

How to replace a failed SCSI disk mirrored with Solstice DiskSuite (command line method) WITHOUT taking Solaris down:

Last document revision: 10/21/02

Author: Tim Lavin

The following assumptions have been made for this example procedure:

- A. Server name= bigserver.company.com
- B. OS = Solaris 7 or 8
- C. DiskSuite versions 4.2 or 4.2.1
- D. The 'bad' disk is hot-swappable.
- E. Metadevice numbers and state database replicas are taken from those used on the actual system.
- F. Actual command line syntax is shown in bold text.

1. Collect output from the following:

```
# metastat  
# metastat -p  
# metadb -i
```

2. To identify the disk to be replaced:

Examine the "metadb -i" output. You should see a "W" in the flags field associated with slice 7 of the disk experiencing write errors. Another indication is to look at the output from the "format" command. Next to the device name the text string "<drive type unknown>" will appear indicating that the disk label cannot be read and therefore it is very likely that the disk has failed. For this example, we will assume the failed disk device is c0t0d0.

3. Delete any metadevice state database replicas that are on the 'bad' disk:

```
# metadb -d c0t0d0s7  
# metadb -i (to make sure they have been deleted)
```

4. State of the submirrors:

The "metastat" command output reports that all submirrors on the bad disk are at a State of "Needs maintenance". This indicates that DiskSuite has automatically disabled the submirrors, so there is no need to "metadetach" the submirrors.

5. Physically replace the failed hot-swappable disk.

6. Partition the new disk:

Easiest way to do this is to copy the partition table from the root mirror (c0t1d0s2) to the new disk (c0t0d0s2) with the following dd command:

```
# dd if=/dev/rdisk/c0t1d0s2 of=/dev/rdisk/c0t0d0s2 count=16
```

Verify the partition table was copied correctly using the format utility. Type "**format**", select the corresponding disk number from the disk selection menu, then type "p", then "p" again to view the partition table. Compare and make sure the partition tables match EXACTLY.

7. Recreate the metadevice state database replicas that were deleted in step 3 from c0t0d0s7:

```
# metadb -a -f -c 3 /dev/dsk/c0t0d0s7
```

```
# metadb -i (verify the creation)
```

8. Re-enable the submirrors:

```
# metareplace -e d1 c0t0d0s0 (d1 is / and c0t0d0s0 is device associated w/ d11 submirror)
```

```
# metareplace -e d2 c0t0d0s1 (d2 is /usr and c0t0d0s1 is device associated w/ d12 submirror)
```

```
# metareplace -e d3 c0t0d0s3 (d3 is /var and c0t0d0s3 is device associated w/ d13 submirror)
```

```
# metareplace -e d5 c0t0d0s5 (d5 is /opt and c0t0d0s5 is device associated w/ d15 submirror)
```

Note: Run "**metastat | grep sync**" to check syncing status. When all syncing is complete, run "**metastat**" again to verify that the "State" of all submirrors is "Okay".

NOTE: If your system does NOT have hot-swappable drives, see the following DiskSuite 4.2 documentation at the docs.sun.com site:

<http://docs.sun.com/?q=Replacing+a+failed+disk&p=/doc/805-5961/6j5kfk4c&a=view>