



Sun StorEdge™ RAID Manager 6.22 and 6.22.1 Upgrade Guide

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 806-7792-14
April 2002, [Revision A](#)

Send comments about this document to: docfeedback@sun.com

Copyright 2002 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303-4900 U.S.A. All rights reserved.

This product or document is 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 Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Solaris, Sun StorEdge, SunVTS, NetBackup, and Ultra Enterprise are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2002 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303-4900 Etats-Unis. Tous droits réservés.

Ce produit ou document est distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Solaris, Sun StorEdge, SunVTS, NetBackup, et Ultra Enterprise sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Contents

1. Contents iii

Preface v

Using UNIX Commands v

Typographic Conventions vi

Shell Prompts vi

Accessing Sun Documentation Online vii

Sun Welcomes Your Comments vii

2. Upgrade Procedures 9

Notice, Restrictions and Prerequisites 10

 Important Notice 10

 Upgrade Restrictions 11

 Upgrade Conditions 12

Preparing for the Upgrade 13

 ▼ To Prepare the Cluster Nodes 13

 ▼ To Obtain the Necessary Software Patches 13

 ▼ To Complete the Upgrade Preparation 15

Setting the System Up For the Upgrade 17

 ▼ To Save Important Files 17

 ▼ To Verify That RAID Modules Are In an Optimal State 18

- ▼ To Save the System State 19
- Removing the System From Service 20
 - ▼ To Stop All Host I/O 20
 - ▼ To Halt RAID Manager 6 Parity Checking 21
 - ▼ To Enable All RAID Controllers For Hard Reset 22
- Upgrading the Host Software 24
 - ▼ To Remove Existing RAID Manager 6.x Patches and Packages 24
 - ▼ To Reboot 26
 - ▼ To Install RAID Manager 6.22 or 6.22.1 and Patches 27
 - ▼ To Edit the `/etc/osa/rmparams` File 28
 - ▼ To Reboot 29
- Upgrading the Module Firmware 31
 - ▼ To Ensure Volumes Are Free After Reboot 31
 - ▼ To Verify Controllers Are Ready For the Firmware Upgrade 31
 - ▼ To Upgrade the Firmware 32
 - ▼ To Restore All RAID Controllers to Soft Reset 34
 - ▼ To Commit the Changes 35
 - ▼ To Verify System Health 36
 - ▼ To Resolve Device Renumbering 37
- Returning the System to Service 38
 - ▼ To Restore VxVM 38
 - ▼ To Uncomment File System Entries in `/etc/vfstab` 38
 - ▼ To Enable Parity Checking 39

Preface

The *Sun StorEdge RAID Manager 6.22 and 6.22.1 Upgrade Guide* describes how to upgrade from Sun StorEdge™ RAID Manager 6.1.1, 6.1.1 Update 1, or 6.1.1 Update 2 to Sun StorEdge RAID Manager 6.22 or 6.22.1. It also covers upgrading from RM 6.22 to RM 6.22.1.

Read this entire document before you attempt to perform the upgrade procedure. Do not deviate from the instructions in this document.

Note – This printed manual may supercede the manual provided on the distribution CD ROM. Please verify that you are using the document with the highest revision number on the title page. For example: 806-7792-*nn*, where *nn* is the revision number of the document.

Using UNIX Commands

This document may not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- AnswerBook2™ online documentation for the Solaris™ operating environment
- *Solaris Handbook for Sun Peripherals*
- Other software documentation that you received with your system

Typographic Conventions

Typeface	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .

Shell Prompts

Shell	Prompt
C shell	<i>machine_name%</i>
C shell superuser	<i>machine_name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Accessing Sun Documentation Online

The `docs.sun.com`SM web site enables you to access a select group of Sun technical documentation on the Web. You can browse the `docs.sun.com` archive or search for a specific book title or subject at:

<http://www.sun.com/products-n-solutions/hardware/docs/index.html>

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can E-mail your comments to Sun at:

`docfeedback@sun.com`

Please include the part number of your document in the subject line of your E-mail.

Upgrade Procedures

This section gives the procedures required to perform the RAID Manager upgrade and provides special notices regarding the upgrade procedure. The topics covered include:

- [“Notice, Restrictions and Prerequisites” on page 10](#)
- [“Preparing for the Upgrade” on page 13](#)
- [“Setting the System Up For the Upgrade” on page 17](#)
- [“Removing the System From Service” on page 20](#)
- [“Upgrading the Host Software” on page 24](#)
- [“Upgrading the Module Firmware” on page 31](#)
- [“Returning the System to Service” on page 38](#)

Notice, Restrictions and Prerequisites

This section contains important information you should read before you begin the upgrade. It covers the following:

- [“Important Notice” on page 10](#)
- [“Upgrade Restrictions” on page 11](#)
- [“Upgrade Conditions” on page 12](#)

Important Notice

This document has been created as a result of bugs and other issues that can cause RAID Manager upgrades to fail and even render a Sun StorEdge A3x00 subsystem inoperative. The main priority of this document is to ensure a safe, reliable upgrade. The upgrade procedure has undergone extensive field testing.

Note – *Follow the procedures in this document exactly.*

It is possible, even likely, that you will encounter instructions within this procedure that can be done better, faster, or even omitted altogether. Resist the temptation to improve upon the procedures in this document; if you alter the procedure, you risk an escalation due to an upgrade gone wrong.

Note – This procedure requires a maintenance window of three to five hours of system downtime, depending on the host boot time and the number of modules to be upgraded. Testers and developers have made every effort to design the procedures so that you complete tasks before a system outage to minimize system downtime.

Upgrade Restrictions

- Do not attempt to make any hardware or software configuration changes that are not mentioned in this document during this upgrade.
- This procedure is for upgrading to RAID Manager 6.22 or 6.22.1 software and firmware only from RAID Manager 6.1.1, 6.1.1 Update 1, 6.1.1 Update 2, and also from RM6.22 to RM 6.22.1.
- This procedure applies only to RAID modules where all physical disks have 40-Mbyte DacStor. Sun StorEdge A3000 arrays upgraded from RAID Manager 6.0 or 6.1 without rebuilt LUNs may have a 2-MByte DacStor. You can get the `dacdiscovery` script to check the DacStor size with the procedure in [Step 3 on page 16](#).

If a drive group is full, additional capacity will be required before you can increase the DacStor to 40 Mbytes. Refer to the *Sun StorEdge RAID Manager 6.22 Release Notes (805-7758-12)* for detailed information on how to address 2-MByte DacStor issues.
- This document does not apply to Sun StorEdge A3500FC systems unless you are upgrading RM 6.22 to RM 6.22.1. RAID Manager 6.22 and 6.22.1 are the only supported versions of RAID Manager for this product.
- This procedure does not apply to products running RAID Manager 6.0 or 6.1 software. To upgrade from RAID Manager 6.0 or 6.1 to 6.22 or 6.22.1, upgrade to RAID Manager 6.1.1 update 2 first (refer to *Raid Manager 6.1.1 Installation and Support Guide*, part number 805-4088). Be aware of differences in DacStor size and use the `dacdiscovery` script when upgrading from 6.0 or 6.1 to 6.1.1.
- This procedure is for use only with single-host, direct-attach Sun StorEdge A3x00 systems, and dual-node, direct-attach Sun StorEdge A3x00 systems running Sun™Cluster 2.x or VCS software. This procedure is not for use with clusters of more than two hosts. If such configurations are encountered, initiate an escalation. Refer to the Sun Cluster documentation for installation details concerning SC 3.x.
- If you are running a data base management system such as ORACLE using raw partition(s) on the A3x00/A1000 array LUNs and you do not have a volume manager such as VERITAS or file system on top of the LUNs, then your configuration can only be upgraded to 6.22.1 and must include the recommended patches. If you have logs or any data going directly to the raw partitions then this restriction applies to you. You must only upgrade to RM 6.22.1 with its patches, which fixes this problem.
- You can not upgrade a system that boots from RAID Manager-controlled LUNs. See [Step 1 on page 15](#) for further information.
- If you are running Solaris 2.5.1, then RM 6.22 is supported, but RM 6.22.1 is not.

Upgrade Conditions

- If you are running VERITAS Volume Manager host-based volume management software on top of RAID Manager 6, you will need to be able to release *all* LUNs under RM 6 from control of the volume management software. This procedure has provided the appropriate commands to do this. You will not be able to perform the RAID controller firmware upgrades required for RAID Manager 6.22x otherwise.
- If you are running the Solstice DiskSuite™ or NetBackup host-based volume management software on top of RAID Manager 6.1.1, initiate an escalation. Upgrading with LUNs under Solstice DiskSuite host-based volume management control is currently not supported.
- If you are running Sun StorEdge Enterprise NetBackup software on the system to be upgraded and you are not already running RM 6.22, contact your service representative. Refer to bug number 4372243.
- If you are using Solaris 2.6 and your `/usr` directory is a partition separate from root, then you must first follow the directions in FIN I0809-1. Otherwise your system may not be able to mount `/usr` as soon as you install RM 6.22.1.

Preparing for the Upgrade

This section contains procedures you should perform well in advance of the actual upgrade. The procedures include:

- [“To Prepare the Cluster Nodes” on page 13](#)
- [“To Obtain the Necessary Software Patches” on page 13](#)

Note – Some of the steps in the following procedures apply only to cluster configurations. Cluster-specific steps are clearly identified. If you are upgrading a single host, follow the directions only for Node A and ignore any directions for Node B.

▼ To Prepare the Cluster Nodes

1. **If your system is running a cluster, note the names of the two nodes or hosts:**

Node A _____

Node B _____

Perform tasks on these two nodes as indicated throughout all the procedures in this guide.

2. **On Node A, create a working upgrade directory in a nonvolatile location.**

For example:

```
# mkdir /var/tmp/RM622_upgrade
```

If your system is running a cluster, repeat this step for Node B.

▼ To Obtain the Necessary Software Patches

Note – If you are running a cluster, complete this entire section first on Node A and repeat on Node B.

1. Obtain the latest packages available for RAID Manager 6.22x, sd, isp, luxadm, and the Kernel Jumbo Patch appropriate for the host operating environment.

This is also a good time to get the latest patches relevant to the Sun StorEdge A3x00 system. The list is available in Early Notifier 20029:

- a. Go to <http://sunsolve.sun.com>.
- b. Click on **Advanced Search**, which is in the left side of the page.
- c. In the **Advanced Search** page, select **Free EarlyNotifier** and click on **NEXT**.
- d. Enter **20029** in the **Document ID: search box** and click **GO**.

TABLE 1 and TABLE 2 list the patches appropriate to each operating environment. The patch revision indicated is the minimum revision level. Use the latest patches.

Note – Patch 105600-19 has been integrated into Kernel Jumbo Patch 105181-25 and subsequent patches.

TABLE 1 Solaris Operating Environment Patches

Patches	Solaris 2.5.1, 11/97* Operating Environment	Solaris 2.6 Operating Environment	Solaris 7 Operating Environment	Solaris 8 Operating Environment
Kernel Jumbo Patch (KJP)	103640- <i>nn</i>	105181- <i>nn</i>	106541- <i>nn</i>	108528- <i>nn</i>
sd	103622- <i>nn</i>	105356- <i>nn</i>	107458- <i>nn</i>	
isp	103934- <i>nn</i>	105600- <i>nn</i>	106924- <i>nn</i>	
luxadm	105310- <i>nn</i>	105375- <i>nn</i>	107473- <i>nn</i>	109529- <i>nn</i>

* Solaris 2.5.1 does not support RM 6.22.1, Fibre Channel, or A3500FC.

TABLE 2 RAID Manager 6.22 Patches

	Solaris 2.5.1 Operating Environment	Solaris 2.6 Operating Environment	Solaris 7 Operating Environment	Solaris 8 Operating Environment
RAID Manager 6.22	108834-07	108834-07	108834-07	108553-07
RAID Manager 6.22.1		112125- <i>nn</i>	112125- <i>nn</i>	112126- <i>nn</i>

2. Copy the packages and patches to the nonvolatile location on Node A that you created in Step 2 on page 13, such as `/var/tmp/RM622_upgrade`.

▼ To Complete the Upgrade Preparation

1. **If you have an encapsulated root device built on a RAID Manager 6.x LUN (even though no RAID Manager 6.x LUNs should be in `rootdg`, as that is not supported), you must unencapsulate the device and move the operating environment image off the RAID Manager 6.x LUN and onto another device.**

You must do this or you will not be able to perform the upgrade to RAID Manager 6.22x. FIN I0619-02 explains the issues and difficulties of booting from an A1000/A3x00/A3500FC device. You should study FIN I0619-02 before attempting to boot from such a device, or attempting to upgrade if you are booting from such a device now.

2. **If you have a Sun Cluster 2.x quorum device built on a RAID Manager 6.x LUN, note which device it is. Please refer to FIN I0520. For further information about quorum devices for this product, contact your service representative.**

This information might be required to restart the Sun Cluster software after the upgrade.

- a. **To determine the quorum device, enter:**

```
# scconf cluster_name -p | grep -i quorum
```

Example output of a serial number is:

```
27000660a
```

- b. **Where 27000660a represents the serial number output from `scconf`, enter:**

```
# finddevices disk | grep -i 27000660a
```

3. If you are using LUNs that were originally created with RAID Manager 6.0 and 6.1, those LUNs have a 2-Mbyte DacStor. To determine which LUNs have a 2-MByte DacStor, use the dacdiscovery script, which is available by following this procedure (In RM 6.22.1, the script is already installed in /etc/raid/bin/dacdiscovery):

- a. Go to `www.sun.com/storage/disk-drives/raid.html`.
- b. Click the `Download` button.
- c. Sign in or register as required.
- d. Accept the license agreement and click `Continue`.
- e. Download the `dacdiscovery` script.

4. To ensure a smooth upgrade, verify that the system device tree is stable.

- a. Make sure that the LUNs are distributed as the customer would expect for normal operation and that you can see all the controllers for each RAID module:

```
# /usr/sbin/osa/lad
```

b. Validate the device tree consistency.

The device ctd's output by the following commands must be consistent:

```
# format
# vxdisk list
# lad
```

If the device ctd's output is not consistent, get whatever help is necessary to fix the inconsistency before you proceed with the upgrade.

5. Obtain the explorer output of your entire system configuration.

If you are running a cluster, repeat this step for Node B.

6. Back up all the data on the RAID Manager 6.1.1 LUNs.

Setting the System Up For the Upgrade

Complete the procedures in this section immediately before you take the system offline to minimize downtime. Begin these procedures 30 minutes before the maintenance window opens. This section covers:

- [“To Save Important Files” on page 17](#)
- [“To Verify That RAID Modules Are In an Optimal State” on page 18](#)
- [“To Save the System State” on page 19](#)

▼ To Save Important Files

Note – If you are running a cluster, perform this entire procedure first on Node A and then repeat the entire procedure on Node B.

1. **Create a subdirectory in the nonvolatile location you created in [Step 2 on page 13](#) and then change to that subdirectory.**

For example:

```
# mkdir /var/tmp/RM622_upgrade/files
# cd /var/tmp/RM622_upgrade/files
```

2. **Back up the following files to the nonvolatile location:**

- /etc/osa/rmparams
- /etc/osa/mnf
- /usr/lib/osa/bin/rmscript
- /etc/path_to_inst
- /kernel/drv/sd.conf
- /kernel/drv/rdriver.conf
- /etc/system
- /etc/vfstab
- /kernel/drv/glm.conf (for PCI systems only)

▼ To Verify That RAID Modules Are In an Optimal State

Note – If you are running a cluster, perform this entire procedure first on Node A and then repeat the entire procedure on Node B.

1. Type:

```
# /usr/sbin/osa/healthck -a
```

2. If any RAID module is not shown as optimal, start the RAID Manager 6 GUI and run the Recovery Application's Recovery Guru to diagnose and correct the problems before proceeding with the upgrade.

The upgrade cannot be performed with any RAID module that is in a sub-optimal state.

▼ To Save the System State

Note – If you are running a cluster, perform this entire procedure first on Node A and then repeat the entire procedure on Node B.

1. Save important RAID Manager 6.x and Volume Manager state information to a nonvolatile location.

For example:

```
# /usr/sbin/osa/lad > /var/tmp/RM622_upgrade/files/lad.out
# diff /usr/lib/osa/rmparams /usr/lib/osa/rmparams.install >
/var/tmp/RM622_upgrade/files/rmparam.diff
# /usr/sbin/vxdg list > /var/tmp/RM622_upgrade/files/vxdg.out
# /usr/sbin/vxdisk list > /var/tmp/RM622_upgrade/files/vxdisk.out
# /usr/sbin/vxprint -hft > /var/tmp/RM622_upgrade/files/vxprint-hft.out
```

2. Save the module profile information.

a. Start the RAID Manager 6 GUI:

```
# /usr/sbin/osa/rm6
```

b. Launch the Configuration Application.

c. Select a specific RAID module from the RAID Module field box.

d. Select the File menu and then click Save Module Profile.

e. From the Save Module Profile dialog, select All and click OK for a complete profile.

f. Save the profile using a unique file name in the nonvolatile subdirectory you created earlier (/var/tmp/RM622_upgrade/files.)

g. Repeat this procedure for each RAID module in the configuration.

Removing the System From Service

This section covers the following procedures:

- [“To Stop All Host I/O” on page 20](#)
- [“To Halt RAID Manager 6 Parity Checking” on page 21](#)
- [“To Enable All RAID Controllers For Hard Reset” on page 22](#)

▼ To Stop All Host I/O

Note – If you are running a cluster, perform this entire procedure first on Node A and then repeat the entire procedure on Node B.

1. **Log in to Node A as a superuser (`root`).**
2. **Verify that any and all host applications that do I/O have been halted.**
3. **Verify that any cluster software has been halted.**
4. **Unmount all file systems that are composed of the RAID Manager 6.x LUNs.**
5. **Comment out any file system entries in `/etc/vfstab` that are RAID Manager 6.x LUNs.**
6. **If any swap devices are built on top of RAID Manager 6.x LUNs, remove them.**
 - a. **Check the swap devices:**

```
# /usr/sbin/swap -l
```

- b. **Remove the swap devices on RAID Manager 6.x LUNs:**

```
# /usr/sbin/swap -d device-name
```

7. **Take offline all disks that are built on top of any RAID Manager 6.x LUNs from volume management control.**

If the system is using VxVM, deport all disk groups that are comprised of RAID Manager 6.x LUNs.

Note – Do not deport the `rootdg` disk group.

a. Deport disk groups:

```
# vxdg deport disk-group-name(s)
```

Note – If this is a Sun Cluster system, the disk groups were already deported when you halted the Sun Cluster software.

b. Verify that they are deported:

```
# vxdg list
```

c. Take each RAID Manager 6.x LUN offline, repeating as necessary for each RAID Manager 6.x LUN:

```
# vxdisk offline cXtXdXs2
```

d. Verify that no disks belonging to RAID Manager 6.x LUNs are still online or in an error state:

```
# vxdisk list | grep -v offline
```

▼ To Halt RAID Manager 6 Parity Checking

Note – If you are running a cluster, perform this entire procedure first on Node A and then repeat the entire procedure on Node B.

1. Start the RAID Manager 6 GUI:

```
# /usr/sbin/osa/rm6
```

2. Select the Maintenance and Tuning Application.

3. Select the Options menu and then click Auto Parity Settings.
4. From the Automatic Parity Check / Repair Settings dialog, make sure Enable Parity Checking With Repair is *not* selected, and then save the setting.
5. Disable any manual parity check by a Solaris operating environment `cron` job.

```
# crontab -l | grep parityck
```

6. If the output is non-null, edit the `crontab` file and comment out the lines returned.
7. Verify that no parity check is currently running:

```
# ps -ef | grep parityck
```

8. If the result is non-null, a parity check is running. Cancel the parity check:

```
# kill -9 parityck-pid(s)
```

▼ To Enable All RAID Controllers For Hard Reset

Note – If you already have RM 6.22 installed, including its firmware, then you must skip this section.

Note – It is not necessary to record the current setting; you will reset the system to the only correct setting later in the procedure. This step prevents controller deadlock during the rest of the procedure.

1. From Node A, set the bit at offset 0x28 in the NVSRAM of all controllers for hard reset:

```
# /etc/raid/bin/nvutil -o 28=2c
```

Note – Ignore the message from `nvutil` regarding “A NVSRAM field has...You must power cycle the affected RAID Module(s),” because a system reboot later in the upgrade procedure produces the same effect.

2. Verify the NVSRAM offset 0x28 is set to 0x2c for all controllers:

```
# /etc/raid/bin/nvutil -o 28
```

Example output:

```
(c4t5d0)lT7a321931 0x28 0x2c  
(c5t4d1)lT93100608 0x28 0x2c
```

Upgrading the Host Software

This section explains how to replace RAID Manager 6.x packages and supporting patches with RAID Manager 6.22 or 6.22.1 packages and supporting patches. The procedures include:

- “To Remove Existing RAID Manager 6.x Patches and Packages” on page 24
- “To Reboot” on page 26
- “To Install RAID Manager 6.22 or 6.22.1 and Patches” on page 27
- “To Edit the /etc/osa/rmparams File” on page 28
- “To Reboot” on page 29

▼ To Remove Existing RAID Manager 6.x Patches and Packages

Note – If you are running a cluster, perform this entire procedure first on Node A and then repeat the entire procedure on Node B.

1. If you are starting from RM 6.22, you should have patch 108834 or 108553 installed. Remove them with `patchrm`. Go to [Step 5 on page 25](#).
2. Determine what revision of relevant RAID Manager 6.1.1 patches exist on the system:

```
# showrev -p | grep 106707
# showrev -p | grep 106513
# showrev -p | grep 106552
# showrev -p | grep 108555
```

Note – 106707 is the patch for Sun Cluster support with RAID Manager 6.1.1 and 106513 is the RAID Manager 6.1.1 jumbo patch.

Note – If you are running RAID Manager 6.1.1 Update 2, `showrev -p` will list 106513-03. You may not be able to remove this patch. This is not a problem, because it was installed as part of the RAID Manager 6.1.1 Update 2 packages and will be removed when you remove those packages.

3. **Remove the patches**, where `xx` is the revision number as indicated by `showrev`:

```
# patchrm patch-number-xx
```

If multiple versions of the same patch exist, remove them in the order most-recent to oldest.

- a. **If the patches were originally installed using the `installpatch` utility that came with the patches, you may see an error. To remove the patch:**

```
# cd /var/sadm/patch/ patch number
```

- b. **Within this directory execute:**

```
# ./backoutpatch patch number
```

4. **Check for the Sun VTS™ diagnostic package and remove it:**

```
# pkginfo -l SUNWvtsse
# pkgrm SUNWvtsse
```

5. **Remove the packages.**



Caution – Do not alter the order of these packages or else the upgrade will not succeed! *Do NOT reboot in this step- see [Step 7 on page 26](#).*

Enter on the command line *in the following order*:

```
# pkgrm SUNWosafw SUNWosamn SUNWosau SUNWosar
```

6. **Install the latest Solaris operating environment patches from [Table 1 on page 14](#) if needed.**

Repeat the following steps for each patch.

- a. Change to the directory where the patches have been copied:

```
# cd /var/tmp/RM622_upgrade
```

- b. For the Solaris operating environment patches only, type:

```
# patchadd patch-number
```



Caution – This step is very important!

7. Ensure a proper reboot:

```
# grep forceload /etc/system
```

- a. If the line `forceload: drv/sd` does not appear, edit the `/etc/system` file so that the file includes:

```
forceload: drv/sd
```

Note – Due to a known issue, the system may panic upon reboot after the removal of the RAID Manager 6.x software if the line `forceload: drv/sd` is not in `/etc/system`.

- b. Review the `/etc/system` file saved during upgrade preparations as a reference, if needed.

▼ To Reboot

Note – The time it takes to reboot depends on your system. It may take up to one hour for each node on a large configuration. Do not skip the reboot. If you skip the reboot, devices will not be reconfigured correctly. Do not power off any of the arrays or other devices at this point.

1. If you are running a cluster, halt Node B and bring it to the `ok` prompt:

```
# init 0
```

2. Reboot Node A:

```
# init 6
```

3. If you are running a cluster, boot Node B after Node A has completely rebooted:

```
ok> boot
```

▼ To Install RAID Manager 6.22 or 6.22.1 and Patches

Note – If you are running a cluster, perform this entire procedure first on Node A and then repeat the entire procedure on Node B.

The instructions assume no other packages or patches are in the directory being used.

1. Go to the directory where the packages and patches have been copied.

For example:

```
# cd /var/tmp/RM622_upgrade
```

2. Use `pkgadd` to install the packages:

```
# pkgadd -d .  
# Which packages do you want to install? all
```

Note – Do not reboot during these steps.

3. Verify that you have the appropriate version of:

- The patches for RAID Manager 6.22 or 6.22.1
- The packages: SUNWosafw, SUNWosamn, SUNWosar, and SUNWosau
- The package SUNWosanv (if you have 6.22.1)

For example, for version 6.22.1:

```
# pkginfo -l packagename
...
VERSION: 06.22, REV=01.54
```

4. Install the RAID Manager 6.22x patches:

```
# patchadd patch-number
```

Repeat for each patch.

▼ To Edit the /etc/osa/rmparams File

Note – If you are running a cluster, perform this entire procedure first on Node A and then repeat the entire procedure on Node B.



Caution – Do not copy the `rmparams` file saved in “[To Save the System State](#)” on [page 19](#) to `/etc/osa/rmparams`. RAID Manager 6.22 introduces new parameters that will be lost if you restore the old file. However, be aware of any prior customizations that must be carried over to the new `rmparams` file. Common customizations include modified values for the parameters `System_MaxLunsPerController` and `System_MaxSCSIid`.

1. **If you are not running a cluster, verify that the value of `Rdac_RetryCount` is 7.**
If it is not, change the value to 7.
2. **If you are running a cluster, change the value of `Rdac_RetryCount` to 1 and `Rdac_NoAltOffline` to TRUE.**
The value for `Rdac_NoAltOffline` is FALSE if you are not using a cluster.
3. **Find the `rmparams` customizations you noted during preparation and modify the new `rmparams` file as needed.**

4. Have the LUN configuration established by running `genscsiconf(1m)`, which will rebuild `sd.conf` and `rdriver.conf`:

```
# /etc/raid/bin/genscsiconf
```

5. If you had previously edited `sd.conf` to contain less targets and hence speed up rebooting times per FIN I0551, you should restore your previously preserved `sd.conf` now.

▼ To Reboot

The time it takes to reboot depends on your system. It may take up to one hour for each node on a large configuration because this will be a reconfiguration reboot.

1. If you are running a cluster, halt Node B and bring it to the `ok` prompt:

```
# init 0
```

2. Reboot Node A:

```
# init 6
```

Note – Leave Node B at the `ok` prompt until directed to do otherwise.

As the reboot finishes, the RAID Manager 6.22 software detects that the RAID controller bootware and firmware are at a down revision level. This is expected behavior; you have not upgraded to the required revisions of RAID Manager 6.22, but you will do so in the next procedure.

You might see the following harmless error messages in the system console when you reboot:

```
Jan 27 05:48:14 e5 unix: WARNING:
/sbus@7,0/QLGC,isp@0,10000/sd@5,0 (sd725):
Jan 27 05:48:14 e5 unix: Error for Command: load/start/stop
Error Level: Retryable
Jan 27 05:48:14 e5 unix: Requested Block: 0
Error Block: 0
Jan 27 05:48:14 e5 unix: Vendor: Symbios
Serial Number:
Jan 27 05:48:14 e5 unix: Sense Key: Not Ready
Jan 27 05:48:14 e5 unix: ASC: 0x4 (<vendor unique code 0x4>),
ASCQ: 0x81, FRU: 0x0
```

If you are rebooting with RAID Manager 6.22.1, you might see additional messages about the NVSRAM, which are described in the *Sun StorEdge RAID Manager 6.22.1 Release Notes* under “Installing and Uninstalling RAID Manager 6.22.1.”

Upgrading the Module Firmware

This section explains how to download the appware and bootware to the Sun StorEdge A3x00 modules. The procedures include:

- [“To Ensure Volumes Are Free After Reboot” on page 31](#)
- [“To Verify Controllers Are Ready For the Firmware Upgrade” on page 31](#)
- [“To Upgrade the Firmware” on page 32](#)
- [“To Restore All RAID Controllers to Soft Reset” on page 34](#)
- [“To Commit the Changes” on page 35](#)
- [“To Verify System Health” on page 36](#)
- [“To Resolve Device Renumbering” on page 37](#)

▼ To Ensure Volumes Are Free After Reboot

- **Check if the host-based volume management or cluster software has resumed operation after rebooting.**

If so, use the same procedures as found in [“Removing the System From Service” on page 20](#) to release RAID Manager-based volumes from cluster and volume management control.

▼ To Verify Controllers Are Ready For the Firmware Upgrade

1. **From Node A, verify all RAID modules are in the optimal state:**

```
# /usr/sbin/osa/healthck -a
```

2. **If any RAID module is not shown as optimal, start the RAID Manager 6 GUI and run the Recovery Application’s Recovery Guru to diagnose and correct the problems before proceeding with the upgrade.**

The upgrade cannot continue with any RAID module in suboptimal state.

Note – If you get a battery alert message, it is ok to proceed with the upgrade.

3. Make sure that the LUNs are distributed as the customer would expect for normal operation and that you can see all the controllers for each RAID module.

From Node A, type:

```
# /usr/sbin/osa/lad
```

▼ To Upgrade the Firmware

Be patient. Downloading the firmware typically takes 5 to 10 minutes or longer for each RAID module.

1. Launch the RAID Manager GUI and start the upgrade.

- a. From Node A, type:

```
# /usr/sbin/osa/rm6
```

- b. Start the Maintenance and Tuning application from the main menu.

- c. If you have both A3x00 RAID Modules and A1000 RAID Modules, you must select the A3x00 RAID Modules one at a time in the “RAID Module:” field and repeat the following steps for each subsequent A3x00 RAID Module. If you have only A3x00 RAID Modules, you should select all the RAID Modules.

Do not upgrade Sun StorEdge A1000 RAID Modules at this point; you will do this after you are done upgrading the A3x00s.

- d. Select the floppy disk icon for firmware download.

- e. Read the IMPORTANT NOTES notice and click OK.

Do not download the NVSRAM yet when prompted because this overrides bit 28.

- f. Select Online in the Select Upgrade Type window.

2. Determine whether you need an intermediate upgrade of the firmware.

- a. Check the current firmware level.

If the firmware level is 02.05.06.32 or later, you do not need to make an intermediate upgrade and you can proceed with [Step 3 on page 33](#).

If the firmware level is not 02.05.06.32 or later, perform an intermediate upgrade.

- b. Select the 02.05.06.32/02.05.06.32 firmware/bootware bundle and click OK.

- c. In the Confirm Upgrade pop-up, click OK to begin downloading the firmware to all of the modules.

Note – Do not exit the RAID Manager 6 GUI application.

When the upgrade is complete, a status window appears and reports success or failure of the upgrade on each module.

d. If any module reports failure to upgrade successfully, invoke the Recovery Guru on that RAID module to diagnose the problem and repeat the firmware download to that particular module or controller.

3. Download the series 03 firmware.

a. Select the latest 03.xx.xx.xx/03.xx.xx.xx bootware and appware bundle and click OK.

If you are loading RAID Manager 6.22.1, select the 03.01.04.xx/03.01.04.00 bootware and appware bundle.

b. In the Confirm Upgrade pop-up, select OK to begin downloading the firmware to all of the modules.

Note – Do not exit the RAID Manager 6 GUI application.

You might see the following harmless error messages in the system console during the firmware download:

```
Jan 27 05:48:14 e5 unix: WARNING:
/sbus@7,0/QLGC,isp@0,10000/sd@5,0 (sd725):
Jan 27 05:48:14 e5 unix: Error for Command: load/start/stop
Error Level: Retryable
Jan 27 05:48:14 e5 unix: Requested Block: 0
Error Block: 0
Jan 27 05:48:14 e5 unix: Vendor: Symbios
Serial Number:
Jan 27 05:48:14 e5 unix: Sense Key: Not Ready
Jan 27 05:48:14 e5 unix: ASC: 0x4 (<vendor unique code 0x4>),
ASCQ: 0x81, FRU: 0x0
```

c. When the upgrade is complete, a status window appears and reports success or failure of the upgrade on each module.

If any module reports failure to upgrade successfully, invoke the Recovery Guru on that RAID module to diagnose the problem, then repeat the firmware download to that module or controller. If necessary, return to [Step c on page 32](#) and repeat for the remaining A3x00 RAID Modules.

4. Verify that the upgrade was successful on all A3x00 RAID Modules.
 - a. From the RAID Manager 6 GUI, launch Configuration Application.
 - b. Select a specific RAID module from the RAID Module field.
 - c. Click the Module Profile option.
 - d. In the Module Profile dialog, click the Controllers option and verify that the bootware and firmware combination is what you selected in [Step a on page 33](#).
If not, repeat the controller upgrade for specific controllers as needed.
 - e. Repeat verification for each RAID module in the configuration.

▼ To Restore All RAID Controllers to Soft Reset

If you are upgrading to RAID Manager 6.22.1, install the new NVSRAM as described in the *Sun StorEdge RAID Manager 6.22.1 Release Notes*, and then upgrade the firmware on any attached A1000s. Otherwise, for RM 6.22, restore all the controllers to soft reset by following this procedure.

1. From Node A, set the bit at offset 0x28 in the NVSRAM for soft reset:

```
# /etc/raid/bin/nvutil -o 28=0c
```

Note – Ignore the message from `nvutil` regarding “A NVSRAM field has... You must power cycle the affected RAID Module(s)”, because a system reset later in the upgrade procedure will produce the same effect.

2. Verify the NVSRAM offset 0x28 is set to 0x2c for all controllers:

```
# /etc/raid/bin/nvutil -o 28
```

Example output:

```
(c4t5d0)1T7a321931 0x28 0x0c  
(c5t4d1)1T93100608 0x28 0x0c
```

3. Verify the health of the NVSRAM:

```
# /etc/raid/bin/nvutil -vf
```

You can now upgrade the firmware for any Sun StorEdge A1000 RAID modules attached to the system. For more information, refer to the procedure in the *Sun StorEdge RAID Manager 6.22 User's Guide*.

▼ To Commit the Changes

Note – There is a bug in `format` that forces you to do an extra reboot if the system is using extended LUN support (more than eight). Rather than rebooting an additional time, use the `/usr/lib/osa/bin/hot_add` utility to discover the extended LUNs.

1. Reboot Node A:

```
# init 6
```

If required, force `format` to see LUNs with IDs greater than 7:

```
# /usr/lib/osa/bin/hot_add
```

Verify that `format` sees more than 8 LUNs:

```
# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
.....
15. c4t4d7 <Symbios-StorEDGEA3000-0301 cyl 748 alt 2 hd 64 sec 64>
/pseudo/rdnexus@4/rdriver@4,7
16. c4t4d8 <Symbios-StorEDGEA3000-0301 cyl 748 alt 2 hd 64 sec 64>
/pseudo/rdnexus@4/rdriver@4,8
17. c4t4d9 <Symbios-StorEDGEA3000-0301 cyl 748 alt 2 hd 64 sec 64>
/pseudo/rdnexus@4/rdriver@4,9
18. c4t4d13 <Symbios-StorEDGEA3000-0301 cyl 748 alt 2 hd 64 sec 64>
/pseudo/rdnexus@4/rdriver@4,d
.....
```

2. If the system is clustered, boot Node B after Node A has been completely rebooted:

```
ok> boot
```

3. Reboot Node B again:

```
# init 6
```

If required, force `format` to see LUNs with ID greater than 7:

```
# /usr/lib/osa/bin/hot_add
```

▼ To Verify System Health

Note – If you are running a cluster, repeat the following steps for Node B after you complete them for Node A.

1. Verify that all RAID modules are in the optimal state:

```
# /usr/sbin/osa/healthck -a
```

2. If any RAID module is not shown as optimal, start the RAID Manager 6 GUI and run the Recovery Application's Recovery Guru to diagnose and correct the problems before proceeding.
3. Make sure that the LUNs are distributed as the customer would expect for normal operation and that you can see all the controllers for each RAID module:

```
# /usr/sbin/osa/lad
```

▼ To Resolve Device Renumbering

Note – If any host configuration changes were made since the initial installation, host devices might have been reordered and renumbered. This is not an issue for most configurations. If the host is running VxVM, it will make the necessary adjustments automatically.

1. **Verify that device renumbering has occurred by comparing the output of the `lad` utility from before and after the upgrade.**

For example:

```
# cd /var/tmp/RM622_upgrade/files
# /usr/sbin/osa/lad > lad.post
# /usr/bin/diff lad.out lad.post
```

If there is no output, the files are the same and there has been no renumbering of the RAID Manager devices. Skip to the section entitled [“Returning the System to Service” on page 38](#).

If renumbering has occurred, use the controller serial number from the `lad` output (1TXXXXXXX) to match up the old and new controller numbers.

2. **If the host is running Sun Cluster and a quorum device is one of the devices that was renumbered, you must re-create the quorum device. See the chapter “Modifying the Sun Cluster Configuration” in the Sun Cluster System Administration Guide for more information on how to do this.**
3. **If the host is running VCS and a quorum device is one of the devices that was renumbered, please refer to the appropriate VCS manual to update this information.**
4. **If the host is directly mounting the RAID Manager devices, modify the `/etc/vfstab` file to reflect the changes in the device renumbering.**

Returning the System to Service

This section provides procedures you must do when you return the system to service. It includes the following procedures:

- [“To Restore VxVM” on page 38](#)
- [“To Uncomment File System Entries in /etc/vfstab” on page 38](#)
- [“To Enable Parity Checking” on page 39](#)

▼ To Restore VxVM

If you are using VxVM, follow these procedures to restore its operation.

1. Import the disk groups:

```
# vxdg import disk-group-name(s)
```

2. Verify that the disks are online and the disk groups are imported:

```
# vxdisk list  
# vxdg list
```

3. Start all volumes in each disk group:

```
# vxvol -g disk-group-name(s) startall
```

▼ To Uncomment File System Entries in /etc/vfstab

1. **Uncomment any file systems entries in /etc/vfstab that comprise RAID Manager 6.22x LUNs.**
2. **If devices were renumbered as a result of the reconfiguration reboot, adjust the file system entries accordingly.**
3. **Mount the file systems.**

▼ To Enable Parity Checking

When you re-enable parity checking, only enable checking without repair. Either select the Without Repair option in the RAID Manager GUI or do not use the `-f` flag with the `parityck` utility when using the `crontab` file.

- **If you disabled parity checking from RAID Manager 6 or `crontab` while performing the procedures in “[To Halt RAID Manager 6 Parity Checking](#)” on [page 21](#), re-enable parity checking.**

Note – If the system is running a cluster, be sure to enable parity checking on just one of the nodes.

The upgrade is complete. You can restart applications.

