



Upgrade by moving volumes

AFF and FAS Controller Upgrade

NetApp

January 14, 2022

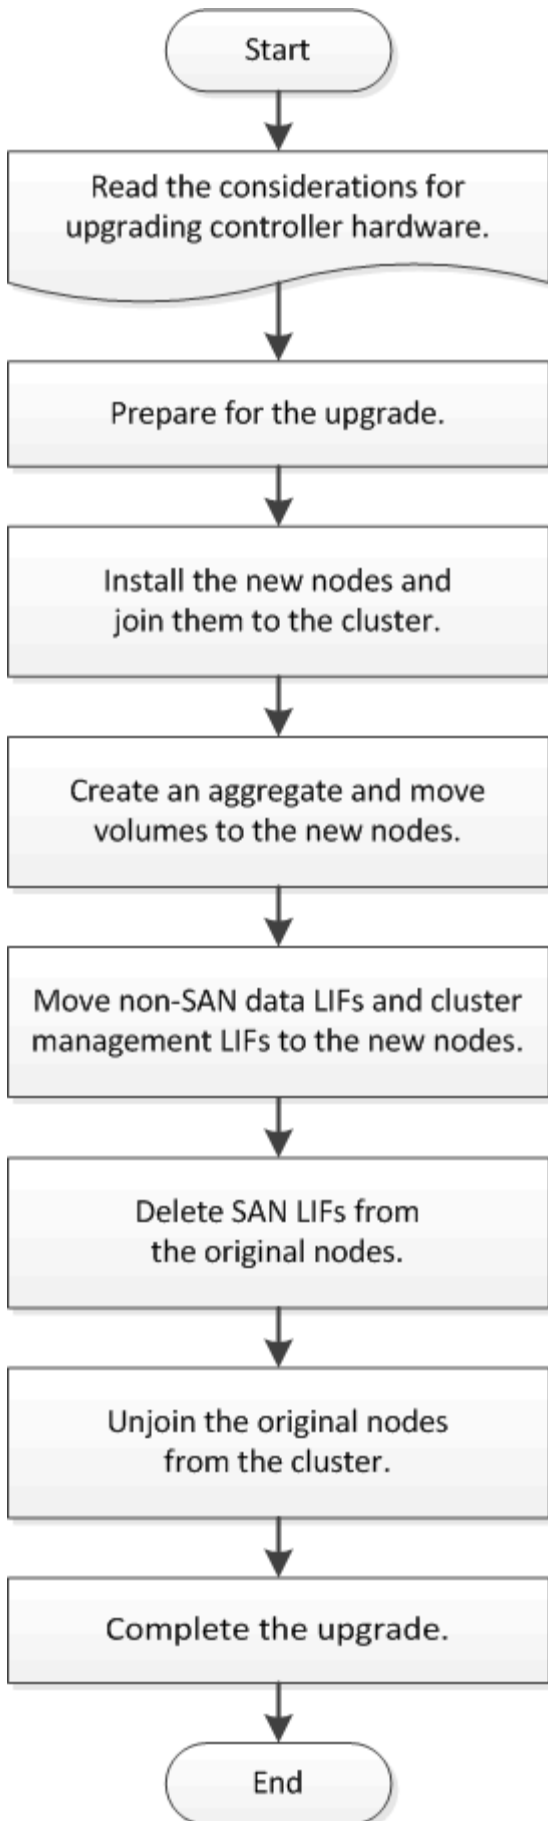
Table of Contents

- Upgrade by moving volumes 1
 - Upgrade controller hardware by moving volumes 1
 - Prepare for the upgrade when moving volumes 3
 - Install the new nodes and join them to the cluster 4
 - Create an aggregate and move volumes to the new nodes. 4
 - Move non-SAN data LIFs and cluster management LIFs to the new nodes 6
 - Move, delete, or create SAN LIFS 7
 - Unjoin the original nodes from the cluster 9
 - Complete the upgrade 10

Upgrade by moving volumes

Upgrade controller hardware by moving volumes

If you are upgrading by moving volumes, you prepare the original nodes and join the new nodes to the cluster. You move volumes to the new nodes, configure LIFs, and unjoin the original nodes from the cluster. Upgrading by moving volumes is a nondisruptive procedure.



Steps

1. [Prepare for the upgrade when moving volumes](#)
2. [Install the new nodes and join them to the cluster](#)
3. [Create an aggregate and move volumes to the new nodes](#)
4. [Move non-SAN data LIFs and cluster management LIFs to the new nodes](#)
5. [Move, delete, or create SAN LIFS](#)
6. [Unjoin the original nodes from the cluster](#)
7. [Complete the upgrade](#)

Prepare for the upgrade when moving volumes

You need to perform a few preparation steps before upgrading controller hardware by moving volumes.

Steps

1. Display the volumes on the original nodes:

```
volume show
```

You use the command output to prepare the list of volumes to move to the new nodes.

2. Display and record license information from the original nodes:

```
system license show
```

3. If you use Storage Encryption on the original nodes and the new nodes have encryption-enabled disks, make sure that the original nodes' disks are correctly keyed:

- a. Display information about self-encrypting disks (SEDs):

```
storage encryption disk show
```

- b. If any disks are associated with a non-manufacture secure ID (non-MSID) key, rekey them to an MSID key:

```
storage encryption disk modify
```

4. If the cluster is currently in a two-node switchless configuration, migrate the cluster to a two-node switched cluster using the type of switch you prefer.

[Migrating to a two-node switched cluster with Cisco cluster switches](#)

[Migrating to a two-node switched cluster with NetApp CN1610 cluster switches](#)

5. Send an AutoSupport message from each original node to inform technical support of the upgrade:

```
system node autosupport invoke -node <node_name> -type all -message "Upgrading  
node_name from <platform_original> to <platform_new>"
```

Install the new nodes and join them to the cluster

You must install the new nodes and join them to the cluster so that you can move volumes from the original nodes.

About this task

When you upgrade controller hardware by moving volumes, both the original nodes and the new nodes must be in the same cluster.

Step

1. Install the new nodes and join them to the cluster:

If the cluster is running...	Follow instructions in...
ONTAP 9.0 or later	Cluster expansion administration
Releases before ONTAP 9.0	Find the Cluster Expansion Express Guide for your version of Data ONTAP 8

Create an aggregate and move volumes to the new nodes

You create at least an aggregate on each of the new nodes to store the volumes you want to move from the original nodes. You must identify an aggregate for each volume and move each volume individually.

Before you begin

Data protection mirror relationships must have been initialized before you can move a volume.

[Find the required data protection procedure.](#)

Steps

1. Create at least one aggregate on each new node:

```
storage aggregate create -aggregate aggr_name -node new_node_name -diskcount integer
```

2. Add the new aggregate to the same SVM as the aggregate on the original node from which you want to move the volumes:

```
vserver add-aggregates
```

Both the new aggregate and the old aggregate from which the volume will be moved must be in the same SVM.

3. Verify that the new aggregate is now assigned to the same SVM as the aggregate on the original node:

```
vserver show -vserversvm_name
```

4. Display information for the volumes that you want to move from the original nodes to the new nodes:

```
volume show -vserver svm_name -node original_node_name
```

You should retain the command output for later reference.

The following example displays volumes on the "vs1" SVM and the "node0" node:

```
cluster::> volume show -vserver vs1 -node node0
Vserver   Volume           Aggregate      State      Type      Size
Available Used%
-----
vs1       clone            aggr1         online    RW        40MB
37.87MB   5%
vs1       voll             aggr1         online    RW        40MB
37.87MB   5%
vs1       vs1root          aggr1         online    RW        20MB
18.88MB   5%
3 entries were displayed.
```

5. Determine an aggregate to which you can move a given volume:

```
volume move target-aggr show -vserver svm_name -volume vol_name
```

The following example shows that the "user_max" volume on the "vs2" SVM can be moved to any of the listed aggregates:

```
cluster::> volume move target-aggr show -vserver vs2 -volume user_max
Aggregate Name   Available Size  Storage Type
-----
aggr2            467.9GB        FCAL
node12a_aggr3   10.34GB        FCAL
node12a_aggr2   10.36GB        FCAL
node12a_aggr1   10.36GB        FCAL
node12a_aggr4   10.36GB        FCAL
5 entries were displayed
```

6. Run a validation check on each volume that you want to move to verify that it can be moved to the specified aggregate:

```
volume move start -vserver svm_name -volume volume_name -destination-aggregate
destination_aggregate_name -perform-validation-only true
```

7. Move the volumes one at a time (advanced privilege level):

```
volume move start -vserver svm_name -volume vol_name -destination-aggregate
destination_aggr_name -cutover-window integer
```

You cannot move the node root volume (vol0). Other volumes, including SVM root volumes, can be moved.

8. Display the outcome of the `volume move` operation to verify that the volumes were moved successfully:

```
volume move show -vserver svm_name -volume vol_name
```

9. If the `volume move` operation does not complete the final phase after multiple attempts, force the move to finish:

```
volume move trigger-cutover -vserver svm_name -volume vol_name -force true
```

Forcing the volume move operation to finish can disrupt client access to the volume that you are moving.

10. Verify that the volumes were moved successfully to the specified SVM and are in the correct aggregate:

```
volume show -vserver svm_name
```

Move non-SAN data LIFs and cluster management LIFs to the new nodes

After you have moved the volumes from the original nodes, you need to migrate the non-SAN data LIFs and cluster-management LIFs from the original nodes to the new nodes.

About this task

You cannot migrate a LIF that is used for copy-offload operations with VMware vStorage APIs for Array Integration (VAI).

Steps

1. From the node where the cluster LIF is hosted, change the home ports for the non-SAN data LIFs from the original nodes to the new nodes:

```
network interface modify -vserver <vserver_name> -lif <lif_name> -home-node <new_node_name> -home-port {<netport|ifgrp>}
```

2. Take one of the following actions:

If you want to migrate...	Then enter...
A specific LIF	<pre>network interface migrate -vserver <vserver_name> -lif <lif_name> -source-node <source_node_name> -destination-node <dest_node_name> -destination -port <dest_port_name></pre>
All the non-SAN data LIFs and cluster-management LIFs	<pre>network interface migrate-all -node <node_name></pre>

The following command migrates a LIF named "datalif1" on the SVM "vs0" to the port "e0d" on "node0b":

```
cluster::> network interface migrate -vserver vs0 -lif datalif1
-destination-node node0b -destination-port e0d
```


The following command migrates all the data and cluster-management LIFs from the current (local) node:

```
cluster::> network interface migrate-all -node local
```

3. Check whether the home node of the cluster-management LIF is on one of the original nodes: `network interface show -lif cluster_mgmt -fields home-node`
4. If the home node of the cluster management LIF is on one of the original nodes, complete the following steps:
 - a. Switch the home node of the cluster-management LIF to one of the new nodes:

```
network interface modify -vserver <cluster_name> -lif cluster_mgmt -home  
-node <new_node_name> -home-port {<netport|ifgrp>}
```

- b. Migrate the cluster-management LIF to one of the new nodes:

```
network interface migrate -vserver <vserver_name> -lif cluster_mgmt  
-destination-node <new_node_name> -destination-port {<netport|ifgrp>}
```

Move, delete, or create SAN LIFS

Move, delete, or create SAN LIFS

Depending on your cluster contents and cluster environment, you need to move, delete, or create SAN LIFs, or re-create deleted SAN LIFs.

- [Considerations for moving SAN LIFS](#)
- [Delete SAN LIFs no longer required from the original nodes](#)
- [Create new SAN LIFs or re-create deleted SAN LIFS](#)

Considerations for moving SAN LIFS

You only need to move the SAN LIFs if you are changing the contents of your cluster, for example, by adding nodes to the cluster or deleting nodes from the cluster. When you move a LIF, you do not have to re-zone your FC fabric or create new iSCSI sessions between the attached hosts of your cluster and the new target interface.

You can move a SAN LIF by using the `network interface modify` command. To move a SAN LIF, you must take the LIF offline, move the LIF to a different home node or port, and then bring it back online in its new location. Asymmetric Logical Unit Access (ALUA) provides redundant paths and automatic path selection as part of any ONTAP SAN solution. Therefore, when the LIF is taken offline for the movement, there is no I/O interruption. The host simply retries and then moves I/O to another LIF.

During the LIF movements, you can nondisruptively perform the following tasks:

- Replace one HA pair of a cluster with an upgraded HA pair in a way that is transparent to the hosts accessing LUN data
- Upgrade a target interface card

- Shift the resources of a storage virtual machine (SVM) from one set of nodes in a cluster to another set of nodes in the same cluster
- When the host server is online, you can move a SAN LUN to a new HA pair without disrupting the host server access to the LUN data

For more information, see the [SAN LIF movement](#) procedure in the *SAN storage management* documentation.

Delete SAN LIFs no longer required from the original nodes

If the cluster is in a SAN environment, you must delete any SAN LIFs that you no longer require from the original nodes before you can unjoin the original nodes from the cluster.

Steps

1. If you have iSCSI initiators, complete the following steps:
 - a. Display a list of active initiators currently connected to an SVM on the original nodes, once for each of the old LIFs:

```
iscsi connection show -vserver <Vserver_name> -lif <old_lif>
```

The following example shows the output of the command with an active initiator connected to SVM vs1:

```
cluster::> iscsi connection show -vserver vs1 -lif data2
          Tpgroup          Conn  Local          Remote          TCP
Recv
Vserver  Name          TSIH  ID    Address          Address          Size
-----  -
vs1      data          9     1    10.229.226.166  10.229.136.188
131400
```

- b. If any initiators are still logged in to an original node, log out of the sessions from your host computer.
2. Display the port set list to determine if any iSCSI or FC LIFs on the original nodes belong to a port set:

```
lun portset show
```

The following example shows output of the `lun portset show` command:

```
cluster:> lun portset show
Virtual
Server  Portset          Protocol  Port Names          Igroups
-----  -
js11    ps0              mixed    LIF1,
          LIF2                igroup1
          ps1              iscsi    LIF3                igroup2
          ps2              fcp      LIF4                -
3 entries were displayed.
```

3. If any iSCSIs or FC LIFs on an original node are members of a port set, remove them from the port set:

```
lun portset remove -vserver <vserver_name> -portset <portset_name> -port-name <lif_name>
```

4. Delete the LIFs on the original nodes:

```
network interface delete -vserver <vserver_name> -lif <lif_name>
```

Create new SAN LIFs or re-create deleted SAN LIFs

Depending on your cluster environment requirements, you might decide to create new SAN LIFs or re-create SAN LIFs that you deleted earlier in this procedure. You can create or re-create SAN LIFs by using the [network interfaces creation](#) procedure in the *Cluster Management Using OnCommand® System Manager* documentation.

Unjoin the original nodes from the cluster

After the volumes have been moved to the new nodes, you unjoin the original nodes from the cluster. When you unjoin a node, the node's configuration is erased and all disks are initialized.

Steps

1. Disable high-availability configuration on the original nodes: `storage failover modify -node <original_node_name> -enabled false`
2. Access the advanced privilege level:

```
set -privilege advanced
```

3. Identify the node that has epsilon:

```
cluster show
```

In the following example, "node0" currently holds epsilon:

```
cluster::*>
Node           Health  Eligibility  Epsilon
-----
node0          true    true         true
node1          true    true         false
node2          true    true         false
node3          true    true         false
```

4. If one of the original nodes holds epsilon, move epsilon to a different node:
 - a. Remove epsilon from the original node: `+ cluster modify -node <original_node_name> -epsilon false`

b. Assign epsilon to a different node: `+ cluster modify -node <new_node_name> -epsilon true`

- From a node that will remain in the cluster, unjoin each original node from the cluster (advanced privilege level):

```
cluster unjoin -node <original_node_name>
```

The system displays a message similar to the following:

```
Warning: This command will unjoin node node_name from the cluster. You
        must unjoin the failover partner as well. After the node is
        successfully unjoined, erase its configuration and initialize
all
        disks by using the "Clean configuration and initialize all
disks (4) "
        option from the boot menu.
        Do you want to continue? {y|n}: y
```

- Enter `y` to continue.

The unjoined node is automatically rebooted and stops at the boot menu.

- From the unjoined node's boot menu, select option **(4) Clean configuration and initialize all disks** to erase the node's configuration and initialize all disks.

The system displays a message similar to the following:

```
Zero disks, reset config and install a new file system?:
This will erase all the data on the disks, are you sure?:
```

- Enter `y` at both prompts.
- If the cluster has only two nodes remaining, configure high availability for the two-node cluster:

```
cluster ha modify -configured true
```

Complete the upgrade

To complete the procedure of upgrading by moving volumes, you need to configure the Service Processor (SP), install new licenses, and set up AutoSupport. You might also need to set up Storage or Volume Encryption and configure the FC or NCA ports.

- Configure the SP on the new nodes as needed:

```
system service-processor network modify
```

- Install new licenses on the new nodes as needed:

```
system license add
```

3. Set up AutoSupport on the new nodes:

```
system node autosupport modify
```

4. From each new node, send a post-upgrade AutoSupport message to technical support:

```
system node autosupport invoke -node node_name -type all -message "node_name  
successfully upgraded from platform_old to platform_new"
```

5. Restore Storage or Volume Encryption functionality by using the appropriate procedure in the [Manage encryption with the CLI](#) content.

Use one of the following procedures, depending on whether you are using onboard or external key management:

- “Restoring onboard key management encryption keys”
- “Restoring external key management encryption keys”

6. If the new nodes have FC ports (onboard or on FC adapters), onboard CNA ports, or a CNA card, configure the FC or CNA ports, enter the following command from the storage system prompt:

```
system node hardware unified-connect modify -node node-name -adapter adapter-  
name -mode {fc|cna} -type {target|initiator}
```

[SAN management with the CLI](#)

You can modify the CNA configuration only when the CNA adapters are offline.

7. Set up a switchless cluster on the new nodes if necessary.

[Migrating to a two-node switched cluster with Cisco cluster switches](#)

[Migrating to a two-node switched cluster with NetApp CN1610 cluster switches](#)

8. As needed, decommission the original systems through the NetApp Support Site to inform NetApp that the systems are no longer in operation and can be removed from support databases:

- a. Log in to the [NetApp Support](#) site.
- b. Click the link **My Installed Systems**.
- c. On the **Installed Systems** page, enter the serial number of the old system in the form and then click **Go!**
- d. On the Decommission Form page, fill out the form and click **Submit**.

Copyright Information

Copyright © 2022 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, 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.