



FLI post-migration

ONTAP FLI

NetApp
July 22, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap-fli/post-migration-performance/post-migration-overview.html> on July 22, 2026. Always check docs.netapp.com for the latest.

Table of Contents

FLI post-migration	1
Learn about ONTAP FLI post-migration	1
Perform ONTAP FLI post-migration clean up	1
Perform ONTAP FLI ESXi VMFSS ATS/CAW remediation	2
Perform ONTAP FLI ESXi post-migration host remediation	3
Post migration - Offline FLI	3
Post migration - Online FLI	5
Perform ONTAP FLI AIX post-migration host remediation	7
Post-Migration - Offline FLI	7
Post-Migration - Online FLI	9
Perform ONTAP FLI HP post-migration host remediation	10
Post-Migration - Offline FLI	10
Post-Migration - Online FLI	11
Perform ONTAP FLI Linux post-migration host remediation	12
Post-migration - Offline FLI	12
Post-migration - Online FLI	12
Perform ONTAP FLI Solaris post-migration host remediation	13
Post-Migration - Offline FLI	13
Post-Migration - Online FLI	14
Perform ONTAP FLI Windows post-migration host remediation	15
Post migration - Offline FLI	15
Post migration - Online FLI	16
ONTAP FLI performance controls	18
iSCSI FLI performance counters	18
ONTAP FLI factors that impact performance	19

FLI post-migration

Learn about ONTAP FLI post-migration

This section describes what to do after the the Foreign LUN Import (FLI) migration is complete. There are general clean up steps, performance validation steps, and steps for each host type (AIX, ESXi, HP, Linux, Solaris, and Windows).

The post-migration steps are important to ensure that the environment is left in a clean and supportable state, and to validate that the migration was successful from both a data integrity and performance perspective.

Perform ONTAP FLI post-migration clean up

After the LUN import is complete and the relationship is deleted, there are several important clean up steps to finalize the migration, ensure data integrity, and remove any temporary configurations that were put in place for the FLI workflow. These steps include confirming the import is complete, verifying data integrity, de-provisioning the foreign LUN from ONTAP discovery, cleaning up backend connectivity, and reviewing logs and documentation.

Following your FLI migration, you should perform the following steps to finalize the migration and ensure a clean state:

- Confirm the LUN import is successfully complete and the data integrity is verified
- Take a volume level Snapshot
- De-provision the LUN from the NetApp storage array initiator.

Depending upon different storage vendor, the steps will vary. Follow the steps recommended by the foreign storage array vendor

Steps

1. Mark the foreign LUN attribute as false (use the foreign disk serial):

```
storage disk modify -serial-number <FOREIGN_SERIAL> -is-foreign false
```

2. Ensure that the foreign storage array LUN is no longer visible in ONTAP:

```
storage disk show -array-name <ARRAY_NAME> -fields disk,serial-number,is-foreign,import-in-progress
```

3. Do the following for FC FLI:
 - a. Delete the zoning between the NetApp storage array initiator and foreign storage array target WWPNs.
 - b. Check on the foreign storage array that NetApp storage array initiators are no longer logged in.
 - c. Convert the FC adapter from initiator to target.

4. Do the following for iSCSI FLI:
 - a. Delete the iSCSI sessions between foreign storage array and NetApp storage array.
 - b. Delete the iSCSI initiator on the intercluster LIF.
 - c. Delete the intercluster LIF if it was created for the foreign LUN import and there is no other application consuming intercluster LIF.
5. Review the logs for errors to verify that your migration was successful.

Import logs are stored in the cluster event log file.

Example

```
DataMig-cmode::*> rows 0; event log show -nodes * -event fli*
7/7/2014 18:37:21    DataMig-cmode-01 INFORMATIONAL
fli.lun.verify.complete: Import verify of foreign LUN 83017542001E of
size 42949672960 bytes from array model DF600F belonging to vendor
HITACHI with NetApp LUN QvChd+EUXoiS is successfully completed.
~~~~~ Output truncated ~~~~~
```

The verification steps to compare the source and destination LUNs are covered in the execute migration phase. The importing and LUN verification steps are covered in the execute migration phase because they are linked to the import job and the foreign LUN.

6. Collect an AutoSupport to document the customer environment, and collect another AutoSupport after completion to record the final configuration:

```
autosupport invoke -node <node_name> -type all -message "migration-
final"
```

Perform ONTAP FLI ESXi VMFSS ATS/CAW remediation

During an online FLI migration, the temporary import relationship does not support VMware ATS or CAW (VAAI Atomic Test and Set or SCSI Compare And Write). If your datastore is VMFS5 and was created on storage that supports ATS, VMware might mark the VMFS header as “ATS-only.” In that state, the datastore can fail to mount when presented through an online FLI relationship.

About this task

- 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 and remount.
- No active I/O can be running against the datastore while changing ATS-only mode. VMware guidance is to evacuate and migrate or shut down VMs and unmount and remount as required to apply the change.

Steps

1. Check the current ATS mode of the datastore.

Run this on an ESXi host for the VMFS volume:

```
vmkfstools -Ph -v1 /vmfs/volumes/<DATASTORE_NAME_OR_UUID>
```

2. Disable ATS-only.



Ensure VMs are migrated off and shut down and the datastore is unmounted from all hosts where required by your change plan. Run on the device path of the first extent LUN (naa.*):

```
vmkfstools --configATSONly 0 /vmfs/devices/disks/naa.<naa_id>:1
```

This changes the volume to allow non-ATS locking (public).

3. Rescan storage adapters (or remount the datastore) so the hosts re-read the updated mode.
4. Re-enable ATS.

After the migration is completed and you want to restore ATS-only behavior:

```
vmkfstools --configATSONly 1 /vmfs/devices/disks/naa.<naa_id>:1
```

5. Rescan and remount again to apply the change.

Perform ONTAP FLI ESXi post-migration host remediation

After migration, ESXi hosts require remediation steps to ensure the newly imported LUNs are properly recognized and functional.

Post migration - Offline FLI

For Offline FLI, bring the LUN online after deleting the lun import and confirming host initiators are connected.

VMs/VMDK	RDM	WFSC	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 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 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 3: Identify datastore details and mount datastores back to ESXi.

VMs/VMDK	RDM	WFSC	VMSC
Step 4: Identify datastore details and mount datastores back to ESXi.		Step 4: Identify datastore details and mount datastores back to ESXi.	Step 4: Power on guest VMs and start applications.
		Step 5: identify the RDM disks used by the cluster and remap them to the WFSC cluster nodes using the same controller IDs and unit numbers (for example, SCSI1:0) to maintain consistency with the previous configuration.	Step 5: Power on guest VMs and start applications
		Step 6: Power on guest VMs and start the cluster application.	



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

Rescan after detaching to confirm clean state

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

Exit maintenance mode after completion

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

Rescan storage on all ESXi hosts mapped to the LUN

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

List devices and filesystems to confirm LUN / filesystem details

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

Identify new NAA IDs if they changed post-migration

```
esxcli storage core device list | grep -i naa
```

Mount VMFS datastore(s) back to ESXi

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

Verify datastore mount

```
esxcli storage filesystem list
```

Add RDM back to VM (done via vSphere UI)

CLI is typically not used; verify backend device visibility instead.

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

Exit maintenance mode (if applicable)

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

Power on guest VMs

(VM power operations are typically done via vCenter, but verification can be done on the ESXi host)

```
vim-cmd vmsvc/getallvms  
vim-cmd vmsvc/power.on <VMID>
```

Post migration - Online FLI

After deleting the import relationship, confirm the new LUN operates independently with CAW/UNMAP/XCOPY/WRITESAME enabled.

Revert VMFS locking mode to ATS-only by unmounting the datastore, setting ATS-only, then remounting. After deleting the LUN import, check if the LUN is set with -space-reserve disabled and -space-allocation enabled. If not, modify it accordingly.

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: Shut down guest VMs gracefully.		Step 2: Shut down guest VMs gracefully.
Step 3: Unmount datastores planned for migration; if shared, unmount from all hosts.		Step 3: Unmount datastores planned for migration; if shared, unmount from all hosts.
Step 4: Check VMFS locking mode (ATS-only vs ATS+SCSI). If ATS-SCSI, change to ATS-Only		Step 4: Check VMFS locking mode (ATS-only vs ATS+SCSI). If ATS-SCSI, change to ATS-Only
Step 5: Identify datastore details and mount datastores back to ESXi.		Step 5: Enable VM-HA / DRS and configure the required heartbeat disks if required.
Step 6: Power on guest VMs and start applications.		Step 6: Power on guest VMs and start applications

VMs/VMDK	RDM	VMSC
Step 7: Validate datastore locking mode and VM/application health		Step 7: Validate datastore locking mode and VM/application health

Rescan after detaching to confirm clean state

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

Exit maintenance mode after completion

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

Rescan storage on all ESXi hosts mapped to the LUN

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

List devices and filesystems to confirm LUN / filesystem details

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

Identify new NAA IDs if they changed post-migration

```
esxcli storage core device list | grep -i naa
```

Mount VMFS datastore(s) back to ESXi

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

Verify datastore mount

```
esxcli storage filesystem list
```

Add RDM back to VM (done via vSphere UI)

CLI is typically not used; verify backend device visibility instead.

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

Exit maintenance mode (if applicable)

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

Power on guest VMs

(VM power operations are typically done via vCenter, # but verification can be done on the ESXi host)

```
vim-cmd vmsvc/getallvms
vim-cmd vmsvc/power.on <VMID>
```

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

```
vmkfstools --configATSOOnly 1 /vmfs/volumes/<DATASTORE_NAME>
```

Perform ONTAP FLI AIX post-migration host remediation

After migration, AIX hosts require remediation steps to ensure the newly imported LUNs are properly recognized and functional.

Post-Migration - Offline FLI

For Offline FLI, bring the LUN online after deleting the lun import and confirming host initiators are connected.

Filesystem	Block device	LVM (FS)	Cluster	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: Rescan storage on all the mapped AIX 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: identify the multipath devices	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: Check for the filesystem details	Step 3: Ensure that the server boots up properly
Step 4: Mount the devices back to the respective folders		Step 4: Activate the VG	Step 4: verify the required cluster configuration and start the cluster application	Step 4: Restart the host to boot from newly mapped LUN that is part of LUN Import

Filesystem	Block device	LVM (FS)	Cluster	Sanboot
		Step 5: List down the LVs present on the VG	Step 5: Validate the cluster status and the application that is consuming.	
		Step 6: Mount the LVMs back to the respective folders		



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

Steps

1. Rescan after detaching to confirm clean state.

```
cfgmgr
```

2. Identify storage devices.

```
lsdev -Cc disk
sanlun lun show all
sanlun lun show all -pv
lsmppio -l hdiskx
```

3. Verify device is attached.

```
lspv |grep hdiskx
```

4. Identify the volume group and filesystems.

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

5. Varyonvg and mount filesystems.

```
varyonvg <vg_name>
mount /mount_point
df -g
```

Post-Migration - Online FLI

After deleting the import relationship, confirm the new LUN operates independently with UNMAP enabled. After deleting the LUN import, check if the LUN is set with `-space-reserve disabled` and `-space-allocation enabled`. If not, modify it accordingly.

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: Rescan storage on all the mapped AIX hosts
		SCSI UNMAP is automatically enabled if the device is thin provisioned.	SCSI UNMAP is automatically enabled if the device is thin provisioned.

Steps

1. Rescan after detaching to confirm clean state.

```
cfgmgr
```

2. Identify storage devices.

```
lsdev -Cc disk
sanlun lun show all
sanlun lun show all -pv
lsmpio -l hdiskx
```

3. Verify device is attached.

```
lspv |grep hdiskx
```

4. Verify UNMAP on the device:

```
chfs -a reclaim=normal /mount_point
lvmstat -v vg_name -r
```

	AIX/PowerVM
Step 1	Capture host information.
Step 2	Unmount the filesystem.
Step 3	Take a backup of the VG.
Step 4	Deactivate the VG.

	AIX/PowerVM
Step 5	Export the VG.

Perform ONTAP FLI HP post-migration host remediation

After migration, HP-UX hosts require remediation steps to ensure the newly imported LUNs are properly recognized and functional.

Post-Migration - Offline FLI

For offline FLI, bring the LUN online after deleting the lun import and confirming host initiators are connected.

Filesystem	Block device	LVM (FS)	Cluster	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: 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: identify the multipath devices	Step 2: From the UEFI shell, verify the device mapping table and boot into hpux from the new Storage device. Set the primary boot device to new Storage LUN
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 discovered properly	Step 3: Check for the filesystem details	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 VG/LVM	Step 4: verify the required cluster configuration and start the cluster application	
		Step 5: Activate the VG	Step 5: Validate the cluster status and the application that is consuming.	
		Step 6: List down the LVs present on the VG		
		Step 7: Mount the LVMs back to the respective folders		

Post-Migration - Online FLI

After deleting the LUN import, check if the LUN is set with `-space-reserve disabled` and `-space-allocation enabled`. If not, modify it accordingly.

For SCSI UNMAP feature to operate correctly, please consult the Vendor documentation to verify the supported software versions. Veritas VxVM + VxFS combination is required to enable this feature.

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: Rescan storage on all the mapped HP-UX hosts

Steps

1. Identify storage devices.

```
ioscan -fnC disk
insf -e
sanlun lun show all
sanlun lun show all -pv
sanlun lun show -wwpn
```

2. Import VG and Online the VG.

```
vgimport /dev/vg_name /dev/disk/disk_name
vgchange -a y /dev/vg_name
```

3. Identify the volume group and filesystems.

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

4. Mount the filesystem.

```
bdf
mount /dev/vg_name/vol_name /mount_point
```

	HP-UX
Step 1	Capture host information.
Step 2	Unmount the filesystem.
Step 3	Take a backup of the VG.

	HP-UX
Step 4	Deactivate the VG.
Step 5	Export the VG.

Perform ONTAP FLI Linux post-migration host remediation

After migration, Linux hosts require remediation steps to ensure the newly imported LUNs are properly recognized and functional.

Post-migration - Offline FLI

For Offline FLI, bring the LUN online after deleting the lun import and confirming host initiators are connected.

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
Step 3: Import the Volume Groups & 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 & add the respective resources for the available LV & filesystems. NOTE: Please note that the newly activated LVM objects have the LV & filesystem structures intact.	
	Step 5: Verify the PCS Cluster status for all the resources are online and available status.	

Post-migration - Online FLI

After deleting the LUN import, check if the LUN is set with `-space-reserve disabled` and `-space-allocation enabled`. If not, modify it accordingly.

Filesystem	Block device	LVM (FS)	Sanboot
Step 1: Rescan storage on all the mapped Linux hosts	Step 1: Rescan storage on all the mapped Linux hosts	Step 1: Rescan storage on all the mapped Linux hosts	Step 1: Rescan storage on all the mapped Linux hosts

Perform ONTAP FLI Solaris post-migration host remediation

After migration, Solaris hosts require remediation steps to ensure the newly imported LUNs are properly recognized and functional.

Post-Migration - Offline FLI

For offline FLI, bring the LUN online after deleting the lun import and confirming host initiators are connected.

Filesystem	Block device	LVM (FS)	Cluster	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: Rescan storage on all the mapped Solaris 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: identify the multipath devices	Step 2: On the SPARC server OBP, configure the new boot device that is presented to this Host HBA
Step 3: Check for the Pool details	Step 3: Identify the multipath device ID and present it to the respective application.	Step 3: check the pool status	Step 3: Check for the filesystem details	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 4: verify the required cluster configuration and start the cluster application	
		Step 5: List down the FS present on the Pool	Step 5: Validate the cluster status and the application that is consuming.	
		Step 6: Mount the FS back to the respective folders		



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

Steps

1. Rescan after attaching.

```
cfgadm -al
```

2. Identify storage devices.

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

3. Verify device is attached.

```
format |grep disk_id
```

4. Import the zpool and FS will automatically mounted.

```
zpool import zpool_name
df -g (to ensure its unmounted)
```

Post-Migration - Online FLI

After deleting the LUN import, check if the LUN is set with -space-reserve disabled and -space-allocation enabled. If not, modify it accordingly.

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: Rescan storage on all the mapped Solaris hosts.



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

Steps

1. Rescan after attaching.

```
cfgadm -al
```

2. Identify storage devices.

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

3. Verify device is attached.

```
`format |grep disk_id`
```

4. Import the zpool and FS will automatically mounted.

```
zpool import zpool_name  
df -g
```

	Solaris
Step 1	Capture host information.
Step 2	Unmount the filesystem.
Step 3	Take a backup of the VG.
Step 4	Deactivate the VG.
Step 5	Export the VG.

Perform ONTAP FLI Windows post-migration host remediation

After migration, Windows hosts require remediation steps to ensure the newly imported LUNs are properly recognized and functional.

Post migration - Offline FLI

For Offline FLI, bring the LUN online after deleting the lun import and confirming host initiators are connected.

Filesystem devices	Physical drives (no filesystem)	WFSC	Hyper V
Step 1: Rescan the storage on all the mapped hosts	Step 1: Rescan the storage on all the mapped hosts	Step 1: Rescan the storage on all the mapped hosts	Step 1: Rescan the disk management, All the Disks Should connect.
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 2: Identify the drives that are listed and online the disks.	Step 2: Select the disks and bring them online, all Should come ONLINE and should get original name and Disk letter.
Step 3: Verify if all the filesystem disks are available and working		Step 3: verify if the all the mapped hosts has the drives available and ONLINE	Step 3: Try to power-on Virtual Machines, Virtual Machine should come online without issue
Step 4: Start the required applications		Step 4: Start the WFSC cluster application and configure the required clustered resources.	

Post migration - Online FLI

After deleting the LUN import, check if the LUN is set with -space-reserve disabled and -space-allocation enabled. If not, modify it accordingly.

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: If the space-allocation is enabled on the target side, then a host reboot is required to reflect the changes.	Step 2: If the space-allocation is enabled on the target side, then a host reboot is required to reflect the changes.

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.

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 FLI performance controls

Foreign LUN Import (FLI) includes controls to reduce the risk of workload starvation and to help manage how much bandwidth the import consumes while production I/O is active.

Throttle the import (cap import throughput)

ONTAP provides lun import throttle to apply a maximum throughput limit to an import relationship. A value of 0 means the import runs at best effort along with user I/O; a non-zero value caps import bandwidth:

To set a maximum import throughput (example: 3MB/s), do the following:

```
set -privilege advanced
lun import throttle -vserver <SVM> -path <LUN_PATH> -max-throughput
-limit 3MB
```

Monitor detailed import progress and throughput

Use lun import show -instance (or -fields) to view detailed status, including current throughput, maximum throughput limit, and any QoS context shown for the import:

```
set -privilege advanced
lun import show -instance
```

Pause (operational throttling)

When online imports share backend resources with host I/O, it can be operationally useful to pause the import to let host I/O recover, then resume from the last checkpoint.

```
set -privilege advanced
lun import pause -vserver <SVM> -path <LUN_PATH>
```

Resume the import

```
set -privilege advanced
lun import resume -vserver <SVM> -path <LUN_PATH>
```

iSCSI FLI performance counters

Discover/confirm available objects (optional)

```
statistics catalog object show | grep -i iswi
```

Start Statistics Collection

```
statistics start -object iswi_session -sample_id import_in_progress
```

Collect iSCSI initiator session counters

```
statistics show -sample-id import_in_progress
```

ONTAP FLI factors that impact performance

The time to complete an import is influenced by multiple variables, including:

- Number of concurrent migrations between the source and destination
- Source array capability and current load
- Destination ONTAP capability and current load
- Amount of host I/O during migration (especially relevant for online imports)
- Frontend fabric and network characteristics, such as bandwidth, fan-in and fan-out, and overall topology
- Other non-FLI ONTAP applications consuming the intercluster LIF resources
- Concurrency, general FLI guidance recommends limiting to no more than 8 concurrent migrations per node for best performance.

Because these variables vary by environment, the most reliable way to estimate performance is to run multiple test migrations of different sizes, then use measured throughput to plan production durations.

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.