



Upgrade by moving volumes

Upgrade controllers

NetApp

February 22, 2024

Table of Contents

- Upgrade by moving volumes 1
 - Workflow 1
 - Prepare for the upgrade when moving volumes 2
 - Install the new nodes and join them to the cluster 2
 - Move Linux iSCSI hosts to new nodes 3
 - Create an aggregate and move volumes to the new nodes 14
 - Move non-SAN data LIFs and cluster-management LIFs to the new nodes 16
 - Move, delete, or create SAN LIFS 18
 - Unjoin the original nodes from the cluster 20
 - Complete the upgrade 21

Upgrade by moving volumes

Workflow

If you are upgrading controller hardware 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.

1

Prepare for the upgrade when moving volumes

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

2

Install the new nodes and join them to the cluster

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

3

Move Linux iSCSI hosts to the new nodes

Before moving iSCSI SAN volumes to new nodes, you create new iSCSI connections and rescan the iSCSI paths to the new nodes.

4

Create an aggregate and move volumes to the new nodes

You create at least one 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

5

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

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

6

Move, delete, or create SAN LIFs

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

7

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.

8

Complete the upgrade

To complete the procedure of upgrading by moving volumes, you 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.

Prepare for the upgrade when moving volumes

You must 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

Move Linux iSCSI hosts to new nodes

Before moving iSCSI SAN volumes to new nodes, you must create new iSCSI connections and rescan the iSCSI paths to the new nodes.

If you do not need to move iSCSI SAN volumes when upgrading by moving volumes, you can skip this procedure and go to [Create an aggregate and move volumes to the new nodes](#).

About this task

- IPv4 interfaces are created when you set up the new iSCSI connections.
- The host commands and examples are specific to Linux operating systems.

Step 1: Set up new iSCSI connections

To move the iSCSI connections, you set up new iSCSI connections to the new nodes.

Steps

1. Create iSCSI interfaces on the new nodes and check ping connectivity from the iSCSI hosts to the new interfaces on the new nodes.

[Create network interfaces](#)

All iSCSI interfaces from the SVM should be reachable by the iSCSI host.

2. On the iSCSI host, identify the existing iSCSI connections from the host to the old node:

```
iscsiadm -m session
```

```
[root@scspr1789621001 ~]# iscsiadm -m session
tcp: [1] 10.230.68.236:3260,1156 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-flash)
tcp: [2] 10.230.68.237:3260,1158 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-flash)
```

3. On the new node, verify the connections from the new node:

```
iscsi session show -vserver <svm-name>
```

```
node_A_1-new:*> iscsi session show -vserver vsa_1
  Tpgroup Initiator Initiator
Vserver Name TSIH Name ISID Alias
-----
-----
vsa_1 iscsi_lf__n1_p1_ 4 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:01
scspr1789621001.gdl.englab.netapp.com
vsa_1 iscsi_lf__n2_p1_ 4 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:02
scspr1789621001.gdl.englab.netapp.com
2 entries were displayed.
```

4. On the new node, list the iSCSI interfaces in ONTAP for the SVM that contains the interfaces:

```
iscsi interface show -vserver <svm-name>
```

```
sti8200mcchtp001htp_siteA:*> iscsi interface show -vserver vsa_1
  Logical Status Curr Curr
Vserver Interface  TPGT Admin/Oper IP Address Node Port Enabled
-----
-----
vsa_1 iscsi_lf__n1_p1_ 1156 up/up 10.230.68.236 sti8200mcc-htp-001 e0g
true
vsa_1 iscsi_lf__n1_p2_ 1157 up/up fd20:8b1e:b255:805e::78c9 sti8200mcc-
htp-001 e0h true
vsa_1 iscsi_lf__n2_p1_ 1158 up/up 10.230.68.237 sti8200mcc-htp-002 e0g
true
vsa_1 iscsi_lf__n2_p2_ 1159 up/up fd20:8b1e:b255:805e::78ca sti8200mcc-
htp-002 e0h true
vsa_1 iscsi_lf__n3_p1_ 1183 up/up 10.226.43.134 sti8200mccip-htp-005 e0c
true
vsa_1 iscsi_lf__n4_p1_ 1188 up/up 10.226.43.142 sti8200mccip-htp-006 e0c
true
6 entries were displayed.
```

5. On the iSCSI host, run discovery on any one of the iSCSI IP addresses on the SVM to discover the new targets:

```
iscsiadm -m discovery -t sendtargets -p iscsi-ip-address
```

Discovery can be run on any IP address of the SVM, including non-iSCSI interfaces.

```
[root@scspr1789621001 ~]# iscsiadm -m discovery -t sendtargets -p
10.230.68.236:3260
10.230.68.236:3260,1156 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
10.226.43.142:3260,1188 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
10.226.43.134:3260,1183 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
10.230.68.237:3260,1158 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
```

6. On the iSCSI host, login to all the discovered addresses:

```
iscsiadm -m node -L all -T node-address -p portal-address -l
```

```
[root@scspr1789621001 ~]# iscsiadm -m node -L all -T iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 -p
10.230.68.236:3260 -l
Logging in to [iface: default, target: iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6, portal:
10.226.43.142,3260] (multiple)
Logging in to [iface: default, target: iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6, portal:
10.226.43.134,3260] (multiple)
Login to [iface: default, target: iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6, portal:
10.226.43.142,3260] successful.
Login to [iface: default, target: iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6, portal:
10.226.43.134,3260] successful.
```

7. On the iSCSI host, verify the login and connections:

```
iscsiadm -m session
```

```
[root@scspr1789621001 ~]# iscsiadm -m session
tcp: [1] 10.230.68.236:3260,1156 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-flash)
tcp: [2] 10.230.68.237:3260,1158 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-flash)
tcp: [3] 10.226.43.142:3260,1188 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-flash)
```

8. On the new node, verify the login and connection with the host:

```
iscsi initiator show -vserver <svm-name>
```

```
sti8200mcchtp001htp_siteA::*> iscsi initiator show -vserver vsa_1
  Tpgroup Initiator
Vserver Name          TSIH Name          ISID
Igroup Name
-----
vsa_1 iscsi_lf__n1_p1_ 4 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:01 igroup_linux
vsa_1 iscsi_lf__n2_p1_ 4 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:02 igroup_linux
vsa_1 iscsi_lf__n3_p1_ 1 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:04 igroup_linux
vsa_1 iscsi_lf__n4_p1_ 1 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:03 igroup_linux
4 entries were displayed.
```

Result

At the end of this task, the host can see all iSCSI interfaces (on the old and new nodes) and is logged in to all those interfaces.

LUNs and volumes are still physically hosted on the old nodes. Because LUNs are reported only on the old node interfaces, the host will show only paths over the old nodes. To see this, run the `sanlun lun show -p` and `multipath -ll -d` commands on the host and examine the command outputs.


```
[root@scspr1789621001 ~]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host vserver
path path /dev/ host vserver
state      type      node      adapter      LIF
-----
up         primary    sdk       host3         iscsi_lf__n2_p1_
up         secondary  sdh       host2         iscsi_lf__n1_p1_
[root@scspr1789621001 ~]# multipath -ll -d
3600a098038304646513f4f674e52774b dm-5 NETAPP ,LUN C-Mode
size=2.0G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
|  `-- 3:0:0:4 sdk 8:160 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
   `-- 2:0:0:4 sdh 8:112 active ready running
```

Step 2: Add the new nodes as reporting nodes

After setting up the connections to the new nodes, you add the new nodes as the reporting nodes.

Steps

1. On the new node, list reporting nodes for LUNs on the SVM:

```
lun mapping show -vserver <svm-name> -fields reporting-nodes -ostype
linux
```

The following reporting nodes are local nodes as LUNs are physically on old nodes node_A_1-old and node_A_2-old.

```
node_A_1-new::*> lun mapping show -vserver vsa_1 -fields reporting-nodes
-ostype linux
vserver path                                igroup      reporting-nodes
-----
vsa_1    /vol/vsa_1_vol1/lun_linux_2  igroup_linux node_A_1-old,node_A_2-
old
.
.
.
vsa_1    /vol/vsa_1_vol9/lun_linux_19 igroup_linux node_A_1-old,node_A_2-
old
12 entries were displayed.
```

2. On the new node, add reporting nodes:

```
lun mapping add-reporting-nodes -vserver <svm-name> -path
/vol/vsa_1_vol*/lun_linux_* -nodes node1,node2 -igroup <igroup_name>
```

```
node_A_1-new::*> lun mapping add-reporting-nodes -vserver vsa_1 -path
/vol/vsa_1_vol*/lun_linux_* -nodes node_A_1-new,node_A_2-new
-igroup igroup_linux
12 entries were acted on.
```

3. On the new node, verify that the newly added nodes are present:

```
lun mapping show -vserver <svm-name> -fields reporting-nodes -ostype
linux vserver path igroup reporting-nodes
```

```
node_A_1-new:*> lun mapping show -vserver vsa_1 -fields reporting-nodes
-ostype linux vserver path igroup reporting-nodes
-----
-----
-----
vsa_1 /vol/vsa_1_voll/lun_linux_2 igroup_linux node_A_1-old,node_A_2-
old,node_A_1-new,node_A_2-new
vsa_1 /vol/vsa_1_voll/lun_linux_3 igroup_linux node_A_1-old,node_A_2-
old,node_A_1-new,node_A_2-new
.
.
.
12 entries were displayed.
```

4. The `sg3-utils` package must be installed on the Linux host. This prevents a `rescan-scsi-bus.sh` utility not found error when you rescan the Linux host for the newly mapped LUNs using the `rescan-scsi-bus` command.

On the host, verify that the `sg3-utils` package is installed:

- For a Debian based distribution:

```
dpkg -l | grep sg3-utils
```

- For a Red Hat based distribution:

```
rpm -qa | grep sg3-utils
```

If required, install the `sg3-utils` package on the Linux host:

```
sudo apt-get install sg3-utils
```

5. On the host, rescan the SCSI bus on the host and discover the newly added paths:

```
/usr/bin/rescan-scsi-bus.sh -a
```

```
[root@stemgr]# /usr/bin/rescan-scsi-bus.sh -a
Scanning SCSI subsystem for new devices
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
  Scanning for device 2 0 0 0 ...
.
.
.
OLD: Host: scsi5 Channel: 00 Id: 00 Lun: 09
  Vendor: NETAPP Model: LUN C-Mode Rev: 9800
  Type: Direct-Access ANSI SCSI revision: 05
0 new or changed device(s) found.
0 remapped or resized device(s) found.
0 device(s) removed.
```

6. On the iSCSI host, list the newly added paths:

```
sanlun lun show -p
```

Four paths are shown for each LUN.

```
[root@stemgr]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host vserver
path path /dev/ host vserver
state  type      node    adapter  LIF
-----
up      primary    sdk     host3     iscsi_lf__n2_p1_
up      secondary  sdh     host2     iscsi_lf__n1_p1_
up      secondary  sdag    host4     iscsi_lf__n4_p1_
up      secondary  sdah    host5     iscsi_lf__n3_p1_
```

7. On the new node, move the volume/volumes containing LUNs from the old nodes to the new nodes.

```

node_A_1-new::*> vol move start -vserver vsa_1 -volume vsa_1_vol1
-destination-aggregate sti8200mccip_htp_005_aggr1
[Job 1877] Job is queued: Move "vsa_1_vol1" in Vserver "vsa_1" to
aggregate "sti8200mccip_htp_005_aggr1". Use the "volume move show
-vserver
vsa_1 -volume vsa_1_vol1" command to view the status of this operation.
node_A_1-new::*> vol move show
Vserver   Volume           State           Move           Phase           Percent-
Complete  Time-To-Complete
-----
-----
vsa_1     vsa_1_vol1    healthy                initializing      -
-

```

8. When the volume move to the new nodes is complete, verify that the volume is online:

```

volume show -state

```

9. The iSCSI interfaces on the new nodes where the LUN now resides are updated as primary paths. If the primary path is not updated after the volume move, run `/usr/bin/rescan-scsi-bus.sh -a` and `multipath -v3` on the host or simply wait for multipath rescanning to take place.

In the following example, the primary path is a LIF on the new node.

```

[root@stemgr]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host vserver
path path /dev/ host vserver
state type node adapter LIF
-----
up      primary  sdag  host4  iscsi_lf__n4_p1_
up      secondary sdk    host3  iscsi_lf__n2_p1_
up      secondary sdh    host2  iscsi_lf__n1_p1_
up      secondary sdah   host5  iscsi_lf__n3_p1_

```

Step 3: Remove reporting nodes and rescan paths

You must remove the reporting nodes and rescan the paths.

Steps

1. On the new node, remove remote reporting nodes (the new nodes) for the Linux LUNs:

```
lun mapping remove-reporting-nodes -vserver <svm-name> -path * -igroup  
<igroup_name> -remote-nodes true
```

In this case, the remote nodes are old nodes.

```
node_A_1-new::*> lun mapping remove-reporting-nodes -vserver vsa_1 -path  
* -igroup igroup_linux -remote-nodes true  
12 entries were acted on.
```

2. On the new node, check reporting nodes for the LUNs:

```
lun mapping show -vserver <svm-name> -fields reporting-nodes -ostype  
linux
```

```
node_A_1-new::*> lun mapping show -vserver vsa_1 -fields reporting-nodes  
-ostype linux
```

vserver	path	igroup	reporting-nodes
vs_a_1	/vol/vsa_1_vol1/lun_linux_2	igroup_linux	node_A_1-new,node_A_2-new
vs_a_1	/vol/vsa_1_vol1/lun_linux_3	igroup_linux	node_A_1-new,node_A_2-new
vs_a_1	/vol/vsa_1_vol2/lun_linux_4	group_linux	node_A_1-new,node_A_2-new
.			
.			
.			

```
12 entries were displayed.
```

3. The `sg3-utils` package must be installed on the Linux host. This prevents a `rescan-scsi-bus.sh` utility not found error when you rescan the Linux host for the newly mapped LUNs using the `rescan-scsi-bus` command.

On the host, verify that the `sg3-utils` package is installed:

- For a Debian based distribution:

```
dpkg -l | grep sg3-utils
```

- For a Red Hat based distribution:

```
rpm -qa | grep sg3-utils
```

If required, install the `sg3-utils` package on the Linux host:

```
sudo apt-get install sg3-utils
```

4. On the iSCSI host, rescan the SCSI bus:

```
/usr/bin/rescan-scsi-bus.sh -r
```

The paths that are removed are the paths from the old nodes.

```
[root@scspr1789621001 ~]# /usr/bin/rescan-scsi-bus.sh -r
Syncing file systems
Scanning SCSI subsystem for new devices and remove devices that have
disappeared
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
sg0 changed: LU not available (PQual 1)
REM: Host: scsi2 Channel: 00 Id: 00 Lun: 00
DEL: Vendor: NETAPP Model: LUN C-Mode Rev: 9800
Type: Direct-Access ANSI SCSI revision: 05
sg2 changed: LU not available (PQual 1)
.
.
.
OLD: Host: scsi5 Channel: 00 Id: 00 Lun: 09
Vendor: NETAPP Model: LUN C-Mode Rev: 9800
Type: Direct-Access ANSI SCSI revision: 05
0 new or changed device(s) found.
0 remapped or resized device(s) found.
24 device(s) removed.
[2:0:0:0]
[2:0:0:1]
.
.
.
```

5. On the iSCSI host, verify that only paths from the new nodes are visible:

```
sanlun lun show -p
```

```
multipath -ll -d
```

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.](#)

- If you are moving iSCSI SAN volumes, verify that you have [created new iSCSI connections](#).



For each storage virtual machine (SVM), NetApp recommends moving all the non-root volumes in the cluster before moving the root volume, and performing this procedure on one SVM at a time.

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 storage virtual machine (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 -vserver svm_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       vol1             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.



If your storage configuration includes volumes with encryption enabled, follow the steps in [Enable encryption on an existing volume with the volume move start command](#) to move these volumes.

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 must 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 (VAAI).

Steps

1. Log in with the cluster-management LIF, and list all LIFs on the original nodes (comma-separated list):

```
network interface show -curr-node <list_of_original_node_names>
```

2. 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|lifgrp>}
```

3. 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> -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
```

4. 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
```

5. 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

Overview

Depending on your cluster contents and cluster environment, you must 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	fc	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 vservice_name -portset portset_name -port-name
lif_name
```

4. Delete the LIFs on the original nodes:

```
network interface delete -vserver vservice_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`
5. 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
```

6. Enter `y` to continue.

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

7. 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?:
```

8. Enter `y` at both prompts.
9. 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 must 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.

1. Configure the SP on the new nodes as needed:

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

2. 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 © 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.