



FLI offline migration

ONTAP FLI

NetApp
January 07, 2026

This PDF was generated from https://docs.netapp.com/us-en/ontap-fli/san-migration/concept_fli_offline_workflow.html on January 07, 2026. Always check docs.netapp.com for the latest.

Table of Contents

- FLI offline migration 1
 - ONTAP FLI offline migration workflow summary 1
 - Prepare hosts for ONTAP FLI offline migration 2
 - Multipath verification for ESXi hosts 3
 - Prepare foreign storage array LUNs for an ONTAP FLI offline migration 12
 - Step 1: Present foreign array source LUNs to ONTAP storage 12
 - Step 2: Discover foreign array source LUNs in ONTAP 12
 - Create the LUN import relationship for an ONTAP FLI offline migration 14
 - Step 1: Identify the source array LUNs as foreign in ONTAP 14
 - Step 2: Create and configure a destination volume 16
 - Step 3: Create the destination LUNs and LUN import relationship 18
 - Import data from a foreign array using ONTAP FLI offline migration 22
 - Verify ONTAP FLI offline migration results 24
 - Remove the LUN import relationship after an ONTAP FLI offline migration 26
 - Perform ONTAP FLI offline post-migration tasks 30

FLI offline migration

ONTAP FLI offline migration workflow summary

A Foreign LUN Import (FLI) data migration is a process that involves several key steps to ensure a successful migration of data from third-party storage arrays to NetApp storage systems. FLI supports offline and online migrations. In an FLI offline migration, the client system is taken offline during data migration from the third party foreign storage array to the NetApp storage system.

Before you begin

- You should complete the [discovery](#), [analysis](#), and [planning](#) phases of the migration process.
- You should [configure your FC adapters for initiator mode](#).
- You should [zone your foreign array target ports with the ONTAP storage initiator ports](#).

The FLI offline workflow includes preparing the hosts and foreign LUNs for import, creating the LUN import relationship, and importing the data.

1

Prepare your host.

Before you perform an FLI offline migration, you should reboot your hosts and verify your host multipathing configuration.

2

Prepare your foreign LUNs.

In the FLI offline migration process, you must take manual steps from your foreign array to present the foreign source LUN to the ONTAP storage system; then you must take additional manual steps to discover the foreign source LUN on the ONTAP storage system.

3

Create the LUN import relationship.

Creating the LUN import relationship for FLI offline migrations includes identifying the source array LUNs as foreign in ONTAP, creating and configuring the destination volume to contain the foreign LUNs, creating destination target LUNs, and finally establishing the import relationship.

4

Import LUN data from the foreign array.

Use FLI to import LUN data from the foreign array.

5

Verify the migration results.

Perform a block-by-block comparison of the source and destination LUNs to verify that the migration is complete and accurate.

6

Remove the LUN import relationship.

After the FLI offline migration is completed, the LUN import relationship can be safely removed.

7

Perform post-migration tasks.

Review logs for errors, verify your host multipathing configuration, and perform application testing to verify that your migration completed successfully.

Prepare hosts for ONTAP FLI offline migration

Before you begin a Foreign LUN Import (FLI) offline migration, you should perform any steps identified in the analysis phase as necessary for host remediation such as installing host attach kits or DSMs. You should also reboot your hosts and verify that host multipathing is configured correctly.

Steps

1. Perform any necessary host remediation steps identified in the [analysis phase](#).
2. Shut down all your open applications.
3. Reboot the host.
4. Review the logs for errors.
5. Verify your host multipathing configuration.
 - For Windows hosts: See [Using Windows Server 2022 with ONTAP](#) for steps to verify your multipath configuration.
 - For Linux hosts: Run the `multipath-ll` command and review the output. All paths should display as active and ready.

Sample output of multipath-ll command

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

Multipath verification for ESXi hosts


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

Steps

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



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



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

Mount Point	Volume Name	Size
UUID	Mounted	Type
Free		
-----	-----	-----
-----	-----	-----

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



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



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

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

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

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

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

```

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

```

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

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



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

## 8. Identify LUN naa to RDM device mapping.



```

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

```

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

```

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

```

10. Record the UUID of the datastore.

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

```

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

```

## 12. Identify the virtual machine hard disks.

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

```

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

```

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

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


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

```

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

```

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



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

```

# esxcli storage nmp device list
naa.60060e801046b96004f2bf4600000014

```

```

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

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

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

```

```

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

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

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

```

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

#### What's next?

Prepare the foreign storage array LUNs for FLI offline migration.

## Prepare foreign storage array LUNs for an ONTAP FLI offline migration

In a Foreign LUN Import (FLI) offline migration, you must take manual steps from your foreign storage array to present the foreign source LUN to your ONTAP storage system; then you must take steps from your ONTAP storage system to discover the foreign LUN.

### Step 1: Present foreign array source LUNs to ONTAP storage

Before you can begin importing data from a foreign storage array LUN using an FLI offline migration, you must present the source LUNs on your foreign storage array to your ONTAP storage system.

#### Steps

1. Log in to the source array.
2. Add the NetApp initiators to the host group created during the plan phase.
3. Select the host LUNs that need to be migrated from available logical LUNs.

Use LUN names for each host mentioned in the source LUNs section of your [Site Survey and Planning worksheet](#).

### Step 2: Discover foreign array source LUNs in ONTAP

After you have presented your foreign array source LUNs to your ONTAP storage system, the LUNs must be discovered in ONTAP before you can create the LUN import relationship.

#### Steps

1. Verify the source LUNs and mapping from the source storage to the destination storage.
2. Log in to the ONTAP storage system through SSH using admin user.
3. Change the mode to advanced.

```
set -privilege advanced
```

4. Enter `y` when asked if you want to continue.
5. Discover the source array on ONTAP. Wait for a few minutes and retry to detect the source array.

```
storage array show
```

The following example shows the discovery of a Hitachi DF600F array.

```
DataMig-ontap::*> storage array show
```

Prefix	Name	Vendor	Model	Options
HIT-1	HITACHI_DF600F_1	HITACHI	DF600F	



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

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

```
storage array config show -array-name <array_name>
```

The following example shows the Hitachi DF600F array discovered through all the initiator ports.

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

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

## What's next?

Create the LUN import relationship.

# Create the LUN import relationship for an ONTAP FLI offline migration

Before you can migrate a LUN from a foreign array to an ONTAP storage system, you must create a LUN import relationship. A LUN import relationship is a persistent pairing between the source and destination storage for the purpose of data import. The source and destination endpoints are LUNs.

Creating the LUN import relationship for Foreign LUN Import (FLI) offline migrations includes identifying the source array LUNs as foreign in ONTAP, creating and configuring the destination volume to contain the foreign LUNs, creating destination target LUNs, and finally establishing the import relationship.

## Before you begin

You should have completed the steps to [prepare your foreign LUNs for FLI offline migration](#).

## Step 1: Identify the source array LUNs as foreign in ONTAP

You'll need to identify the source array LUNs as foreign LUNs in ONTAP before you begin your FLI offline migration.

### Steps

1. List the source LUNs mapped from the foreign array; then verify the disk properties and paths.

```
storage disk show -array-name <array_name> -fields disk, serial-number,
container-type, owner, path-lun-in-use-count, import-in-progress, is-
foreign
```

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

The following example shows the source LUNs from the Hitachi DF600F array.



```
DataMig-ontap:*> storage disk show -array-name HITACHI_DF600F_1 -fields
disk, serial-number, container-type, owner, path-lun-in-use-count,
import-in-progress, is-foreign
```

```
disk      owner is-foreign container-type import-in-progress path-lun-in-
use-count serial-number
-----
```

```
-----
HIT-1.2   -      false      unassigned      false      0,0,0,0,0,0,0,0
83017542001E
HIT-1.3   -      false      unassigned      false      0,0,0,0,0,0,0,0
83017542000E
HIT-1.14  -      false      unassigned      false      0,0,0,0,0,0,0,0
830175420019
3 entries were displayed.
```

## 2. Use the serial number to mark the source LUN as foreign in ONTAP:

```
storage disk set-foreign-lun -serial-number <lun_serial_number> -is
-foreign true
```

The following example marks the source LUNs from the Hitachi DF600F array as foreign.

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

## 3. Verify the source LUN is marked as foreign.

```
storage disk show -array-name <array_name> -fields disk, serial-number,
container-type, owner,import-in-progress, is-foreign
```

The following example shows the source LUNs from the Hitachi DF600F array marked as foreign.

```
DataMig-ontap::*> storage disk show -array-name HITACHI_DF600F_1 -fields
disk, serial-number, container-type, owner,import-in-progress, is-
foreign
```

```
disk      owner is-foreign container-type import-in-progress serial-
number
-----
-----
HIT-1.2   -      true      foreign      false      83017542001E
HIT-1.3   -      true      foreign      false      83017542000E
HIT-1.4   -      true      foreign      false      83017542000F
3 entries were displayed.
```

## Step 2: Create and configure a destination volume

Before you create the LUN import relationship for an FLI offline migration, you must create a volume on your ONTAP storage system to contain the LUNs you will import from your foreign array.

### About this task

Beginning with ONTAP 9.17.1, data migration of foreign LUNs using FLI offline migration is supported with ASA r2 systems. ASA r2 systems vary from other ONTAP systems (ASA, AFF, and FAS) in the implementation of its storage layer. In ASA r2 systems, volumes are automatically created when a storage unit (LUN or namespace) is created. Therefore, you do not need to create a volume before creating the LUN import relationship. You can skip this step if you are using an ASA r2 system.

Learn more about [ASA r2 systems](#).

### Steps

1. Create a destination volume.

```
volume create -vserver <SVM_name> -volume <volume_name> -aggregate
<aggregate> -size <volume_size> -snapshot-policy default
```

The following example creates a volume named `winvol` on the `aggr1` aggregate with a size of 100 GB.

```
DataMig-ontap::*> vol create -vserver datamig winvol aggr1 -size 100g
```

2. Disable the default Snapshot policy on each volume.

```
volume modify -vserver <SVM_name> -volume <volume_name> -snapshot-policy
none
```

If default Snapshot copies exist prior to FLI migration, the volume needs additional space to store changed data.

The following example disables the default Snapshot policy on the `winvol` volume.

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

```
Warning: You are changing the Snapshot policy on volume winvol to none.  
Any Snapshot copies on this volume from the previous policy will not be  
deleted by
```

```
    this new Snapshot policy.
```

```
Do you want to continue? {y|n}: y
```

```
Volume modify successful on volume winvol of Vserver datamig.
```

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

```
vol modify -vserver <SVM_name> -volume * -fractional-reserve 0  
-snapshot-policy none
```

The following example sets the `fractional-reserve` option to 0 and the Snapshot policy to `none` for all the volumes in the `datamig` SVM.

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

4. Verify your volume settings.

```
volume show -vserver <SVM_name> -volume * -fields fractional-  
reserve,snapshot-policy
```

The `fractional-reserve` and `snapshot-policy` settings should be 0 and `none`, respectively.

5. Delete any existing Snapshot copies.

```
set advanced; snap delete -vserver <SVM_name> -volume <volume_name>  
-snapshot * -force true
```



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

### Step 3: Create the destination LUNs and LUN import relationship

For FLI offline migration, the destination LUNs on your ONTAP storage system, must be created and mapped to an igroup; then they must be offlined before creating the LUN import relationship.

#### About this task

Beginning with ONTAP 9.17.1, data migration of foreign LUNs using FLI offline migration is supported with [ASA r2 systems](#). ASA r2 systems vary from other ONTAP systems (ASA, AFF, and FAS) in the implementation of its storage layer. In ASA r2 systems, volumes are automatically created when a storage unit (LUN or namespace) is created. Each volume contains only one storage unit. Therefore, for ASA r2 systems, you do not need to include the volume name in the `-path` option when creating the LUN; you should include the storage unit path instead.

#### Steps

1. Create destination LUNs.

```
lun create -vserver <SVM_name> -path <volume_path|storage_unit_path>
-ostype <os_type> -foreign-disk <serial_number>
```

The following example creates LUNs on the `datamig` SVM with the specified paths and foreign disk serial numbers. The `-ostype` option specifies the operating system type of the LUN.

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

Created a LUN of size 40g (42949672960)

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

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

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



The `lun create` command detects the LUN size and alignment based on partition offset and creates the LUN accordingly with foreign-disk option. Some I/O will always appear be partial writes and will therefore look misaligned. Examples of this would be database logs.

2. Verify the size and source LUN of the newly created LUNs.

```
lun show -vserver <SVM_name> -fields vserver, path, state, mapped, type,
size
```

The following example shows the LUNs created in the `datamig` SVM with their paths, states, mapped status, types, and sizes.

```
DataMig-ontap::*> lun show -vserver datamig
```

Vserver	Path	State	Mapped	Type
Size				
-----	-----	-----	-----	-----
datamig	/vol/esxvol/bootlun	online	unmapped	vmware
20GB				
datamig	/vol/esxvol/linuxrdmvlun	online	unmapped	linux
2GB				
datamig	/vol/esxvol/solrdmplun	online	unmapped	solaris
2GB				
datamig	/vol/winvol/gdrive	online	unmapped	windows_2008
3GB				

4 entries were displayed.

3. If you are running ONTAP 9.15.1 or later, disable space allocation for the newly created LUNs.

Space allocation is enabled by default for newly created LUNs in ONTAP 9.15.1 and later.

```
lun modify -vserver <vserver_name> -volume <volume_name> -lun <lun_name>  
-space-allocation disabled
```

4. Verify that space allocation is disabled.

```
lun show -vserver <vserver_name> -volume <volume_name> -lun <lun_name>  
-fields space-allocation
```

5. Create a host igroup of protocol FCP and add host initiators.

```
lun igroup create -ostype <os_type> -protocol fcp -vserver <SVM_name>  
-igroup <igroup_name> -initiator <initiator_wwpn1>,<initiator_wwpn2>
```

Find initiator WWPNS from storage groups section of your Site Survey planning worksheet.

The following example creates igroups for the destination LUNs with the specified operating system types and initiators.

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

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

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



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

## 6. Map the destination LUNs to an igroup.

```
lun map -vserver <SVM_name> -path <volume_path|storage_unit_path>
-igroup <igroup_name> -lun-id <lun_id>
```

The following example maps the destination LUNs to their respective igroups with the specified paths and LUN IDs.

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

## 7. Offline the destination LUNs.

```
lun offline -vserver <SVM_name> -path <volume_path|storage_unit_path>
```

The following example offlines the destination LUNs in the datamig SVM.

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

8. Create the LUN import relationship between the destination and source LUNs.

```
lun import create -vserver <SVM_name> -path  
<volume_path|storage_unit_path> -foreign-disk <serial_number>
```

The following example creates the LUN import relationship for the destination LUNs in the `datamig` SVM with their respective paths and foreign disk serial numbers.

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

9. Verify the LUN import relationship is created.

```
lun import show -vserver <SVM_name> -fields vserver, foreign-disk, path,  
operation, admin-state, operational-state, percent-complete
```

The following example shows the LUN import relationship created for the destination LUNs in the `datamig` SVM with their respective foreign disks and paths.

```

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

```

### What's next?

[Import the data from the foreign LUNs to the ONTAP LUNs.](#)

### Related information

- [Learn more about unaligned I/O.](#)
- [Learn more about enabling space allocation for SAN protocols.](#)

## Import data from a foreign array using ONTAP FLI offline migration

After you have created the LUN import relationship between the source and destination LUNs for an FLI offline migration, you can import the data from the foreign array to the ONTAP storage system.

Beginning with ONTAP 9.17.1, data migration of foreign LUNs using FLI offline migration is supported with ASA r2 systems. ASA r2 systems vary from other ONTAP systems (ASA, AFF, and FAS) in the implementation of its storage layer. In ASA r2 systems, volumes are automatically created when a storage unit (LUN or namespace) is created. Each volume contains only one storage unit. Therefore, for ASA r2 systems, you do not need to include the volume name in the `-path` option when creating the LUN; you should include the storage unit path instead.

Learn more about [ASA r2 systems](#).

### Steps

1. Start the data import from the foreign LUNs to the ONTAP LUNs.



```
lun import start -vserver <SVM_name> -path  
<volume_path|storage_unit_path>
```

This example shows the command to start the data import for LUNs named **bootlun**, **fdrive** and **gdrive** in the **winvol** volume and the **datamig** SVM.

```
DataMig-ontap::*> lun import start -vserver datamig -path  
/vol/winvol/bootlun  
  
DataMig-ontap::*> lun import start -vserver datamig -path  
/vol/winvol/fdrive  
  
DataMig-ontap::*> lun import start -vserver datamig -path  
/vol/winvol/gdrive
```

## 2. Monitor the import progress.

```
lun import show -vserver <SVM_name> -fields vserver, foreign-disk, path,  
admin-state, operational-state, percent-complete, imported-blocks,  
total-blocks, estimated-remaining-duration
```

You can compare the progress you are seeing here with the migration performance estimates that you developed after performing your test migrations.

This example shows the command to monitor the import progress for the **datamig** SVM.

```
DataMig-ontap::*> lun import show -vserver datamig -fields vserver,
foreign-disk, path, admin-state, operational-state, percent-complete,
imported-blocks, total-blocks, , estimated-remaining-duration
```

vserver	foreign-disk	path	admin-state	operational-state	percent-complete	imported-blocks	total-blocks	estimated-remaining-duration
datamig 83017542000E	/vol/winvol/fdrive		started	completed	100	4194304	4194304	-
datamig 83017542000F	/vol/winvol/gdrive		started	completed	100	6291456	6291456	-
datamig 830175420010	/vol/linuxvol/bootlun		started	in_progress	35107077	41943040	00:00:48	83

3 entries were displayed.

### 3. Verify that the data import is completed successfully.

```
lun import show -vserver <SVM_name> -fields vserver, foreign-disk, path,
admin-state, operational-state, percent-complete, imported-blocks,
total-blocks, , estimated-remaining-duration
```

This example shows the command to verify the import status for the **datamig** SVM.

```
DataMig-ontap::*> lun import show -vserver datamig -fields vserver,
foreign-disk, path, admin-state, operational-state, percent-complete,
imported-blocks, total-blocks, , estimated-remaining-duration
```

The **operational status** shows as **completed** when the import job is completed successfully.

### What's next?

[Verify the migration results.](#)

## Verify ONTAP FLI offline migration results

After a LUN is migrated from the foreign array to your ONTAP storage system, FLI can perform a block-by-block comparison of the source and destination LUNs to verify that the migration is complete and accurate. A migration verification takes approximately the same amount of time as (or slightly more than) the migration.

A migration verification is not required, but is highly recommended.

### About this task

Beginning with ONTAP 9.17.1, data migration of foreign LUNs using FLI offline migration is supported with [ASA r2 systems](#). ASA r2 systems vary from other ONTAP systems (ASA, AFF, and FAS) in the implementation of its storage layer. In ASA r2 systems, volumes are automatically created when a storage unit (LUN or namespace) is created. Each volume contains only one storage unit. Therefore, for ASA r2 systems, you do not need to include the volume name in the `-path` option when creating the LUN; you should include the storage unit path instead.

### Steps

1. Start the LUN migration verification.

```
lun import verify start -vserver <SVM_name> -path  
<volume_path|storage_unit_path>
```

This example shows the command to start the LUN migration verification for LUNs named **bootlun**, **fdrive** and **gdrive** in the **winvol** volume and the **datamig** SVM.

```
DataMig-ontap::*> lun import verify start -vserver datamig -path  
/vol/winvol/bootlun
```

```
DataMig-ontap::*> lun import verify start -vserver datamig -path  
/vol/winvol/fdrive
```

```
DataMig-ontap::*> lun import verify start -vserver datamig -path  
/vol/winvol/gdrive
```

2. Monitor the verification status.

```
lun import show -vserver <SVM_name> -fields vserver, foreign-disk, path,  
admin-state, operational-state, percent-complete, imported-blocks,  
total-blocks, estimated-remaining-duration
```

This example shows the command to monitor the verification status for the **datamig** SVM.

```
DataMig-ontap::*> lun import show -vserver datamig -fields vservers,
foreign-disk, path, admin-state, operational-state, percent-complete,
imported-blocks, total-blocks, , estimated-remaining-duration
```

vservers	foreign-disk	path	admin-state	operational-state	percent-complete	imported-blocks	total-blocks	estimated-remaining-duration
datamig 83017542000E	/vol/winvol/fdrive	started	in_progress	57	-	4194304	00:01:19	
datamig 83017542000F	/vol/winvol/gdrive	started	in_progress	40	-	6291456	00:02:44	
datamig 830175420010	/vol/linuxvol/bootlun	started	in_progress	8	-	41943040	00:20:29	

3 entries were displayed.

You can execute the same command to track the progress of the verification. The **operational-state** displays a status of **completed** when the verification job is completed successfully.

### 3. Stop the LUN verification.

```
lun import verify stop -vserver <SVM_name> -path
<volume_path|storage_unit_path>
```

This example shows the command to stop the LUN verification for the **datamig** SVM.

```
DataMig-ontap::*> lun import verify stop -vserver datamig -path
/vol/esxvol/winrdmplug
```

The LUN import verification must be explicitly stopped before bringing the LUN back online. Otherwise, the `lun online` command fails. This step must be performed manually even if the status shows that the verification is completed.

#### What's next?

[Remove the LUN import relationship.](#)

## Remove the LUN import relationship after an ONTAP FLI offline migration

After the Foreign LUN Import (FLI) offline migration is completed, the LUN import

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

Beginning with ONTAP 9.17.1, data migration of foreign LUNs using FLI offline migration is supported with [ASA r2 systems](#). ASA r2 systems vary from other ONTAP systems (ASA, AFF, and FAS) in the implementation of its storage layer. In ASA r2 systems, volumes are automatically created when a storage unit (LUN or namespace) is created. Each volume contains only one storage unit. Therefore, for ASA r2 systems, you do not need to include the volume name in the `-path` option when creating the LUN; you should include the storage unit path instead.

## Steps

1. Delete the import relationship to remove the data import jobs.

```
lun import delete -vserver <SVM_name> -path  
<volume_path|storage_unit_path>
```

This example shows the command to delete the import relationship for LUNs named **bootlun**, **fdrive** and **gdrive** in the **winvol** volume and the **datamig** SVM.

```
DataMig-ontap::*> lun import delete -vserver datamig -path  
/vol/winvol/bootlun  
  
DataMig-ontap::*> lun import delete -vserver datamig -path  
/vol/winvol/fdrive  
  
DataMig-ontap::*> lun import delete -vserver datamig -path  
/vol/winvol/gdrive
```

2. Verify that the import jobs are deleted.

```
lun import show -vserver <SVM_name>
```

This example shows the command to verify that the import jobs are deleted for the **datamig** SVM.

```
DataMig-ontap::*> lun import show -vserver datamig  
There are no entries matching your query.
```

3. Mark the foreign LUN attribute to false.

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

This example shows the command to mark the foreign LUN attribute to `false` for the LUNs named

**bootlun**, **fdrive** and **gdrive** in the **winvol** volume and the **datamig** SVM.

```
DataMig-ontap::*> storage disk modify { -serial-number 83017542001E }
-is-foreign false

DataMig-ontap::*> storage disk modify { -serial-number 83017542000E }
-is-foreign false

DataMig-ontap::*> storage disk modify { -serial-number 83017542000F }
-is-foreign false
```

4. Verify the foreign LUNs are marked as false.

```
storage disk show -array-name <array_name> -fields disk, serial-number,
container-type, owner, import-in-progress, is-foreign
```

This example shows the command to verify the foreign LUNs are marked as false on the **HITACHI_DF600F_1** array.

```
DataMig-ontap::*> storage disk show -array-name HITACHI_DF600F_1 -fields
disk, serial-number, container-type, owner,import-in-progress, is-
foreign
```

disk	owner	is-foreign	container-type	import-in-progress	serial-number
HIT-1.2	-	false	unassigned	false	83017542001E
HIT-1.3	-	false	unassigned	false	83017542000E
HIT-1.4	-	false	unassigned	false	83017542000F

3 entries were displayed.

5. Bring the destination LUNs online.

```
lun online -vserver <SVM_name> -path <volume_path|storage_unit_path>
```

This example shows the command to bring the destination LUNs online for LUNs named **bootlun**, **fdrive** and **gdrive** in the **winvol** volume and the **datamig** SVM.

```
DataMig-ontap::*> lun online -vserver datamig -path /vol/winvol/bootlun

DataMig-ontap::*> lun online -vserver datamig -path /vol/winvol/fdrive

DataMig-ontap::*> lun online -vserver datamig -path /vol/winvol/gdrive
```

6. Verify the LUNs are online.

```
lun show -vserver <SVM_name>
```

This example shows the command to verify the LUNs are online for the **datamig** SVM.

```
DataMig-ontap::*> lun show -vserver datamig
```

Vserver	Path	State	Mapped	Type
Size				
-----	-----	-----	-----	-----
datamig	/vol/esxvol/bootlun	online	mapped	vmware
20GB				
datamig	/vol/esxvol/linuxrdmvlun	online	mapped	linux
2GB				
datamig	/vol/esxvol/solrdmplun	online	mapped	solaris
2GB				
3 entries were displayed.				

7. Optionally, view the event log to verify the migration results.

```
event log show -event fli*
```

This example shows sample output of the command to view the event log for FLI migration results.

```
DataMig-ontap::*> event log show -event fli*
```

```
7/7/2014 18:37:21 DataMig-ontap-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.
7/7/2014 18:37:15 DataMig-ontap-01 INFORMATIONAL
fli.lun.verify.complete: Import verify of foreign LUN 830175420015 of
size 42949672960 bytes from array model DF600F belonging to vendor
HITACHI with NetApp LUN QvChd+EUXoiX is successfully completed.
7/7/2014 18:02:21 DataMig-ontap-01 INFORMATIONAL
fli.lun.import.complete: Import of foreign LUN 83017542000F of size
3221225472 bytes from array model DF600F belonging to vendor HITACHI is
successfully completed. Destination NetApp LUN is QvChd+EUXoiU.
```

### What's next?

[Perform post-migration tasks for an FLI offline migration.](#)

## Perform ONTAP FLI offline post-migration tasks

Any outstanding server remediation not performed earlier is performed during post-migration.

The third-party software is removed, NetApp software is installed and configured, and then the host is brought up accessing the LUNs on NetApp. See the topic *Host remediation* for examples of post-migration remediation for specific host types.

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



## 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.