

A Complete Guide to Replace Failed SAS / SCSI Disk and Mirrored with SVM

1. Display existing configured and attached disk in system. In this sample, c1t1d0 will be a failed disk.

```
# cfdm -al
Ap_Id                Type      Receptacle  Occupant    Condition
c1                   scsi-bus  connected   configured  unknown
c1::dsk/c1t0d0       disk      connected   configured  unknown
c1::dsk/c1t1d0       disk      connected   configured  unknown
```

2. View current mirror group settings which are d10 and d20 and sub-mirrors are d11,d12,d21,22.

```
# metastat -c
d20                m  517MB  d22  d21
   d22              s  517MB  c1t1d0s1
   d21              s  517MB  c1t0d0s1
d10                m  8.2GB  d12  d11
   d12              s  8.2GB  c1t1d0s0
   d11              s  8.2GB  c1t0d0s0
```

3. View existing metadb configurations. Current db is lay on s7 in both disks.

```
# metadb
      flags          first blk      block count
a m p lu0          16             8192        /dev/dsk/c1t0d0s7
a   p lu0          8208            8192        /dev/dsk/c1t0d0s7
a   p lu0          16400           8192        /dev/dsk/c1t0d0s7
a   p lu0          16             8192        /dev/dsk/c1t1d0s7
a   p lu0          8208            8192        /dev/dsk/c1t1d0s7
a   p lu0          16400           8192        /dev/dsk/c1t1d0s7
```

4. Detach current metadb on failed disk (c1t1d0).

```
# metadb -d /dev/dsk/c1t1d0s7
# metadb
      flags          first blk      block count
a m p lu0          16             8192        /dev/dsk/c1t0d0s7
a   p lu0          8208            8192        /dev/dsk/c1t0d0s7
a   p lu0          16400           8192        /dev/dsk/c1t0d0s7
#
```

5. Current settings for bootpath is on the failed disk (c1t1d0).

```
# eeprom
ata-dma-enabled=1
atapi-cd-dma-enabled=0
ttyb-rtts-dtr-off=false
ttyb-ignore-cd=true
ttya-rtts-dtr-off=false
ttya-ignore-cd=true
ttyb-mode=9600,8,n,1,-
ttya-mode=9600,8,n,1,-
lba-access-ok=1
prealloc-chunk-size=0x2000
bootpath=/pci@0,0/pci15ad,1976@10/sd@1,0:a
keyboard-layout=US-English
console=text
```

6. This is how to view disk ID that will be used in setting up the bootpath.

```
# echo|format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
  0. clt0d0 <VMware,-VMware Virtual -1.0 cyl 1432 alt 2 hd 255 sec 63>
    /pci00,0/pci15ad,1976@10/sd00,0
  1. clt1d0 <VMware,-VMware Virtual -1.0 cyl 1432 alt 2 hd 255 sec 63>
    /pci00,0/pci15ad,1976@10/sd01,0
Specify disk (enter its number): Specify disk (enter its number):
```

7. Set up new bootpath to be bootable on another existing disk (c1t0d0).

```
# eeprom bootpath=/pci00,0/pci15ad,1976@10/sd00,0:a
# eeprom
ata-dma-enabled=1
atapi-cd-dma-enabled=0
ttyb-rts-dtr-off=false
ttyb-ignore-cd=true
ttya-rts-dtr-off=false
ttya-ignore-cd=true
ttyb-mode=9600,8,n,1,-
ttya-mode=9600,8,n,1,-
lba-access-ok=1
prealloc-chunk-size=0x2000
bootpath=/pci00,0/pci15ad,1976@10/sd00,0:a
keyboard-layout=US-English
console=text
```

8. Detach existing sub-mirrors that configured on the failed disk from current mirrored group. D22 is removed from d20.

```
# metadetach d20 d22
d20: submirror d22 is detached
# metastat -c
d20          m  517MB d21
   d21       s  517MB c1t0d0s1
d10          m  8.2GB d11
   d11       s  8.2GB c1t0d0s0
d22          s  517MB c1t1d0s1
d12          s  8.2GB c1t1d0s0
```

9. Remove sub-mirrors from the settings.

```
# metaclear d12
d12: Concat/Stripe is cleared
# metaclear d22
d22: Concat/Stripe is cleared
# metastat -c
d20          m  517MB d21
   d21       s  517MB c1t0d0s1
d10          m  8.2GB d11
   d11       s  8.2GB c1t0d0s0
```

10. Unconfigure the failed disk, you may skip this if the output was error because it might be the bug issue or existing failed disk having some issue that prevent it to be unconfigured (busy).

```
# cfgadm -c unconfigure c1::dsk/clt1d0
cfgadm: Hardware specific failure: failed to unconfigure SCSI device: I/O error
```

11. Replace new disk

12. Verify replaced disk and below example is **c1t2d0**.

```
# cfgadm -al
Ap_Id                Type      Receptacle  Occupant  Condition
c1                   scsi-bus  connected   configured unknown
c1::dsk/clt0d0       disk      connected   configured unknown
c1::dsk/clt1d0       disk      connected   configured unknown
c1::dsk/clt2d0       disk      connected   configured unknown
```

13. Verify and view new replaced disk information using “format” utility

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
  0. clt0d0 <VMware,-VMware Virtual -1.0 cyl 1432 alt 2 hd 255 sec 63>
    /pci0,0/pci15ad,1976@10/sd@0,0
  1. clt1d0 <VMware,-VMware Virtual -1.0 cyl 1432 alt 2 hd 255 sec 63>
    /pci0,0/pci15ad,1976@10/sd@1,0
  2. clt2d0 <VMware,-VMware Virtual -1.0 cyl 1433 alt 2 hd 255 sec 63>
    /pci0,0/pci15ad,1976@10/sd@2,0
Specify disk (enter its number): 2
```

14. Format new replaced disk using “fdisk” tool

```
FORMAT MENU:
  disk      - select a disk
  type      - select (define) a disk type
  partition - select (define) a partition table
  current   - describe the current disk
  format    - format and analyze the disk
  fdisk     - run the fdisk program
  repair    - repair a defective sector
  label     - write label to the disk
  analyze   - surface analysis
  defect    - defect list management
  backup    - search for backup labels
  verify    - read and display labels
  save      - save new disk/partition definitions
  inquiry   - show vendor, product and revision
  volname   - set 8-character volume name
  !<cmd>    - execute <cmd>, then return
  quit

format> fdisk
```

15. Accept to format as SOLARIS system partition by issue keyword "y".

```
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
format>
```

16. Save this new settings by label command.

```
format> label
Ready to label disk, continue? y

format> verify
Warning: Primary label on disk appears to be different from
current label.

Warning: Check the current partitioning and 'label' the disk or use the
'backup' command.

Primary label contents:

Volume name = <          >
ascii name = <VMware,-VMware Virtual -1.0 cyl 1432 alt 2 hd 255 sec 63>
pcyl       = 1434
ncyl       = 1432
acyl       = 2
bcyl       = 0
nhead     = 255
nsect     = 63
Part      Tag      Flag      Cylinders      Size      Blocks
0 unassigned  wm        0              0          (0/0/0)      0
1 unassigned  wm        0              0          (0/0/0)      0
2 backup     wu        0 - 1431      10.97GB     (1432/0/0) 23005080
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         7.84MB     (1/0/0)      16065
9 unassigned  wm        0              0          (0/0/0)      0

format> q
#
```

17. Copy existing partition table on old disk unit (c1t0d0) into new replaced disk.

```
# prtvtoc /dev/rdsk/c1t0d0s2 | fmthard -s - /dev/rdsk/c1t2d0s2
fmthard: New volume table of contents now in place.
#
```

## 18. Verify replaced disk having the same settings in partition table compared to old unit disk.

```
format> verify
Primary label contents:
Volume name = <          >
ascii name = <VMware,-VMware Virtual -1.0  cyl 1432 alt 2 hd 255 sec 63>
pcyl      = 1434
ncyl      = 1432
acyl      = 2
bcyl      = 0
nhead     = 255
nsect     = 63
Part      Tag      Flag      Cylinders      Size      Blocks
0         root      wm        359 - 1431     8.22GB    (1073/0/0) 17237745
1         swap      wu         1 - 66         517.72MB  (66/0/0)   1060290
2         backup   wm         0 - 1431     10.97GB   (1432/0/0) 23005080
3         unassigned wu         0              0         (0/0/0)    0
4         unassigned wu         0              0         (0/0/0)    0
5         unassigned wu         0              0         (0/0/0)    0
6         unassigned wu         0              0         (0/0/0)    0
7         home      wm         67 - 358      2.24GB    (292/0/0)  4690980
8         boot      wu         0 - 0          7.84MB    (1/0/0)    16065
9         unassigned wu         0              0         (0/0/0)    0
format> q
#
```

## 19. Create metadb on the replaced disk with same slice number (s7)

```
# metadb -a -f -c3 /dev/dsk/c1t2d0s7
# metadb
      flags      first blk      block count
a m p lu0      16             8192          /dev/dsk/c1t0d0s7
a p lu0      8208           8192          /dev/dsk/c1t0d0s7
a p lu0     16400          8192          /dev/dsk/c1t0d0s7
a u          16             8192          /dev/dsk/c1t2d0s7
a u          8208           8192          /dev/dsk/c1t2d0s7
a u          16400          8192          /dev/dsk/c1t2d0s7
#
```

20. Set up new meta devices on replaced disk with proper sub-mirrors d12 and d22. Verify it using `metastat -c` to display mirror group settings.

```
# metastat -c
d20          m  517MB d21
   d21       s  517MB c1t0d0s1
d10          m  8.2GB d11
   d11       s  8.2GB c1t0d0s0
# metainit -f d12 1 1 c1t2d0s0
d12: Concat/Stripe is setup
# metattach d10 d12
d10: submirror d12 is attached
# metastat -c
d20          m  517MB d21
   d21       s  517MB c1t0d0s1
d10          m  8.2GB d11 d12 (resync-0%)
   d11       s  8.2GB c1t0d0s0
   d12       s  8.2GB c1t2d0s0
# metainit -f d22 1 1 c1t2d0s1
d22: Concat/Stripe is setup
# metattach d20 d22
d20: submirror d22 is attached
# metastat -c
d20          m  517MB d21 d22
   d21       s  517MB c1t0d0s1
   d22       s  517MB c1t2d0s1
d10          m  8.2GB d11 d12 (resync-2%)
   d11       s  8.2GB c1t0d0s0
   d12       s  8.2GB c1t2d0s0
#
```

21. Install grub boot settings into new replaced disk (c1t2d0) so this replaced disk will be bootable.

```
# /sbin/installgrub /boot/grub/stage1 /boot/grub/stage2 /dev/rdsd/c1t2d0s0
stage1 written to partition 0 sector 0 (abs 16065)
stage2 written to partition 0, 274 sectors starting at 50 (abs 16115)
#
```

22. Verify mirror re-syncing and that's mean the data in old disk unit is now copying over into new replaced disk.

```
# metastat -c
d20          m  517MB d21 d22
   d21       s  517MB c1t0d0s1
   d22       s  517MB c1t2d0s1
d10          m  8.2GB d11 d12 (resync-26%)
   d11       s  8.2GB c1t0d0s0
   d12       s  8.2GB c1t2d0s0
#
```

23. The data syncing is verified to be finished when the “resync” is disappeared.

```
# metastat -c
d20          m  517MB d21 d22
  d21        s  517MB c1t0d0s1
  d22        s  517MB c1t2d0s1
d10          m  8.2GB d11 d12
  d11        s  8.2GB c1t0d0s0
  d12        s  8.2GB c1t2d0s0
# █
```

## Reference

**Solaris Volume Manager (SVM) SPARC How to Replace a Failed, SCSI Disk, Mirrored with SVM** (Doc ID 1469821.1) (JANUARY 24, 2017). Available online at:

[https://support.oracle.com/knowledge/Sun%20Microsystems/1469821\\_1.html](https://support.oracle.com/knowledge/Sun%20Microsystems/1469821_1.html)