

Preparing the RAID Module as a Boot Device



© 1999 Sun Microsystems, Inc.
901 San Antonio Road, Palo Alto, California 94303 U.S.A

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

TRADEMARKS

Sun, Sun Microsystems, the Sun logo, Sun StorEdge, Starfire, Ultra, Ultra Enterprise, Sun Enterprise, Solaris, and Solstice DiskSuite are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. SUN MICROSYSTEMS, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS PUBLICATION AT ANY TIME.



Please
Recycle

Contents

Supported Platforms and Solaris Operating Environment Versions	1
To Make Your HW RAID Module Bootable.....	2
Setting Up and Verifying Alternate Boot Paths.....	5
Last-Minute Notes	6
RM6 6.22 Bootability with Sun StorEdge A3x00s	9

Preparing the RAID Module as a Boot Device



If you want to use the RAID Module as a boot device (that is, install the Solaris™ Operating Environment and the RM6 software on the RAID Module), you must prepare the module to be a boot device before installing any software.

Examples of when you would want to use your RAID Module as your boot device:

- New Installation
- You want HW RAID Protection on the Solaris Operating Environment, instead of doing it through software-based RAID.

Supported Platforms and Solaris Operating Environment Versions

- Sun StorEdge™ A1000:
 - Ultra™ 1, Ultra 10, Sun Enterprise™ x500, Sun Enterprise 450 with Solaris Operating Environment 2.6HW3 or 2.6 5/98;
 - Sun Enterprise 250, Ultra 60 with Solaris Operating Environment 2.5.1 SWHP
- Sun StorEdge A3000:
 - Ultra Enterprise 4000, Ultra Enterprise 5000, Ultra Enterprise 6000, Sun Enterprise 450 with Solaris Operating Environment 2.6HW3 or 2.6 5/98



- Sun StorEdge A3500:
Ultra Enterprise™ 4500, Ultra Enterprise 5500, Ultra Enterprise 6500,
Sun Enterprise 450 with Solaris Operating Environment 2.6HW3 or
2.6 5/98
- Sun StorEdge A3500: Sun Enterprise 10000 with Solaris Operating
Environment 2.5.1 HW3 Server Edition or 2.6 5/98

Note – For the Sun Enterprise 10000, this also includes the 1x2 Sonoma Lite.

Cautions

- The Solaris Operating Environment requires the boot device to be LUN 0. For the best results, make sure LUN 0 uses RAID Level 1, 3 or 5 to enable data protection in the drive group.
- If you use the RAID Module as a boot device, you must never delete LUN 0 or reset the configuration — you will lose your boot device completely.
- Using RAID 3 or RAID 5 in setting up your RAID Module as a boot device requires a minimum of three drives. RAID 1 requires a minimum of two drives and can be incremented in two's.
- Using RAID 0 in setting up your boot device is not supported.
- In the Independent Controller Configuration, only one host system can boot because only one controller owns LUN 0.
- If you make a RAID Module boot device that has a single controller (Sun StorEdge A1000), dual controllers on the same SCSI bus, or the Independent Controller Configuration, you do not have RDAC failover protection in the event a controller fails or has path problems. Any data path failure problem with the boot-path controller could result in the loss of your boot device.
- Never place your boot RAID Module boot device under Veritas VM or Solstice DiskSuite™ control. Doing so may corrupt the Solaris Operating Environment and require it to be reinstalled.

To Make Your HW RAID Module Bootable

- 1. Backup *all* data on your HW RAID Module before beginning procedure.**
- 2. Install LUN 0 on your HW RAID Device. If this is a new installation, you might want to make sure that your default LUN 0 from the factory is the size that you want before proceeding.**

3. **Boot cdrom or install the Solaris Operating Environment through JumpStart onto LUN 0 on your HW RAID device. Let the Solaris Operating Environment installation program set your eeprom to boot off your RAID Module. After OS installation, let it reboot off your RAID Module. The OS install includes any and all patches for RM6 6.1.1 Update 1.**
4. **Install RM6 6.1.1 Update 1.**
 - a. **Patch 106513-xx**
 - b. **Patch 106552-xx**
5. **Edit the /usr/lib/osa/rmparams file and make the variable Rdac_SupportDisabled TRUE.**
6. **Boot -r.**
7. **Edit the rmparams file again and make Rdac_SupportDisabled FALSE.**
8. **Run the command /etc/init.d/rdacctrl config.**
9. **Edit the /etc/system file and add the following entry:**

```
rootdev: /pseudo/rdnexus@0/rdriver@4,0:a
```

The rdnexus and rdriver numbers are based on an entry in the /kernel/drv/rdriver.conf file. For example:

```
name="rdriver" module=1 lun=0 target=4 parent="/pseudo/rdnexus@0"
dev_a=0x800028 dev_b=0x800188;
```

Look at the "target" number for the rdriver number.

For systems with more than one RAID device, the correct module should be the *first* instance of lun=0, target=5 from the bottom of the file. In that line, you should see the correct rdnexus@<n> number. In the example code below, the device is booting from lun=0, target=5, and rdnexus@2.

Pay close attention to the rdnexus@<n> value and target value. Using the wrong values could result in a crashed system.

```
name="rdriver" module=1 lun=0 target=5 parent="/pseudo/rdnexus@3"
dev_a=0x800190 dev_b=0x800020;
name="rdriver" module=1 lun=1 target=4 parent="/pseudo/rdnexus@0"
dev_a=0x800fa0 dev_b=0x800590;
name="rdriver" module=1 lun=2 target=4 parent="/pseudo/rdnexus@0" unbound=1;
```



```
name="rdriver" module=1 lun=2 target=5 parent="/pseudo/rdnexus@3" unbound=1;
name="rdriver" module=1 lun=3 target=4 parent="/pseudo/rdnexus@0" unbound=1;
name="rdriver" module=1 lun=3 target=5 parent="/pseudo/rdnexus@3" unbound=1;
name="rdriver" module=1 lun=4 target=4 parent="/pseudo/rdnexus@0" unbound=1;
name="rdriver" module=1 lun=4 target=5 parent="/pseudo/rdnexus@3" unbound=1;
name="rdriver" module=1 lun=5 target=4 parent="/pseudo/rdnexus@0" unbound=1;
name="rdriver" module=1 lun=5 target=5 parent="/pseudo/rdnexus@3" unbound=1;
name="rdriver" module=1 lun=6 target=4 parent="/pseudo/rdnexus@0" unbound=1;
name="rdriver" module=1 lun=6 target=5 parent="/pseudo/rdnexus@3" unbound=1;
name="rdriver" module=1 lun=7 target=4 parent="/pseudo/rdnexus@0" unbound=1;
name="rdriver" module=1 lun=7 target=5 parent="/pseudo/rdnexus@3" unbound=1;
Correct module: name="rdriver" module=2 lun=0 target=5 parent="/pseudo/rdnexus@2"
                 dev_a=0x800118 dev_b=0xffffffff;
name="rdriver" module=2 lun=1 target=5 parent="/pseudo/rdnexus@2"
                 dev_a=0x800c58 dev_b=0xffffffff;
name="rdriver" module=2 lun=2 target=5 parent="/pseudo/rdnexus@2"
                 dev_a=0x800c60 dev_b=0xffffffff;
name="rdriver" module=2 lun=3 target=5 parent="/pseudo/rdnexus@2"
                 dev_a=0x800c68 dev_b=0xffffffff;
name="rdriver" module=2 lun=4 target=5 parent="/pseudo/rdnexus@2"
                 dev_a=0x800c70 dev_b=0xffffffff;
name="rdriver" module=2 lun=5 target=5 parent="/pseudo/rdnexus@2"
                 dev_a=0x800c78 dev_b=0xffffffff;
name="rdriver" module=2 lun=6 target=5 parent="/pseudo/rdnexus@2"
                 dev_a=0x800c80 dev_b=0xffffffff;
name="rdriver" module=2 lun=7 target=5 parent="/pseudo/rdnexus@2"
                 dev_a=0x800c88 dev_b=0xffffffff;
name="rdriver" class="scsi" target=0 lun=0
name="rdriver" class="scsi" target=0 lun=1
```

10. Boot -r.

You are all set to boot off your HW RAID Device.

When you run the rdacctl config command, you should see the following:



```
area51-29# /etc/init.d/rdacctrl config
Array Monitor stopped
Re-generating rdriver.conf file ...
The NVSRAM settings of controller c0t5d0s0(1T82018106) are correct.
The NVSRAM settings of controller c1t4d0s0(1T70110025) are correct.
The NVSRAM settings of controller c4t4d0s0(1T81410920) are correct.
The NVSRAM settings of controller c5t5d0s0(1T65172824) are correct.
nvutil command succeeded.
```

```
Array Monitor initiated
RDAC daemons initiated
```

What is happening is the RDAC is regenerating its configuration files.

Setting Up and Verifying Alternate Boot Paths

- 1. Do a probe-scsi-all, there should be a SCSI ID listed for each RAID Controller on the host. Record this information down, including the full device path.**
- 2. Boot up off your RAID Module and then bring up RM6.**
- 3. Double click on the Recovery icon and select the RAID Module that is your boot device and verify that your RAID Module is in Optimal Mode.**
- 4. Select Options -> Manual Recovery -> Controllers Pairs.**
- 5. Highlight the controller that is the default boot device and select Place Offline.**
- 6. When the controller is offline, select Recovery Guru. You should see a “data path failure” for the module. Do not fix the problem at this time.**
- 7. Select Module Profile and confirm all LUNs are now on the alternate controller.**
- 8. Bring your host down to the “ok” prompt.**
- 9. Based on the information from the probe-scsi-all, use OBP’s nvalias command to create a alias to boot off of. Once this is completed, go to the next step.**



10. Boot alias -r.

11. Once the host is up, bring up RM6. Select the Recovery icon, and select the Module that is your default boot device. You should get a data path failure. Select Fix and follow the instructions provided. When this is done, run the Recovery Guru again and make sure your RAID Module that you are booting off of is in Optimal Mode.

12. Double click the Maintenance/Tuning icon. Select the RAID Module that is your boot device and select LUN balancing. Verify that your default boot path owns LUN 0.

Now you have your primary and alternate boot paths set up. Note that this is only for the Sun StorEdge A3000 and A3500. The Sun StorEdge A1000 and Independent Controller Configuration do not have alternate paths to boot off of.

Note – All HW RAID controllers used for boot device must have firmware 2.5.2.11 and RM6 6.1.1 Update 1.

Last-Minute Notes

- During the boot process, if the HW RAID Controller that contains LUN 0 is missing or damaged, the system will not boot. Replace this controller and the system can boot off the HW RAID controllers.
- In your LUN 0 disk set, the first disk that contains the beginning of the boot blocks is pulled or failed during the early portion of the boot process. the system may hang or fail to boot. This occurred on the Sun StorEdge A1000 with the Ultra 10 in one of the test cases.
- Before doing the probe-scsi-all to get your device paths, it is recommended that you set auto-boot to false and then do a reset-all before you do a probe-scsi-all. Otherwise you risk hanging your host from OBP.
- When doing an initial boot off of an Ultra10, it will fail with a “trap 3”. Refer to bugids 4166678 and 1251360. The workaround is to boot again after getting this error message.
- When booting a Sun StorEdge A3x00/A1000 off of a Sun Enterprise 450, you see the following messages:

```
Fatal SCSI error  at script address 258 Unexpected disconnect
Drive not ready
Fatal SCSI error  at script address 258 Unexpected disconnect
```



```
Fatal SCSI error at script address 258 Unexpected disconnect
Boot device: /pci@4,4000/scsi@2,1/disk@4,0:a File and args:
Fatal SCSI error at script address 258 Unexpected disconnect
Fatal SCSI error at script address 258 Unexpected disconnect
Drive not ready
Fatal SCSI error at script address 258 Unexpected disconnect
Fatal SCSI error at script address 258 Unexpected disconnect
Fatal SCSI error at script address 258 Unexpected disconnect
Fatal SCSI error at script address 258 Unexpected disconnect
Drive not ready
Fatal SCSI error at script address 258 Unexpected disconnect
Fatal SCSI error at script address 258 Unexpected disconnect
Fatal SCSI error at script address 258 Unexpected disconnect
Fatal SCSI error at script address 258 Unexpected disconnect
Drive not ready
Fatal SCSI error at script address 258 Unexpected disconnect
Fatal SCSI error at script address 258 Unexpected disconnect
```

```
Fatal SCSI error at script address 258 Unexpected disconnect
Fatal SCSI error at script address 258 Unexpected disconnect
Drive not ready
Fatal SCSI error at script address 258 Unexpected disconnect
Fatal SCSI error at script address 258 Unexpected disconnect
Fatal SCSI error at script address 258 Unexpected disconnect
Fatal SCSI error at script address 258 Unexpected disconnect
Drive not ready
Fatal SCSI error at script address 258 Unexpected disconnect
Fatal SCSI error at script address 258 Unexpected disconnect
```

What is happening in this case is that the Sun StorEdge A1000 being booted off of has the only disk that is going to respond to the host this early in the boot process. It doesn't matter if you had any LUNs created or not on the HW RAID Module. You will see these messages when you boot off of this PCI host and it will not cause any problems with your configuration.



- When you do a probe-scsi-all on your host and you have any of the HW RAID Modules attached, you see something as shown below:

```
ok probe-scsi-all
/pci@4,4000/scsi@2,1
Target 4
Unit 0  Disk      SYMBIOS RSM Array 2000  0205
Unit 1  Device type 20      SYMBIOS RSM Array 2000  0205
Unit 2  Device type 20      SYMBIOS RSM Array 2000  0205
Unit 3  Device type 20      SYMBIOS RSM Array 2000  0205
Unit 4  Device type 20      SYMBIOS RSM Array 2000  0205
Unit 5  Device type 20      SYMBIOS RSM Array 2000  0205
Unit 6  Device type 20      SYMBIOS RSM Array 2000  0205
Unit 7  Device type 20      SYMBIOS RSM Array 2000  0205
```

```
/pci@4,4000/scsi@2
```

```
/pci@1f,4000/scsi@4,1
```

```
Target 5
Unit 0  Disk      SYMBIOS RSM Array 2000  0205
Unit 1  Device type 20      SYMBIOS RSM Array 2000  0205
Unit 2  Device type 20      SYMBIOS RSM Array 2000  0205
Unit 3  Device type 20      SYMBIOS RSM Array 2000  0205
Unit 4  Device type 20      SYMBIOS RSM Array 2000  0205
Unit 5  Device type 20      SYMBIOS RSM Array 2000  0205
Unit 6  Device type 20      SYMBIOS RSM Array 2000  0205
Unit 7  Device type 20      SYMBIOS RSM Array 2000  0205
```

What is being shown in the above example is that your HW RAID Module has one configured LUN:

```
Unit 0  Disk      SYMBIOS RSM Array 2000  0205
```

The following lines means that there is no configured LUN and there is no problem with the host or HW RAID Module:

```
Unit 1  Device type 20      SYMBIOS RSM Array 2000  0205
```

```
Unit 2 Device type 20 SYMBIOS RSM Array 2000 0205
Unit 3 Device type 20 SYMBIOS RSM Array 2000 0205
```

Note if you had a Sun StorEdge™ A3500 and you did a probe-scsi-all or a format, the giveaway that you had a A3500 in your configuration is shown below:

- c1t0d0 <Symbios-StorEDGEA3500-9300 cyl 4319 alt 2 hd 64 sec 64>/pseudo/rdnexus@1/rdriver@0,0

For the Sun StorEdge A1000, it would be the same as above, except for the 1000 in place of the 3500.

- When doing the ./rdacctrl config command, do not interrupt this process. You can possibly cause problems with booting off your HW RAID Module.
- When doing the ./rdacctrl config command on the Ultra 1, you will see the following:

```
dilbert35# /etc/init.d/rdacctrl config
        Array Monitor stopped
        RDAC daemons stopped
        Re-generating rdriver.conf file ...
        Unable to unload RDAC driver - reboot to make new
configuration take effect
        The NVSRAM settings of controller c0t0d0s0(1T73235624)
are correct
        nvutil command succeeded.

        Array Monitor initiated
        RDAC daemons initiated
```

The “Unable to unload RDAC driver” are normal and there’s no problem.

RM6 6.22 Bootability with Sun StorEdge A3x00s

The current procedure for booting off your Sun StorEdge A3x00 does not work for the next generation RAID Manager software 6.22. Follow the steps below to boot with RM6 6.22.



The supported platforms and versions of the Solaris Operating Environment for HW RAID Bootability are:

- Sun StorEdge A3000:
Ultra Enterprise 4000, Ultra Enterprise 5000, Ultra Enterprise 6000,
Sun Enterprise 450 with Solaris Operating Environment 2.6 5/98 or
Solaris 7 Operating Environment
- Sun StorEdge A3500:
Ultra Enterprise 4500, Ultra Enterprise 5500, Ultra Enterprise 6500
with Solaris Operating Environment 2.6 5/98 or Solaris 7 Operating
Environment
- Sun StorEdge A3500:
Sun Enterprise 10000 with Solaris Operating Environment 2.6 5/98

Note – The supported platforms for PACS/GA are:

PACS: Starfire™ server; Sunfire™ server

GA: Starfire server; Sunfire server; Sun Enterprise 450

You need an Optimal LUN 0, which may be either a RAID 1, 3, or 5. A RAID 0 LUN for bootability is not supported.

To boot with RM6 6.22:

- 1. Install OS (let suninstall set your eeprom to the boot device).**
Select "autoboot" at the end of suninstall.
- 2. Install any patch clusters and/or RM6 patches.**
 - a. Install RM6 6.22**
 - b. Edit the rmparms file for 16 lun support, if needed. Then run "genscsiconf".**
 - c. Edit sd.conf, if needed. This is to help speed up reboots. Otherwise on reboots, the host will timeout for every non-existent LUN.**
 - d. Run /etc/init.d/rdacctrl config.**
 - e. Edit the /etc/system file and add the rootdev entry. Refer to step 9 of "To Make Your HW RAID Module Bootable" on page 3.**
 - f. Boot -r.**

The initial phase of bootability is only for initial installs due to the following outstanding issues:

- 4240638 -> Solaris 7 Operating Environment 32-bit upgrade to Solaris 7 Operating Environment 64-bit doesn't work
- 4240577 -> Solaris 7 Operating Environment suninstall doesn't work when upgrading from 2.6 5/98
- 4240583 -> OBP, probe-fcal-all fails to see any FCAL A3x00 LUNs and because of this, you cannot boot off of FC A3x00

Notes

When you run the `rdacctrl` command, you see one of two informational messages:

```
area51-44# ./rdacctrl config
drvconfig: Driver (ssd) failed to attach
RDAC daemons initiated
ID[RAIDarray.rdaemon.1001] RDAC Resolution Daemon locked in memory
WARNING: set_drivers: No arrays were found that are controlled by
either the sd or the ssd driver
area51-44#
```

Or:

```
area51-44# /etc/init.d/rdacctrl config
RDAC daemons initiated
ID[RAIDarray.rdaemon.1001] RDAC Resolution Daemon locked in memory
area51-44#
```





Sun Microsystems Computer Company
A Sun Microsystems, Inc. Business
901 San Antonio Road
Palo Alto, CA 94303 USA
650 960-1300
FAX 650 969-9131
<http://www.sun.com>

Sales Offices

Argentina: +54-1-317-5600
Australia: +61-2-9844-5000
Austria: +43-1-60563-0
Belgium: +32-2-716-7911
Brazil: +55-11-5181-8988
Canada: +905-477-6745
Chile: +56-2-638-6364
Colombia: +571-622-1717
Commonwealth of Independent States:
+7-502-935-8411
Czech/Slovak Republics:
+42-2-205-102-33
Denmark: +45-44-89-49-89
Estonia: +372-6-308-900
Finland: +358-9-525-561
France: +33-01-30-67-50-00
Germany: +49-89-46008-0
Greece: +30-1-680-6676
Hong Kong: +852-2802-4188
Hungary: +36-1-202-4415
Iceland: +354-563-3010
India: +91-80-559-9595
Ireland: +353-1-8055-666
Israel: +972-9-956-9250
Italy: +39-39-60551
Japan: +81-3-5717-5000
Korea: +822-3469-0114
Latin America/Caribbean:
+1-650-688-9464
Latvia: +371-755-11-33
Lithuania: +370-729-8468
Luxembourg: +352-491-1331
Malaysia: +603-264-9988
Mexico: +52-5-258-6100
Netherlands: +31-33-450-1234
New Zealand: +64-4-499-2344
Norway: +47-2218-5800
People's Republic of China:
Beijing: +86-10-6849-2828
Chengdu: +86-28-678-0121
Guangzhou: +86-20-8777-9913
Shanghai: +86-21-6247-4068
Poland: +48-22-658-4535
Portugal: +351-1-412-7710
Russia: +7-502-935-8411
Singapore: +65-438-1888
South Africa: +2711-805-4305
Spain: +34-1-596-9900
Sweden: +46-8-623-90-00
Switzerland: +41-1-825-7111
Taiwan: +886-2-514-0567
Thailand: +662-636-1555
Turkey: +90-212-236 3300
United Arab Emirates:
+971-4-366-333
United Kingdom: +44-1-276-20444
United States: +1-800-821-4643
Venezuela: +58-2-286-1044
Worldwide Headquarters:
+1-650-960-1300