



Solaris host remediation

ONTAP 7-Mode Transition

NetApp
August 29, 2024

This PDF was generated from https://docs.netapp.com/us-en/ontap-7mode-transition/san-host/task_preparing_to_transition_solaris_host_data_luns_with_zfs_file_system.html on August 29, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Solaris host remediation 1
 - Transitioning Solaris host data LUNs with ZFS file systems 1
 - Transitioning Solaris host data LUNs with Sun Volume Manager 12

Solaris host remediation

If you are using the 7-Mode Transition Tool (7MTT) to move from ONTAP operating in 7-Mode to clustered ONTAP in a SAN environment, you must perform a series of steps on your Solaris host before and after transition to avoid transition complications.

The following scenarios are not supported for any transition workflow (support for copy-based or copy-free transitions):

- Transition of SAN boot LUNs

You can set up a SAN boot LUN to work in a Veritas dynamic multipathing (DMP) environment or a Solaris MPxIO environment by running the Solaris Host Utilities and using the FC protocol. The method you use to set up a SAN boot LUN can vary, depending on your volume manager and file system.

[Solaris Host Utilities 6.2 Installation and Setup Guide](#)

- Solaris host clustering transition
- Veritas configuration

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/rdisk/c0t60A98000383035356C2447384D396550d0s2  scsi_vhci0 FCP
6g            7
fas8040-shu01    /vol/zfs/zfs1
/dev/rdisk/c0t60A98000383035356C2447384D39654Ed0s2  scsi_vhci0 FCP
6g            7
fas8040-shu01    /vol/ufs/ufs2
/dev/rdisk/c0t60A98000383035356C2447384D39654Ad0s2  scsi_vhci0 FCP
5g            7
fas8040-shu01    /vol/ufs/ufs1
/dev/rdisk/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**.



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/rdisk/c5t600A0980383030444D2B466542485935d0s2 scsi_vhci0 FCP
6g                C
vs_5                /vol/zfs/zfs1
/dev/rdisk/c5t600A0980383030444D2B466542485934d0s2 scsi_vhci0 FCP
6g                C
vs_5                /vol/ufs/ufs2
/dev/rdisk/c5t600A0980383030444D2B466542485937d0s2 scsi_vhci0 FCP
5g                C
vs_5                /vol/ufs/ufs1
/dev/rdisk/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)/
host          lun          device
vserver(Cmode)  lun-pathname  filename
adapter  protocol  size    mode
-----
-----
vs_sru17_5      /vol/zfs/zfs2
/dev/rdisk/c5t600A0980383030444D2B466542485935d0s2 scsi_vhci0 FCP
6g            C
vs_sru17_5      /vol/zfs/zfs1
/dev/rdisk/c5t600A0980383030444D2B466542485934d0s2 scsi_vhci0 FCP
6g            C
vs_sru17_5      /vol/ufs/ufs2
/dev/rdisk/c5t600A0980383030444D2B466542485937d0s2 scsi_vhci0 FCP
5g            C
vs_sru17_5      /vol/ufs/ufs1
/dev/rdisk/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
```

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)/
host          lun          device
vserver(Cmode) lun-pathname filename
adapter      protocol    size    mode
-----
fas8040-shu01      /vol/zfs/zfs2
/dev/rdisk/c0t60A98000383035356C2447384D396550d0s2 scsi_vhci0 FCP
6g          7
fas8040-shu01      /vol/zfs/zfs1
/dev/rdisk/c0t60A98000383035356C2447384D39654Ed0s2 scsi_vhci0 FCP
6g          7
fas8040-shu01      /vol/ufs/ufs2
/dev/rdisk/c0t60A98000383035356C2447384D39654Ad0s2 scsi_vhci0 FCP
5g          7
fas8040-shu01      /vol/ufs/ufs1
/dev/rdisk/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:

```
metainport -s set-name
```

```
# metainport -s svm
Drives in regular diskset including disk
c0t60A98000383035356C2447384D39654Ad0:
    c0t60A98000383035356C2447384D39654Ad0
    c0t60A98000383035356C2447384D396548d0
More info:
    metainport -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/rdisk/c5t600A0980383030444D2B466542485935d0s2 scsi_vhci0 FCP
6g                C
vs_5                /vol/zfs/zfs1
/dev/rdisk/c5t600A0980383030444D2B466542485934d0s2 scsi_vhci0 FCP
6g                C
vs_5                /vol/ufs/ufs2
/dev/rdisk/c5t600A0980383030444D2B466542485937d0s2 scsi_vhci0 FCP
5g                C
vs_5                /vol/ufs/ufs1
/dev/rdisk/c5t600A0980383030444D2B466542485936d0s2 scsi_vhci0 FCP
5g                C
```

7. Verify that the Sun Volume Manager planned to test is available for import:

```
metainport -r -v
```

```
# metainport -r -v
Import: metainport -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:

```
metainport -s set-name disk-id
```

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

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

9. Check whether the metaset is available:

metaset

10. Run the file system check:

fsck -F ufs /dev/md/svm/rdisk/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/rdisk/c5t600A0980383030444D2B466542485935d0s2  scsi_vhci0 FCP
6g            C
vs_sru17_5      /vol/zfs/zfs1
/dev/rdisk/c5t600A0980383030444D2B466542485934d0s2  scsi_vhci0 FCP
6g            C
vs_sru17_5      /vol/ufs/ufs2
/dev/rdisk/c5t600A0980383030444D2B466542485937d0s2  scsi_vhci0 FCP
5g            C
vs_sru17_5      /vol/ufs/ufs1
/dev/rdisk/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

```
# metainport -s svm
Drives in regular diskset including disk
c0t60A98000383035356C2447384D39654Ad0:
    c0t60A98000383035356C2447384D39654Ad0
    c0t60A98000383035356C2447384D396548d0
More info:
    metainport -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/rdisk/d1
```

```
# fsck -F ufs /dev/md/svm/rdisk/d1
** /dev/md/svm/rdisk/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

Copyright information

Copyright © 2024 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.