Clustered Data ONTAP® 8.3

Volume Move Express Guide

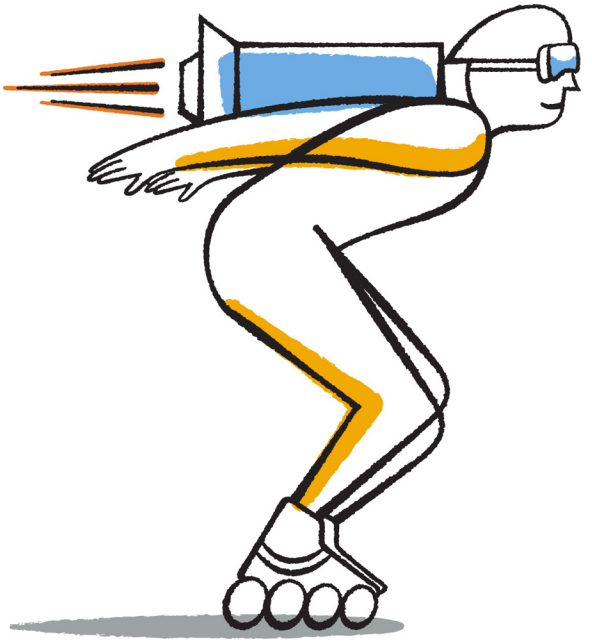

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-09992_A0
April 2015

Contents

Deciding whether to use this guide 4
Volume move workflow 6
Planning the method and timing of a volume move 7
Moving a volume using System Manager 8
Moving a volume using the CLI (advanced) 10
Selecting the destination aggregate 11
Adding LUN reporting nodes before moving a volume 13
Starting the volume move 14
Monitoring the volume move 15
Performing the volume move cutover 16
Updating LUN reporting nodes after moving a volume 17
Updating NDMP backup after moving a volume 17
Where to find additional information 19
Copyright information 20
Trademark information 21
How to send comments about documentation and receive update notification 22
Index 23

Deciding whether to use this guide

You can use this guide to nondisruptively move a volume from one node to another node within the same Storage Virtual Machine (SVM) in a Data ONTAP 8.3 cluster.

Requirements for using this guide

Before you use this guide, you must meet the following requirements:

- The cluster is running Data ONTAP 8.3.
- You have cluster administrator privileges.
- You want to use best practices, not explore every available option.
- You do not want to read a lot of conceptual background.
- You can use either OnCommand System Manager or the Data ONTAP command-line interface.
- You know which volume you want to move.
  For help deciding which volumes to move, you can use OnCommand Performance Manager.
- The volume that will be moved is a data volume, not a root 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.

If this guide is not suitable for your situation, you should see the following documentation instead:

- *Clustered Data ONTAP 8.3 Logical Storage Management Guide*
- *NetApp Knowledgebase*
  The Knowledgebase contains instructions on moving the root volume.
- *NetApp Documentation: OnCommand Workflow Automation (current releases)*
  OnCommand Workflow Automation enables you to run prepackaged workflows that automate management tasks such as the workflows described in Express Guides.

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.

Clustered Data ONTAP 8.3 High-Availability Configuration Guide

- If you want to move only a LUN–without its containing volume–you can use the LUN move mechanism.
  Clustered Data ONTAP 8.3 SAN Administration Guide

Volume move workflow

Before moving a volume, you should select a method and plan the timing. In most cases, you can use the simple System Manager procedure. If your situation requires it, you should use the more advanced CLI method. After the move, you might have to update the NDMP backup configuration.

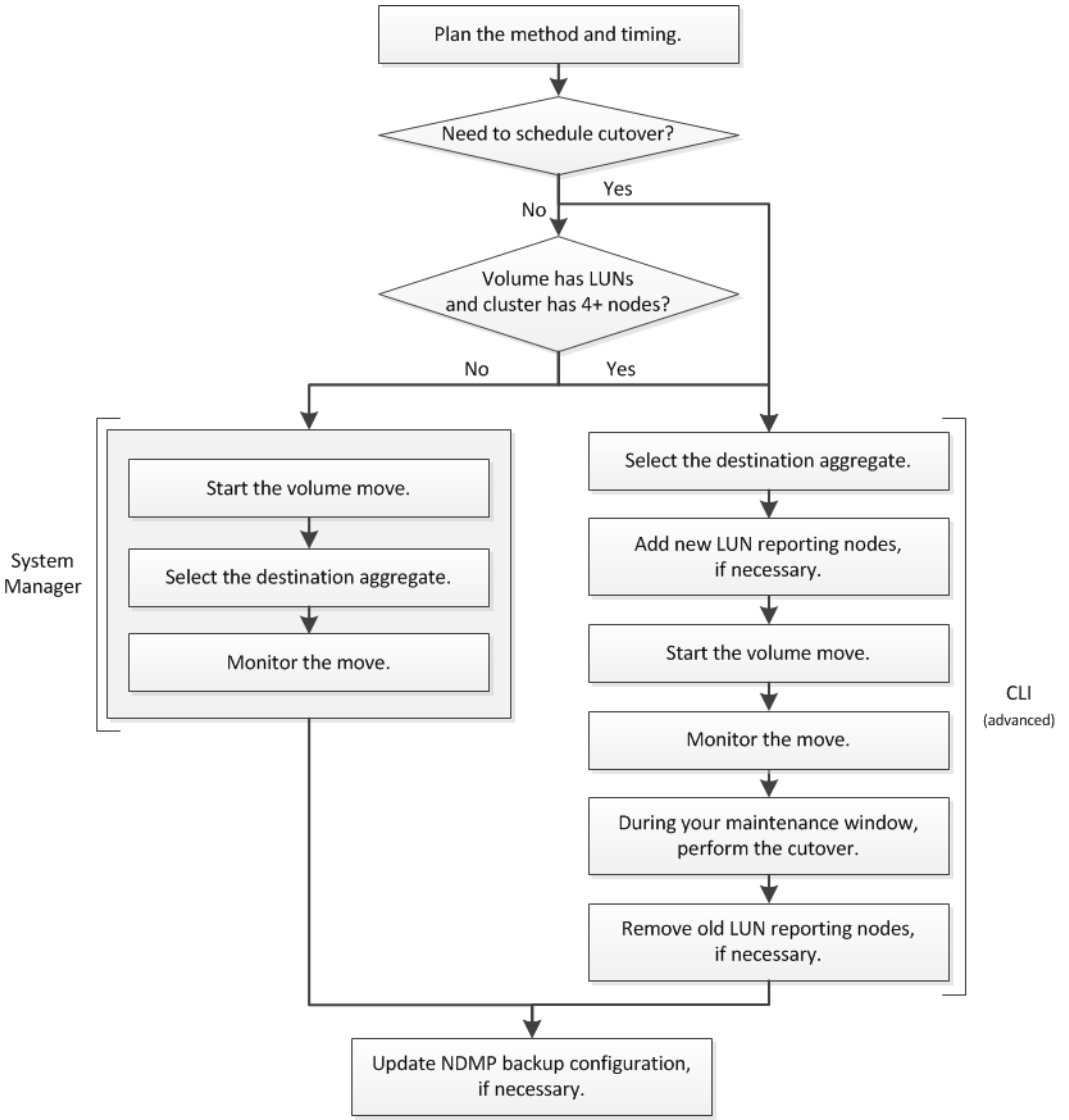

Planning the method and timing of a volume move

In most situations, you can use System Manager to move a volume. If you need to control the final cutover or update LUN reporting nodes, you must follow an advanced procedure in the command-line interface (CLI). You can also optionally plan the timing of a volume move.

Steps

1. Decide whether you have to use the CLI because you require a manual cutover:

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

   - If your company's standard practice requires you to control when changes occur in the storage system, use the CLI to 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.
   - If it is acceptable for the cutover to occur automatically, use System Manager.
     Although you can use the CLI to move a volume with automatic cutover, it is easier to use System Manager.

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

2. Decide whether you have to use the CLI because you must update LUN reporting nodes:

   - If the volume does not contain LUNs or if the cluster contains only two nodes, use System Manager.
   - If the volume contains LUNs and the cluster contains four or more nodes, use the CLI so that you can update LUN reporting nodes if the volume moves to a different HA pair.

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, Data ONTAP throttles the resources available to the volume move operation.
   - If you want the move to occur as quickly as possible, 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.

Moving a volume using System Manager

You can easily move a volume in System Manager by selecting a volume, starting the move, and then optionally monitoring the resulting job. When using System Manager, a volume move operation finishes automatically.

Steps

1. Select the volume:
    a. Expand the **Storage Virtual Machines** hierarchy in the left navigation pane.
    b. Select **Storage** > **Volumes**.
    c. In the **Volumes** window, select the volume you want to move.
    d. Click **Move**.

    The Move Volume dialog box is displayed.

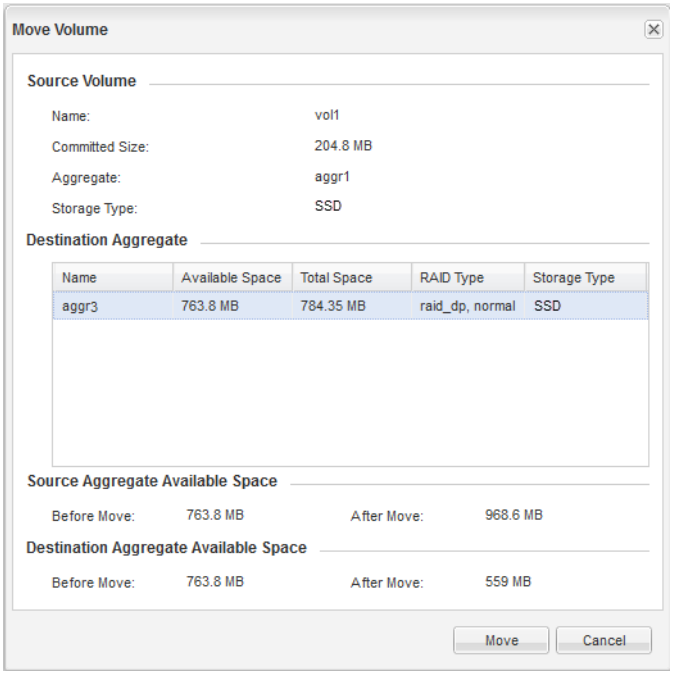

2. Select the destination aggregate and start the move:

   a. Review 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. For example, if the goal is to alter performance characteristics of the volume, you can focus on aggregates with the desired storage type.

   b. Select a destination aggregate from the list.

   c. Review the available space on the source and destination aggregates before the move and after the move.

   d. Click **Move**, and then click **Move** again to confirm the action.

   e. When a new **Move Volume** dialog box is displayed, leave it open so that you can use it for the next step.

3. Optional: Monitor the volume move job:

   a. In the new **Move Volume** dialog box, click the link to the **Job ID** of the volume move job.

      If the dialog box is already closed, you can expand the **Cluster** hierarchy, and then select **Diagnostics > Jobs**.

      The Jobs window is displayed, with the newest jobs at the bottom of the list.

   b. Locate the Volume Move job, and review 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.

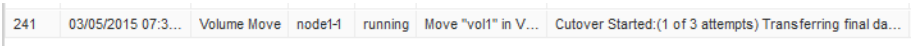

   c. Continue to monitor the job either by remaining on the **Jobs** window and clicking **Refresh** or by periodically returning to the **Jobs** window.

      The job status is `Complete: Successful` when the volume move finishes.

Moving a volume using the CLI (advanced)

When using the CLI to move a volume within a maintenance window or with updates to LUN reporting nodes, you select a destination aggregate, optionally update LUN reporting nodes, and start the move. After monitoring the operation until it is ready, you then trigger the cutover and optionally update LUN reporting nodes.

Steps

1. Selecting the destination aggregate on page 11
2. Adding LUN reporting nodes before moving a volume on page 13
3. Starting the volume move on page 14
4. Monitoring the volume move on page 15
5. Performing the volume move cutover on page 16
6. Updating LUN reporting nodes after moving a volume on page 17

Selecting the destination aggregate

Before you start moving a volume through the command-line interface, you must identify the aggregate to which you will move the volume. The `volume move target-aggr show` command identifies all possible destination aggregates for you to evaluate for your specific situation.

Steps

1. Identify the volume's SVM and its current aggregate (known as the *source* aggregate) by using the `volume show` command with the `-volume` and `-instance` parameters.

   Example

   ```
   cluster1::> volume show -volume vol1 -instance

                                   Vserver Name: SVM1
                                    Volume Name: vol1
                                Aggregate Name: aggr1
                                               ...
   ```

2. Display all aggregates that are potential destinations for this volume by using the `volume move target-aggr show` command.

   The output includes only data aggregates with enough available capacity for the volume.

   Example

   ```
   cluster1::> volume move target-aggr show -vserver SVM1 -volume vol1
   Aggregate Name   Available Size   Storage Type
   --------------   --------------   ------------
   aggr3            763.8GB          SSD
   aggr4            345.5GB          SSD
   ```

3. Select a possible destination aggregate from the list.

4. Ensure that the destination aggregate is assigned to the SVM:

   a. Determine whether the SVM has aggregates assigned to it by using the `vserver show` command with the `-vserver` and `-fields aggr-list` parameters.

      If the list of aggregates is empty, aggregate assignment is not a factor in choosing a destination aggregate.

      Example

      The following output displays an SVM with an empty list of aggregates:

      ```
      cluster1::> vserver show -vserver SVM1 -fields aggr-list
      vserver aggr-list
      ------- ---------
      SVM1    -
      ```

b. If the `aggr-list` contains aggregates and the destination aggregate is not among them, select a different aggregate or use the `vserver add-aggregates` command to assign more aggregates to the SVM.

5. Ensure that the destination aggregate is on the desired node:

a. Identify the node of the source aggregate by using the `storage aggregate show` command with the `-aggregate` and `-fields node` parameters.

Example

```
cluster1::> storage aggregate show -aggregate aggr1 -fields node
aggregate node
--------- -------------
aggr1     node-1-1
```

b. Identify the node of the destination aggregate by using the `storage aggregate show` command with the `-aggregate` and `-fields node` parameters.

Example

```
cluster1::> storage aggregate show -aggregate aggr3 -fields node
aggregate node
--------- -------------
aggr3     node-1-2
```

c. If the aggregate is not on the desired node—for example, if the destination aggregate is on the same node as the source aggregate—select a different aggregate.

6. If the cluster contains more than one disk type, ensure that the destination aggregate is of the desired disk type.

a. Identify the storage type of the source aggregate by using the `storage aggregate show` command with the `-aggregate` and `-fields aggregate-type` parameters.

Example

```
cluster1::> storage aggregate show -aggregate aggr1 -fields aggregate-type
aggregate storage-type
--------- ------------
aggr1     SSD
```

b. Identify the disk type of the destination aggregate by using the `storage aggregate show` command with the `-fields aggregate-type` parameter.

Example

```
cluster1::> storage aggregate show -aggregate aggr3 -fields aggregate-type
aggregate storage-type
--------- ------------
aggr3     SSD
```

c. If the aggregate does not have the desired storage type, select a different aggregate.

For example, if the goal is to alter performance characteristics of the volume, you must select the aggregate with the desired disk type.

Adding LUN reporting nodes before moving a volume

If the volume that you plan to move contains LUNs and the destination aggregate is on another HA pair, you should add the new HA pair to the Selective LUN Map (SLM) reporting-nodes list. This ensures that optimized LUN paths are maintained after the volume is moved.

About this task

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

Steps

1. Add the destination node and its partner node to the volume's reporting-nodes list by using the `lun mapping add-reporting-nodes` command with the `-destination-aggregate` parameter set to the aggregate that you selected as the destination.

Example

```
cluster1::> lun mapping add-reporting-nodes -vserver SVM1 -volume vol1 -lun * -igroup ig1
-destination-aggregate aggr3
```

2. Verify that the LUN map contains four nodes by using the `lun mapping show -fields reporting-nodes` command.

Example

```
cluster1::> lun mapping show -vserver SVM1 -volume vol1 -fields reporting-nodes
vserver  path        igroup   reporting-nodes
-------- ----------- -------- ------------------------------------------
SVM1     /vol/vol1   ig1      cluster1-1,cluster1-2,cluster1-3,cluster1-4
```

If the cluster contains only four nodes and all four are included as reporting-nodes, the `reporting-nodes` field displays a dash (-).

3. Rescan from the host to discover the newly added paths.

4. Add the new paths to your MPIO configuration.

Starting the volume move

You can start the process of moving a volume to another aggregate in a way that defers the final cutover phase until a maintenance window of your choice.

Before you begin

If the volume is a data protection mirror, the mirror relationship must be initialized.

Steps

1. Verify that the volume move is permitted by the system by using the `volume move start` command with the `-perform-validation-only` parameter.

   The command verifies that no current operations prevent the start of the volume move operation. For example, it is not possible to move a volume that is undergoing a SnapMirror resync or Snapshot restore operation. If mutually exclusive operations are running, you must wait for them to finish before continuing.

2. Set the privilege level to advanced by using the `set -privilege advanced` command.

   Advanced privilege is required so that you can use the `-cutover-action` parameter that defers cutover until it is manually triggered. Without the `-cutover-action` parameter, the cutover occurs automatically.

3. Start the volume move operation by using the `volume move start` command with the `-cutover-action wait` parameter.

   #### Example

   ```
   cluster1::*> volume move start -vserver SVM1 -volume vol1 -destination-aggregate aggr3 -
   cutover-action wait

   Warning: When the transfer is complete, trigger the cutover to avoid loss of
            space savings from efficiency. Do you want to proceed?
             {y|n}: y
   [Job 153] Job is queued: Move "vol1" in Vserver "SVM1" to aggregate "aggr3". Use the
   "volume move show -vserver SVM1 -volume vol1" command to view the status of this
   operation.
   ```

4. Return the privilege level to admin by using the `set -privilege admin` command.

Monitoring the volume move

After starting a volume move operation, you should use the `volume move show` command to monitor its progress and identify when the move is ready for manual cutover.

Steps

1. View information about the volume move operation by using the `volume move show` command.

   After an **`initializing`** phase, the move operation enters a **`replicating`** phase.

   Example

   ```
   cluster1::> volume move show
   Vserver   Volume     State    Move Phase Percent-Complete Time-To-Complete
   --------- ---------- -------- ---------- ---------------- ----------------
   SVM1      vol1       healthy  replicating
                                            62%              Tue Mar 10 15:33:16 20
   ```

2. Monitor the progress of the Percent Complete field by repeatedly using the `volume move show` command.

   Example

   ```
   cluster1::> volume move show
   Vserver   Volume     State    Move Phase Percent-Complete Time-To-Complete
   --------- ---------- -------- ---------- ---------------- ----------------
   SVM1      vol1       healthy  replicating
                                            93%              Tue Mar 10 15:33:16 20
   ```

   Tip: You can use the date and time in the Time-To-Complete field as an indicator of when to monitor the operation.

3. When the Move Phase changes to **`cutover_hard_deferred`**, schedule the cutover at a convenient time, such as a maintenance window.

   Example

   ```
   cluster1::> volume move show
   Vserver   Volume     State    Move Phase Percent-Complete Time-To-Complete
   --------- ---------- -------- ---------- ---------------- ----------------
   SVM1      vol1       alert    cutover_hard_deferred
                                            -                -
   ```

   The move operation is ready for you to manually trigger the cutover. In the meantime, any changes made to the volume are periodically replicated to the destination aggregate.

Performing the volume move cutover

When the volume move operation is in a deferred cutover state and you are ready to finish the volume move operation, you can manually trigger the cutover to the new aggregate.

Steps

1. Finish the volume move operation by using the `volume move trigger-cutover` command.

   Example

   ```
   cluster1::> volume move trigger-cutover -vserver SVM1 -volume vol1
   ```

   After the current transfer is complete, cutover begins.

2. Verify that the cutover has started by using the `volume move show` command.

   Example

   The following output shows an operation that is in the **`cutover`** phase and then the **`finishing`** phase.

   ```
   cluster1::> volume move show  -vserver SVM1 -volume vol1
   Vserver   Volume     State    Move Phase Percent-Complete Time-To-Complete
   --------- ---------- -------- ---------- ---------------- ----------------
   SVM1      vol1       healthy  cutover    -                -
   ```

   ```
   cluster1::> volume move show  -vserver SVM1 -volume vol1
   Vserver   Volume     State    Move Phase Percent-Complete Time-To-Complete
   --------- ---------- -------- ---------- ---------------- ----------------
   SVM1      vol1       healthy  finishing    -                -
   ```

3. Verify that the entire operation is finished by repeatedly using the `volume move show` command with the `-vserver`, `-volume`, and `-instance` parameters until the state is **`done`**.

   Example

   The following output shows an operation that is 100% complete, in a **`completed`** move phase, and in a **`done`** state.

   ```
   cluster1::>  volume move show -vserver SVM1 -volume vol1 -instance

                   Vserver Name: SVM1
                    Volume Name: vol1
                                ...
               Detailed Status: Successful
            Percentage Complete: 100%
                                ...
                     Move Phase: completed
                     Move State: done
   ```

If the operation fails, no action is required; the volume remains on the source aggregate. You can review the `Detailed Status` field for information about the cause.

Updating 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 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 remote nodes from the reporting-nodes list by using the `lun mapping remove-reporting-nodes` command with the `-remote-nodes` parameter.

   Example

   ```
   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.

   Example

   ```
   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.

Updating 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 SVM-scoped NDMP mode, you can skip this procedure.
- You must perform only one of these actions, not both.

Choices

- 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.

Where to find additional information

If you want more information about administering clustered Data ONTAP, you can see the Data ONTAP 8 product library.

Related information

NetApp Documentation: Data ONTAP 8 (current releases)

Copyright information

Copyright © 1994–2015 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at *http://www.netapp.com/us/legal/netapptmlist.aspx*.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

How to send comments about documentation and receive update notification

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email to *doccomments@netapp.com*. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

aggregates
  destination for moving volume *11*
audience
  for the guide *4*

B

backup
  NDMP, updating after volume move *17*

C

CLI
  choosing versus System Manager *7*
  moving volumes with *10*
command-line interface
  See CLI
comments
  how to send feedback about documentation *22*
cutover
  deciding between manual and automatic *7*
  deferring when starting a volume move *14*
  performing manually *16*

D

destination
  aggregates, selecting *11*
documentation
  additional information *19*
  how to receive automatic notification of changes to *22*
  how to send feedback about *22*

E

express guides
  additional documentation *19*
  requirements for using the volume move guide *4*
  volume move workflow *6*

F

feedback
  how to send comments about documentation *22*
FlexVol volume move
  See volume move
flowcharts
  volume move workflow *6*

I

information
  how to send feedback about improving documentation *22*

L

LUN map
  See Selective LUN Mapping
LUNs
  adding reporting node before volume move *13*
  choosing System Manager or CLI for moving volumes with *7*

M

monitoring
  volume move using System Manager *8*
  volume move using the CLI *15*
moving
  volumes using System Manager *8*
  volumes using the CLI *10*
  volumes, workflow for *6*
  with deferred cutover *14*

N

NDMP backup
  updating after volume move *17*
nodes
  adding LUN reporting node before volume move *13*
nondisruptive volume move
  See volume move

P

planning
- destination of volume move *11*
- timing and method of volume move *7*

S

Selective LUN Mapping
- adding nodes before volume move *13*
- removing nodes after volume move *17*

SLM
- *See* Selective LUN Mapping

suggestions
- how to send feedback about documentation *22*

System Manager
- choosing versus CLI *7*
- moving volumes with *8*

T

target
- *See* destination

timing
- deciding for volume move operation *7*

twitter
- how to receive automatic notification of documentation changes *22*

V

vol move
- *See* volume move

volume move
- adding LUN reporting nodes using the CLI *13*
- choosing System Manager or CLI *7*
- cutover decision *7*
- cutover using the CLI *16*
- deciding the timing *7*
- increasing speed *7*
- monitoring using System Manager *8*
- monitoring using the CLI *15*
- planning *7*
- removing LUN reporting nodes using the CLI *17*
- selecting destination aggregate using the CLI *11*
- starting through the CLI *14*
- updating NDMP backup after *17*
- using the CLI *10*
- workflow for *6*

volumes
- moving using System Manager *8*
- moving using the CLI *10*
- moving, workflow for *6*
- requirements for using this guide to move *4*

W

workflows
- volume move *6*