



Volume move management

System Manager Classic

NetApp
May 08, 2024

This PDF was generated from <https://docs.netapp.com/us-en/ontap-sm-classic/volume-move/index.html> on May 08, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Volume move management 1
 - Volume move overview 1
 - Volume move workflow 2

Volume move management

Volume move overview

Using ONTAP System Manager classic interface with ONTAP 9.7 or earlier, you can nondisruptively move a data volume from one node to another node within the same storage virtual machine (SVM) in an ONTAP 9 cluster.

Requirements for using these procedures

Before you use perform these procedures, ensure that the following conditions are met:

- The cluster is running ONTAP 9.
- You have cluster administrator privileges.
- You want to use best practices, not explore every available option.
- You want to use the *Classic* System Manager UI for ONTAP 9.7 and earlier releases, not the ONTAP System Manager UI for ONTAP 9.7 and later.

For some tasks, you must use the ONTAP command-line interface (CLI).

- You know which volume you want to move.

For help in deciding which volumes to move, you can use Active IQ Unified Manager (formerly OnCommand Unified Manager).

- The volume that will be moved is a data volume.
- Any new or repurposed hardware is fully installed and already has aggregates.
- If the cluster has LUNs, all nodes have two paths per LUN.
- Flow control is not enabled on cluster network ports.
- For volumes containing namespaces, the cluster is running ONTAP 9.6 or later.

Volume move is not supported for NVMe configurations running ONTAP 9.5.

Other ways to do this in ONTAP

To complete this task using...	Refer to...
The redesigned System Manager (available with ONTAP 9.7 and later)	Manage volumes
The ONTAP command-line interface	Logical storage management

Alternatives to volume move

Before moving volumes, you should evaluate whether the following approaches are better suited to your situation:

- If you want to nondisruptively upgrade a controller in place, you can consider using aggregate relocation (ARL), which does not require physical data movement.

[High Availability](#)


- If you want to move only a LUN without its containing volume, you can use the LUN move process.

[SAN administration](#)

Volume move workflow

You can move a volume by using the ONTAP System Manager *classic* interface with ONTAP 9.7 or earlier.

Before moving a volume, you should select a method for the volume move operation and plan the timing of the operation. After the move, you might have to update the NDMP backup configuration.



Other ways to do this in ONTAP

To complete this workflow using...	Refer to...
The redesigned System Manager (available with ONTAP 9.7 and later)	Manage volumes
The ONTAP command-line interface	Logical storage management

Plan the method and timing of a volume move

You can use the ONTAP System Manager *classic* interface with ONTAP 9.7 or earlier to move a volume and to decide whether to have a manual cutover. If you need to update LUN reporting nodes, you must follow an advanced procedure in the command-line interface (CLI). Optionally, you can also plan the timing of a volume move.

About this task

The source volume of a SnapMirror or SnapVault relationship can be moved while the volume is being mirrored. SnapMirror services encounter a brief pause during the cutover phase of the volume move job.

The destination volume can also be moved. In the iterative phase, SnapMirror or SnapVault updates and volume move operations run concurrently. When evaluating whether a cutover is possible in the cutover phase, priority between the cutover and SnapMirror or SnapVault updates is determined on a first-come, first-served basis. Until the first operation finishes, other operations are blocked.

Steps

1. Decide whether you require a manual cutover.

Cutover is the moment at which the move operation finishes and ONTAP starts serving data from the volume on the new aggregate. The cutover can occur automatically or you can trigger the cutover manually.

If your company's standard practice requires you to control when changes occur in the storage system, you can manually perform the final cutover of the move operation during a maintenance window.

A cutover does not require an outage, but you can use a maintenance window to control *when* it occurs.



The volume move operation is nondisruptive, regardless of whether you choose automatic or manual cutover.

2. If the volume contains LUNs and the cluster contains four or more nodes, use the CLI to update the LUN reporting nodes if the volume moves to a different HA pair.

If the volume does not contain LUNs or if the cluster contains only two nodes, you can skip this step.

3. **Optional:** Plan a time using the following considerations:

- A volume move operation might take more time than expected because moves are designed to occur nondisruptively in the background in a manner that preserves client access and overall system performance.

For example, ONTAP throttles the resources that are available to the volume move operation.

- If you want the move to occur as quickly as possible, you must select a time with less cluster activity, especially the following activities:
 - I/O operations on the volume
 - Jobs using background resources, for example, when controller CPU usage is less than 50 percent
 - Jobs using the cluster interconnect
- A move cannot be started while the volume is affected by the following operations: volume offline, restrict, or destroy; SnapMirror resync, break, or restore; and Snapshot restore.

You must wait for any of these specific operations to finish before you can start the move.

- While the volume move operation occurs, a MetroCluster switchback cannot occur, although a switchover can occur.
- MetroCluster switchbacks are blocked when volume move operations are in progress for volumes belonging to the switched over site. Switchbacks are not blocked when volume move operations are in progress for volumes local to the surviving site.
- Forced MetroCluster switchovers can occur when volume move operations are in progress.

Related information

[Verifying LUN reporting nodes after moving a volume](#)

Move a volume using System Manager

Using the ONTAP System Manager *classic* interface with ONTAP 9.7 or earlier, you can move a volume by selecting a volume and the destination aggregate, starting the volume move operation, and optionally monitoring the volume move job. When using System Manager, a volume move operation finishes automatically.

Before you begin

You should have reviewed the available space on the source aggregate and destination aggregate before the volume move operation and after the volume move operation.

About this task

A volume move operation is supported only within the same cluster. In addition, you should note that the aggregate you are moving the volume to and the aggregate you are moving the volume from must be in the same storage virtual machine (SVM). A volume move does not disrupt client access.

Steps

1. Navigate to the **Volumes** window.
2. Select the volume that you want to move, and then click **Actions > Move**.
3. Select the destination aggregate, and then start the volume move operation:
 - a. Select a destination aggregate from the list of possible aggregates, which includes only the aggregates that have the required capacity.

You should review the available space, total space, RAID type, and storage type of the aggregates. For example, if the goal is to alter the performance characteristics of the volume, you can focus on aggregates with the desired storage type.

- b. Click **Move**, and then click **Move** again to confirm that you want to proceed with the volume move operation.

When the Move Volume dialog box is displayed, leave the dialog box open if you want to monitor the volume move job.

4. **Optional:** Monitor the volume move job:

- a. In the **Move Volume** dialog box, click the link to the **Job ID** of the volume move job.
- b. Locate the volume move job, and then review the information in the **Status** column.

The job can be in any one of several phases, such as transferring the initial baseline of data or starting a cutover attempt.

241	03/05/2015 07:3...	Volume Move	node1-1	running	Move "vol1" in V...	Cutover Started:(1 of 3 attempts) Transferring final da...
-----	--------------------	-------------	---------	---------	---------------------	--

- c. Click **Refresh** in the **Jobs** window to view the updated job status.

241	03/05/2015 07:3...	Volume Move	node1-1	success	Move "vol1" in V...	Complete: Successful [0]
-----	--------------------	-------------	---------	---------	---------------------	--------------------------

The job status changes to **Complete: Successful** when the volume move operation finishes.

5. If the volume move job enters the cutover deferred phase, perform a manual cutover.

- a. From the **Volumes** window, select the volume for which you initiated the volume move job.
- b. Initiate cutover for the volume:

If you are running...	Perform these steps...
ONTAP 9.3 or later	<ol style="list-style-type: none"> 1. Expand the volume and click the Show More Details link to view more information about the volume. 2. In the Overview tab, click Cutover.
ONTAP 9.2 or earlier	In the Volume Move Details tab, click Cutover .

- c. In the **Cutover** dialog box, click **Advanced Options**.
- d. Specify the cutover action and the cutover duration.

- e. Click **OK**.
6. Repeat [Step 4](#).

Verify LUN reporting nodes after moving a volume

You can add LUN reporting nodes to help maintain optimized LUN paths.

If the volume that you move contains LUNs, and the destination aggregate is on another high-availability (HA) pair, ONTAP automatically adds an HA pair to the Selective LUN Map reporting-nodes list

Before you begin

Two LIFs must be configured: one LIF on the destination node and the other LIF on the HA partner of the destination node.

About this task

This procedure is required only if you move a volume from one HA pair to a different HA pair. If you move a volume to a different node of the same HA pair—for example, if you have a two-node cluster or a MetroCluster configuration—you can skip this procedure.

Steps

1. Verify that the destination node and its partner node are in the reporting-nodes list of the volume. If the nodes are not in the reporting-nodes list, add the destination node and its partner node to the reporting-nodes list of the volume:

```
lun mapping add-reporting-nodes
```

2. Rescan from the host to discover the newly added paths.
3. Add the new paths to your MPIO configuration.
4. Remove the previous LUN owner and its partner node from the reporting-nodes list:

```
lun mapping remove-reporting-nodes -remote-nodes -vserver vservice_name -path  
lun_path -igroup igroup_name
```

5. Rescan the host to verify removal of old paths.

See your host documentation for specific steps to rescan your hosts.

Update LUN reporting nodes after moving a volume

If the volume that you moved contained LUNs and the volume is now on a different HA pair, you should remove all remote nodes from the Selective LUN Map (SLM) reporting-nodes list. The LUN map then contains only the owner node and its HA partner, which ensures that only optimized LUN paths are used.

About this task

This procedure is necessary only if you moved the volume from its HA pair to a different HA pair. If the volume is on a different node of the same HA pair, you can skip this procedure.

Steps

1. Remove all of the remote nodes from the reporting-nodes list by using the `lun mapping remove-`

reporting-nodes command with the `-remote-nodes` parameter.

```
cluster1::> lun mapping remove-reporting-nodes -vserver SVM1 -volume  
vol1 -igroup ig1 -remote-nodes true
```

2. Verify that the LUN map contains only the owner node and its partner by using the `lun mapping show` command with the `-fields reporting-nodes` parameter.

```
cluster1::> lun mapping show -vserver SVM1 -volume vol1 -fields  
reporting-nodes  
vserver  path          igroup    reporting-nodes  
-----  -  
SVM1     /vol/vol1    ig1      cluster1-3,cluster1-4
```

3. Remove stale device entries for the host operating system.
4. Rescan from the host to refresh the host's available paths.

See your host documentation for specific steps to rescan your hosts.

Update NDMP backup after moving a volume

If the volume that you moved was previously backed up to tape using NDMP in a specific configuration, after moving the volume, you can perform one of the following actions to ensure the volume continues to be backed up successfully: create a baseline or migrate the backup LIF to the node containing the moved volume.

About this task

- This procedure is necessary only if the backup application does not support the cluster-aware backup (CAB) extension and the backup process uses node-scoped NDMP.

If the backup application supports CAB and it is configured to use the SVM-scoped NDMP mode, you can skip this procedure.

- You must perform only one of these actions, not both.

Procedure

- From the backup application, create a new baseline.
- Identify the LIF that is configured for the backup process, and then migrate the LIF to the node where the volume now resides.

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.