




tech tips



How to replace a disk with Solaris Volume Manager

The following tip is to explain how to replace a failed disk in solaris system using the Solaris Volume Manager in Solaris 10

For this example, consider the following:

Device Relocation Information:

Device Reloc Device ID

c1t0d0 Yes id1,sd@SSEAGATE_ST318203LSUN18G_LRA287950000W0270V5A

c1t1d0 Yes id1,sd@SFUJITSU_MAG3182L_SUN18G_01504189_____

```
# metastat -p
d300 -m d310 d320 1
d310 1 1 c3t1d0s0
d320 1 1 c3t2d0s0
d3 -m d13 d23 1
d13 1 1 c1t0d0s3
d23 1 1 c1t1d0s3
d5 -m d15 d25 1
d15 1 1 c1t0d0s5
d25 1 1 c1t1d0s5
d1 -m d11 d21 1
d11 1 1 c1t0d0s1
d21 1 1 c1t1d0s1
d0 -m d10 d20 1
d10 1 1 c1t0d0s0
d20 1 1 c1t1d0s0
```

The failed disk for this example will be c1t1d0 which is a mirrored copy of c1t0d0.

The first thing we need to do is determine whether this disk has any meta databases installed on it. To do this type:

```
# metadb -i
      flags      first blk      block count
a m p lu0 16      8192      /dev/dsk/c1t0d0s6
a   p lu0 8208     8192      /dev/dsk/c1t0d0s6
a   p lu0 16      8192      /dev/dsk/c1t0d0s7
a   p lu0 8208     8192      /dev/dsk/c1t0d0s7
a   p lu0 16      8192      /dev/dsk/c1t1d0s6
a   p lu0 8208     8192      /dev/dsk/c1t1d0s6
a   p lu0 16      8192      /dev/dsk/c1t1d0s7
a   p lu0 8208     8192      /dev/dsk/c1t1d0s7
```

As you can see, there are 2 metadb's on the failed disk. Slice 6 and 7. We can go ahead and remove these.

```
# metadb -d /dev/dsk/c1t1d0s6 /dev/dsk/c1t1d0s7
```

The meta databases should now be deleted on the failed disk. You can check by running 'metadb -i' again.

```
# metadb -i
      flags      first blk      block count
a m p lu0      16            8192          /dev/dsk/c1t0d0s6
a   p lu0      8208          8192          /dev/dsk/c1t0d0s6
a   p lu0      16            8192          /dev/dsk/c1t0d0s7
a   p lu0      8208          8192          /dev/dsk/c1t0d0s7
```

Next we need to unconfigure the device to do this run a 'cfgadm -al'

```
# cfgadm -al | grep c1t1d0
c1::dsk/c1t1d0  disk      connected    configured    unknown
# cfgadm -c unconfigure c1::dsk/c1t1d0
```

The disk should now be unconfigured from the system

```
# cfgadm -al | grep c1t1d0
c1::dsk/c1t1d0  disk      connected    unconfigured    unknown
```

You can go ahead with the physical replace of the drive.

Now, we will need to configure the drive.

```
# cfgadm -c configure c1::dsk/c1t1d0
```

To check it has been configured run:

```
# cfgadm -al | grep c1t1d0
c1::dsk/c1t1d0  disk      connected    configured    unknown
```

We will now need to copy the disk labels across from the working disk and copy them to the new disk.

The prtvtoc will copy the existing layout of partitions and sizes of the working disk and fmthard will format the new disk using that information.

```
# prtvtoc /dev/rdisk/c1t0d0s2 | fmthard -s - /dev/rdisk/c1t1d0s2
```

The next step is to run metareplace on all the meta devices which contained the failed disk c1t1d0.

To identify these disk you can run 'metastat -p' again

```
# metastat -p

d300 -m d310 d320 1
d310 1 1 c3t1d0s0
d320 1 1 c3t2d0s0
d3 -m d13 d23 1
d13 1 1 c1t0d0s3
d23 1 1 c1t1d0s3
d5 -m d15 d25 1
d15 1 1 c1t0d0s5
d25 1 1 c1t1d0s5
d1 -m d11 d21 1
d11 1 1 c1t0d0s1
d21 1 1 c1t1d0s1
d0 -m d10 d20 1
d10 1 1 c1t0d0s0
d20 1 1 c1t1d0s0
```

As you can see from the above, c1t1d0 had mirrors in d0, d1, d3 and d5.

Do the following:

```
# metareplace -e d0 c1t1d0s0
# metareplace -e d1 c1t1d0s1
# metareplace -e d3 c1t1d0s3
# metareplace -e d5 c1t1d0s5
```

The disks should now be resync, when this is complete all the metadevices should be in an OK state.