change the schedule to run on different days of the week or month, select **Schedule Type** and change the schedule from the options provided.
8. **Optional:** Select **Recurring Schedule** to repeat the snapshot schedule indefinitely.

9. **Optional:** Enter or modify the name for the new snapshot in the **New Snapshot Name** field.

If you leave the field blank, the system uses the time and date of the snapshot's creation as the name.

10. **Optional:** Select the **Include Snapshots in Replication When Paired** check box to ensure that the snapshots are captured in replication when the parent volume is paired.

11. To change the retention setting, select from the following options:

- Click **Keep Forever** to retain the snapshot on the system indefinitely.
- Click **Set Retention Period** and use the date spin boxes to select a length of time for the system to retain the snapshot.

12. Click **Save Changes**.

Copy a snapshot schedule

You can copy a schedule and maintain its current attributes.

1. Click **Data Protection > Schedules**.
2. Click the Actions icon for the schedule you want to copy.
3. In the resulting menu, click **Make a Copy**.

The **Create Schedule** dialog box appears, populated with the current attributes of the schedule.

4. **Optional:** Enter a name and updated attributes for the new schedule.
5. Click **Create Schedule**.

Delete a snapshot schedule

You can delete a snapshot schedule. After you delete the schedule, it does not run any future scheduled snapshots. Any snapshots that were created by the schedule remain on the storage system.

1. Click **Data Protection > Schedules**.
2. Click the **Actions** icon for the schedule you want to delete.
3. In the resulting menu, click **Delete**.
4. Confirm the action.

Perform remote replication between clusters running NetApp Element software

For clusters running Element software, real-time replication enables the quick creation of remote copies of volume data. You can pair a storage cluster with up to four other storage clusters. You can replicate volume data synchronously or asynchronously from either cluster in a cluster pair for failover and failback scenarios.

The replication process includes these steps:

Replicating volume data on clusters running Element software



- Plan for replication
- Pair clusters
- Pair volumes
- Validate the replication
- Remove the replication pair

- [Plan cluster and volume pairing for real-time replication](#)
- [Pair clusters for replication](#)
- [Pair volumes](#)
- [Validate volume replication](#)
- [Delete a volume relationship after replication](#)
- [Manage volume relationships](#)

Plan cluster and volume pairing for real-time replication

Real-time remote replication requires that you pair two storage clusters running Element software, pair volumes on each cluster, and validate replication. After replication completes, you should delete the volume relationship.

What you'll need

- You must have cluster administrator privileges to one or both clusters being paired.
- All node IP addresses on both management and storage networks for paired clusters are routed to each other.
- MTU of all paired nodes must be the same and be supported end-to-end between clusters.
- Both storage clusters should have unique cluster names, MVIPs, SVIPs., and all node IP addresses.
- The difference between Element software versions on the clusters is no greater than one major version. If the difference is greater, one of the clusters must be upgraded to perform data replication.



WAN Accelerator appliances have not been qualified by NetApp for use when replicating data. These appliances can interfere with compression and deduplication if deployed between two clusters that are replicating data. Be sure to fully qualify the effects of any WAN Accelerator appliance before you deploy it in a production environment.

Find more information

- [Pair clusters for replication](#)
- [Pair volumes](#)
- [Assign a replication source and target to paired volumes](#)

Pair clusters for replication

You must pair two clusters as a first step to using real-time replication functionality. After you pair and connect two clusters, you can configure active volumes on one cluster to be continuously replicated to a second cluster, providing continuous data protection (CDP).

What you'll need

- You must have cluster administrator privileges to one or both clusters being paired.
- All node MIPs and SIPs are routed to each other.
- Less than 2000 ms of round-trip latency between clusters.
- Both storage clusters should have unique cluster names, MVIPs, SVIPs, and all node IP addresses.
- The difference between Element software versions on the clusters is no greater than one major version. If the difference is greater, one of the clusters must be upgraded to perform data replication.



Cluster pairing requires full connectivity between nodes on the management network. Replication requires connectivity between the individual nodes on the storage cluster network.

You can pair one cluster with up to four other clusters for replicating volumes. You can also pair clusters within the cluster group with each other.

Find more information

[Network port requirements](#)

Pair clusters using MVIP or a pairing key

You can pair a source and target cluster using the MVIP of the target cluster if there is cluster administrator access to both clusters. If cluster administrator access is only available on one cluster in a cluster pair, a pairing key can be used on the target cluster to complete the cluster pairing.

1. Select one of the following methods to pair clusters:
 - Pair clusters using MVIP: Use this method if there is cluster administrator access to both clusters. This method uses the MVIP of the remote cluster to pair two clusters.
 - Pair clusters using a pairing key: Use this method if there is cluster administrator access to only one of the clusters. This method generates a pairing key that can be used on the target cluster to complete the cluster pairing.

Find more information

- [Pair clusters using MVIP](#)
- [Pair clusters using a pairing key](#)

Pair clusters using MVIP

You can pair two clusters for real-time replication by using the MVIP of one cluster to establish a connection with the other cluster. Cluster administrator access on both of clusters is required to use this method. The cluster administrator user name and

password is used to authenticate cluster access before the clusters can be paired.

1. On the local cluster, select **Data Protection > Cluster Pairs**.
2. Click **Pair Cluster**.
3. Click **Start Pairing** and click **Yes** to indicate that you have access to the remote cluster.
4. Enter the remote cluster MVIP address.
5. Click **Complete pairing on remote cluster**.

In the **Authentication Required** window, enter the cluster administrator user name and password of the remote cluster.

6. On the remote cluster, select **Data Protection > Cluster Pairs**.
7. Click **Pair Cluster**.
8. Click **Complete Pairing**.
9. Click the **Complete Pairing** button.

Find more information

- [Pair clusters using a pairing key](#)
- [Pairing clusters using MVIP \(video\)](#)

Pair clusters using a pairing key

If you have cluster administrator access to a local cluster but not the remote cluster, you can pair the clusters using a pairing key. A pairing key is generated on a local cluster and then sent securely to a cluster administrator at a remote site to establish a connection and complete the cluster pairing for real-time replication.

1. On the local cluster, select **Data Protection > Cluster Pairs**.
2. Click **Pair Cluster**.
3. Click **Start Pairing** and click **No** to indicate that you do not have access to the remote cluster.
4. Click **Generate Key**.



This action generates a text key for pairing and creates an unconfigured cluster pair on the local cluster. If you do not complete the procedure, you will need to manually delete the cluster pair.

5. Copy the cluster pairing key to your clipboard.
6. Make the pairing key accessible to the cluster administrator at the remote cluster site.



The cluster pairing key contains a version of the MVIP, user name, password, and database information to permit volume connections for remote replication. This key should be treated in a secure manner and not stored in a way that would allow accidental or unsecured access to the user name or password.



Do not modify any of the characters in the pairing key. The key becomes invalid if it is modified.

7. On the remote cluster, select **Data Protection > Cluster Pairs**.
8. Click **Pair Cluster**.
9. Click **Complete Pairing** and enter the pairing key in the **Pairing Key** field (paste is the recommended method).
10. Click **Complete Pairing**.

Find more information

- [Pair clusters using MVIP](#)
- [Pairing clusters using a cluster pairing key \(video\)](#)

Validate the cluster pair connection

After the cluster pairing has completed, you might want to verify the cluster pair connection to ensure replication success.

1. On the local cluster, select **Data Protection > Cluster Pairs**.
2. In the **Cluster Pairs** window, verify that the cluster pair is connected.
3. **Optional:** Navigate back to the local cluster and the **Cluster Pairs** window and verify that the cluster pair is connected.

Pair volumes

After you have established a connection between clusters in a cluster pair, you can pair a volume on one cluster with a volume on the other cluster in the pair. When a volume pairing relationship is established, you must identify which volume is the replication target.

You can pair two volumes for real-time replication that are stored on different storage clusters in a connected cluster pair. After you pair two clusters, you can configure active volumes on one cluster to be continuously replicated to a second cluster, providing continuous data protection (CDP). You can also assign either volume to be the source or target of the replication.

Volume pairings are always one-to-one. After a volume is part of a pairing with a volume on another cluster, you cannot pair it again with any other volume.

What you'll need

- You have established a connection between clusters in a cluster pair.
- You have cluster administrator privileges to one or both clusters being paired.

Steps

1. [Create a target volume with read or write access](#)
2. [Pair volumes using a volume ID or pairing key](#)
3. [Assign a replication source and target to paired volumes](#)

Create a target volume with read or write access

The replication process involves two endpoints: the source and the target volume. When

you create the target volume, the volume is automatically set to read/write mode to accept the data during the replication.

1. Select **Management > Volumes**.
2. Click **Create Volume**.
3. In the Create a New Volume dialog box, enter the Volume Name.
4. Enter the total size of the volume, select a block size for the volume, and select the account that should have access to the volume.
5. Click **Create Volume**.
6. In the Active window, click the Actions icon for the volume.
7. Click **Edit**.
8. Change the account access level to Replication Target.
9. Click **Save Changes**.

Pair volumes using a volume ID or pairing key

The pairing process involves pairing two volumes by using either a volume ID or a pairing key.

1. Pair volumes by selecting one of the following methods:
 - Using a volume ID: Use this method if you have cluster administrator access to both clusters on which volumes are to be paired. This method uses the volume ID of the volume on the remote cluster to initiate a connection.
 - Using a pairing Key: Use this method if you have cluster administrator access to only the source cluster. This method generates a pairing key that can be used on the remote cluster to complete the volume pair.



The volume pairing key contains an encrypted version of the volume information and might contain sensitive information. Only share this key in a secure manner.

Find more information

- [Pair volumes using a volume ID](#)
- [Pair volumes using a pairing key](#)

Pair volumes using a volume ID

You can pair a volume with another volume on a remote cluster if you have cluster administrator credentials for the remote cluster.

What you'll need

- Ensure that the clusters containing the volumes are paired.
- Create a new volume on the remote cluster.



You can assign a replication source and target after the pairing process. A replication source or target can be either volume in a volume pair. You should create a target volume that contains no data and has the exact characteristics of the source volume, such as size, block size setting for the volumes (either 512e or 4k), and QoS configuration. If you assign an existing volume as the replication target, the data on that volume will be overwritten. The target volume can be greater or equal in size to the source volume, but it cannot be smaller.

- Know the target Volume ID.

Steps

1. Select **Management > Volumes**.
2. Click the **Actions** icon for the volume you want to pair.
3. Click **Pair**.
4. In the **Pair Volume** dialog box, select **Start Pairing**.
5. Select **I Do** to indicate that you have access to the remote cluster.
6. Select a **Replication Mode** from the list:
 - **Real-time (Asynchronous)**: Writes are acknowledged to the client after they are committed on the source cluster.
 - **Real-time (Synchronous)**: Writes are acknowledged to the client after they are committed on both the source and target clusters.
 - **Snapshots Only**: Only snapshots created on the source cluster are replicated. Active writes from the source volume are not replicated.
7. Select a remote cluster from the list.
8. Choose a remote volume ID.
9. Click **Start Pairing**.

The system opens a web browser tab that connects to the Element UI of the remote cluster. You might be required to log on to the remote cluster with cluster administrator credentials.

10. In the Element UI of the remote cluster, select **Complete Pairing**.
11. Confirm the details in **Confirm Volume Pairing**.
12. Click **Complete Pairing**.

After you confirm the pairing, the two clusters begin the process of connecting the volumes for pairing. During the pairing process, you can see messages in the **Volume Status** column of the **Volume Pairs** window. The volume pair displays `PausedMisconfigured` until the volume pair source and target are assigned.

After you successfully complete the pairing, it is recommended that you refresh the Volumes table to remove the **Pair** option from the **Actions** list for the paired volume. If you do not refresh the table, the **Pair** option remains available for selection. If you select the **Pair** option again, a new tab opens and because the volume is already paired, the system reports a `StartVolumePairing Failed: xVolumeAlreadyPaired` error message in the **Pair Volume** window of the Element UI page.

Find more information

- [Volume pairing messages](#)

- [Volume pairing warnings](#)
- [Assign a replication source and target to paired volumes](#)

Pair volumes using a pairing key

If you do not have cluster admin credentials for a remote cluster, you can pair a volume with another volume on a remote cluster using a pairing key.

What you'll need

- Ensure that the clusters containing the volumes are paired.
- Ensure that there is a volume on the remote cluster to use for the pairing.



You can assign a replication source and target after the pairing process. A replication source or target can be either volume in a volume pair. You should create a target volume that contains no data and has the exact characteristics of the source volume, such as size, block size setting for the volumes (either 512e or 4k), and QoS configuration. If you assign an existing volume as the replication target, the data on that volume will be overwritten. The target volume can be greater or equal in size to the source volume, but it cannot be smaller.

Steps

1. Select **Management > Volumes**.
2. Click **Actions** icon for the volume you want to pair.
3. Click **Pair**.
4. In the **Pair Volume** dialog box, select **Start Pairing**.
5. Select **I Do Not** to indicate that you do not have access to the remote cluster.
6. Select a **Replication Mode** from the list:
 - **Real-time (Asynchronous)**: Writes are acknowledged to the client after they are committed on the source cluster.
 - **Real-time (Synchronous)**: Writes are acknowledged to the client after they are committed on both the source and target clusters.
 - **Snapshots Only**: Only snapshots created on the source cluster are replicated. Active writes from the source volume are not replicated.
7. Click **Generate Key**.



This action generates a text key for pairing and creates an unconfigured volume pair on the local cluster. If you do not complete the procedure, you will need to manually delete the volume pair.

8. Copy the pairing key to your computer's clipboard.
9. Make the pairing key accessible to the cluster admin at the remote cluster site.



The volume pairing key should be treated in a secure manner and not used in a way that would allow accidental or unsecured access.



Do not modify any of the characters in the pairing key. The key becomes invalid if it is modified.

10. In the remote cluster Element UI, select **Management > Volumes**.
11. Click the Actions icon for the volume you want to pair.
12. Click **Pair**.
13. In the **Pair Volume** dialog box, select **Complete Pairing**.
14. Paste the pairing key from the other cluster into the **Pairing Key** box.
15. Click **Complete Pairing**.

After you confirm the pairing, the two clusters begin the process of connecting the volumes for pairing. During the pairing process, you can see messages in the **Volume Status** column of the **Volume Pairs** window. The volume pair displays `PausedMisconfigured` until the volume pair source and target are assigned.

After you successfully complete the pairing, it is recommended that you refresh the Volumes table to remove the **Pair** option from the **Actions** list for the paired volume. If you do not refresh the table, the **Pair** option remains available for selection. If you select the **Pair** option again, a new tab opens and because the volume is already paired, the system reports a `StartVolumePairing Failed: xVolumeAlreadyPaired` error message in the **Pair Volume** window of the Element UI page.

Find more information

- [Volume pairing messages](#)
- [Volume pairing warnings](#)
- [Assign a replication source and target to paired volumes](#)

Assign a replication source and target to paired volumes

After volumes are paired, you must assign a source volume and its replication target volume. A replication source or target can be either volume in a volume pair. You can also use this procedure to redirect data sent to a source volume to a remote target volume should the source volume become unavailable.

What you'll need

You have access to the clusters containing the source and target volumes.

Steps

1. Prepare the source volume:
 - a. From the cluster that contains the volume you want to assign as source, select **Management > Volumes**.
 - b. Click the **Actions** icon for the volume you want to assign as source and click **Edit**.
 - c. In the **Access** drop-down list, select **Read/Write**.



If you are reversing source and target assignment, this action will cause the volume pair to display the following message until a new replication target is assigned:
`PausedMisconfigured`

Changing access pauses volume replication and causes the transmission of data to cease. Be sure that you have coordinated these changes at both sites.

- d. Click **Save Changes**.
2. Prepare the target volume:
 - a. From the cluster that contains the volume you want to assign as target, select **Management > Volumes**.
 - b. Click the Actions icon for the volume you want to assign as target and click **Edit**.
 - c. In the **Access** drop-down list, select **Replication Target**.



If you assign an existing volume as the replication target, the data on that volume will be overwritten. You should use a new target volume that contains no data and has the exact characteristics of the source volume, such as size, 512e setting, and QoS configuration. The target volume can be greater or equal in size to the source volume, but it cannot be smaller.

- d. Click **Save Changes**.

Find more information

- [Pair volumes using a volume ID](#)
- [Pair volumes using a pairing key](#)

Validate volume replication

After a volume is replicated, you should ensure that the source and target volumes are active. When in an active state, volumes are paired, data is being sent from the source to the target volume, and the data is in sync.

1. From both clusters, select **Data Protection > Volume Pairs**.
2. Verify that the volume status is Active.

Find more information

[Volume pairing warnings](#)

Delete a volume relationship after replication

After replication completes and you no longer need the volume pair relationship, you can delete the volume relationship.

1. Select **Data Protection > Volume Pairs**.
2. Click the **Actions** icon for the volume pair you want to delete.
3. Click **Delete**.
4. Confirm the message.

Manage volume relationships

You can manage volume relationships in many ways, such as pausing replication, reversing volume pairing, changing the mode of replication, deleting a volume pair, or deleting a cluster pair.

Find more information

- [Pause replication](#)
- [Change the mode of replication](#)
- [Delete volume pairs](#)

Pause replication

You can manually pause replication if you need to stop I/O processing for a short time. You might want to pause replication if there is a surge in I/O processing and you want to reduce the processing load.

1. Select **Data Protection > Volume Pairs**.
2. Click the Actions icon for the volume pair.
3. Click **Edit**.
4. In the **Edit Volume Pair** pane, manually pause the replication process.



Pausing or resuming volume replication manually causes the transmission of data to cease or resume. Be sure that you have coordinated these changes at both sites.

5. Click **Save Changes**.

Change the mode of replication

You can edit volume pair properties to change the replication mode of the volume pair relationship.

1. Select **Data Protection > Volume Pairs**.
2. Click the Actions icon for the volume pair.
3. Click **Edit**.
4. In the **Edit Volume Pair** pane, select a new replication mode:
 - **Real-time (Asynchronous)**: Writes are acknowledged to the client after they are committed on the source cluster.
 - **Real-time (Synchronous)**: Writes are acknowledged to the client after they are committed on both the source and target clusters.
 - **Snapshots Only**: Only snapshots created on the source cluster are replicated. Active writes from the source volume are not replicated. **Attention**: Changing the mode of replication changes the mode immediately. Be sure that you have coordinated these changes at both sites.
5. Click **Save Changes**.

Delete volume pairs

You can delete a volume pair if you want to remove a pair association between two volumes.

1. Select **Data Protection > Volume Pairs**.
2. Click the Actions icon for the volume pair you want to delete.
3. Click **Delete**.
4. Confirm the message.

Delete a cluster pair

You can delete a cluster pair from the Element UI of either of the clusters in the pair.

1. Click **Data Protection > Cluster Pairs**.
2. Click the Actions icon for a cluster pair.
3. In the resulting menu, click **Delete**.
4. Confirm the action.
5. Perform the steps again from the second cluster in the cluster pairing.

Cluster pair details

The Cluster Pairs page on the Data Protection tab provides information about clusters that have been paired or are in the process of being paired. The system displays pairing and progress messages in the Status column.

- **ID**

A system-generated ID given to each cluster pair.

- **Remote Cluster Name**

The name of the other cluster in the pair.

- **Remote MVIP**

The management virtual IP address of the other cluster in the pair.

- **Status**

Replication status of the remote cluster

- **Replicating Volumes**

The number of volumes contained by the cluster that are paired for replication.

- **UUID**

A unique ID given to each cluster in the pair.

Volume pair details

The Volume Pairs page on the Data Protection tab provides information about volumes that have been paired or are in the process of being paired. The system displays pairing and progress messages in the Volume Status column.

- **ID**

System-generated ID for the volume.

- **Name**

The name given to the volume when it was created. Volume names can be up to 223 characters and contain a-z, 0-9, and dash (-).

- **Account**

Name of the account assigned to the volume.

- **Volume Status**

Replication status of the volume

- **Snapshot Status**

Status of the snapshot volume.

- **Mode**

The client write replication method. Possible values are as follows:

- Async
- Snapshot-Only
- Sync

- **Direction**

The direction of the volume data:

- Source volume icon (➔) indicates data is being written to a target outside the cluster.
- Target volume icon (←) indicates data is being written to the local volume from an outside source.

- **Async Delay**

Length of time since the volume was last synced with the remote cluster. If the volume is not paired, the value is null.

- **Remote Cluster**

Name of the remote cluster on which the volume resides.

- **Remote Volume ID**

Volume ID of the volume on the remote cluster.

- **Remote Volume Name**

Name given to the remote volume when it was created.

Volume pairing messages

You can view volume pairing messages during the initial pairing process from the Volume Pairs page under the Data Protection tab. These messages can display on both source and target ends of the pair in the Replicating Volumes list view.

- **PausedDisconnected**

Source replication or sync RPCs timed out. Connection to the remote cluster has been lost. Check network connections to the cluster.

- **ResumingConnected**

The remote replication sync is now active. Beginning the sync process and waiting for data.

- **ResumingRRSync**

A single helix copy of the volume metadata is being made to the paired cluster.

- **ResumingLocalSync**

A double helix copy of the volume metadata is being made to the paired cluster.

- **ResumingDataTransfer**

Data transfer has resumed.

- **Active**

Volumes are paired and data is being sent from the source to the target volume and the data is in sync.

- **Idle**

No replication activity is occurring.

Volume pairing warnings

The Volume Pairs page on the Data Protection tab provides these messages after you pair volumes. These messages can display on both source and target ends of the pair (unless otherwise indicated) in the Replicating Volumes list view.

- **PausedClusterFull**

Because the target cluster is full, source replication and bulk data transfer cannot proceed. The message displays on the source end of the pair only.

- **PausedExceededMaxSnapshotCount**

The target volume already has the maximum number of snapshots and cannot replicate additional snapshots.

- **PausedManual**

Local volume has been manually paused. It must be unpaused before replication resumes.

- **PausedManualRemote**

Remote volume is in manual paused mode. Manual intervention required to unpause the remote volume before replication resumes.

- **PausedMisconfigured**

Waiting for an active source and target. Manual intervention required to resume replication.

- **PausedQoS**

Target QoS could not sustain incoming IO. Replication auto-resumes. The message displays on the source end of the pair only.

- **PausedSlowLink**

Slow link detected and stopped replication. Replication auto-resumes. The message displays on the source end of the pair only.

- **PausedVolumeSizeMismatch**

Target volume is not the same size as the source volume.

- **PausedXCOPY**

A SCSI XCOPY command is being issued to a source volume. The command must complete before replication can resume. The message displays on the source end of the pair only.

- **StoppedMisconfigured**

A permanent configuration error has been detected. The remote volume has been purged or unpaired. No corrective action is possible; a new pairing must be established.

Use SnapMirror replication between Element and ONTAP clusters (Element UI)

You can create SnapMirror relationships from the Data Protection tab in the NetApp Element UI. SnapMirror functionality must be enabled to see this in the user interface.

IPv6 is not supported for SnapMirror replication between NetApp Element software and ONTAP clusters.

[NetApp video: SnapMirror for NetApp HCI and Element Software](#)

Systems running NetApp Element software support SnapMirror functionality to copy and restore Snapshot copies with NetApp ONTAP systems. The primary reason for using this technology is disaster recovery of NetApp HCI to ONTAP. Endpoints include ONTAP, ONTAP Select, and Cloud Volumes ONTAP. See TR-4641 NetApp HCI Data Protection.

[NetApp Technical Report 4641: NetApp HCI Data Protection](#)

Find more information

- [Building your Data Fabric with NetApp HCI, ONTAP, and Converged Infrastructure](#)
- [Replication between NetApp Element Software and ONTAP \(ONTAP CLI\)](#)

SnapMirror overview

Systems running NetApp Element software support SnapMirror functionality to copy and restore snapshots with NetApp ONTAP systems.

Systems running Element can communicate directly with SnapMirror on ONTAP systems 9.3 or higher. The NetApp Element API provides methods to enable SnapMirror functionality on clusters, volumes, and snapshots. Additionally, the Element UI includes all necessary functionality to manage SnapMirror relationships between Element software and ONTAP systems.

You can replicate ONTAP originated volumes to Element volumes in specific use cases with limited functionality. For more information, see [Replication between Element software and ONTAP \(ONTAP CLI\)](#).

Enable SnapMirror on the cluster

You must manually enable SnapMirror functionality at the cluster level through the NetApp Element UI. The system comes with SnapMirror functionality disabled by default, and it is not automatically enabled as part of a new installation or upgrade. Enabling the SnapMirror feature is a one-time configuration task.

SnapMirror can only be enabled for clusters running Element software used in conjunction with volumes on a NetApp ONTAP system. You should enable SnapMirror functionality only if your cluster is connected for use with NetApp ONTAP volumes.

What you'll need

The storage cluster must be running NetApp Element software.

Steps

1. Click **Clusters > Settings**.
2. Find the cluster-specific settings for SnapMirror.
3. Click **Enable SnapMirror**.



Enabling SnapMirror functionality permanently changes the Element software configuration. You can disable the SnapMirror feature and restore the default settings only by returning the cluster to the factory image.

4. Click **Yes** to confirm the SnapMirror configuration change.

Enable SnapMirror on the volume

You must enable SnapMirror on the volume in the Element UI. This allows replication of data to specified ONTAP volumes. This is permission from the administrator of the cluster running NetApp Element software for SnapMirror to control a volume.

What you'll need

- You have enabled SnapMirror in the Element UI for the cluster.
- A SnapMirror endpoint is available.
- The volume must be 512e block size.
- The volume is not participating in remote replication.
- The volume access type is not Replication Target.



You can also set this property when creating or cloning a volume.

Steps

1. Click **Management > Volumes**.
2. Click the **Actions** icon for the volume you want to enable SnapMirror for.
3. In the resulting menu, select **Edit**.
4. In the **Edit Volume** dialog box, select the check box **Enable SnapMirror**.
5. Click **Save Changes**.

Create a SnapMirror endpoint

You must create a SnapMirror endpoint in the NetApp Element UI before you can create a relationship.

A SnapMirror endpoint is an ONTAP cluster that serves as a replication target for a cluster running Element software. Before you create a SnapMirror relationship, you first create a SnapMirror endpoint.

You can create and manage up to four SnapMirror endpoints on a storage cluster running Element software.



If an existing endpoint was originally created using the API and credentials were not saved, you can see the endpoint in the Element UI and verify its existence, but it cannot be managed using the Element UI. This endpoint can then only be managed using the Element API.

For details about API methods, see [Manage storage with the Element API](#).

What you'll need

- You should have enabled SnapMirror in the Element UI for the storage cluster.
- You know the ONTAP credentials for the endpoint.

Steps

1. Click **Data Protection > SnapMirror Endpoints**.
2. Click **Create Endpoint**.
3. In the **Create a New Endpoint** dialog box, enter the cluster management IP address of the ONTAP system.
4. Enter the ONTAP administrator credentials associated with the endpoint.
5. Review additional details:
 - **LIFs**: Lists the ONTAP intercluster logical interfaces used to communicate with Element.
 - **Status**: Shows the current status of the SnapMirror endpoint. Possible values are: connected, disconnected, and unmanaged.

6. Click **Create Endpoint**.

Create a SnapMirror relationship

You must create a SnapMirror relationship in the NetApp Element UI.



When a volume is not yet enabled for SnapMirror and you select to create a relationship from the Element UI, SnapMirror is automatically enabled on that volume.

What you'll need

SnapMirror is enabled on the volume.

Steps

1. Click **Management > Volumes**.
2. Click the **Actions** icon for the volume that is to be a part of the relationship.
3. Click **Create a SnapMirror Relationship**.
4. In the **Create a SnapMirror Relationship** dialog box, select an endpoint from the **Endpoint** list.
5. Select if the relationship will be created using a new ONTAP volume or an existing ONTAP volume.
6. To create a new ONTAP volume in the Element UI, click **Create new volume**.
 - a. Select the **Storage Virtual Machine** for this relationship.
 - b. Select the **Aggregate** from the drop-down list.
 - c. In the **Volume Name Suffix** field, enter a suffix.



The system detects the source volume name and copies it to the **Volume Name** field. The suffix you enter appends the name.

- d. Click **Create Destination Volume**.
7. To use an existing ONTAP volume, click **Use existing volume**.
 - a. Select the **Storage Virtual Machine** for this relationship.
 - b. Select the volume that is the destination for this new relationship.
 8. In the **Relationship Details** section, select a policy. If the selected policy has keep rules, the Rules table displays the rules and associated labels.
 9. **Optional:** Select a schedule.

This determines how often the relationship creates copies.
 10. **Optional:** In the **Limit Bandwidth to** field, enter the maximum amount of bandwidth that can be consumed by data transfers associated with this relationship.
 11. Review additional details:
 - **State:** Current relationship state of the destination volume. Possible values are:
 - uninitialized: The destination volume has not been initialized.
 - snapmirrored: The destination volume has been initialized and is ready to receive SnapMirror updates.
 - broken-off: The destination volume is read/write and snapshots are present.

- **Status:** Current status of the relationship. Possible values are idle, transferring, checking, quiescing, quiesced, queued, preparing, finalizing, aborting, and breaking.
- **Lag Time:** The amount of time in seconds that the destination system lags behind the source system. The lag time must be no more than the transfer schedule interval.
- **Bandwidth Limit:** The maximum amount of bandwidth that can be consumed by data transfers associated with this relationship.
- **Last Transferred:** Timestamp of the last transferred snapshot. Click for further information.
- **Policy Name:** The name of the ONTAP SnapMirror policy for the relationship.
- **Policy Type:** Type of ONTAP SnapMirror policy selected for the relationship. Possible values are:
 - `async_mirror`
 - `mirror_vault`
- **Schedule Name:** Name of the pre-existing schedule on the ONTAP system selected for this relationship.

12. To not initialize at this time, ensure that the **Initialize** check box is not selected.



Initialization can be time-consuming. You might want to run this during off-peak hours. Initialization performs a baseline transfer; it makes a snapshot copy of the source volume, then transfers that copy and all the data blocks it references to the destination volume. You can initialize manually or use a schedule to start the initialization process (and subsequent updates) according to the schedule.

13. Click **Create Relationship**.

14. Click **Data Protection > SnapMirror Relationships** to view this new SnapMirror relationship.

SnapMirror relationship actions

You can configure a relationship from the SnapMirror Relationships page of the Data Protection tab. The options from the Actions icon are described here.

- **Edit:** Edits the policy used or schedule for the relationship.
- **Delete:** Deletes the SnapMirror relationship. This function does not delete the destination volume.
- **Initialize:** Performs the first initial baseline transfer of data to establish a new relationship.
- **Update:** Performs an on-demand update of the relationship, replicating any new data and Snapshot copies included since the last update to the destination.
- **Quiesce:** Prevents any further updates for a relationship.
- **Resume:** Resumes a relationship that is quiesced.
- **Break:** Makes the destination volume read-write and stops all current and future transfers. Determine that clients are not using the original source volume, because the reverse resync operation makes the original source volume read-only.
- **Resync:** Reestablishes a broken relationship in the same direction before the break occurred.
- **Reverse Resync:** Automates the necessary steps to create and initialize a new relationship in the opposite direction. This can be done only if the existing relationship is in a broken state. This operation will not delete the current relationship. The original source volume reverts to the most recent common Snapshot copy and resynchronizes with the destination. Any changes that are made to the original source volume since the last successful SnapMirror update are lost. Any changes that were made to, or new data written

into the current destination volume is sent back to the original source volume.

- **Abort:** Cancels a current transfer in progress. If a SnapMirror update is issued for an aborted relationship, the relationship continues with the last transfer from the last restart checkpoint that was created before the abort occurred.

SnapMirror labels

A SnapMirror label serves as a marker for transferring a specified snapshot according to the retention rules of the relationship.

Applying a label to a snapshot marks it as a target for SnapMirror replication. The role of the relationship is to enforce the rules upon data transfer by selecting the matching labeled snapshot, copying it to the destination volume, and ensuring the correct number of copies are kept. It refers to the policy to determine the keep count and the retention period. The policy can have any number of rules and each rule has a unique label. This label serves as the link between the snapshot and the retention rule.

It is the SnapMirror label that indicates which rule is applied for the selected snapshot, group snapshot, or schedule.

Add SnapMirror labels to snapshots

SnapMirror labels specify the snapshot retention policy on the SnapMirror endpoint. You can add labels to snapshots and group snapshots.

You can view available labels from an existing SnapMirror relationship dialog box or the NetApp ONTAP System Manager.



When you add a label to a group snapshot, any existing labels to individual snapshots are overwritten.

What you'll need

- SnapMirror is enabled on the cluster.
- The label you want to add already exists in ONTAP.

Steps

1. Click **Data Protection > Snapshots** or **Group Snapshots** page.
2. Click the **Actions** icon for the snapshot or group snapshot you want to add a SnapMirror label to.
3. In the **Edit Snapshot** dialog box, enter text in the **SnapMirror Label** field. The label must match a rule label in the policy applied to the SnapMirror relationship.
4. Click **Save Changes**.

Add SnapMirror labels to snapshot schedules

You can add SnapMirror labels to snapshot schedules to ensure that a SnapMirror policy is applied. You can view available labels from an existing SnapMirror relationship dialog box or the NetApp ONTAP System Manager.

What you'll need

- SnapMirror must be enabled at the cluster level.

- The label you want to add already exists in ONTAP.

Steps

1. Click **Data Protection > Schedules**.
2. Add a SnapMirror label to a schedule in one of the following ways:

Option	Steps
Creating a new schedule	<ol style="list-style-type: none"> a. Select Create Schedule. b. Enter all other relevant details. c. Select Create Schedule.
Modifying existing schedule	<ol style="list-style-type: none"> a. Click the Actions icon for the schedule you want to add a label to and select Edit. b. In the resulting dialog box, enter text in the SnapMirror Label field. c. Select Save Changes.

Find more information

[Create a snapshot schedule](#)

Disaster recovery using SnapMirror

In the event of a problem with a volume or cluster running NetApp Element software, use the SnapMirror functionality to break the relationship and failover to the destination volume.



If the original cluster has completely failed or is non-existent, contact NetApp Support for further assistance.

Perform a failover from an Element cluster

You can perform a failover from the Element cluster to make the destination volume read/write and accessible to hosts on the destination side. Before you perform a failover from the Element cluster, you must break the SnapMirror relationship.

Use the NetApp Element UI to perform the failover. If the Element UI is not available, you can also use ONTAP System Manager or ONTAP CLI to issue the break relationship command.

What you'll need

- A SnapMirror relationship exists and has at least one valid snapshot on the destination volume.
- You have a need to failover to the destination volume due to unplanned outage or planned event at the primary site.

Steps

1. In the Element UI, click **Data Protection > SnapMirror Relationships**.

2. Find the relationship with the source volume that you want to failover.
3. Click the **Actions** icon.
4. Click **Break**.
5. Confirm the action.

The volume on the destination cluster now has read-write access and can be mounted to the application hosts to resume production workloads. All SnapMirror replication is halted as a result of this action. The relationship shows a state of broken-off.

Perform a failback to Element

When the issue on the primary side has been mitigated, you must resynchronize the original source volume and fail back to NetApp Element software. The steps you perform vary depending on whether the original source volume still exists or whether you need to failback to a newly created volume.

Find more information

- [Perform a failback when source volume still exists](#)
- [Perform a failback when source volume no longer exists](#)
- [SnapMirror failback scenarios](#)

SnapMirror failback scenarios


The SnapMirror disaster recovery functionality is illustrated in two failback scenarios. These assume the original relationship has been failed over (broken).

The steps from the corresponding procedures are added for reference.




In the examples here, R1 = the original relationship in which the cluster running NetApp Element software is the original source volume (Element) and ONTAP is the original destination volume (ONTAP). R2 and R3 represent the inverse relationships created through the reverse resync operation.

The following image shows the failback scenario when the source volume still exists:



The following image shows the failback scenario when the source volume no longer exists:



Find more information

- [Perform a failback when source volume still exists](#)
- [Perform a failback when source volume no longer exists](#)

Perform a failback when source volume still exists

You can resynchronize the original source volume and fail back using the NetApp Element UI. This procedure applies to scenarios where the original source volume still exists.

1. In the Element UI, find the relationship that you broke to perform the failover.
2. Click the Actions icon and click **Reverse Resync**.
3. Confirm the action.



The Reverse Resync operation creates a new relationship in which the roles of the original source and destination volumes are reversed (this results in two relationships as the original relationship persists). Any new data from the original destination volume is transferred to the original source volume as part of the reverse resync operation. You can continue to access and write data to the active volume on the destination side, but you will need to disconnect all hosts to the source volume and perform a SnapMirror update before redirecting back to the original primary.

4. Click the Actions icon of the inverse relationship that you just created and click **Update**.

Now that you have completed the reverse resync and ensured that there are no active sessions connected to the volume on the destination side and that the latest data is on the original primary volume, you can perform the following steps to complete the failback and reactivate the original primary volume:

5. Click the Actions icon of the inverse relationship and click **Break**.
6. Click the Actions icon of the original relationship and click **Resync**.



The original primary volume can now be mounted to resume production workloads on the original primary volume. The original SnapMirror replication resumes based on the policy and schedule configured for the relationship.

7. After you confirm that the original relationship status is “snapmirrored”, click the Actions icon of the inverse relationship and click **Delete**.

Find more information

[SnapMirror failback scenarios](#)

Perform a failback when source volume no longer exists

You can resynchronize the original source volume and fail back using the NetApp Element UI. This section applies to scenarios in which the original source volume has been lost but the original cluster is still intact. For instructions about how to restore to a new cluster, see the documentation on the NetApp Support Site.

What you'll need

- You have a broken-off replication relationship between Element and ONTAP volumes.
- The Element volume is irretrievably lost.
- The original volume name shows as NOT FOUND.

Steps

1. In the Element UI, find the relationship that you broke to perform the failover.

Best Practice: Make note of the SnapMirror policy and schedule details of the original broken-off relationship. This information will be required when recreating the relationship.

2. Click the **Actions** icon and click **Reverse Resync**.
3. Confirm the action.



The Reverse Resync operation creates a new relationship in which the roles of the original source volume and the destination volume are reversed (this results in two relationships as the original relationship persists). Because the original volume no longer exists, the system creates a new Element volume with the same volume name and volume size as the original source volume. The new volume is assigned a default QoS policy called sm-recovery and is associated with a default account called sm-recovery. You will want to manually edit the account and QoS policy for all volumes that are created by SnapMirror to replace the original source volumes that were destroyed.

Data from the latest snapshot is transferred to the new volume as part of the reverse resync operation. You can continue to access and write data to the active volume on the destination side, but you will need to disconnect all hosts to the active volume and perform a SnapMirror update before reinstating the original primary relationship in a later step. After you complete the reverse resync and ensure that there are no active sessions connected to the volume on the destination side and that the latest data is on the original primary volume, continue with the following steps to complete the failback and reactivate the original primary volume:

4. Click the **Actions** icon of the inverse relationship that was created during the Reverse Resync operation and click **Break**.
5. Click the **Actions** icon of the original relationship, in which the source volume does not exist, and click **Delete**.
6. Click the **Actions** icon of the inverse relationship, which you broke in step 4, and click **Reverse Resync**.
7. This reverses the source and destination and results in a relationship with the same volume source and volume destination as the original relationship.
8. Click the **Actions** icon and **Edit** to update this relationship with the original QoS policy and schedule settings you took note of.
9. Now it is safe to delete the inverse relationship that you reverse resynced in step 6.

Find more information

[SnapMirror failback scenarios](#)

Perform a transfer or one-time migration from ONTAP to Element

Typically, when you use SnapMirror for disaster recovery from a SolidFire storage cluster running NetApp Element software to ONTAP software, Element is the source and ONTAP the destination. However, in some cases the ONTAP storage system can serve as the source and Element as the destination.

- Two scenarios exist:
 - No previous disaster recovery relationship exists. Follow all the steps in this procedure.

- Previous disaster recovery relationship does exist, but not between the volumes being used for this mitigation. In this case, follow only steps 3 and 4 below.

What you'll need

- The Element destination node must have been made accessible to ONTAP.
- The Element volume must have been enabled for SnapMirror replication.

You must specify the Element destination path in the form `hostip:/lun/<id_number>`, where `lun` is the actual string "lun" and `id_number` is the ID of the Element volume.

Steps

1. Using ONTAP, create the relationship with the Element cluster:

```
snapmirror create -source-path SVM:volume|cluster://SVM/volume
-destination-path hostip:/lun/name -type XDP -schedule schedule -policy
policy
```

```
cluster_dst::> snapmirror create -source-path svm_1:volA_dst
-destination-path 10.0.0.11:/lun/0005 -type XDP -schedule my_daily
-policy MirrorLatest
```

2. Verify that the SnapMirror relationship was created by using the ONTAP `snapmirror show` command.

See information about creating a replication relationship in the ONTAP documentation and for complete command syntax, see the ONTAP man page.

3. Using the `ElementCreateVolume` API, create the target volume and set the target volume access mode to SnapMirror:

Create an Element volume using the Element API

```
{
  "method": "CreateVolume",
  "params": {
    "name": "SMTargetVolumeTest2",
    "accountID": 1,
    "totalSize": 100000000000,
    "enable512e": true,
    "attributes": {},
    "qosPolicyID": 1,
    "enableSnapMirrorReplication": true,
    "access": "snapMirrorTarget"
  },
  "id": 1
}
```

4. Initialize the replication relationship using the ONTAP `snapmirror initialize` command:

```
snapmirror initialize -source-path hostip:/lun/name  
-destination-path SVM:volume|cluster://SVM/volume
```

Replication between NetApp Element software and ONTAP (ONTAP CLI)

Replication between NetApp Element software and ONTAP overview (ONTAP CLI)

You can ensure business continuity on an Element system by using SnapMirror to replicate snapshot copies of an Element volume to an ONTAP destination. In the event of a disaster at the Element site, you can serve data to clients from the ONTAP system, and then reactivate the Element system when service is restored.

Beginning with ONTAP 9.4, you can replicate snapshot copies of a LUN created on an ONTAP node back to an Element system. You might have created a LUN during an outage at the Element site, or you might be using a LUN to migrate data from ONTAP to Element software.

You should work with Element to ONTAP backup if the following apply:

- You want to use best practices, not explore every available option.
- You want to use the ONTAP command-line interface (CLI), not System Manager or an automated scripting tool.
- You are using iSCSI to serve data to clients.

If you require additional SnapMirror configuration or conceptual information, see [Data protection overview](#).

About replication between Element and ONTAP

Beginning with ONTAP 9.3, you can use SnapMirror to replicate snapshot copies of an Element volume to an ONTAP destination. In the event of a disaster at the Element site, you can serve data to clients from the ONTAP system, then reactivate the Element source volume when service is restored.

Beginning with ONTAP 9.4, you can replicate snapshot copies of a LUN created on an ONTAP node back to an Element system. You might have created a LUN during an outage at the Element site, or you might be using a LUN to migrate data from ONTAP to Element software.

Types of data protection relationship

SnapMirror offers two types of data protection relationship. For each type, SnapMirror creates a snapshot copy of the Element source volume before initializing or updating the relationship:

- In a *disaster recovery (DR)* data protection relationship, the destination volume contains only the snapshot copy created by SnapMirror, from which you can continue to serve data in the event of a catastrophe at the primary site.
- In a *long-term retention* data protection relationship, the destination volume contains point-in-time snapshot copies created by Element software, as well as the snapshot copy created by SnapMirror. You might want

to retain monthly snapshot copies created over a 20-year span, for example.

Default policies

The first time you invoke SnapMirror, it performs a *baseline transfer* from the source volume to the destination volume. The *SnapMirror policy* defines the contents of the baseline and any updates.

You can use a default or custom policy when you create a data protection relationship. The *policy type* determines which snapshot copies to include and how many copies to retain.

The table below shows the default policies. Use the `MirrorLatest` policy to create a traditional DR relationship. Use the `MirrorAndVault` or `Unified7year` policy to create a unified replication relationship, in which DR and long-term retention are configured on the same destination volume.

Policy	Policy Type	Update behavior
MirrorLatest	async-mirror	Transfer the snapshot copy created by SnapMirror.
MirrorAndVault	mirror-vault	Transfer the snapshot copy created by SnapMirror and any less recent snapshot copies made since the last update, provided they have SnapMirror labels “daily” or “weekly”.
Unified7year	mirror-vault	Transfer the snapshot copy created by SnapMirror and any less recent snapshot copies made since the last update, provided they have SnapMirror labels “daily”, “weekly”, or “monthly”.



For complete background information on SnapMirror policies, including guidance on which policy to use, see [Data protection overview](#).

Understanding SnapMirror labels

Every policy with the “mirror-vault” policy type must have a rule that specifies which snapshot copies to replicate. The rule “daily”, for example, indicates that only snapshot copies assigned the SnapMirror label “daily” should be replicated. You assign the SnapMirror label when you configure Element snapshot copies.

Replication from an Element source cluster to an ONTAP destination cluster

You can use SnapMirror to replicate snapshot copies of an Element volume to an ONTAP destination system. In the event of a disaster at the Element site, you can serve data to clients from the ONTAP system, then reactivate the Element source volume when service is restored.

An Element volume is roughly equivalent to an ONTAP LUN. SnapMirror creates a LUN with the name of the Element volume when a data protection relationship between Element software and ONTAP is initialized. SnapMirror replicates data to an existing LUN if the LUN meets the requirements for Element to ONTAP replication.

Replication rules are as follows:

- An ONTAP volume can contain data from one Element volume only.
- You cannot replicate data from an ONTAP volume to multiple Element volumes.

Replication from an ONTAP source cluster to an Element destination cluster

Beginning with ONTAP 9.4, you can replicate snapshot copies of a LUN created on an ONTAP system back to an Element volume:

- If a SnapMirror relationship already exists between an Element source and an ONTAP destination, a LUN created while you are serving data from the destination is automatically replicated when the source is reactivated.
- Otherwise, you must create and initialize a SnapMirror relationship between the ONTAP source cluster and the Element destination cluster.

Replication rules are as follows:

- The replication relationship must have a policy of type “async-mirror”.

Policies of type “mirror-vault” are not supported.

- Only iSCSI LUNs are supported.
- You cannot replicate more than one LUN from an ONTAP volume to an Element volume.
- You cannot replicate a LUN from an ONTAP volume to multiple Element volumes.

Prerequisites

You must have completed the following tasks before configuring a data protection relationship between Element and ONTAP:

- The Element cluster must be running NetApp Element software version 10.1 or later.
- The ONTAP cluster must be running ONTAP 9.3 or later.
- SnapMirror must have been licensed on the ONTAP cluster.
- You must have configured volumes on the Element and ONTAP clusters that are large enough to handle anticipated data transfers.
- If you are using the “mirror-vault” policy type, a SnapMirror label must have been configured for the Element snapshot copies to be replicated.



You can only perform this task in the [Element software web UI](#) or using the [API methods](#).

- You must have ensured that port 5010 is available.
- If you foresee that you might need to move a destination volume, you must have ensured that full-mesh connectivity exists between the source and destination. Every node on the Element source cluster must be able to communicate with every node on the ONTAP destination cluster.

Support details

The following table shows support details for Element to ONTAP backup.


Resource or feature	Support details
---------------------	-----------------

SnapMirror	<ul style="list-style-type: none"> • The SnapMirror restore feature is not supported. • The <code>MirrorAllSnapshots</code> and <code>XDPDefault</code> policies are not supported. • The “vault” policy type is not supported. • The system-defined rule “all_source_snapshots” is not supported. • The “mirror-vault” policy type is supported only for replication from Element software to ONTAP. Use “async-mirror” for replication from ONTAP to Element software. • The <code>-schedule</code> and <code>-prefix</code> options for <code>snapmirror policy add-rule</code> are not supported. • The <code>-preserve</code> and <code>-quick-resync</code> options for <code>snapmirror resync</code> are not supported. • Storage efficiency is not preserved. • Fan-out and cascade data protection deployments are not supported.
ONTAP	<ul style="list-style-type: none"> • ONTAP Select is supported beginning with ONTAP 9.4 and Element 10.3. • Cloud Volumes ONTAP is supported beginning with ONTAP 9.5 and Element 11.0.
Element	<ul style="list-style-type: none"> • Volume size limit is 8 TiB. • Volume block size must be 512 bytes. A 4K byte block size is not supported. • Volume size must be a multiple of 1 MiB. • Volume attributes are not preserved. • Maximum number of snapshot copies to be replicated is 30.
Network	<ul style="list-style-type: none"> • A single TCP connection is allowed per transfer. • The Element node must be specified as an IP address. DNS hostname lookup is not supported. • IPspaces are not supported.
SnapLock	SnapLock volumes are not supported.
FlexGroup	FlexGroup volumes are not supported.
SVM DR	ONTAP volumes in an SVM DR configuration are not supported.
MetroCluster	ONTAP volumes in a MetroCluster configuration are not supported.

Workflow for replication between Element and ONTAP

Whether you are replicating data from Element to ONTAP or from ONTAP to Element, you need to configure a job schedule, specify a policy, and create and initialize the relationship. You can use a default or custom policy.

The workflow assumes that you have completed the prerequisite tasks listed in [Prerequisites](#). For complete background information on SnapMirror policies, including guidance on which policy to use, see [Data protection overview](#).



Enable SnapMirror in Element software

Enable SnapMirror on the Element cluster

You must enable SnapMirror on the Element cluster before you can create a replication relationship. You can only perform this task in the Element software web UI or using the [API method](#).

Before you begin

- The Element cluster must be running NetApp Element software version 10.1 or later.
- SnapMirror can only be enabled for Element clusters used with NetApp ONTAP volumes.

About this task

The Element system comes with SnapMirror disabled by default. SnapMirror is not automatically enabled as part of a new installation or upgrade.



Once enabled, SnapMirror cannot be disabled. You can only disable the SnapMirror feature and restore the default settings by returning the cluster to the factory image.

Steps

1. Click **Clusters > Settings**.
2. Find the cluster-specific settings for SnapMirror.
3. Click **Enable SnapMirror**.

Enable SnapMirror on the Element source volume

You must enable SnapMirror on the Element source volume before you can create a replication relationship. You can only perform this task in the Element software web UI or using the [ModifyVolume](#) and [ModifyVolumes](#) API methods.


Before you begin

- You must have enabled SnapMirror on the Element cluster.
- The volume block size must be 512 bytes.
- The volume must not be participating in Element remote replication.
- The volume access type must not be “Replication Target”.

About this task

The procedure below assumes the volume already exists. You can also enable SnapMirror when you create or clone a volume.

Steps

1. Select **Management > Volumes**.
2. Select the  button for the volume.
3. In the drop-down menu, select **Edit**.
4. In the **Edit Volume** dialog, select **Enable SnapMirror**.
5. Select **Save Changes**.

Create a SnapMirror endpoint

You must create a SnapMirror endpoint before you can create a replication relationship. You can only perform this task in the Element software web UI or using the [SnapMirror API methods](#).

Before you begin

You must have enabled SnapMirror on the Element cluster.

Steps

1. Click **Data Protection > SnapMirror Endpoints**.
2. Click **Create Endpoint**.
3. In the **Create a New Endpoint** dialog, enter the ONTAP cluster management IP address.
4. Enter the user ID and password of the ONTAP cluster administrator.
5. Click **Create Endpoint**.

Configure a replication relationship

Create a replication job schedule

Whether you are replicating data from Element to ONTAP or from ONTAP to Element, you need to configure a job schedule, specify a policy, and create and initialize the relationship. You can use a default or custom policy.

You can use the `job schedule cron create` command to create a replication job schedule. The job schedule determines when SnapMirror automatically updates the data protection relationship to which the schedule is assigned.

About this task

You assign a job schedule when you create a data protection relationship. If you do not assign a job schedule, you must update the relationship manually.

Step

1. Create a job schedule:

```
job schedule cron create -name job_name -month month -dayofweek day_of_week
-day day_of_month -hour hour -minute minute
```

For `-month`, `-dayofweek`, and `-hour`, you can specify `all` to run the job every month, day of the week, and hour, respectively.

Beginning with ONTAP 9.10.1, you can include the Vserver for your job schedule:

```
job schedule cron create -name job_name -vserver Vserver_name -month month
-dayofweek day_of_week -day day_of_month -hour hour -minute minute
```

The following example creates a job schedule named `my_weekly` that runs on Saturdays at 3:00 a.m.:


```
cluster_dst::> job schedule cron create -name my_weekly -dayofweek
"Saturday" -hour 3 -minute 0
```

Customize a replication policy

Create a custom replication policy

You can use a default or custom policy when you create a replication relationship. For a custom unified replication policy, you must define one or more *rules* that determine which snapshot copies are transferred during initialization and update.

You can create a custom replication policy if the default policy for a relationship is not suitable. You might want to compress data in a network transfer, for example, or modify the number of attempts SnapMirror makes to transfer snapshot copies.

About this task

The *policy type* of the replication policy determines the type of relationship it supports. The table below shows the available policy types.

Policy type	Relationship type
async-mirror	SnapMirror DR
mirror-vault	Unified replication

Step

1. Create a custom replication policy:

```
snapmirror policy create -vserver SVM -policy policy -type async-
mirror|mirror-vault -comment comment -tries transfer_tries -transfer-priority
low|normal -is-network-compression-enabled true|false
```

For complete command syntax, see the man page.

Beginning with ONTAP 9.5, you can specify the schedule for creating a common snapshot copy schedule for SnapMirror Synchronous relationships by using the `-common-snapshot-schedule` parameter. By default, the common snapshot copy schedule for SnapMirror Synchronous relationships is one hour. You can specify a value from 30 minutes to two hours for the snapshot copy schedule for SnapMirror Synchronous relationships.

The following example creates a custom replication policy for SnapMirror DR that enables network compression for data transfers:

```
cluster_dst::> snapmirror policy create -vserver svml -policy
DR_compressed -type async-mirror -comment "DR with network compression
enabled" -is-network-compression-enabled true
```

The following example creates a custom replication policy for unified replication:

```
cluster_dst::> snapmirror policy create -vserver svml -policy my_unified
-type mirror-vault
```

After you finish

For “mirror-vault” policy types, you must define rules that determine which snapshot copies are transferred during initialization and update.

Use the `snapmirror policy show` command to verify that the SnapMirror policy was created. For complete command syntax, see the man page.

Define a rule for a policy

For custom policies with the “mirror-vault” policy type, you must define at least one rule that determines which snapshot copies are transferred during initialization and update. You can also define rules for default policies with the “mirror-vault” policy type.

About this task

Every policy with the “mirror-vault” policy type must have a rule that specifies which snapshot copies to replicate. The rule “bi-monthly”, for example, indicates that only snapshot copies assigned the SnapMirror label “bi-monthly” should be replicated. You assign the SnapMirror label when you configure Element snapshot copies.

Each policy type is associated with one or more system-defined rules. These rules are automatically assigned to a policy when you specify its policy type. The table below shows the system-defined rules.

System-defined rule	Used in policy types	Result
sm_created	async-mirror, mirror-vault	A snapshot copy created by SnapMirror is transferred on initialization and update.
daily	mirror-vault	New snapshot copies on the source with the SnapMirror label “daily” are transferred on initialization and update.
weekly	mirror-vault	New snapshot copies on the source with the SnapMirror label “weekly” are transferred on initialization and update.
monthly	mirror-vault	New snapshot copies on the source with the SnapMirror label “monthly” are transferred on initialization and update.

You can specify additional rules as needed, for default or custom policies. For example:

- For the default `MirrorAndVault` policy, you might create a rule called “bi-monthly” to match snapshot copies on the source with the “bi-monthly” `SnapMirror` label.
- For a custom policy with the “mirror-vault” policy type, you might create a rule called “bi-weekly” to match snapshot copies on the source with the “bi-weekly” `SnapMirror` label.

Step

1. Define a rule for a policy:

```
snapmirror policy add-rule -vserver SVM -policy policy_for_rule -snapmirror
-label snapmirror-label -keep retention_count
```

For complete command syntax, see the man page.

The following example adds a rule with the `SnapMirror` label `bi-monthly` to the default `MirrorAndVault` policy:

```
cluster_dst::> snapmirror policy add-rule -vserver svm1 -policy
MirrorAndVault -snapmirror-label bi-monthly -keep 6
```

The following example adds a rule with the `SnapMirror` label `bi-weekly` to the custom `my_snapvault` policy:

```
cluster_dst::> snapmirror policy add-rule -vserver svm1 -policy
my_snapvault -snapmirror-label bi-weekly -keep 26
```

The following example adds a rule with the `SnapMirror` label `app_consistent` to the custom `Sync` policy:

```
cluster_dst::> snapmirror policy add-rule -vserver svm1 -policy Sync
-snapmirror-label app_consistent -keep 1
```

You can then replicate snapshot copies from the source cluster that match this `SnapMirror` label:

```
cluster_src::> snapshot create -vserver vs1 -volume voll -snapshot
snapshot1 -snapmirror-label app_consistent
```

Create a replication relationship

Create a relationship from an Element source to an ONTAP destination

The relationship between the source volume in primary storage and the destination volume in secondary storage is called a *data protection relationship*. You can use the `snapmirror create` command to create a data protection relationship from an Element source to an ONTAP destination, or from an ONTAP source to an Element destination.

You can use SnapMirror to replicate snapshot copies of an Element volume to an ONTAP destination system. In the event of a disaster at the Element site, you can serve data to clients from the ONTAP system, then reactivate the Element source volume when service is restored.

Before you begin

- The Element node containing the volume to be replicated must have been made accessible to ONTAP.
- The Element volume must have been enabled for SnapMirror replication.
- If you are using the “mirror-vault” policy type, a SnapMirror label must have been configured for the Element snapshot copies to be replicated.



You can only perform this task in the [Element software web UI](#) or using the [API methods](#).

About this task

You must specify the Element source path in the form `<hostip:>/lun/<name>`, where “lun” is the actual string “lun” and `name` is the name of the Element volume.

An Element volume is roughly equivalent to an ONTAP LUN. SnapMirror creates a LUN with the name of the Element volume when a data protection relationship between Element software and ONTAP is initialized. SnapMirror replicates data to an existing LUN if the LUN meets the requirements for replicating from Element software to ONTAP.

Replication rules are as follows:

- An ONTAP volume can contain data from one Element volume only.
- You cannot replicate data from an ONTAP volume to multiple Element volumes.

In ONTAP 9.3 and earlier, a destination volume can contain up to 251 snapshot copies. In ONTAP 9.4 and later, a destination volume can contain up to 1019 snapshot copies.

Step

1. From the destination cluster, create a replication relationship from an Element source to an ONTAP destination:

```
snapmirror create -source-path <hostip:>/lun/<name> -destination-path  
<SVM:volume>|<cluster://SVM/volume> -type XDP -schedule schedule -policy  
<policy>
```

For complete command syntax, see the man page.

The following example creates a SnapMirror DR relationship using the default `MirrorLatest` policy:

```
cluster_dst::> snapmirror create -source-path 10.0.0.11:/lun/0005  
-destination-path svm_backup:volA_dst -type XDP -schedule my_daily  
-policy MirrorLatest
```

The following example creates a unified replication relationship using the default `MirrorAndVault` policy:

```
cluster_dst:> snapmirror create -source-path 10.0.0.11:/lun/0005
-destination-path svm_backup:volA_dst -type XDP -schedule my_daily
-policy MirrorAndVault
```

The following example creates a unified replication relationship using the `Unified7year` policy:

```
cluster_dst::> snapmirror create -source-path 10.0.0.11:/lun/0005
-destination-path svm_backup:volA_dst -type XDP -schedule my_daily
-policy Unified7year
```

The following example creates a unified replication relationship using the custom `my_unified` policy:

```
cluster_dst::> snapmirror create -source-path 10.0.0.11:/lun/0005
-destination-path svm_backup:volA_dst -type XDP -schedule my_daily
-policy my_unified
```

After you finish

Use the `snapmirror show` command to verify that the SnapMirror relationship was created. For complete command syntax, see the man page.

Create a relationship from an ONTAP source to an Element destination

Beginning with ONTAP 9.4, you can use SnapMirror to replicate snapshot copies of a LUN created on an ONTAP source back to an Element destination. You might be using the LUN to migrate data from ONTAP to Element software.

Before you begin

- The Element destination node must have been made accessible to ONTAP.
- The Element volume must have been enabled for SnapMirror replication.

About this task

You must specify the Element destination path in the form `<hostip:>/lun/<name>`, where “lun” is the actual string “lun” and `name` is the name of the Element volume.

Replication rules are as follows:

- The replication relationship must have a policy of type “`async-mirror`”.

You can use a default or custom policy.

- Only iSCSI LUNs are supported.
- You cannot replicate more than one LUN from an ONTAP volume to an Element volume.
- You cannot replicate a LUN from an ONTAP volume to multiple Element volumes.

Step

1. Create a replication relationship from an ONTAP source to an Element destination:

```
snapmirror create -source-path <SVM:volume>|<cluster://SVM/volume>  
-destination-path <hostip:>/lun/<name> -type XDP -schedule schedule -policy  
<policy>
```

For complete command syntax, see the man page.

The following example creates a SnapMirror DR relationship using the default `MirrorLatest` policy:

```
cluster_dst::> snapmirror create -source-path svm_1:volA_dst  
-destination-path 10.0.0.11:/lun/0005 -type XDP -schedule my_daily  
-policy MirrorLatest
```

The following example creates a SnapMirror DR relationship using the custom `my_mirror` policy:

```
cluster_dst::> snapmirror create -source-path svm_1:volA_dst  
-destination-path 10.0.0.11:/lun/0005 -type XDP -schedule my_daily  
-policy my_mirror
```

After you finish

Use the `snapmirror show` command to verify that the SnapMirror relationship was created. For complete command syntax, see the man page.

Initialize a replication relationship

For all relationship types, initialization performs a *baseline transfer*: it makes a snapshot copy of the source volume, then transfers that copy and all the data blocks it references to the destination volume.

Before you begin

- The Element node containing the volume to be replicated must have been made accessible to ONTAP.
- The Element volume must have been enabled for SnapMirror replication.
- If you are using the “mirror-vault” policy type, a SnapMirror label must have been configured for the Element snapshot copies to be replicated.



You can only perform this task in the [Element software web UI](#) or using the [API methods](#).

About this task

You must specify the Element source path in the form `<hostip:>/lun/<name>`, where “lun” is the actual string “lun” and `name` is the name of the Element volume.

Initialization can be time-consuming. You might want to run the baseline transfer in off-peak hours.

If initialization of a relationship from an ONTAP source to an Element destination fails for any reason, it will continue to fail even after you have corrected the problem (an invalid LUN name, for example). The workaround is as follows:



1. Delete the relationship.
2. Delete the Element destination volume.
3. Create a new Element destination volume.
4. Create and initialize a new relationship from the ONTAP source to the Element destination volume.

Step

1. Initialize a replication relationship:

```
snapmirror initialize -source-path <hostip:>/lun/<name> -destination-path  
<SVM:volume|cluster://SVM/volume>
```

For complete command syntax, see the man page.

The following example initializes the relationship between the source volume 0005 at IP address 10.0.0.11 and the destination volume `volA_dst` on `svm_backup`:

```
cluster_dst::> snapmirror initialize -source-path 10.0.0.11:/lun/0005  
-destination-path svm_backup:volA_dst
```

Serve data from a SnapMirror DR destination volume

Make the destination volume writeable

When disaster disables the primary site for a SnapMirror DR relationship, you can serve data from the destination volume with minimal disruption. You can reactivate the source volume when service is restored at the primary site.

You need to make the destination volume writeable before you can serve data from the volume to clients. You can use the `snapmirror quiesce` command to stop scheduled transfers to the destination, the `snapmirror abort` command to stop ongoing transfers, and the `snapmirror break` command to make the destination writeable.

About this task

You must specify the Element source path in the form `<hostip:>/lun/<name>`, where “lun” is the actual string “lun” and `name` is the name of the Element volume.

Steps

1. Stop scheduled transfers to the destination:

```
snapmirror quiesce -source-path <hostip:>/lun/<name> -destination-path  
<SVM:volume>|<cluster://SVM/volume>
```

For complete command syntax, see the man page.

The following example stops scheduled transfers between the source volume 0005 at IP address 10.0.0.11 and the destination volume `volA_dst` on `svm_backup`:

```
cluster_dst::> snapmirror quiesce -source-path 10.0.0.11:/lun/0005
-destination-path svm_backup:volA_dst
```

2. Stop ongoing transfers to the destination:

```
snapmirror abort -source-path <hostip:>/lun/<name> -destination-path
<SVM:volume>|<cluster://SVM/volume>
```

For complete command syntax, see the man page.

The following example stops ongoing transfers between the source volume 0005 at IP address 10.0.0.11 and the destination volume `volA_dst` on `svm_backup`:

```
cluster_dst::> snapmirror abort -source-path 10.0.0.11:/lun/0005
-destination-path svm_backup:volA_dst
```

3. Break the SnapMirror DR relationship:

```
snapmirror break -source-path <hostip:>/lun/<name> -destination-path
<SVM:volume>|<cluster://SVM/volume>
```

For complete command syntax, see the man page.

The following example breaks the relationship between the source volume 0005 at IP address 10.0.0.11 and the destination volume `volA_dst` on `svm_backup` and the destination volume `volA_dst` on `svm_backup`:

```
cluster_dst::> snapmirror break -source-path 10.0.0.11:/lun/0005
-destination-path svm_backup:volA_dst
```

Configure the destination volume for data access

After making the destination volume writeable, you must configure the volume for data access. SAN hosts can access the data from the destination volume until the source volume is reactivated.

1. Map the Element LUN to the appropriate initiator group.
2. Create iSCSI sessions from the SAN host initiators to the SAN LIFs.
3. On the SAN client, perform a storage re-scan to detect the connected LUN.

Reactivate the original source volume

You can reestablish the original data protection relationship between the source and destination volumes when you no longer need to serve data from the destination.

About this task

The procedure below assumes that the baseline in the original source volume is intact. If the baseline is not intact, you must create and initialize the relationship between the volume you are serving data from and the original source volume before performing the procedure.

You must specify the Element source path in the form `<hostip:>/lun/<name>`, where “lun” is the actual string “lun” and `name` is the name of the Element volume.

Beginning with ONTAP 9.4, snapshot copies of a LUN created while you are serving data from the ONTAP destination are automatically replicated when the Element source is reactivated.

Replication rules are as follows:

- Only iSCSI LUNs are supported.
- You cannot replicate more than one LUN from an ONTAP volume to an Element volume.
- You cannot replicate a LUN from an ONTAP volume to multiple Element volumes.

Steps

1. Delete the original data protection relationship:

```
snapmirror delete -source-path <SVM:volume>|<cluster://SVM/volume>  
-destination-path <hostip:>/lun/<name> -policy <policy>
```

For complete command syntax, see the man page.

The following example deletes the relationship between the original source volume, 0005 at IP address 10.0.0.11, and the volume you are serving data from, `volA_dst` on `svm_backup`:

```
cluster_dst::> snapmirror delete -source-path 10.0.0.11:/lun/0005  
-policy MirrorLatest -destination-path svm_backup:volA_dst
```

2. Reverse the original data protection relationship:

```
snapmirror resync -source-path <SVM:volume>|<cluster://SVM/volume>  
-destination-path <hostip:>/lun/<name> -policy <policy>
```

For complete command syntax, see the man page.

Although `resync` does not require a baseline transfer, it can be time-consuming. You might want to run the `resync` in off-peak hours.

The following example reverses the relationship between the original source volume, 0005 at IP address 10.0.0.11, and the volume you are serving data from, `volA_dst` on `svm_backup`:

```
cluster_dst::> snapmirror resync -source-path svm_backup:volA_dst
-destination-path 10.0.0.11:/lun/0005 -policy MirrorLatest
```

3. Update the reversed relationship:

```
snapmirror update -source-path <SVM:volume>|<cluster://SVM/volume>
-destination-path <hostip:>/lun/<name>
```

For complete command syntax, see the man page.



The command fails if a common snapshot copy does not exist on the source and destination. Use `snapmirror initialize` to re-initialize the relationship.

The following example updates the relationship between the volume you are serving data from, `volA_dst` on `svm_backup`, and the original source volume, `0005` at IP address `10.0.0.11`:

```
cluster_dst::> snapmirror update -source-path svm_backup:volA_dst
-destination-path 10.0.0.11:/lun/0005
```

4. Stop scheduled transfers for the reversed relationship:

```
snapmirror quiesce -source-path <SVM:volume>|<cluster://SVM/volume>
-destination-path <hostip:>/lun/<name>
```

For complete command syntax, see the man page.

The following example stops scheduled transfers between the volume you are serving data from, `volA_dst` on `svm_backup`, and the original source volume, `0005` at IP address `10.0.0.11`:

```
cluster_dst::> snapmirror quiesce -source-path svm_backup:volA_dst
-destination-path 10.0.0.11:/lun/0005
```

5. Stop ongoing transfers for the reversed relationship:

```
snapmirror abort -source-path <SVM:volume>|<cluster://SVM/volume> -destination
-path <hostip:>/lun/<name>
```

For complete command syntax, see the man page.

The following example stops ongoing transfers between the volume you are serving data from, `volA_dst` on `svm_backup`, and the original source volume, `0005` at IP address `10.0.0.11`:

```
cluster_dst::> snapmirror abort -source-path svm_backup:volA_dst
-destination-path 10.0.0.11:/lun/0005
```

6. Break the reversed relationship:

```
snapmirror break -source-path <SVM:volume>|<cluster://SVM/volume> -destination  
-path <hostip:>/lun/<name>
```

For complete command syntax, see the man page.

The following example breaks the relationship between the volume you are serving data from, `volA_dst` on `svm_backup`, and the original source volume, `0005` at IP address `10.0.0.11`:

```
cluster_dst::> snapmirror break -source-path svm_backup:volA_dst  
-destination-path 10.0.0.11:/lun/0005
```

7. Delete the reversed data protection relationship:

```
snapmirror delete -source-path <SVM:volume>|<cluster://SVM/volume>  
-destination-path <hostip:>/lun/<name> -policy <policy>
```

For complete command syntax, see the man page.

The following example deletes the reversed relationship between the original source volume, `0005` at IP address `10.0.0.11`, and the volume you are serving data from, `volA_dst` on `svm_backup`:

```
cluster_src::> snapmirror delete -source-path svm_backup:volA_dst  
-destination-path 10.0.0.11:/lun/0005 -policy MirrorLatest
```

8. Reestablish the original data protection relationship:

```
snapmirror resync -source-path <hostip:>/lun/<name> -destination-path  
<SVM:volume>|<cluster://SVM/volume>
```

For complete command syntax, see the man page.

The following example reestablishes the relationship between the original source volume, `0005` at IP address `10.0.0.11`, and the original destination volume, `volA_dst` on `svm_backup`:

```
cluster_dst::> snapmirror resync -source-path 10.0.0.11:/lun/0005  
-destination-path svm_backup:volA_dst
```

After you finish

Use the `snapmirror show` command to verify that the SnapMirror relationship was created. For complete command syntax, see the man page.

Update a replication relationship manually

You might need to update a replication relationship manually if an update fails because of a network error.

About this task

You must specify the Element source path in the form `<hostip:>/lun/<name>`, where “lun” is the actual string “lun” and `name` is the name of the Element volume.

Steps

1. Update a replication relationship manually:

```
snapmirror update -source-path <hostip:>/lun/<name> -destination-path  
<SVM:volume>|<cluster://SVM/volume>
```

For complete command syntax, see the man page.



The command fails if a common snapshot copy does not exist on the source and destination. Use `snapmirror initialize` to re-initialize the relationship.

The following example updates the relationship between the source volume 0005 at IP address 10.0.0.11 and the destination volume `volA_dst` on `svm_backup`:

```
cluster_src::> snapmirror update -source-path 10.0.0.11:/lun/0005  
-destination-path svm_backup:volA_dst
```

Resynchronize a replication relationship

You need to resynchronize a replication relationship after you make a destination volume writeable, after an update fails because a common Snapshot copy does not exist on the source and destination volumes, or if you want to change the replication policy for the relationship.

About this task

Although `resync` does not require a baseline transfer, it can be time-consuming. You might want to run the `resync` in off-peak hours.

You must specify the Element source path in the form `<hostip:>/lun/<name>`, where “lun” is the actual string “lun” and `name` is the name of the Element volume.

Step

1. Resync the source and destination volumes:

```
snapmirror resync -source-path <hostip:>/lun/<name> -destination-path  
<SVM:volume>|<cluster://SVM/volume> -type XDP -policy <policy>
```

For complete command syntax, see the man page.

The following example resyncs the relationship between the source volume 0005 at IP address 10.0.0.11 and the destination volume `volA_dst` on `svm_backup`:

```
cluster_dst::> snapmirror resync -source-path 10.0.0.11:/lun/0005
-policy MirrorLatest -destination-path svm_backup:volA_dst
```

Back up and restore volumes

You can back up and restore volumes to other SolidFire storage, as well as secondary object stores that are compatible with Amazon S3 or OpenStack Swift.

When you restore volumes from OpenStack Swift or Amazon S3, you need manifest information from the original backup process. If you are restoring a volume that was backed up on a SolidFire storage system, no manifest information is required.

Find more information

- [Back up a volume to an Amazon S3 object store](#)
- [Back up a volume to an OpenStack Swift object store](#)
- [Back up a volume to a SolidFire storage cluster](#)
- [Restore a volume from backup on an Amazon S3 object store](#)
- [Restore a volume from backup on an OpenStack Swift object store](#)
- [Restore a volume from backup on a SolidFire storage cluster](#)

Back up a volume to an Amazon S3 object store

You can back up volumes to external object stores that are compatible with Amazon S3.

1. Click **Management > Volumes**.
2. Click the Actions icon for the volume you want to back up.
3. In the resulting menu, click **Backup to**.
4. In the **Integrated Backup** dialog box under **Backup to**, select **S3**.
5. Select an option under **Data Format**:
 - **Native**: A compressed format readable only by SolidFire storage systems.
 - **Uncompressed**: An uncompressed format compatible with other systems.
6. Enter a hostname to use to access the object store in the **Hostname** field.
7. Enter an access key ID for the account in the **Access Key ID** field.
8. Enter the secret access key for the account in the **Secret Access Key** field.
9. Enter the S3 bucket in which to store the backup in the **S3 Bucket** field.
10. Enter a nametag to append to the prefix in the **Nametag** field.
11. Click **Start Read**.

Back up a volume to an OpenStack Swift object store

You can back up volumes to external object stores that are compatible with OpenStack

Swift.

1. Click **Management > Volumes**.
2. Click the Actions icon for the volume to back up.
3. In the resulting menu, click **Backup to**.
4. In the **Integrated Backup** dialog box under **Backup to**, select **Swift**.
5. Select a data format under **Data Format**:
 - **Native**: A compressed format readable only by SolidFire storage systems.
 - **Uncompressed**: An uncompressed format compatible with other systems.
6. Enter a URL to use to access the object store in the **URL** field.
7. Enter a user name for the account in the **Username** field.
8. Enter the authentication key for the account in the **Authentication Key** field.
9. Enter the container in which to store the backup in the **Container** field.
10. **Optional**: Enter a name tag to append to the prefix in the **Nametag** field.
11. Click **Start Read**.

Back up a volume to a SolidFire storage cluster

You can back up volumes residing on a cluster to a remote cluster for storage clusters running Element software.

Ensure that the source and target clusters are paired.

See [Pair clusters for replication](#).

When backing up or restoring from one cluster to another, the system generates a key to be used as authentication between the clusters. This bulk volume write key allows the source cluster to authenticate with the destination cluster, providing a level of security when writing to the destination volume. As part of the backup or restore process, you need to generate a bulk volume write key from the destination volume before starting the operation.

1. On the destination cluster, **Management > Volumes**.
2. Click the Actions icon for the destination volume.
3. In the resulting menu, click **Restore from**.
4. In the **Integrated Restore** dialog box, under **Restore from**, select **SolidFire**.
5. Select an option under **Data Format**:
 - **Native**: A compressed format readable only by SolidFire storage systems.
 - **Uncompressed**: An uncompressed format compatible with other systems.
6. Click **Generate Key**.
7. Copy the key from the **Bulk Volume Write Key** box to your clipboard.
8. On the source cluster, go to **Management > Volumes**.
9. Click the Actions icon for the volume to back up.
10. In the resulting menu, click **Backup to**.

11. In the **Integrated Backup** dialog box under **Backup to**, select **SolidFire**.
12. Select the same option you selected earlier in the **Data Format** field.
13. Enter the management virtual IP address of the destination volume's cluster in the **Remote Cluster MVIP** field.
14. Enter the remote cluster user name in the **Remote Cluster Username** field.
15. Enter the remote cluster password in the **Remote Cluster Password** field.
16. In the **Bulk Volume Write Key** field, paste the key you generated on the destination cluster earlier.
17. Click **Start Read**.

Restore a volume from backup on an Amazon S3 object store

You can restore a volume from a backup on an Amazon S3 object store.

1. Click **Reporting > Event Log**.
2. Locate the backup event that created the backup you need to restore.
3. In the **Details** column for the event, click **Show Details**.
4. Copy the manifest information to your clipboard.
5. Click **Management > Volumes**.
6. Click the Actions icon for the volume you want to restore.
7. In the resulting menu, click **Restore from**.
8. In the **Integrated Restore** dialog box under **Restore from**, select **S3**.
9. Select the option that matches the backup under **Data Format**:
 - **Native**: A compressed format readable only by SolidFire storage systems.
 - **Uncompressed**: An uncompressed format compatible with other systems.
10. Enter a hostname to use to access the object store in the **Hostname** field.
11. Enter an access key ID for the account in the **Access Key ID** field.
12. Enter the secret access key for the account in the **Secret Access Key** field.
13. Enter the S3 bucket in which to store the backup in the **S3 Bucket** field.
14. Paste the manifest information into the **Manifest** field.
15. Click **Start Write**.

Restore a volume from backup on an OpenStack Swift object store

You can restore a volume from a backup on an OpenStack Swift object store.

1. Click **Reporting > Event Log**.
2. Locate the backup event that created the backup you need to restore.
3. In the **Details** column for the event, click **Show Details**.
4. Copy the manifest information to your clipboard.
5. Click **Management > Volumes**.
6. Click the Actions icon for the volume you want to restore.

7. In the resulting menu, click **Restore from**.
8. In the **Integrated Restore** dialog box under **Restore from**, select **Swift**.
9. Select the option that matches the backup under **Data Format**:
 - **Native**: A compressed format readable only by SolidFire storage systems.
 - **Uncompressed**: An uncompressed format compatible with other systems.
10. Enter a URL to use to access the object store in the **URL** field.
11. Enter a user name for the account in the **Username** field.
12. Enter the authentication key for the account in the **Authentication Key** field.
13. Enter the name of the container in which the backup is stored in the **Container** field.
14. Paste the manifest information into the **Manifest** field.
15. Click **Start Write**.

Restore a volume from backup on a SolidFire storage cluster

You can restore a volume from a backup on a SolidFire storage cluster.

When backing up or restoring from one cluster to another, the system generates a key to be used as authentication between the clusters. This bulk volume write key allows the source cluster to authenticate with the destination cluster, providing a level of security when writing to the destination volume. As part of the backup or restore process, you need to generate a bulk volume write key from the destination volume before starting the operation.

1. On the destination cluster, click **Management > Volumes**.
2. Click the Actions icon for the volume you want to restore.
3. In the resulting menu, click **Restore from**.
4. In the **Integrated Restore** dialog box, under **Restore from**, select **SolidFire**.
5. Select the option that matches the backup under **Data Format**:
 - **Native**: A compressed format readable only by SolidFire storage systems.
 - **Uncompressed**: An uncompressed format compatible with other systems.
6. Click **Generate Key**.
7. Copy the **Bulk Volume Write Key** information to the clipboard.
8. On the source cluster, click **Management > Volumes**.
9. Click the Actions icon for the volume you want to use for the restore.
10. In the resulting menu, click **Backup to**.
11. In the **Integrated Backup** dialog box, select **SolidFire** under **Backup to**.
12. Select the option that matches the backup under **Data Format**.
13. Enter the management virtual IP address of the destination volume's cluster in the **Remote Cluster MVIP** field.
14. Enter the remote cluster user name in the **Remote Cluster Username** field.
15. Enter the remote cluster password in the **Remote Cluster Password** field.
16. Paste the key from your clipboard into the **Bulk Volume Write Key** field.

17. Click **Start Read**.

Configure custom Protection Domains

For Element clusters that contain more than two storage nodes, you can configure custom Protection Domains for each node. When you configure custom Protection Domains, you must assign all nodes in the cluster to a domain.



When you assign Protection Domains, a data sync between nodes begins, and some cluster operations are unavailable until the data sync completes. After a custom Protection Domain is configured for a cluster, when you add a new storage node, you cannot add drives for the new node until you assign a Protection Domain for the node and allow the data sync to complete. Visit the [Protection Domains documentation](#) to learn more about Protection Domains.



For a custom Protection Domain scheme to be useful for a cluster, all storage nodes within each chassis must be assigned to the same custom Protection Domain. You need to create as many custom Protection Domains as is needed for this to be the case (the smallest possible custom Protection Domain scheme is three domains). As a best practice, configure an equal number of nodes per domain and try to ensure that each node assigned to a particular domain is of the same type.

Steps

1. Click **Cluster > Nodes**.
2. Click **Configure Protection Domains**.

In the **Configure Custom Protection Domains** window, you can see the currently configured Protection Domains (if any) as well as Protection Domain assignments for individual nodes.

3. Enter a name for the new custom Protection Domain, and click **Create**.

Repeat this step for all new Protection Domains you need to create.

4. For each node in the **Assign Nodes** list, click the dropdown in the **Protection Domain** column and select a Protection Domain to assign to that node.



Ensure that you understand your node and chassis layout, the custom Protection Domain scheme you have configured, and the scheme's effects on data protection before applying the changes. If you apply a Protection Domain scheme and immediately need to make changes, it could be some time before you can do so because of the data sync that happens once a configuration is applied.

5. Click **Configure Protection Domains**.

Result

Depending on the size of your cluster, the data sync data between domains could take some time. After the data sync is complete, you can view the custom Protection Domain assignments on the **Cluster > Nodes** page, and the Element web UI dashboard shows the protection status of the cluster in the **Custom Protection Domain Health** pane.

Possible errors

Here are some errors you might see after applying a custom Protection Domain configuration:

Error	Description	Resolution
SetProtectionDomainLayout Failed: ProtectionDomainLayout would leave NodeID {9} unusable. Default and non-default names cannot both be used together.	A node does not have a Protection Domain assigned.	Assign a Protection Domain to the node.
SetProtectionDomainLayout Failed: Protection domain type 'custom' splits Protection Domain type 'chassis'.	A node in a multi-node chassis is assigned a Protection Domain that is different from other nodes in the chassis.	Ensure that all nodes in the chassis are assigned the same Protection Domain.

Find more information

- [Custom Protection Domains](#)
- [Manage storage with the Element API](#)

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.