



# Übung: RAID-1

Solaris Volume Manager Administration Guide

<http://docs.oracle.com/cd/E19253-01/816-4520/index.html>

Chapter 10 RAID-1 (Mirror) Volumes (Overview)

<http://docs.oracle.com/cd/E19253-01/816-4520/about-mirrors-2/index.html>

## 1. Erstellen des RAID-1

<http://docs.oracle.com/cd/E19253-01/816-4520/6manpiejl/index.html>

```
root@ssm1417-067:~ # metainit -f d51 1 1 c1t1d0s0
d51: Concat/Stripe is setup
```

```
root@ssm1417-067:~ # metainit -f d52 1 1 c1t2d0s0
d52: Concat/Stripe is setup
```

```
root@ssm1417-067:~ # metainit d50 -m d51 d52
metainit: d50: WARNING: This form of metainit is not recommended.
The submirrors may not have the same data.
Please see ERRORS in metainit(1M) for additional information.
d50: Mirror is setup
```

## 2. Initialisieren des Dateisystems auf dem erstellten RAID

```
root@ssm1417-067:~ # newfs /dev/md/dsk/d50
newfs: construct a new file system /dev/md/rdisk/d50: (y/n)? y
/dev/md/rdisk/d50:      4153344 sectors in 1014 cylinders of 128 tracks, 32 sectors
      2028.0MB in 45 cyl groups (23 c/g, 46.00MB/g, 11264 i/g)
super-block backups (for fsck -F ufs -o b=#) at:
 32, 94272, 188512, 282752, 376992, 471232, 565472, 659712, 753952, 848192,
3298432, 3392672, 3486912, 3581152, 3675392, 3769632, 3863872, 3958112,
4052352, 4146592
```

## 3. Einhängen und Speicherbelegung ausgeben

```
root@ssm1417-067:~ # mount /dev/md/dsk/d50 /mnt
```

```
root@ssm1417-067:~ # df -h /mnt
Filesystem      size  used  avail capacity  Mounted on
/dev/md/dsk/d50  1.9G  2.0M  1.9G    1%    /mnt
```

## 4. Aufräumen

```
root@ssm1417-067:~ # umount /mnt
```

```
root@ssm1417-067:~ # metaclear d50
d50: Mirror is cleared
```

```
root@ssm1417-067:~ # metaclear d51
d51: Concat/Stripe is cleared
```

```
root@ssm1417-067:~ # metaclear d52
d52: Concat/Stripe is cleared
```

```
root@ssm1417-067:~ # metadb -d -f c1t1d0s7
root@ssm1417-067:~ # metadb -d -f c1t2d0s7
```

# ZFS

## Oracle Solaris ZFS Administration Guide

[http://docs.oracle.com/cd/E23823\\_01/html/819-5461/index.html](http://docs.oracle.com/cd/E23823_01/html/819-5461/index.html)

### 1. Einfacher ZFS Pool

*Pool erstellen und zweite Platte als Mirror hinzufügen*

```
root@ssm1417-067:~ # zpool create -f my_pool c1t1d0
root@ssm1417-067:~ # zpool attach -f my_pool c1t1d0 c1t2d0
```

```
root@ssm1417-067:~ # zpool status my_pool
pool: my_pool
state: ONLINE
scan: resilvered 83.5K in 0h0m with 0 errors on Thu Nov 21 10:36:56 2013
config:
```

NAME	STATE	READ	WRITE	CKSUM
my_pool	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0
c1t1d0	ONLINE	0	0	0
c1t2d0	ONLINE	0	0	0

errors: No known data errors

*Zweite Platte entfernen (Spiegel aufbrechen) und normal hinzufügen (RAID-0)*

```
root@ssm1417-067:~ # zpool detach my_pool c1t2d0
root@ssm1417-067:~ # zpool add my_pool c1t2d0
```

```
root@ssm1417-067:~ # zpool status my_pool
pool: my_pool
state: ONLINE
scan: resilvered 83.5K in 0h0m with 0 errors on Thu Nov 21 10:36:56 2013
config:
```

NAME	STATE	READ	WRITE	CKSUM
my_pool	ONLINE	0	0	0
c1t1d0	ONLINE	0	0	0
c1t2d0	ONLINE	0	0	0

errors: No known data errors

*Alle wichtigen Grunddaten abfragen:*

```
root@ssm1417-067:~ # zfs get all my_pool
NAME      PROPERTY          VALUE                SOURCE
my_pool   type              filesystem           -
my_pool   creation          Thu Nov 21 10:36 2013 -
my_pool   used              104K                 -
my_pool   available         3.91G                -
my_pool   referenced        31K                   -
my_pool   compressratio     1.00x                -
my_pool   mounted           yes                   -
my_pool   quota             none                  default
my_pool   reservation       none                  default
my_pool   recordsize        128K                  default
my_pool   mountpoint         /my_pool              default
my_pool   sharenfs          off                   default
my_pool   checksum          on                    default
my_pool   compression       off                   default
my_pool   atime             on                    default
my_pool   devices           on                    default
my_pool   exec              on                    default
my_pool   setuid            on                    default
my_pool   readonly         off                   default
my_pool   zoned             off                   default
my_pool   snapdir           hidden                default
my_pool   aclinherit        restricted             default
my_pool   canmount          on                    default
my_pool   shareiscsi        off                   default
my_pool   xattr             on                    default
my_pool   copies            1                     default
my_pool   version           5                     -
my_pool   utf8only         off                   -
my_pool   normalization     none                  -
my_pool   casesensitivity   sensitive              -
my_pool   vscan             off                   default
my_pool   nbmand            off                   default
my_pool   sharesmb          off                   default
my_pool   refquota         none                  default
my_pool   refreservation   none                  default
my_pool   primarycache     all                   default
my_pool   secondarycache   all                   default
my_pool   usedbysnapshots  0                     -
my_pool   usedbydataset     31K                   -
my_pool   usedbychildren   73.5K                 -
my_pool   usedbyrefreservation 0                     -
my_pool   logbias          latency               default
my_pool   sync              standard              default
my_pool   rstchown         on                    default
```

## 2. Übung: ZFS Dateisysteme und Attribute

*Einige ZFS Dateisysteme anlegen, „Home“ Verzeichnisse auf max. 50MB beschränken.*

```
root@ssm1417-067:~ # zfs create my_pool/home
root@ssm1417-067:~ # zfs create my_pool/home/my_home
root@ssm1417-067:~ # zfs create my_pool/home/your_home
root@ssm1417-067:~ # zfs create my_pool/home/his_home
root@ssm1417-067:~ # zfs create my_pool/home/her_home
root@ssm1417-067:~ # zfs create my_pool/home/our_home
root@ssm1417-067:~ # zfs list
```

NAME	USED	AVAIL	REFER	MOUNTPOINT
my_pool	354K	3.91G	32K	/my_pool
my_pool/home	190K	3.91G	35K	/my_pool/home
my_pool/home/her_home	31K	3.91G	31K	/my_pool/home/her_home
my_pool/home/his_home	31K	3.91G	31K	/my_pool/home/his_home
my_pool/home/my_home	31K	3.91G	31K	/my_pool/home/my_home
my_pool/home/our_home	31K	3.91G	31K	/my_pool/home/our_home
my_pool/home/your_home	31K	3.91G	31K	/my_pool/home/your_home

```
root@ssm1417-067:~ # zfs set quota=50m my_pool/home
root@ssm1417-067:~ # zfs list
```

NAME	USED	AVAIL	REFER	MOUNTPOINT
my_pool	570K	3.91G	32K	/my_pool
my_pool/home	192K	49.8M	37K	/my_pool/home
my_pool/home/her_home	31K	49.8M	31K	/my_pool/home/her_home
my_pool/home/his_home	31K	49.8M	31K	/my_pool/home/his_home
my_pool/home/my_home	31K	49.8M	31K	/my_pool/home/my_home
my_pool/home/our_home	31K	49.8M	31K	/my_pool/home/our_home
my_pool/home/your_home	31K	49.8M	31K	/my_pool/home/your_home

*In „my\_home“ mind. 15MB reservieren:*

```
root@ssm1417-067:~ # zfs set reservation=15m my_pool/home/my_home
root@ssm1417-067:~ # zfs list
```

NAME	USED	AVAIL	REFER	MOUNTPOINT
my_pool	15.3M	3.89G	32K	/my_pool
my_pool/home	15.2M	34.8M	37K	/my_pool/home
my_pool/home/her_home	31K	34.8M	31K	/my_pool/home/her_home
my_pool/home/his_home	31K	34.8M	31K	/my_pool/home/his_home
my_pool/home/my_home	31K	49.8M	31K	/my_pool/home/my_home
my_pool/home/our_home	31K	34.8M	31K	/my_pool/home/our_home
my_pool/home/your_home	31K	34.8M	31K	/my_pool/home/your_home

*In allen Verzeichnissen ausser „my\_home“ sinkt der verfügbare Platz um 15MB, da diese ja fix für „my\_home“ reserviert wurden.*

*Aufräumen:*

```
root@ssm1417-067:~ # zfs destroy my_pool/home/your_home
root@ssm1417-067:~ # zfs destroy my_pool/home/our_home
root@ssm1417-067:~ # zfs destroy my_pool/home/my_home
root@ssm1417-067:~ # zfs destroy my_pool/home/his_home
root@ssm1417-067:~ # zfs destroy my_pool/home/her_home
root@ssm1417-067:~ # zfs destroy my_pool/home
```

```
root@ssm1417-067:~ # zpool destroy my_pool
```

## Persönliche Lernkontrolle

1. Welches sind die Unterschiede zwischen der Erzeugung eines RAID-1 ohne und eines RAID-1 mit Resync?

*Bei einem Resync werden die Daten von der bestehenden Disk0 auf die zusätzliche Disk1 synchronisiert (übernommen). Bei einem Erzeugen ohne Resync werden die Daten gelöscht.*

2. Was passiert bei einem RAID-1 Resync?

*Die Daten werden auf beide Disks gespiegelt / kopiert.*

3. Welche Werkzeuge werden für das Management von ZFS hauptsächlich benötigt?

*zpool*

*zfs*