



ONTAP FLI phases

ONTAP FLI

NetApp
July 22, 2026

Table of Contents

- ONTAP FLI phases. 1
 - Learn about ONTAP FLI phases 1
 - ONTAP Foreign LUN Import (FLI) Discovery phase 1
 - Prepare for an iSCSI migration with ONTAP FLI 1
 - ONTAP FLI discovery phase tasks 2
 - Recommended ONTAP FLI discovery deliverables 3
 - ONTAP Foreign LUN Import (FLI) Analyze phase 3
 - ONTAP FLI analyze phase tasks 4
 - Use the IMT during the ONTAP FLI analyze phase 4
 - ONTAP FLI host changes and rollback strategy 5
 - ONTAP FLI gap analysis 5
 - Plan and prepare 6
 - ONTAP Foreign LUN Import (FLI) Plan and Prepare phase 6
 - ONTAP Foreign LUN Import (FLI) test migration and cutover planning 8
 - Host remediation 15
 - ONTAP Foreign LUN Import (FLI) Execute phase 33
 - ONTAP FLI execute phase tasks 33
 - ONTAP Foreign LUN Import (FLI) Verify phase 35
 - ONTAP FLI verify phase tasks 35

ONTAP FLI phases

Learn about ONTAP FLI phases

This section describes the end-to-end methodology for executing Foreign LUN Import (FLI) migrations—with specific emphasis on iSCSI backend FLI—using a structured, phased approach. It outlines the objectives, inputs, validation checkpoints, and deliverables across the Discovery, Analyze, Plan & Prepare, Execute, and Verify phases to ensure a predictable, supportable, and low-risk migration from a foreign storage array to ONTAP.

The workflow emphasizes upfront data collection and interoperability verification, host and fabric readiness, explicit iSCSI backend design constraints (IC LIF usage, explicit target portals/IQNs, full-mesh connectivity, and scale limits), and disciplined change management to protect application availability and rollback options. By combining IMT-driven support validation, targeted remediation planning, test migrations, and rigorous post-migration cleanup, the approach ensures that migrations complete successfully while resulting in a fully supported, well-documented, and operationally clean post-migration state for both FC- and iSCSI-based FLI scenarios.

ONTAP Foreign LUN Import (FLI) Discovery phase

The ONTAP FLI discovery phase is the first step in the migration process. It is critical to gather all necessary information about the source environment, connectivity, and target ONTAP system to ensure a smooth migration. The discovery phase sets the foundation for planning, preparing, executing, and verifying the migration.

Prepare for an iSCSI migration with ONTAP FLI

The ONTAP FLI discovery phase gathers the technical details that will be required to complete the migration:

- Plan the migration workflow (online vs offline)
- Execute host remediation safely
- Validate connectivity and discovery of the foreign LUN on ONTAP before creating the import relationship

For structured, repeatable data collection across hosts, NetApp controllers, and fabrics, you can use NetApp Active IQ OneCollect. OneCollect is designed to collect information from storage, hosts, fabrics, and switches for migration planning and validation.

For FC FLI, the discovery phase must capture the foreign array target port WWPNs.

For iSCSI FLI, the discovery phase must:

- Capture the foreign array target portal IPs and target IQNs
- Verify IC LIF reachability
- Plan full mesh iSCSI sessions (each HA node connecting to each target portal) to ensure the foreign LUN is discovered through all available paths

ONTAP FLI discovery phase tasks

Each component has a set of tasks to collect data or validate connectivity.

Component	Tasks
Host	• Identify the host OS version, multipathing stack and configuration, and any special consumption model, such as filesystem, raw device, or cluster disk, to drive host remediation and cutover steps. • Capture LUN inventory, including the size, identifiers or serial numbers, partition layouts or offsets, and how each LUN is consumed by applications for planning and rollback. • For FC frontend environments, record host initiator WWPNs. These are used for zoning and masking, if applicable to the overall migration and cutover design. • For iSCSI frontend environments, if used post-migration, record host initiator IQNs and target portal IPs that the host will use for login after cutover.
Fabric and networking	• For FC environments: ◦ Document current zoning (to and from the host and source array zones) so it can be safely removed or updated after cutover. ◦ Capture the switch identifiers and ports involved. • For iSCSI environments: ◦ Document VLAN, subnet, and routing path between ONTAP IC LIFs and foreign array target portals ◦ Confirm there are no firewalls or ACLs blocking communication.
NetApp storage array (destination)	• Confirm the ONTAP version and ECV readiness for iSCSI FLI.  Initiator creation is blocked if the cluster is not at the supported level. • Identify which nodes will participate, such as the LUN-owning node and its HA partner, and record the intercluster LIFs available on those nodes. Plan for redundancy; plan for IC LIFs, initiators, and sessions on both HA nodes. • Record IC LIF IPspace and subnet and verify they can reach foreign target portals. Avoid creating new IC LIFs in a different subnet if existing IC LIFs are already used elsewhere in the cluster. • Validate destination-side constraints you must enforce during planning, such as the import max size, 512B block-size requirement, and session and import scaling limits relevant to iSCSI FLI.

Component	Tasks
Foreign Storage Array (Source)	- Identify the foreign array target portal IPs and the target IQNs; iSCSI discovery sessions are not supported for iSCSI FLI, so these must be explicitly known up front. - Identify the specific LUNs mapped to the source host, including LUN size and block size. Ensure the LUN is stable, not expanding, and suitable for import planning. - Record host access and masking objects such as host and initiator records, so you can safely unmap or remap during cutover and validate rollback readiness. - Plan and run foreign-array qualification.  Field qualification is required for iSCSI FLI; do not assume FC-qualified arrays are automatically qualified for iSCSI.

Recommended ONTAP FLI discovery deliverables

- A single inventory sheet mapping: Host → Source LUN(s) → Foreign target portal/IQN → ONTAP IC LIFs/nodes → Destination LUN(s) for clean execution and rollback.
- A OneCollect package (or equivalent) captured before change execution to reduce triage time if issues occur.

Related information

Tools

ONTAP Foreign LUN Import (FLI) Analyze phase

The analyze phase determines what must be fixed before you move into detailed migration planning and execution. The main outcome of this phase is to ensure every host, and the full SAN stack it depends on (OS, volume manager, HBAs, multipathing, switch firmware, and the target ONTAP release), can be brought into a supported post migration configuration. This is done by comparing each host's current state against supported combinations in the NetApp Interoperability Matrix (IMT), then defining a clear target configuration per host (OS level, MPIO/ALUA behavior, HBA driver/firmware, and any required host tooling) and producing a gap analysis that lists the exact changes required to reach that target.

Schedule and review the gap analysis to avoid remediation changes, which can impact application compatibility and outage planning. In many environments, some changes, such as OS patching or HBA updates, are safer to complete ahead of the migration window. However, multipathing stack changes are often delayed until cutover or post migration because altering MPIO behavior can affect supportability with the current array and complicate rollback.

For iSCSI FLI, the analyze phase should also capture iSCSI specific readiness:

- Confirm the host's iSCSI multipath design remains supported post cutover
- Ensure the foreign array is qualified for iSCSI backend use, since backend transport differences can change interoperability behavior.

ONTAP FLI analyze phase tasks

The analyze phase is host-centric, but it should also validate the full end-to-end support chain (host to fabric to ONTAP and back again) and confirm the FLI backend interoperability position.

Component	Tasks
Host	• Compare the current host configuration to the IMT-supported configurations and determine the target post migration configuration: ◦ OS level ◦ Multipathing approach ◦ HBA model, driver, and firmware ◦ Required host tooling • Create a per-host gap analysis with the required: ◦ Hotfixes and patches ◦ OS updates ◦ HBA driver and firmware upgrades ◦ Multipath changes ◦ List any required utilities
Fabric (FC/iSCSI network)	• Validate the switch and fabric support (model and firmware) in the end-to-end view and record required upgrades if the fabric is not supported for the target ONTAP and host combination.
NetApp storage array (destination)	• Confirm the target ONTAP platform and release is compatible with the intended FLI workflow, including any applicable AFF enablement expectations, and record any prerequisites that influence the target design.
Foreign storage array (source)	• Validate whether the source array is covered by FLI backend interoperability. If interoperability is not listed in the IMT for ONTAP 9.9.1 and later releases, use the SAN LUN Migrate App to help qualify whether it might be supported.

Use the IMT during the ONTAP FLI analyze phase

The Interoperability Matrix Tool (IMT) is a repository of configurations and a guided way to narrow to a supported end to end result configurations where NetApp products interoperate with third party components. IMT can contain both supported and certified configurations:

- Supported configurations are qualified by NetApp.
- Certified configurations are qualified by a third party to work with NetApp components.

Check the IMT for:

- Capture IMT actions in your worksheet: Document IMT-recommended software, patches, and upgrades for both hosts and switches in your planning worksheet.
- Start broad, then refine: Begin with static criteria (ONTAP release, protocol, HA/cluster mode), then add the host OS, volume manager, and HBA details using your site survey as a guide.
- Avoid over-filtering: If your search becomes too specific and returns no results, relax filters, review multiple valid configurations, and choose the best supported match.
- Normalize HBA identifiers: HBAs may show up as OEM part numbers, cross-reference them before entering them into the IMT.
- Check every host: Run the IMT supportability check for each host included in the migration scope.
- Do two interoperability checks:
 - backend FLI support: Confirm the source array backend scenario is supported.
 - End-to-end post-migration support: Confirm support for the ONTAP, host OS, multipath, HBAs, and the switch fabric combination.
- Practical IMT workflow:
 - Search the source array model, then select “Foreign LUN Import (FLI) backend Interoperability” and choose the ONTAP platform and ONTAP release.
 - For host compatibility, use “Build end-to-end view with ONTAP SAN host” to validate the OS, multipath, and the HBA.
 - For switch compatibility, use “Build end-to-end view for SAN-Switch” from the SAN host view to validate fabric and switch support.

If the foreign array is not listed in IMT (ONTAP 9.9.1 and later)

Beginning with ONTAP 9.9.1, if the foreign array is not listed as supported in IMT, you can use the SAN LUN Migrate Tool to help qualify whether the array might be supported for FLI. The app is downloaded from the NetApp Support Site and run it from a Linux host with block access (FC or iSCSI) to the source array LUN. The results are reviewed with engineering. The Import from the Foreign Array can be field qualified. For iSCSI FLI, the Import from the Foreign Array is always field qualified.

ONTAP FLI host changes and rollback strategy

The analyze phase should explicitly decide what will change and when, because host changes can introduce risk and can affect support of the current storage as well. This is important if the host stack requires upgrade before it can be supported with ONTAP.

- Host updates (patches, HBA firmware and driver) can be safer when done ahead of time where possible but still require coordination with maintenance windows and application teams.
- Changes to the multipathing stack (MPIO, ALUA, and third-party multipath) can affect the current array supportability, so they are often deferred until the migration event or even until after cutover.
- Delaying MPIO reconfiguration can make rollback simpler if you must revert to the original storage during the outage window.

ONTAP FLI gap analysis

The gap analysis summarizes the current environment versus the NetApp-recommended target configuration and lists upgrades or remediation required to achieve a supported post migration state.

A complete target configuration per host typically includes:

- OS configuration level and required patches or hotfixes
- Multipathing approach and required settings (MPIO and ALUA behavior)
- HBA and CNA model, driver and firmware version
- Host utilities and tooling requirements relevant to your environment, if applicable

The gap analysis should be reviewed promptly because remediation actions might require downtime windows and might interact with application support requirements.

Plan and prepare

ONTAP Foreign LUN Import (FLI) Plan and Prepare phase

The plan and prepare phase turns the discovery and analysis outputs into an executable migration plan and confirms that the environment is ready for the migration window. Most of the project effort typically sits here because this is where you finalize the cutover approach, online or offline, validate end to end connectivity, and prove the workflow through one or more test migrations before moving to production.

For iSCSI FLI, this phase must also validate the iSCSI backend design constraints. The ONTAP iSCSI software initiator uses IC LIFs, requires explicit target portal IP and target IQN, no discovery sessions, and is designed to run within defined scale limits (imports per node, sessions per LIF/node, initiator per LIF, maximum import size).

ONTAP FLI plan and prepare phase results

- Finalize the remediation plan using the host gap analysis, including what changes are done pre migration, including what changes are done during the cutover window, such as MPIO changes.
- Build the source-to-destination LUN mapping, source LUN identifiers and serials to the destination ONTAP LUN names and paths, including rollback mapping.
- Validate backend connectivity from ONTAP to the foreign array using IC LIF reachability and a full mesh session topology where each HA node establishes sessions to each target portal. This is the prerequisite for reliable foreign LUN discovery.
- Perform one or more test migrations to confirm configuration correctness and to gather throughput samples that help estimate production durations and staffing effort.

ONTAP FLI plan and prepare phase tasks

The following are the ONTAP FLI plan and prepare tasks for each component.

Component	Tasks
Host	• Finalize the host remediation plan using the gap analysis, including whether the change should happen pre migration, during cutover, or post migration. • Choose the migration mode, online or offline, based on downtime tolerance and workload criticality; align remediation timing accordingly. • Validate and document the host-side rediscovery and remediation steps for the intended use case, standard LUN, clustered disk, RDM/VMFS datastore, etc., to be used during the migration window.
Fabric / Network	• FC environments: Back up existing zone sets and plan zone changes required for cutover and post migration cleanup. • iSCSI backend environments: Validate end to end IP reachability between ONTAP IC LIFs and foreign array target portal IPs (routing/VLAN/firewall incl. TCP/3260) and ensure redundancy. • Do not add new IC LIFs on a different subnet solely to reach foreign targets when existing IC LIFs are already in use on another subnet (treat as unsupported planning).
NetApp storage array (Destination ONTAP)	• Confirm ONTAP cluster readiness for iSCSI FLI (supported release/ECV) and that migration-only iSCSI backend configuration will be removed after migration (important for upgrade/revert readiness). • Identify and confirm IC LIFs on the LUN owning node and HA partner and plan resiliency (commonly two IC LIFs per node on different ports). • Create software iSCSI initiators (one initiator per IC LIF) and establish full mesh sessions (each HA node for each foreign target portal and IQN) and verify sessions are healthy. • Discover foreign LUNs in ONTAP using the storage disk show command, validate path health and basic properties including the 512B requirement, and then mark the foreign LUNs as foreign. • Provision a test destination LUN with size parity using the -foreign -disk parameter, map it using the planned front end protocol (FC or iSCSI), and execute at least one test migration to validate the end to end workflow and capture throughput. • Ensure the plan respects iSCSI FLI constraints (no discovery sessions, no MCS, ERL0 only, import and session limits, and maximum import size).

Component	Tasks
Foreign storage array (Source)	• Collect target portal IPs and target IQNs, the sessions are created explicitly and discovery sessions are not used. • Create a vendor-side access object (host/host-group equivalent) for ONTAP initiators and map a small test LUN (commonly ~1 GB) only to ONTAP initiators for qualification and testing. • Verify the foreign target sees ONTAP initiator logins and sessions and that the correct LUNs are mapped for import. • Ensure the production host is not simultaneously presented the same source LUN during cutover to avoid dual presentation. • Complete foreign array qualification using supported qualification tools and approach prior to production imports.

ONTAP Foreign LUN Import (FLI) test migration and cutover planning

Run one or more test migrations to confirm that arrays, the network and fabric, ONTAP configuration, and host remediation steps are correct, and to collect throughput samples that can be extrapolated for production scheduling.

A good test migration validates:

- Backend readiness: IC LIF reachability, iSCSI sessions stability, and foreign LUN discovery is consistent across paths.
- Workflow correctness: Import relationship creation/start/show/complete (and optional verify) behaves as expected under the selected mode (online/offline).
- Host readiness: Host rescan/remediation steps work reliably for the target use case (e.g., ESXi VMFS/RDM, Windows standalone/WFC).
- Operational estimates: Throughput observed in test runs provides the most realistic basis for estimating production migration duration and is generally better than any generic benchmark.

Test migration example using Hitachi AMS2100

The following is an example test migration using a Hitachi AMS2100 as the foreign array. Depending on the arrays involved, host operating systems, and other variables, your steps may be different.

You can use the following example as a general guide to the steps required to perform test migrations. NetApp recommends performing test migrations as early as possible in order to find and have as much time as possible to resolve any issues brought to light by the tests. A test migration should be performed on all combinations of source and destination arrays before proceeding with the production migrations.

Steps

1. Create a 2 GB test LUN on the source array.
 - a. Log in to Hitachi Storage Navigator Modular as `system`.
 - b. Select **AMS 2100 array**.
 - c. Click **Show and Configure Array**.
 - d. Log in using `root`.

- e. Expand **Groups** and select **Logical Units**.
 - f. Select **Create LU** to create the test LUN.
 - g. Create a test LUN of 2 GB.
 - h. Click **OK**.
 - i. Skip the LUN assignment here and proceed by clicking **Close**.
 - j. Verify the test LUN was created.
2. Assign the test LUN to the host group created for ONTAP initiators.
 - a. Expand **Groups** and select **Logical Units**.
 - b. Select **Host Groups** to map the test LUN to the **cDOT_FLI** host group.
 - c. Select host group **cDOT_FLI** created in the previous step and click **Edit Host Group**.
 - d. Choose the ports for the host group. In this example we choose **0a, 0e, 1a, 1e**. Select the **Forced Set to All Selected Ports** option.
 - e. In the HDS Storage Navigator, edit the host group.
 - f. Click **Logical Units** and add the test LUN.
 - g. Click **OK** to map the LUN.
 - h. Select **Yes, I Have Read the Above Warning and Want to Edit Host Group** and click **Confirm**.
 - i. Verify host group creation and click **Close**.
 3. Verify the test LUN and mapping from the source storage to destination storage and perform Foreign LUN Import (FLI) import.
 - a. Log in to the ONTAP storage through SSH using the `admin` user.
 - b. Change the mode to **Advanced**.

```
DataMig-cmode::> set -privilege advanced
```

- c. Enter 'y' when asked if you want to continue with advanced commands.
- d. Discover the source array on ONTAP. Wait for a few minutes and retry to detect the source array.

```
storage array show
```

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.

For example, the DataMig-cmode cluster initiator ports 0a and 0b of ONTAP are connected to Cisco ports 4/9 and 4/11. To reset port 4/9 on the Cisco switch:

```
conf t
interface fc4/
shutdown
no shutdown
exit
exit
```

Resetting one port is usually enough. Check the array list and LUN paths after resetting one port.

- e. Verify the source array is discovered through all the initiator ports:

storage array config show -array-name HITACHI_DF600F_1

Node Initiator	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				

- f. List the test LUN mapped from Hitachi storage and verify the disk properties and paths:

storage disk show -array-name HITACHI_DF600F_1 -instance

```
Disk: HIT-1.1
Container Type: unassigned
Owner/Home: - / -
DR Home: -
Stack ID/Shelf/Bay: - / - / -
LUN: 0
Array: HITACHI_DF600F_1
Vendor: HITACHI
Model: DF600F
Serial Number: 83017542001A
UID:
48495441:43484920:38333031:37353432:30303236:00000000:00000000:000000
00:00000000:00000000
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
-----
-----
-----
DataMig-cmode-01 0a 0 DM-Cisco9506-1:4-9 DM-
Cisco9506-1:2-24 AO INU 50060e801046b968 2 2 Gb/S 0 0
DataMig-cmode-01 0b 0 DM-Cisco9506-2:4-9 DM-
Cisco9506-2:2-24 AO INU 50060e801046b96c 2 2 Gb/S 0 0
DataMig-cmode-01 0b 0 DM-Cisco9506-2:4-9 DM-
Cisco9506-2:1-14 AO INU 50060e801046b964 1 2 Gb/S 0 0
DataMig-cmode-01 0a 0 DM-Cisco9506-1:4-9 DM-
Cisco9506-1:1-14 AO INU 50060e801046b960 1 2 Gb/S 0 0
DataMig-cmode-02 0a 0 DM-Cisco9506-1:4-11 DM-
Cisco9506-1:2-24 AO INU 50060e801046b968 2 2 Gb/S 0 0
DataMig-cmode-02 0b 0 DM-Cisco9506-2:4-11 DM-
Cisco9506-2:2-24 AO INU 50060e801046b96c 2 2 Gb/S 0 0
DataMig-cmode-02 0b 0 DM-Cisco9506-2:4-11 DM-
Cisco9506-2:1-14 AO INU 50060e801046b964 1 2 Gb/S 0 0
DataMig-cmode-02 0a 0 DM-Cisco9506-1:4-11 DM-
Cisco9506-1:1-14 AO INU 50060e801046b960 1 2 Gb/S 0 0

Errors:
-

DataMig-cmode::*>

```

4. Mark the source LUN as foreign using the serial number:

```
storage disk set-foreign-lun { -serial-number 83017542001A } -is-foreign true
```

a. Verify the source LUN is marked as foreign:

```
storage disk show -array-name HITACHI_DF600F_1
```

b. List all foreign arrays and their serial numbers:

```
storage disk show -container-type foreign -fields serial-number
```



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

5. Create a destination volume:

```
vol create -vserver datamig flivol aggr1 -size 10g
```

6. Create a test LUN using a foreign LUN:

```
lun create -vserver datamig -path /vol/flivol/testlun1 -ostype linux -foreign-disk 83017542001A
```

- a. List the test LUN and verify the size of the LUN with the source LUN:

```
lun show
```



For FLI offline migration, the LUN must be online to map it to an igroup and then must be offline before creating the LUN import relationship.

7. Create test igroup of protocol FCP without adding any initiators:

```
lun igroup create -vserver datamig -igroup testig1 -protocol fcp -ostype linux
```

8. Map the test LUN to the test igroup:

```
lun map -vserver datamig -path /vol/flivol/testlun1 -igroup testig1
```

9. Offline the test LUN:

```
lun offline -vserver datamig -path /vol/flivol/testlun1
```

10. Create import relationship with test LUN and foreign LUN:

```
lun import create -vserver datamig -path /vol/flivol/testlun1 -foreign-disk 83017542001A
```

11. Start the migration (import):

```
lun import start -vserver datamig -path /vol/flivol/testlun1
```

- a. Monitor the import progress:

```
lun import show -vserver datamig -path /vol/flivol/testlun1
```

- b. Check the import job is completed successfully:

```
lun import show -vserver datamig -path /vol/flivol/testlun1
```

```

vserver foreign-disk  path                operation admin
operational percent
                                in progress state state
complete
-----
-----
datamig 83017542001A  /vol/flivol/testlun1
                                import    started
                                completed
100

```

12. Start the verify job to compare source and destination LUNs. Monitor the verify progress:

```
lun import verify start -vserver datamig -path /vol/flivol/testlun1
```

```

DataMig-cmode::*> lun import show -vserver datamig -path
/vol/flivol/testlun1
vserver foreign-disk  path                operation admin operational
percent
                                in progress state state
complete
-----
-----
datamig 83017542001A  /vol/flivol/testlun1
                                verify    started
                                in_progress
44

```

a. Check the verify job is complete without any errors:

```
lun import show -vserver datamig -path /vol/flivol/testlun1
```

```

vserver foreign-disk  path                operation admin
operational percent
                                in progress state state
complete
-----
-----
datamig 83017542001A  /vol/flivol/testlun1
                                verify    started
                                completed
100

```

13. Delete the import relationship to remove the migration job:

```
lun import delete -vserver datamig -path /vol/flivol/testlun1
```

```
lun import show -vserver datamig -path /vol/flivol/testlun1
```

14. Unmap the test LUN from the test igroup:

```
lun unmap -vserver datamig -path /vol/flivol/testlun1 -igroup testig1
```

15. Online the test LUN:

```
lun online -vserver datamig -path /vol/flivol/testlun1
```

16. Mark the foreign LUN attribute to false:

```
storage disk modify -serial-number 83017542001A -is-foreign false
```



Do not remove the host group created on source storage with ONTAP initiator ports. The same host group is reused during migrations from that source array.

17. Remove test LUN from source storage.

- Log in to Hitachi Storage Navigator Modular as a system.
- Select **AMS 2100 array** and click **Show and Configure Array**.
- Log in using `root`.
- Select **Groups**, then select **Host Groups**.
- Select **cDOT_FLI Igroup** and click **Edit Host Group**.
- In the Edit Host Group window, select all target ports chosen to map the test LUN and select **Forced Set to All Selected Ports**.
- Select the **Logical Units** tab.
- Select the test LUN from the **Assigned Logical Units** window.
 - Select **Remove** to remove the LUN mapping.
 - Click **OK**.

18. Do not remove the host group and continue deleting the test LUN.

- Select **Logical Units**.
- Select the test LUN created in the previous step.
- Click **Delete LUN**.
- Click **Confirm** to delete the test LUN.

19. Delete the test LUN on the destination storage.

- Log in to the ONTAP storage through SSH using `admin` user.
- Offline the test LUN on the NetApp storage system:

```
lun offline -vserver datamig -path /vol/flivol/testlun1
```



Make sure you do not select another host LUN.

- Destroy the test LUN on the NetApp storage system:

```
lun destroy -vserver datamig -path /vol/flivol/testlun1
```

d. Offline the test volume on the NetApp storage system:

```
vol offline -vserver datamig -volume flivol
```

e. Destroy the test volume on the NetApp storage system:

```
vol destroy -vserver datamig -volume flivol
```

Host remediation

Recommended ONTAP FLI approaches by operating system and configuration

Align the host operating system and storage configuration with the recommended Foreign LUN Import (FLI) approach before selecting a migration method.

Intent

The table below helps align host operating system and storage configuration with the recommended Foreign LUN Import (FLI) approach before selecting a migration method. For each host type and usage model, it indicates whether Online FLI, Offline FLI, or either is appropriate. The Notes / Rationale column highlights key technical constraints—such as clustering semantics, SCSI 3 persistent reservations, and boot dependencies—that directly influence whether host I/O can safely coexist with an active import relationship. This guidance should be used early in planning to avoid unsupported workflows and to set realistic expectations around outage windows and remediation timing.

Key Guidance

- Clustering almost always implies Offline FLI due to SCSI 3 Persistent Reservation requirements.
- Standalone hosts and non clustered virtualized workloads are generally flexible and can use either Online or Offline FLI.
- SAN boot use cases are technically supported in both modes but are often operationally simpler to execute as Offline FLI.
- Final mode selection should factor in application downtime tolerance, rollback requirements, and operational complexity, not just protocol support.

Detailed Guidance

Windows

Configuration	Recommendation	Reasons if any
NTFS Volume	Online / Offline FLI	Supported in both modes; choice depends on application downtime tolerance.
Windows Failover Cluster (WFSC)	Offline FLI	SCSI 3 Persistent Reservations are required by WFSC and are not supported while a LUN is in an active import relationship.

Configuration	Recommendation	Reasons if any
Hyper V	Online / Offline FLI	Standalone Hyper V supported; clustered Hyper V follows WFSC guidance.
SAN Boot	Online / Offline FLI	Requires careful cutover planning; typically treated as offline unless downtime is well controlled.

Linux (RHEL/SLES)

Configuration	Recommendation	Reasons if any
Filesystem / RAW Disk	Online / Offline FLI	Supported in both modes with standard multipathing.
LVM	Online / Offline FLI	Supported; ensure correct host discovery and PV/VG consistency during cutover.
Pacemaker Cluster	Online / Offline FLI	Supported, but cluster behavior must be validated during cutover.
SAN Boot	Online / Offline FLI	Usually operationally simpler as offline, though technically supported online.

Solaris

Configuration	Recommendation	Reasons if any
Standalone	Online / Offline FLI	Either mode supported depending on outage tolerance.
Oracle Solaris Cluster (OSC / Sun Cluster)	Offline FLI	Cluster requires exclusive SCSI behaviors not supported during active import.
SAN Boot	Online / Offline FLI	Treat similarly to other OS SAN boot scenarios.

AIX

Configuration	Recommendation	Reasons if any
Standalone	Online / Offline FLI	Supported in both modes.
PowerHA (HACMP)	Offline FLI	Clustered disks require reservation semantics incompatible with online import.
Virtualized (PowerVM NPIV/VFC, vSCSI - no clustering)	Online / Offline FLI	Supported when not part of a cluster.
Virtualized (PowerVM with clustering)	Offline FLI	Same clustering limitations as physical PowerHA.

Configuration	Recommendation	Reasons if any
SAN Boot	Online / Offline FLI	Operational preference often offline.

HP-UX

Configuration	Recommendation	Reasons if any
Standalone	Online / Offline FLI	Supported in both modes.
Host Clustering (MC/ServiceGuard)	Offline FLI	Cluster disk semantics require offline migration.
SAN Boot	Online / Offline FLI	Usually planned as offline in practice.

VMware ESXi

Configuration	Recommendation	Reasons if any
VMDK (Standard Datastores)	Online / Offline FLI	Either mode supported; online minimizes VM downtime.
RDM (Standard)	Online / Offline FLI	Supported; ensure consistent host rescans during cutover.
VMSC / VMHA	Online / Offline FLI	Supported with careful coordination.
WFSC (Clustered VMDK)	Offline FLI	Windows clustering limitations apply.
WFSC (RDM)	Offline FLI	SCSI 3 PR required, online import unsupported.

ONTAP FLI AIX host remediation

Before running the ONTAP FLI migration on an AIX host, complete these pre-migration steps. There are separate steps for offline and online migrations.

Pre-Migration - Offline FLI

These steps are common for both online and offline FLI. Before unprovisioning the source LUN from the host, ensure the following steps are completed.

Filesystem	Block device	LVM (FS)	Cluster	Sanboot
Step 1: Capture LUN details for filesystems in scope for migration.	Step 1: Capture LUN details used as raw block devices (no filesystem) in scope for migration.	Step 1: Capture LUNs mapped to the host with LVM configuration that are in scope for migration.	Step 1: Capture overall cluster details and the disks involved in the migration.	Step 1: Capture SAN-booted LUN details and associated partition information.

Filesystem	Block device	LVM (FS)	Cluster	Sanboot
Step 2: Unmount the filesystem.	Step 2: Remove the SCSI block devices associated with the multipath device	Step 2: Unmount the LVM device if it is already mounted.	Step 2: Shut down cluster services gracefully as per the recommended procedure.	Step 2: Verify the boot device list from OS as well as from the IPL Initial Program Load/SMS menu. Ensure all devices are referenced by their device name properly.
Step 3: Remove the SCSI block devices associated with the multipath device.	Step 3: Rescan all the devices.	Step 3: Take a backup of the VG and deactivate/varyoff the VG. eg: varyoffvg vg_name	Step 3: Remove the SCSI block devices associated with the multipath device. eg: rmdev -Rdl hdiskx	
Step 4: Rescan the host to make sure no stale entries of the existing device mapping.	Step 4: Rescan the host to make sure no stale entries of the existing device mapping.	Step 4: Rescan the host to make sure no stale entries of the existing device mapping. eg: cfgmgr	Step 4: Rescan the host to ensure there are no stale entries in the existing device mapping.	



Refer to OS Vendor Documentation for Latest and updated command list.

Steps

1. Identify storage devices.

```
lsdev -Cc disk
sanlun lun show all (This works only for NetApp storage)
sanlun lun show all -pv
lsmpio -l hdiskx
```

2. Identify the volume group and filesystems.

```
lspv
lsvg <vg_name>
lsvg -l <vg_name>
lsvg -p <vs_name>
```

3. Varyoffvg and unmount filesystems.

```
df -g (to get the mount points)
umount /mount_point
varyoffvg <vg_name>
df -g (to ensure its unmounted)
```

4. Remove the scsi block devices.

```
rmdev -Rdl hdiskx
```

5. Verify device is detached.

```
lspv |grep hdiskx
```

6. Rescan after detaching to confirm clean state.

```
cfgmgr
```

Pre-Migration - Online FLI

Below are specific steps for online FLI but these steps must be preceded by steps in the Pre-Migration - Offline FLI section above.

The following steps must be carried out after the NetApp LUN is provisioned to the Host.

At this point the lun import is already created, lun is mapped to the host igroup and there is either an iSCSI or FC Initiator logged in from the Host.

For offline FLI, the NetApp LUN remains in the offline state.

Filesystem	Block device	LVM (FS)	Sanboot
Step 1: Rescan storage on all the mapped AIX hosts.	Step 1: Rescan storage on all the mapped AIX hosts.	Step 1: Rescan storage on all the mapped AIX hosts.	Step 1: Power on the server.
Step 2: Identify the multipath devices. eg: lsdev -Cc disk; lsmPIO -l hdiskx	Step 2: Identify the multipath devices. eg: lsdev -Cc disk; lsmPIO -l hdiskx	Step 2: Identify the multipath devices. eg: lsdev -Cc disk; lsmPIO -l hdiskx	Step 2: Verify the boot device list from the IPL Initial Program Load/SMS menu will list the new LUNS referencing the devices list coming from new Storage device. Set the boot order to the new device list path.
Step 3: Check for the filesystem details.	Step 3: Identify the multipath device ID and present it to the respective application.	Step 3: Ensure the new LUNs are discovered.	Step 3: Ensure that the server boots up properly.
Step 4: Mount the devices back to the respective folders.		Step 4: The VGs should be activated automatically.	

Filesystem	Block device	LVM (FS)	Sanboot
		Step 5: Verify the VG, LV and Filesystems are active and mountable.	



Refer to OS Vendor Documentation for Latest and updated command list.

Steps

1. Identify storage devices

```
lsdev -Cc disk
sanlun lun show all (This works only for NetApp storage)
sanlun lun show all -pv
lsmpio -l hdiskx
```

2. Identify the volume group and filesystems.

```
lspv
lsvg <vg_name>
lsvg -l <vg_name>
lsvg -p <vs_name>
```

3. Varyoffvg and unmount filesystems.

```
df -g (to get the mount points)
umount /mount_point
varyoffvg <vg_name>
df -g (to ensure its unmounted)
```

4. Remove the scsi block devices.

```
rmdev -Rdl
```

5. Verify device is detached.

```
lspv |grep hdiskx
```

6. Rescan after detaching to confirm clean state.

```
cfgmgr
```

ONTAP FLI ESXi host remediation

Before running the ONTAP FLI migration on an ESXi host, complete these pre-migration steps. There are separate steps for offline and online migrations.

Pre-Migration - Offline FLI

This group of steps are common for both Online and Offline FLI. Before unprovisioning the source LUN from the host, ensure that the following steps have been completed.

VMs/VMDK	RDM	WSFC	VMSC
Step 1: Shut down guest VMs gracefully.	Step 1: For VMs using RDMs, Identify the backend mapping with nmp device and remove RDM disks gracefully.	Step 1: If WFSC is present, shut down the cluster gracefully and follow Windows remediation steps to remove disks.	Step 1: If VM-HA or VMSC is configured, disable HA, release heartbeat disk(s) if present, and adjust cluster configuration as needed.
Step 2: Unmount datastores planned for migration; if shared, unmount from all hosts.	Step 2: Detach relevant NMP devices involved in the migration.	Step 2: Shut down guest VMs gracefully.	Step 2: Unmount datastores planned for migration; if shared, unmount from all hosts.
Step 3: Check VMFS locking mode (ATS-only vs ATS+SCSI). If ATS-only, change to SCSI (required only for Online FLI).		Step 3: Unmount datastores planned for migration; if shared, unmount from all hosts. a	Step 3: Detach relevant NMP devices involved in the migration.
Step 4: Detach relevant NMP devices involved in the migration.		Step 4: Detach relevant NMP devices involved in the migration.	

Key points to know

- The **ATS-only** setting is stored in the VMFS header, which resides on the first extent of the datastore. If the datastore spans multiple extents, only the first extent needs remediation.
- If the datastore is shared across multiple ESXi hosts, updating it from one host is typically sufficient; other hosts pick up the change after a rescan/remount.
- No active I/O can be running against the datastore while changing ATS-only mode. VMware guidance is to evacuate/migrate or shut down VMs and unmount/remount as required to apply the change.

ESX Example Commands: Sync Block

Put Host in Maintenance Mode

```
esxcli system maintenanceMode set --enable true
```

Identify storage devices and datastores

```
esxcli storage core device list
esxcli storage filesystem list
```

Check VMFS ATS locking mode

```
vmkfstools --queryfs -h /vmfs/volumes/<DATASTORE_NAME>
```

Unmount the datastore (repeat on all ESXi hosts if shared)

```
esxcli storage filesystem unmount -l <DATASTORE_NAME>
```

Detach the backing device (NMP / naa device)

```
esxcli storage core device set -d <naa.ID> --state=off
```

Verify device is detached

```
esxcli storage core device list -d <naa.ID>
```

Rescan after detaching to confirm clean state

```
esxcli storage core adapter rescan --all
```

Exit maintenance mode after completion

```
esxcli system maintenanceMode set --enable false
```

Pre-Migration - Online FLI

Below are specific steps of online FLI but these steps must be preceded by steps in the Pre-Migration Offline FLI.

The following steps must be carried out after the NetApp LUN is provisioned to the Host.

At this point the lun import is already created, lun is mapped to the host igroup and there is either iSCSI or FC Initiator logged in from the Host.

In case of offline FLI the NetApp LUN remains in offline state.

VMs/VMDK	RDM	VMSC
Step 1: Rescan storage on all mapped ESXi hosts.	Step 1: Rescan storage on all mapped ESXi hosts.	Step 1: Rescan storage on all mapped ESXi hosts.
Step 2: List volume/device details to confirm filesystem details on the LUN.	Step 2: Identify new NAA IDs and add RDMs back to the respective VMs.	Step 2: List volume/device details to confirm filesystem details on the LUN.
Step 3: Identify datastore details and mount datastores back to ESXi.	Step 3: Power on guest VMs and start applications.	Step 3: Identify datastore details and mount datastores back to ESXi.
Step 4: Power on guest VMs and start applications.		Step 4: Power on guest VMs and start applications.
		Step 5: Power on guest VMs and start applications



Refer to OS Vendor Documentation for latest and updated command list.

Put Host in Maintenance Mode

```
esxcli system maintenanceMode set --enable true
```

Identify storage devices and datastores

```
esxcli storage core device list
esxcli storage filesystem list
```

Check VMFS ATS locking mode

```
vmkfstools --queryfs -h /vmfs/volumes/<DATASTORE_NAME>
```

Unmount the datastore (repeat on all ESXi hosts if shared)

```
esxcli storage filesystem unmount -l <DATASTORE_NAME>
```

If ATS-only is detected, convert to ATS+SCSI

```
vmkfstools --configATSOOnly 0 /vmfs/volumes/<DATASTORE_NAME>
```

Detach the backing device (NMP / naa device)

```
esxcli storage core device set -d <naa.ID> --state=off
```

Verify device is detached

```
esxcli storage core device list -d <naa.ID>
```

Rescan after detaching to confirm clean state

```
esxcli storage core adapter rescan --all
```

ONTAP FLI HP host remediation

Before running the ONTAP FLI migration on an HP host, complete these pre-migration steps. There are separate steps for offline and online migrations.

Pre-Migration - Offline FLI

These steps are common for both online and offline FLI. Before unprovisioning the source LUN from the host, ensure the following steps are completed.

Filesystem (LVM & Non-LVM)	Host Cluster (SCSI3 PR)	SANBOOT
Step 1: Collect end-to-end Block Device configuration: multipath health/paths, LVM layout on the boot LUN, and UEFI settings confirming the correct boot device and boot parameters.	Step 1: Capture overall cluster details and the disks involved in the migration.	Step 1: Collect end-to-end SAN-boot configuration: multipath health/paths, LVM layout on the boot LUN, and UEFI settings confirming the correct boot device and boot parameters.
Step 2: Make sure the /etc/fstab entries for the respective data devices (LVM/Non-LVM) are proper.	Step 2: Shut down cluster services gracefully as per the recommended procedure.	Step 2: Make sure the /etc/fstab entries for SAN Boot device including the data devices (lvm/nonlvm) are proper

Filesystem (LVM & Non-LVM)	Host Cluster (SCSI3 PR)	SANBOOT
Step 3: Ensure all the operations to the target mount points are paused. Ensure there are no processes still accessing the blockdevice . If required drain/kill the open process handles.	Step 3: Remove the SCSI block devices associated with the multipath device.	Step 3: From the UEFI shell, verify the device mapping table and boot into hpux from the Storage device
Step 4: Umount the filesystem & perform the LVM volume group configuration backup. Deactivate the VG & export the VG	Step 4: Rescan the host to ensure there are no stale entries in the existing device mapping.	Step 4: Make sure primary boot device is set to Storage device so that it will automatically boot on subsequent reboots
Step 5: Remove all the stale SCSI device files.		Step 5: Reboot the host to ensure the host can come up fine the above changes. Once the host comes up fine, understand the host pre remediation for FLI migration is ready.
Step 6: Rescan the sci bus so that the host has the updated device list eg: ioscan -fnC disk		Step 6: Now shutdown the sanboot host, to proceed with ONTAP FLI workflow.

Pre-Migration - Online FLI

Below are specific steps of Online FLI but these steps must be preceded by steps in the Pre-Migration Offline FLI The following steps must be carried out after the NetApp LUN is provisioned to the Host.

Filesystem	Block device	LVM (FS)	Sanboot
Step 1: Rescan storage on all the mapped HP-UX hosts	Step 1: Rescan storage on all the mapped HP-UX hosts.	Step 1: Rescan storage on all the mapped HP-UX hosts	Step 1: PowerON the server
Step 2: identify the multipath devices	Step 2: identify the multipath devices	Step 2: identify the multipath devices	Step 2: From the UEFI shell, verify the device mapping table and boot into hpux from the new Storage device
Step 3: Check for the filesystem details	Step 3: Identify the multipath device ID and present it to the respective application.	Step 3: Ensure the new LVM device files are properly created	Step 3: modify the primary boot device to new Storage lun Restart the host to boot from newly mapped LUN that is part of LUN Import
Step 4: Mount the devices back to the respective folders		Step 4: Import the VG/LVM	
		Step 5: Activate the VG	
		Step 6: List down the LVs present on the VG	

Filesystem	Block device	LVM (FS)	Sanboot
		Step 7: Mount the LVMs back to the respective folders	



Refer to OS Vendor Documentation for Latest and updated command list.

Steps

1. Identify storage devices.

```
ioscan -fnC disk
sanlun lun show all (Works on Netapp Storage)
sanlun lun show all -pv
sanlun lun show -wwpn
```

2. Identify the volume group and filesystems.

```
vgdisplay -v /dev/vg_name
vgdisplay /dev/vg_name
lvdisplay /dev/vg_name/lv_name
```

3. Unmount filesystem and export VG.

```
bdf
umount /mount_point
vgchange -a n /dev/vg_name
vgexport /dev/hpux_vg_online
```

4. Rescan after detaching to confirm clean state.

```
insf -e
ioscan -fnC disk
```

ONTAP FLI Linux host remediation

Before running the ONTAP FLI migration on a Linux host, complete these pre-migration steps. There are separate steps for offline and online migrations.

Pre-Migration - Offline FLI

These steps are common for both online and offline FLI. Before unprovisioning the source LUN from the host, ensure the following steps are completed.

Filesystem (LVM & Non-LVM)	Pacemaker with LVM (SCSI3 PR)	SANBOOT
Step 1: Collect end-to-end Block Device configuration: multipath health/paths, LVM layout on the block device and GRUB/bootloader settings confirming the correct boot device and boot parameters.	Step 1: Collect end-to-end Block Device configuration: multipath health/paths, LVM layout on the block device and GRUB/bootloader settings confirming the correct boot device and boot parameters.	Step 1: Collect end-to-end SAN-boot configuration: multipath health/paths, LVM layout on the boot LUN, and GRUB/bootloader settings confirming the correct boot device and boot parameters.
Step 2: Make sure the <code>/etc/fstab</code> entries for the respective data devices (LVM/Non-LVM) are referenced by their FS UUID.	Step 2: Disable the cluster resource groups for the respective LUNs being migrated. Once disabled, delete the resource groups.	Step 2: Make sure the <code>/etc/fstab</code> entries for SAN Boot device including the data devices (lvm/nonlvm) are referenced by their FS UUID.
Step 3: Ensure all the operations to the target mount points are paused. Ensure there are no processes still accessing the block device. If required, drain/kill the open process handles.	Step 3: Perform the LVM VG configuration backup to a file for the respective resource groups that we just disabled/deleted. Followed by Deactivation of the VGs.	Step 3: Ensure that you have updated the <code>GRUB_CMDLINE_LINUX</code> line with the correct UUID references. Recreate a new GRUB settings based on the UUID device references.
Step 4: Umount the filesystem and perform the LVM volume group configuration backup. Deactivate the VG and export the VG	Step 4: Remove the source LUNs from the PCS stonith configuration.	Step 4: Validate the changes with the new grub settings and recreate the new initrd image to reflect changes with SAN Boot Device.
Step 5: Flush the multipath configuration for the device and remove all the stale SCSI device files.	Step 5: Flush the multipath configuration for the device and remove all the stale SCSI device files. NOTE: If any lun persistent reservations is seen, clear them and ensure no PR keys are	Step 5: Reboot the host to ensure the host can come up fine the above changes. Once the host comes up fine, understand the host pre remediation for FLI migration is ready.
Step 6: Reload the multipath configuration so that the host has the updated device list. If required, perform the host SCSI Bus rescan.	Step 6: Rescan the host to ensure there are no stale entries in the existing device mapping.	Step 6: Now shutdown the sanboot host, to proceed with ONTAP FLI workflow.

Pre-Migration - Online FLI

Below are specific steps of online FLI but these steps must be preceded by steps in the pre-migration offline FLI. The following steps must be carried out after the NetApp LUN is provisioned to the host. These steps must be carried out after the NetApp LUN is provisioned to the host. And it is applicable only the online FLI.

Filesystem (LVM & Non-LVM)	Pacemaker with LVM (SCSI3 PR)	SANBOOT
Step 1: Rescan storage on all the mapped Linux hosts.	Step 1: Rescan storage on all the mapped Linux hosts.	Step 1: PowerON the server.
Step 2: Discover the multipath devices.	Step 2: Discover the multipath devices. Ensure the migrated LUNs are added to PCS Stonith Configuration.	Step 2: On the server HBA BIOS or UEFI configure the new boot device that is presented to this Host HBA.

Filesystem (LVM & Non-LVM)	Pacemaker with LVM (SCSI3 PR)	SANBOOT
Step 3: Import the Volume Groups and perform the VG activation.	Step 3: Perform VG Configuration Backup restore for those newly discovered VG/LVM, followed by VG activation.	Step 3: Restart the host to boot from newly mapped LUN that is part of LUN Import.
Step 4: Mount the LVM config on to the respective mount points.	Step 4: Create the new PCS resource groups and add the respective resources for the available LV and filesystems. NOTE: Please note that the newly activated LVM objects have the LV and filesystem structures intact.	
	Step 5: Verify the PCS Cluster status for all the resources are online and available status.	

ONTAP FLI Solaris host remediation

Before running the ONTAP FLI migration on a Solaris host, complete these pre-migration steps. There are separate steps for offline and online migrations.

Pre-Migration - Offline FLI

These steps are common for both online and offline FLI. Before unprovisioning the source LUN from the host, ensure the following steps are completed.

Filesystem (ZFS/Zpool)	OSC (Oracle Sun Cluster) (SCSI3 PR) with Zpool	SANBOOT
Step 1: Collect end-to-end Block Device configuration: multipath health/paths, Zpool layout on the block device and OBP (ok prompt) settings confirming the correct boot device and boot parameters.	Step 1: Collect end-to-end Block Device configuration: multipath health/paths, LVM layout on the block device and OBP settings confirming the correct boot device and boot parameters.	Step 1: Collect end-to-end SAN-boot configuration: multipath health/paths, RPool layout on the boot LUN, and OBP settings confirming the correct boot device and boot parameters.
Step 2: Ensure all the operations to the target mount points are paused. Ensure there are no processes still accessing the block device . If required drain/kill the open process handles.	Step 2: Disable the cluster resource groups for the respective LUNs being migrated.	Step 2: Reboot the host to ensure the host can come up fine the above changes. Once the host comes up fine, understand the host pre remediation for FLI migration is ready.
Step 3: Ensure the Pool is not in use Unmount any filesystems using the pool:	Step 3: Ensure the pool is not in use	Step 3: Now shutdown the sanboot host, to proceed with ONTAP FLI workflow.
Step 4: Export the pool Verify export	Step 4: Remove the source LUNs from the OSC configuration.	

Filesystem (ZFS/Zpool)	OSC (Oracle Sun Cluster) (SCSI3 PR) with Zpool	SANBOOT
Step 5: If required, perform the host SCSI Bus rescan. eg: devfsadm	Step 5: Remove all the stale SCSI device files if they exist NOTE: If any lun persistent reservations is seen, clear them and ensure no PR keys are present	
	Step 6: Rescan the host to ensure there are no stale entries in the existing device mapping.	



Refer to OS Vendor Documentation for Latest and updated command list.

Steps

1. Identify storage devices.

```
echo | format
sanlun lun show all (This is only applicable for NetApp devices)
mpathadm list lu
mpathadm show lu lu_name
luxadm display node wwn/ disk path
```

2. Identify the zpool and filesystems.

```
zpool status <zpool_name>
df -h
```

3. Unmount and export filesystems.

```
zfs umount /mount_point
zpool export <zpool_name>
df -h (to ensure its unmounted)
```

4. Remove the scsi block devices.

```
rm /dev/rdisk/disk_ID
```

5. Verify device is detached.

```
format | grep disk_id
```

6. Rescan after detaching to confirm clean state.

```
cfgadm -al
```

Pre-Migration - Online FLI

Below are specific steps of online FLI but these steps must be preceded by steps in the pre-migration - offline FLI. The following steps must be carried out after the NetApp LUN is provisioned to the Host.

Filesystem	Block device	LVM (FS)	Sanboot
Step 1: Rescan storage on all the mapped Solaris hosts.	Step 1: Rescan storage on all the mapped Solaris hosts.	Step 1: Rescan storage on all the mapped Solaris hosts.	Step 1: Power on the server.
Step 2: Identify the multipath storage devices.	Step 2: Identify the multipath devices.	Step 2: Identify the multipath devices.	Step 2: On the SPARC server OBP, configure the new boot device that is presented to this Host HBA. Modify the nvalias to represent new disk path.
Step 3: Check for the filesystem details.	Step 3: Identify the multipath device ID and present it to the respective application.	Step 3: Verify the Pool status.	Step 3: Restart the host to boot from newly mapped LUN that is part of LUN Import.
Step 4: Mount the devices back to the respective folders.		Step 4: Import the Pool.	
		Step 5: Check the pool status.	
		Step 6: Mount the pool/ZFS.	



Refer to OS Vendor Documentation for Latest and updated command list.

Steps

1. Identify storage devices.

```
format
sanlun lun show all
mpathadm list lu
mpathadm show lu
```

2. Identify the zpool and filesystems.

```
zpool status <zpool_name>
df -h
```

3. Varyoffvg and unmount filesystems.

```
zfs umount /mount_point
zpool export <zpool_name>
df -h (to ensure its unmounted)
```

4. Remove the scsi block devices.

```
rm /dev/rdisk/disk_ID
```

5. Verify device is detached.

```
format | grep disk_id
```

6. Rescan after detaching to confirm clean state.

```
cfgadm -al
```

ONTAP FLI Windows host remediation

Before running the ONTAP FLI migration on a Windows host, complete these pre-migration steps. There are separate steps for offline and online migrations.

Pre-Migration - Offline FLI

These steps apply to both Offline and Online FLI. Complete the following host-side preparation before unmapping the source LUNs from the Windows host.

Filesystem devices	Physical drives (no filesystem)	WSFC (WFC)	Hyper V
Step 1: Identify the filesystem disks that panned for migration	Step 1: Identify the physical disks that panned for migration	Step 1: Identify the Cluster disks that panned for migration	Step 1: Stop all the application on the Virtual Machines and Shutdown the Virtual Machines.
Step 2: Shutdown the applications gracefully	Step 2 Shutdown the applications gracefully	Step 2 Shutdown the applications gracefully	Step 2: Offline the disk from the host.
Step 3: Offline the disks		Step 3: Shutdown the cluster and stop the cluster services on all the cluster nodes	Step 3: Unmap all the LUNs from the host.

Filesystem devices

Steps

1. Identify the filesystem disks that planned for migration.

```
Get-Disk | Format-Table Number, FriendlyName, Size, PartitionStyle
```

2. Shutdown the applications gracefully.

```
Stop-Service -Name <ServiceName>
```

3. Offline the disks.

```
Set-Disk -Number <DiskNumber> -IsOffline $true
```

Physical drives (no filesystem)

Steps

1. Identify the physical disks that planned for migration.

```
Get-PhysicalDisk | Format-Table FriendlyName, Size, MediaType
```

2. Shutdown the applications gracefully.

```
Stop-Service -Name <ServiceName>
```

WSFC (Windows Server Failover Clustering)

Steps

1. Identify the Cluster disks that planned for migration.

```
Get-ClusterResource | Where-Object ResourceType -eq "Physical Disk"
```

2. Shutdown the applications gracefully.

```
Stop-ClusterResource -Name <ResourceName>
```

3. Shutdown the cluster and stop the cluster services on all the cluster nodes.

```
Stop-Cluster -Name <ClusterName>
```

Hyper V

Steps

1. Stop all the application on the Virtual Machines and Shutdown the Virtual Machines.

```
Stop-VM -Name <VM_Name> -Force
```

2. Offline the disk from the host.

```
Set-Disk -Number <DiskNumber> -IsOffline $true
```

3. Unmap all the LUNs from the host from the storage array side.

4. Make sure the Disks are removed properly from the Host.

```
Get-Disk | Where-Object OperationalStatus -ne "Online"
```

Pre-Migration - Online FLI

Below are specific steps of online FLI but these steps must be preceded by steps in the pre-migration - offline FLI.

The following steps must be carried out after the NetApp LUN is provisioned to the host.

At this point the lun import is already created, lun is mapped to the host igroup and there is either iSCSI or FC Initiator logged in from the host.

For offline FLI, the NetApp LUN remains in the offline state.

Filesystem devices	Physical drives (no filesystem)
Step 1: Rescan the storage on all the mapped hosts.	Step 1: Rescan the storage on all the mapped hosts.
Step 2: Identify the drives that are listed and online the disks.	Step 2: add the physical disk to the required application and start.
Step 3: Verify if all the filesystem disks are available and working.	
Step 4: Start the required applications.	

Filesystem devices

Steps

1. Rescan the storage on all the mapped hosts.

Update-HostStorageCache

2. Identify the drives that are listed and online the disks.

```
Get-Disk | Where-Object IsOffline -eq $true  
Set-Disk -Number <DiskNumber> -IsOffline $false
```

3. Verify if all the filesystem disks are available and working.

Get-Volume | Format-Table DriveLetter, FileSystemLabel, Size, SizeRemaining

4. Start the required applications.

Start-Service -Name <ServiceName>

Physical drives (no filesystem)

Steps

1. Rescan the storage on all the mapped hosts.

Update-HostStorageCache

2. Add the physical disk to the required application and start.

Application-specific configuration

WSFC

Steps

1. Rescan the storage on all the mapped hosts.

Update-HostStorageCache

2. Identify the drives that are listed and online the disks.

```
Get-Disk | Where-Object IsOffline -eq $true  
Set-Disk -Number <DiskNumber> -IsOffline $false
```

3. Verify if all the mapped hosts has the drives available and online.

```
Get-ClusterResource | Where-Object ResourceType -eq "Physical Disk"
```

4. Start the WSFC cluster application and configure the required clustered resources.

```
Start-Cluster -Name <ClusterName>  
Start-ClusterResource -Name <ResourceName>
```

Hyper V

Steps

1. Rescan the disk management, All the Disks Should connect.

Update-HostStorageCache

2. Select the disks and bring them online, all Should come ONLINE and should get original name and disk letter.

```
Get-Disk | Where-Object IsOffline -eq $true  
Set-Disk -Number <DiskNumber> -IsOffline $false
```

3. Try to power-on Virtual Machines, Virtual Machine should come online without issue.

```
Start-VM -Name <VM_Name>
```

ONTAP Foreign LUN Import (FLI) Execute phase

The execute phase is where the migration is performed. It covers the end-to-end tasks required to run either an offline or online Foreign LUN Import and to transition host access from the foreign array to the ONTAP array.

ONTAP FLI execute phase tasks

The following are the ONTAP FLI execute tasks for each component.

Component	Tasks
Host	• Review host event and system logs to catch issues early and reduce risk before any major reconfiguration. • Reboot hosts as appropriate to ensure stability prior to early cutover or remediation. • After cutover, validate host access. All expected devices are present, multipathing is healthy, and applications and file systems come up cleanly.
iSCSI FLI and the network	• Confirm IC LIF reachability to the foreign array target portals and ensure backend connectivity remains stable throughout import. • Validate iSCSI sessions remain healthy because they drive foreign LUN discovery and path availability.
Fabric	• Ensure required zoning and masking is in place so ONTAP initiator ports can access the foreign array LUNs for discovery and import. • After cutover, update zoning as needed so hosts access the destination ONTAP LUNs and prepare for cleanup later.
NetApp storage array (Destination)	• Confirm foreign LUNs are discovered in ONTAP before creating import jobs. • Create the import relationship and start import jobs after discovery is confirmed. • Monitor import progress and handle exceptions. Use pause and resume as needed to regulate copy streams, especially if host latency increases. • Cutover handling: ◦ Offline FLI cutover occurs after import completes. ◦ Online FLI host cutover occurs earlier while import continues in the background.
Foreign storage array (Source)	• Ensure the source LUNs intended for migration are correctly presented for ONTAP discovery and import (FC masking or iSCSI target visibility depending on transport). • Ensure that the Source LUNs are not accessed via the host during the import. This is critical to ensure that write operations do not occur from multiple sources. • During online cutover, coordinate unmapping and remapping actions as required to prevent dual presentation and ensure the host transitions cleanly to ONTAP.

For both FC and iSCSI-backend FLI, execution success depends heavily on maintaining stable backend connectivity, by ensuring FC/iSCSI session health and full path redundancy to the foreign LUN. Because the host I/O (online) and import copy can compete for the same backend resources, the import may be throttled operationally (for example, using pause and resume) if needed to protect host latency, and migrations are typically scheduled during non-peak windows.

ONTAP Foreign LUN Import (FLI) Verify phase

The verify phase confirms that the migration completed as planned and focuses on post-migration cleanup. This includes validating application behavior, confirming correct host access and paths to the ONTAP LUNs, collecting migration logs for auditing and troubleshooting, and removing temporary connectivity and configuration used only for the migration.

As part of cleanup, you remove source-side initiator records and zoning that are no longer needed, and you document the final configuration so operations teams have a clear “as-built” state.

For iSCSI FLI, the verify phase should also include removing the software iSCSI initiator sessions and initiators created on intercluster LIFs. If the intercluster LIFs were created only for the foreign LUN import, even those LIFs must be deleted.

ONTAP FLI verify phase tasks

The following are the ONTAP FLI verify tasks for each component.

Component	Tasks
Host	• Perform application-level validation against the migrated LUNs. (file system consistency check, checksum error, dbcc checkdb (MS SQL), rman (Oracle)) • Confirm all expected file systems, datastores, cluster disks are accessible and stable after migration and remediation. • Confirm the IO Path Count and Path Type are on expected lines.
FC FLI fabric	• Remove host-to-source storage zones from the active zone set, or equivalent access rules, if not removed already during the early cutover. • Remove source-to-destination connectivity zones created specifically for migration. • Back up the updated post-migration zone set for rollback and audit purposes.
iSCSI FLI network	• Remove iSCSI FLI backend connectivity objects created for migration. • The iSCSI initiator sessions must be deleted and the iSCSI Initiator Object created on the intercluster LIFs must be deleted.

Component	Tasks
NetApp storage array (Destination)	• Verify post-migration storage correctness, including LUN alignment where required by your host or use case. • Collect FLI migration import logs for audit and troubleshooting. Trigger AutoSupport after the migration is complete. • Remove any manually created Snapshot copies used solely for rollback during the migration window. • Remove source-array configuration artifacts from ONTAP that are no longer required, such as source array objects and foreign-disk association details. Confirm the same using the storage disk show command. • Document the final ONTAP configuration, including the LUN paths, mappings, igroups, initiators, and host access model.
Foreign storage array (Source)	• Remove destination-storage (NetApp) initiator groups and host groups created to present source LUNs to ONTAP during migration. • Remove host initiator groups and host records related to the original host mapping, when decommissioning source access. • Ensure migrated source LUNs are no longer presented in a way that could cause dual presentation to the host. This is critical and must be ensured throughout the migration.

Copyright information

Copyright © 2026 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.