NetApp

Transitioning Solaris host data LUNs with ZFS file systems

ONTAP 7-Mode Transition

NetApp
October 09, 2025

Table of Contents

Transitioning Solaris host data LUNs with ZFS file systems

If you transition a Solaris host data LUN with ZFS file system from Data ONTAP operating in 7-Mode to clustered Data ONTAP using the 7-Mode Transition Tool (7MTT), you must perform specific steps before and after transition to remediate transition issues on the host.

Preparing to transition Solaris host data LUNs with ZFS file system

Before you transition Solaris host LUNs with ZFS file systems from Data ONTAP operating in 7-Mode to clustered Data ONTAP, you must gather information you need for the transition process.

This applies to copy-based transitions and copy-free transitions.

Steps

1. On the 7-Mode controller, identify the name of the LUN to be transitioned:

   lun show

   ```
     fas8040-shu01> lun show
           /vol/ufs/ufs1                    5g (5368709120)    (r/w, online,
   mapped)
           /vol/ufs/ufs2                    5g (5368709120)    (r/w, online,
   mapped)
           /vol/zfs/zfs1                    6g (6442450944)    (r/w, online,
   mapped)
           /vol/zfs/zfs2                    6g (6442450944)    (r/w, online,
   mapped)
   ```

2. On the host, locate the SCSI device file name for the LUN:

   sanlun lun show

   The SCSI device file name is located in the `device filename` column.

```
# sanlun lun show
controller(7mode)/                              device
host                    lun
vserver(Cmode)          lun-pathname    filename
adapter     protocol    size    mode
------------------------------------------------------------------------
---------------------------------------------------
fas8040-shu01           /vol/zfs/zfs2
/dev/rdsk/c0t60A98000383035356C2447384D396550d0s2 scsi_vhci0 FCP
6g          7
fas8040-shu01           /vol/zfs/zfs1
/dev/rdsk/c0t60A98000383035356C2447384D39654Ed0s2 scsi_vhci0 FCP
6g          7
fas8040-shu01           /vol/ufs/ufs2
/dev/rdsk/c0t60A98000383035356C2447384D39654Ad0s2 scsi_vhci0 FCP
5g          7
fas8040-shu01           /vol/ufs/ufs1
/dev/rdsk/c0t60A98000383035356C2447384D396548d0s2 scsi_vhci0 FCP
5g          7
```

3. List the zpool:

   zpool list

4. Record the zpool and get the disks associated with the zpool:

   zpool status *pool-name*

```
# zpool list
NAME     SIZE  ALLOC   FREE  CAP  HEALTH  ALTROOT
n_pool  11.9G  2.67G  9.27G  22%  ONLINE  -

# zpool status
  pool: n_pool
 state: ONLINE
 scan: none requested
config:

        NAME                                       STATE    READ WRITE
CKSUM
        n_pool                                     ONLINE      0     0
0
            c0t60A98000383035356C2447384D396550d0  ONLINE      0     0
0
            c0t60A98000383035356C2447384D39654Ed0  ONLINE      0     0
0

errors: No known data errors
```

5. List and record the ZFS datasets within a ZFS storage pool:

zfs list

```
# zfs list
NAME           USED  AVAIL  REFER  MOUNTPOINT
n_pool        2.67G  9.08G   160K  /n_pool
n_pool/pool1  1.50G  2.50G  1.50G  /n_pool/pool1
n_pool/pool2  1.16G  2.84G  1.16G  /n_pool/pool2
```

Testing data LUNs on Solaris hosts with ZFS file system before the cutover phase of copy-based transitions

If you are using the 7-Mode Transition Tool (7MTT) 2.2 or later and Data ONTAP 8.3.2 or later to transition your Solaris host ZFS data LUNs, you can test your transitioned clustered Data ONTAP LUNs to verify that you can mount your MPIO device before the cutover phase.

- Your source host with ZFS data LUNs need to be offline before starting the test phase transition.

  See *Oracle Doc ID 1316472.1: LUN Copy Is Not Supported While ZFS Zpool Is Online* for details.

- Your new clustered Data ONTAP LUNs must be mapped to your test host and your LUNs must be ready for

transition.

- Exporting the zpool on the production host causes application disruption; all I/O operations should be stopped prior to 7-Mode LUN.

You should maintain hardware parity between the test host and the source host, and you should perform the following steps on the test host.

Your clustered Data ONTAP LUNs are in read/write mode during testing. They convert to read-only mode when testing is complete and you are preparing for the cutover phase.

1. On the production (source) host, export the zpool:

   #zpool export *pool-name*

   ```
   # zpool export n_pool

   # zpool import
     pool: n_pool
       id: 5049703405981005579
    state: ONLINE
   action: The pool can be imported using its name or numeric identifier.
   config:

           n_pool                                      ONLINE
             c0t60A98000383035356C2447384D396550d0   ONLINE
             c0t60A98000383035356C2447384D39654Ed0   ONLINE
   ```

2. After the baseline data copy is complete, select **Test Mode** in the 7MTT user interface (UI).

3. In the 7MTT UI, click **Apply Configuration**.

   > (i) After this step, you can set the application back to online and start I/O operations to 7-Mode LUNs. Subsequent steps do not cause any application disruptions.

4. On the production host, import the zpool:

   #zpool import *pool-name*

   ```
   # zpool import n_pool
   ```

5. On the test host, rescan your new clustered Data ONTAP LUNs:

   a. Identify the FC host ports (type fc-fabric):
      #cfgadm -l

   b. Unconfigure the first fc-fabric port:
      #cfgadm -c unconfigure c1

   c. Configure the 1st fc-fabric port:
      #cfgadm -c unconfigure c2

d. Repeat the steps for other fc-fabric ports.

  e. Display information about the host ports and their attached devices:
     ```
     # cfgadm -al
     ```

  f. Reload the driver:
     ```
     # devfsadm -Cv

     # devfsadm -i iscsi
     ```

6. Verify that your clustered Data ONTAP LUNs are present:

   #sanlun lun show

   ```
   # sanlun lun show
   controller(7mode)/                              device
   host                        lun
   vserver(Cmode)        lun-pathname       filename
   adapter     protocol    size     mode
   ----------------------------------------------------------------------
   ----------------------------------------------------
   vs_5              /vol/zfs/zfs2
   /dev/rdsk/c5t600A0980383030444D2B466542485935d0s2 scsi_vhci0 FCP
   6g        C
   vs_5              /vol/zfs/zfs1
   /dev/rdsk/c5t600A0980383030444D2B466542485934d0s2 scsi_vhci0 FCP
   6g        C
   vs_5              /vol/ufs/ufs2
   /dev/rdsk/c5t600A0980383030444D2B466542485937d0s2 scsi_vhci0 FCP
   5g        C
   vs_5              /vol/ufs/ufs1
   /dev/rdsk/c5t600A0980383030444D2B466542485936d0s2 scsi_vhci0 FCP
   5g        C
   ```

7. Verify that the zpool planned to test is available for import:

   #zpool import

```
# zpool import
  pool: n_pool
    id: 5049703405981005579
 state: ONLINE
action: The pool can be imported using its name or numeric identifier.
config:

        n_pool                                 ONLINE
          c5t600A0980383030444D2B466542485935d0  ONLINE
          c5t600A0980383030444D2B466542485934d0  ONLINE
```

8. Import the zpool using the pool name or the pool ID:

    ◦ **#zpool import pool-name**

    ◦ **#zpool import pool-id**

    ```
    #zpool import n_pool
    ```

    +

    ```
    #zpool import 5049703405981005579
    ```

9. Verify that the ZFS datasets are mounted:

    ◦ **zfs list**

    ◦ **df -ah**

    ```
    # zfs list
    NAME           USED   AVAIL   REFER   MOUNTPOINT
    n_pool         2.67G  9.08G    160K   /n_pool
    n_pool/pool1   1.50G  2.50G   1.50G   /n_pool/pool1
    n_pool/pool2   1.16G  2.84G   1.16G   /n_pool/pool2
    ```

10. Perform testing as needed.

11. Shut down the test host.

12. In the 7MTT UI, click **Finish Test**.

If your clustered Data ONTAP LUNs are to be remapped to your source host, you must prepare your source host for the cutover phase. If your clustered Data ONTAP LUNs are to remain mapped to your test host, no further steps are required on the test host.

Preparing for cutover phase when transitioning Solaris host data LUNs with ZFS file systems

If you are transitioning a Solaris host data LUN with ZFS file system from Data ONTAP operating in 7-Mode to clustered Data ONTAP, you must perform certain steps before entering the cutover phase.

If you are using an FC configuration, fabric connectivity and zoning to the clustered Data ONTAP nodes must be established.

If you are using an iSCSI configuration, the iSCSI sessions to the clustered Data ONTAP nodes must be discovered and logged in.

For copy-based transitions, perform these steps before initiating the Storage Cutover operation in the 7-Mode Transition Tool (7MTT).

For copy-free transitions, perform these steps before initiating the Export & Halt 7-Mode operation in 7MTT.

Steps

1. Stop I/O on all mount points.
2. Shut down each application accessing the LUNs according to the recommendations of the application vendor.
3. Export the zpool:

   zpool export *pool-name*

   ```
   # zpool export n_pool
   ```

4. Verify that the zpools are exported:

   - Should list the zpool that is exported: + **zpool import**
   - Should not list the zpool that is exported: + **zpool list**

```
# zpool export n_pool

# zpool list
no pools available

# zpool import
  pool: n_pool
    id: 5049703405981005579
 state: ONLINE
action: The pool can be imported using its name or numeric identifier.
config:

        n_pool                                       ONLINE
          c0t60A98000383035356C2447384D396550d0  ONLINE
          c0t60A98000383035356C2447384D39654Ed0  ONLINE
```

Mounting Solaris host LUNs with ZFS file systems after transition

After transitioning Solaris host LUNs with ZFS file systems from Data ONTAP operating in 7-Mode to clustered Data ONTAP, you must mount the LUNs.

For copy-based transitions, you perform these steps after completing the Storage Cutover operation in the 7-Mode Transition Tool (7MTT).

For copy-free transitions, you perform these steps after the Import Data & Configuration operation in the 7MTT is complete.

1. Discover your new clustered Data ONTAP LUNs by rescanning the host.

   a. Identify the FC Host Ports (type fc-fabric):
      #cfgadm –l

   b. Unconfigure the 1st fc-fabric port:
      #cfgadm –c unconfigure c1

   c. Unonfigure the second fc-fabric port:
      #cfgadm –c unconfigure c2

   d. Repeat the steps for other fc-fabric ports.

   e. Verify that the information about the host ports and their attached devices is correct:
      # cfgadm –al

   f. Reload the driver:
      # devfsadm –Cv # devfsadm –i iscsi

2. Verify that your clustered Data ONTAP LUNs have been discovered:

   sanlun lun show The **lun-pathname** values for the clustered Data ONTAP LUNs should be the same as the **lun-pathname** values for the 7-Mode LUNs prior to transition. The mode column should

display "C" instead of "7".

```
# sanlun lun show
controller(7mode)/                      device
host                    lun
vserver(Cmode)        lun-pathname      filename
adapter    protocol    size    mode
--------------------------------------------------------------------
----------------------------------------------------
vs_sru17_5            /vol/zfs/zfs2
/dev/rdsk/c5t600A0980383030444D2B466542485935d0s2 scsi_vhci0 FCP
6g       C
vs_sru17_5            /vol/zfs/zfs1
/dev/rdsk/c5t600A0980383030444D2B466542485934d0s2 scsi_vhci0 FCP
6g       C
vs_sru17_5            /vol/ufs/ufs2
/dev/rdsk/c5t600A0980383030444D2B466542485937d0s2 scsi_vhci0 FCP
5g       C
vs_sru17_5            /vol/ufs/ufs1
/dev/rdsk/c5t600A0980383030444D2B466542485936d0s2 scsi_vhci0 FCP
5g       C
```

3. Check for zpools that are available to import:

   zpool import

```
# zpool import
  pool: n_vg
    id: 3605589027417030916
 state: ONLINE
action: The pool can be imported using its name or numeric identifier.
config:

        n_vg                                 ONLINE
          c0t600A098051763644575D445443304134d0  ONLINE
          c0t600A098051757A46382B445441763532d0  ONLINE
```

4. Import the zpools that were used for transition by pool name or using the pool ID:

   ◦ **zpool import *pool-name***

   ◦ **zpool import *pool-id***

```
# zpool list
no pools available

# zpool import
  pool: n_pool
    id: 5049703405981005579
 state: ONLINE
action: The pool can be imported using its name or numeric identifier.
config:

        n_pool                                    ONLINE
          c0t60A98000383035356C2447384D396550d0  ONLINE
          c0t60A98000383035356C2447384D39654Ed0  ONLINE

# zpool import n_pool
```

+

```
# zpool import 5049703405981005579

[59] 09:55:53 (root@sunx2-shu04) /tmp
# zpool list
NAME     SIZE   ALLOC   FREE  CAP  HEALTH  ALTROOT
n_pool  11.9G  2.67G  9.27G  22%  ONLINE  -
```

5. Check whether the zpool is online by doing one of the following:

   ◦ **zpool status**

   ◦ **zpool list**

```
# zpool status
  pool: n_pool
 state: ONLINE
 scan: none requested
config:

        NAME                                        STATE      READ WRITE
CKSUM
        n_pool                                      ONLINE        0      0
0
            c0t60A98000383035356C2447384D396550d0  ONLINE        0      0
0
            c0t60A98000383035356C2447384D39654Ed0  ONLINE        0      0
0

errors: No known data errors
```

+

```
# zpool list
NAME     SIZE  ALLOC   FREE  CAP  HEALTH  ALTROOT
n_pool  11.9G  2.67G  9.27G  22%  ONLINE  -
```

6. Verify the mount points by using one of the following commands:

   ◦ **zfs list**

   ◦ **df -ah**

```
# zfs list
NAME           USED  AVAIL  REFER  MOUNTPOINT
n_pool        2.67G  9.08G   160K  /n_pool
n_pool/pool1  1.50G  2.50G  1.50G  /n_pool/pool1
n_pool/pool2  1.16G  2.84G  1.16G  /n_pool/pool2

#df -ah
n_pool                  12G   160K   9.1G    1%   /n_pool
n_pool/pool1           4.0G   1.5G   2.5G   38%   /n_pool/pool1
n_pool/pool2           4.0G   1.2G   2.8G   30%   /n_pool/pool2
```