

ZPOOL: Add a mirror to a concat zpool

Article Number: 188 | Rating: Unrated | Last Updated: Sun, Jun 3, 2018 9:17 AM

ZPOOL: Add a mirror to a concat zpool

```
# Tested on RHEL 6 & 7

# Even if ZFS can use individual slices or partitions, it is
recommended to use whole disks.

# Having a concat zpool like following one

POOLNAME=my_pool

zpool status $POOLNAME
    pool: my_pool
    state: ONLINE
    scan: none requested
    config:

              NAME        STATE      READ WRITE CKSUM
    my_pool      ONLINE            0     0      0
                  sdb      ONLINE            0     0      0
                  sdc      ONLINE            0     0      0

    errors: No known data errors

zpool list
    NAME      SIZE   ALLOC   FREE    CAP   DEDUP  HEALTH  ALTROOT
    my_pool  19.9G  104K  19.9G    0%  1.00x  ONLINE   -
```

```
# we would like to add a mirror

DEVICE01_ORIG=/dev/sdb
DEVICE02_ORIG=/dev/sdc

DEVICE01_MIRR=/dev/sdd
DEVICE02_MIRR=/dev/sde

# We add one by one the new devices to the existing ones to form the
new mirrors

zpool attach $POOLNAME $DEVICE01_ORIG $DEVICE01_MIRR

# If you have an error like this one:

    invalid vdev specification
    use '-f' to override the following errors:
    /dev/sdd does not contain an EFI label but it may contain
partition
    information in the MBR.

# you should use '-f' option to create the pool - first ensure that
disk(s) are the
# right one(s):

# zpool attach -f $POOLNAME $DEVICE01_ORIG $DEVICE01_MIRR

# Check

zpool status $POOLNAME
    pool: my_pool
    state: ONLINE
    scan: resilvered 36.5K in 0h0m with 0 errors on Tue Sep  1
```

```
15:18:06 2015
```

```
config:
```

NAME	STATE	READ	WRITE	CKSUM	
my_pool	ONLINE	0	0	0	
mirror-0	ONLINE	0	0	0	
sdb	ONLINE	0	0	0	
sdd	ONLINE	0	0	0	<----
sdc	ONLINE	0	0	0	

```
errors: No known data errors
```

```
zpool attach $POOLNAME $DEVICE02_ORIG $DEVICE02_MIRR
```

```
# If you have an error like this one:
```

```
invalid vdev specification
use '-f' to override the following errors:
/dev/sde does not contain an EFI label but it may contain
partition
information in the MBR.
```

```
# you should use '-f' option to create the pool - first ensure that
disk(s) are the
# right one(s):
```

```
# zpool attach -f $POOLNAME $DEVICE02_ORIG $DEVICE02_MIRR
```

```
# Check
```

```
zpool status $POOLNAME
```

```
pool: my_pool
state: ONLINE
scan: resilvered 64K in 0h0m with 0 errors on Tue Sep 1 15:18:45
2015
```

```
config:
```

NAME	STATE	READ	WRITE	CKSUM
my_pool	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0
sdb	ONLINE	0	0	0
sdd	ONLINE	0	0	0 <----
mirror-1	ONLINE	0	0	0
sdc	ONLINE	0	0	0
sde	ONLINE	0	0	0 <----

```
errors: No known data errors
```

```
zpool list
```

NAME	SIZE	ALLOC	FREE	CAP	DEDUP	HEALTH	ALTROOT
my_pool	19.9G	116K	19.9G	0%	1.00x	ONLINE	-

```
# -----
-----  
# Adding a mirror to a concat zpool using disks of different sizes  
# -----  
-----
```

```
# I will use two asymmetric disk partitions with the aim of giving an example:
```

```
lvmdiskscan | egrep "sdb1|sdc1"  
/dev/sdb1 [ 5.00 GiB]  
/dev/sdc1 [ 10.00 GiB]
```

```
# 1.- If we try to mirror a pool using a disk that is smaller than
```

```
the existing one, this
# is what happens:

zpool create my_pool sdc1

zpool list my_pool
  NAME      SIZE  ALLOC   FREE    CAP  DEDUP  HEALTH  ALTROOT
  my_pool  9.94G  77.5K  9.94G    0%  1.00x  ONLINE   -
  
zpool attach my_pool sdc1 sdb1
  cannot attach sdb1 to sdcl: device is too small

# No problem so far: System prevents us from creating a mirror in a
wrong way; that's all

# 2.- When adding a mirror using a disk that is larger than the
existing one, we have to
# use the "autoexpand" property, that must be enabled for the whole
size of the new disk
# to be used instead of limiting the pool to the original size. This
will happens once
# the original disk has been removed from the pool, otherwise we
would have an "asymmetric"
# mirror.

# Create a simple pool

zpool create my_pool sdb1

zpool list my_pool
  NAME      SIZE  ALLOC   FREE    CAP  DEDUP  HEALTH  ALTROOT
  my_pool  4.97G   109K  4.97G    0%  1.00x  ONLINE   -
  
# Enable "autoexpand"
```

```
zpool set autoexpand=on my_pool

# Add the mirror (new disk bigger than existing one)

zpool attach my_pool sdb1 sdc1

# Size remains constant...

zpool list my_pool
  NAME      SIZE  ALLOC   FREE    CAP  DEDUP  HEALTH  ALTROOT
  my_pool  4.97G  159K  4.97G    0%  1.00x  ONLINE   -
# ... until I detach the original disk. At that moment the pool is
# "autoexpanded"

zpool detach my_pool sdb1

zpool list my_pool
  NAME      SIZE  ALLOC   FREE    CAP  DEDUP  HEALTH  ALTROOT
  my_pool  9.97G  168K  9.97G    0%  1.00x  ONLINE   -
# If we had forgotten to enable auto-expansion before detaching the
original disk, pool's
# size is not expanded even if we change "autoexpand" property
afterwards. To get to
# expand the pool we have to run a "zpool online" command on the new
disk(s)

zpool create my_pool sdb1

zpool attach -f my_pool sdb1 sdc1
```

```
zpool list my_pool
  NAME      SIZE  ALLOC   FREE    CAP  DEDUP  HEALTH  ALTROOT
  my_pool  4.97G  156K  4.97G    0%  1.00x  ONLINE  -
zpool detach my_pool sdb1

zpool list my_pool
  NAME      SIZE  ALLOC   FREE    CAP  DEDUP  HEALTH  ALTROOT
  my_pool  4.97G  156K  4.97G    0%  1.00x  ONLINE  -
# Ooops, I forgot to enable the autoexpand property, I will do it
now...

zpool get autoexpand my_pool
  NAME      PROPERTY     VALUE   SOURCE
  my_pool  autoexpand  off     default

zpool set autoexpand=on my_pool

# No way, it doesn't work:

zpool list my_pool
  NAME      SIZE  ALLOC   FREE    CAP  DEDUP  HEALTH  ALTROOT
  my_pool  4.97G  172K  4.97G    0%  1.00x  ONLINE  -
# No problem, let's run an "online" on the new disk (even if it is
already online)

zpool online my_pool sdc1

# and voilà!, the pool has been expanded:

zpool list my_pool
  NAME      SIZE  ALLOC   FREE    CAP  DEDUP  HEALTH  ALTROOT
  my_pool  9.97G  153K  9.97G    0%  1.00x  ONLINE  -
```

Posted - Sun, Jun 3, 2018 9:17 AM. This article has been viewed 3972 times.

Online URL: <http://kb.ictbanking.net/article.php?id=188>