



FLI offline: Preparation for cutover

ONTAP

NetApp
September 20, 2021

Table of Contents

- FLI offline: Preparation for cutover 1
 - Rebooting the host to validate the system state 1
 - Verifying host LUN path and multipath configuration verification 1
 - Prepare hosts for FLI offline migration 12
 - Presenting source LUNs to ONTAP storage during FLI 12
 - Verifying source LUNs on destination storage for offline FLI 12
 - Configuring migration jobs 14

FLI offline: Preparation for cutover

During pre-migration of foreign LUN imports (FLIs), validate and verify the hosts and source LUN paths. After the host reboots, it is shut down in preparation for the migration.

When migration and remediation are complete, hosts can be brought up connected to the new destination storage, and the applications can be verified by end users.

Rebooting the host to validate the system state

Migration hosts are rebooted prior to making any changes to their configuration. Before proceeding with migration, verify that the system is in a known good state.

To verify that the server configuration is persistent and pristine across reboots, complete the following steps:

Steps

1. Shut down all your open applications.
2. Reboot the host.
3. Review the logs for errors.

Verifying host LUN path and multipath configuration verification

Prior to any migrations, verify that multipathing is correctly configured and working properly. All available paths to LUNs should be active.


Multipath verification for Windows hosts

As part of the Foreign LUN Import (FLI) process, you should verify that multipath is configured and functioning correctly on your hosts.

Complete the following steps for Windows hosts.

Steps

1. Open Disk Management.
 - a. On the Windows desktop, click **Start**.
 - b. In the Start Search field, type `diskmgmt.msc`.
 - c. In the Programs list, click `diskmgmt`.
2. Right-click each disk for which you want to verify the multiple paths and then click **Properties**.
3. On the MPIO tab, in the Select the MPIO policy list, click all the paths that are active.



To verify multipathing using the command line, complete the following steps:

4. Open Windows command prompt.
5. Run `mpclaim.exe --v c:\multipathconfig.txt` to capture multipath configuration.

Multipath verification for Linux hosts

As part of the Foreign LUN Import (FLI) process, you should verify that multipath is configured and functioning correctly on your hosts.

Complete the following steps for Linux hosts.

Step

1. To verify that DM-MP multipath is configured and functioning correctly on a Linux host, run the following commands:
`multipath -ll`

```

mpath2 (360060e801046b96004f2bf4600000012) dm-6 HITACHI,DF600F
[size=2.0G][features=0][hwhandler=0][rw]
\_ round-robin 0 [prio=1][active]
  \_ 0:0:1:2 sdg 8:96 [active][ready]
  \_ 1:0:1:2 sdo 8:224 [active][ready]
\_ round-robin 0 [prio=0][enabled]
  \_ 0:0:0:2 sdc 8:32 [active][ready]
  \_ 1:0:0:2 sdk 8:160 [active][ready]
mpath1 (360060e801046b96004f2bf4600000011) dm-5 HITACHI,DF600F
[size=2.0G][features=0][hwhandler=0][rw]
\_ round-robin 0 [prio=1][active]
  \_ 0:0:0:1 sdb 8:16 [active][ready]
  \_ 1:0:0:1 sdj 8:144 [active][ready]
\_ round-robin 0 [prio=0][enabled]
  \_ 0:0:1:1 sdf 8:80 [active][ready]
  \_ 1:0:1:1 sdn 8:208 [active][ready]
mpath0 (360060e801046b96004f2bf4600000010) dm-0 HITACHI,DF600F
[size=20G][features=0][hwhandler=0][rw]
\_ round-robin 0 [prio=1][active]
  \_ 0:0:1:0 sde 8:64 [active][ready]
  \_ 1:0:1:0 sdm 8:192 [active][ready]
\_ round-robin 0 [prio=0][enabled]
  \_ 0:0:0:0 sda 8:0 [active][ready]
  \_ 1:0:0:0 sdi 8:128 [active][ready]
mpath3 (360060e801046b96004f2bf4600000013) dm-7 HITACHI,DF600F
[size=3.0G][features=0][hwhandler=0][rw]
\_ round-robin 0 [prio=1][active]
  \_ 0:0:0:3 sdd 8:48 [active][ready]
  \_ 1:0:0:3 sdl 8:176 [active][ready]
\_ round-robin 0 [prio=0][enabled]
  \_ 0:0:1:3 sdh 8:112 [active][ready]
  \_ 1:0:1:3 sdp 8:240 [active][ready]
[root@dm-rx200s6-22 ~]#

```


Multipath verification for ESXi hosts

As part of the Foreign LUN Import (FLI) process, you should verify that multipath is configured and functioning correctly on your hosts.


Complete the following steps for ESXi hosts.

Steps

1. Determine ESXi and virtual machine using VMware vSphere Client.



2. Determine SAN LUNs to be migrated using vSphere Client.



3. Determine VMFS and RDM (vfat) volumes to be migrated: `esxcli storage filesystem list`

```

Mount Point                                Volume Name
UUID                                       Mounted Type      Size
Free
-----
/vmfs/volumes/538400f6-3486df59-52e5-00262d04d700  BootLun_datastore
538400f6-3486df59-52e5-00262d04d700      true  VMFS-5  13421772800
12486443008
/vmfs/volumes/53843dea-5449e4f7-88e0-00262d04d700  VM_datastore
53843dea-5449e4f7-88e0-00262d04d700      true  VMFS-5  42681237504
6208618496
/vmfs/volumes/538400f6-781de9f7-c321-00262d04d700
538400f6-781de9f7-c321-00262d04d700      true  vfat    4293591040
4269670400
/vmfs/volumes/c49aad7f-afbab687-b54e-065116d72e55
c49aad7f-afbab687-b54e-065116d72e55      true  vfat    261853184
77844480
/vmfs/volumes/270b9371-8fbedc2b-1f3b-47293e2ce0da
270b9371-8fbedc2b-1f3b-47293e2ce0da      true  vfat    261853184
261844992
/vmfs/volumes/538400ef-647023fa-edef-00262d04d700
538400ef-647023fa-edef-00262d04d700      true  vfat    299712512
99147776
~ #

```



In case of VMFS with extends \(\spanned VMFS\), all LUNs that are part of the span should be migrated. To show all the extends in the GUI, go to Configuration \> Hardware \> Storage and click datastore to select the Properties link.



Post-migration, while adding them back to storage, you will see multiple LUN entries with the same VMFS label. In this scenario you should ask the customer to select only the entry marked as head.

4. Determine the LUN and size to be migrated: `esxcfg-scsidevs -c`

```
Device UID                               Device Type           Console Device
Size      Multipath PluginDisplay Name
mpx.vmhba36:C0:T0:L0                     CD-ROM
/vmfs/devices/cdrom/mpx.vmhba36:C0:T0:L0          0MB      NMP
Local Optiarc CD-ROM (mpx.vmhba36:C0:T0:L0)
naa.60060e801046b96004f2bf4600000014   Direct-Access
/vmfs/devices/disks/naa.60060e801046b96004f2bf4600000014  20480MB  NMP
HITACHI Fibre Channel Disk (naa.60060e801046b96004f2bf4600000014)
naa.60060e801046b96004f2bf4600000015   Direct-Access
/vmfs/devices/disks/naa.60060e801046b96004f2bf4600000015  40960MB  NMP
HITACHI Fibre Channel Disk (naa.60060e801046b96004f2bf4600000015)
~~~~~ Output truncated ~~~~~
~ #
```

5. Identify raw device mapping (RDM) LUNs to be migrated.

6. Find RDM devices: `find /vmfs/volumes -name *-rdm*`

```
/vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Windows2003/Windows2003_1-rdmp.vmdk
/vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Windows2003/Windows2003_2-rdm.vmdk
/vmfs/volumes/53843dea-5449e4f7-88e0-00262d04d700/Linux/Linux_1-rdm.vmdk
/vmfs/volumes/53843dea-5449e4f7-88e0-00262d04d700/Solaris10/Solaris10_1-
rdmp.vmdk
```

7. Remove `-rdmp` and `-rdm` from preceding output and run the `vmkfstools` command to find vml mapping and RDM type.

```

# vmkfstools -q /vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Windows2003/Windows2003_1.vmdk
vmkfstools -q /vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Windows2003/Windows2003_1.vmdk
Disk /vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Windows2003/Windows2003_1.vmdk is a Passthrough Raw Device
Mapping
Maps to: vml.020002000060060e801046b96004f2bf4600000016444636303046
~ # vmkfstools -q /vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Windows2003/Windows2003_2.vmdk
Disk /vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Windows2003/Windows2003_2.vmdk is a Non-passthrough Raw
Device Mapping
Maps to: vml.020003000060060e801046b96004f2bf4600000017444636303046
~ # vmkfstools -q /vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Linux/Linux_1.vmdk
Disk /vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Linux/Linux_1.vmdk is a Non-passthrough Raw Device Mapping
Maps to: vml.020005000060060e801046b96004f2bf4600000019444636303046
~ # vmkfstools -q /vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Solaris10/Solaris10_1.vmdk
Disk /vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Solaris10/Solaris10_1.vmdk is a Passthrough Raw Device
Mapping
Maps to: vml.020004000060060e801046b96004f2bf4600000018444636303046
~ #

```



Passthrough is RDM with physical \(\RDMP\), and nonpassthrough is RDM with virtual \(\RDMV\). VMs with virtual RDMs and VM Snapshot copies will break after migration due to VM Snapshot delta vmdk pointing to an RDM that has a stale naa ID. So before migration, ask the customer to remove all Snapshot copies in such VMs. Right-click VM and click the Snapshot --> Snapshot Manager Delete All button. Refer to NetApp KB 3013935 for details about hardware-accelerated locking for VMware on NetApp storage.

8. Identify LUN naa to RDM device mapping.


```

~ # esxcfg-scsidevs -u | grep
vml.020002000060060e801046b96004f2bf4600000016444636303046
naa.60060e801046b96004f2bf4600000016
vml.020002000060060e801046b96004f2bf4600000016444636303046
~ # esxcfg-scsidevs -u | grep
vml.020003000060060e801046b96004f2bf4600000017444636303046
naa.60060e801046b96004f2bf4600000017
vml.020003000060060e801046b96004f2bf4600000017444636303046
~ # esxcfg-scsidevs -u | grep
vml.020005000060060e801046b96004f2bf4600000019444636303046
naa.60060e801046b96004f2bf4600000019
vml.020005000060060e801046b96004f2bf4600000019444636303046
~ # esxcfg-scsidevs -u | grep
vml.020004000060060e801046b96004f2bf4600000018444636303046
naa.60060e801046b96004f2bf4600000018
vml.020004000060060e801046b96004f2bf4600000018444636303046
~ #

```

9. Determine virtual machine configuration: `esxcli storage filesystem list | grep VMFS`

```

/vmfs/volumes/538400f6-3486df59-52e5-00262d04d700  BootLun_datastore
538400f6-3486df59-52e5-00262d04d700      true  VMFS-5  13421772800
12486443008
/vmfs/volumes/53843dea-5449e4f7-88e0-00262d04d700  VM_datastore
53843dea-5449e4f7-88e0-00262d04d700      true  VMFS-5  42681237504
6208618496
~ #

```

10. Record the UUID of the datastore.

11. Make a copy of `/etc/vmware/hostd/vmlInventory.xml` and note the contents of file and vmx config path.

```

~ # cp /etc/vmware/hostd/vmInventory.xml
/etc/vmware/hostd/vmInventory.xml.bef_mig
~ # cat /etc/vmware/hostd/vmInventory.xml
<ConfigRoot>
  <ConfigEntry id="0001">
    <objID>2</objID>
    <vmxCfgPath>/vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Windows2003/Windows2003.vmx</vmxCfgPath>
  </ConfigEntry>
  <ConfigEntry id="0004">
    <objID>5</objID>
    <vmxCfgPath>/vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Linux/Linux.vmx</vmxCfgPath>
  </ConfigEntry>
  <ConfigEntry id="0005">
    <objID>6</objID>
    <vmxCfgPath>/vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Solaris10/Solaris10.vmx</vmxCfgPath>
  </ConfigEntry>
</ConfigRoot>

```

12. Identify the virtual machine hard disks.

This information is required post-migration to add the removed RDM devices in order.

```

~ # grep fileName /vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Windows2003/Windows2003.vmx
scsi0:0.fileName = "Windows2003.vmdk"
scsi0:1.fileName = "Windows2003_1.vmdk"
scsi0:2.fileName = "Windows2003_2.vmdk"
~ # grep fileName /vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Linux/Linux.vmx
scsi0:0.fileName = "Linux.vmdk"
scsi0:1.fileName = "Linux_1.vmdk"
~ # grep fileName /vmfs/volumes/53843dea-5449e4f7-88e0-
00262d04d700/Solaris10/Solaris10.vmx
scsi0:0.fileName = "Solaris10.vmdk"
scsi0:1.fileName = "Solaris10_1.vmdk"
~ #

```

13. Determine RDM device, virtual machine mapping, and compatibility mode.

14. Using the preceding information, note the RDM mapping to device, virtual machine, compatibility mode, and order.


You will need this information later, when adding RDM devices to the VM.

```

Virtual Machine -> Hardware -> NAA -> Compatibility mode
Windows2003 VM -> scsi0:1.fileName = "Windows2003_1.vmdk" ->
naa.60060e801046b96004f2bf4600000016
-> RDM Physical
Windows2003 VM -> scsi0:2.fileName = "Windows2003_2.vmdk" ->
naa.60060e801046b96004f2bf4600000017
-> RDM Virtual
Linux VM -> scsi0:1.fileName = "Linux_1.vmdk" ->
naa.60060e801046b96004f2bf4600000019 -> RDM Virtual
Solaris10 VM -> scsi0:1.fileName = "Solaris10_1.vmdk" ->
naa.60060e801046b96004f2bf4600000018 -> RDM Physical

```

15. Determine multipath configuration.
16. Obtain multipath settings for your storage in the vSphere Client:
 - a. Select an ESX or ESXi host in the vSphere Client and click the Configuration tab.
 - b. Click **Storage**.
 - c. Select a datastore or mapped LUN.
 - d. Click **Properties**.
 - e. In the Properties dialog box, select the desired extent, if necessary.
 - f. Click **Extent Device > Manage Paths** and obtain the paths in the Manage Path dialog box.



17. Obtain LUN multipathing information from the ESXi host command line:
 - a. Log in to the ESXi host console.
 - b. Run `esxcli storage nmp device list` to get multipath information.

```

# esxcli storage nmp device list
naa.60060e801046b96004f2bf4600000014

```

```
Device Display Name: HITACHI Fibre Channel Disk
(naa.60060e801046b96004f2bf4600000014)
Storage Array Type: VMW_SATP_DEFAULT_AA
Storage Array Type Device Config: SATP VMW_SATP_DEFAULT_AA does
not support device configuration.
Path Selection Policy: VMW_PSP_RR
Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useANO=0; lastPathIndex=3:
NumIOsPending=0,numBytesPending=0}
Path Selection Policy Device Custom Config:
Working Paths: vmhba2:C0:T1:L0, vmhba2:C0:T0:L0, vmhba1:C0:T1:L0,
vmhba1:C0:T0:L0
Is Local SAS Device: false
Is Boot USB Device: false
```

```
naa.60060e801046b96004f2bf4600000015
Device Display Name: HITACHI Fibre Channel Disk
(naa.60060e801046b96004f2bf4600000015)
Storage Array Type: VMW_SATP_DEFAULT_AA
Storage Array Type Device Config: SATP VMW_SATP_DEFAULT_AA does
not support device configuration.
Path Selection Policy: VMW_PSP_RR
Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useANO=0; lastPathIndex=0:
NumIOsPending=0,numBytesPending=0}
Path Selection Policy Device Custom Config:
Working Paths: vmhba2:C0:T1:L1, vmhba2:C0:T0:L1, vmhba1:C0:T1:L1,
vmhba1:C0:T0:L1
Is Local SAS Device: false
Is Boot USB Device: false
```

```
naa.60060e801046b96004f2bf4600000016
Device Display Name: HITACHI Fibre Channel Disk
(naa.60060e801046b96004f2bf4600000016)
Storage Array Type: VMW_SATP_DEFAULT_AA
Storage Array Type Device Config: SATP VMW_SATP_DEFAULT_AA does
not support device configuration.
Path Selection Policy: VMW_PSP_RR
Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useANO=0; lastPathIndex=1:
NumIOsPending=0,numBytesPending=0}
Path Selection Policy Device Custom Config:
Working Paths: vmhba2:C0:T1:L2, vmhba2:C0:T0:L2, vmhba1:C0:T1:L2,
vmhba1:C0:T0:L2
Is Local SAS Device: false
Is Boot USB Device: false
```

```
naa.60060e801046b96004f2bf4600000017
  Device Display Name: HITACHI Fibre Channel Disk
(naa.60060e801046b96004f2bf4600000017)
  Storage Array Type: VMW_SATP_DEFAULT_AA
  Storage Array Type Device Config: SATP VMW_SATP_DEFAULT_AA does
not support device configuration.
  Path Selection Policy: VMW_PSP_RR
  Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useANO=0; lastPathIndex=1:
NumIOsPending=0,numBytesPending=0}
  Path Selection Policy Device Custom Config:
  Working Paths: vmhba2:C0:T1:L3, vmhba2:C0:T0:L3, vmhba1:C0:T1:L3,
vmhba1:C0:T0:L3
  Is Local SAS Device: false
  Is Boot USB Device: false
```

```
naa.60060e801046b96004f2bf4600000018
  Device Display Name: HITACHI Fibre Channel Disk
(naa.60060e801046b96004f2bf4600000018)
  Storage Array Type: VMW_SATP_DEFAULT_AA
  Storage Array Type Device Config: SATP VMW_SATP_DEFAULT_AA does
not support device configuration.
  Path Selection Policy: VMW_PSP_RR
  Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useANO=0; lastPathIndex=1:
NumIOsPending=0,numBytesPending=0}
  Path Selection Policy Device Custom Config:
  Working Paths: vmhba2:C0:T1:L4, vmhba2:C0:T0:L4, vmhba1:C0:T1:L4,
vmhba1:C0:T0:L4
  Is Local SAS Device: false
  Is Boot USB Device: false
```

```
naa.60060e801046b96004f2bf4600000019
  Device Display Name: HITACHI Fibre Channel Disk
(naa.60060e801046b96004f2bf4600000019)
  Storage Array Type: VMW_SATP_DEFAULT_AA
  Storage Array Type Device Config: SATP VMW_SATP_DEFAULT_AA does
not support device configuration.
  Path Selection Policy: VMW_PSP_RR
  Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useANO=0; lastPathIndex=1:
NumIOsPending=0,numBytesPending=0}
  Path Selection Policy Device Custom Config:
  Working Paths: vmhba2:C0:T1:L5, vmhba2:C0:T0:L5, vmhba1:C0:T1:L5,
vmhba1:C0:T0:L5
```

```
Is Local SAS Device: false
Is Boot USB Device: false
```

Prepare hosts for FLI offline migration

The FLI offline execution phase includes the preparation of migration hosts.

In many instances it may be possible to have performed this remediation prior to this step. If not, then this is where you would perform any host remediation such as installing host attach kits or DSMs. From the analysis phase, you will have a gap list of items that need to be performed on each host in order for that host to be in a supported configuration using ONTAP. Depending on the type of migration being performed, either the host would be remediated and then rebooted (online FLI/7-Mode to ONTAP FLI), or it would be remediated and then shut down pending the completion of the migration process (offline FLI).

Presenting source LUNs to ONTAP storage during FLI

As part of the offline FLI process, you must present your source LUNs to your ONTAP storage.

Steps

1. Log in to the source array.
2. Add the NetApp initiators to the host group created during the plan phase.
3. Select the host LUNs that need to be migrated from available logical LUNs. Use LUN names for each host mentioned in the source LUNs section of your Site Survey and Planning worksheet.

Verifying source LUNs on destination storage for offline FLI

As part of the offline Foreign LUN Import process, you must verify your source LUNs on your destination storage.

Steps

1. Verify the source LUNs and mapping from source storage to destination storage.
2. Log in to the ONTAP storage through SSH using admin user.
3. Change the mode to Advanced: `set -privilege advanced`
4. Enter `y` when asked if you want to continue.
5. Discover the source array on ONTAP. Wait for a few minutes and retry to detect the source array. `storage array show`

```
DataMig-cmode::*> storage array show
Prefix                               Name      Vendor      Model Options
-----
HIT-1                                HITACHI_DF600F_1  HITACHI     DF600F
```



When the storage array is discovered for the first time, ONTAP might not show the array by discovering automatically. Use the following instructions to reset the switch port where ONTAP initiator ports are connected.

6. Verify the source array is discovered through all the initiator ports.

```
DataMig-cmode::*> storage array config show -array-name HITACHI_DF600F_1
```

Node	LUN Group	LUN Count	Array Name	Array Target Port
DataMig-cmode-01	0	1	HITACHI_DF600F_1	50060e801046b960
0a				50060e801046b964
0b				50060e801046b968
0a				50060e801046b96c
0b				
DataMig-cmode-02	0	1	HITACHI_DF600F_1	50060e801046b960
0a				50060e801046b964
0b				50060e801046b968
0a				50060e801046b96c
0b				

7. List the source LUNs mapped from Hitachi storage. Verify the disk properties and paths.

You should see the number of paths expected based on your cabling (at least two paths for each source controller). You should also check the event log after masking the array LUNs.

```

DataMig-cmode:*> storage disk show -array-name HITACHI_DF600F_1 -fields
disk, serial-number, container-type, owner, path-lun-in-use-count,
import-in-progress, is-foreign
disk      owner is-foreign container-type import-in-progress path-lun-in-
use-count serial-number
-----
-----
HIT-1.2 -   false      unassigned      false          0,0,0,0,0,0,0,0
83017542001E
HIT-1.3 -   false      unassigned      false          0,0,0,0,0,0,0,0
83017542000E
HIT-1.14 -  false      unassigned      false          0,0,0,0,0,0,0,0
830175420019
3 entries were displayed.

DataMig-cmode:*>

```

Configuring migration jobs

The FLI offline workflow requires configuration of the source LUN and destination LUNs.

Steps

1. For FLI migration, the source LUN needs to be marked as foreign. Mark the source LUNs as foreign using the serial number.

```

DataMig-cmode:*> storage disk set-foreign-lun { -serial-number
83017542001E }
                -is-foreign true
DataMig-cmode:*> storage disk set-foreign-lun { -serial-number
83017542000E }
                -is-foreign true
DataMig-cmode:*> storage disk set-foreign-lun { -serial-number
83017542000F }
                -is-foreign true

```

2. Verify the source LUN is marked as foreign.


```
DataMig-cmode::*> storage disk show -array-name HITACHI_DF600F_1 -fields
disk, serial-number, container-type, owner,import-in-progress, is-
foreign
disk      owner is-foreign container-type import-in-progress serial-
number
-----
-----
HIT-1.2  -      true      foreign      false      83017542001E
HIT-1.3  -      true      foreign      false      83017542000E
HIT-1.4  -      true      foreign      false      83017542000F
3 entries were displayed.
```

3. Create destination volumes.

```
DataMig-cmode::*> vol create -vserver datamig winvol aggr1 -size 100g
[Job 5606] Job succeeded: Successful
```

4. Disable default Snapshot policy on each volume. If default Snapshot copies exist prior to FLI migration, the volume needs additional space to store changed data.

```
DataMig-cmode::> volume modify -vserver datamig -volume winvol -snapshot
-policy none

Warning: You are changing the Snapshot policy on volume winvol to none.
Any Snapshot copies on this volume from the previous policy will not be
deleted by
      this new Snapshot policy.
Do you want to continue? {y|n}: y
Volume modify successful on volume winvol of Vserver datamig.
```

5. Set fraction_reserveoption for each volume to 0 and set the Snapshot policy to none.

```
DataMig-cmode::> vol modify -vserver datamig -volume * -fractional
-reserve 0 -snapshot-policy none
Volume modify successful on volume winvol of Vserver datamig.
```

6. Check your volume settings.

```
DataMig-cmode::> vol show -vserver datamig -volume * -fields fractional-
reserve,snapshot-policy
vservervolumesnapshot-policyfractional-reserve
-----
datamig datamig_rootnone0%
datamigwinvolnone0%
Volume modify successful on volume winvol of Vserver datamig.
```

7. Delete any existing Snapshot copies.

```
DataMig-cmode::> set advanced; snap delete -vserver datamig -vol winvol
-snapshot * -force true
1 entry was acted on.
```



FLI migration modifies every block of the target LUN. If default or other Snapshot copies exist on a volume prior to FLI migration, the volume gets filled up. Changing the policy and removing any existing Snapshot copies before FLI migration are required. Snapshot policy can be set again post-migration.



The LUN create command detects the size and alignment based on partition offset and creates the LUN accordingly with foreign-disk option. For a review, see the NetApp Knowledgebase article **What is an unaligned I/O?** Also note that some I/O will always appear be partial writes and will therefore look misaligned. Examples of this would be database logs.

[What is an unaligned I/O?](#)

8. Create destination LUNs using foreign LUN.

```
DataMig-cmode::*> lun create -vserver datamig -path /vol/winvol/bootlun
-ostype windows_2008 -foreign-disk 83017542001E

Created a LUN of size 40g (42949672960)

Created a LUN of size 20g (21474836480)
DataMig-cmode::*> lun create -vserver datamig -path
/vol/linuxvol/lvmlun1 -ostype linux -foreign-disk 830175420011

Created a LUN of size 2g (2147483648)
DataMig-cmode::*> lun create -vserver datamig -path /vol/esxvol/bootlun
-ostype vmware -foreign-disk 830175420014

Created a LUN of size 20g (21474836480)
```

9. List the destination LUNs and verify the size of LUN with source LUN.

```
DataMig-cmode::*> lun show -vserver datamig
Vserver    Path                               State   Mapped   Type
Size
-----
-----
datamig    /vol/esxvol/bootlun                online  unmapped vmware
20GB
datamig    /vol/esxvol/linuxrdmvlun           online  unmapped linux
2GB
datamig    /vol/esxvol/solrdmplun              online  unmapped solaris
2GB
datamig    /vol/winvol/gdrive                  online  unmapped windows_2008
3GB
4 entries were displayed.

DataMig-cmode::*>
```



For FLI offline migration, the LUN must be mapped to the igroup and then be offlined before creating the LUN import relationship.

10. Create host igroup of protocol FCP and add initiators. Find initiator WWPNs from storage groups section of your Site Survey planning worksheet.

```
DataMig-cmode::*> lun igroup create -ostype windows -protocol fcp
-vserver datamig -igroup dm-rx200s6-21 -initiator
21:00:00:24:ff:30:14:c4,21:00:00:24:ff:30:14:c5

DataMig-cmode::*> lun igroup create -ostype linux -protocol fcp -vserver
datamig -igroup dm-rx200s6-22 -initiator
21:00:00:24:ff:30:04:85,21:00:00:24:ff:30:04:84

DataMig-cmode::*> lun igroup create -ostype vmware -protocol fcp
-vserver datamig -igroup dm-rx200s6-20 -initiator
21:00:00:24:ff:30:03:ea,21:00:00:24:ff:30:03:eb
```



Use the same LUN ID as source. Refer to source LUNS section of your Site Survey planning worksheet.

11. Map the destination LUNs to igroup.

```
DataMig-cmode:*> lun map -vserver datamig -path /vol/winvol/bootlun
-igroup dm-rx200s6-21 -lun-id 0
DataMig-cmode:*> lun map -vserver datamig -path /vol/linuxvol/bootlun
-igroup dm-rx200s6-22 -lun-id 0
DataMig-cmode:*> lun map -vserver datamig -path /vol/esxvol/bootlun
-igroup dm-rx200s6-20 -lun-id 0
```

12. Offline the destination LUNs.

```
DataMig-cmode:*> lun offline -vserver datamig -path /vol/esxvol/bootlun
DataMig-cmode:*> lun offline -vserver datamig -path
/vol/esxvol/linuxrdmvlun
DataMig-cmode:*> lun offline -vserver datamig -path
/vol/esxvol/solrdmplun
```

13. Create import relationship with destination LUN and source LUN.

```
DataMig-cmode:*> lun import create -vserver datamig -path
/vol/winvol/bootlun -foreign-disk 83017542001E
DataMig-cmode:*> lun import create -vserver datamig -path
/vol/linuxvol/ext3lun -foreign-disk 830175420013
DataMig-cmode:*> lun import create -vserver datamig -path
/vol/esxvol/linuxrdmvlun -foreign-disk 830175420018
DataMig-cmode:*> lun import create -vserver datamig -path
/vol/esxvol/solrdmplun -foreign-disk 830175420019
```

14. Verify the import job creation.

```

DataMig-cmode::*> lun import show -vserver datamig
vserver foreign-disk path operation admin operational
percent
in progress state state
complete
-----
-----
datamig 83017542000E /vol/winvol/fdrive import stopped
stopped
0
datamig 83017542000F /vol/winvol/gdrive import stopped
stopped
0
datamig 830175420010 /vol/linuxvol/bootlun
import stopped
stopped
0
3 entries were displayed.

```

Copyright Information

Copyright © 2021 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.