



Perform remote replication between clusters running NetApp Element software

Element Software

NetApp
October 01, 2024


Table of Contents

- Perform remote replication between clusters running NetApp Element software 1
 - Plan cluster and volume pairing for real-time replication 1
 - Pair clusters for replication 2
 - Pair volumes 4
 - Validate volume replication 9
 - Delete a volume relationship after replication 10
 - Manage volume relationships 10

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:



- [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.

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.