

How to migrate the root pool

- I recently wanted to migrate the root pool to a new device. This turned out to be easy to do, using existing facilities. The original root pool was on an old 80-gig disk. This system also had a data pool on a newer 1 TB disk. Here's what the `format` command showed for them:

```
0. c2t0d0 <Unknown-Unknown-0001 cyl 9726 alt 2 hd 255 sec 63>
   /pci@0,0/pci1043,8389@11/disk@0,0
1. c2t2d0 <ATA-ST31000524AS-JC4B-931.51GB>
   /pci@0,0/pci1043,8389@11/disk@2,0I
```

- wanted to migrate the root pool to a new SSD. The `format` command was available to prepare the SSD. I could use the `zpool` command to create the pool on that new device, and `beadm` and `installgrub` to perform the migration. That part worked out nicely. I had to use a variety of commands to complete the migration.
- Add the SSD: Just shut down the computer, install the SSD hardware, and boot the system. Here's the new output from `format`:

```
0. c2t0d0 <Unknown-Unknown-0001 cyl 9726 alt 2 hd 255 sec 63>
   /pci@0,0/pci1043,8389@11/disk@0,0
1. c2t1d0 <ATA-SanDiskSDSSDP06-0 cyl 9966 alt 2 hd 224 sec 56>
   /pci@0,0/pci1043,8389@11/disk@1,0
2. c2t2d0 <ATA-ST31000524AS-JC4B-931.51GB>
   /pci@0,0/pci1043,8389@11/disk@2,0
```

- Prepare the SSD: Create the fdisk partition within `format`:

```
format> fdisk
No fdisk table exists. The default partition for the disk is:
  a 100% "SOLARIS System" partition
Type "y" to accept the default partition, otherwise type "n" to edit the
partition table.
y
```

- Create the slice:

```
partition> 0
Part      Tag      Flag      Cylinders      Size      Blocks
0 unassigned  wm        0              0          (0/0/0)      0
Enter partition id tag[unassigned]: root
Enter partition permission flags[wm]:
Enter new starting cyl[1]: 3
Enter partition size[0b, 0c, 3e, 0.00mb, 0.00gb]: $
partition> p
Current partition table (unnamed):
Total disk cylinders available: 9965 + 2 (reserved cylinders)
Part      Tag      Flag      Cylinders      Size      Blocks
0         root     wm        3 - 9964       59.59GB   (9962/0/0) 124963328
1 unassigned  wm        0              0          (0/0/0)      0
2 backup     wu        0 - 9964       59.61GB   (9965/0/0) 125000960
3 unassigned  wm        0              0          (0/0/0)      0
4 unassigned  wm        0              0          (0/0/0)      0
5 unassigned  wm        0              0          (0/0/0)      0
6 unassigned  wm        0              0          (0/0/0)      0
7 unassigned  wm        0              0          (0/0/0)      0
8 boot       wu        0 - 0          6.12MB    (1/0/0)     12544
9 unassigned  wm        0              0          (0/0/0)      0
partition> l
Ready to label disk, continue? y
```

- Get the root pool version:

```
# zpool get all rpool
NAME      PROPERTY      VALUE      SOURCE
rpool    size          74G        -
...
rpool    version      28         local
```

- Try to create the new root pool, with a new pool name:

```
# zpool create -o version=28 rpool1 c2t1d0s0
invalid vdev specification
use '-f' to override the following errors:
/dev/dsk/c2t1d0s0 overlaps with /dev/dsk/c2t1d0s2
```

- Try again with the force option:

```
# zpool create -f -o version=28 rpool1 c2t1d0s0
# zpool list
NAME      SIZE  ALLOC   FREE  EXPANDSZ   CAP  DEDUP  HEALTH  ALTROOT
dpool     928G  85.6G   842G   2.50M      9%   1.00x  ONLINE  -
rpool     74G   7.05G   66.9G   -           9%   1.00x  ONLINE  -
rpool1    59.5G  108K    59.5G   -           0%   1.00x  ONLINE  -
```

- Create the BE, on the new device with a new name:

```
# beadm create -p rpool1 oi_151a6x
WARNING: menu.lst file /rpool1/boot/grub/menu.lst does not exist,
generating a new menu.lst file
Created successfully
```

- Verify that it exists:

```
# beadm list
BE      Active Mountpoint Space Policy Created
oi_151a6 NR / 5.98G static 2012-09-13 16:33
oi_151a6x R - 4.15G static 2013-06-06 15:55
openindiana - - 13.5M static 2012-09-13 08:55
```

- Install the boot blocks:

```
# installgrub -m /boot/grub/stage1 /boot/grub/stage2 /dev/rdisk/c2t1d0s0
Updating master boot sector destroys existing boot managers (if any).
continue (y/n)?y
stage2 written to partition 0, 277 sectors starting at 50 (abs 12594)
stage1 written to partition 0 sector 0 (abs 12544)
stage1 written to master boot sector
```

- Change the BIOS boot order by shutting the system down and entering the BIOS setup. Then put the SSD first in the boot order and reboot.
- At this point, I upgraded to oi_151a7. This confirmed that the new root pool was functional. Here's the initial boot environment:

```
# beadm list
BE      Active Mountpoint Space Policy Created
oi_151a6 R - 6.01G static 2012-09-13 16:33
oi_151a6x NR / 4.33G static 2013-06-06 15:55
openindiana - - 13.5M static 2012-09-13 08:55
```

- Upgrade:

```
# pkg image-update --be-name oi_151a7
WARNING: The boot environment being modified is not the active one. Changes made in the active BE will
not be reflected on the next boot.
Packages to update: 895
Create boot environment: Yes
Create backup boot environment: No
...
A clone of oi_151a6x exists and has been updated and activated.
On the next boot the Boot Environment oi_151a7 will be
mounted on '/'. Reboot when ready to switch to this updated BE.
```

- Check the BEs again:

```
# beadm list
BE      Active Mountpoint Space Policy Created
oi_151a6 R - 6.01G static 2012-09-13 16:33
oi_151a6x N / 101K static 2013-06-06 15:55
oi_151a7 R - 5.31G static 2013-06-06 16:56
openindiana - - 13.5M static 2012-09-13 08:55
```

- Shut down OS:

```
# init 5
updating //platform/i86pc/boot_archive
updating //platform/i86pc/amd64/boot_archive
```

- Press the 'Power' button to reboot. Confirm that the upgrade was successful. Notice that there are still two active boot environments:

```
$ beadm list
BE      Active Mountpoint Space Policy Created
oi_151a6  R      -           6.01G static 2012-09-13 16:33
oi_151a6x -      -           16.8M static 2013-06-06 15:55
oi_151a7  NR     /           5.33G static 2013-06-06 16:56
openindiana -      -           13.5M static 2012-09-13 08:55
```

- Some of the old root pool is still in use. My home directory was on rpool/export/home/mills . To simplify this migration, I decided to move it to the data pool. First, create new filesystems on the data pool:

```
# zfs create dpool/export
# zfs create dpool/export/home
```

- My home directory in the /etc/passwd file was automounted to /home/mills from /export/home/mills . The first thing I did was to copy it to /dpool/export/home/mills using 'cpio'. Then I edited /etc/passwd to change my home directory to /dpool/export/home/mills . After that change, it was no longer automounted. After a reboot, I confirmed that the old root pool was no longer needed for my home directory:

```
# zfs unmount rpool/export/home/mills
# zfs unmount rpool/export/home
# zfs unmount rpool/export
```

- Still, there are a few pieces left:

```
# zfs list | egrep 'dump|swap'
rpool/dump      895M 65.5G 895M -
rpool/swap      952M 65.8G 637M -
```

- To move the dump device, first get the properties of the old one:

```
$ zfs get all rpool/dump | egrep 'SOURCE|local'
NAME      PROPERTY      VALUE      SOURCE
rpool/dump volsize       895M      local
rpool/dump checksum      off        local
rpool/dump compression  off        local
rpool/dump refreservation none        local
rpool/dump dedup       off        local
```

- Create another one on rpool1:

```
# zfs create -o checksum=off -o compression=off -o refreservation=none -o dedup=off -V 895M rpool1/dump
```

- Try to move it:

```
# dumpadm -d /dev/zvol/dsk/rpool1/dump
dumpadm: dump device /dev/zvol/dsk/rpool1/dump is too small to hold a system dump
dump size 1812297728 bytes, device size 938475520 bytes
# dumpadm
Dump content: kernel pages
Dump device: /dev/zvol/dsk/rpool/dump (dedicated)
Savecore directory: /var/crash/ati
Savecore enabled: no
Save compressed: on
```

- Expand the volume and try again:

```
# dumpadm -d /dev/zvol/dsk/rpool1/dump
```

```
Dump content: kernel pages
Dump device: /dev/zvol/dsk/rpool1/dump (dedicated)
Savecore directory: /var/crash/ati
Savecore enabled: no
Save compressed: on
```

- Now, get the properties of the old swap device:

```
$ zfs get all rpool/swap | egrep 'SOURCE|local'
NAME      PROPERTY      VALUE      SOURCE
rpool/swap volsize       895M      local
rpool/swap refreservation 952M      local
```

- Create a new one on rpool1:

```
# zfs create -o refreservation=952M -V 895M rpool1/swap
```

- Move the swap device by editing /etc/vfstab:

```
o Move the swap device by editing /etc/vfstab:
root@ati:/etc# cp -p vfstab vfstab-
root@ati:/etc# ex vfstab
root@ati:/etc# diff vfstab- vfstab
l2c12
< /dev/zvol/dsk/rpool/swap          -          -          swap          -          no          -
---
> /dev/zvol/dsk/rpool1/swap        -          -          swap          -          no          -
```

- Reboot and confirm that rpool is no longer used:

```
# dumpadm
Dump content: kernel pages
Dump device: /dev/zvol/dsk/rpool1/dump (dedicated)
Savecore directory: /var/crash/ati
Savecore enabled: no
Save compressed: on
# swap -l
swapfile          dev      swaplo    blocks    free
/dev/zvol/dsk/rpool1/swap 96,2      8 1832952 1832952
# beadm list
Active Mountpoint Space Policy Created
oi_151a6          R        -          6.01G static 2012-09-13 16:33
oi_151a6x         -        -          16.8M static 2013-06-06 15:55
oi_151a7          NR       /          5.34G static 2013-06-06 16:56
openindiana      -        -          13.5M static 2012-09-13 08:55
# zpool list
NAME      SIZE  ALLOC   FREE  EXPANDSZ   CAP  DEDUP  HEALTH  ALTROOT
dpool    928G  85.6G  842G    2.50M      9%  1.00x  ONLINE  -
rpool    74G   6.19G  67.8G    -          8%  1.00x  ONLINE  -
rpool1   59.5G 7.17G  52.3G    -         12%  1.00x  ONLINE  -
```

- Export the pool and observe the result:

```
# zpool export rpool
# zpool list
NAME      SIZE  ALLOC   FREE  EXPANDSZ   CAP  DEDUP  HEALTH  ALTROOT
dpool    928G  85.6G  842G    2.50M      9%  1.00x  ONLINE  -
rpool1   59.5G 7.18G  52.3G    -         12%  1.00x  ONLINE  -
# zfs list
NAME                USED  AVAIL  REFER  MOUNTPOINT
dpool                85.6G  828G   24K    /dpool
dpool/export         83.8G  828G   22K    /dpool/export
dpool/export/home    83.8G  828G  83.8G  /dpool/export/home
dpool/opt            1.82G  828G   1.82G  /dpool/opt
dpool/opt/local      21K    828G   21K    /dpool/opt/local
rpool1               8.10G  50.5G  36.5K  /rpool1
rpool1/ROOT          5.17G  50.5G   31K    legacy
rpool1/ROOT/oi_151a6x 16.8M  50.5G  4.33G  /
rpool1/ROOT/oi_151a7  5.16G  50.5G  4.27G  /
rpool1/dump          2.00G  50.5G  2.00G  -
rpool1/swap          952M   51.4G   16K    -
# getent passwd mills
mills:x:107:10:Gary Mills:/dpool/export/home/mills:/bin/ksh
# beadm list
BE      Active Mountpoint Space Policy Created
oi_151a6x -        -          16.8M static 2013-06-06 15:55
oi_151a7 NR       /          5.34G static 2013-06-06 16:56
```

- I could have resumed automounting my home directory by changing the mount point of dpool/export to /export, but I decided to leave it the way it was.
- Here's another upgrade, just to confirm that the new root pool was correct:

```

# pkg image-update --be-name oi_151a8
  Packages to remove: 16
  Packages to install: 6
  Packages to update: 879
  Create boot environment: Yes
Create backup boot environment: No
DOWNLOAD
Completed          PKGS      FILES    XFER (MB)
                   901/901  22745/22745  566.2/566.2
PHASE              ACTIONS
Removal Phase     13844/13844
Install Phase     12382/12382
Update Phase      23637/23637
PHASE              ITEMS
Package State Update Phase 1780/1780
Package Cache Update Phase  895/895
Image State Update Phase   2/2
...
root@ati:~# beadm list
BE      Active Mountpoint Space Policy Created
oi_151a6x -      -      16.8M static 2013-06-06 15:55
oi_151a7 N      /      11.4M static 2013-06-06 16:56
oi_151a8 R      -      8.76G static 2013-08-11 16:12
# bootadm list-menu
the location for the active GRUB menu is: /rpool1/boot/grub/menu.lst
default 2
timeout 30
0 oi_151a6x
1 oi_151a7
2 oi_151a8
# init 5

```

- Press the power switch to reboot. The upgrade was successful, completing the migration to a new device.