NetApp

Migration process overview

ONTAP FLI

NetApp
January 07, 2026

Table of Contents

Migration process overview

Migration process overview

The FLI migration process is a five-phase methodology that applies to any data migration: discover, analyze, plan, execute, and verify.

These phases provide a general framework to help identify where common tasks are performed throughout the migration process. The charts in this section show the tasks that can be performed in parallel in each of the four major components: host, fabric, destination storage, and source storage.

Discover phase workflow

The discover phase of the migration process focuses on collecting information used for host remediation and creating migration plans in the later steps. The collection of most information is automated using a data collection tool, such as OneCollect.

The following figure shows the discover phase workflow.
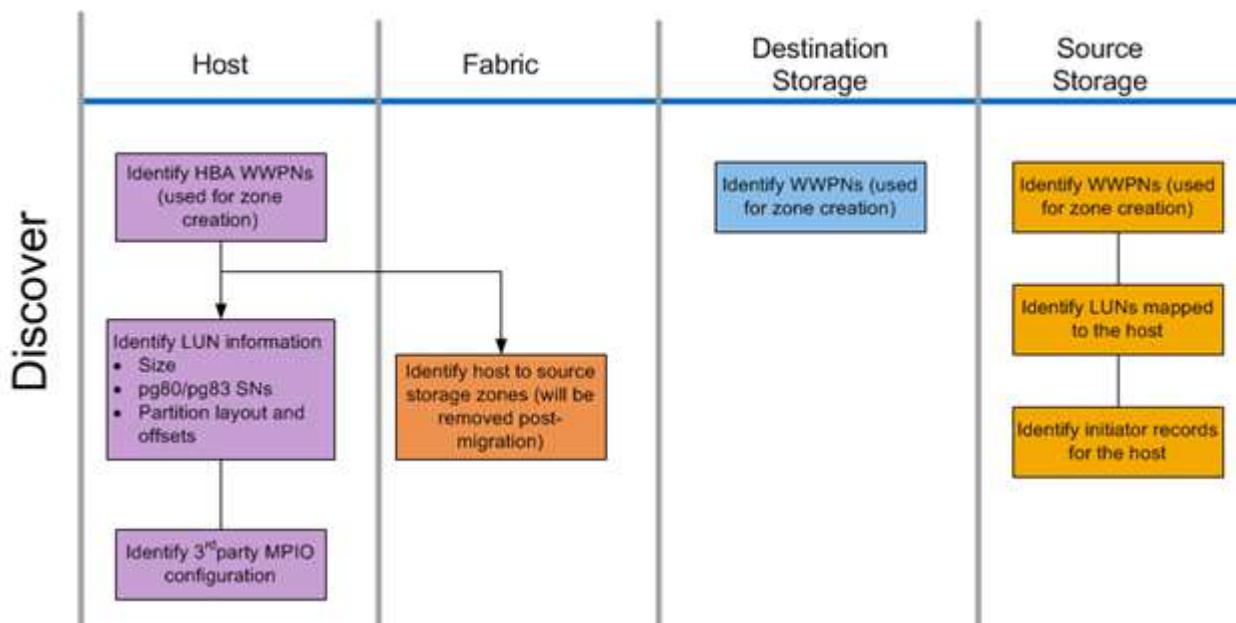
The discover phase tasks are listed in the following table.

Component	Tasks
Host	1. Identify HBA WWPNs (used for zone creation). 2. Identify LUN information (size, serial numbers, partition layouts, and offsets). 3. Identify third-party MPIO configuration, host operating system, HBA/CNA models and firmware, and so on.
Fabric	Identify host to source storage zones. (These are removed post-migration).
Destination storage	Identify the WWPNs for the ports that will be used for initiator/target usage.
Source storage	1. Identify WWPNs (used for zone creation). 2. Identify LUNs mapped to the host. 3. Identify initiator records for the host.

Analyze phase workflow

The analyze phase focuses on items that must be addressed before migration planning. Host configuration specifics that fall outside of the Interoperability Matrix must be identified.

For each host, a target configuration (post-migration) is identified, and a gap analysis is performed to identify specific components that are not supported. The host analysis should be reviewed immediately upon completion. Required updates might break compatibility with applications running on each host.

Usually, required host changes are not made until the actual migration event. This is due to the common need to schedule maintenance windows, but it is often less risky to do host changes in advance where possible, such as system patching and host bus adapter (HBA) updates. In addition, system updates are frequently done in coordination with application updates utilizing the same maintenance events. Typically, any changes made to the multipath I/O (MPIO) configuration before migration will affect the support of the current storage as well. For example, removing PowerPath from a host and reconfiguring it to use native MPIO and Asymmetric Logical Unit Access (ALUA) on Linux might not be supported by the current storage configuration.

Delaying MPIO reconfiguration until after the migration simplifies the process for rollback if required.

The planning phase tasks are listed in the following table.

Component	Tasks
Host	1. Perform a gap analysis for each host. Identify required hot fixes/patches, OS updates, HBA driver, and firmware upgrades required to match the selected target configuration on the NetApp IMT. In addition, requirements for other NetApp software to be installed on this host (SnapDrive®, SnapManager®) should be taken into consideration. 2. Determine a target configuration (post-migration) for each host (OS configuration, MPIO, HBA details, Host Utility Kit version). 3. Determine additional NetApp product requirements (SnapDrive, SnapManager).

Related information

NetApp Interoperability

Plan phase workflow

The planning phase of the data migration process focuses on tasks required to create detailed migration plans and to make sure that everything is ready for the actual migration. The majority of migration work is the planning performed during this phase.

The planning phase is where you develop a remediation plan using the host gap analysis information gathered in the analysis phase. Use the host remediation information while planning. After end-to-end connectivity is verified, a test migration is performed to make sure everything is properly configured before beginning the production migration.

The following figure shows the plan workflow.

The planning phase tasks are listed in the following table.

Component	Tasks
Fabric	1. Back up the existing zonesets. 2. Zone the source storage to destination storage.

Component	Tasks
Destination storage	1. Connect initiator ports to fabric. 2. Discover source storage and test LUN. Mark the source LUN as foreign. 3. Create test LUN using foreign LUN. 4. Create test igroup and map test LUN. 5. Offline test LUN. 6. Create import relationship and start test migration. 7. Delete import relationship and unmap test LUN. 8. Remove the test LUN.
Source storage	1. Create host group for destination storage using initiator port WWPNs. 2. Create test LUN (1GB). 3. Assign (map/mask) test LUN to destination storage host group. 4. Remove the test LUN.

FLI supported configurations

The FLI environment must be deployed in a supported manner to ensure proper operation and support. As engineering qualifies new configurations, the list of supported configurations will change. Refer to the NetApp Interoperability Matrix to verify support for specific configurations.

ONTAP 8.3 and later are the only supported destination storage. Migrations to third-party storage are not supported.

For a list of supported source storage arrays, switches, and firmware, see the Interoperability Matrix. The data migration program will provide support for the configurations in the NetApp Interoperability Matrix.

Once the import is complete and all LUNs have been migrated to NetApp controllers, ensure that all configurations are supported.

Related information

NetApp Interoperability Matrix Tool

Execute phase workflows

The execute phase focuses on the LUN migration tasks for performing an FLI offline or online migration.

The host event logs are reviewed in order to find and correct any problems and reduce risk. The hosts are rebooted to make sure that there are no underlying issues with the hosts before major reconfiguration occurs.

After the source LUNs are visible on the destination storage, migration jobs can be created and executed. After migrations are complete (FLI offline) or the FLI LUN relationship is established (FLI online), the host is zoned to the destination storage. New LUNs are mapped, and host remediation can begin for drivers, multipath software, and any other updates that have been identified in the analyze phase.

Offline migration workflow

The offline migration workflow is performed at the execute phase of the migration process. The offline workflow diagram shows the tasks that are performed on the host, the fabric, the destination storage and the source storage.
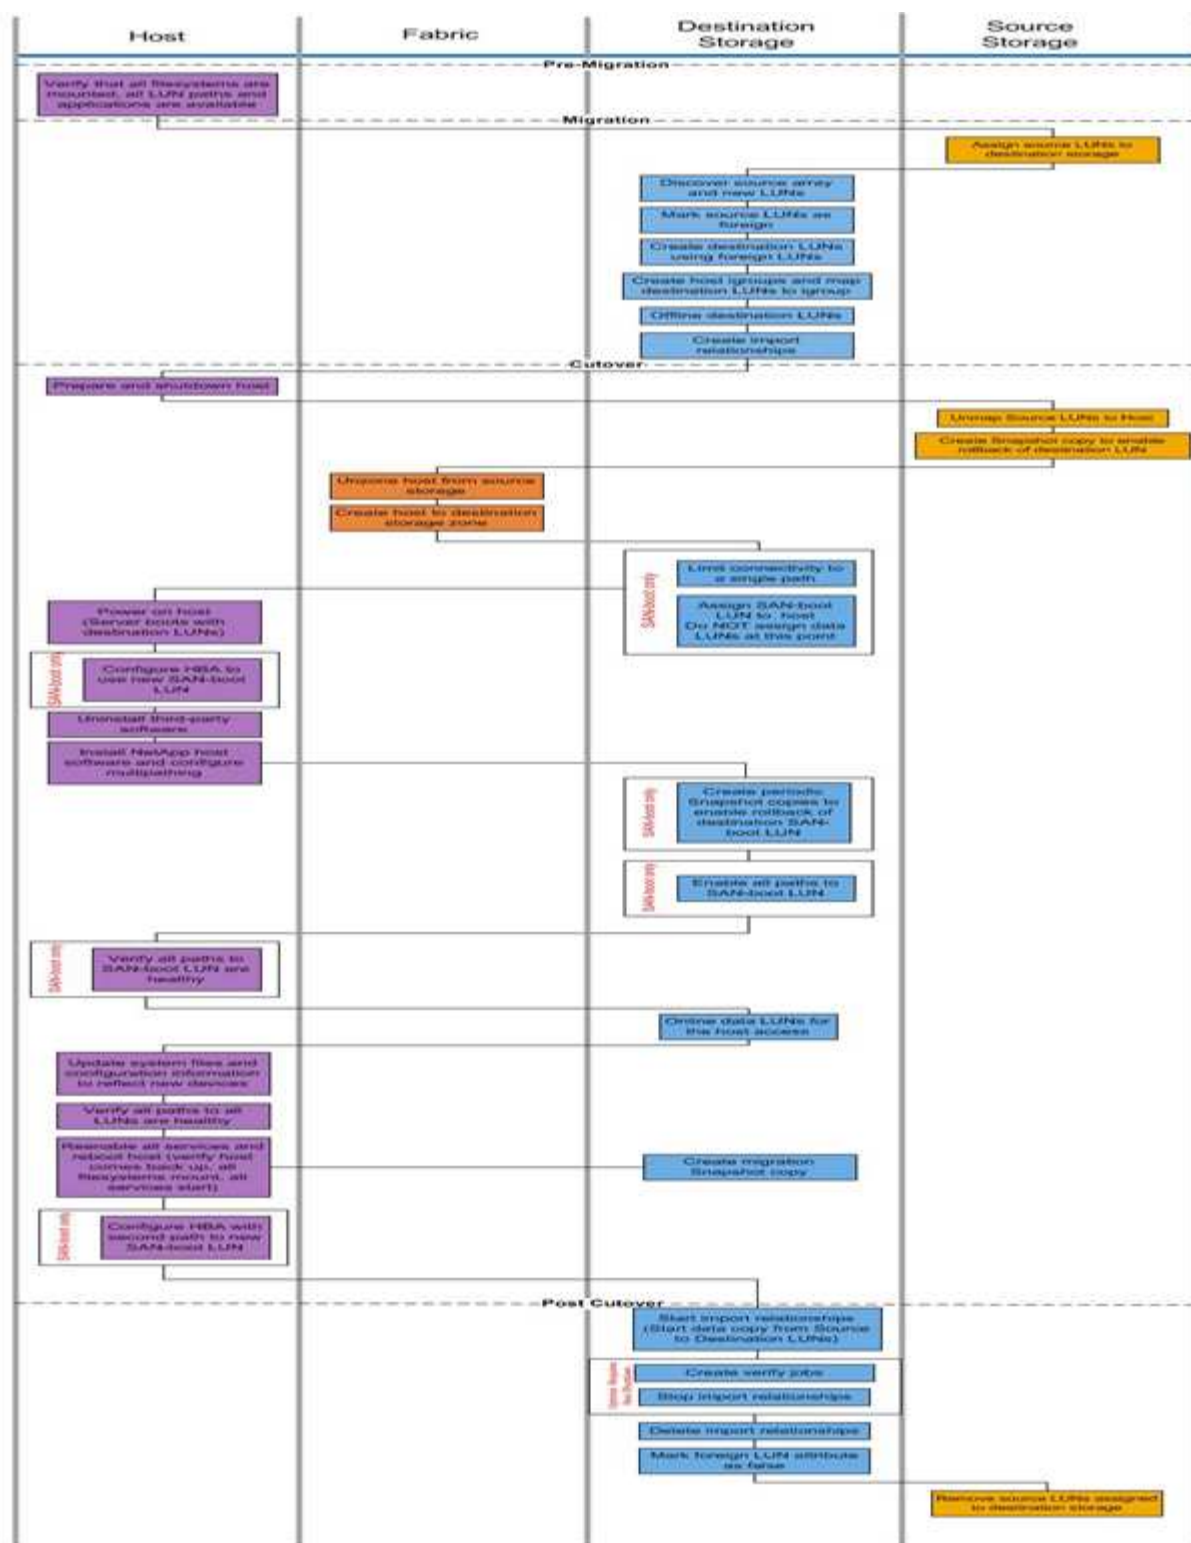
The offline workflow tasks are listed in the following table.

Component	Tasks
Host	1. Reboot the host and verify that all file systems mount, all LUN paths are available, and services start.
	2. Prepare and shut down the host.
	3. After the migration completes, power on the host.
	4. Configure the HBA to use new SAN boot LUN (SAN boot only).
	5. Uninstall third-party MPIO.
	6. Install NetApp host software and configure multipathing.
	7. Verify all paths to SAN boot LUN are healthy (SAN boot only).
	8. Update system files and configuration to reflect new devices.
	9. Verify all paths to all LUNs are healthy.
	10. Reenable all services and reboot host (verify host comes back up, all file systems mount, all services start).
	11. Configure the HBA with the second path to new SAN boot LUN (SAN boot only).
Fabric	1. Unzone the host from the source storage.
	2. Create host to destination storage zone.

Component	Tasks
Destination storage	1. Discover source array and new LUNs. 2. Mark source LUNs as foreign. 3. Create destination LUNs using foreign LUNs. 4. Create host initiator igroups and map destination LUNs to igroup.migration Snapshot copy. 5. Offline destination LUNs. 6. Create import relationships and start import jobs. 7. Create verify jobs (optional). 8. Delete import relationships. 9. Mark foreign LUN attribute to false. 10. Online destination LUNs. 11. Create a Snapshot® copy to enable rollback of destination LUN. 12. Limit connectivity to a single path (SAN boot only). 13. Assign SAN boot LUN to host; do not assign data LUNs at this point (SAN boot only). 14. Verify all host ports are logged in. 15. Create periodic Snapshot copies to enable rollback of destination SAN boot LUN (SAN boot only). 16. Enable all paths to SAN boot LUN (SAN boot only). 17. Assign data LUNs to host. 18. Create a final Snapshot copy.
Source storage	1. Assign source LUNs to destination storage. 2. Remove the source LUNs assigned to destination storage.

Online migration workflow

The online migration workflow is performed at the execute phase of the migration process. The online workflow diagram shows the tasks that are performed on the host, the fabric, the destination storage and the source storage.

The online workflow tasks are listed in the following table.

| Component	Tasks
Host	1. Verify that all file systems are mounted and all LUN paths and applications are available. 2. **Optional:** If the LUNs being imported are for ESX, review and follow the instructions in Appendix A: ESX CAW/ATS Remediation. 3. Prepare and shut down the hosts. 4. Power on the hosts with destination LUNs. 5. Configure the HBA to use new SAN boot LUN (SAN boot only). 6. Uninstall third-party MPIO. 7. Install NetApp host software and configure multipathing. 8. Verify all paths to SAN boot LUN are healthy (SAN boot only). 9. Update system files and configuration to reflect new devices. 10. Verify all paths to all LUNs are healthy. 11. Reenable all services and reboot host (verify host comes back up, all file systems mount, all services start). 12. Configure the HBA with the second path to new SAN boot LUN (SAN boot only).
Fabric	1. Unzone the host from the source storage. 2. Create host to destination storage zone.

Component	Tasks
Destination storage	1. Discover source array and new LUNs. 2. Mark source LUNs as foreign. 3. Create destination LUNs using foreign LUNs. 4. Create host initiator igroups and map destination LUNs to igroup. 5. Offline the destination LUNs. 6. Remove hosts from source array LUN masking (igroups). 7. Create import relationships and start import jobs. 8. Perform host step 4 earlier (remap hosts to new LUN locations). 9. Limit connectivity to a single path (SAN boot only). 10. Assign SAN boot LUN to host; do not assign data LUNs at this point (SAN boot only). 11. Create periodic Snapshot copies to enable rollback of destination SAN boot LUN (SAN boot only). 12. Enable all paths to SAN boot LUN (SAN boot only). 13. Online destination LUNs. 14. Create a Snapshot copy to enable rollback of destination LUN. 15. Start import relationships (start data copy from source to destination LUNs). 16. Create verify jobs and stop import relationships (optional). 17. Delete import relationships. 18. Mark foreign LUN attribute to false.
Source storage	1. Assign source LUNs to destination storage. 2. Unmap source LUNs to host. 3. Create Snapshot copy to enable rollback of destination LUN. 4. Remove the source LUNs assigned to destination storage.

Verify phase workflow

The verify phase of the migration process focuses on the post-migration cleanup and confirming the accuracy of the execution of the migration plan. Initiator records on the

source storage and the zone between the source and destination zone are removed.

The following figure shows the verify phase workflow.

The verify phase tasks are listed in the following table.

Component	Tasks
Host	Customer performs application testing.
Fabric	1. Remove host to source storage zones from the zoneset. 2. Remove source storage to destination zones. 3. Backup the new zoneset.
Destination storage	1. Verify LUN alignment. 2. Collect FLI import logs from the Event log. 3. Remove manually created Snapshot copies used during the migration process. 4. Remove source array. 5. Document final configuration.
Source storage	1. Remove host storage initiator groups. 2. Remove destination storage initiator groups.

Discover phase data collection procedures

The discover phase collects customer environment information necessary for successful migration planning and execution.

Use Active IQ OneCollect in the data collection phase. For complete details, see the Active IQ OneCollect documentation.

Analyze phase IMT best practices

Analyze phase IMT best practices

The analyze phase focuses on items that must be addressed before proceeding with the migration activities. The host configuration information must be compared to supported configurations documented in the NetApp Interoperability Matrix (IMT).

The IMT is a web-based utility that enables searching for information about configurations for NetApp products that work with third-party products and components qualified by NetApp. The IMT contains both supported and certified NetApp configurations. Supported configurations are those qualified by NetApp. Certified configurations are those qualified by a third-party company to work with NetApp components.

IMT best practices

- Enter the NetApp IMT recommendations for required software and upgrades into the Switches and Hosts section of your planning worksheet.
- Start by entering static information, such as ONTAP OS, protocol, and CF mode, into the IMT. Then, using the site survey as a filter guide, enter host OS, volume manager, and HBA information.
- Do not be so specific as to have no results returned; it is better to view multiple returned results and choose the best fit.
- Host HBAs are sometimes reported on the OEM part number and will need to be cross-referenced before they are entered into the IMT.
- Check each host against the IMT for supportability.

Related information

NetApp Interoperability

FLI interoperability and support criteria

The FLI Interoperability Matrix (IMT) is a distinctly different interoperability tool designed to better support the NetApp-qualified source arrays that work with FLI.

Prior to performing a Foreign LUN Import, two areas of interoperability need to be checked:

- Check that FLI is supported. You can do this by reviewing the FLI IMT.
- Check that the complete end-to-end configuration, after the import(s) are complete, is a supported configuration. This is done from the FAS/AFF IMT.

In addition, check these three criteria against the target ONTAP version:

- The source storage platform model and microcode version.

- The SAN switch model and microcode version.

- The NetApp controller, the customer environment (switches, HBAs, firmware, server hardware, etc.) and SAN-attached clients that mount the LUNs after migration.

If any of these three components are not supported, some remediation might be necessary to ensure full success and support during and after the migration process.

Related information

[NetApp Interoperability](#)

Checking supported configurations for FLI using the IMT

You should use the Interoperability Matrix Tool (IMT) to find information about configurations for NetApp products that work with third-party products and components qualified by NetApp.

> ⓘ Beginning in ONTAP 9.9.1, if your array is not listed as supported in the IMT, you can use the SAN LUN Migrate App on the NetApp Support Site to determine if your array might be supported.

Steps

1. Go to the Interoperability Matrix Tool.

2. Search for the array model.

3. Select the solution **Foreign LUN Import (FLI) Back-end Interoperability**.

4. Select the **FAS model** and **ONTAP version** to determine the supported configurations.

5. For front-end supported host configurations, click **build end to end view with ONTAP SAN host**.

6. For switch-supported configurations, click **build end to end view for SAN-Switch** from the **ONTAP SAN host** tab.

Related information

[NetApp Interoperability](#)

Checking supported configurations for FLI using the SAN LUN Migrate App

Beginning in ONTAP 9.9.1, you can use the SAN LUN Migrate App to qualify a foreign source array for FLI. The SAN LUN Migrate App can be used when the desired foreign array is not listed in the FLI IMT.

Steps

1. Go to the NetApp Support Site.

2. Under **Filter by category**, select **Migration**.

3. Under **SAN LUN Migration**, click **Download App**.

4. Run the app from a FC or iSCSI Linux host that has block access to the source array.

If the foreign source array can be qualified, a green check mark is displayed. If the foreign source array cannot be qualified, a red X is displayed.

Enablement for non-supported LUNs

Enablement for non-supported LUNs

It is important to verify that the host OS, HBA, switch, and ONTAP array for your source array and your final configuration are all listed as supported in the Interoperability Matrix.

The following sections provide information for these use cases:

- Importing iSCSI LUNs as FC LUNs
- Moving migrated LUNs to AFF platforms

Related information

[NetApp Interoperability Matrix Tool](#)

Importing non-FC LUNs

Only FC LUNs are supported by FLI. However, there is a workaround that allows you to import iSCSI LUNs. Because you will be importing the iSCSI LUNs as FC LUNs, unlike other FLI online 7-Mode to ONTAP workflows, the disruption window would span this entire workflow:

Because you will be importing the iSCSI LUNs as FC LUNs, unlike other FLI online 7-Mode to ONTAP workflows, the disruption window would span this entire workflow.

Steps
1. On the source array, you will need to unmap the desired iSCSI LUN from its iSCSI igroup.
2. On the source array, map the LUN to a FC igroup, making sure that the destination array WWPNs have been added to the igroup.
3. Import the LUN.
4. After the LUN has been imported, you can create a new iSCSI igroup and add the hosts to the igroup.
5. On the hosts, rescan for LUNs.

   Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Related information

[NetApp Interoperability Matrix Tool](#)

Using Foreign LUN Import to import LUNs into AFF

Beginning with ONTAP 9.1, AFF supports FLI. You can use FLI to import LUNs from other arrays directly into ONTAP clusters.

Beginning with ONTAP 8.3.2, AFF can support FLI with an approved Process Variance Request (PVR). Contact your NetApp account team to get the PVR submitted for approval. On approval, the submitter, usually a NetApp System Engineer, will receive an approval letter with instruction for enabling FLI functionality.

For versions of ONTAP software previous to 8.3.2, you will need to stage FLI imports to a non-AFF HA pair on the same cluster with the AFF. After the migration has been completed, you can then use non-disruptive operations (NDO) such as vol or LUN move to move the migrated LUNs to AFF. If your AFF cluster doesn't have any non-AFF nodes, talk to your account team about the possibility of borrowing swing gear to facilitate this.

Gap analysis report
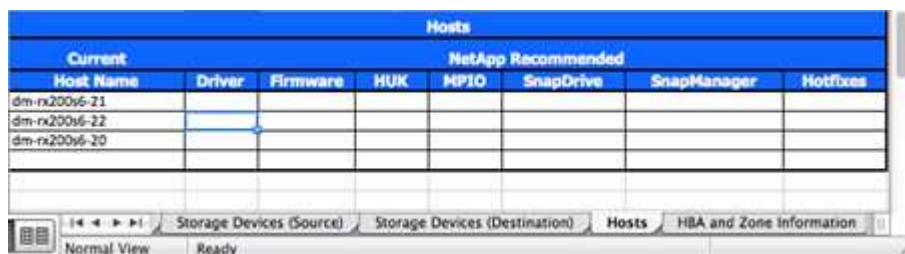
The gap analysis is a report of the customer's current and NetApp-recommended environment. It presents all recommended upgrades to the customer's environment that will need to take place post-migration.

The target configuration (post-migration) includes details for each host (OS configuration, MPIO, HBA details, Host Utility Kit version, and so on). Information about additional NetApp-required products, such as SnapDrive and SnapManager, is also available.

The changes required are typically not made until the actual migration event, because of the usual need to schedule maintenance windows. Typically, any changes made to the MPIO configuration before migration will affect the support of the current storage as well.

The completed NetApp Recommended section in the Hosts section of your Site Survey and Planning worksheet will serve as the gap analysis report. The gap analysis must be completed for every host included in the migration project. The completed gap analysis report must be reviewed with the customer.

The following is an example Gap analysis report.

Plan and prepare phase procedures

Plan and prepare phase procedures

The FLI planning phase focuses on the tasks required to create detailed migration plans and prepare the customer environment for the actual migration. One or more test migrations are performed during this phase to verify the installation and setup of the foreign LUN import.

The following are tasks to complete during the plan phase:

- Create a mapping of source and destination LUNs by entering the storage mapping information for each storage array in the LUN Details section of your Site Survey and Planning worksheet.
- Wire the source storage into the fabric based on the planning information.
- Configure the switch zones.
- Perform one or more test migrations to verify the installation and setup.

Wiring best practices for FLI migration

To configure ONTAP storage for FLI migration, you should wire the source storage into the fabric based on your planning information and recommended best practices.

The following wiring best practices are recommended when configuring ONTAP storage for FLI migration.

- Use dual fabrics for redundancy.
- Use at least two initiators and two target ports from each destination storage for FLI migration.
- Do not zone destination storage initiator ports with the host. Initiator ports of ONTAP are used to zone with target ports of source storage.

The following is an example of wiring for source storage and destination storage in a production fabric.
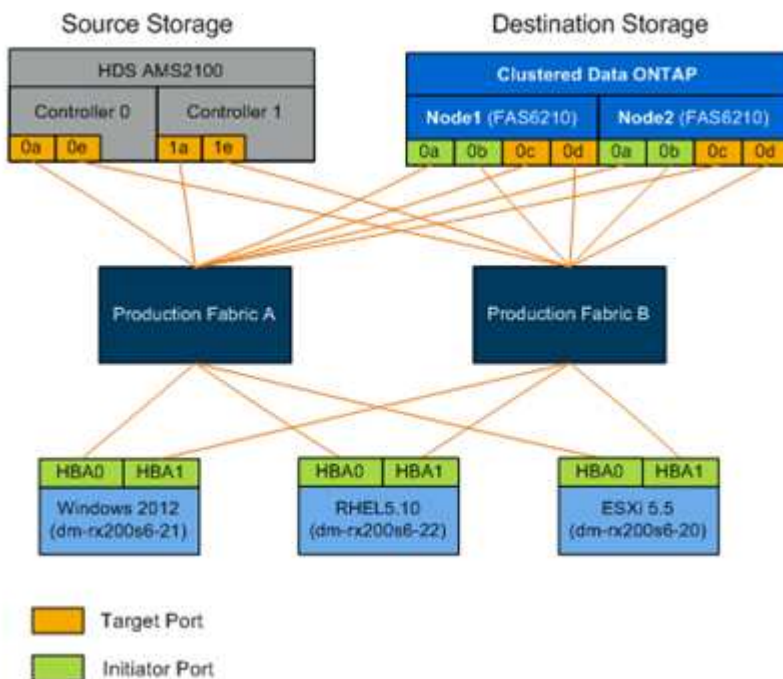
Configuring switch zones

You must create required zones on the SAN switches for connecting the source storage to the destination storage.

Steps

1. Back up the existing zonesets on each switch in the production and migration fabric.

2. Zone the source storage and destination storage as shown below.
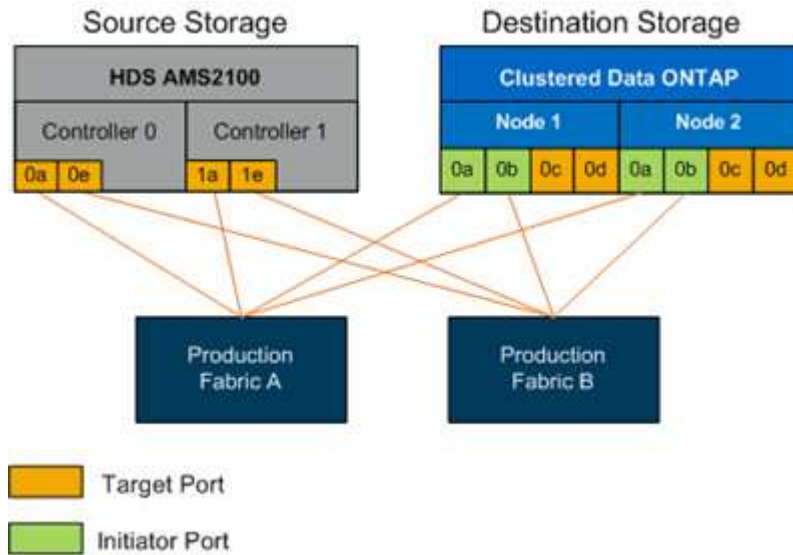
3. Create the zone and add it to the zoneset in production fabric A.

   The following is an example of the production zone, production fabric A for zone ZONE_AMS2100_cDOT_Initiator_fabA.

WWPN	Zone Members
50:06:0e:80:10:46:b9:60	AMS2100 Ctrl 0 Port 0a
50:06:0e:80:10:46:b9:68	AMS2100 Ctrl 1 Port 1a
50:0a:09:80:00:d3:51:59	ONTAP Node 1 Port 0a
50:0a:09:80:00:e7:81:04	ONTAP Node 2 Port 0a

4. Activate the zoneset in fabric A.

5. Create the zone and add it to the zoneset in production fabric B.

   The following is an example of the production zone, production fabric A for zone ZONE_AMS2100_cDOT_Initiator_fabB.

WWPN	Zone Members
50:06:0e:80:10:46:b9:64	AMS2100 Ctrl 0 Port 0e
50:06:0e:80:10:46:b9:6c	AMS2100 Ctrl 1 Port 1e
50:0a:09:80:00:d3:51:59	ONTAP Node 1 Port 0b
50:0a:09:80:00:e7:81:04	ONTAP Node 2 Port 0b

6. Activate the zoneset in production fabric B.

How to configure source arrays

Consult the array documentation for the source array in order to add a host entry for the initiator ports (LUN masking, igroup in NetApp parlance). This information can be retrieved from the Storage Groups section of your Site Survey and Planning worksheet.

Migration tests

You should perform one or more test migrations in order to verify that your arrays, switches, and hosts are properly configured and also in order to get several samples that can be extrapolated from to determine migration durations and levels of effort.

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.

To perform a test migration, complete the following steps:

Steps

1. Create a 2 GB test LUN on the source array.
2. Log in to Hitachi Storage Navigator Modular as system.
3. Select AMS 2100 array.

4. Click **Show and Configure Array**.

5. Log in using root.

6. Expand **Groups** and select **Logical Units**.

7. Select **Create LU** to create the test LUN.

8. Create a test LUN of 2 GB.

9. Click **OK**.

10. Skip the LUN assignment here and proceed by clicking **Close**.

11. Verify LUN 0026 is created.

12. Expand **Groups**and select **Logical Units**.

13. Select **Host Groups** to map the test LUN to the cDOT_FLI host group.

14. Select host group cDOT_FLI created in the previous step and click **Edit Host Group**.

15. 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.

16. Click **Logical Units** and add the test LUN LUN0026.

17. Click **OK** to map the LUN.

18. Select **Yes, I Have Read the Above Warning and Want to Edit Host Group** and click **Confirm**.

19. Verify host group creation and click **Close**.

20. Verify the test LUN and mapping from the source storage to destination storage and perform Foreign LUN Import (FLI) import.

21. Log in to the ONTAP storage through SSH using admin user.

22. Change the mode to Advanced. `DataMig-cmode::> set -privilege advanced`

23. Enter y when asked if you want to continue with advanced commands.

24. Discover the source array on ONTAP. Wait for a few minutes and retry to detect the source array.`storage`

```
array show
```

a. 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 fc4/9
shutdown
no shutdown
exit
exit
```

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

25. Verify the source array is discovered through all the initiator ports: `storage array config show -array-name HITACHI_DF600F_1`

```
               LUN   LUN
Node           Group Count    Array Name         Array Target Port
Initiator
-----------    ----- -----    ---------------------------
-----------------------
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
```

26. 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:
48495441:43484920:38333031:37353432:30303236:00000000:00000000:00000000:
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
------------------ ---------  -----  --------------------
------------------- --- ---  ---------------------  ------  -------
----------- ------------
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::*>
```

27. Mark the source LUN as foreign using the serial number: `storage disk set-foreign-lun {
    -serial-number 83017542001A } -is-foreign true`

28. Verify the source LUN is marked as foreign: `storage disk show -array-name HITACHI_DF600F_1`

29. List all foreign arrays and their serial numbers: `storage disk show -container-type foreign
    -fields serial-number`

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

30. Create a destination volume: `vol create -vserver datamig flivol aggr1 -size 10g`

31. Create a test LUN using a foreign LUN: `lun create -vserver datamig -path
    /vol/flivol/testlun1 -ostype linux -foreign-disk 83017542001A`

32. List the test LUN and verify the size of the LUN with the source LUN: `lun show`

> (i) 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.

33. Create test igroup of protocol FCP without adding any initiators: `lun igroup create -vserver
    datamig -igroup testig1 -protocol fcp -ostype linux`

34. Map the test LUN to the test igroup: `lun map -vserver datamig -path /vol/flivol/testlun1
    -igroup testig1`

35. Offline the test LUN: `lun offline -vserver datamig -path /vol/flivol/testlun1`

36. Create import relationship with test LUN and foreign LUN: `lun import create -vserver datamig
    -path /vol/flivol/testlun1 -foreign-disk 83017542001A`

37. Start the migration (import): `lun import start -vserver datamig -path
    /vol/flivol/testlun1`

38. Monitor the import progress: `lun import show -vserver datamig -path
    /vol/flivol/testlun1`

39. 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
```

40. 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
```

41. 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
```

42. 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`

43. Unmap the test LUN from the test igroup: `lun unmap -vserver datamig -path /vol/flivol/testlun1 -igroup testig1`

44. Online the test LUN: `lun online -vserver datamig -path /vol/flivol/testlun1`

45. 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.

46. Remove test LUN from source storage.

   a. Log in to Hitachi Storage Navigator Modular as a system.

   b. Select AMS 2100 array and click **Show and Configure Array**.

   c. Log in using root.

   d. Select **Groups**, then select **Host Groups**.

   e. Select *cDOT_FLI Igroup* and click **Edit Host Group**.

   f. In the **Edit Host Group** window, select all target ports chosen to map the test LUN and select **Forced Set to All Selected Ports**.

   g. Select the **Logical Units** tab.

   h. Select the test LUN from the **Assigned Logical Units** window.

   i. Select **Remove** to remove the LUN mapping.

   j. Click OK.

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

   l. Select Logical Units.

   m. Select the test LUN created in the previous step (LUN0026).

   n. Click **Delete LUN**.

   o. Click **Confirm** to delete the test LUN.

47. Delete the test LUN on the destination storage.

   a. Log in to the ONTAP storage through SSH using admin user.

   b. 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.

   c. 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`