

StorageWorks RAID Array 200
Online Management Utility
for Digital UNIX

User's Guide

AA-Q6TGD-TE

Digital Equipment Corporation
Maynard, Massachusetts

Fourth Printing, April 1995

The information in this document is subject to change without notice and should not be construed as a commitment by Digital Equipment Corporation.

Digital Equipment Corporation assumes no responsibility for any errors that might appear in this document.

The software, if any, described in this document is furnished under a license and may be used or copied only in accordance with the terms of such license. No responsibility is assumed for the use or reliability of software or equipment that is not supplied by Digital Equipment Corporation or its affiliated companies.

Copyright © 1995 by Digital Equipment Corporation. All rights reserved.

Printed in the U.S.A.

The following are trademarks of Digital Equipment Corporation:

AXP, Digital UNIX, StorageWorks, SWXCR, and the Digital logo.

The following are third-party trademarks:

Microsoft and MS-DOS are registered trademarks and Windows is a trademark of Microsoft Corporation.

Intel, Intel486, and i486 are trademarks of Intel Corporation.

All other trademarks and registered trademarks are the property of their respective holders.

Contents

Revision Record	vii
About This Guide	ix
1 Installing the RAID Array 200 Online Management Utility for Digital UNIX	
1.1 Introduction.....	1-1
1.2 Before Installing the Utility	1-2
1.3 Files Added to Your System During Installation.....	1-2
1.4 Loading the Utility	1-4
1.4.1 Verifying the Installation.....	1-7
1.4.2 Post Installation.....	1-8
1.5 Accessing the Array.....	1-9
1.5.1 Creating Device Special Files	1-9
1.5.1.1 Determining Controller Numbers.....	1-9
1.5.1.2 Determining Device Special File Names.....	1-10
1.5.1.3 Invoking MAKEDEV	1-11
1.5.2 Reconfiguring the Kernel.....	1-12
1.5.3 Preparing Logical RAID Drives for Access by Digital UNIX Filesystem	1-14
1.5.3.1 Creating the Partitions on a Logical RAID Drive Using Disklabel	1-14
1.5.3.2 Creating a Filesystem on a Logical RAID Drive	1-16
1.5.3.3 Mounting the Filesystem	1-16
2 Using the RAID Array 200 Online Management Utility for Digital UNIX Command Line	
2.1 Monitoring the RAID Subsystem.....	2-1
2.2 Command Flags for swxcrmon	2-3
2.2.1 Notification of Events to the System Console	2-3
2.2.2 Notification of Events to Log Files	2-3
2.2.3 Notification of Events to Mail	2-4
2.2.4 Frequency Status Update	2-5
2.2.5 Events Reported by the Monitor	2-6
2.2.6 Repeating Notification Events	2-7
2.3 Performing Parity Check and Repair on the RAID Subsystem	2-7

2	Using the RAID Array 200 Online Management Utility for Digital UNIX Command Line (cont.)	
2.4	Command Qualifiers for swxrcrcheck	2-8
2.4.1	Repairing the Logical RAID Drive Parity	2-8
2.4.2	Notification of Bad Blocks	2-8
2.5	Altering Rebuild Rate on the RAID Subsystem.....	2-10
2.5.1	Viewing the Rebuild Rate	2-10
2.5.2	Setting the Rebuild Rate	2-11
3	Command Line Messages	
3.1	Monitor Process Messages.....	3-1
3.2	Parity Check Messages	3-8
4	Using the RAID Array 200 Online Management Utility for Digital UNIX Graphical User Interface (GUI)	
4.1	Overview of the RAID Array 200 Online Management Utility for Digital UNIX Graphical User Interface (GUI).....	4-1
4.2	Accessing the Online Management Utility for Digital UNIX GUI.....	4-2
4.3	Using the Online Management GUI.....	4-4
4.3.1	Using the Mouse.....	4-4
4.3.2	Using Windows	4-5
4.3.3	Using Menus	4-6
4.3.4	Using Help	4-6
4.3.5	Exiting from the Online Management GUI	4-6
4.4	Viewing the Subsystem from the Online Management Utility.....	4-7
4.4.1	Selecting a Subsystem	4-7
4.4.2	Viewing Disk Drives	4-10
4.4.3	Viewing Logical RAID Drives.....	4-14
4.5	Printing the Configuration	4-17
5	Maintaining the Subsystem with the GUI	
5.1	Maintaining the Subsystem.....	5-1
5.2	Replacing a Failed Drive	5-2
5.2.1	Hot Spare	5-3
5.2.2	Hot Swap.....	5-3
5.2.3	Marking a Drive as Failed	5-4
5.2.3.1	Rebuilding a Drive	5-6
5.2.3.2	Marking a Drive as Optimal	5-8

5 Maintaining the Subsystem with the GUI (cont.)	
5.3 Making a Hot Spare.....	5-9
5.4 Setting Rebuild Rate.....	5-10
5.5 Viewing the Bad Block Table	5-11
5.6 Checking Parity	5-12
Index	Index-1
Readers Comment Form	Inside Back Cover
Figures	
4-1 The Base Window of the Online Management Utility for Digital UNIX GUI, Displaying a 1-Channel Configuration	4-2
4-2 The Base Window of the Online Management Utility for Digital UNIX GUI, Displaying a 3-Channel Configuration	4-3
4-3 The Adapter Selected Buttons.....	4-7
4-4 The Subsystem Displayed after Selecting EISA Slot 8	4-8
4-5 The Subsystem Displayed after Selecting PCI Slot 6	4-9
4-6 The Disk Drive Matrix for a 1-Channel Subsystem.....	4-10
4-7 The Disk Drive Matrix for a 3-Channel Subsystem.....	4-12
4-8 The Logical Drive Table.....	4-14
4-9 Printing the Configuration	4-17
5-1 Fail Drive Confirmation Window	5-5
5-2 The Rebuild Drive Confirmation Window	5-6
5-3 The Rebuild Drive Confirmation Window	5-7
5-4 The Rebuild Completed Message.....	5-7
5-5 The Make Optimal Confirmation Window	5-9
5-6 Set Rebuild Rate Window.....	5-10
5-7 The Parity Repair Option Window	5-12
5-8 The Parity Progress Window	5-13
5-9 Parity Check Completes Message	5-13
5-10 Parity Check Status Report Appears When Parity Check Completes	5-14
5-11 Parity Errors Detected Message	5-14

Tables

1-1	Command Line Files Added During Installation	1-3
1-2	GUI Files Added During Installation	1-3
2-1	Summary of swxcrmon Flags.....	2-2
2-2	Summary of swxcrcheck Flags	2-7
3-1	Monitor Process Messages.....	3-1
3-2	Parity Check Messages	3-8
4-1	Mouse Functions	4-4
4-2	Window Elements	4-5
4-3	Drive Matrix Information	4-11
4-4	Disk Drive Status.....	4-13
4-5	Logical Drive Table Information	4-15
4-6	Logical Drive Status	4-16

Revision Record

This Revision Record provides a concise publication history of this manual. It lists the manual revision levels, release dates, and reasons for the revisions.

The following revision history lists all revisions of this publication and their effective dates. The publication part number is included in the Revision Level column, with the last entry denoting the latest revision. This publication supports StorageWorks RAID Array 200 Subsystems in the Digital UNIX operating system environment.

Revision Level	Date	Summary of Changes
AA-Q6TGA-TE	August 1994	Original release. This release supports the StorageWorks RAID Array 200 Online Management Utility for Digital UNIX command line only.
AA-Q6TGB-TE		
AA-Q6TGC-TE	January 1995	This release supports the StorageWorks RAID Array 200 Online Management Utility for Digital UNIX command line and Graphical User Interface (GUI).
AA-Q6TGD-TE	April 1995	This release reflects the change of DEC OSF/1 to Digital UNIX.

About This Guide

This section identifies the audience of this guide and describes the contents (chapter by chapter) and structure. In addition, this section includes a list of associated documents and the conventions used in this guide.

The StorageWorks RAID Array 200 Online Management Utility for Digital UNIX User's Guide

This guide describes how to install the StorageWorks RAID Array 200 Online Management Utility for Digital UNIX and how to use the RAID Array 200 Online Management Utility for Digital UNIX command line and Graphical User Interface (GUI).

Intended Audience

This guide supports users of the StorageWorks RAID Array 200 Subsystems in a Digital UNIX environment. You should be familiar with StorageWorks RAID terminology, Digital UNIX and the procedures for adding a device to your system.

Organization

This guide contains the following:

Chapter 1: Installing the RAID Array 200 Online Management Utility for Digital UNIX

Installing the RAID Array 200 Online Management Utility for Digital UNIX lists the files added to your system by this installation, and discusses running the *setld* utility to install the command line and graphical user interface; the procedures for adding a device to your system; and preparing logical RAID drives for access by the Digital UNIX filesystem.

Chapter 2: Using the RAID Array 200 Online Management Utility for Digital UNIX Command Line

Using the RAID Array 200 Online Management Utility for Digital UNIX Command Line discusses monitoring the RAID subsystem through the SWXCR monitor, methods of event notification; and using the command line to run parity check and repair, and to set the rebuild rate.

Chapter 3: Command Line Messages

Command Line Messages describes the informational, and error messages displayed by the command line monitor and parity check processes.

Chapter 4: Using the RAID Array 200 Online Management Utility for Digital UNIX Graphical User Interface (GUI)

Using the RAID Array 200 Online Management Utility for Digital UNIX Graphical User Interface (GUI) describes how to access and use the GUI.

Chapter 5: Maintaining the Subsystem with the GUI

Maintaining the Subsystem with the GUI discusses how to use the GUI to: replace a failed drive, including: failing, rebuilding and marking a drive optimal; create a hot spare, set rebuild rate, view bad block table, and check parity.

Conventions

This guide uses the following conventions:

Style	Meaning
boldface monospace type	To be input by the user
plain monospace type	Screen text
<i>italic type</i> screens,	For emphasis, manual titles, utilities, menus, and filenames

Special Notices

This guide uses in the following to emphasize specific information.

WARNING

WARNING indicates the presence of a hazard that can cause personal injury if the hazard is not avoided.

CAUTION

CAUTION indicates the presence of a hazard that might damage hardware or corrupt software.

NOTE

Notes provide additional information.

Related Documentation

Refer to the *StorageWorks RAID Array 200 Subsystems Controller Installation and Standalone Configuration Utility User's Guide* for information on the controller and on configuring your subsystem.

1

Installing the Storage Works RAID Array 200 Online Management Utility for Digital UNIX

After you install and configure the RAID Subsystem, you can install the StorageWorks RAID Array 200 Online Management Utility for Digital UNIX. This software consists of the SWXCR Monitor utility, SWXCR Parity Check utility, SWXCR Rebuild Rate utility, and SWXCR Graphical User Interface (GUI).

1.1 Introduction

The RAID Array 200 Online Management Utility for Digital UNIX consists of a command line and a Graphical User Interface (GUI).

The command line provides:

- Monitoring
- Parity check and repair
- Rebuild control (of redundant logical RAID drives)
- Viewing help

The GUI utility provides:

- Management of the RAID subsystem
- Rebuild control (of redundant logical RAID drives)
- Parity Check

1.2 Before Installing the Utility

Prior to installing the RAID Array 200 Online Management Utility for Digital UNIX on your subsystem, you should have completed the following:

- Run the EISA Configuration Utility (for EISA controllers only)
- Set controller options and parameters
- Installed the controller and subsystem
- Connected the subsystem to the controller
- Configured logical RAID drives on the subsystem with the standalone configuration utility

In addition, we recommend that you back up your system before you attempt to install the utility for Digital UNIX.

1.3 Files Added to Your System During Installation

The installation of the online management command line and Graphical User Interface (GUI) utilities for Digital UNIX create the files listed in Tables 1-1 and 1-2 on your system.

Table 1-1 Command Line Files Added During Installation

File	Description
/usr/opt/swxcr/bin/swxcrmon	Monitor utility executable
/usr/opt/swxcr/man/man8/swxcrmon.8	swxcrmon manpage
/usr/opt/swxcr/.X11Startup	Sample shell script
/usr/opt/swxcr/S99swxcr_monitor	Sample rc3 script shell
/usr/bin/swxcrmon	Symbolic link
/usr/share/man/man8/swxcrmon.8	Symbolic link
/usr/opt/swxcr/bin/swxcrcheck	Parity check utility executable
/usr/opt/swxcr/bin/swxcrate	Rebuild rate utility executable
/usr/share/man/man8/swxcrcheck.8	swxcrcheck manpage
/usr/share/man/man8/swxcrate.8	swxcrate manpage
/usr/bin/swxcrcheck	Symbolic link
/usr/bin/swxcrate	Symbolic link
/usr/share/man/man8/swxcrcheck.8	Symbolic link
/usr/share/man/man8/swxcrate.8	Symbolic link

Table 1-2 GUI Files Added During Installation

/usr/opt/swxcrmgr/bin/swxcrmgr	Management utility executable
/usr/bin/swxcrmgr	Symbolic link
/usr/opt/swxcrmgr/lib/swxcrmgr	GUI Resource file
/usr/lib/x11/app-defaults/swxcrmgr	Symbolic link
/usr/opt/swxcrmgr/man/man8/swxcr.8	swxcr manpage
/usr/opt/swxcrmgr/man/man8/swxcrmgr.8	swxcrmgr manpage

1.4 Loading the Utility

To load the online management utility for Digital UNIX command line and GUI, you need to run the *setld* utility.

The procedure for loading the online management utility for Digital UNIX takes about 5 minutes to complete.

To load and configure the online management utility for Digital UNIX, follow these steps:

1. Load the distribution diskette into your floppy drive.
2. Mount the device. Type this command:

```
# mount -o ro /dev/fd0c mount_point
```

where *mount_point* is a valid mount point on your system, for example, */mnt*. (Make sure that */mnt* is not in use.)

3. From the Bourne shell root account, type this command to run *setld* and install the utilities:

```
# setld -l /mount_point
```

The installation procedure requires you specify and confirm whether you want to install the command line, the GUI or both. The installation then does the following:

- Checks for the necessary file system space.
- Displays copyright information.
- Displays the approximate installation time (5 minutes) and disk space requirements.
- Installs the StorageWorks RAID Array 200 Online Management Utility for Digital UNIX V1.0.0.
- Creates symbolic links.

The following is an example of an installation.

The subsets listed below are optional:

There may be more optional subsets than can be presented on a single screen. If this is the case, you can choose subsets screen by screen or all at once on the last screen. All of the choices you make will be collected for your confirmation before any subsets are installed.

- 1) StorageWorks RAID Array 200 Management Utility for DEC OSF/1 V1.0.0
- 2) StorageWorks SWXCR Utility for Digital UNIX V1.2.0

Or you may choose one of the following options:

- 3) ALL of the above
- 4) CANCEL selections and redisplay menus
- 5) EXIT without installing any subsets

Enter your choices or press RETURN to redisplay menus.

Choices (for example, 1 2 4-6): 3

You are installing the following optional subsets:

StorageWorks RAID Array 200 Online Management Utility
for Digital UNIX V1.0.0

StorageWorks SWXCR Utility for Digital UNIX V1.2.0

Is this correct? (y/n):y

Checking file system space required to install selected subsets:

File system space checked OK.

*

© Digital Equipment Corporation 1994. All rights reserved.

Restricted Rights: Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of DFARS 252.227-7013, or in FAR 52.227-19, or in FAR 52.227-14 Alt. III, as applicable.

This software is proprietary to and embodies the confidential technology of Digital Equipment Corporation. Possession, use or copying of this software and media is authorized only pursuant to a valid written license from Digital or an authorized sublicensor.

*

StorageWorks RAID Array 200 Online Management Utility for Digital UNIX

Press return to continue...

The installation of the StorageWorks RAID Array 200 Management Utility for Digital UNIX V1.0.0 will take about 5 minutes and requires approximately 850 Kbytes of disk space both during and after installation.

StorageWorks RAID Array 200 Management Utility for Digital UNIX V1.0.0

Copying from . (disk)

Verifying

*

© Digital Equipment Corporation 1994. All rights reserved.

Restricted Rights: Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of DFARS 252.227-7013, or in FAR 52.227-19, or in FAR 52.227-14 Alt. III, as applicable.

This software is proprietary to and embodies the confidential technology of Digital Equipment Corporation. Possession, use or copying of this software and media is authorized only pursuant to a valid written license from Digital or an authorized sublicensor.

*

Press return to continue...

The installation of the StorageWorks SWXCR Utility for Digital UNIX V1.2.0 will take about 5 minutes and requires approximately 450 Kbytes of disk space both during and after installation.

StorageWorks SWXCR Utility for Digital UNIX V1.2.0

Copying from . (disk)

Verifying

Configuring "StorageWorks RAID Array 200 Management Utility for Digital UNIX V1.0.0" (SWXCRMGR100)

Now creating links from /usr/bin files to kit executable files.

```
Now creating links from /usr/lib/X11/app-defaults files to
kit resource files.

Now creating links from /usr/share/man/man8 files to kit
man pages.

The StorageWorks RAID Array 200 Management Utility for
Digital UNIX Installation is now complete.

Configuring "StorageWorks SWXCR Utility for Digital UNIX
V1.2.0" (SWXCR120)

Now creating links from /usr/bin files to kit executable
files.

Now creating links from /usr/share/man/man8 files to kit
man pages.

*****
**

The SWXCR Utility will create log files in /usr/opt/swxcr,
unless the SWXCR_LOGPATH environmental variable is
defined. If you wish to use an alternate directory for
log files, you must create that directory and set
SWXCR_LOGPATH to the desired path.

*****
**

The StorageWorks SWXCR Utility for Digital UNIX Installation
is now complete.
```

1.4.1 Verifying the Installation

You can verify that all files were installed in the correct directories by typing the following commands:

```
# setld -v SWXCR120
# setld -v SWXCRMGR100
```

This executes the installation verification procedure (IVP).

1.4.2 Post Installation

After the installation completes, you have the option to create a log file directory in Korn/Bourne shells as follows:

In Korn/Bourne shells enter: **SWXCR_LOGPATH=log_file.directory**
export SWXCR_LOGPATH

Then, you can issue **env** command to examine defined path variable

In csh enter: **setenv SWXCR_LOGPATH log_file.directory**

For example, **setenv SWXCR_LOGPATH /logfiles**

Then, if you define the SWXCR_LOGPATH environment variable as the path to the log file directory, the SWXCR utility writes any log files to that specified directory. If you do not define the SWXCR_LOGPATH environment variable, any log files are written to /usr/opt/swxcr.

To start the SWXCR Monitor automatically when you start up the system, place an rc3 command script containing the swxcrmon command in the /sbin/rc3.d directory. This command script then causes the SWXCR Monitor utility to be invoked when entering a multiuser run level. Refer to the rc3 manpages for more information about the multiuser run level. There is a sample command script provided in /usr/opt/swxcr/S99swxcr_monitor. You can copy this file directly to /sbin/rc3.d/S99swxcr_monitor. However, the command flags in this script and the xcr controller name should be modified as desired. See Chapter 2 for more information about monitor commands. See Section 1.5, Accessing the Array to determine the xcr controller name.

If you are running the system in a Windows environment and want to start up the SWXCR monitor automatically, you should place a .X11Startup command script in the home directory of the system manager account. For example, if the system manager uses the root account the command script will have to be placed in /.X11Startup. This command script will only be run if the system manager logs into the system using the Windows interface. A sample script is provided by the installation in /usr/opt/swxcr/.X11Startup and should be modified as desired.

1.5 Accessing the Array

When you first installed the subsystem, you used the standalone utility to create logical RAID drives. In order to access these drives through the Digital UNIX operating system, you need to do the following as described in the following Section 1.5.1 through 1.5.3:

- Create device special files
- Reconfigure the kernel
- Prepare logical RAID drives for access by Digital UNIX filesystem

1.5.1 Creating Device Special Files

Create block and character device special files in */dev* for each of the logical drives that you defined.

You need the following information to create the device special files:

- Controller numbers
- Device special file names

1.5.1.1 Determining Controller Numbers

The operating system assigns controller numbers at boot after scanning the slots in the system for the SWXCR controller. The operating system scans in lowest to highest slot number order, first the EISA slots, then the PCI slots. The SWXCR controllers appear as *xcrn* devices, where *n* is the controller number that the operating system assigns. Digital UNIX assigns *xcr0* to the first controller it encounters, *xcr1* to the second, and so on.

For example, suppose you have 3 SWXCR controllers in a system, an EISA controller in slot 1, an EISA controller in slot 7, and a PCI controller in slot 4. The operating system will assign *xcr0* to the EISA controller in slot 1, *xcr1* to the EISA controller in slot 7, and *xcr2* to the PCI controller in slot 4.

View the console log in */usr/var/adm/messages* to see the controller numbers that the operating system assigned to the controllers in your system at boot time. Record the controller numbers on a piece of paper in order to have them available to determine the device special file names.

1.5.1.2 Determining Device Special File Names

The names of the block and character device special files, which appear in */dev* as: */dev/r x* (block devices) and */dev/r r x* (character devices), include the device number, x , which you need to calculate.

Calculate the device number using the following equation:

controller number * 8 + logical RAID drive number

(8 is a constant, representing the total number of logical RAID drives it is possible to have on a controller)

Example 1: Suppose you have a single EISA controller in your system, and you created a single logical RAID drive on it using the standalone utility. Digital UNIX will assign the name *xcr0* to that controller on boot. The logical RAID drive number will be 0, because you created only one logical RAID drive (logical RAID drives are numbered from 0-7 in ascending order for each controller).

Therefore:

controller number you will use for the calculation in this case is 0

logical RAID drive number will be 0.

Calculate the device number as:

controller number * 8 + logical RAID drive number, or $0 * 8 + 0$

Then:

device number = 0.

block device special file = *re0*.

character device special file = *rre0*.

Example 2: Assume you have a PCI controller in your system in addition to the EISA controller configuration described in Example 1. The EISA controller is *xcr0*, and the PCI controller is *xcr1*. You created 2 logical RAID drives on your PCI controller, logical RAID drive 0 and logical RAID drive 1. You created 2 logical RAID drives, so you will need to calculate two device numbers for the PCI controller. Therefore:
controller number you will use for the calculation in this case is *1*
logical RAID drive numbers will be *0* and *1*

Calculate the device numbers as:
controller number * 8 + logical RAID drive number, or $1 * 8 + 0$ and $1 * 8 + 1$
Then:
device numbers = *8* and *9*.
device special file names = *re8* and *re9*.
character device special file names = *rre8* and *rre9*.

Now, calculate the device numbers and determine the device and character block special file names for every logical drive on your subsystem. Then go to section 1.5.1.3 to create the device special files.

1.5.1.3 Invoking MAKEDEV

After you determine the device special file names, to create the device special files, follow these steps:

1. Invoke the MAKEDEV utility in */dev* using the name of the device special file you want to create. For example, type: **`/dev/MAKEDEV re0`**
2. Repeat Step 1 to create the device special file for each of the device special file names you determined.

Go to Section 1.5.2 to reconfigure the kernel.

1.5.2 Reconfiguring the Kernel

Reconfigure the kernel so that your system will recognize and display the logical RAID drives on the console during boot. The device special files must be created using the steps detailed in Sections 1.5.1.1 to 1.5.1.3 prior to reconfiguring the kernel.

To add device entries to the configuration file in */sys/conf*, follow these steps:

1. Change to */sys/conf*, type: # **cd /sys/conf**
2. Save the original configuration file (the configuration file is usually the system name in capital letters. In this example, the configuration file is */sys/conf/NODENM*), type: # **cp NODENM NODENM.old**
3. Open the configuration file with an editor and check for any entries for the SWXCR controllers in your system appearing as follows:

The EISA RAID Controller will appear in the configuration file similar to:
controller xcr0 at eisa0 slot 4 vectr xcrintr

The PCI RAID controller will appear in the configuration file similar to:
controller xcr1 at pci0 slot 4 vectr xcrintr

NOTE

If you configured your controller with the genvmux kernel and no entry for it appears in your custom kernel configuration file, then you will need to add the entry in the configuration file using the format for either the EISA or PCI RAID controller shown above. View the console log file in */usr/var/adm/messages* to see how the controllers were assigned at boot time.

4. Add all the device special files that you created in */dev* to the configuration file as follows:

Type: **device disk re0 at xcr0 drive 0**

where *re0* is the block device special file name, *xcr0* is the controller, and drive *0* is the device special file number.

For example, suppose we have an EISA controller in slot 7 and a PCI controller in slot 3. We have defined 2 logical RAID drives on each controller, logical RAID drives 0 and 1 on each.

The entries in the configuration file would look as follows:

```
controller xcr0 at eisa0 slot 7 vectr xcrintr
device disk re0 at xcr0 drive 0
device disk re1 at xcr0 drive 1
controller xcr1 at pci0 slot 3 vectr xcrintr
device disk re8 at xcr1 drive 8
device disk re9 at xcr1 drive 9
```

5. Once you type all the device special file entries in the configuration file, type:
doconfig -c NODENM
where *NODENM* is the configuration file name
The system will copy the current configuration file to a backup file, and ask you if you want to edit the configuration file.
6. Select No to edit the configuration file. Digital UNIX creates a new kernel */sys/NODENM/vmunix*
7. Save the old version of *vmunix* in the root directory (*/*) and copy this new *vmunix* kernel into *vmunix*, type:
cp vmunix vmunix.old
mv /sys/NODENM/vmunix /vmunix
The next time the system is rebooted, it will display the configured logical RAID drives on the console.
You have completed reconfiguring your kernel.

1.5.3 Preparing Logical RAID Drives for Access by Digital UNIX Filesystem

Digital UNIX treats a logical RAID drive much like a SCSI device; therefore, to prepare your logical RAID drive for access by the Digital UNIX filesystem you will need to do the following:

- Create the partitions on the logical RAID drive using `disklabel`
- Create a filesystem on the logical RAID drive
- Mount the filesystem to copy files to it

1.5.3.1 Creating the Partitions on a Logical RAID Drive Using Disklabel

You create the partitions on a logical RAID drive by issuing a `disklabel` command. The `disklabel` partitions the logical RAID drive for access by the Digital UNIX filesystem. Digital UNIX defines only partitions `a`, `b`, `c`, and `g` for the SWXCR array controller. In addition, the size of the logical RAID drive must be at least 300 MB to support the `a` and `b` partition. If you try to label a logical RAID drive with less than 300 MB you will get partition errors. Refer to the *re* and *disktab* man pages for more information about `disklabel`.

To create the read/write partitions on a logical RAID drive using the standard partition default sizes, enter the following:


```
# disklabel -rw devicepartition SWXCR
```

→ *device* is the character device special filename as defined in Section 1.5.1.2

→ *partition* is the letter representing the partition, a, b, c, or g on which you want to create a filesystem (usually c for the whole drive).

→ *SWXCR* defines this logical RAID drive as attached to the SWXCR array controller. Always use the device label *SWXCR*, regardless of what drive types make up the logical RAID drive.

For example, to create partitions on character device, *rre0*, enter:

```
# disklabel -rw rre0 SWXCR
```

To view a logical RAID drive partition, enter: # **disklabel -r device**

1.5.3.2 Creating a Filesystem on a Logical RAID Drive

Use the *newfs* command to create a filesystem on a logical RAID drive the same way that you would create a filesystem on a disk device by entering the following:

```
# newfs devicepartition
```

→ *device* is the block device name as defined in Section 1.5.1.2

→ *partition* is the letter representing the partition, a, b, c, or g on which you want to create a filesystem (usually c for the whole drive)

For example, to create a filesystem on partition c of block device, re0c, enter the following:

```
# newfs re0c
```

This creates a filesystem on the c partition of logical RAID drive 0 of the *xcr0* controller.

1.5.3.3 Mounting the Filesystem

To access the logical RAID drive, mount it as a device filesystem to a mount point. For example:

```
# mount /dev/re0c /mnt
```

To view the mounted filesystem, enter: *df*

The logical RAID drive is now accessible to the filesystem much as a disk device would be. The filesystem can not see the RAID functionality and number of physical devices attached to the SWXCR array controller. This device appears as a single logical RAID drive or "disk" to the user as viewed by the filesystem.

To copy files to this device, use the *cp* command as follows:

```
# cp filename /mnt
```

→ *filename* is the name of the file to be copied

→ */mnt* is the mount point for the filesystem

This completes the installation of the StorageWorks RAID Array 200 Online Management Utility for Digital UNIX. Go on to Chapter 2 to start the monitor, and then to Chapter 4, if you want to start the graphical user interface.

Using the RAID Array 200 Online Management Utility for Digital UNIX Command Line

The RAID Array 200 Online Management Utility for Digital UNIX provides the SWXCR Monitor online utility, to notify you of subsystem status changes, the SWXCR parity check utility, to perform parity check and repair functions on RAID logical drives, and the SWXCR Rebuild Rate utility, to modify the rebuild and parity check rate for a given controller.

2.1 Monitoring the RAID Subsystem

The monitor function scans the array for nonoptimal device status and reports the status to you. You must invoke a separate monitor process for each SWXCR RAID controller you have on your subsystem. If you do not start a monitor process, you are not notified of events that occur on your subsystem.

To invoke a monitor process, ensure that your \$PATH environment variable path includes /usr/bin. Then type the swxcrmon command as follows:

```
# swxcrmon ctlrname
```

where *ctlrname* is the name of the xcr controller. Otherwise, you can type the command as follows:

```
# /usr/bin/swxcrmon ctlrname
```

For example, if you specify xcr0 as the controller name, the monitor process monitors all logical RAID drives configured on the xcr0 controller.

Refer to Post Installation Section 1.4.2 of Chapter 1 for more information about starting the SWXCR Monitor utility automatically.

The swxcrmon command creates a separate background process.

Table 2–1 is a summary of the swxcrmon command flags. Refer to the following sections for a detailed description of the flags.

Table 2–1 Summary of swxcrmon Flags

Action	Flag
Send notification of subsystem events to the system console.	-r
Write the event text to a log file.	-l
Scan for status once every <i>n</i> minutes, where <i>n</i> is an interval of 1 to 10 minutes.	-i 1 ...10
Report an event only the first time it occurs. If -n is not specified, then events are reported every time they are encountered.	-n
Send a mail message notifying users of subsystem events. Send to users listed in the user-created file /usr/opt/swxcr/swxcrmon.maillist if it exists, or root if no file is present.	-m

When you issue the swxcrmon command, the Monitor utility displays the following text on successful startup:

SWXCR Monitor Utility V1.2

©Digital Equipment Corporation 1994. All rights reserved.

2.2 Command Flags for `swxcrmon`

The following sections contain detailed descriptions of the `swxcrmon` command flags.

2.2.1 Notification of Events to the System Console

You can specify whether notification of subsystem events are sent to the system console.

```
# swxcrmon xcrn -r
```

If neither `-l` nor `-r` is specified, then events are sent to the system console.

2.2.2 Notification of Events to Log Files

You can specify that events are written to the product-specific log file using the `-l` flag as follows:

```
# swxcrmon xcrn -l
```

If you do not specify the `-l` flag, no log file is created.

When the StorageWorks RAID Array 200 Online Management Utility for Digital UNIX is installed, this directory is created:

```
/usr/opt/swxcr
```

If you specify the `-l` flag when you start the monitor process, a log file is created in the `/usr/opt/swxcr` directory. However, if you define `SWXCR_LOGPATH`, the log file is created in the directory specified by `SWXCR_LOGPATH`. The log file name is as follows:

```
dd-Mon-hh:mm_XCRn.log
```

where *dd* is the current date, *Mon* is the current month, *hh* is the current hour, *mm* is the current minute, and *XCRn* is the name of the specified controller. For example, the log file for the XCR0 controller might be `05-May-18:42_XCR0.log`.

2.2.3 Notification of Events to Mail

You have the option to have the monitor notify a specified list of users of subsystem events by means of the Mail utility. To send notification of events to a group of users, enter the command as follows:

```
$ swxcrmon xcrn -m
```

This will search for the file /usr/opt/swxcr/swxcrmon.maillist. The file, which you create, contains a list of users who are to receive mail messages. If this file is not found, mail will be sent to the root user on the system. The default (when -m is not used) is not to send mail. A sample swxcrmon.maillist file follows:

```
root
lucy
bob@emubird
peter@stone.dec.com
```

It is important to note that Mail is sent from the account of the user who starts the monitor process.

When notification of events by mail is specified, events are collected and sent out in one (or more) Mail message(s). There is not one Mail message per event. If there is at least one event to report in a given scan of the array, at least one Mail message is sent when that array scan is complete. The number of events and the length of the event text determine the number of Mail messages sent. More than one Mail message may be sent (each with unique events listed in the order in which the events were encountered).

If you specify the -m flag, the swxcrmon command parses the swxcrmon.maillist file, and sends a test message to each user on the list. If the test message fails, an error message is displayed and the monitor process is not created. Note that specifying an invalid user (one that does not exist) will not necessarily cause an error. In this case, mail may be returned to the account that started the monitor. If you notice this, you should edit the swxcrmon.maillist file (this can be done while the monitor is running) to correct the problem. An example of the initial test message which is sent is as follows:

```
Received: by abcdef.dec.com; id AA01658; Tue, 12 Jul 1994 09:53:01 -0400
Date: Tue, 12 Jul 1994 09:53:01 -0400
From: system PRIVILEGED account <root>
Message-Id: <9407121353.AA01658@abcdef.dec.com>
Apparently-To: root
```

Subj: SWXCR XCR0 Event Notification from node abcdef.dec.com
This is a test message from SWXCRMON, the StorageWorks SWXCR RAID array subsystem monitor process. You will be receiving MAIL notification of events.
The following is an example of a Mail message sent from the swxcrmon process.

Received: by abcdef.dec.com; id AA01674; Tue, 12 Jul 1994 09:53:02 -0400
Date: Tue, 12 Jul 1994 09:53:02 -0400
From: system PRIVILEGED account <root>
Message-Id: <9407121353.AA01674@abcdef.dec.com>
Apparently-To: root
Subj: SWXCR XCR0 Event Notification from node abcdef.dec.com
Logical RAID drive 0 is degraded.
Logical RAID drive 1 is degraded.
The hard disk at channel 0, target 2 is failed.

2.2.4 Frequency of Status Updates

You can specify the frequency (in minutes) that the disk array is scanned for status. Use the `-i` flag as follows:

```
# swxcrmon xcrn -i n
```

where *n* is a value of 1 to 10 minutes. If the `-i` flag is not specified, the default is a 5-minute interval. The interval is the time between completion of one array scan and the start of the next array scan. Messages from the monitor may not be exactly *n* minutes apart.

2.2.5 Events Reported by the Monitor

The monitor reports the following conditions:

- Logical RAID drive status changes.
- SCSI device errors.
- SCSI device status changes.
- Rebuild status.
- SCSI device insertion or removal.
- Shelf events or errors.

The monitor process may not report on every event, depending on the following:

- The length of time specified by the `-i` flag. Refer to the *Frequency of Status Updates* Section 2.2.4 of this chapter for detailed information on this flag.

You can use the `-i` flag to specify the time between array scan operations. If you specify a relatively short time interval, then the array is scanned frequently. The more frequent the scan operations, the more likely it is that an event is captured and reported. If you specify a longer time interval, status from one event can be overwritten by status from a subsequent event.

For example, if you remove a hard disk and replace it with another disk in the same target and channel, the disk removal event status is overwritten by the disk insertion event status. If you specify a short array scan interval, it is more likely that both the disk removal and the disk insertion events are reported.

- During a rebuild, the SWXCR monitor process reports on each of the affected logical RAID drives. If you have a relatively small logical RAID drive and the array scan interval is relatively long, the rebuild of the logical RAID drive might not be reported. In this example, the rebuild of the logical RAID drive may have started after an array scan operation and completed before the next array scan. If this happens, then a rebuild in progress event is not reported for that logical RAID drive.

2.2.6 Repeating Notification of Events

By default, an event is reported every time it is encountered (during an array scan) until it is either fixed or cleared. To specify that events are reported only once, use the `-n` flag as follows:

```
# swxcrmon xcrn -n
```

Now the event is reported only once, even if the same condition is encountered on the next array scan.

2.3 Performing Parity Check and Repair on the RAID Subsystem

The parity check function reads the data from the logical RAID drive and computes the expected parity information for each series of blocks. It compares the computed parity to the stored parity information on the drives.

Use this function after a power failure or a system crash to verify the integrity of parity information on logical RAID drives. This utility is only available for redundant RAID levels (for example, RAID 1, RAID 0+1, and RAID 5).

To invoke the parity check utility, ensure that your `$PATH` environment variable path includes `/usr/bin`. Then, type the `swxcrcheck` command as follows:

```
$ swxcrcheck devname
```

where *devname* specifies the device name for the logical RAID drive on the controller.

For example, if you specify *re0* as the device name, the parity is checked on logical RAID drive *0*.

Table 2–2 summarizes the default actions of this command. Refer to the following sections for a detailed description of the qualifiers.

Table 2–2 Summary of swxcrcheck Flags

Action	Flag
Repair any inconsistent parity information.	-r
Display a log of bad block information.	-l
Write a copy of the bad block log to a file	-o

2.4 Command Qualifiers for `swxcrcheck`

The following sections contain detailed descriptions of the `swxcrcheck` command flags. In each example below, the *ren* indicates logical RAID drive *ren*, for example, `re0`.

2.4.1 Repairing the Logical RAID Drive Parity

To specify parity repair use the `swxcrcheck` command with the `-r` flag as follows:

```
$ swxcrcheck ren -r
```

Parity check will automatically repair any inconsistencies that are found. If you do not specify `-r`, repair does not take place.

2.4.2 Notification of Bad Blocks

To display a list of bad blocks found, use the `swxcrcheck` command with the `-l` flag as follows:

```
$ swxcrcheck ren -l
```

Otherwise, parity check does not display a list of the bad block information.

To save a log of bad block information to a file, use the `-o` flag as follows:

```
$ swxcrcheck ren -o filename
```

When the StorageWorks RAID Array 200 Online Management Utility for Digital UNIX is installed, the following directory is created:

```
/usr/opt/swxcr
```

If you specify the `-o` flag when the parity check utility is started, a log file is created in the `/usr/opt/swxcr` directory unless you have previously defined `SWXCR_LOGPATH`; in this case, the log file is created in the directory specified by `SWXCR_LOGPATH`. If the filename argument is specified and includes a directory path, the file is created in the directory path specified in the filename. Note that if a directory path is specified in the filename, the filename must also be included (it is not defaulted).

If you specify the `-o` flag without a filename argument, the filename is as follows:

```
dd-Mon-hh:mm_ren.log
```

where *dd* is the current date, *Mon* is the current month, *hh* is the current hour, *mm* is the current minute, and *ren* is the name of the specified device. For example, the log file for the `re0` device might be `06-Jun-17:19_re0.log`.

If both the `-l` and the `-o` flags are specified, then a log is written to the terminal and to the output log file. If the specified output file cannot be opened, the log is written to the terminal instead.

The following is a sample output from a parity check process with the `-l` flag specified. In this case, bad blocks were found. The list specifies the block number where a bad block was found and the number of bad blocks found at that address.

```
$ swxcrcheck re2 -l
```

```
© Digital Equipment Corporation 1994. All rights reserved.
```

```
Parity check complete.
```

```
The following bad blocks were found:
```

	Block Number	# Bad Blocks	
	-----	-----	
②	000000576	000000016	①
	0000000992	000000016	
	0000001472	000000016	
	0000001984	000000016	
	0000002176	000000016	
	0000002208	000000016	
	0000002976	000000016	
	0000003680	000000016	

1. The controller transfers data to the disks in multiple-block chunks. In this example, the number of blocks is 16 (or 8 Kbytes). If a bad block is found, the whole chunk is marked bad, not just the one bad block. As a result, all errors show a size of 16 blocks.
2. Note that the numbers in the left column are all divisible by the chunk size (in this case, 16). This is because the controller marks an entire chunk as bad whenever an error is found. In this example, if block 580 is a bad block, then this represents the chunk from 576-591. If more than one error is found in a chunk, the controller reports that chunk as bad only once. Note that if a set of bad blocks spans across two chunks (for example, 591-592), the controller reports both chunks as bad (block 576 for 32 rather than 16) blocks. Note also that the list may not be in strict numerical order due to the way the controller processes commands.

2.5 Altering Rebuild Rate on the RAID Subsystem

The rebuild rate function allows you to view, or select the rate at which a rebuild or parity check will complete. A rebuild or parity check will slow down system I/O performance and system I/O will slow down rebuild or parity check performance. Therefore, you can use the rebuild rate function to raise or lower the priority of a rebuild or parity check and thus speed up or slow down its performance.

NOTE

Before invoking the rebuild rate utility, ensure that your \$PATH environment variable includes /usr/bin.

2.5.1 Viewing the Rebuild Rate

To view the rebuild rate, invoke the rebuild rate utility by entering the command as follows:

```
$ swxcrate devname
```

where *devname* is the name of the controller, *xcr0*, for example.

The utility displays a message such as the following:

```
Rebuild Rate is currently set to 50.
```

2.5.2 Setting the Rebuild Rate

You can change the rebuild rate while a rebuild or parity check is in progress. If you find that either process takes longer than you would like, you can increase the rebuild rate, or if you find that system I/O performance suffers because of the parity check or rebuild, you can lower the rate.

To set the rebuild rate, invoke the rebuild rate utility by entering the command as follows:

```
$ swxcrrate devname rate
```

where *devname* is the name of the controller, *xcr0*, for example, and *rate* is the rate you wish to set. Acceptable values for *rate* are any integer between 0 and 50. The slowest rate is 0 (all priority is given to I/O and least priority to the rebuild or check). The fastest rate is 50 (equal priority is given to rebuild or check). The utility displays a message such as:

```
Rebuild Rate set to 50 (was 40).
```

Command Line Messages

3.1 Monitor Process Messages

The monitor process displays the general informational and error messages described in Table 3–1 through the reporting mechanisms you set up.

Table 3–1 Monitor Process Messages

Message	Meaning
Controller checksum error; monitor process terminated.	The monitor encountered a checksum error when trying to read the controller configuration. In this case, the controller did not return any configuration to use so the monitor process will terminate with an error.
NVRAM and EEPROM configurations do not match.	The monitor process determined, when reading the configuration from EEPROM on the controller, that the NVRAM and EEPROM configurations do not match. The monitor process was able to continue because the EEPROM configuration was transferred, but the user should run the configuration utility to determine which configuration is correct. (It is not clear whether the EEPROM configuration or the NVRAM configuration is the correct one).

Table 3-1 Monitor Process Messages (continued)

Message	Meaning
NVRAM configuration does not match.	The monitor process determined, when reading the configuration from EEPROM, that the EEPROM configuration and the actual configuration do not match. This is because of a startup error, such as a storage enclosure was not found, because it's power was not turned on, or a device failed. Check your storage enclosure, or run the configuration utility to determine what caused the startup error.
Replacement table full. Monitor can no longer report on rebuilds. Monitor process is terminating!!!	<p>The controller maintains a list of changes to the configuration caused by standby replacements. The monitor process attempts to clear this list after processing a rebuild start, but it may not be able to clear the table if the SWXCR configuration continually changes. While the monitor makes repeated attempts to clear the table, it is possible (but highly unlikely) that the monitor cannot do so.</p> <p>In this case, the controller can continue to function, but the monitor process can no longer correctly report on standby replacements. Therefore, if the list is full and the monitor process cannot clear it, the monitor process will terminate. Reboot your system.</p>

Table 3-1 Monitor Process Messages (continued)

Message	Meaning
Internal log structures getting full. OR Internal log structures are full!	The SWXCR controller maintains an internal log of changes in device states within a session (power-cycle or reset to the next power-cycle or reset). This log is a fixed-size list, and when this list is full, the monitor process reports the second message. When the list has less than three entries left, the monitor process reports the first message. If you see either message, you must shut down the system, cycle power, and reboot. This clears the internal log structure.
The hard disk at channel <chn>, target <tgt> is being rebuilt. This affects the status of logical RAID drive(s) <drive-list>.	A failed hard disk at the specified channel and target is being rebuilt. <chn> is the channel number; <tgt> is the SCSI target ID. The <drive-list> shows the logical RAID drives to be rebuilt as part of the rebuild of the specified hard disk.

Table 3-1 Monitor Process Messages (continued)

Message	Meaning
<p>The hard disk at channel <old-chn>, target <old-chn> is being replaced by the hard disk at channel <new-chn>, target <new-tgt>. This affects the status of logical RAID drive(s) <drive-list>.</p>	<p>When a hot spare rebuild is in progress, the monitor reports this message. It indicates the channel and target of both the hard disk that failed and the hot spare hard disk which is replacing it.</p> <p>The <old-chn> is the channel number of the hard disk being replaced; <old-tgt> is the SCSI target id of the hard disk being replaced; <new-chn> is the channel number of the hot spare device; <new-tgt> is the SCSI target id of the hot spare device.</p> <p>The <drive-list> shows the logical RAID drive numbers that have blocks on the disk being replaced. Each of these logical RAID drives becomes critical and is rebuilt as part of the rebuild of the hard disk. The monitor reports on the progress of the rebuild, including the rebuild of each logical RAID drive.</p> <p>If the rebuild is successful, the monitor reports that each of the logical RAID drives is now on line. If a rebuild operation fails, the monitor reports that the logical RAID drives are still degraded.</p>
<p>Logical RAID drive <drive-num> is <drive-status></p>	<p>Monitor is reporting on a logical RAID drive status. If -nonrepeat is specified, only changes in logical RAID drive status are reported. Otherwise, any changes in status are reported, and any logical RAID drive status which is not 'optimal' is reported. <drive-num> is the logical RAID drive unit number and <drive-status> is one of the following: "Optimal," "Degraded," "Failed," or "Unknown."</p>

Table 3-1 Monitor Process Messages (continued)

Message	Meaning
<p>The hard disk at channel <chn>, target <tgt> had a <err-type> error.</p> <p>OR</p> <p>The hard disk at channel <chn>, target <tgt> had <count> <err-type> errors</p>	<p>A SCSI device at the specified channel and target had one of the specified type of errors or the device had the specified number of the specified type of error. The number of errors is the increment since the last report of the specified error type <u>not</u> the total accumulation of that type of error on the hard disk. <chn> is the channel number and <tgt> is the SCSI target id. <count> is the number of errors that occurred; <err-type> is one of the following: "parity," "soft," "hard," or "miscellaneous."</p>
<p>The hard disk at channel <chn>, target <tgt> is <dev-sts>.</p>	<p>Monitor is reporting on the status of the hard disk at the specified channel and target. SCSI disk status changes are reported only once, regardless of the presence or absence of the -nonrepeat flag. <chn> is the channel number and <tgt> is the SCSI target id. <dev-sts> is one of the following: "failed, a hot spare," "write only," or "optimal."</p>
<p>A new hard disk has been inserted at channel <chn>, target <tgt>.</p> <p>OR</p> <p>The hard disk at channel <chn>, target <tgt> has been removed.</p>	<p>These two messages notify the user of insertion and removal of drives. <chn> is the channel number and <tgt> is the SCSI target id. These messages are reported only once, regardless of the presence or absence of the -nonrepeat flag.</p>
<p>The hard disk at channel <chn>, target <tgt> is not being automatically replaced by a hot spare. Please replace the device and rebuild it.</p>	<p>The hard disk (whose status will have already been reported) has failed or been removed and is not being automatically replaced by a hot spare. User action is required to replace the drive. <chn> is the channel number and <tgt> is the SCSI target id.</p>

Table 3-1 Monitor Process Messages (continued)

Message	Meaning
<p>Logical RAID drive <drive-num> is rebuilding.</p> <p>OR</p> <p>Logical RAID drive <drive-num> is rebuilding (<percent>% complete).</p>	<p>When a logical RAID drive starts rebuilding, The monitor displays this message. If -norepeat is specified, the first message displays. Otherwise, the second message displays, and includes the percent of that disk that has been rebuilt. <drive-num> is the logical RAID drive unit number. <percent> is the percentage complete.</p>
<p>The rebuild (on hard disk at channel <chn>, target <tgt>) has successfully completed.</p> <p>OR</p> <p>The rebuild (on hard disk at channel <chn>, target <tgt>) completed with some errors.</p> <p>OR</p> <p>The rebuild (on hard disk at channel <chn>, target <tgt>) failed because the new drive failed.</p> <p>OR</p> <p>The rebuild (on hard disk at channel <chn>, target <tgt>) failed because the source drive failed.</p> <p>OR</p> <p>The rebuild (on hard disk at channel <chn>, target <tgt>) failed because of bad blocks in the source drive(s).</p>	<p>These messages are used to report on rebuild status. When a hard disk is rebuilt, the monitor reports on each affected system drive as it is being rebuilt. When the disk is completely rebuilt, or the disk rebuild fails, one of these messages is used to indicate the completion status. If possible, the failure reason is indicated. <chn> is the channel number and <tgt> is the SCSI target id.</p>

Table 3-1 Monitor Process Messages (continued)

Message	Meaning
Write cache error has occurred.	Some write data in the controller cache was not successfully updated to disk due to disk failure.
Shelf failure on channel <chn>.	A shelf failure occurred on the specified channel. <chn> is the channel number. A shelf failure may indicate that a power failure occurred on the shelf.
Shelf error on channel <chn>.	A shelf error occurred on the specified channel. <chn> is the channel number. A shelf error may indicate that a drive is not seated properly or that a blower or redundant power supply failed.
Monitor processing terminating. Fatal error encountered: <err_text>.	The monitor process encountered the specified fatal error. User action depends on the error specified in <err_text>.

3.2 Parity Check Messages

The parity check utility displays the informational and error messages described in Table 3–2 through the reporting mechanisms you set up.

Table 3–2 Parity Check Messages

Message	Meaning
The bad block table is full. There may be more bad blocks on the specified system drive than those specified above.	The bad block table is a fixed size. If too many bad block are found, the table fills up and some bad blocks may not get reported. This message displays at the end of the /LOG display.
Invalid logical RAID drive specified.	The device name specified in the command line does not correspond to a logical RAID drive.
Error in parity check command occurred.	Controller returned unknown status from the Parity Check command.
Parity check cannot be performed on RAID 0 (nonredundant) logical RAID drive. OR Parity check cannot be performed on JBOD (nonredundant) logical RAID drive.	The logical RAID drive specified is a nonredundant RAID level (for example, RAID 0) and therefore parity check cannot be performed.

Table 3–2 Parity Check Messages (cont.)

Message	Meaning
One or more SCSI disks in this logical RAID drive has failed.	Parity check cannot be performed because one or more SCSI disks in the logical RAID drive is dead; parity checks can be performed only on logical RAID drives in the optimal state. The disk should be replaced and the logical RAID drive rebuilt.
A rebuild or parity check is already in progress.	Only one rebuild or parity check at a given time is allowed by the controller. This error indicates that either a rebuild or parity check is currently in progress. The user will have to wait until the function is complete, and then reissue the swxcrcheck command.

Using the RAID Array 200 Online Management Utility for Digital UNIX Graphical User Interface

This chapter presents the functionality available for monitoring and maintaining StorageWorks RAID Array 200 Subsystems with the Online Management Utility Graphical User Interface (GUI) for Digital UNIX

4.1 Overview of the RAID Array 200 Online Management Utility for Digital UNIX Graphical User Interface (GUI)

The RAID Array 200 Online Management Utility for Digital UNIX Graphical User Interface (GUI) provides the following functionality to maintain a RAID Array 200 Subsystem in a Digital UNIX environment:

- View the status of the disk drives in the array
- View the logical drives
- Check and repair parity of redundant logical drives
- View bad blocks of a logical drive
- Manually replace a drive after a disk failure using rebuild (for redundant configurations), or make optimal (for non-redundant configurations)
- Set the rebuild rate
- Create hot spares

This chapter describes the GUI and how to use it to monitor the status of a RAID Array 200 subsystem. Chapter 5 describes how to use the GUI to manage the array.

4.2 Accessing the Online Management Utility for Digital UNIX GUI

To start the Online Management GUI, type the following: `swxcrmgr`

The online management GUI base window appears similar to the one shown in Figure 4–1 for a 1-channel subsystem, or similar to the one shown in Figure 4–2 for a 3-channel subsystem.

Figure 4–1 The Base Window of the Online Management Utility for Digital UNIX GUI, Displaying a 1–Channel Configuration

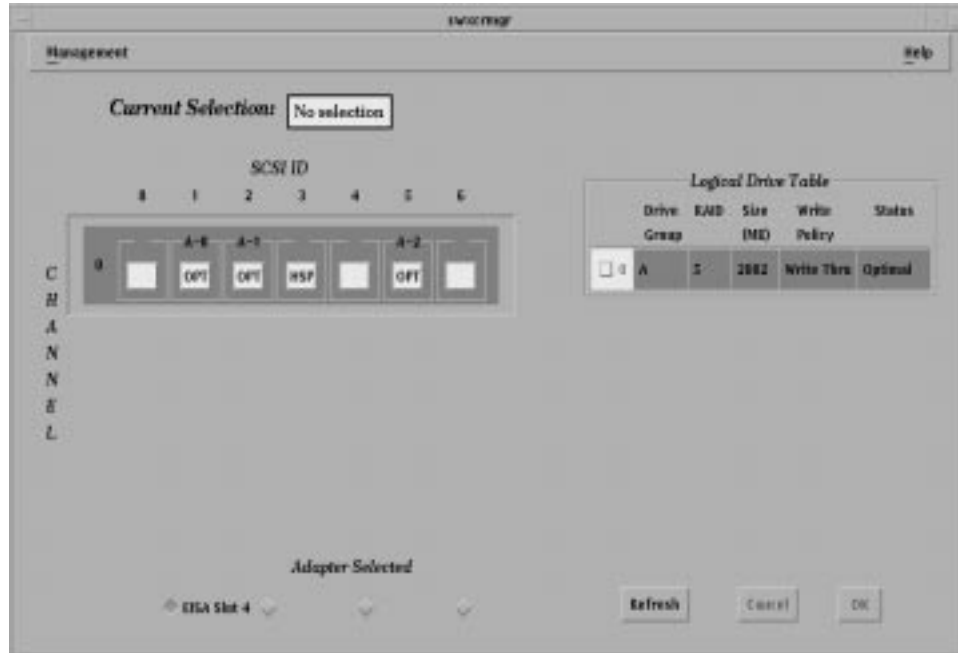
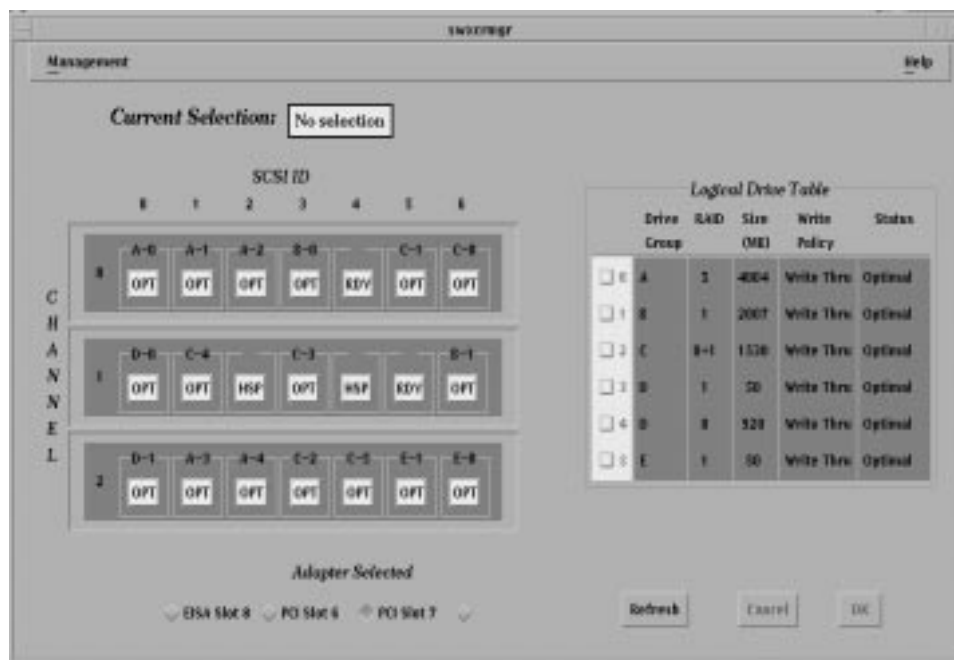


Figure 4–2 The Base Window of the Online Management Utility for Digital UNIX GUI, Displaying a 3–Channel Configuration



The base window displays:

- Menus
- A matrix of the disk drives installed in your subsystem
- A table of the logical drives configured on your subsystem
- The current function selected
- The current adapter selected
- Action buttons

4.3 Using the Online Management GUI

The online management GUI is compliant with the Motif style guidelines. It is a window and menu driven environment designed to be used with a mouse.

4.3.1 Using the mouse

Table 4-1 describes the functions of each of the mouse buttons for a right-handed, three-button mouse.

Table 4-1 Mouse Functions

To:	Press:
Select an item	Left mouse button For example, double click on a drive icon to display drive parameters
Access a menu	Left mouse button For example, point to a menu button on the screen, and click the left mouse button to open the menu and display its options

4.3.2 Using Windows

Table 4–2 describes the elements that may appear on the GUI windows.

Table 4–2 Window Elements

Elements	Description
Menu	Point to a menu on the screen and click your left mouse button to display menu items
Title	Displays name of application you accessed, StorageWorks(TM) RAID Array 200 Management Utility for Digital UNIX(TM) V1.0.0 Copyright Digital Equipment Corporation 1994
Resize corner	Allows you to change window size from any corner of window NOTE: The contents of the base window are NOT resizeable
Control area	Displays menus
Action Buttons	Point and click to an action button to start or cancel the selected function, or to refresh the information displayed in the window

4.3.3 Using Menus

From the online management GUI base window, you can access the following menus:

- *Management*
- *Help*

The *Management* window provides the following items in the order shown:

Default Rebuild Rate

Rebuild Drive

Parity Check

Fail Drive

Make Optimal

Make Hot Spare

Bad Block Table

Print Configuration

Exit

The *Help* window provides the following items in the order shown:

On Menu Options

About this utility

To view the items on a menu, point to the menu and click your left mouse button.

4.3.4 Using Help

To get help while using the online management for Digital UNIX GUI base window, select *Help* from the *Help* menu.

4.3.5 Exiting from the Online Management GUI

When you finish using the online management GUI, choose the *Exit* item from the *Management* menu. The *Exit swxcrmgr* window appears. Click *OK* to exit, or click *Cancel* to remain in the utility.

4.4 Viewing the Subsystem from the Online Management GUI

A subsystem consists of a controller and the storage enclosures attached to it. The operating condition of a RAID Array 200 subsystem depends upon the status of each individual disk drive installed in the storage enclosures, and the status of each of the logical drives configured with those disk drives. You can view the status of the disk drives and logical drives in each of the storage enclosures attached to each of your controllers from the base window.

To monitor a subsystem, do the following:

- Select a controller
- View the Disk Drives
- View the Logical Drives

4.4.1 Selecting a Subsystem

The base window lists the controllers installed in your system under *Adapter Selected*. The base window displays controllers by the system expansion slot, [Expansion Board Type] [Slot] [Slot Number], in which you installed them, as shown in Figure 4-3.

To view a subsystem, click on the slot number of a controller. The base window then displays the logical drive information for the subsystems connected to the selected controller (adapter), as shown in Figures 4-4 and 4-5.

Figure 4-3 The Adapter Selected Buttons




Figure 4-4 The Subsystem Displayed after Selecting EISA Slot 8

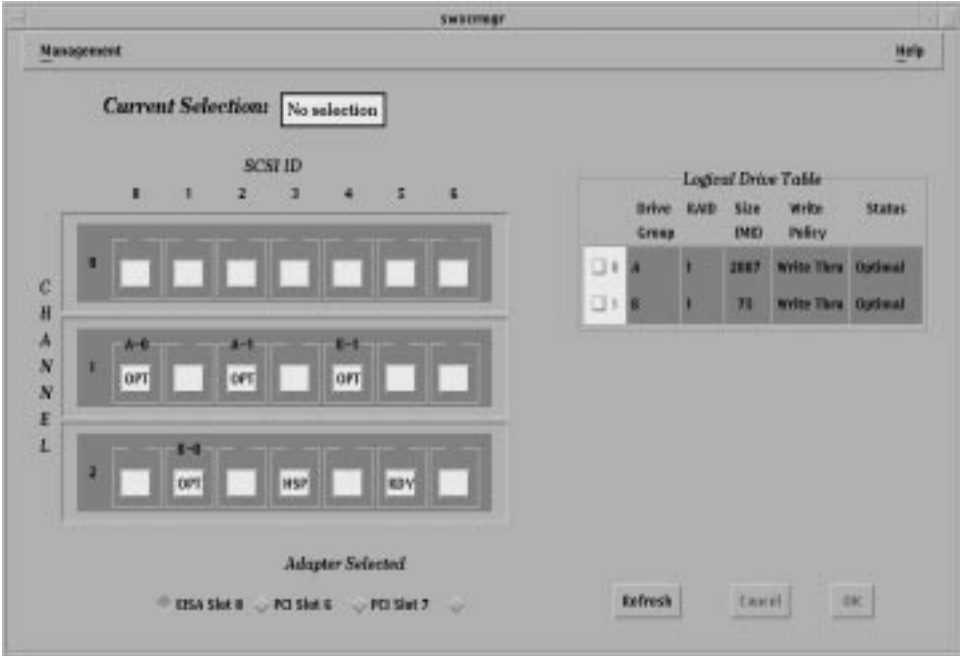
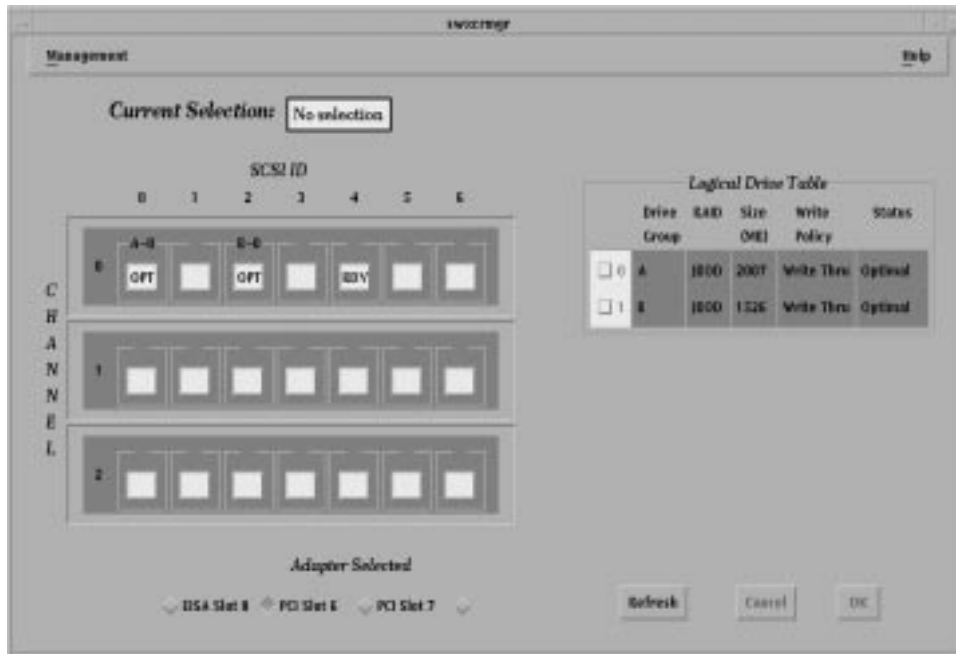


Figure 4-5 The Subsystem Displayed after Selecting PCI Slot 6



4.4.2 Viewing Disk Drives

The GUI displays disk drives in a disk drive matrix on the base window, similar to the drive matrixes shown in Figure 4-6 for a 1-channel subsystem and Figure 4-7 for a 3-channel subsystem.

The drive matrix displays all of the drives connected to each channel of the controller installed in the selected slot.

Table 4-3 describes all the information the drive matrix displays, except status. Table 4-4 describes disk drive status.

Figure 4-6 The Disk Drive Matrix for a 1-Channel Subsystem




Table 4-3 Drive Matrix Information

Element	Description
Channel	<p>Each channel on your controller is an independent SCSI bus supporting up to 7 devices. Additional channels provide you with additional disk capacity (see SCSI ID).</p> <p>The RAID controller can access only one drive per channel at a time, but can simultaneously access up to 3-channels at a time. If you have a 3-channel controller, you can divide up the disk drives in your subsystem on different channels, and thus increase the subsystem performance by allowing the controller to simultaneously access one disk per each channel.</p> <p>You can configure drive groups on one channel of a controller, or across channels of a controller.</p>
SCSI ID	<p>Each drive on the channel (an independent SCSI bus) must have a unique SCSI ID. You can have up to 7 disk drives per channel. If you want disk capacity above seven drives, you need additional channels.</p>
Disk Drive Icons	<p>Each drive installed in your subsystem(s) is represented in the drive matrix by a drive icon. The drive icon displays the status of the drive (Table 4-4 describes each drive status), and changes color to indicate when you have selected this drive.</p>

Table 4-3 Drive Matrix Information (cont.)

Element	Description
Drive Group Designator	The drive group designator is a letter that appears at the top of the drive icon to identify each drive that belongs to the same drive group. You can easily determine drives that are not configured into a drive group, because they do not have drive group designators.
Drive Number	The drive number appears at the top of the drive icon. The online management GUI numbers each drive, so that you can determine easily how many drives belong to a drive group. Drives that are not configured into a drive group do not require a drive number.

Figure 4-7 The Disk Drive Matrix for a 3-Channel Subsystem




Table 4-4 Disk Drive Status

Status Code	Status	Description
RDY	Ready	A disk drive installed in your subsystem, but not configured into a drive group or as a hot spare All disk drives have this status when you first install the subsystem
HSP	Hot Spare	A drive that you leave unused in your subsystem in the event that one of the disks of a logical drive fails If one of the disks of a logical drive fails, the subsystem automatically begins using the hot spare drive in place of the failed drive
OPT	Optimal	A disk drive configured into a drive group A drive with this status is in good working condition
FLD	Failed	A disk drive configured into a drive group that the subsystem can not access because of some mechanical or other type of failure You need to replace a drive with this status See Chapter 3 for information on how and when to replace and rebuild failed drives
WOL	Write Only	A disk drive that the system is rebuilding The data may not be good on a drive being rebuilt; therefore, the subsystem can not read data from a disk that is being rebuilt; the subsystem can only write to it

4.4.3 Viewing Logical RAID Drives

The GUI displays logical RAID drives in a *Logical Drive Table* on the base window, similar to the *Logical Drive Table* shown in Figure 4–8.

The *Logical Drive Table* displays all of the drives connected to the channel or channels of the controller in the selected slot.

Table 4–5 describes all the information the *Logical Drive Table* displays, except status. Table 4–6 describes logical drive status.

Figure 4–8 The Logical Drive Table

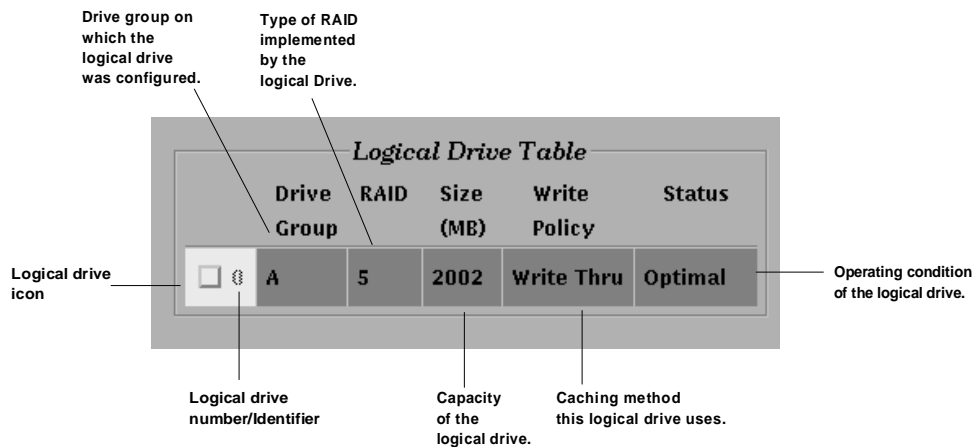


Table 4-5 Logical Drive Table Information

Element	Description
Logical Drive Icon	The GUI represents each logical drive configured on your subsystem by a logical drive icon in the <i>Logical Drive Table</i> .
Logical Drive Number	Each logical drive configured on your subsystem has a unique number/identifier between 0 and 7. The controller assigns the logical drive number when you create the logical drive with the standalone utility.
Drive Group	The <i>Logical Drive Table</i> lists the drive group used to configure the logical drive. This way, you can quickly locate in the drive matrix the drive group associated with a logical drive.
RAID	Indicates the RAID level of the logical drive. The RAID level determines how the logical drive stores data.
Size	Indicates the storage capacity of the logical drive. The storage capacity of the drive group and the RAID level used for the logical drive determine the storage capacity of the logical drive.
Write Policy	Indicates the caching policy selected for this logical drive, write thru, or write back. See the <i>StorageWorks RAID Array 200 Subsystem Controller Installation and Standalone Configuration Utility User's Guide</i> to learn about the caching options.
Status	The status of the logical drive is dependent upon the status of the disk drives in the drive group used to configure the logical drive. See Table 4-6 for a description of logical drive status.

Table 4-6 Logical Drive Status

Status	Description
Optimal	All of the disk drives in the drive group associated with this logical drive have a status of optimal. The logical drive is in good operating condition.
Degraded	The controller can not access a disk drive in the drive group on which this logical drive was configured. This status appears for RAID 5, RAID 0 + 1, and RAID 1 logical drives only, since these RAID levels have redundancy and can be rebuilt. See Section 5.2 on replacing a failed drive.
Dead	The controller can not access two or more disk drives in the drive group associated with a RAID 5, RAID 0 +1, or RAID 1 logical drive. You must restore data from backup media. or The controller can not access a disk drive in the drive group associated with a RAID 0 logical drive, or JBOD configuration. You replace the failed drive and restore data from backup media. See Section 5.2 on replacing a failed drive.


4.5 Printing the Configuration

Use the Print Configuration option to output configuration information to a text file suitable for printing. The file will contain information such as number of drive groups, and number, type, and size of logical drives.

To Print the Configuration, follow these steps:

1. Select *Print Configuration* from the *Management* menu. A window appears, similar to the one shown in Figure 4–9, containing a field for the name of the file to save the information to. The default is SWXCR.asc, which will appear in the current directory (/) from where the GUI was invoked.
2. Type a filename in the field (or accept the default), and click OK in the window. The GUI creates a file for the configuration information.

Figure 4–9 The Print Current Configuration Window



Note

The GUI will overwrite old versions of the file; therefore, when printing multiple configurations, make sure to save each one under separate file names.

5

Maintaining the Subsystem with the GUI

This chapter presents the functionality available for monitoring and maintaining StorageWorks RAID Array 200 Subsystems with the Online Management Utility for Digital UNIX GUI.

5.1 Maintaining the Subsystem

The RAID Array 200 Online Management Utility for Digital UNIX GUI allows you to do the following to maintain a RAID Array 200 subsystem:

- Replace failed drives in redundant and nonredundant configurations
- Create hot spares
- Set the rebuild rate
- Check and repair parity for redundant configurations
- View bad block table

5.2 Replacing a Failed Drive

One of the major benefits of a RAID subsystem is the capability to continue accessing the data stored on the subsystem in the event of a single disk failure. Although redundant RAID configurations can withstand disk failures, most can withstand only a single disk failure. For many redundant RAID configurations, if two disks fail at the same time, all the information stored on the logical RAID drives with which those disks are associated is lost.

If a drive experiences severe enough errors, the controller will automatically fail the drive. However, if you know or suspect that one of the drives in your storage enclosure is bad (some hints include a high rate of errors, or degraded system performance) you will want to replace it with another drive as soon as possible.

Your RAID Array 200 subsystem and RAID Array 200 Online Management for Digital UNIX GUI provides you with the following methods for replacing failed drives:

- Hot spare
- Hot swap
- Mark as failed/Rebuild (redundant configurations only)
- Mark as failed/Make optimal (non-redundant configurations only)

Sections 5.2.1 to 5.2.3 describe when and how to use each of the methods for replacing a failed drive.

5.2.1 Hot Spare

You can use this method to replace a drive if:

- The drive you want to replace contains only redundant RAID configurations (RAID 1, 5, 0+1)
- You defined one or more hot spare drives

See section 5.3 for information on how to create a hot spare drive.

To replace a drive with a hot spare, remove the bad or suspect drive from the storage enclosure and the RAID controller automatically begins to recreate the data stored on the removed disk onto the hot spare drive disk.

5.2.2 Hot Swap

You can use this method to replace a drive if:

- The drive you want to replace contains only redundant RAID configurations (RAID 1, 5, 0+1)
- Your storage enclosure supports hot swaps (StorageWorks storage enclosures all support hot swap capability)
- You enabled the Fault Management option when first configuring the subsystem (see the Controller Installation and Standalone Configuration Utility User's Guide for information on how to enable this option)

To replace a drive using hot swap, remove the bad or suspect drive from the storage enclosure, and wait 20 seconds (so the controller recognizes the drive change) before inserting a new drive in the same slot from which you removed the bad/suspect drive. The RAID controller automatically begins to recreate the data stored on the removed disk onto the disk drive inserted into the storage enclosure in its place.

5.2.3 Marking a drive as Failed

Use the Mark as Failed option in conjunction with the Rebuild or Mark as Optimal options to replace a drive if:

- The hot spare or hot swap options are not supported on your subsystem
- To replace a drive from a non-redundant configuration (RAID 0, JBOD)
- To replace a drive, which is logging errors


NOTE

Use caution when failing drives. If the drive you fail is part of a non-redundant logical RAID drive (RAID 0 or JBOD), all the data on that logical RAID drive will be lost. Backup a drive before you mark it failed.

To replace a drive using Mark as Failed, follow these steps:

1. Select *Fail Drive* from the *Management* menu. *Fail Drive* appears as the Current Selection at the top of the base window.
2. Click on any drive icon with a status of *OPT* or *WOL*, and then click *OK*.
3. A window appears, as shown in Figure 5-1 for you to confirm that you want to fail the selected drive. Click *OK*. The GUI displays FLD status under the icon for that drive. The amber light on the drive flashes, becomes solid, and the drive spins down.

Figure 5–1 Fail Drive Confirmation Window



4. Locate the disk with the solid amber light in the storage enclosure and remove it.
5. Insert a replacement disk into the same slot of the storage enclosure from which you removed the disk that you failed.
If the drive that you failed contained a RAID 1, 5, or 0 + 1 redundant configuration go to section 5.2.3.1 to rebuild the drive. If the drive you failed contained RAID 0, or JBOD non-redundant configurations go to Section 5.2.3.2 to mark the drive optimal and restore data from backup.

5.2.3.1 Rebuilding a Drive

Use Rebuild a Drive when replacing a failed drive which contained a RAID 1, 5, or 0 + 1 redundant configuration. The rebuild operation reconstructs the data which the failed drive had contained onto the replacement drive.

To rebuild a drive, after replacing a drive you failed in the storage enclosure with a drive in good working order, follow these steps:

1. Select *Rebuild Drive* from the *Management* menu. *Rebuild Drive* appears as the *Current Selection* at the top of the base window.
2. Click on the icon of the drive with a status of FLD (failed) that you replaced, and then click *OK*. A window appears, similar to the one shown in Figure 5-2, for you to confirm that you want to rebuild the selected drive. Click *OK*.

Figure 5-2 The Rebuild Drive Confirmation Window



3. A window appears, similar to the one shown in Figure 5-3, with a slider bar showing the progress of the rebuild on the current logical drive. The controller rebuilds each logical RAID drive associated with this disk drive. A message appears, as shown in Figure 5-4 indicating when the controller finishes rebuilding all of the logical RAID drives. Click *OK* and the base window appears.

Figure 5-3 The Rebuild Drive Confirmation Window


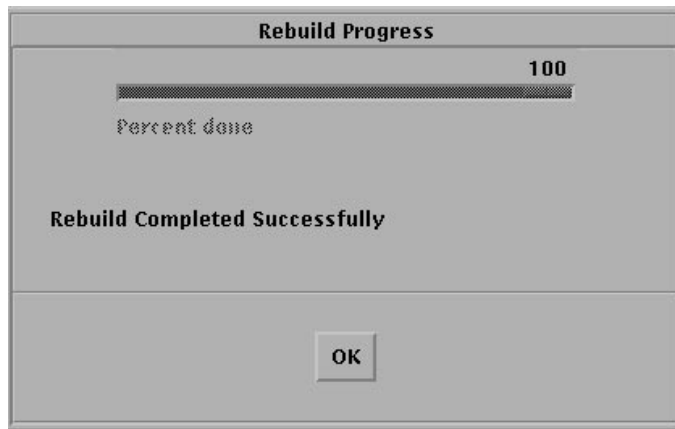


Figure 5-4 The Rebuild Completed Message



Note

The status of the drive changes to WOL (write-only) during the rebuild. Upon completion of the rebuild, the status of the drive and all logical drives associated with it will change to optimal.

5.2.3.2 Marking a Drive as Optimal

Use the Mark as Optimal option when replacing a failed drive which contained RAID 0, or JBOD non-redundant configurations. Mark as Optimal spins up a drive and changes the status of the drive to optimal.

Note

Mark as Optimal does not fix any problems with a drive; it merely makes the system use the drive. If a drive is prone to fail, it will fail again. Also, data will not be written to the drive when it is made optimal. If you find errors with a drive, fail it and replace it.

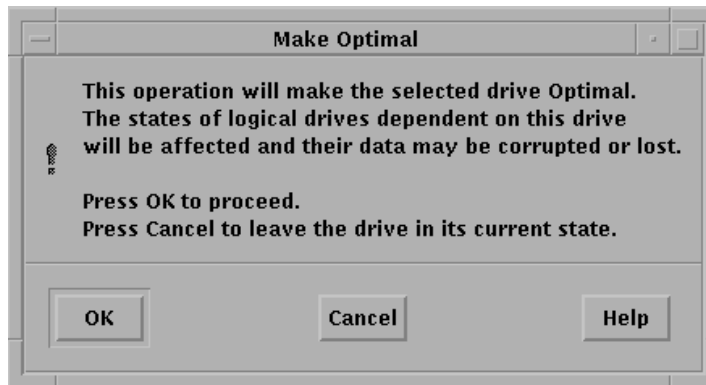
You can only mark optimal a drive with a status of FLD. If you have a drive with a status of WOL, you must first change its status to FLD, before you can mark it optimal. See Section 5.2.3 to fail a drive.

To make drives with a status of HSP or RDY optimal, you must use the stand-alone configuration utility. See the *StorageWorks RAID Array 200 Subsystem Controller Installation and Standalone Configuration Utility User's Guide*.

To make a drive optimal, after replacing a suspect/bad drive in the storage enclosure with a drive in good working order, follow these steps:

1. Select *Make Optimal* from the *Management* menu. *Make Optimal* appears in *Current Selection* at the top of the base window.
2. Click on the icon of the drive with a status of FLD (failed) that you replaced, and then click *OK*. A window appears, similar to the one shown in Figure 5-5, for you to confirm that you want to mark the selected drive as optimal. Click *OK*. The status of the drive changes to OPT. The amber light on the drive, lit when the drive is FLD, will go off when the drive becomes optimal. The status of all the logical drives contained in this drive also change to optimal.

Figure 5-5 The Make Optimal Confirmation Window



5.3 Making a Hot Spare

Make one or more hot spare disk drives on which the RAID controller can automatically start to rebuild information when a drive fails which contained a RAID 1, 5, or 0 + 1 redundant configuration.

You can make a hot spare out of any disk drive, which is not in a drive group. Use a disk drive with a capacity equal to or greater than the largest disk capacity of any drive in a group.

To make a hot spare, follow these steps:

1. Select *Make Hot Spare* from the Management menu. *Make Hot Spare* appears in *Current Selection* at the top of the base window.
2. Click on any drive icon with a status of RDY (ready), and then click *OK*. The drive status changes to HSP (Hot Spare). The controller will now use this drive as a hot spare, if a drive configured into a drive group fails.

If you decided that you no longer want to use a disk as a hot spare, you can return its status to ready. To return a hot spare disk to ready status, follow these steps:

1. Select *Make Hot Spare* from the Management menu. *Make Hot Spare* appears in *Current Selection* at the top of the base window.
2. Click on any drive icon with a status of HSP (Hot Spare), and then click *OK*. The drive status changes to RDY (ready). The drive is now available for you to use in a drive group (through the standalone configuration utility), or to remove from the subsystem.

5.4 Setting Rebuild Rate

Use the Default Rebuild Rate option to select how quickly a rebuild or parity check will complete.

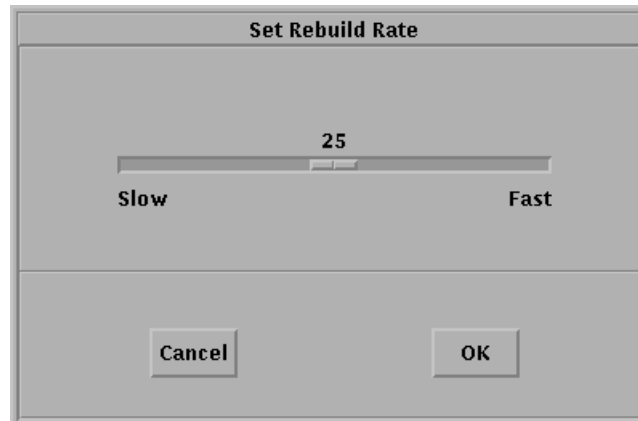
A rebuild or parity check will slow down system I/O performance and system I/O will slow down rebuild or parity check performance. Therefore, you can use the rebuild rate function to raise or lower the priority of a rebuild or parity check and thus speed up or slow down its performance.

You can assign a rebuild rate value between 0 and 50. The slowest rate is 0 (most priority is given to I/O and least to the rebuild). The fastest rate is 50 (more priority given to the rebuild). 50 is the default rebuild rate.

To change the rebuild and parity check rate, follow these steps:

1. Select *Default Rebuild Rate* from the *Management* menu. A pop-up window will appear with the current rebuild rate displayed on a slider bar as shown in Figure 5–6.

Figure 5–6 Set Rebuild Rate Window



2. To exit without changing the rate, click the Cancel button in the pop-up window. To change the rate, drag the slider bar with the mouse to the value desired, and then click OK. The base window appears.

Note that you can change the rebuild rate while a rebuild or parity check is in progress. If you find that the rebuild or check is taking longer than desired, you can increase the rebuild rate, or if you find that system I/O performance is suffering because of the rebuild or check, you can lower the rate.

5.5 Viewing the Bad Block Table

The controller keeps track of two bad block tables, one for rebuild and one for write-back cache.

The *rebuild bad block table* tracks the location of areas on a physical disk which are inaccessible during a rebuild, or contain incorrect data during a parity check.

The *write-back bad block table* tracks the location of areas on a physical disk where the controller encounters errors writing data.

The entries in either bad block table show the logical drive containing the error, the location (starting logical block address) of the error, and the length (number of sectors) of the error. Remember this table displays the logical block addresses for the logical drive, not the physical disk.

The GUI automatically displays the bad block table upon completion of a parity check/repair if it found errors. At other times, to view the bad block table, follow these steps:

1. Select *Bad Block Table* from the *Management* menu. The *Bad Block Table* window appears, as shown in Figure 5–10.
2. Select either *Rebuild* or *Write Back* from the button on the right side of the table. A list will appear (with scroll bar if necessary) showing any errors by logical drive, starting block, and number of blocks bad.
3. Click *OK* in the *Bad Block Table* window to close it when you finish viewing.

5.6 Checking Parity

Run Parity Check to validate parity information for logical drives with parity. Logical drives on the subsystem with parity include RAID 1, RAID 0 + 1, and RAID 5 logical drives.

The Parity Check detects any internal inconsistencies between data and parity information, and optionally makes repairs.

To run Parity Check, follow these steps:

1. Select *Parity Check* from the *Management* menu. The *Parity Repair Option* Window appears, as shown in Figure 5–7.

Figure 5–7 The Parity Repair Option Window



2. Select either the *Repair Parity* or *No Repair* mode by clicking on the appropriate button, and then clicking *OK*.
3. Click on the Logical drive for which you want to check parity and click *OK*. The *Parity Check Progress* Window appears, as shown in Figure 5–8, indicating parity check started, and percent completed.
4. If the utility detected no errors, when parity check completes, a screen similar to the one shown in Figure 5–9 appears. Click *OK* to close the *Parity Check Progress* window.
If the utility detects errors, when parity check completes, a bad block table appears as shown in Figure 5–10.

Figure 5–8 The Parity Progress Window




Figure 5–9 Parity Check Completes Message


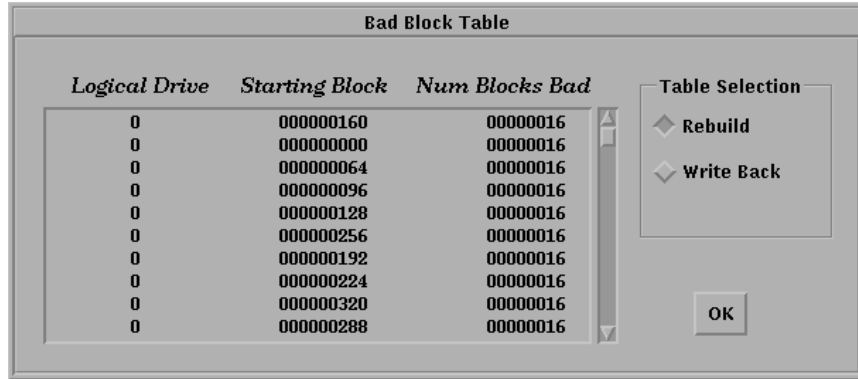
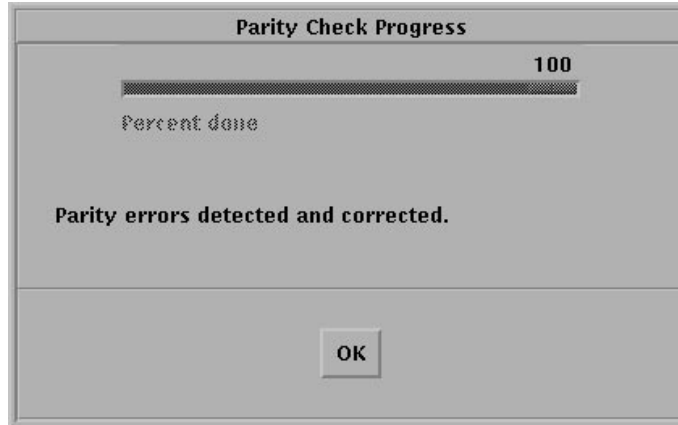


Figure 5–10 Parity Status Report Appears When Parity Check Completes



5. Click *OK* to close the bad block table after viewing it. A *Parity Check Progress* window appears, as shown in Figure 5–11.

Figure 5–11 Parity Errors Detected Message



6. Click *OK* to close the *Parity Check Progress* window.

A

accessing array, 1–9
action buttons, 4–3
adapter selected, 4–7

B

bad block
 notification of, 2–8
 table, 5–11

C

channel, 4–11
command line, 1–1
 files, 1–3
 loading, 1–4
 monitoring, 2–1
 parity check and repair, 2–7
configuration
 printing, 4–17
 redundant, 5–3, 5–5
controller
 numbers, 1–9
 choosing, 4–7
current function selected, 4–3

D

Digital UNIX
 preparing logical RAID drives for
 access by, 1–14
devices
 create special files for, 1–9
 MAKEDEV, 1–11
 mounting, 1–4
device special files
 creating, 1–9, 1–11
 determining controller numbers, 1–9
 determining names, 1–10
disk drives
 hot spare, 5–3, 5–9
 icons, 4–11
 marking as failed, 5–4
 matrix, 4–12
 number, 4–12
 rebuilding, 5–6
 replacing, failed, 5–2
 SCSI id, 4–11
 status, 4–13
 viewing, 4–10
disklabel, 1–14
drive group designator, 4–12

E

expansion slot, 4–7

F

filesystem, 1–14, 1–16
flags, 2–2
function selected, 4–3

G

graphical user interface (GUI), 1–1
 accessing, 4–2
 exiting, 4–6
 files, 1–3
 help, 4–6
 loading, 1–4
 menus, 4–6
 parity check, 5–12
 overview, 4–1
 using, 4–4
 window elements, 4–5

H

hot spare, 5–3
 making, 5–9
hot swap, 5–2, 5-3, 5-4

I

Installing the utility, 1–2, 1–4
 sample, 1–5

L

logical RAID drive
 creating filesystem on, 1–16
 creating partitions on, 1–14
 drive group, 4–15
 icon, 4–15
 number, 4–15
 preparing for access by Digital UNIX
 filesystem, 1–14
 size, 4–15
 status, 4–15, 4–16
 viewing, 4–14
 write policy, 4–15

M

marking drive
 failed, 5–4
 optimal, 5–8
MAKEDEV, 1–11
monitoring, 2–1
 invoking monitor, 2–1
 messages, 3–1
 status update, 2–5
Motif, 4–4
mount device, 1–4
mounting filesystem, 1–16

N

newfs, 1–16
notification of events
 by monitor, 2–6
 repeating, 2–7
 to log files, 2–3
 to mail, 2–4
 to system console, 2–3

P

parity

- check, 5–12
- check and repair, 2–7
- flags, 2–7
- messages, 3–8
- notification of bad blocks, 2–8
- repairing, 2–8, 5–12

partitions, 1–14 to 1–16

printing configuration, 4–17

R

RAID, 4–15, 5–3, 5–5

rebuilding, 5–6

rebuild rate

- viewing, 2–10, 5–10
- setting, 2–11, 5–10

reconfiguring the kernel, 1–12

replacing failed drive 5–2

S

setld utility, 1–4

- l, 1–4
- v, 1–7

StorageWorks RAID Array 200 Online Management utility for Digital UNIX

- files, 1–3
- GUI overview, 4–1
- loading, 1–4
- installing, 1–2
- post installation, 1–8
- verifying installation, 1–7

subsystem

- 1-channel, 4–2, 4–10
- 3-channel, 4–3, 4–12
- maintaining, 5–1
- monitoring, 2–1
- performing parity check, 2–7
- selecting, 4–7
- viewing from GUI, 4–7

swxcr.asc, 4–17

swxcrcheck, 1–3, 2–8

swxcrmgr, 1–3

swxcrmon, 1–3, 2–1

- automatic starting, 1–8
- flags, 2–2
- messages, 3–1
- status update frequency, 2–5

swxcrrate, 1–3, 2–10, 2–11

V

viewing

- bad block table, 5–11
- disk drives, 4–10
- logical RAID drive, 4–14
- subsystem, 4–7

verifying installation, 1–7

X

x11startup, 1–8

Reader's Comments

Manual Order Number:

Digital is committed to providing the best products and services. Since our manuals are important components of our products, we value your comments, corrections, and suggestions for improvements. Please take a few minutes to fill out and return this form. Attach additional sheets, if necessary. Thank you.

Manual Rating	Excellent	Good	Fair	Poor
Accuracy (correct presentation of facts)	[]	[]	[]	[]
Completeness (adequate information)	[]	[]	[]	[]
Clarity (easy to understand)	[]	[]	[]	[]
Organization (logical sequence of information)	[]	[]	[]	[]
Layout (easy to follow subject matter)	[]	[]	[]	[]
Indexing (easy to locate desired information)	[]	[]	[]	[]

Errors Noted (include page, paragraph, table or figure number) _____

Most-Liked Features: _____

Least-Liked Feature _____

Suggestions for Improvement _____

Return Address: Digital Equipment Corporation Customer Research Response Center ATTN: Nan Andrews Amish 334 South Street, SHR3-2/W3 Shrewsbury, MA 01545	Name	Phone	
	Title		
	Company		
	Street Address		
	Mail Stop		
	City	State	ZIP
	Country (if other than USA)		

