



Data protection and disaster recovery

System Manager Classic

NetApp
April 09, 2024

This PDF was generated from <https://docs.netapp.com/us-en/ontap-sm-classic/peering/index.html> on April 09, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Data protection and disaster recovery 1
 - Cluster and SVM peering configuration 1
 - Volume disaster recovery 11
 - Volume disaster recovery preparation 23
 - Volume backup using SnapVault 32
 - Volume restore management using SnapVault 40

Data protection and disaster recovery

Cluster and SVM peering configuration

Cluster and SVM peering overview

Cluster administrators can create authenticated peer relationships between clusters and SVMs to enable the clusters to communicate with each other so that data is replicated between volumes in different clusters. You can perform the procedures using the ONTAP System Manager *classic* interface, which is available with ONTAP 9.7 and earlier ONTAP 9 releases.

Use the ONTAP System Manager *classic* interface to create cluster peer relationships and SVM peer relationships if the following apply:

- You are working with clusters running ONTAP 9.7 or earlier ONTAP 9 releases.
- You want cluster peering relationships that are authenticated.
- You want to use best practices, not explore every available option.
- You want to use System Manager, not the ONTAP command-line interface (CLI) or an automated scripting tool.

Other ways to do this in ONTAP

ONTAP System Manager in ONTAP 9.3 simplifies the way that you configure peer relationships between clusters and between SVMs. The cluster peering procedure and SVM peering procedure can be used for all ONTAP 9 versions. You should use the appropriate procedure for your version of ONTAP.

To perform these tasks with...	Refer to...
The redesigned System Manager (available with ONTAP 9.7 and later)	<ul style="list-style-type: none">• Cluster management with System Manager
The ONTAP command-line interface (CLI)	<ul style="list-style-type: none">• Cluster and SVM peering overview with the CLI <p>Use the command-line interface to set up cluster peering relationships and SVM peering relationships.</p> <ul style="list-style-type: none">• Network management <p>Use the command-line interface to configure subnets, intercluster LIFs, routes, firewall policies, and other networking components</p>

Prerequisites for cluster peering

Before you set up cluster peering using the ONTAP System Manager *classic* interface with ONTAP 9.7 or earlier, you should confirm that the connectivity, port, IP address,

subnet, firewall, and cluster-naming requirements are met.

Connectivity requirements

Every intercluster LIF on the local cluster must be able to communicate with every intercluster LIF on the remote cluster.

Although it is not required, it is typically simpler to configure the IP addresses used for intercluster LIFs in the same subnet. The IP addresses can reside in the same subnet as data LIFs, or in a different subnet. The subnet used in each cluster must meet the following requirements:

- The subnet must have enough IP addresses available to allocate to one intercluster LIF per node.

For example, in a six-node cluster, the subnet used for intercluster communication must have six available IP addresses.

Each node must have an intercluster LIF with an IP address on the intercluster network.

Intercluster LIFs can have an IPv4 address or an IPv6 address.



ONTAP 9 enables you to migrate your peering networks from IPv4 to IPv6 by optionally allowing both protocols to be present simultaneously on the intercluster LIFs. In earlier releases, all intercluster relationships for an entire cluster were either IPv4 or IPv6. This meant that changing protocols was a potentially disruptive event.

Port requirements

You can use dedicated ports for intercluster communication, or share ports used by the data network. Ports must meet the following requirements:

- All ports that are used to communicate with a given remote cluster must be in the same IPspace.

You can use multiple IPspaces to peer with multiple clusters. Pair-wise full-mesh connectivity is required only within an IPspace.

- The broadcast domain that is used for intercluster communication must include at least two ports per node so that intercluster communication can fail over from one port to another port.

Ports added to a broadcast domain can be physical network ports, VLANs, or interface groups (ifgrps).

- All ports must be cabled.
- All ports must be in a healthy state.
- The MTU settings of the ports must be consistent.

Firewall requirements

Firewalls and the intercluster firewall policy must allow the following protocols:

- ICMP service
- TCP to the IP addresses of all the intercluster LIFs over the ports 10000, 11104, and 11105
- Bidirectional HTTPS between the intercluster LIFs

Although HTTPS is not required when you set up cluster peering using the CLI, HTTPS is required later if you use ONTAP System Manager to configure data protection.

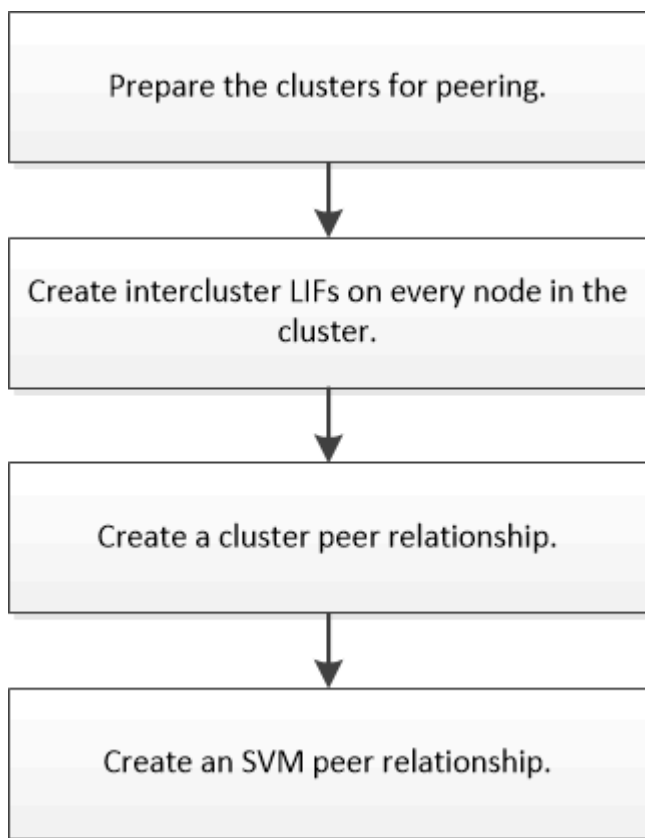
The default `intercluster` firewall policy allows access through the HTTPS protocol and from all IP addresses (0.0.0.0/0). You can modify or replace the policy if necessary.

Related information

[Data protection](#)

Cluster and SVM peering workflow

You can set up a peering relationship by using the ONTAP System Manager with ONTAP 9.7 or earlier. Setting up a peering relationship involves preparing each cluster for peering, creating intercluster logical interfaces (LIFs) on each node of each cluster, setting up a cluster peer relationship, and then setting up an SVM peering relationship.



If you are running ONTAP 9.2 or earlier, you create an SVM peering relationship while creating a data protection relationship between the source volume and the destination volume.

Prepare for cluster peering

Before creating a cluster peering relationship using the ONTAP System Manager *classic* interface with ONTAP 9.7 or earlier, you must verify that the time on each cluster is synchronized with an external Network Time Protocol (NTP) server, and determine the subnets, ports, and passphrases that you want to use.

Steps

1. If you are running ONTAP 9.2 or earlier, determine the passphrase that you want to use for each cluster peer relationship.

The passphrase must include at least eight characters.

For the relationship between...	The passphrase is...
Cluster A and Cluster B	

Beginning with ONTAP 9.3, you can generate the passphrase from the remote cluster while creating the cluster peer relationship.

Creating a cluster peer relationship (Beginning with ONTAP 9.3)

2. Identify the subnets, IP addresses, and ports that you will use for intercluster LIFs.

By default, the IP address is automatically selected from the subnet. If you want to specify the IP address manually, you must ensure that the IP address either is already available in the subnet or can be added to the subnet later. Information about subnets is available in the Network tab.

Create a table similar to the following table to record information about the clusters. The following table assumes that each cluster has four nodes. If a cluster has more than four nodes, add rows for the additional information.

	Cluster A	Cluster B
Subnet (ONTAP 9.2 or earlier)		
IP address (Beginning with ONTAP 9.3, optional for ONTAP 9.2 or earlier)		
Node 1 port		
Node 2 port		
Node 3 port		
Node 4 port		

Configure peer relationships (Beginning with ONTAP 9.3)

A peer relationship defines the network connections that enable clusters and SVMs to exchange data securely. Beginning with ONTAP 9.3, until ONTAP 9.7, you can use the ONTAP System Manager *classic* interface to perform a simplified method to configure peer relationships between clusters and between SVMs.

Create intercluster LIFs (Beginning with ONTAP 9.3)

Beginning with ONTAP 9.3, until ONTAP 9.7, you can use the ONTAP System Manager *classic* interface to create intercluster logical interfaces (LIFs), which enable the cluster network to communicate with a node. You must create an intercluster LIF within each IPspace that will be used for peering, on each node in each cluster for which you want to create a peer relationship.

About this task

For example, if you have a four-node cluster that you want to peer with cluster X over IPspace A, and peer with cluster Y over IPspace Y, then you need a total of eight intercluster LIFs; Four on IPspace A (one per node), and four on IPspace Y (one per node).

You must perform this procedure on both clusters for which you want to create a peer relationship.

Steps

1. Click **Configuration > Advanced Cluster Setup**.
2. In the **Setup Advanced Cluster Features** window, click **Proceed** next to the **Cluster Peering** option.
3. Select an IPspace from the **IPspace** list.
4. Enter the IP address, port, network mask, and gateway details of each node.

	IPspace	IP Address	Port	Netmask	Gateway (Optional)
sti50-vsim-ucs103a	Default	10.53.32.1	e0d	255.255.240.0	
sti50-vsim-ucs103b		10.53.32.2	e0d		

5. Click **Submit and Continue**.

What to do next

You should enter the cluster details in the Cluster Peering window to continue with cluster peering.

Create a cluster peer relationship (Beginning with ONTAP 9.3)

Beginning with ONTAP 9.3, until ONTAP 9.7, you can use the ONTAP System Manager *classic* interface create a cluster peer relationship between two clusters by providing a system-generated passphrase and the IP addresses of the intercluster LIFs of the remote cluster.

About this task

Beginning with ONTAP 9.6, cluster peering encryption is enabled by default on all newly created cluster peering relationships. Cluster peering encryption must be enabled manually for peering relationship created prior to upgrading to ONTAP 9.6. Cluster peering encryption is not available for clusters running ONTAP 9.5 or earlier. Therefore, both clusters in the peering relationship must be running ONTAP 9.6 in order to enable cluster peering encryption.

Cluster peering encryption uses the Transport Security Layer (TLS) to secure cross-cluster peering

communications for ONTAP features such as SnapMirror and FlexCache.

Steps


- 1. In the **Target Cluster Intercluster LIF IP addresses** field, enter the IP addresses of the intercluster LIFs of the remote cluster.
- 2. Generate a passphrase from the remote cluster.
 - a. Specify the management address of the remote cluster.
 - b. Click **Management URL** to launch ONTAP System Manager on the remote cluster.
 - c. Log in to the remote cluster.
 - d. In the **Cluster Peers** window, click **Generate Peering Passphrase**.
 - e. Select the IPspace, validity of the passphrase, and SVM permissions.

You can allow all of the SVMs or selected SVMs for peering. When a SVM peer request is generated, the permitted SVMs are automatically peered with the source SVMs without requiring you to accept the peer relationship from the remote SVMs.

- f. Click **Generate**.

The passphrase information is displayed.

Generate Peering Passphrase



Passphrase generated successfully

Use the following information for peering based on the IPspace "Default":

Intercluster LIF IP Address

172.21.91.12

Passphrase

QS7k+laFYJzdV9UMPXvHgWd

Passphrase Validity

Valid Until Mon Nov... America/New_Y

SVM Permissions

All

Email passphrase details

Copy passphrase details

Done

- g. Click **Copy passphrase details** or **Email passphrase details**.

- h. Click **Done**.
3. In the source cluster, enter the generated passphrase that you obtained in [Step 2](#).
4. Click **Initiate Cluster Peering**.

The cluster peer relationship is successfully created.

5. Click **Continue**.

What to do next

You should specify the SVM details in the SVM Peering window to continue with the peering process.

Create SVM peer relationships

Beginning with ONTAP 9.3, until ONTAP 9.7, you can use the ONTAP System Manager *classic* interface to create SVM peer relationships. The storage virtual machine (SVM) peering enables you to establish a peer relationship between two SVMs for data protection.

Steps

1. Select the initiator SVM.
2. Select the target SVM from the list of permitted SVMs.
3. Click **Initiate SVM Peering**.
4. Click **Continue**.

What to do next

You can view the intercluster LIFs, cluster peer relationship, and SVM peer relationship in the Summary window.

Configure peer relationships (ONTAP 9.2 and earlier)

Using the ONTAP System Manager *classic* interface with ONTAP 9.2 or an earlier ONTAP 9 release, you can create SVM peer relationships.

A peer relationship defines network connections that enable clusters and SVMs to exchange data securely. You must create a cluster peer relationship before you can create an SVM peer relationship.

Create intercluster interfaces on all nodes (ONTAP 9.2 or earlier)

Using the ONTAP System Manager *classic* interface with ONTAP 9.2 or an earlier ONTAP 9 release, you can create intercluster LIFs that will be used for peering.

Clusters communicate with each other through logical interfaces (LIFs) that are dedicated to intercluster communication. You must create an intercluster LIF within each IPspace that will be used for peering. The LIFs must be created on each node in each cluster for which you want to create a peer relationship.

Before you begin

You must have identified the subnet and ports, and optionally the IP addresses, that you plan to use for the intercluster LIFs.

About this task

You must perform this procedure on both clusters for which you want to create a peer relationship. For example, if you have a four-node cluster that you want to peer with cluster X over IPspace A, and peer with cluster Y over IPspace Y, then you need a total of eight intercluster LIFs; Four on IPspace A (one per node), and four on IPspace Y (one per node).

Steps

1. Create an intercluster LIF on one node of the source cluster:

- a. Navigate to the **Network Interfaces** window.
- b. Click **Create**.

The Create Network Interface dialog box is displayed.

- c. Enter a name for the intercluster LIF.

You can use “icl01” for the intercluster LIF on the first node, and “icl02” for the intercluster LIF on the second node.

- d. Select **Intercluster Connectivity** as the interface role.
- e. Select the IPspace.
- f. In the **Add Details** dialog box, select **Using a subnet** from the **Assign IP Address** drop-down list, and then select the subnet that you want to use for intercluster communication.

By default, the IP address is automatically selected from the subnet after you click **Create**. If you do not want to use the IP address that is automatically selected, you must manually specify the IP address that the node uses for intercluster communication.

- g. If you want to manually specify the IP address that the node uses for intercluster communication, select **Use this IP Address**, and type the IP address.

You must ensure that the IP address that you want to use either is already available in the subnet or can be added to the subnet later.

- h. In the **Ports** area, click the node that you are configuring, and select the port that you want to use for this node.
- i. If you decided not to share ports for intercluster communication with data communication, confirm that the selected port displays “0” in the **Hosted Interface Count** column.

Create Network Interface

Specify the following details to add a new network interface for data and management access of the chosen SVM.

Name:

Interface Role: ☐ Serves Data ☒ Intercluster Connectivity

SVM:

Protocol Access: ☐ CIFS ☐ iSCSI ☐ NFS ☐ FC/FCoE

Management Access: ☐ Enable Management Access

Subnet:

☒ The IP address is selected from this subnet. ☐ Use this IP Address:

This IP address will be added to the chosen subnet if the address is not already present in the subnet available range.

Port:


Ports or Adapters	Hosted Interface Count	Speed
clusterA-node1		
e0c	3	1000 Mbps
e0d	0	1000 Mbps
e0e	0	1000 Mbps

j. Click **Create**.

2. Repeat [Step 1](#) for each node in the cluster.


Each node in the cluster has an intercluster LIF.

3. Make a note of the IP addresses of the intercluster LIFs so that you can use them later when you create peer relationships with other clusters:

- In the **Network Interfaces** window, in the **Role** column, click , clear the **All** check box, and then select **Intercluster**.

The Network Interfaces window displays only intercluster LIFs.

- Note down the IP addresses that are listed in the **IP Addresses/WWPN** column, or leave the **Network Interfaces** window open so that you can retrieve the IP addresses later.

You can click the column display icon  to hide the columns that you do not want to view.

Results

All of the nodes in each cluster have intercluster LIFs that can all communicate with each other.

Create a cluster peer relationship (ONTAP 9.2 or earlier)

Using the ONTAP System Manager *classic* interface with ONTAP 9.2 or an earlier ONTAP 9 release, you can create a cluster peer relationship between two clusters by

entering a predetermined passphrase and the IP addresses of the intercluster LIFs of the remote cluster, and then verifying that the relationship was created successfully.

Before you begin

- You must know the IP addresses of all of the intercluster LIFs of the clusters that you want to peer.
- You must know the passphrase that you will use for each peer relationship.

About this task

You must perform this procedure on each cluster.

Steps

1. From the source cluster, create a cluster peer relationship with the destination cluster.

- a. Click the **Configurations** tab.
- b. In the **Cluster Settings** pane, click **Cluster Peers**.
- c. Click **Create**.

The **Create Cluster Peer** dialog box is displayed.

- d. In the **Details of the remote cluster to be peered** area, specify the passphrase that both peers will use to ensure an authenticated cluster peer relationship.
- e. Enter the IP addresses of all of the intercluster LIFs of the destination cluster (one per node) separated by commas.

Create Cluster Peer

For a cluster to communicate with another cluster in a peer relationship, enter a passphrase and the intercluster IP addresses of the peer cluster.
[Tell me more about cluster peering](#)

Details of the local cluster

Cluster Name: clusterA

Intercluster IP Addresses:

clusterA-node1	10.53.52.120
clusterA-node2	10.53.52.121

Details of the remote cluster to be peered

Passphrase: *****

Intercluster IP Addresses: 10.238.14.33,10.238.14.36

- f. Click **Create**.

The authentication status is “pending” because only one cluster has been configured.

2. Switch to the destination cluster, and then create a cluster peer relationship with the source cluster:

- a. Click the **Configurations** tab.
- b. In the **Cluster Settings** pane, click **Cluster Peers**.
- c. Click **Create**.

The Create Cluster Peer dialog box is displayed.

- d. In the **Details of the remote cluster to be peered** area, specify the same passphrase that you specified in [Step 1d](#) and the IP addresses of the intercluster LIFs of the source cluster, and then click

Create.

Create Cluster Peer

For a cluster to communicate with another cluster in a peer relationship, enter a passphrase and the intercluster IP addresses of the peer cluster.
[Tell me more about cluster peering](#)

Details of the local cluster

Cluster Name: clusterB

Intercluster IP Addresses:

clusterB-node1	10.238.14.33
clusterB-node2	10.238.14.36

Details of the remote cluster to be peered

Passphrase:

Intercluster IP Addresses:

10.53.52.120,10.53.52.121

- From the **Cluster Peers** window of the destination cluster, confirm that the source cluster is “available” and that the authentication status is “OK”.

Cluster Peers

'Availability' and 'Authentication Status' information might be stale for up to several minutes.

Create Modify Passphrase Modify Peer Network Parameters Delete Refresh

Peer Cluster	Availability	Authentication Status
clusterA	available	ok

You might have to click **Refresh** to view the updated information.

The two clusters are in a peer relationship.

- Switch to the source cluster, and confirm that the destination cluster is “available” and that the authentication status is “OK”.

You might have to click **Refresh** to view the updated information.

What to do next

Create an SVM peer relationship between the source and destination SVMs while creating a data protection relationship between the source volume and the destination volume.

[Volume backup using SnapVault](#)

[Volume disaster recovery preparation](#)

Volume disaster recovery

Volume disaster recovery overview

You can quickly activate a destination volume after a disaster and then reactivate the source volume in ONTAP using the ONTAP System Manager classic interface (ONTAP 9.7 and earlier).

Use this procedure if you want to perform volume-level disaster recovery in the following way:


- You are working with clusters running ONTAP 9.
- You are a cluster administrator.
- You have configured the SnapMirror relationship following [Volume disaster recovery preparation](#)
- The cluster administrator of the source cluster has declared that the data in the source volume is unavailable due to events such as virus infection leading to data corruption or accidental deletion of data.
- You want to use System Manager, not the ONTAP command-line interface or an automated scripting tool.
- You want to use the System Manager classic interface for ONTAP 9.7 and earlier releases, not the ONTAP System Manager UI for ONTAP 9.7 and later.
- You want to use best practices, not explore every available option.
- You do not want to read a lot of conceptual background.

Other ways to do this in ONTAP

To perform these tasks with...	See this content...
The redesigned System Manager (available with ONTAP 9.7 and later)	Serve data from a SnapMirror destination
The ONTAP command line interface	Activate the destination volume

Volume disaster recovery workflow

The volume disaster recovery workflow includes activating the destination volume, configuring the destination volume for data access, and reactivating the original source volume.



Additional information is available to help you to manage the volume-level disaster recovery relationships and provides other methods of disaster recovery to protect the availability of your data resources.

- [Volume backup using SnapVault](#)

Describes how to quickly configure backup vault relationships between volumes that are located in different ONTAP clusters.

- [Volume restore management using SnapVault](#)

Describes how to quickly restore a volume from a backup vault in ONTAP.

Activate the destination volume

When the source volume is unable to serve data due to events such as data corruption, accidental deletion or an offline state, you must activate the destination volume to provide data access until you recover the data on the source volume. Activation involves stopping future SnapMirror data transfers and breaking the SnapMirror relationship.

Verify the status of the source volume

When the source volume is unavailable, you must verify that the source volume is offline and then identify the destination volume that must be activated for providing data access.

About this task

You must perform this task from the **source** cluster.

Steps

1. Navigate to the **Volumes** window.
2. Select the source volume, and then verify that the source volume is offline.
3. Identify the destination volume in the SnapMirror relationship.
 - Beginning with ONTAP 9.3: Double-click the source volume to view the details, and then click **PROTECTION** to identify the destination volume in the SnapMirror relationship and the name of the SVM that contains the volume.
- ONTAP 9.2 or earlier: Click the **Data Protection** tab at the bottom of the Volumes page to identify the destination volume in the SnapMirror relationship and the name of the SVM that contains the volume.

Volume: vol_mirror_src

[Back to All volumes](#) [Edit](#) [Delete](#) [Actions](#) [Refresh](#)

Overview Snapshots Copies **Data Protection** Storage Efficiency Performance

Health	Destination SVM	Destination Volume	Destination Clu...	Relationship...	Transfer S...	Type	Lag Time	Policy
	svm2	vol_mirror_src_dst	cluster2	Snapmirrored	Idle	Version-Resilient ...	45 min(s)	Mirror/Snap...

Create	Edit	Delete	Status	Snapshot Copies	Resize	Storage Efficiency	Move	Storage QoS	Protect	Re...
Name	Aggregate	Status	Thin Pro...	% Used	Availabl...	Total Sp...	Storage ...	Is Volu...	Encrypted	
svm1_svm1_root...	aggr2	Online	No	5	970.48 MB	1 GB	Disabled	No	No	
svm1_vol123_vault	aggr2	Online	No	5	121.35 MB	128.02 MB	Enabled	No	No	
Vol1	aggr3	Offline	-NA-	-NA-	-NA-	-NA-	Disabled	No	No	
svm2_root	aggr1	Online	No	5	971.12 MB	1 GB	Disabled	No	No	
Destination St...	Destination Vo...	Is Healthy	Relationship St...	Transfer Status	Type	Lag Time	Policy			
svm1	vol1	Yes	Snapmirrored	Idle	Mirror	7 day(s) 12 hr(s)...	DPDefault			
Details	Space Allocation	Snapshot Copies	Storage Efficiency	Data Protection	Volume Move Det	Performance				

Break the SnapMirror relationship

You must quiesce and break the SnapMirror relationship to activate the destination volume. After quiescing, future SnapMirror data transfers are disabled.

Before you begin

The destination volume must be mounted on the destination SVM namespace.

About this task


You must perform this task from the **destination** cluster.

Steps

- Depending on the System Manager version that you are running, perform one of the following steps:
 - ONTAP 9.4 or earlier: Click **Protection > Relationships**.
 - Beginning with ONTAP 9.5: Click **Protection > Volume Relationships**.
- Select the SnapMirror relationship between the source and the destination volumes.
- Click **Operations > Quiesce** to disable future data transfers.
- Select the confirmation check box, and then click **Quiesce**.

The quiesce operation might take some time; you must not perform any other operation on the SnapMirror relationship until the transfer status is displayed as *Quiesced*.

- Click **Operations > Break**.
- Select the confirmation check box, and then click **Break**.



The SnapMirror relationship is in Broken Off state.

Source Sto...	Source Vol...	Destinatio...	Destinatio...	Is Healthy	Relationship	Transfer St...	Relationship	Lag Time	Policy Name	Policy Type
svm1	svm1_root	svm1_svm1_r...	svm2	Yes	Snapmirrored	Idle	Mirror	26 min(s)	DPDefault	Asynchronous
svm1	vol1	svm1_vol1_m...	svm2	Yes	Broken Off	Idle	Mirror	None	DPDefault	Asynchronous

Source Location:	svm1.vol1	Is Healthy:	Yes	Transfer Status:	Idle
Destination Location:	svm2:svm1_vol1_mirror	Relationship State:	Broken Off	Current Transfer Type:	None
Source Cluster:	cluster-1	Network Compression Ratio:	Not Applicable	Current Transfer Error:	None
Destination Cluster:	cluster-1			Last Transfer Error:	None
Transfer Schedule:	hourly			Last Transfer Type:	Update
Data Transfer Rate:	Unlimited			Latest Snapshot Timestamp:	02/22/2017 13:05:00
Lag Time:	None			Latest Snapshot Copy:	snapmirror.9b4dae7c-e6d0-11e6-b44a-00a09b1a1bda_2149622820.2017-1

Verify the destination volume status

After breaking the SnapMirror relationship, you must verify that the destination volume has read/write access and that the destination volume settings match the settings of the source volume.

About this task

You must perform this task from the **destination** cluster.

Steps

1. Navigate to the **Volumes** window.
2. Select the destination volume from the **Volumes** list, and then verify that the destination volume type is `rw`, which indicates read/write access.
3. Verify that the volume settings such as thin provisioning, deduplication, compression, and autogrow on the destination volume match the settings of the source volume.

You can use the volume settings information that you noted after creating the SnapMirror relationship to verify the destination volume settings.

4. If the volume settings do not match, modify the settings on the destination volume as required:

- a. Click **Edit**.
- b. Modify the general settings, storage efficiency settings, and advanced settings for your environment, as required.
- c. Click **Save and Close**.

Edit Volume

General | Storage Efficiency | Advanced

Name:

Security style: ▼

☒ Configure UNIX permissions (Optional)

	Read	Write	Execute
Owner	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Group	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Others	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>

☒ Thin Provisioned

When a volume is thin provisioned, space for the volume is not allocated in advance. Instead, space is allocated as data is written to the volume. The unused aggregate space is available to other thin provisioned volumes and LUNs.

[Tell me more about Thin Provisioning](#)

Save Save and Close Cancel

- d. Verify that the columns in the **Volumes** list are updated with the appropriate values.
5. Enable Snapshot copy creation for the destination volume.
- a. Depending on your ONTAP version, navigate to the **Configure Volume Snapshot Copies** page in one of the following ways:

Beginning with ONTAP 9.3: Select the destination volume, and then click **Actions > Manage Snapshots > Configure**.

ONTAP 9.2 or earlier: Select the destination volume, and then click **Snapshot Copies > Configure**.
 - b. Select the **Enable scheduled Snapshot Copies** check box, and then click **OK**.

Configure Volume Snapshot Copies

? Snapshot Reserve (%):

☒ Make Snapshot directory (.snapshot) visible
Visibility of .snapshot directory on this volume at the client mount points.

☒ Enable scheduled Snapshot Copies

Snapshot Policies and Schedules

Select a Snapshot policy that has desired schedules for Snapshot copies:

Snapshot Policy:

Schedules of Selected Snapshot Policy:

Schedul...	Retained S...	Schedule	SnapMirror Label
hourly	6	Advance cron - {Minu...	-
daily	2	Daily - Run at 0 hour 1...	daily
weekly	2	On weekdays - Sund...	weekly

Current Timezone: US/Pacific

[Tell me more about Snapshot configurations](#)

OK Cancel

Configure the destination volume for data access

After activating the destination volume, you must configure the volume for data access. NAS clients and SAN hosts can access the data from the destination volume until the source volume is reactivated.

About this task

You must perform this task from the **destination** cluster.

Procedure

- NAS environment:
 - a. Mount the NAS volumes to the namespace using the same junction path that the source volume was mounted to in the source SVM.
 - b. Apply the appropriate ACLs to the CIFS shares at the destination volume.
 - c. Assign the NFS export policies to the destination volume.
 - d. Apply the quota rules to the destination volume.
 - e. Redirect clients to the destination volume by performing the necessary steps such as changing the DNS name resolution.
 - f. Remount the NFS and CIFS shares on the clients.

- SAN environment:
 - a. Map the LUNs to the appropriate initiator group to make the LUNs in the volume available to the SAN clients.
 - b. For iSCSI, create iSCSI sessions from the SAN host initiators to the SAN LIFs.
 - c. On the SAN client, perform a storage re-scan to detect the connected LUNs.

What to do next

You should resolve the problem that caused the source volume to become unavailable. You must bring the source volume back online when possible, and then resynchronize and reactivate the source volume.

Related information

[ONTAP 9 Documentation Center](#)

Reactivate the source volume

When the source volume becomes available, you must resynchronize the data from the destination volume to the source volume, update any modifications after the resynchronization operation, and activate the source volume.

Resynchronize the source volume

When the source volume is online, you must resynchronize the data between the destination volume and the source volume to replicate the latest data from the destination volume.


Before you begin

The source volume must be online.

About this task

You must perform the task from the **destination** cluster.


The following image shows that the data is replicated from the active destination volume to the read-only source volume:



Steps

1. Depending on the System Manager version that you are running, perform one of the following steps:
 - ONTAP 9.4 or earlier: Click **Protection > Relationships**.
 - Beginning with ONTAP 9.5: Click **Protection > Volume Relationships**.
2. Select the SnapMirror relationship between the source and destination volumes.
3. Make a note of the transfer schedule and the policy configured for the SnapMirror relationship.

4. Click **Operations > Reverse Resync**.
5. Select the confirmation check box, and then click **Reverse Resync**.



Beginning with ONTAP 9.3, the SnapMirror policy of the relationship is set to `MirrorAllSnapshots` and the mirror schedule is set to `None`.

If you are running ONTAP 9.2 or earlier, the SnapMirror policy of the relationship is set to `DPDefault` and the mirror schedule is set to `None`.

6. On the source cluster, specify a SnapMirror policy and schedule that match the protection configuration of the original SnapMirror relationship:
 - a. Depending on the System Manager version that you are running, perform one of the following steps:
 - ONTAP 9.4 or earlier: Click **Protection > Relationships**.
 - Beginning with ONTAP 9.5: Click **Protection > Volume Relationships**.
 - b. Select the SnapMirror relationship between the resynchronized source volume and the destination volume, and then click **Edit**.
 - c. Select the SnapMirror policy and schedule, and then click **OK**.

Update the source volume

After resynchronizing the source volume, you might want to ensure that all the latest changes are updated on the source volume before activating the source volume.

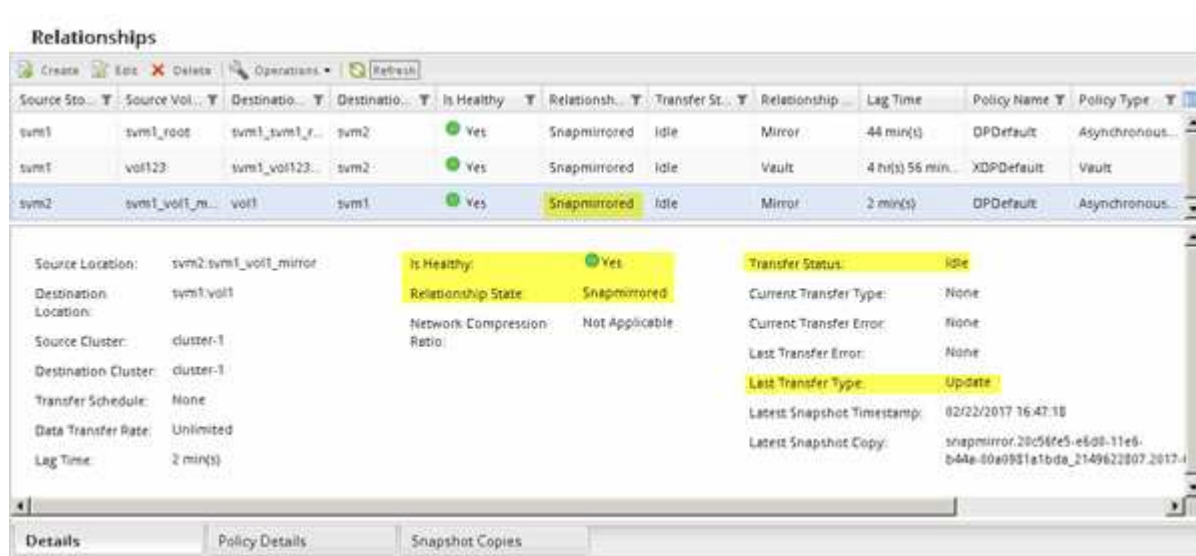
About this task

You must perform this task from the **source** cluster.

Steps

1. Depending on the System Manager version that you are running, perform one of the following steps:

- ONTAP 9.4 or earlier: Click **Protection > Relationships**.
 - Beginning with ONTAP 9.5: Click **Protection > Volume Relationships**.
2. Select the SnapMirror relationship between the source and the destination volumes, and then click **Operations > Update**.
 3. Perform an incremental transfer from the recent common Snapshot copy between the source and destination volumes.
 - Beginning with ONTAP 9.3: Select the **As per policy** option.
 - ONTAP 9.2 or earlier: Select the **On demand** option.
 4. **Optional:** Select **Limit transfer bandwidth to** in order to limit the network bandwidth used for transfers, and then specify the maximum transfer speed.
 5. Click **Update**.
 6. Verify that the transfer status is **Idle** and last transfer type is **Update** in the **Details** tab.




Reactivate the source volume

After resynchronizing the data from the destination volume to the source volume, you must activate the source volume by breaking the SnapMirror relationship. You should then resynchronize the destination volume to protect the reactivated source volume.

About this task


Both the break and reverse resync operations are performed from the **source** cluster.

The following image shows that the source and destination volumes are read/write when you break the SnapMirror relationship. After the reverse resync operation, the data is replicated from the active source volume to the read-only destination volume.




Steps

1. Depending on the System Manager version that you are running, perform one of the following steps:
 - ONTAP 9.4 or earlier: Click **Protection > Relationships**.
 - Beginning with ONTAP 9.5: Click **Protection > Volume Relationships**.
2. Select the SnapMirror relationship between the source and the destination volumes.
3. Click **Operations > Quiesce**.
4. Select the confirmation check box, and then click **Quiesce**.
5. Click **Operations > Break**.
6. Select the confirmation check box, and then click **Break**.



7. Click **Operations > Reverse Resync**.
8. Select the confirmation check box, and then click **Reverse Resync**.



Beginning with ONTAP 9.3, the SnapMirror policy of the relationship is set to `MirrorAllSnapshots` and the SnapMirror schedule is set to `None`.

If you are running ONTAP 9.2 or earlier, the SnapMirror policy of the relationship is set to `DPDefault` and the SnapMirror schedule is set to `None`.

9. Navigate to the source volume in the volumes page, and verify that the SnapMirror relationship you created is listed and the relationship state is `Snapmirrored`.
10. On the destination cluster, specify a SnapMirror policy and schedule that match the protection configuration of the original SnapMirror relationship for the new SnapMirror relationship:
 - a. Depending on the System Manager version that you are running, perform one of the following steps:
 - ONTAP 9.4 or earlier: Click **Protection > Relationships**.
 - Beginning with ONTAP 9.5: Click **Protection > Volume Relationships**.
 - b. Select the SnapMirror relationship between the reactivated source and the destination volumes, and then click **Edit**.
 - c. Select the SnapMirror policy and schedule, and then click **OK**.

Results

The source volume has read/write access and is protected by the destination volume.

Volume disaster recovery preparation

Volume disaster recovery preparation overview

You can quickly protect a source volume on a peered ONTAP cluster in preparation for disaster recovery. You should use this procedure if you want to configure and monitor SnapMirror relationships between peered clusters for volume disaster recovery and do

not need a lot of conceptual background for the tasks.

SnapMirror provides scheduled asynchronous, block-level data protection. SnapMirror replicates Snapshot copies and can replicate NAS or SAN volumes on which deduplication, data compression, or both are run, including volumes containing qtrees and LUNs. SnapMirror configuration information is stored in a database that ONTAP replicates to all the nodes in the cluster.

Use this procedure if you want to create SnapMirror relationships for volume-level disaster recovery in the following way:

- You are working with clusters running ONTAP 9.
- You are a cluster administrator.
- You have configured the cluster peer relationship and the SVM peer relationship.

[Cluster and SVM peering configuration](#)


- You have enabled the SnapMirror license on both the source and the destination clusters.
- You want to use default policies and schedules, and not create custom policies.
- You want to use best practices, not explore every available option (ONTAP 9.7 and earlier).

Other ways to do this in ONTAP

To perform these tasks with...	Refer to...
The redesigned System Manager (available with ONTAP 9.7 and later)	Prepare for mirroring and vaulting
The ONTAP command line interface	Create a cluster peer relationship (ONTAP 9.3 and later)

Volume disaster recovery preparation workflow

Preparing volumes for disaster recovery involves verifying the cluster peer relationship, creating the SnapMirror relationship between volumes residing on peered clusters, setting up the destination SVM for data access, and monitoring the SnapMirror relationship periodically.



Additional documentation is available to help you activate the destination volume to test the disaster recovery setup or when a disaster occurs. You can also learn more about how to reactivate the source volume after the disaster.

[Volume disaster recovery](#)

+ Describes how to quickly activate a destination volume after a disaster and then reactivate the source volume in ONTAP.

Verify the cluster peer relationship and SVM peer relationship

Before you set up a volume for disaster recovery, you must verify that the source and destination clusters are peered and are communicating with each other through the peer relationship.

Procedure

- If you are running ONTAP 9.3 or later, perform the following steps to verify the cluster peer relationship and SVM peer relationship:
 - a. Click **Configuration > Cluster Peers**.
 - b. Verify that the peered cluster is authenticated and is available.

<div> + Create Edit Delete Refresh Manage SVM Permissions </div>					
<input checked="" type="checkbox"/> Peer Cluster	Availability	Authentication Status	Local Cluster IPspace	Peer Cluster Intercluster IP Addresses	Last Updated Time
<input checked="" type="checkbox"/> cluster2	Available	OK	Default	10.237.213.119,10.237.213.127	Nov 27, 2017, 2:13 PM

- c. Click **Configuration > SVM Peers**.

- d. Verify that the destination SVM is peered with the source SVM.
- If you are running ONTAP 9.2 or earlier, perform the following steps to verify the cluster peer relationship and SVM peer relationship:
 - a. Click the **Configurations** tab.
 - b. In the **Cluster Details** pane, click **Cluster Peers**.
 - c. Verify that the peered cluster is authenticated and available.

'Availability' and 'Authentication Status' information might be stale for up to several minutes.

Create	Modify Passphrase	Modify Peer Network Parameters	Delete	Refresh
Peer Cluster	Availability	Authentication Status		
cluster-1	available	ok		

- d. Click the **SVMs** tab and select the source SVM.
- e. In the **Peer Storage Virtual Machines** area, verify the destination SVM is peered with the source SVM.

If you do not see any peered SVM in this area, you can create the SVM peer relationship when creating the SnapMirror relationship.

Creating the SnapMirror relationship (ONTAP 9.2 or earlier)

Create the SnapMirror relationship (Beginning with ONTAP 9.3)

You must create a SnapMirror relationship between the source volume on one cluster and the destination volume on the peered cluster for replicating data for disaster recovery.

Before you begin

- The destination aggregate must have available space.
- Both the clusters must be configured and set up appropriately to meet the requirements of your environment for user access, authentication, and client access.

About this task

You must perform this task from the **source** cluster.

Steps

1. Click **Storage > Volumes**.
2. Select the volume for which you want to create a mirror relationship, and then click **Actions > Protect**.
3. In the **Relationship Type** section, select **Mirror** from the **Relationship Type** drop-down list.
4. In the **Volumes: Protect Volumes** page, provide the following information:
 - a. Select **Mirror** as the relationship type.
 - b. Select the destination cluster, destination SVM, and the suffix for the name of the destination volume.

Only peered SVMs and allowed SVMs are listed under destination SVMs.

- c. Click .

- d. In the **Advanced Options** dialog box, verify that `MirrorAllSnapshots` is set as the protection policy.

`DPDefault` and `MirrorLatest` are the other default protection policies that are available for `SnapMirror` relationships.

- e. Select a protection schedule.

By default, the `hourly` schedule is selected.

- f. Verify that **Yes** is selected for initializing the `SnapVault` relationship.

All of the data protection relationships are initialized by default. Initializing the `SnapMirror` relationship ensures that the destination volume has a baseline to start protecting the source volume.

- g. Click **Apply** to save the changes.

Advanced Options ✕

Protection Policy MirrorAllSnapshots ▼

SnapMirror Labels	Retention Count
sm_created	1
all_source_snapshots	1

Protection Schedule hourly ▼

Every hour at 05 minute(s)

i Initialize Protection ☒ Yes ☐ No

i SnapLock for SnapVault SnapVault SnapLock for SnapVault is not supported for the selected destination or the selected relationship type.

i FabricPool There is no FabricPool assigned to the destination SVM.

Apply

5. Click **Save** to create the `SnapMirror` relationship.

6. Verify that the relationship status of the `SnapMirror` relationship is in the `Snapmirrored` state.

- a. Navigate to the **Volumes** window, and then select the volume that the volume for which you created the `SnapMirror` relationship.
- b. Double-click the volume to view the volume details, and then click **PROTECTION** to view the data protection status of the volume.

Volume: vol_mirror_src

Overview Snapshots Copies Data Protection Storage Efficiency Performance

Health	Destination SVM	Destination Volume	Destination Clu...	Relationship...	Transfer S...	Type	Lag Time	Policy
	svm2	vol_mirror_src_dst	clv00872	Snapmirrored	Idle	Version-Flexible ...	None	MirrorAllSnap...

What to do next

You must make a note of the settings for the source volume such as thin provisioning, deduplication, compression, and autogrow. You can use this information to verify the destination volume settings when you break the SnapMirror relationship.

Create the SnapMirror relationship (ONTAP 9.2 or earlier)

You must create a SnapMirror relationship between the source volume on one cluster and the destination volume on the peered cluster for replicating data for disaster recovery.

Before you begin

- You must have the cluster administrator user name and password for the destination cluster.
- The destination aggregate must have available space.
- Both the clusters must be configured and set up appropriately to meet the requirements of your environment for user access, authentication, and client access.

About this task

You must perform this task from the **source** cluster.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. Click the **Volumes** tab.
4. Select the volume for which you want to create a mirror relationship, and then click **Protect**.

The Create Protection Relationship window is displayed.

5. In the **Relationship Type** section, select **Mirror** from the **Relationship Type** drop-down list.
6. In the **Destination Volume** section, select the peered cluster.
7. Specify the SVM for the destination volume:

If the SVM is...	Then...
Peered	Select the peered SVM from the list.
Not peered	<div>a. Select the SVM.</div> <div>b. Click Authenticate.</div> <div>c. Enter the cluster administrator’s credentials of the peered cluster, and then click Create.</div>

8. Create a new destination volume:
 - a. Select the **New Volume** option.
 - b. Use the default volume name or specify a new volume name.
 - c. Select the destination aggregate.

Destination Volume

Cluster: cluster.1

Storage Virtual Machine: svm2(peered) [Browse...](#)

Volume: ☒ New Volume ☐ Select Volume

Volume name: svm1_svm1_root_mirror

Aggregate: aggr2 [Browse...](#)
387.19 GB available (of 390.21 GB)

Space Reserve (optional): Default

9. In the **Configuration Details** section, select **MirrorAllSnapshots** as the mirror policy.

DPDefault and MirrorLatest are the other default mirror policies that are available for SnapMirror relationships.

10. Select a protection schedule from the list of schedules.
11. Ensure that the **Initialize Relationship** check box is selected, and then click **Create**.

Initializing the SnapMirror relationship ensures that the destination volume has a baseline to start protecting the source volume.

Configuration Details

Mirror Policy: MirrorAllSnapshots [Browse...](#) [Create Policy](#)
SnapMirror labels: sm_created

Schedule: ☒ hourly [Browse...](#) [Create Schedule](#)
Every hour at 05 minute(s)
☐ None

☒ Initialize Relationship

The relationship is initialized by starting a baseline transfer of data from the source volume to the destination volume.

The initialization operation might take some time. The Status section shows the status of each job.

Create Protection Relationship

Source Volume

Cluster: cluster-1
Storage Virtual Machine: svm1
Volume: svm1_root { Used space 844 KB }

Destination Volume

Cluster: cluster-1
Storage Virtual Machine: svm2
Volume: svm1_svm1_root_mirror

Configuration Details

Mirror Policy: DPDefault
Schedule: hourly

Status

Create volume	✓ Completed successfully
Create relationship	✓ Completed successfully
Initialize relationship	✓ Started successfully

12. Verify the relationship status of the SnapMirror relationship:

- Select the volume for which you created the SnapMirror relationship from the **Volumes** list, and then click **Data Protection**.
- In the **Data Protection** tab, verify that the SnapMirror relationship that you created is listed and that the relationship state is Snapmirrored.

Destination Storage Virtual Mach...	Destination Volume	Is Healthy	Relationship State	Transfer Status	Type	Lag Time	Policy
svm2	svm1_svm1_root_mirror	Yes	Snapmirrored	Idle	Mirror	13 mins	DPDefault

What to do next

You must make a note of the settings for the source volume such as thin provisioning, deduplication, compression, and autogrow. You can use this information to verify the destination volume settings when you break the SnapMirror relationship.

Set up the destination SVM for data access

You can minimize data access disruption when activating the destination volume by setting up required configurations such as LIFs, CIFS shares, and export policies for the NAS environment, and LIFs and initiator groups for the SAN environment on the SVM containing the destination volume.

About this task

You must perform this task on the **destination** cluster for the SVM containing the destination volume.

Procedure

- NAS environment:
 - a. Create NAS LIFs.
 - b. Create CIFS shares with the same share names that were used on the source.
 - c. Create appropriate NFS export policies.
 - d. Create appropriate quota rules.
- SAN environment:
 - a. Create SAN LIFs.
 - b. **Optional:** Configure portsets.
 - c. Configure initiator groups.
 - d. For FC, zone the FC switches to enable the SAN clients to access the LIFs.

What to do next

If any changes were made on the SVM containing the source volume, you must replicate the changes manually on the SVM containing the destination volume.

Related information

[ONTAP 9 Documentation Center](#)

Monitor the status of SnapMirror data transfers

You should periodically monitor the status of the SnapMirror relationships to ensure that the SnapMirror data transfers are occurring as per the specified schedule.

About this task

You must perform this task from the **destination** cluster.

Steps

1. Depending on the System Manager version that you are running, perform one of the following steps:
 - ONTAP 9.4 or earlier: Click **Protection > Relationships**.
 - Beginning with ONTAP 9.5: Click **Protection > Volume Relationships**.
2. Select the SnapMirror relationship between the source and the destination volumes, and then verify the status in the **Details** bottom tab.

The Details tab displays the health status of the SnapMirror relationship and shows the transfer errors and lag time.

- The Is Healthy field must display **Yes**.

For most SnapMirror data transfer failures, the field displays **No**. In some failure cases, however, the field continues to display **Yes**. You must check the transfer errors in the Details section to ensure that no data transfer failure occurred.

- The Relationship State field must display **Snapmirrored**.
- The Lag Time must be no more than the transfer schedule interval.

For example, if the transfer schedule is hourly, then the lag time must not be more than an hour.

You should troubleshoot any issues in the SnapMirror relationships.

[NetApp Technical Report 4015: SnapMirror Configuration and Best Practices for ONTAP 9.1, 9.2](#)

Source Location:	source_SVM/Vol1	Is Healthy:	Yes	Transfer Status:	Idle
Destination Location:	dest_SVM:source_SVM_Vol1	Relationship State:	Snapshotred	Current Transfer Type:	None
Source Cluster:	cluster-2	Network Compression Ratio:	Not Applicable	Current Transfer Error:	None
Destination Cluster:	cluster-1			Last Transfer Error:	None
Transfer Schedule:	hourly			Last Transfer Type:	Initialize
Data Transfer Rate:	Unlimited			Latest Snapshot Timestamp:	09/16/2014 23:42:24
Lag Time:	None			Latest Snapshot Copy:	snapmirror.3e21ed5f-31a3-11e4-8dc7-005056974d2d_2147484686.2014-09-16_233529

Volume backup using SnapVault

Volume backup using SnapVault overview

You can quickly configure SnapVault backup relationships between volumes that are located in different clusters. The SnapVault backup contains a set of read-only backup copies, which are located on a destination volume that you can use for restoring data when data is corrupted or lost.

Use this procedure if you want to create SnapVault backup relationships for volumes in the following way:

- You are working with clusters running ONTAP 9.
- You are a cluster administrator.
- You have configured the cluster peer relationship and the SVM peer relationship.

[Cluster and SVM peering configuration](#)

- You must have enabled either the SnapMirror or SnapVault license, after all of the nodes in the cluster have been upgraded to the same version of ONTAP 9.
- You want to use default protection policies and schedules, and not create custom policies.
- You do not want to back up data for a single file or LUN restore.
- You want to use best practices, not explore every available option.
- You do not want to read a lot of conceptual background.
- You want to use System Manager, not the ONTAP command-line interface or an automated scripting tool.
- You want to use the System Manager classic interface for ONTAP 9.7 and earlier releases, not the ONTAP System Manager UI for ONTAP 9.7 and later.

If these assumptions are not correct for your situation, or if you want more conceptual background information, you should see the following resource:

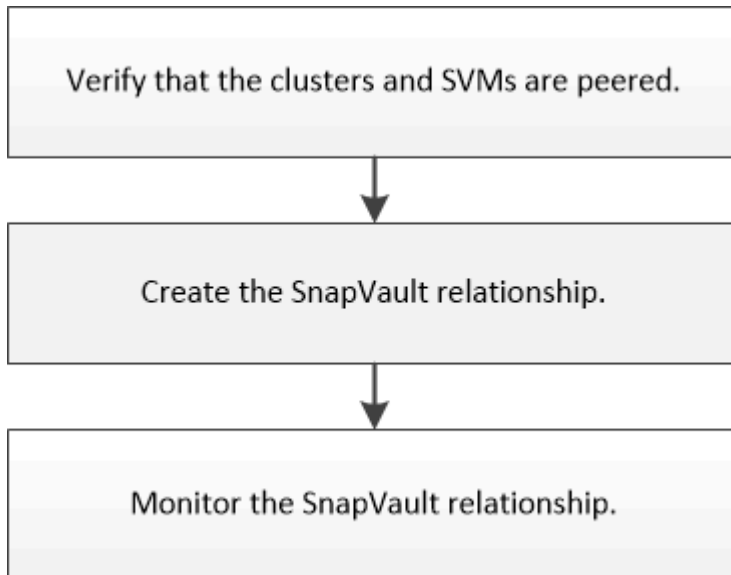
[NetApp Technical Report 4183: SnapVault Best Practices](#)

Other ways to do this in ONTAP

To perform these tasks with...	See this content...
The redesigned System Manager (available with ONTAP 9.7 and later)	Configure mirrors and vaults
The ONTAP command line interface	Create a replication relationship

SnapVault backup configuration workflow

Configuring a SnapVault backup relationship includes verifying the cluster peer relationship, creating the SnapVault relationship between the source and the destination volumes, and monitoring the SnapVault relationship.



Additional documentation is available to help you restore data from a destination volume to test the backed-up data or when the source volume is lost.

- [Volume restore management using SnapVault](#)

Describes how to quickly restore a volume from a SnapVault backup in ONTAP

Verify cluster peer relationship and SVM peer relationship

Before you set up a volume for data protection by using SnapVault technology, you must verify that the source cluster and destination cluster are peered and are communicating with each other through the peer relationship. You must also verify that the source SVM and destination SVM are peered and are communicating with each other through the peer relationship.

About this task

You must perform this task from the **source** cluster.

Procedure


- If you are running ONTAP 9.3 or later, perform the following steps to verify the cluster peer relationship and SVM peer relationship:
 - a. Click **Configuration > Cluster Peers**.
 - b. Verify that the peered cluster is authenticated and is available.



The screenshot shows a table with columns: Peer Cluster, Availability, Authentication Status, Local Cluster IPspace, Peer Cluster Intercluster IP Addresses, and Last Updated Time. A single row is visible for 'cluster2' with status 'Available' and 'OK'.

Peer Cluster	Availability	Authentication Status	Local Cluster IPspace	Peer Cluster Intercluster IP Addresses	Last Updated Time
cluster2	Available	OK	Default	10.237.213.119, 10.237.213.127	Nov 27, 2017, 2:13 PM

- c. Click **Configuration > SVM Peers**.
 - d. Verify that the destination SVM is peered with the source SVM.
- If you are running ONTAP 9.2 or earlier, perform the following steps to verify the cluster peer relationship and SVM peer relationship:
 - a. Click the **Configurations** tab.
 - b. In the **Cluster Details** pane, click **Cluster Peers**.
 - c. Verify that the peered cluster is authenticated and available.



The screenshot shows a table with columns: Peer Cluster, Availability, and Authentication Status. A single row is visible for 'cluster-1' with status 'available' and 'ok'. Above the table is a toolbar with buttons: Create, Modify Passphrase, Modify Peer Network Parameters, Delete, and Refresh. A message at the top states: 'Availability and Authentication Status information might be stale for up to several minutes.'

Peer Cluster	Availability	Authentication Status
cluster-1	available	ok

- d. Click the **SVMs** tab and select the source SVM.
- e. In the **Peer Storage Virtual Machines** area, verify the destination SVM is peered with the source SVM.

If you do not see any peered SVM in this area, you can create the SVM peer relationship when creating the SnapVault relationship.

Creating the SnapVault relationship (ONTAP 9.2 or earlier)

Create a SnapVault relationship (Beginning with ONTAP 9.3)

You must create a SnapVault relationship between the source volume on one cluster and the destination volume on the peered cluster to create a SnapVault backup.

Before you begin

- You must have the cluster administrator user name and password for the destination cluster.
- The destination aggregate must have available space.

About this task

You must perform this task from the **source** cluster.

Steps


1. Click **Storage > Volumes**.
2. Select the volume that you want to back up, and then click **Actions > Protect**.

You can also select multiple source volumes, and then create SnapVault relationships with a single destination volume.

3. In the **Volumes: Protect Volumes** page, provide the following information:
 - a. Select **Vault** from the **Relationship Type** drop-down list.
 - b. Select the destination cluster, destination SVM, and the suffix for the destination volume.

Only peered SVMs and permitted SVMs are listed under destination SVMs.

The destination volume is automatically created. The name of the destination volume is the source volume name appended with the suffix.

- c. Click .
- d. In the **Advanced Options** dialog box, verify that the **Protection Policy** is set as `XDPDefault`.
- e. Select the **Protection Schedule**.

By default, the `daily` schedule is selected.

- f. Verify that **Yes** is selected for initializing the SnapVault relationship.

All data protection relationships are initialized by default.

- g. Click **Apply** to save the changes.

Advanced Options



Protection Policy XDPDefault

SnapMirror Labels	Retention Count
daily	7
weekly	52

Protection Schedule daily

Every Night at 0:10 AM

i Initialize Protection ☒ Yes ☐ No

i SnapLock for SnapVault There are no SnapLock aggregates assigned to the destination SVM.

i FabricPool There is no FabricPool assigned to the destination SVM.

Apply

- In the **Volumes: Protect Volumes** page, click **Validate** to verify whether the volumes have matching SnapMirror labels.
- Click **Save** to create the SnapVault relationship.
- Verify that the status of the SnapVault relationship is in the `Snapmirrored` state.
 - Navigate to the **Volumes** window, and then select the volume that is backed up.
 - Expand the volume and click **PROTECTION** to view the data protection status of the volume.

Volumes on SVM All SVMs

Volume: vol_src [Back to All volumes](#) [Edit](#) [Delete](#) [Actions](#) [Refresh](#)

Overview Snapshots Copies **Data Protection** Storage Efficiency Performance

Health	Destination SVM	Destination Volume	Destination Clu...	Relationsh...	Transfer S...	Type	Lag Time	Policy
	vsl	vol_src_dst	cluster1	Snapmirrored	idle	Vault	29 min(s)	XDPDefault

Create the SnapVault relationship (ONTAP 9.2 or earlier)

You must create a SnapVault relationship between the source volume on one cluster and the destination volume on the peered cluster to create a SnapVault backup.

Before you begin

- You must have the cluster administrator user name and password for the destination cluster.
- The destination aggregate must have available space.

About this task

You must perform this task from the **source** cluster.

Steps

1. Click **Storage > SVMs**.
2. Select the SVM, and then click **SVM Settings**.
3. Click the **Volumes** tab.
4. Select the volume that you want to back up, and then click **Protect**.
5. In the **Create Protection Relationship** dialog box, select **Vault** from the **Relationship Type** drop-down list.
6. In the **Destination Volume** section, select the peered cluster.
7. Specify the SVM for the destination volume:

If the SVM is...	Then...
Peered	Select the peered SVM from the list.
Not peered	<ol style="list-style-type: none">a. Select the SVM.b. Click Authenticate.c. Enter the cluster administrator's credentials of the peered cluster, and then click Create.

8. Create a new destination volume:
 - a. Select the **New Volume** option.
 - b. Use the default volume name or enter a new volume name.
 - c. Select the destination aggregate.
 - d. Ensure that the **Enable dedupe** check box is selected.

Destination Volume

Cluster:

Storage Virtual Machine:

Volume: ☒ New Volume ☐ Select Volume

Volume name:

Aggregate:

☒ Enable dedupe 70.13 GB available (of 70.14 GB)

9. In the **Configuration Details** section, select `XDPDefault` as the protection policy.
10. Select a protection schedule from the list of schedules.
11. Ensure that the **Initialize Relationship** check box is selected to transfer the base Snapshot copy, and then click **Create**

Configuration Details

Vault Policy: XDPDefault Browse... [Create Policy](#)
 Snapshot with labels matching: daily, weekly

Schedule: ☒ weekly Browse... [Create Schedule](#)
 Every Sun at 0:15 am
☐ None

☒ Initialize Relationship

The wizard creates the relationship with the specified vault policy and schedule. The relationship is initialized by starting a baseline transfer of data from the source volume to the destination volume.

The Status section shows the status of each job.

Create Protection Relationship

Source Volume

Cluster: cluster-1
 Storage Virtual Machine: svm1
 Volume: vol_2 { Used space 292 KB }

Destination Volume

Cluster: cluster-1
 Storage Virtual Machine: vs0
 Volume: svm1_vol_2_vault

Configuration Details

Vault Policy: XDPDefault
 Schedule: weekly

Status

Create volume	✓ Completed successfully
Enable dedupe	✓ Completed successfully
Create relationship	✓ Completed successfully
Initialize relationship	✓ Started successfully

Ok

12. Verify that the relationship status of the SnapVault relationship is in the **Snapmirrored** state.
 - a. Select the volume from the Volumes list, and then click **Data Protection**.
 - b. In the **Data Protection** bottom tab, verify that the SnapMirror relationship you created is listed and the

relationship state is Snapmirrored and type is Vault.

Name	Aggregate	Status	Thin Provi...	% Used	Available ...	Total Space	Storage Et...	Is Volume ...	Encrypted
svm1_root	aggr1	Online	No	5	979.56 MB	1 GB	Disabled	No	No
svm2_svm1_...	aggr2	Online	No	5	121.36 MB	128.92 MB	Enabled	No	No
vol1	aggr2	Online	No	0	1017.7 MB	1 GB	Disabled	No	No
vol123	aggr1	Online	Yes	5	1.9 GB	2 GB	Disabled	Yes	No

Destination Stora...	Destination Volu...	Is Healthy	Relationship State	Transfer Status	Type	Lag Time	Policy
svm2	svm1_vol123_vault	Yes	Snapmirrored	Idle	Vault	4 hr(s) 21 min(s)	XDPDefault

DetailsSpace AllocationSnapshot CopiesStorage EfficiencyData ProtectionVolume Move DataPerformance

Monitor the SnapVault relationship

You should periodically monitor the status of the SnapVault relationships to ensure that the data is backed up on the destination volume per the specified schedule.

About this task

You must perform this task from the **destination** cluster.

Steps

- Depending on the System Manager version that you are running, perform one of the following steps:
 - ONTAP 9.4 or earlier: Click **Protection > Relationships**.
 - Beginning with ONTAP 9.5: Click **Protection > Volume Relationships**.
- Select the SnapVault relationship between the source and the destination volumes, and then verify the status in the **Details** bottom tab.

The health status of the SnapVault relationship, any transfer errors, and the lag time are displayed:

- The Is Healthy field must display **Yes**.

For most data transfer failures, the field displays **No**. In some failure cases, however, the field continues to display **Yes**. You must check the transfer errors in the Details section to ensure that no data transfer failure occurred.

- The Relationship State field must display **Snapmirrored**.
- The Lag Time must be not more than the transfer schedule interval.

For example, if the transfer schedule is daily, then the lag time must not be more than a day.

You should troubleshoot any issues in the SnapVault relationships. The troubleshooting procedures for SnapMirror relationships are also applicable to SnapVault relationships.

[NetApp Technical Report 4015: SnapMirror Configuration and Best Practices for ONTAP 9.1, 9.2](#)

Relationships										
Create	Edit	Delete	Operations	Refresh						
Source St...	Source V...	Destinati...	Destinati...	Is Healthy	Relations...	Transfer...	Relationshi...	Lag Time	Policy Na...	Policy Type
svm1	svm1_root	svm1_svm1...	svm2	Yes	Snapmirror...	Idle	Mirror	33 min(s)	DPDefault	Asynchronous Mirr...
svm1	vol123	svm1_vol12...	svm2	Yes	Snapmirror...	Idle	Vault	4 hr(s) 28 m...	XDPDefault	Vault

Source Location:	svm1:vol123	Is Healthy:	Yes	Transfer Status:	Idle
Destination Location:	svm2:svm1_vol123_vault	Relationship State:	Snapmirrored	Current Transfer Type:	None
Source Cluster:	cluster-1	Network Compression Ratio:	Not Applicable	Current Transfer Error:	None
Destination Cluster:	cluster-1			Last Transfer Error:	None
Transfer Schedule:	daily			Last Transfer Type:	Update
Data Transfer Rate:	Unlimited			Latest Snapshot Timestamp:	02/28/2017 00:10:00
Lag Time:	4 hr(s) 28 min(s)			Latest Snapshot Copy:	daily.2017-02-28_0010

Volume restore management using SnapVault

Volume restore using SnapVault overview

You can quickly restore a volume from a SnapVault backup in ONTAP when there is a data loss.

Use this procedure if you want to restore from the vault backup in the following way:

- You are working with clusters running ONTAP 9.
- You are a cluster administrator.
- You have configured the vault relationship following the procedure described in [Volume backup using SnapVault](#)
- You do not want to perform a single file or LUN restore.
- You want to use best practices, not explore every available option.
- You do not want to read a lot of conceptual background.
- You want to use the System Manager classic interface for ONTAP 9.7 and earlier releases, not the ONTAP System Manager UI for ONTAP 9.7 and later.

If these assumptions are not correct for your situation, or if you want more conceptual background information, you should see the following resource:


[NetApp Technical Report 4183: SnapVault Best Practices](#)

Other ways to do this in ONTAP

To perform these tasks with...	See this content...
The redesigned System Manager (available with ONTAP 9.7 and later)	Restore a volume from an earlier Snapshot copy
The ONTAP command line interface	Restore the contents of a volume from a SnapMirror destination

Volume restore workflow

When your source volume is unavailable or data is corrupted, you can perform a restore from a SnapVault backup. Restoring a volume from a SnapVault backup involves selecting the SnapVault destination volume, restoring either to a new volume or existing volume, and verifying the restore operation.



Additional information is available to help you to manage the SnapVault backup relationships and to use other methods of data protection to protect the availability of your data resources.

- [Volume disaster recovery preparation](#)

Describes how to quickly configure a destination volume on a different ONTAP cluster in preparation for disaster recovery.

- [Volume disaster recovery](#)

Describes how to quickly activate a destination volume from a different ONTAP cluster after a disaster, as well as how to restore the SnapMirror relationship to its original state by reactivating the source volume after its recovery.

Identify the SnapVault backup destination volume

You must identify the SnapVault backup destination volume from which you want to restore data when the data in the source volume is corrupted or lost.

About this task

You must perform this task from the **source** cluster.

Steps

1. Enter the URL `https://IP-address-of-cluster-management-LIF` in a web browser and log in to System Manager using your cluster administrator credential.
2. Navigate to the **Volumes** window.
3. Identify the destination volume in the SnapVault relationship and the name of the SVM that contains the volume:
 - ONTAP 9.3 or later: Double-click the volume to view the details, and then click **PROTECTION**.
 - ONTAP 9.2 or earlier: Click the **Data Protection** tab at the bottom of the Volumes window.

Restore data from a SnapVault backup

After selecting the SnapVault backup destination volume, you must perform the restore operation either to a new volume to test the backed-up data or to an existing volume to restore the lost or corrupted data.

About this task

You must perform this task from the **destination** cluster.

Steps

1. Depending on the System Manager version that you are running, perform one of the following steps:
 - ONTAP 9.4 or earlier: Click **Protection > Relationships**.
 - Beginning with ONTAP 9.5: Click **Protection > Volume Relationships**.
2. Select the SVM that contains the SnapVault backup destination volume, and then click **Operations > Restore**.
3. In the **Restore** dialog box, restore the data to the original source volume or a new volume:

If you want to restore to...	Then...
The original source volume	Select Source volume .

If you want to restore to...	Then...
A new volume	<ol style="list-style-type: none"> Select Other volume. Select the peered cluster and the peered SVM for the volume. Select a peered SVM from the list. If the SVM is not peered, create the SVM peer relationship: <ol style="list-style-type: none"> Select the SVM. Click Authenticate. Enter the cluster administrator's credentials of the peered cluster, and then click Create. Select New Volume. If you want to change the default name, displayed in the format <code>destination_SVM_name_destination_volume_name_restore</code>, specify a new name and select the containing aggregate for the volume. Select the Enable dedupe check box.

Restore to

☐ Source volume
☒ Other volume

Cluster:

Storage Virtual Machine:

Volume: ☒ New Volume ☐ Select Volume

Volume name:
Aggregate:

☒ Enable dedupe
517.22 GB available (of 520.28 GB)

- Select either the latest Snapshot copy or select a specific Snapshot copy that you want to restore.
- Select the **OK to restore the volume from the Snapshot copy** check box.
- Select the **Enable Network Compression** check box to compress the data that is being transferred during the restore operation.
- Click **Restore**.

During the restore process, the volume being restored is changed to read-only. After the restore operation finishes, the temporary relationship is removed and the restored volume is changed to read/write.

The image shows a 'Configuration' dialog box with the following elements:

- Configuration** (Title bar)
- Radio button selection:
 - ☒ Latest Snapshot copy: "daily.2014-09-11_0010" 09/11/2014 00:10:00
 - ☐ Select Snapshot copy: [Text field] [Browse...]
- Checkboxes:
 - ☒ OK to restore the volume from the Snapshot copy
 - ☒ Enable Network Compression
- Buttons at the bottom right: [Restore] [Cancel]

8. Click **OK** in the message box.

Verify the restore operation

After performing the restore operation from the SnapVault backup destination volume, you must verify the status of the restore operation on the source cluster.

About this task

You must perform this task from the **source** cluster.

Steps

1. Navigate to the **Volumes** window.
2. Select the source volume in the volumes list and perform one of the following actions, depending on your ONTAP version:
 - Beginning with ONTAP 9.3: Double-click the source volume to view the details, and then click **PROTECTION** to identify the destination volume in the SnapMirror relationship and the name of the SVM that contains the volume.
 - ONTAP 9.2 or earlier: Click the **Data Protection** bottom tab to identify the destination volume in the SnapMirror relationship and the name of the SVM that contains the volume. The Type field displays `Restore temporarily`. After the restore operation is completed, the field displays `Vault`.

You should troubleshoot any issues in the SnapVault relationships. The troubleshooting procedures for SnapMirror relationships are also applicable to SnapVault relationships.

[NetApp Technical Report 4015: SnapMirror Configuration and Best Practices for ONTAP 9.1, 9.2](#)

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.