FLI online workflow

ONTAP FLI

NetApp
May 27, 2025

Table of Contents

FLI online workflow

FLI online workflow

This is the second of the four FLI workflow examples, which covers FLI online migration. The source array in this example is an EMC VNX5500.

the online workflow has the following tasks:

1. Preparing the LUN path

2. Performing a disruptive cutover

3. Importing the source LUN data

4. Verifying migration results

5. Cleaning up the FLI online migration

6. FLI online post-migration tasks

> (i) If the NetApp controller destination is a MetroCluster, do NOT use the online workflow. If a site failover occurred during an active online import, write pass-throughs to the source array could fail which would lead to a verification failure and potential data loss. If the destination is MetroCluster, use the FLI offline process.

Rebooting hosts

You have the option to reboot hosts prior to starting this workflow in order to verify that the host is in a known good state.

Before rebooting your host, it would also be a good time to take a Snapshot copy in order to facilitate a revert if needed later. 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.

Supported host operating systems for FLI online

The FLI online workflow can be used for LUNs connected to hosts running one of the following operating systems starting with the versions listed.

1. Microsoft (all versions of the servers listed are supported):

    ◦ Windows Server 2008 and later (includes Windows Server failover clusters)

    ◦ Microsoft Hyper-V server 2008 and later

    ◦ Windows Server 2012 and later (includes Windows Server 2012 cluster)

- Microsoft Hyper-V Server 2012 and later

2. VMware:

- All ESXi 5.x and later releases

3. Linux:

- Red Hat Enterprise Linux (RHEL) 5.x and later

Consider using the FLI offline workflow for host operating systems not on the preceding list.

> ⓘ If the NetApp controller destination is a MetroCluster, do NOT use the online workflow. If a site
> failover occurred during an active online import, write pass-throughs to the source array could
> fail which would lead to a verification failure and potential data loss. If the destination is
> MetroCluster, use the FLI offline process, regardless of the host OS.

Verify host LUN path and multipath configuration

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

Prepare hosts for FLI online migration

The FLI online execution phase includes the preparation of migration hosts in order to be in a supported configuration.

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, or simply remediated.

FLI online: Preparing the LUN path

To prepare for FLI online migration, you verify the host and source LUN paths and other details.

Steps

1. In ONTAP, change the privilege level to `advanced`.

```
cluster::> set adv

Warning: These advanced commands are potentially dangerous; use them
only when directed to do so by NetApp personnel.
Do you want to continue? {y|n}: y
```

2. Verify that the source array can be seen on the destination controller.

```
      cluster::*> storage array show
Prefix                            Name    Vendor          Model Options
-------- -------------------------- -------- ----------------
----------
DGC-1                         DGC_LUNZ_1      DGC                LUNZ
1 entries were displayed.
```

3. Display source LUN details.

```
cluster::*> storage array config show -array-name DGC_LUNZ_1 -instance

         Controller Name: ontaptme-fc-cluster-01
               LUN Group: 0
       Array Target Ports: 500601643ea067da
               Initiator: 0c
              Array Name: DGC_LUNZ_1
   Target Side Switch Port: stme-5010-3:2-1
Initiator Side Switch Port: stme-5010-3:2-3
      Number of array LUNs: 1

         Controller Name: ontaptme-fc-cluster-01
               LUN Group: 0
       Array Target Ports: 500601653ea067da
               Initiator: 0d
              Array Name: DGC_LUNZ_1
   Target Side Switch Port: stme-5010-4:2-1
Initiator Side Switch Port: stme-5010-4:2-3
      Number of array LUNs: 1
~~~~~~~~~~~ output truncated for readability ~~~~~~~~~~~~~~~~~~
8 entries were displayed.
```

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

```
cluster::*> storage array config show -array-name DGC_LUNZ_1
            LUN   LUN
Node          Group Count                           Array Name        Array Target
Port Initiator
------------ ----- ----- ----------------------------
---------------------- ---------
ontaptme-fc-cluster-01
                0     1                             DGC_LUNZ_1
500601643ea067da        0c

500601653ea067da        0d

5006016c3ea067da        0c

5006016d3ea067da        0d
ontaptme-fc-cluster-02
                0     1                             DGC_LUNZ_1
500601643ea067da        0c

500601653ea067da        0d

5006016c3ea067da        0c

5006016d3ea067da        0d
8 entries were displayed.
```

ⓘ The word wrapping in the following output has no significance.

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

```
cluster::*> storage disk show -array-name DGC_LUNZ_1 -instance
                Disk: DGC-1.9
      Container Type: unassigned
          Owner/Home: -  / -
            DR Home: -
   Stack ID/Shelf/Bay: -  / -  / -
                 LUN: 0
               Array: DGC_LUNZ_1
              Vendor: DGC
               Model: VRAID
       Serial Number: 600601603F103100662E70861000E511
                 UID:
60060160:3F103100:662E7086:1000E511:00000000:00000000:00000000:00000000:
00000000:00000000
                 BPS: 512
       Physical Size: -
            Position: present
Checksum Compatibility: block
           Aggregate: -
                Plex: -
Paths:
                          LUN  Initiator Side       Target Side
Link
Controller        Initiator    ID  Switch Port        Switch Port
Acc Use  Target Port              TPGN    Speed     I/O KB/s
IOPS
------------------ ---------  -----  --------------------
------------------ --- ---  ----------------------  ------  -------
----------- ------------
ontaptme-fc-cluster-02
                 0c            0  stme-5010-3:2-4      stme-5010-
3:2-2      AO  INU  5006016c3ea067da           2   4 Gb/S
0          0
ontaptme-fc-cluster-02
                 0d            0  stme-5010-4:2-4      stme-5010-
4:2-2      AO  INU  5006016d3ea067da           2   4 Gb/S
0          0
ontaptme-fc-cluster-02
                 0d            0  stme-5010-4:2-4      stme-5010-
4:2-1      ANO RDY  500601653ea067da           1   4 Gb/S
0          0

Errors:
-
```

6. View the source LUN.

```
cluster::*> storage disk show -array-name DGC_LUNZ_1
                       Usable             Disk    Container   Container
Disk                   Size Shelf Bay Type    Type         Name
Owner
--------------- ---------- ----- --- ------- ----------- ---------
--------
DGC-1.9                     -      -    - LUN      unassigned  -         -
```

7. Mark the source LUN as foreign.

```
cluster::*> storage disk set-foreign-lun -is-foreign true -disk DGC-1.9
```

8. Verify the source LUN is marked as foreign.

```
cluster::*> storage disk show -array-name DGC_LUNZ_1
                       Usable             Disk    Container   Container
Disk                   Size Shelf Bay Type    Type         Name
Owner
--------------- ---------- ----- --- ------- ----------- ---------
--------
DGC-1.9
```

9. Serial numbers are used in FLI LUN import commands. List all foreign LUNs and their serial numbers.

```
cluster::*> storage disk show -container-type foreign -fields serial-
number
disk    serial-number
------- -------------------------------
DGC-1.9 600601603F103100662E70861000E511
```

10. Create a destination volume.

```
cluster::*> vol create -vserver fli -volume fli_vol -aggregate aggr1
-size 2t
[Job 13888] Job succeeded: Successful
```

11. Verify volume.

```
cluster::*> vol show -vserver fli
Vserver    Volume        Aggregate    State      Type      Size
Available Used%
---------  ------------  ------------  ---------  ----  ----------
----------  -----
fli        fli_root      aggr1         online     RW         1GB
972.6MB    5%
fli        fli_vol       aggr1         online     RW         2TB
1.90TB     5%
2 entries were displayed.
```

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

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

13. Check your volume settings.

```
cluster::*> vol show -vserver datamig -volume * -fields fractional-
reserve,snapshot-policy
vservervolumesnapshot-policyfractional-reserve
--------------------------------------------------
datamigdatamig_rootnone0%
datamigwinvolnone0%
Volume modify successful on volume winvol of Vserver datamig.
```

14. Delete any existing Snapshot copies.

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

(i) FLI migration modifies every block of the target LUNs. 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 is required. Snapshot policy can be set again post-migration.

(i) 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 of I/O misalignment, review NetApp Knowledgebase article **What is an unaligned I/O**? Also note that some I/O will always appear to be partial writes and will therefore look misaligned. Examples of this would be database logs.

15. Create the target LUN. The `LUN create` command detects the size and alignment based on partition offset and creates the LUN accordingly with the foreign-disk argument.

```
cluster::*> lun create -vserver fli -path /vol/fli_vol/OnlineFLI_LUN
-ostype windows_2008 -foreign-disk 600601603F103100662E70861000E511

Created a LUN of size 1t (1099511627776)
```

16. Verify new LUN.

```
cluster::*> lun show -vserver fli
Vserver    Path                                 State    Mapped    Type
Size
---------  -----------------------------  -------  --------  --------
--------
fli        /vol/fli_vol/OnlineFLI_LUN        online   unmapped windows_2008
1TB
```

17. Create an igroup of protocol FCP with host initiators.

```
cluster::*> igroup create -vserver fli -igroup FLI -protocol fcp -ostype
windows -initiator 10:00:00:00:c9:e6:e2:79
```

18. Verify that the host logs in on all paths to the new igroup.

```
cluster::*> igroup show -vserver fli -igroup FLI
    Vserver name: fli
     Igroup name: FLI
        Protocol: fcp
      OS Type: Windows
Portset Binding Igroup: -
    Igroup UUID: 5c664f48-0017-11e5-877f-00a0981cc318
          ALUA: true
     Initiators: 10:00:00:00:c9:e6:e2:77 (logged in)
10:00:00:00:c9:e6:e2:79 (logged in)
```

19. Offline the destination LUN.

```
cluster::*> lun offline -vserver fli -path /vol/fli_vol/OnlineFLI_LUN

Warning: This command will take LUN "/vol/fli_vol/OnlineFLI_LUN" in
Vserver "fli" offline.
Do you want to continue? {y|n}: y
```

20. Map the destination LUN to the igroup.

```
cluster::*> lun map -vserver fli -path /vol/fli_vol/OnlineFLI_LUN
-igroup FLI
```

21. Create import relationship between new LUN and foreign LUN.

```
cluster::*> lun import create -vserver fli -path
/vol/fli_vol/OnlineFLI_LUN -foreign-disk
600601603F103100662E70861000E511
```

FLI online: Performing a disruptive cutover

This example gives you the general steps for performing a disruptive cutover for the FLI online migration process.

For host remediation walkthrough of Windows, Linux, and ESXi, see the follow-on sections of this guide as well as the host operating system and host-attach kit documentation.

Steps

1. On the foreign array, display the storage group to which the source LUN is mapped.

   See the vendor documentation for the appropriate commands.

2. If the LUNs that are being imported are for an ESXi host, review and follow the instructions for the *ESXi CAW/ATS remediation* topic.

3. Unmap the source LUN from the hosts.

   (i)   The disruption window begins here.

   The disruption begins immediately after the `unmap` command is executed. Generally, the disruption window can be measured in minutes. The disruption window is the time taken to repoint the host at the new NetApp target and to scan for LUNs.

   You must ensure that this is the only LUN that is mapped to this igroup, because removing the host (initiator) from the igroup affects other LUNs that are mapped to the igroup. See the vendor documentation for the appropriate commands.

4. Verify that the host initiators are no longer present.

5. On the ONTAP cluster, bring the destination LUN online, and verify that it is mapped.

```
cluster::*> lun online -vserver fli -path /vol/fli_vol/OnlineFLI_LUN
```

6. Verify that the LUN is online.

```
cluster::*> lun show -vserver fli
Vserver    Path                                    State    Mapped    Type
Size
---------  ------------------------------  -------  --------  --------
--------
fli        /vol/fli_vol/OnlineFLI_LUN      online   mapped    windows_2008
1TB
```

7. Rescan the disks on the host, find the LUN on the ONTAP target, and then verify that the DSM has claimed the LUN.

> **(i)** The disruption window ends here.

8. Verify that you can see all expected paths, and check your event logs to verify that no errors exist.

At this point, the disruptive part of this migration is complete, unless there are outstanding host remediation tasks (identified during your analysis and planning phases) that are disruptive.

The LUNs are online and mapped, and the hosts are now mounting the new ONTAP-hosted LUN. Reads are passed through the ONTAP array to the source LUN, and writes are written to both the new ONTAP-hosted LUN and the original source LUN. The source LUN and destination LUN will stay synchronized until the migration is complete and the LUN relationship has been broken.

FLI online: Importing the data

These are the steps to import the data from the source LUN to the destination LUN.

Steps

1. Start the migration import.

```
cluster::*> lun import start -vserver fli -path
/vol/fli_vol/OnlineFLI_LUN
```

2. Display FLI status.

```
cluster::*> lun import start -vserver fli -path
/vol/fli_vol/OnlineFLI_LUN
```

FLI online: Verifying migration results

A verify job is optional, but recommended. It is a block-by-block comparison of the source and destination LUNs. Verify jobs take almost the same or slightly more time than migration time.

Start the verify job to compare source and destination LUNs. Monitor the verify progress. The LUNs being verified need to be offline for the duration of the verification session. The verification session can potentially be lengthy because it is a block-by-block comparison between source and destination LUNs. Although verification is not required, it is a good idea to verify a subset of the LUNs imported/migrated to feel comfortable about the import process. These verifications would be in addition to those performed during the test/pilot migrations.

> (i)   This process is disruptive.

> (i)   The LUN import verify must be explicitly stopped before bringing the LUN back online. Otherwise, the LUN online fails. See the following CLI output.

Steps

1. Offline the LUNs to be verified. <The disruption window begins here>

   ```
   cluster::*> lun offline -vserver fli_72C -path /vol/flivol/72Clun1
   Warning: This command will take LUN "/vol/flivol/72Clun1" in Vserver
   "fli_72C" offline.
   Do you want to continue? {y|n}: y
   ```

2. Start LUN verify.

   ```
   lun import verify start -vserver fli_72C -path /vol/flivol/72Clun1
   ```

3. Display LUN verify status.

   ```
   ontaptme-fc-cluster::*> lun import show -vserver fli_72C -path
   /vol/flivol/72Clun1
   vserver foreign-disk   path                     operation admin operational
   percent
                                                   in progress state state
   complete
   ----------------------------------------------------------------------
   -------
   fli_72C D0i1E+G8Wg6m   /vol/flivol/72Clun1 verify     started

   9
   ```

4. Stop LUN verification. This step needs to be performed manually even if the status shows that verify is complete.

```
lun import verify stop -vserver fli_72C -path /vol/flivol/72Clun1
```

5. Online the LUN upon completion of verification. <The disruption window ends here>

```
lun online -vserver fli_72C -path /vol/flivol/72Clun1
```

Cleaning up the FLI online migration

After the FLI online migration completes, you remove the LUN import relationship.

When ready, the LUN import relationship can be safely removed because the host is now accessing the new NetApp array for all I/O to the new ONTAP LUN, and the source LUN is no longer in use.

Step

1. Delete the LUN import relationship.

```
lun import delete -vserver fli_72C -path /vol/flivol/72Clun1
```

FLI online post-migration tasks

Any server remediation not performed pre-migration, is performed during post-migration.

Any third-party software is removed. NetApp software is installed and configured. See host remediation for examples of post-migration remediation for specific host types.

Review logs for errors, check pathing, and perform any application testing to verify that your migration completed cleanly and successfully.