NetApp

Transitioning Solaris host data LUNs with Sun Volume Manager

ONTAP 7-Mode Transition

NetApp
October 09, 2025

Table of Contents

Transitioning Solaris host data LUNs with Sun Volume Manager

If you transition a Solaris host data LUN with Solaris Volume Manager 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 LUNs with Sun Volume Manager

Before you transition Solaris host data LUNs with Sun Volume Manager from ONTAP operating in 7-Mode to clustered ONTAP, you must gather information you need for the transition process.

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

Steps

1. Display your LUNs to identify the name of the LUNs 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 device file name for the LUN:

   #sanlun lun show

   The device file name is listed 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
```

3. List and record the SVM, and then get the disks associated with the SVM:

metaset

metaset -s *set-name*

```
# metaset
Set name = svm, Set number = 1
Host                    Owner
  Solarisx2-shu04          Yes
Drive                                        Dbase
/dev/dsk/c0t60A98000383035356C2447384D39654Ad0   Yes
/dev/dsk/c0t60A98000383035356C2447384D396548d0   Yes
```

```
# metastat -s svm
svm/d2: Concat/Stripe
    Size: 10452992 blocks (5.0 GB)
    Stripe 0:
        Device                                           Start Block
Dbase    Reloc
        /dev/dsk/c0t60A98000383035356C2447384D396548d0s0          0
No       Yes

svm/d1: Concat/Stripe
    Size: 10452992 blocks (5.0 GB)
    Stripe 0:
        Device                                           Start Block
Dbase    Reloc
        /dev/dsk/c0t60A98000383035356C2447384D39654Ad0s0          0
No       Yes

Device Relocation Information:
Device                                               Reloc   Device ID
/dev/dsk/c0t60A98000383035356C2447384D396548d0     Yes
id1,sd@n60a98000383035356c2447384d396548
/dev/dsk/c0t60A98000383035356C2447384D39654Ad0     Yes
id1,sd@n60a98000383035356c2447384d39654a
```

4. List and record the mount points:

df –ah

```
# df -ah
Filesystem              size    used   avail capacity   Mounted on
/dev/md/svm/dsk/d1      4.9G    1.5G    3.4G     31%      /d1
/dev/md/svm/dsk/d2      4.9G    991M    3.9G     20%      /d2
```

Testing data LUNs on Solaris hosts with Sun Volume Manager 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 can continue to run I/O to your source 7-Mode LUNs during testing.

Your source host with Sun Volume Manager data LUNs needs to be offline before starting the test phase transition.

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

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.

Steps

1. On the production host, disable the disk sets:

   metaset -s svm -t

   metaset -s svm -A disable

   metaset -s svm -r

   metaset -s svm -P

   metaset

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**.
4. In the production host, import the disk sets:

   metaimport -s _set-name_

```
# metaimport -s svm
Drives in regular diskset including disk
c0t60A98000383035356C2447384D39654Ad0:
  c0t60A98000383035356C2447384D39654Ad0
  c0t60A98000383035356C2447384D396548d0
More info:
  metaimport -r -v c0t60A98000383035356C2447384D39654Ad0


[22] 04:51:29 (root@sunx2-shu04) /
# metastat -s svm
svm/d2: Concat/Stripe
    Size: 10452992 blocks (5.0 GB)
    Stripe 0:
        Device                                        Start Block
Dbase   Reloc
        /dev/dsk/c0t60A98000383035356C2447384D396548d0s0        0
No      Yes

svm/d1: Concat/Stripe
    Size: 10452992 blocks (5.0 GB)
    Stripe 0:
        Device                                        Start Block
Dbase   Reloc
        /dev/dsk/c0t60A98000383035356C2447384D39654Ad0s0        0
No      Yes

Device Relocation Information:
Device                                          Reloc   Device ID
/dev/dsk/c0t60A98000383035356C2447384D396548d0   Yes
id1,sd@n60a98000383035356c2447384d396548
/dev/dsk/c0t60A98000383035356C2447384D39654Ad0   Yes
id1,sd@n60a98000383035356c2447384d39654a
```

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 first fc-fabric port:
      #cfgadm –c unconfigure c2

   d. Repeat the steps for the 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 Sun Volume Manager planned to test is available for import:

metaimport –r –v

```
# metaimport -r -v
Import: metaimport -s <newsetname> c5t600A0980383030444D2B466542485937d0
Device                                        offset        length replica
flags
c5t600A0980383030444D2B466542485937d0              16          8192       a m
luo
c5t600A0980383030444D2B466542485936d0              16          8192       a
luo
```

8. Import the metaset with a new name:

metaimport –s *set-name disk-id*

Disk-id is obtained from the `metaimport --r --v` command.

```
# metaimport -s svm c5t600A0980383030444D2B466542485937d0
Drives in regular diskset including disk
c5t600A0980383030444D2B466542485937d0:
  c5t600A0980383030444D2B466542485937d0
  c5t600A0980383030444D2B466542485936d0
More info:
  metaimport -r -v c5t600A0980383030444D2B466542485937d0
```

9. Check whether the metaset is available:

   metaset

10. Run the file system check:

    fsck -F ufs /dev/md/svm/rdsk/d1

11. Use the mount command to manually mount.

12. Perform testing as needed.

13. Shut down the test host.

14. 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 Sun Volume Manager data LUNs

If you are transitioning a Solaris host data LUN with Sun Volume Manager 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.

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. Unmount all of the mount points:

umount *mount_point*

```
#umount /d1
#umount /d2
```

4. Perform the following operations on metaset:

metaset –s *set-name* **–A disable**

metaset –s *set-name* –r

metaset –s *set-name* **–P**

```
metaset -s n_vg -A disable
metaset -s n_vg -r
metaset -s n_vg -P
```

Mounting Solaris host LUNs with Solaris Volume Manager after transition

After transitioning Solaris host LUNs with Solaris Volume Manager from ONTAP operating in 7-Mode to clustered 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.

Steps

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

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

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

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

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

   e. Verify the host ports and their attached devices:
      # cfgadm –al

   f. Reload the driver:
      # devfsadm –Cv

```
# devfsadm –i iscsi
```

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

sanlun lun show

- The `lun-pathname` values for the clustered 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. Import disk sets into existing Solaris Volume Manager configurations, using the same disk set name:

metaimport -s *set-name*

```
# metaimport -s svm
Drives in regular diskset including disk
c0t60A98000383035356C2447384D39654Ad0:
  c0t60A98000383035356C2447384D39654Ad0
  c0t60A98000383035356C2447384D396548d0
More info:
  metaimport -r -v c0t60A98000383035356C2447384D39654Ad0

# metastat -s svm
svm/d2: Concat/Stripe
    Size: 10452992 blocks (5.0 GB)
    Stripe 0:
        Device                                          Start Block
Dbase    Reloc
        /dev/dsk/c0t60A98000383035356C2447384D396548d0s0        0
No       Yes

svm/d1: Concat/Stripe
    Size: 10452992 blocks (5.0 GB)
    Stripe 0:
        Device                                          Start Block
Dbase    Reloc
        /dev/dsk/c0t60A98000383035356C2447384D39654Ad0s0        0
No       Yes

Device Relocation Information:
Device                                             Reloc   Device ID
/dev/dsk/c0t60A98000383035356C2447384D396548d0     Yes
id1,sd@n60a98000383035356c2447384d396548
/dev/dsk/c0t60A98000383035356C2447384D39654Ad0     Yes
id1,sd@n60a98000383035356c2447384d39654a
```

4. Run file system check:

   fsck -F ufs /dev/md/svm/rdsk/d1

```
# fsck -F ufs /dev/md/svm/rdsk/d1
** /dev/md/svm/rdsk/d1
** Last Mounted on /d1
** Phase 1 - Check Blocks and Sizes
** Phase 2 - Check Pathnames
** Phase 3a - Check Connectivity
** Phase 3b - Verify Shadows/ACLs
** Phase 4 - Check Reference Counts
** Phase 5 - Check Cylinder Groups
3 files, 1573649 used, 3568109 free (13 frags, 446012 blocks, 0.0%
fragmentation)
```

5. Manually mount each of the devices using the `mount` command.

```
# /sbin/mount -F ufs -o largefiles /dev/md/svm/dsk/d1   /d1
   # /sbin/mount -F ufs -o largefiles /dev/md/svm/dsk/d2   /d2
```

6. Verify the mount point:

df –ah