



white papers



# **Mirroring the Boot Drive with Solstice Disk Suite/Solaris Volume Manager**

**Brad Isbell, Seeds of Genius  
June 2006**

**Revision 1.1 6/04/2006  
Edition: June 2006**



**All rights reserved.**

**This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation.**

**No part of this product or document may be reproduced in any form by any means without prior written authorization of Seeds of Genius and its licensors, if any.**

**Abstract:**

To ensure the highest uptime possible for the Sun servers, it is important to mirror the boot disk drive. This document was written as a quick reference guide to perform this essential administrative task. Since many Seeds of Genius customer have purchased the SunFire V240 system this will be the reference platform for this article. The same principles apply to mirroring the boot drive of any host running the Solaris OE.

**Requirements:**

If using Solaris 9 or 10, the mirroring software is included within the operating systems. If using Solaris 8 you will need to have the Solaris 8 DVD or CD Software 2 of 2.

**Procedures:**

We will be creating a RAID 1 (mirror) using the first 2 disks on a SunFire V240 system. A SunFire V240 will have the following logical device files for the four internal disk drives.

```
/dev/dsk/c1t0d0s2    <==  Boot drive
/dev/dsk/c1t1d0s2    <==  Disk drive to mirror boot drive
/dev/dsk/c1t2d0s2
/dev/dsk/c1t3d0s2
```

The third and fourth disk can be used to hold extra replicas of the metadvice state database. This is explained further in step 4 below.

For the sake of example we will assume the following partitioning on the boot drive. Please refer to your **/etc/vfstab** after installation to confirm your settings.

```
c1t0d0s0  /
c1t0d0s1  swap
c1t0d0s3  /usr
```

1. **Installing Solaris Operating Environment.** Install Solaris to preferred specifications following these rules:

- During the installation make sure to leave a small partition to create the metadvice state database on. This is a slice of disk that will hold all the metadata for all RAID configurations. Generally this will be done on slice 7, 50MB will be plenty of space.

During installation you can create a 50MB filesystem mounted to **/META** on slice 7. Once the install completes you can *umount /META* and then modify **/etc/vfstab** to remove this entry for future system reboots. To keep the / directory clean, remove the **/META** directory.

- Install the entire distribution with OEM support.

2. **Install Solstice Disk Suite on Solaris 8 systems.** If you are using Solaris 8 you will need to install Disk Suite. If you have Solaris 9 or Solaris 10 installed you can skip to step 3.

Insert the Solaris 8 DVD after the installation is complete and perform the following.

```
# cd
/cdrom/cdrom0/s0/Solaris_8/EA/products/DiskSuite_4.2.1/sparc/
Packages
# pkadd -d.
```

Answer yes to all the questions and then reboot.

3. **Partition second disk drive to match the partitioning of the boot Drive.**

```
# prtvtoc /dev/rdisk/c1t0d0s2 > /tmp/vtoc
# fmthard -s /tmp/vtoc /dev/rdisk/c1t1d0s2
```

4. **Create Meta-State Databases.** It is generally good practice to create several replicas of the metadvice state databas for each slice. In the following example 3 state database replicas will be created on each disk's slice 7.

```
# metadb -a -f -c3 c1t0d0s7 c1t1d0s7
```

Confirm that the state databases have been created.

```
# metadb
```

Due to the rule that you must have over half (half +1) of the state databases accessible to boot the system. You will want to create a small slice on a third drive for another set of three state database replicas. If you lose one of these two disks all you will have is half of the replicas available, so the system will not be quickly booted without further administration.

For example, you can create a small slice at c1t2d0s7 and then create more state databases.

```
# metadb -a -c3 c1t2d0s7
```

5. **Create logical devices for mirroring.** Under Solstice Disk Suite/Solaris Volume Manager you can only mirror striped logical devices. You can not create a mirror using physical devices. So you will first have to setup each slice as a logical stripe of 1 unit. You will then mirror these logical devices.

Create a logical stripe of 1 for each slice of disk that will be a part of each mirror. You must use the -f (force) option for the slices that contain

mounted file systems.

```
# metainit -f d10 1 1 c1t0d0s0
# metainit -f d11 1 1 c1t0d0s1
# metainit -f d13 1 1 c1t0d0s3
```

Create the logical stripes on the second drive.

```
# metainit d20 1 1 c1t1d0s0
# metainit d21 1 1 c1t1d0s1
# metainit d23 1 1 c1t1d0s3
```

Confirm the creation of these devices.

```
# metastat
```

6. **Create initial mirrors.** This step will only set up one half of each mirror. The second half will be attached in the next step.

```
# metainit d0 -m d10
# metaroot d0
```

The metaroot command will modify the **/etc/vfstab** file and the **/etc/system** file for you so that you will boot from the **/dev/md/dsk/d0** device.

Create Mirrors for **swap** and **/usr**.

```
# metainit d1 -m d11
# metainit d3 -m d13
```

Modify the **/etc/vfstab** file to reflect these new meta-devices. The entries should look something like the following.

```
/dev/md/dsk/d1 - - swap - no -
/dev/md/dsk/d3 /dev/md/rdisk/d3 /usr ufs 1 no -
```

Reboot in order to get the **/** and **/usr** file systems to be remounted as the md (metadisk) devices

```
# init 6
```

After the reboot check to make sure that the filesystems are mounted as **/dev/md/dsk** devices.

```
# df -k
```

7. **Attach second half of mirrors.**

```
# metattach d0 d20
# metattach d1 d21
# metattach d3 d23
```

Run the metastat command to track the progress of the syncing of the mirrors.

```
# metastat
```

8. **Setup system PROM.** You will need to modify the system PROM in order to boot from the second half of the mirror if the first half fails.

```
# eeprom boot-device="disk disk1 net"
```

9. **TEST.** Wait until the mirrors are finished syncing.

```
# reboot -- "disk1"
```

**If you need any help please contact the Seeds of Genius support center at [support@seedsofgenius.com](mailto:support@seedsofgenius.com).**