



Azure administration

Cloud Volumes ONTAP

NetApp
February 27, 2026

Table of Contents

- Azure administration 1
 - Change the Azure VM type for Cloud Volumes ONTAP 1
 - Override CIFS locks for Cloud Volumes ONTAP HA pairs in Azure 1
 - Use an Azure Private Link or service endpoints for Cloud Volumes ONTAP systems 2
 - Overview 3
 - Disable Azure Private Links and use service endpoints instead 3
 - Work with Azure Private Links 4
 - Move an Azure resource group for Cloud Volumes ONTAP in Azure console 7
 - Segregate SnapMirror traffic in Azure 7
 - About SnapMirror traffic segregation in Azure 7
 - Step 1: Create an additional NIC and attach to the destination VM 8
 - Step 2: Create a new IPspace, broadcast domain, and intercluster LIF for the new NIC 10
 - Step 3: Verify cluster peering between the source and destination systems 10
 - Step 4: Create SVM peering between the source and destination system 11
 - Step 5: Create a SnapMirror replication relationship between the source and destination system 12

Azure administration

Change the Azure VM type for Cloud Volumes ONTAP

You can choose from several VM types when you launch Cloud Volumes ONTAP in Microsoft Azure. You can change the VM type at any time if you determine that it is undersized or oversized for your needs.

About this task

- Automatic giveback must be enabled on a Cloud Volumes ONTAP HA pair (this is the default setting). If it isn't, then the operation will fail.

[ONTAP 9 Documentation: Commands for configuring automatic giveback](#)

- Changing the VM type can affect Microsoft Azure service charges.
- The operation restarts Cloud Volumes ONTAP.

For single-node systems, I/O is interrupted.

For HA pairs, the change is nondisruptive. HA pairs continue to serve data.



NetApp Console changes one node at a time by initiating takeover and waiting for give back. NetApp's Quality Assurance team tested both writing and reading files during this process and didn't see any issues on the client side. As connections changed, some retries were observed on the I/O level, but the application layer overcame the rewiring of NFS/CIFS connections.

Steps

1. On the **Systems** page, select the system.
2. On the Overview tab, click the Features panel and then click the pencil icon next to **VM type**.

If you are using a node-based pay-as-you-go (PAYGO) license, you can optionally choose a different license and VM type by clicking the pencil icon next to **License type**.

3. Select a VM type, select the check box to confirm that you understand the implications of the change, and then click **Change**.

Result

Cloud Volumes ONTAP reboots with the new configuration.

Override CIFS locks for Cloud Volumes ONTAP HA pairs in Azure

The organization or account admin can enable a setting in the NetApp Console that prevents issues with Cloud Volumes ONTAP storage giveback during Azure maintenance events. When you enable this setting, Cloud Volumes ONTAP vetoes CIFS locks and resets active CIFS sessions.

About this task

Microsoft Azure schedules periodic maintenance events on its virtual machines. When a maintenance event occurs on a Cloud Volumes ONTAP HA pair, the HA pair initiates storage takeover. If there are active CIFS sessions during this maintenance event, the locks on CIFS files can prevent storage giveback.

If you enable this setting, Cloud Volumes ONTAP will veto the locks and reset the active CIFS sessions. As a result, the HA pair can complete storage giveback during these maintenance events.



This process might be disruptive to CIFS clients. Data that is not committed from CIFS clients could be lost.

Before you begin

You need to create a Console agent before you can change the Console settings. [Learn how.](#)

Steps

1. From the left navigation pane, go to **Administration > Agents**.
2. Click the **...** icon for the Console agent that manages your Cloud Volumes ONTAP system.
3. Select **Cloud Volumes ONTAP Settings**.

The screenshot shows the NetApp Console interface for managing agents. The top navigation bar includes the NetApp logo, 'Console', and dropdown menus for 'Organization' (NetAppNew) and 'Project' (Project-1). The left sidebar has 'Agents' and 'Overview' options. The main content area displays a table of agents with columns for Name, Location, Status, and Deployment Type. A context menu is open for the console agent, with 'Cloud Volumes ONTAP Settings' highlighted by a red rectangle.

Name	Location	Status (1)	Deployment Type
AWSAgent	US East (N. Virginia)	Active	aws
...5678	eastus	Active	...
...tAWS	US East (N. Virginia)	Active	...

4. Under **Azure**, click **Azure CIFS locks for Azure HA systems**.
5. Click the checkbox to enable the feature and then click **Save**.

Use an Azure Private Link or service endpoints for Cloud Volumes ONTAP systems

Cloud Volumes ONTAP uses an Azure Private Link for connections to its associated storage accounts. If needed, you can disable Azure Private Links and use service endpoints instead.

Overview

By default, the NetApp Console enables an Azure Private Link for connections between Cloud Volumes ONTAP and its associated storage accounts. An Azure Private Link secures connections between endpoints in Azure and provides performance benefits.

If required, you can configure Cloud Volumes ONTAP to use service endpoints instead of an Azure Private Link.

With either configuration, the Console always limits network access for connections between Cloud Volumes ONTAP and storage accounts. Network access is limited to the VNet where Cloud Volumes ONTAP is deployed and the VNet where the Console agent is deployed.

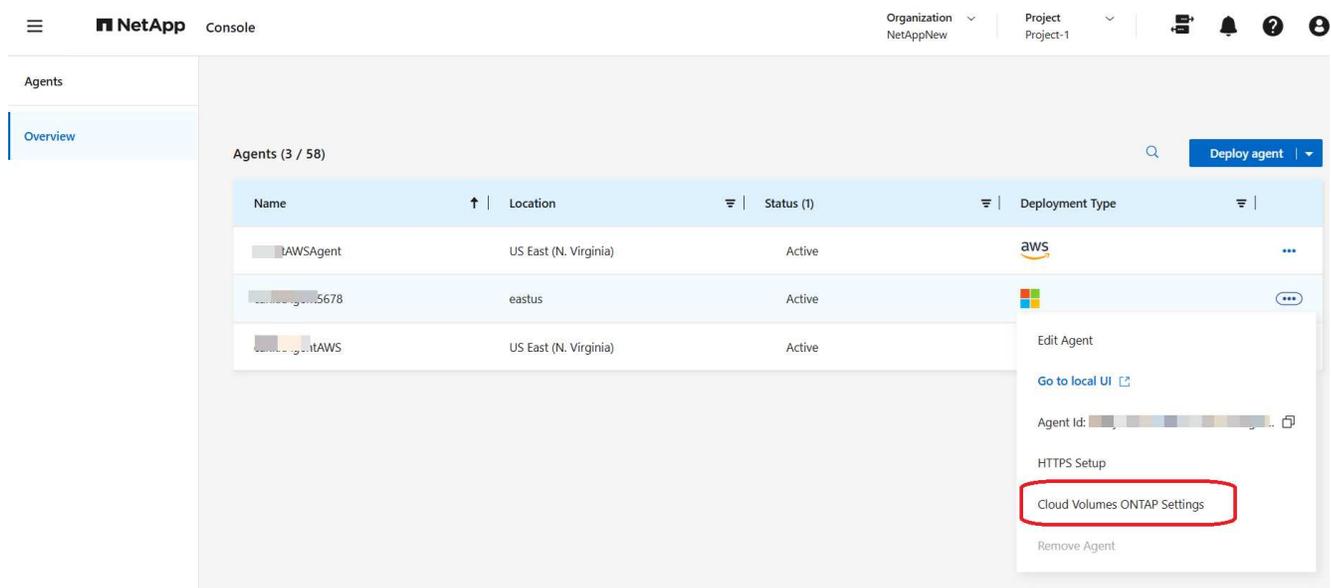
Disable Azure Private Links and use service endpoints instead

If required by your business, you can change a setting in the Console so that it configures Cloud Volumes ONTAP to use service endpoints instead of an Azure Private Link. Changing this setting applies to new Cloud Volumes ONTAP systems that you create. Service endpoints are only supported in [Azure region pairs](#) between the Console agent and Cloud Volumes ONTAP VNets.

The Console agent should be deployed in the same Azure region as the Cloud Volumes ONTAP systems that it manages, or in the [Azure region pair](#) for the Cloud Volumes ONTAP systems.

Steps

1. From the left navigation pane, go to **Administration > Agents**.
2. Click the **...** icon for the Console agent that manages your Cloud Volumes ONTAP system.
3. Select **Cloud Volumes ONTAP Settings**.



4. Under **Azure**, click **Use Azure Private Link**.
5. Deselect **Private Link connection between Cloud Volumes ONTAP and storage accounts**.
6. Click **Save**.

After you finish

If you disabled Azure Private Links and the Console agent uses a proxy server, you must enable direct API

traffic.

[Learn how to enable direct API traffic on the Console agent](#)

Work with Azure Private Links

In most cases, there's nothing that you need to do to set up Azure Private links with Cloud Volumes ONTAP. The Console manages Azure Private Links for you. But if you use an existing Azure Private DNS zone, then you'll need to edit a configuration file.

Requirement for custom DNS

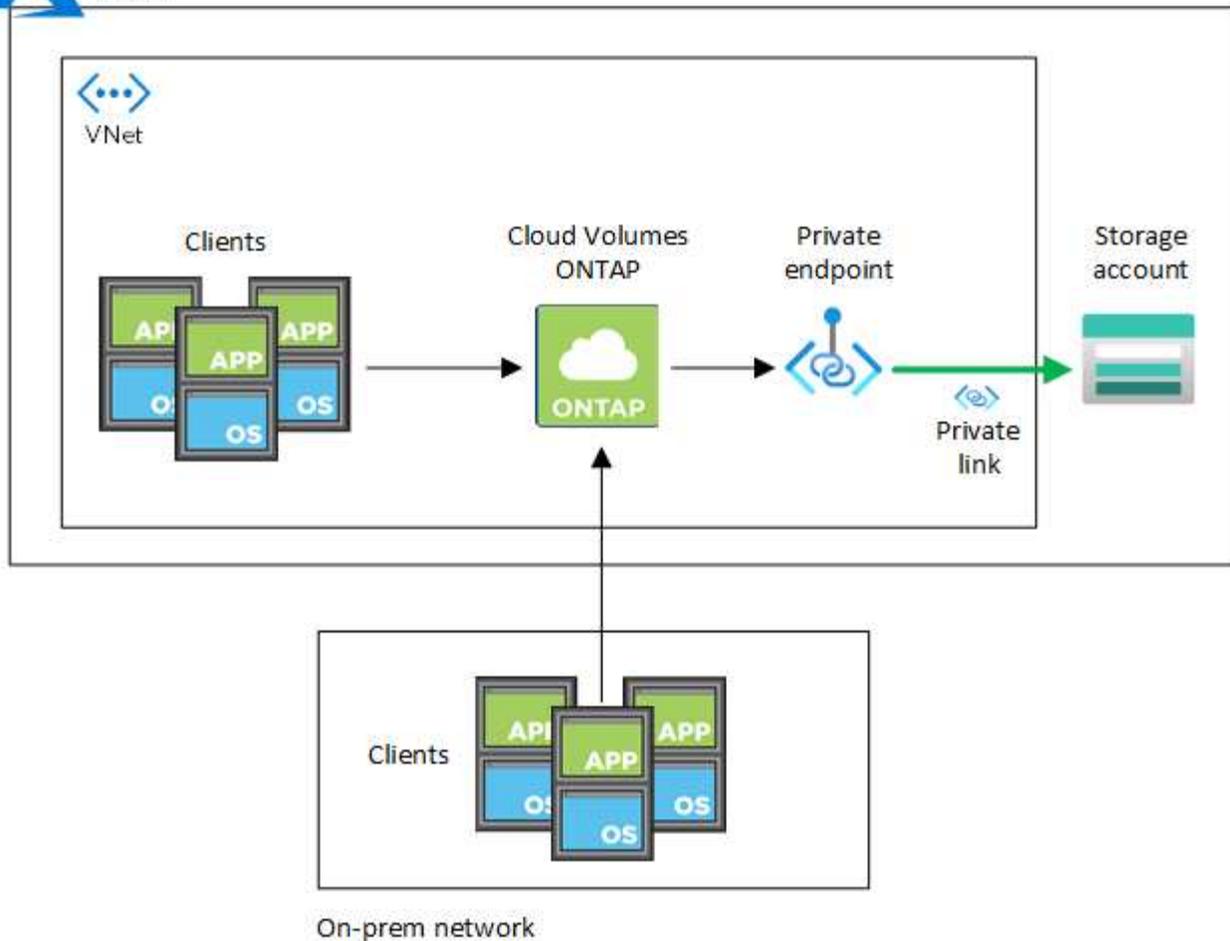
Optionally, if you work with custom DNS, you need to create a conditional forwarder to the Azure private DNS zone from your custom DNS servers. To learn more, refer to [Azure's documentation on using a DNS forwarder](#).

How Private Link connections work

When the Console deploys Cloud Volumes ONTAP in Azure, it creates a private endpoint in the resource group. The private endpoint is associated with storage accounts for Cloud Volumes ONTAP. As a result, access to Cloud Volumes ONTAP storage travels through the Microsoft backbone network.

Client access goes through the private link when clients are within the same VNet as Cloud Volumes ONTAP, within peered VNets, or in your on-premises network when using a private VPN or ExpressRoute connection to the VNet.

Here's an example that shows client access over a private link from within the same VNet and from an on-premises network that has either a private VPN or ExpressRoute connection.



If the Console agent and Cloud Volumes ONTAP systems are deployed in different VNets, then you must set up VNet peering between the VNet where the Console agent is deployed and the VNet where the Cloud Volumes ONTAP systems are deployed.

Provide details about your Azure Private DNS

If you use [Azure Private DNS](#), then you need to modify a configuration file on each Console agent. Otherwise, the Console can't set the Azure Private Link connection between Cloud Volumes ONTAP and its associated storage accounts.

Note that the DNS name must match Azure DNS naming requirements [as shown in Azure documentation](#).

Steps

1. SSH to the Console agent host and log in.
2. Navigate to the `/opt/application/netapp/cloudmanager/docker_occm/data` directory.
3. Edit `app.conf` by adding the `user-private-dns-zone-settings` parameter with the following keyword-value pairs:

```
"user-private-dns-zone-settings" : {
  "resource-group" : "<resource group name of the DNS zone>",
  "subscription" : "<subscription ID>",
  "use-existing" : true,
  "create-private-dns-zone-link" : true
}
```

The `subscription` keyword is required only if the private DNS zone is in a different subscription than that of the Console agent.

4. Save the file and log off the Console agent.

A reboot isn't required.

Enable rollback on failures

If the Console fails to create an Azure Private Link as part of specific actions, it completes the action without the Azure Private Link connection. This can happen when creating a new system (single node or HA pair), or when the following actions occur on an HA pair: creating a new aggregate, adding disks to an existing aggregate, or creating a new storage account when going above 32 TiB.

You can change this default behavior by enabling rollback if the Console fails to create the Azure Private Link. This can help to ensure that you're fully compliant with your company's security regulations.

If you enable rollback, the Console stops the action and rolls back all resources that were created as part of the action.

You can enable rollback through the API or by updating the `app.conf` file.

Enable rollback through the API

Step

1. Use the `PUT /occm/config` API call with the following request body:

```
{ "rollbackOnAzurePrivateLinkFailure": true }
```

Enable rollback by updating `app.conf`

Steps

1. SSH to the host of the Console agent and log in.
2. Navigate to the following directory: `/opt/application/netapp/cloudmanager/docker_occm/data`
3. Edit `app.conf` by adding the following parameter and value:

```
"rollback-on-private-link-failure": true
```

4. Save the file and log off the Console agent.

A reboot isn't required.

Move an Azure resource group for Cloud Volumes ONTAP in Azure console

Cloud Volumes ONTAP supports Azure resource groups moves but the workflow happens in the Azure console only.

You can move a Cloud Volumes ONTAP system from one resource group to a different resource group in Azure within the same Azure subscription. Moving resource groups between different Azure subscriptions is not supported.

Steps

1. Remove the Cloud Volumes ONTAP system. Refer to [Removing Cloud Volumes ONTAP systems](#).
2. Execute the resource group move in the Azure console.

To complete the move, refer to [Move resources to a new resource group or subscription in Microsoft Azure's documentation](#).

3. On the **Systems** page, discover the system.
4. Look for the new resource group in the information for the system.

Result

The system and its resources (VMs, disks, storage accounts, network interfaces, snapshots) are in the new resource group.

Segregate SnapMirror traffic in Azure

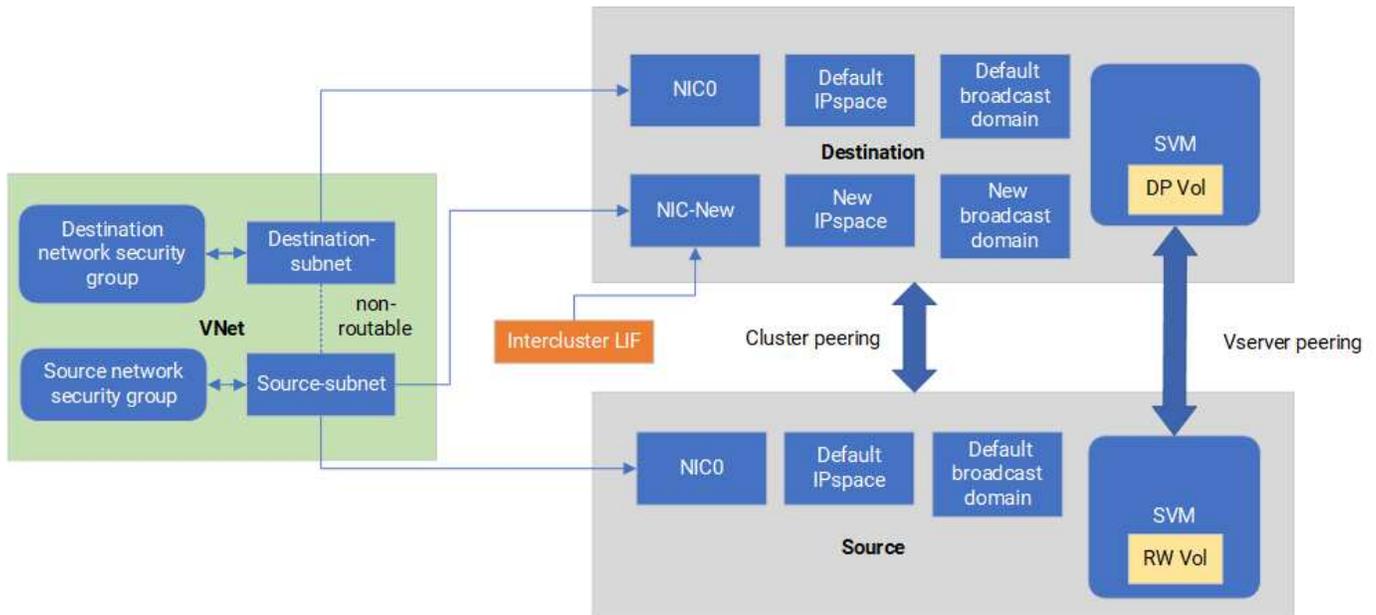
With Cloud Volumes ONTAP in Azure, you can segregate SnapMirror replication traffic from data and management traffic. To segregate SnapMirror replication traffic from your data traffic, you'll add a new network interface card (NIC), an associated intercluster LIF and a non-routable subnet.

About SnapMirror traffic segregation in Azure

By default, the NetApp Console configures all NICs and LIFs in a Cloud Volumes ONTAP deployment on the same subnet. In such configurations, SnapMirror replication traffic and data and management traffic use the same subnet. Segregating SnapMirror traffic leverages an additional subnet that isn't routable to the existing subnet used for data and management traffic.

Figure 1

The following diagrams show the segregation of SnapMirror replication traffic with an additional NIC, an associated intercluster LIF and a non-routable subnet in a single node deployment. An HA pair deployment differs slightly.



Before you begin

Review the following considerations:

- You can only add a single NIC to a Cloud Volumes ONTAP single node or HA-pair deployment (VM instance) for SnapMirror traffic segregation.
- To add a new NIC, the VM instance type you deploy must have an unused NIC.
- The source and destination clusters should have access to the same Virtual Network (VNet). The destination cluster is a Cloud Volumes ONTAP system in Azure. The source cluster can be a Cloud Volumes ONTAP system in Azure or an ONTAP system.

Step 1: Create an additional NIC and attach to the destination VM

This section provides instructions for how to create an additional NIC and attach it to the destination VM. The destination VM is the single node or HA-pair system in Cloud Volumes ONTAP in Azure where you want to set up your additional NIC.

Steps

1. In the ONTAP CLI, stop the node.

```
dest::> halt -node <dest_node-vm>
```

2. In the Azure portal, check that the VM (node) status is stopped.

```
az vm get-instance-view --resource-group <dest-rg> --name <dest-vm>
--query instanceView.statuses[1].displayStatus
```

3. Use the Bash environment in Azure Cloud Shell to stop the node.
 - a. Stop the node.

```
az vm stop --resource-group <dest_node-rg> --name <dest_node-vm>
```

- b. Deallocate the node.

```
az vm deallocate --resource-group <dest_node-rg> --name <dest_node-vm>
```

4. Configure network security group rules to make the two subnets (source cluster subnet and destination cluster subnet) non-routable to each other.

- a. Create the new NIC on the destination VM.
- b. Look up the subnet ID for the source cluster subnet.

```
az network vnet subnet show -g <src_vnet-rg> -n <src_subnet> --vnet -name <vnet> --query id
```

- c. Create the new NIC on the destination VM with the subnet ID for the source cluster subnet. Here you enter the name for the new NIC.

```
az network nic create -g <dest_node-rg> -n <dest_node-vm-nic-new> --subnet <id_from_prev_command> --accelerated-networking true
```

- d. Save the privateIPaddress. This IP address, <new_added_nic_primary_addr>, is used to create an intercluster LIF in [broadcast domain](#), [intercluster LIF for the new NIC](#).

5. Attach the new NIC to the VM.

```
az vm nic add -g <dest_node-rg> --vm-name <dest_node-vm> --nics <dest_node-vm-nic-new>
```

6. Start the VM (node).

```
az vm start --resource-group <dest_node-rg> --name <dest_node-vm>
```

7. In the Azure portal, go to **Networking** and confirm that the new NIC, e.g. nic-new, exists and accelerated networking is enabled.

```
az network nic list --resource-group azure-59806175-60147103-azure-rg --query "[].{NIC: name, VM: virtualMachine.id}"
```

For HA-pair deployments, repeat the steps for the partner node.

Step 2: Create a new IPspace, broadcast domain, and intercluster LIF for the new NIC

A separate IPspace for intercluster LIFs provides logical separation between networking functionality for replication between clusters.

Use the ONTAP CLI for the following steps.

Steps

1. Create the new IPspace (`new_ipspace`).

```
dest::> network ipspace create -ipspace <new_ipspace>
```

2. Create a broadcast domain on the new IPspace (`new_ipspace`) and add the `nic-new` port.

```
dest::> network port show
```

3. For single-node systems, the newly added port is `e0b`. For HA-pair deployments with managed disks, the newly added port is `e0d`. For HA-pair deployments with page blobs, the newly added port is `e0e`. Use the node name not the VM name. Find the node name by running `node show`.

```
dest::> broadcast-domain create -broadcast-domain <new_bd> -mtu 1500  
-ipspace <new_ipspace> -ports <dest_node-cot-vm:e0b>
```

4. Create an intercluster LIF on the new broadcast-domain (`new_bd`) and on the new NIC (`nic-new`).

```
dest::> net int create -vserver <new_ipspace> -lif <new_dest_node-ic-  
lif> -service-policy default-intercluster -address  
<new_added_nic_primary_addr> -home-port <e0b> -home-node <node> -netmask  
<new_netmask_ip> -broadcast-domain <new_bd>
```

5. Verify creation of the new intercluster LIF.

```
dest::> net int show
```

For HA-pair deployments, repeat the steps for the partner node.

Step 3: Verify cluster peering between the source and destination systems

This section provides instructions for how to verify peering between the source and destination systems.

Use the ONTAP CLI for the following steps.

Steps

1. Verify that the intercluster LIF of the destination cluster can ping the intercluster LIF of the source cluster. Because the destination cluster executes this command, the destination IP address is the intercluster LIF IP address on the source.

```
dest::> ping -lif <new_dest_node-ic-lif> -vserver <new_ipspace>
-destination <10.161.189.6>
```

2. Verify that the intercluster LIF of the source cluster can ping the intercluster LIF of the destination cluster. The destination is the IP address of the new NIC created on the destination.

```
src::> ping -lif <src_node-ic-lif> -vserver <src_svm> -destination
<10.161.189.18>
```

For HA-pair deployments, repeat the steps for the partner node.

Step 4: Create SVM peering between the source and destination system

This section provides instructions for how to create SVM peering between the source and destination system.

Use the ONTAP CLI for the following steps.

Steps

1. Create cluster peering on the destination using the source intercluster LIF IP address as the `-peer-addr`s. For HA pairs, list the source intercluster LIF IP address for both nodes as the `-peer-addr`s.

```
dest::> cluster peer create -peer-addr <10.161.189.6> -ip-space
<new_ipspace>
```

2. Enter and confirm the passphrase.
3. Create cluster peering on the source using the destination cluster LIF IP address as the `peer-addr`s. For HA pairs, list the destination intercluster LIF IP address for both nodes as the `-peer-addr`s.

```
src::> cluster peer create -peer-addr <10.161.189.18>
```

4. Enter and confirm the passphrase.
5. Check that the cluster peered.

```
src::> cluster peer show
```

Successful peering shows **Available** in the availability field.

6. Create SVM peering on the destination. Both source and destination SVMs should be data SVMs.

```
dest::> vserver peer create -vserver <dest_svm> -peer-vserver <src_svm>
-peer-cluster <src_cluster> -applications snapmirror``
```

7. Accept SVM peering.

```
src::> vserver peer accept -vserver <src_svm> -peer-vserver <dest_svm>
```

8. Check that the SVM peered.

```
dest::> vserver peer show
```

Peer state shows **peered** and peering applications shows **snapmirror**.

Step 5: Create a SnapMirror replication relationship between the source and destination system

This section provides instructions for how to create a SnapMirror replication relationship between the source and destination system.

To move an existing SnapMirror replication relationship, you must first break the existing SnapMirror replication relationship before you create a new SnapMirror replication relationship.

Use the ONTAP CLI for the following steps.

Steps

1. Create a data protected volume on the destination SVM.

```
dest::> vol create -volume <new_dest_vol> -vserver <dest_svm> -type DP
-size <10GB> -aggregate <aggr1>
```

2. Create the SnapMirror replication relationship on the destination which includes the SnapMirror policy and schedule for the replication.

```
dest::> snapmirror create -source-path src_svm:src_vol -destination
-path dest_svm:new_dest_vol -vserver dest_svm -policy
MirrorAllSnapshots -schedule 5min
```

3. Initialize the SnapMirror replication relationship on the destination.

```
dest::> snapmirror initialize -destination-path <dest_svm:new_dest_vol>
```

4. In the ONTAP CLI, validate the SnapMirror relationship status by running the following command:

```
dest::> snapmirror show
```

The relationship status is `Snapmirrored` and the health of the relationship is `true`.

5. Optional: In the ONTAP CLI, run the following command to view the actions history for the SnapMirror relationship.

```
dest::> snapmirror show-history
```

Optionally, you can mount the source and destination volumes, write a file to the source, and verify the volume is replicating to the destination.

Copyright information

Copyright © 2026 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.