# StorageWorks RAID Array 200 Online Management Utility for OpenVMS AXP

## User's Guide

AA–Q6WVB–TE

**Digital Equipment Corporation**
**Maynard, Massachusetts**

**Second Printing, April 1995**

# *Contents*

## 2 Using the RAID Array 200 Online Management Utility for OpenVMS Command Line



## 3 Command Line Messages



## 4 Using the RAID Array 200 Online Management Utility for OpenVMS Graphical User Interface

## Figures



## Tables

**Tables** (cont.)

# 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 OpenVMS AXP operating system environment.

| Revision Level | Date | Summary of Changes |
|---|---|---|
| AA–Q6WVA–TE | August 1994 | Original release. This release supports the StorageWorks RAID Array 200 Online Management Utility for OpenVMS AXP command line only. |
| AA–Q6WVB–TE | April 1995 | This release adds supports for the StorageWorks RAID Array 200 Online Management Utility for OpenVMS AXP Graphical User Interface (GUI).<br><br>**About Figures In This Revision**<br><br>In this revision of the guide figures displaying 3-channel controllers do not reflect the current version of the GUI. Compare these figures with the figures showing 1-channel controllers, to view how the title and menu bars actually appear on the current version of the GUI. |

# 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 OpenVMS User's Guide

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

## Intended Audience

This guide supports users of the StorageWorks RAID Array 200 Subsystems in an OpenVMS environment on an AXP system. You should be familiar with StorageWorks RAID terminology, OpenVMS 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 OpenVMS**

Installing the RAID Array 200 Online Management Utility for OpenVMS describes features of the utility; lists supported systems and files added to your system by this installation; loading the utility onto boot devices in systems running OpenVMS V6.1 and V6.2; loading the utility onto data devices in systems running OpenVMS V6.1 and V6.2; and preparing a logical RAID drive (used as a system device or boot device) for shadowing on a system running OpenVMS V6.2.

**Chapter 2: Using the RAID Array 200 Online Management Utility for OpenVMS Command Line**

Using the RAID Array 200 Online Management Utility for OpenVMS Command Line discusses monitoring the RAID subsystem through the SWXCR monitor; and using the command line to perform parity check and repair, and to work with drives to: make hotspare drives, fail drives, rebuild drives and make drives optimal.

**Chapter 3: Messages**

Messages describe the informational, and error messages displayed by the command line monitor, parity check and work-with-drive processes.

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

Using the RAID Array 200 Online Management Utility for OpenVMS V6.2 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 |
| *italic type* | For emphasis, manual titles, utilities, menus, screens, and filenames |

## Figures

In this revision of the guide figures displaying 3-channel controllers do not reflect the current version of the GUI. Compare these figures with the figures showing
1-channel controllers, to view how the title and menu bars actually appear on the current version of the GUI.

## 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 Online Management Utility for OpenVMS*

*This chapter presents information about the systems supported, the kit contents, loading the online management utility for OpenVMS onto systems running OpenVMS V6.1 and V6.2, and preparing logical RAID drives on systems running OpenVMS V6.2.*

## 1.1 Introduction

The RAID Array 200 Online Management Utility for OpenVMS varies for systems running OpenVMS V6.1 and V6.2.

The online management utility offers command line and extended command line utilities for systems running OpenVMS V6.1 and V6.2.

The command line provides:

- OpenVMS AXP driver for the controller
- Monitoring
- Parity check and repair

The extended command line provides the following working-with-drives commands:

- Making a hot spare
- Failing a drive
- Rebuilding a drive
- Viewing help

The online management utility offers a Graphical User Interface (GUI) only for systems running OpenVMS V6.2.

The GUI utility provides:

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

In addition, OpenVMS V6.2 offers the capability to shadow logical RAID drives.

## 1.2 Before Installing the Utility

Prior to installing the RAID Array 200 Online Management Utility for OpenVMS onto your system, 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 OpenVMS

## 1.3 Contents of the Kit

The contents of the kit varies for systems running OpenVMS V6.1 and V6.2.

Table 1–1 lists and describes the contents of the online management utility files on the OpenVMS V6.1 distribution kit CD-ROM.

Table 1–2 lists and describes the contents of the online management utility files on the OpenVMS 6.1 update kit floppy diskette.

Table 1–3 lists and describes the contents of the online management utility files on the OpenVMS V6.2 distribution kit CD-ROM.

Table 1–4 lists and describes the contents of the online management utility files on the OpenVMS V6.2 update kit floppy diskette.

**Table 1–1 Online Management Utility Files Contained on the OpenVMS V6.1 Distribution Kit CD-ROM**

| File | Description |
|---|---|
| SWXCR_E$STARTUP.COM | Startup command file for auto-configuration |
| SWXCR_ECONFIG.COM | Utility command file for configuration |
| DEC$SW$DRDRIVER.EXE | OpenVMS device driver for SWXCR-E |
| DEC$SW$ICBM_09.EXE | Auto-configuration module for SWXCR-E |
| SWXCR_EWRITE_3PB.COM | Utility command file for SWXCR-E boot support |
| SWXCR$MON_INIT.EXE | Initialization executable for status monitor |
| SWXCR$MONITOR.EXE | Status monitor executable |
| SWXCR$PCHECK.EXE | Parity check executable |
| SWXCR$CLD.CLD | Command-line definitions for monitor/check |
| SWXCR$HELP.HLP | Help text for monitor/check |
| SWXCR_E.RELEASE_NOTES | Release Notes |

**Table 1–2 Online Management Utility Files Contained on the OpenVMS V6.1 Upgrade Kit Floppy Diskette**

| File | Description |
|---|---|
| SWXCR$UTIL.EXE | Executable for extended utilities (working-with-drive commands) |
| SWXCR$CLD.CLD | Command-line definitions for monitor/check |
| SWXCR$HELP.HLP | Help text for monitor/check |
| SWXCR_E.RELEASE_NOTES | Release Notes |

**Table 1–3 Online Management Utility Files Contained on the OpenVMS V6.2 Distribution Kit CD-ROM**

| File | Description |
|---|---|
| SWXCR_E$STARTUP.COM | Startup command file for auto-configuration |
| SWXCR_ECONFIG.COM | Utility command file for configuration |
| DR$INIT.COM | Command file for forced-error initialization (support for system-disk shadowing) |
| DR$INIT.EXE | Executable for forced-error initialization (support for system-disk shadowing) |
| SWXCR$CLD.CLD | Command-line definitions for monitor/check |
| SWXCR$HELP.HLP | Help text for monitor/check |
| SWXCR_E.RELEASE_NOTES | Release Notes |

**Table 1–4 Online Management Utility Files Contained on the OpenVMS V6.2 Upgrade Kit Floppy Diskette**

| File | Description |
|---|---|
| SWXCR$UTIL.EXE | Executable for extended utilities (working-with-drive commands) |
| SWXCR$MON_INIT.EXE | Initialization executable for status monitor |
| SWXCR$MONITOR.EXE | Status monitor executable |
| SWXCR$PCHECK.EXE | Parity check executable |
| SWXCR$RRATE.EXE | Executable for SET/SHOW rebuild rate |
| SWXCR$MGR.EXE | Executable for GUI-based status/config utility |
| SWXCR$MGR.DAT | Resource file for SWXCRMGR.EXE |
| SWXCR$CLD.CLD | Command-line definitions for monitor/check |
| SWXCR$HELP.HLP | Help text for monitor/check |
| SWXCR_E.RELEASE_NOTES | Release Notes |

## 1.4 Loading the Online Management Utility

The procedure for loading the online management utility for OpenVMS AXP varies depending on the following:

- Version of OpenVMS (on which you plan to load it)
- Whether you will load it onto a system disk or data device
- Whether you will load the command line only, or the command line and the extended command line utility

See Table 1–5 to determine which section of this chapter to see to load the online management utilities for OpenVMS.

**Table 1-5 Installation Roadmap**

| To load the utilities onto: | See: |
|---|---|
| OpenVMS V6.1 Boot Device | Section 1.4.1.1 |
| OpenVMS V6.1 Data Device | Section 1.4.1.2 |
| OpenVMS V6.2 Boot Device | Section 1.4.2.1 |
| OpenVMS V6.2 Data Device | Section 1.4.2.2 |

*Boot device* indicates the storage device in your subsystem where you located the operating system. You boot your system from this device.

*Data device* indicates a nonbootable storage device in your subsystem. A data device does not contain your operating system. You cannot boot your system from this device.

### 1.4.1 Systems Running OpenVMS V6.1

The online management utility offers only the command line components in the OpenVMS V6.1 environment.

#### 1.4.1.1 Boot Device Installation -- OpenVMS V6.1

To load the online management utility onto a RAID Array 200 Subsystem used as an OpenVMS V6.1 boot device, you need to do the following:

- Identify devices for OpenVMS
- Install OpenVMS
- Install the command line utility
- Optionally, install the extended command line utilities

##### *1.4.1.1.1 Identify Devices for OpenVMS V6.1*

To identify devices for OpenVMS V6.1, follow these steps:

1. Refer to the *StorageWorks RAID Array 200 Subsystem Controller Installation and Standalone Configuration Utility User's Guide* for information on installing the RAID controller and configuring a logical RAID drive.
2. Once the RAID controller's installation/configuration has been successfully completed, boot the system from the OpenVMS AXP CD-ROM V6.1 distribution media.
3. To identify devices for OpenVMS, highlight option 2 from the OpenVMS Installation menu and press the Enter key. A DCL ($) prompt appears.
4. At the DCL ($) prompt, type: `$ @[SWXCR_E$KIT]DRCONNECT`
   Then press the Enter key. OpenVMS displays the following prompt:
   `Enter the slot number for the SWXCR-Ex controller:`
5. Enter the slot number for your controller (displayed when you ran the EISA Configuration Utility to install the controller) and press the Enter key. OpenVMS displays the following prompt:
   `Enter the IRQ for the SWXCR-Ex controller [11]:`
6. Press the Enter key to accept 11 as the IRQ, or type another number for the IRQ (displayed when you ran the EISA Configuration Utility to install the controller) and press the Enter key.
   This completes the device identification for OpenVMS.
7. To return to the OpenVMS Installation menu, type: `$ LOGOUT`

#### *1.4.1.1.2 Installing OpenVMS V6.1*

To install OpenVMS V6.1, choose option 1 from the OpenVMS Installation menu. Specify the DRAn: as the target installation device. Follow the series of prompts and provide the appropriate information. When the OpenVMS installation completes, return to the Installation menu.

**NOTE**

The DRAn: (n=0-7) (target installation) device must be a logical RAID drive that you have previously configured using the Standalone Utility.

#### *1.4.1.1.3 Invoking the PCSI Load Procedure*

To invoke the PCSI load procedure, which loads the online management utility, follow these steps:

1. Choose option 2 from the Installation menu. A DCL ($) prompt appears.
2. Invoke the PCSI load procedure, type:
   $ **@[SWXCR_E$KIT]INVOKE_PCSI**
   The following prompt: `Enter target device name:`
3. Type the name of the target system disk (DRAn:) and press the Enter key. PCSI load procedure start messages appear, then PCSI displays the prompt:
   `Do you want all the default values for this product? [Yes]`
4. Press the Enter key to accept default values. PCSI displays the prompt:
   `Do you want to view the values? [No]`
5. Press the Enter key and PCSI loads the online management utility.
6. Return to the OpenVMS Installation menu when PCSI completes, type:
   $ **LOGOUT**
7. Highlight option 3 and press the Enter key to shut down the system.
8. Reboot the system directly from the DRAn: device.

This completes the installation of the online management utility command line onto a system disk. Go to Section 1.4.1.3 if you want to load the extended command line utilities.

**NOTE**

If you copy the contents of a system disk to your subsystem using, for example, BACKUP/IMAGE, then you must repeat Steps 5 through 8 listed under the *Data Device Installation* section in this chapter identifying the second disk as the target device.

### 1.4.1.2 Data Device Installation -- OpenVMS V6.1

To load the online management utility onto a data device, follow these steps:

1. Boot OpenVMS AXP from your system disk.
2. Log into your system account.
3. Insert the OpenVMS AXP CD V6.1 into your CD-ROM drive.
4. Mount the volume, type: `$ MOUNT DEVICE:/OV=ID`
   where *DEVICE* is the CD-ROM drive
5. Invoke the PCSI load procedure, type:
   `$ @DEVICE:[SWXCR_E$KIT]INVOKE_PCSI.COM`
   where *DEVICE* is the CD-ROM drive
   PCSI displays the following prompt: `Enter target devicename:`
6. Enter the name of the target system disk (DKn:). PCSI loads the utility.
7. When the PCSI load procedure completes, add the following line to your system's startup file to ensure that the system recognizes the controller each time you boot the system: `$ @SYS$MANAGER:SWXCR_E$STARTUP`
8. Reboot from the OpenVMS system disk to access the SWXCR DRAn: devices.

This completes the installation of the online management command line utility onto a data device. Go to Section 1.4.1.3 if you want to load the extended command line utilities.

#### 1.4.1.3 Extended Command Line Utilities -- OpenVMS V6.1

To load the extended command line utilities onto a RAID Array 200 Subsystem used as either a system disk or data device, follow these steps:

1. Insert the OpenVMS Extended Utilities diskette into your system's floppy drive.
2. Mount the floppy diskette, type: `$ MOUNT/OV=ID DVA0`
3. Load the extended command line utilities, type:
   `$ @DVA0:[SWXCR_E$KIT]INSTALL.COM`
   The install utility loads the extended utilities.
4. Unmount the floppy diskette and remove it from the floppy drive.

This completes the installation of the online management utility for OpenVMS V6.1.

### 1.4.2 Systems Running OpenVMS V6.2

The online management utility offers the command line and Graphical User Interface (GUI) components in the OpenVMS V6.2 environment.

#### 1.4.2.1 Boot Device Installation -- OpenVMS V6.2

To load the online management utility onto a RAID Array 200 Subsystem used as an OpenVMS V6.2 boot device, you need to do the following:

- Optionally, initialize device for shadowing, if you plan to shadow it
- Install OpenVMS 6.2 to a logical RAID drive
- Install the command line utility component
- Optionally, install the extended command line and GUI components

##### *1.4.2.1.1 Initializing a Boot Device for Shadowing*

OpenVMS offers device shadowing. If you plan to shadow your boot device, got to Section 1.5 prior to installing OpenVMS V6.2 and the online management utility onto your boot device.

#### *1.4.2.1.2 Installing OpenVMS V6.2*

To install OpenVMS V6.2 onto a logical RAID drive, follow these steps:

1. Refer to the *StorageWorks RAID Array 200 Subsystem Controller Installation and Standalone Configuration Utility User's Guide* for information on installing the RAID controller and configuring a logical RAID drive.
2. Once the RAID controller's installation/configuration has been successfully completed, boot the system from the OpenVMS AXP CD-ROM V6.2 distribution media.
3. To install OpenVMS V6.2, choose option 1 from the OpenVMS Installation menu. Specify DRAn: as the target installation device.

> **NOTE**
>
> The target installation device: DRAn: (n=0-7) must be a logical RAID drive that you have previously configured using the Standalone Utility.

   Follow the series of prompts and provide the appropriate information.
4. The OpenVMS installation menu appears when the installation completes.

#### *1.4.2.1.3 Invoking the PCSI Load Procedure*

Invoke the PCSI load procedure to load the online management utility.

To invoke the PCSI load procedure, follow these steps:

> **NOTE**
>
> When upgrading from OpenVMS 6.1 to OpenVMS 6.2, select Yes to remove old SWXCR files, at the install program prompt to remove previous version.

1. Boot the system from the OpenVMS AXP CD-ROM V6.2 distribution media, if you have not already done so in Section 1.4.2.1.2.
2. Choose option 4 from the OpenVMS V6.2 Installation menu. A DCL ($) prompt appears.
3. Invoke the PCSI load procedure, type:
   `$ @[SWXCR$KIT]INVOKE_PCSI`
   The following prompt: `Enter target device name:`
4. Type the name of the target system disk (DRAn:) and press the Enter key. PCSI load procedure start messages appear, then PCSI displays the prompt:
   `Do you want all the default values for this product? [Yes]`
5. Press the Enter key to accept default values. PCSI displays the prompt:
   `Do you want to view the values? [No]`
6. Press the Enter key and PCSI loads the online management utility.
7. Return to the OpenVMS Installation menu when PCSI completes, type:
   `$ LOGOUT`
8. Highlight option 5 and press the Enter key to shut down the system.
9. Reboot the system directly from the DRAn: device.

If you plan to load the extended command line and the GUI components of the utility, go on to Section 1.4.2.3. Otherwise, this completes the installation of the RAID array 200 Online Management Utility for OpenVMS onto a system disk.

### 1.4.2.2 Data Device Installation -- OpenVMS V6.2

To load the online management utility onto a RAID Array 200 Subsystem used as an OpenVMS V6.2 data device, you need to do the following:

- Optionally, initialize device for shadowing, if you plan to shadow it
- Install the command line utility component
- Optionally, install the extended command line and GUI components

#### *1.4.2.2.1 Initializing a Data Device for Shadowing*

OpenVMS offers device shadowing. If you plan to shadow your data device, go to Section 1.6 prior to loading data and the online management utility onto your data device.

#### *1.4.2.2.2 Invoking the PCSI Load Procedure*

To load the online management utility command line component onto your system, follow these steps:

1. Log into your system account.
2. Insert the online management utility floppy diskette into the floppy diskette drive.
3. Mount the volume, type: `$ MOUNT DEVICE:/OV=ID`
   where DEVICE is the CD-ROM drive
4. Invoke the PCSI load procedure, type:
   `$ @DEVICE:[SWXCR$KIT]INVOKE_PCSI.COM`
   where DEVICE is the floppy diskette drive
   PCSI displays the following prompt: `Enter target devicename:`
5. Enter the name of the target system disk (DKn:). PCSI loads the utility.

If you plan to load the extended command line and the GUI components of the utility, go on to Section 1.4.2.3. Otherwise, this completes the installation of the RAID array 200 Online Management Utility for OpenVMS onto a data disk.

#### 1.4.2.3 Extended Command Line and GUI -- OpenVMS V6.2

To load the online management utility extended command line and GUI components onto a RAID Array 200 Subsystem used as either a system disk or data device, follow these steps:

1. Insert the OpenVMS Extended Utilities diskette into your system's floppy drive.
2. Mount the floppy diskette, type: `$ MOUNT/OV=ID DVA0`
3. Load the extended command line utilities, type: `$ @DVA0:[SWXCR$KIT]INVOKE_PCSI`
   The install utility loads the extended utilities.
4. Unmount the floppy diskette and remove it from the floppy drive.

This completes the installation of online management utility for OpenVMS 6.2.

## 1.5 Shadowing a Logical RAID Drive

To shadow a logical RAID drive, you must reserve a portion of the drive for forced-error information. Shadowing uses approximately one 512-byte block for each 2 Megabytes of disk capacity. Therefore, once shadowed, your drive will appear to have less than its actual, raw capacity.

Use the DR$INIT utility found on the extended utilities diskette to reserve a portion of the Logical RAID drive for the forced error information. Back up the data on any logical RAID drive you intend to prepare for shadowing, because the DR$INIT utility re-initializes the drive, destroying any data on it.

How you prepare a logical RAID drive for shadowing depends on whether you plan to use the logical RAID drive as a system or data drive.

To prepare a logical RAID drive used as a system drive for shadowing, see Section 1.5.1.

To prepare a logical RAID drive used as a data drive for shadowing, after you load the online management utility for OpenVMS, go to Section 1.5.2.

### 1.5.1 Preparing a Logical RAID Drive as System Drive for Shadowing -- OpenVMS V6.2

To prepare a Logical RAID drive used as a system drive for shadowing, follow these steps:

1. Boot your system from the OpenVMS AXP distribution CD-ROM.
2. When the system boots, select the *Execute DCL commands and procedures* option from the *Installation* menu.
3. Invoke the DR$INIT utility, at the DCL prompt, type:

   $ **@DKxx:[SWXCR$KIT]DR$INIT DRyy:**

   where *DKxx* specifies the device name of your CD-ROM
   where DRyy specifies the Logical RAID drive on which you will load your system files.

   The utility will destroy any data already on this drive, and prompts you to go ahead with the procedure.
4. Specify to continue with this procedure. The utility prompts you for a disk label.
5. Enter a disk label (each disk in a shadow set must have the same label) and press the Enter key. The DR$INIT utility runs.
6. When the DR$INIT utility completes, logout; at the DCL prompt type:

   $ **LOGOUT**

   The installation menu appears.
7. Go to Section 1.4.2.1 to complete the online management utility installation onto an OpenVMS V6.2 boot device.

### 1.5.2 Preparing a Logical RAID Drive as Data Drive for Shadowing -- OpenVMS V6.2

Prior to preparing for volume shadowing a logical RAID drive used as a data device, install the online management utility for OpenVMS as described in Section 1.4.2.

To prepare a Logical RAID drive used as a data drive for shadowing, follow these steps:

1. Boot from your system disk.
2. Invoke the DR$INIT utility, type:

   $ @SYS$MANAGER:DR$INIT DRxx:

   where DRxx specifies the Logical RAID drive on which you will load your system files.

   The utility prompts you to reinitialize the drive.
3. Select yes and press the Enter key to reinitialize the drive. The utility prompts you for a disk label for this logical RAID drive.
4. Type a disk label and press the Enter key. The utility initializes the drive. When the utility completes, you can then use the drive for a shadow set.
5. Go to Section 1.4.2.2 to complete the online management utility installation onto an OpenVMS V6.2 data device.

# 2

# *Using the RAID Array 200 Online Management Utility for OpenVMS Command Line*

*This chapter presents information about invoking the monitor utility and invoking the parity utility to check and repair the subsystem.*

## 2.1 Introduction

The StorageWorks RAID Array 200 Online Management Utility for OpenVMS AXP Command Line component provides the following functionality to maintain a RAID Array 200 Subsystem in an OpenVMS V6.1 or V6.2 environment:

- Monitor to notify you of subsystem status changes
- Parity Check and Repair to check the consistency of a logical RAID drive

In addition the online management utility provides the following working-with-drives functionality, if you loaded the extended command line utilities:

- Making a hot spare drive
- Failing a drive
- Making a drive optimal
- Rebuilding a drive
- Help

Use these utilities while the RAID subsystem is in use to monitor and maintain its operation.

## 2.2 Monitoring the RAID Subsystem

The monitor function scans the controller for status and error messages and reports these messages to you. You must start a monitor process to receive notification of events that occur on your subsystem, and you must invoke a separate monitor process for each RAID controller you have on your subsystem.

Section 2.2.1 describes how to invoke a monitor process.

Section 2.2.2 describes the qualifiers you can use with the monitor process.

### 2.2.1 Invoking a Monitor Process

A couple of things to remember before invoking a monitor process:

- The monitor process will scan a controller only if it has at least one logical RAID drive configured.
- Mount a logical RAID device prior to invoking a monitor process, otherwise the monitor process will lock the logical RAID drive and you can not access that logical RAID drive unless you specify /SYSTEM or /SHARE as one of the qualifiers with the OpenVMS MOUNT command.

To invoke a monitor process, follow these steps:

1. View your RAID devices, type the OpenVMS show devices command as follows: **$ SHOW DEV DRA**
   OpenVMS assigns *dr* as the SWXCR controller name. The SWXCR assigns each controller it finds a different letter starting with *a*. Therefore the operating system assigns the first SWXCR controller it finds the devicename *dra*, the second *drb*, and so on.
2. Type the SWXCR command as follows:
   $ **SWXCR MONITOR devname/x**
   where devname specifies the name of your SWXCR controller, and */x* indicates any of the qualifiers listed in Table 2–1 that you want to specify when invoking this command. You do not need to specify a qualifier.
   The MONITOR command creates a separate, detached process -- a monitor process, called, SWXCR$MON_DRx, where x is the controller letter of the device specified in the SWXCR MONITOR command. A monitor process assumes the UIC, privileges, and Mail account of the process that issued the SWXCR MONITOR command.

OpenVMS does not allow you to create two processes with the same name within the same UIC group. Therefore, if you attempt to start two monitor processes for a single controller, the MONITOR command may fail with the SS$_DUPLNAM status condition.
When the monitor command successfully creates a monitor process, OpenVMS displays a message, which states the identification of the created process as in this example:

```
$ SWXCR MONITOR DRA0/INTERVAL=1
Identification of created monitor process is 000015E.
```

The monitor process displays status change messages at OPCOM by default. You can control where the monitor displays messages by entering the appropriate qualifier to turn-off messages to OPCOM, or to send messages via e-mail, or to a log file, as described in Section 2.2.2 Qualifiers for the MONITOR Command.

### 2.2.2 Qualifiers for the MONITOR Command

The MONITOR command offers the qualifiers listed in Table 2–1. More detailed descriptions of each of these qualifiers follows Table 2–1.

**Table 2–1 MONITOR Command Qualifiers**

| Qualifier | Syntax | Default |
|---|---|---|
| Notification of events to OPCOM | /OPCOM<br>/OPCOM<br>/NOOPCOM | /OPCOM |
| Notification of events to log files | /NOLOG<br>/LOG | /NOLOG |
| Notification of events to mail | /NOMAIL<br>/MAIL | /NOMAIL |
| Frequency of Status Updates (in minutes) | /INTERVAL = x | /INTERVAL = 5 |
| Repeating notification of events | /NOREPEAT<br>/REPEAT | /NOREPEAT |

***/OPCOM***
***/NOOPCOM***

The utility automatically sends notification of subsystem events to OPCOM by default. Use /NOOPCOM to stop the utility from sending notification of subsystem events to OPCOM. Use /OPCOM to once again start the utility sending notification of subsystem events to OPCOM.

***/LOG***
***/NOLOG***

Use /LOG if you want the monitor to write notification of subsystem events to a product-specific log file.

The utility creates a log file in the directory SYS$COMMON:[SWXCR$LOG] (created by the load utility procedure) with a filename as follows:
SWXCR$MON_'*device-name*'.LOG
where *device-name* is the *drx* device name.
For example, the log file for all *dra* devices would be:
SYS$COMMON:[SWXCR$LOG]SWXCR$MON_DRA.LOG

By default the monitor does not write notification to a product-specific log file.

***/MAIL***
***/NOMAIL***

Use /MAIL if you want the monitor to send notification of subsystem events through mail to a list of users, or to the system account on your local node. By default, the monitor process does not send mail.

To send notification of events to a user (or a group of users), enter the SWXCR MONITOR command with the /MAIL qualifier as follows:

```
$ SWXCR MONITOR DRA0/MAIL=(node::user[,node::user[,...]])
```

If you specify the /MAIL qualifier with no user name, the system sends the mail to the SYSTEM account on your local node.

**NOTE**

The utility sends Mail from the account of the user who starts the monitor process.

When you specify notification of event to mail, the utility collects events that the monitor process encounters during a scan and sends them out in one (or more) Mail message(s) listed in the order in which the monitor process encountered them. The number of mail messages the utility sends depends upon the number of events and the length of the event text.

If you specify the /MAIL qualifier, the SWXCR MONITOR command parses the mailing list specified, and sends a test message to each user on the list. If the test message fails (for example, because that user name does not exist) an error message displays and the monitor process is not created.

An example of the initial test message displays as follows:

```
From:   ABCDE::SYSTEM "SWXCR DRA Monitor Process from
        node ABCDE" 6-APR-1994 3:25:03.56
To:     SYSTEM
CC:
Subj:   SWXCR DRA Event Notification from node ABCDE
This is a test message from SWXCR$MON_DRA, the
StorageWorks SWXCR RAID array subsystem monitor
process. You will be receiving MAIL notification of
events.
```

An example of a Mail message sent from the SWXCR monitor process follows:

```
From:   ABCDE::SYSTEM "SWXCR DRA Monitor Process from
        node ABCDE"  6-APR-199413:26:06.58
To:     SYSTEM
CC:
Subj:   SWXCR DRA Event Notification from node ABCDE
The hard disk at channel 0, target 0 has been removed.
Logical RAID drive 0 is degraded.
Logical RAID drive 1 is degraded.
Logical RAID drive 2 is degraded.
The hard disk at channel 0, target 0 is being replaced
by the hard disk at channel 0, target 1. This affects
the status of logical RAID drives 0, 1, and 2.
Logical RAID drive 0 is rebuilding (24% complete).
```

***/INTERVAL***

Use /INTERVAL to change the scan frequency interval in minutes. You can specify a scan frequency of anywhere between 1 and 10 minutes. The monitor scans the controller for status and error messages at five minute intervals by default. The monitor may not scan at exactly the specified interval due to controller activity.

The controller stores a limited number of messages. When the controller reaches its message capacity, it overwrites older messages. Therefore, if you specify a relatively short scan interval, the more likely the monitor will see and report an event. If you specify a longer scan interval, the controller can overwrite a given event before the monitor can see and report it.

To change the scan frequency, enter the SWXCR MONITOR command with the /INTERVAL qualifier as follows:

$ **SWXCR MONITOR DRA0 /INTERVAL=x**
where *x* is a value of 1 to 10 minutes.

***/REPEAT***
***/NOREPEAT***

Use /REPEAT to specify for the monitor to report a recurring event only the first time or every time it occurs. By default the utility reports an event the first time it occurs only.

The monitor always reports device errors and hard disk insertion or removal events only once, even if you start repeat notification of events. Controller device events are one of these types:

- Parity error
- Hard error
- Soft error
- Miscellaneous error

### 2.2.3 Events Reported by the Monitor

Events reported by the monitor include the following subsystem conditions:

- Logical RAID drive status changed from optimal to degraded. Because a drive in a logical RAID drive with a redundant RAID level (RAID 1, RAID 0 + 1, or RAID 5) has a status of failed.
- Logical RAID drive status changed from optimal or degraded to failed. Because one disk in a logical RAID drive with a nonredundant level has a status of failed and the logical RAID drive cannot be accessed, or two or more drives in a logical RAID Drive with a redundant have a status of failed and the logical RAID drive cannot be accessed.
- Controller errors (parity error, hard error, soft error, or miscellaneous error).
- Controller failure.
- A rebuild (or standby rebuild) is in progress (including the percentage complete and affected logical RAID drives).

The monitor process may not report on events due to the following:

- The length of time specified by the /INTERVAL qualifier may cause the monitor to miss events (if too short) or may cause the events to be overwritten in the buffer (if too long), since the buffer has a limited capacity.

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

  Refer to Section 2.2.2, Frequency of Status Updates, for detailed information on this qualifier.

- A small logical RAID drive was rebuilt.

  For example, 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 monitor might not report the rebuild of the logical RAID drive because the rebuild may have started after an array scan operation and completed before the next array scan. However, the monitor will report the completion of the rebuild.

### 2.2.4 Adding the MONITOR Command to your System Startup File

You can add the MONITOR command to your system startup file to automatically start the monitor process when you boot your system.

To add the Monitor command to your system startup file, follow these steps:

1. Edit *Sys$start*, type:

   **EDIT SYSTARTUP_VMS.COM**

2. Add a line for each monitor process you want to start, type, for example:

   $ **@ SYS$MANAGER:SWXCR MON/INTER=3/LOG DRA**

   $ **@ SYS$MANAGER:SWXCR MON/INTER=3/LOG DRB**

   where the /INTER=3 qualifier determines the number of minutes between scan intervals ( see Section 2.2.2)

   where the /LOG qualifier indicates notification of events to log files ( see Section 2.2.2)

   where, DRA and DRB indicate each of the controllers that you want to monitor. (You must specify the device name of each controller you want to monitor.)

### 2.2.5 Stopping a Monitor Process

To stop end a monitor process, follow these steps:

1. Enter the OpenVMS show system command as follows, type:
   $ **SHOW SYS**
2. Locate the process ID (PID) of the monitor process.
3. Enter the OpenVMS stop command as follows, type:
   $ **STOP/id=xxx**
   where xxx is the monitor process PID

## 2.3 Performing Parity Check and Repair on the RAID Subsystem

Perform a parity check after a power failure or a system crash to verify the integrity of parity information on logical RAID drives with redundant RAID levels: RAID 1, RAID 0+1, and RAID 5.

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. Optionally, you can specify for the check function to repair stored parity information.

Section 2.3.1 describes how to invoke a check process.

Section 2.3.2 describes the qualifiers you can use with the check process.

### 2.3.1 Invoking the Parity Check Utility

To invoke the parity check utility, type the SWXCR command as follows:

$ **SWXCR CHECK *devname/x***

- where *devname* specifies the device name for the logical RAID drive on the controller. For example entering the devicename, *dra1*, would cause the utility to check the parity for logical RAID drive *1* on controller *a*.
- where */x* indicates any of the qualifiers listed in Table 2–2 that you want to specify when invoking this command. You do not need to specify a qualifier.

### 2.3.2 Qualifiers for the CHECK Command

The CHECK command offers the qualifiers listed in Table 2–2. More detailed descriptions of each of these qualifiers follows Table 2–2.

**Table 2–2 CHECK Command Qualifiers**

| Qualifier | Syntax | Default |
|---|---|---|
| Repairing logical RAID drive parity | /NOREPAIR<br>/REPAIR | /NOREPAIR |
| Notification of bad blocks to terminal | /NOLOG<br>/LOG | /NOLOG |
| Notification of bad blocks to log or other file | /OUTPUT<br>/OUTPUT = filename | /NOOUTPUT |

***/REPAIR***
***/NOREPAIR***

Use /REPAIR if you want the CHECK command to also repair logical RAID drive parity inconsistencies.

By default the CHECK command does not repair parity inconsistencies.

***/NOLOG***
***/LOG***

Use /LOG for the utility to display a list of bad blocks found by the check command to your terminal screen.

By default the utility does not display a list of bad blocks.

Sample output from a parity check process with the /LOG qualifier specified follows.

Parity check complete.

The following bad blocks were found:

| Block Number | # Bad Blocks |
|---|---|
| 0000000576[2] | 0000000016[1] |
| 0000000992 | 0000000016 |
| 0000001472 | 0000000016 |
| 0000001984 | 0000000016 |
| 0000002176 | 0000000016 |
| 0000002208 | 0000000016 |
| 0000002976 | 0000000016 |
| 0000003680 | 0000000016 |

In this case, the utility found bad blocks. The list specifies the logical RAID drive number, the block number where a bad block was found, and the number of bad blocks found at that address.

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 commands are processed by the controller.

***/OUTPUT***
***/OUTPUT = filename***

Use /OUPUT to save a log of bad block information to:

SYS$COMMON:[SWXCR$LOG]SWXCR$DR*xn*_CHECK.LOG
where *x* is the controller letter and *n* is the unit number specified in the *devname* parameter of the SWXCR CHECK command.

Use /OUTPUT = filename to save a log of bad block information to the filename specified.

If you specify both the /LOG and the /OUTPUT qualifiers, then a log the utility writes a log to the terminal and to the output log file. If the utility cannot open the specified output file, it writes the log to the terminal instead. Specifying the /LOG and /OUTPUT=tt: qualifiers produce the same effect.

## 2.4 Working With Drives

If you loaded the extended command line utilities, you can to do the following with the disk drives in your subsystem from the command line:

- Make a hot spare drive
- Fail a drive
- Rebuild a drive
- Make a drive optimal

### 2.4.1 Understanding Parameters

You need to specify the following parameters to issue each of the working-with-drive-commands. These parameters define the path to the drive that you want to act upon:

| | |
|---|---|
| devicename | The name of your SWXCR controller. The operating system assigns *dr* as the SWXCR controller name. The SWXCR assigns each controller it finds a different letter starting with a. Therefore the operating system assigns the first SWXCR controller it finds the devicename *dra*, the second *drb*, and so on. |
| channel | The SCSI bus by which drives are connected to the controller. The controller is connected to the system through the EISA or PCI bus. Enter the channel number to which the drive you want to act upon is connected to the controller. |
| target | Each drive connected to the controller has a unique target id from 0-7. Specify which drive you want to act upon by typing the target id of that drive. |

### 2.4.2 Making Hot Spare Drives

Use the *swxcr spare* command to make a hot spare of any drive not configured into a drive group. Hot spare drives remain unused in your subsystem unless a drive fails from a redundant logical RAID drive. In that case, the controller will rebuild the data from the failed drive onto one of the hot spare drives.

To make a drive a hot spare, type, for example: **`swxcr spare dra 2 1`** where *dra* is the device name of your first SWXCR controller, *2* is the channel number to which the physical drive that you want to make a spare is connected, and *1* is the target id of the same physical drive.

Or, type: **`swxcr spare`**

and the SWXCR utility prompts you as follows:

`_Device:` type the device name of the controller to which the drive that you want to make a hot spare is connected and press the Enter key.

`_Channel:` type the channel number to which the physical drive that you want to make a hot spare is connected and press the Enter key.

`_Target:` type the target id of the physical drive that you want to make a hot spare and press the Enter key.

View the monitor process for the SWXCR controller that you specified for notification of the result of this command.

### 2.4.3 Failing a Drive

Use the *swxcr fail* command to fail a drive that you want to remove from your storage enclosure.

**WARNING**

If you fail:

- a drive from a RAID 0, or JBOD nonredundant logical RAID drive configuration, the data stored on that drive will no longer be available to the system
- more than one drive at a time in a redundant configuration, your data will no longer be available to the system
- the drive on which you stored your system files, your system will no longer function

To fail a drive, type, for example: **`swxcr fail dra 2 1`**
where *dra* is the device name of your first SWXCR controller, *2* is the channel number to which the physical drive that you want to fail is connected, and *1* is the target id of the same physical drive.

Or, type: **`swxcr fail`**

and the SWXCR utility prompts you as follows:

`_Device:` type the device name of the controller to which the drive that you want to fail is connected and press the Enter key.

`_Channel:` type the channel number to which the physical drive that you want to fail is connected and press the Enter key.

`_Target:` type the target id of the physical drive that you want to fail and press the Enter key.

View the monitor process for the SWXCR controller that you specified for notification of the result of this command.

### 2.4.4 Rebuilding a Drive

Use the *swxcr rebuild* command to rebuild a failed drive, which is part of a redundant RAID 1, RAID 0+1, or RAID 5 logical RAID drive.

To rebuild a drive, follow these steps:

1. Remove the failed drive from its slot.
2. Insert a drive of equal or greater capacity to the failed drive into the slot from which you removed the failed drive.
3. Type, for example: **`swxcr rebuild dra 2 1`**
   where dra is the device name of your first SWXCR controller, 2 is the channel number to which the physical drive that you want to rebuild is connected, and 1 is the target id of the same physical drive.

Or, type: **`swxcr rebuild`**

and the SWXCR utility prompts you as follows:

`_Device:` type the device name of the controller to which the drive that you want to rebuild is connected and press the Enter key.

`_Channel:` type the channel number to which the physical drive that you want to rebuild is connected and press the Enter key.

`_Target:` type the target id of the physical drive that you want to rebuild and press the Enter key.

View the monitor process for the SWXCR controller that you specified for notification of the result of this command.

### 2.4.5 Making a Drive Optimal

Use the *swxcr mkopt* command to mark a drive optimal that you replaced in the system.

> **NOTE**
>
> Use the *mkopt* command for troubleshooting purposes only and not to rebuild a degraded redundant logical RAID drive

To make a drive optimal, type, for example: **`swxcr mkopt dra 2 1`**
where *dra* is the device name of your first SWXCR controller, *2* is the channel number to which the physical drive that you want to make optimal is connected, and *1* is the target id of the same physical drive.

Or, type: **`swxcr mkopt`**

and the SWXCR utility prompts you as follows:

`_Device:` type the device name of the controller to which the drive that you want to make optimal is connected and press the Enter key.

`_Channel:` type the channel number to which the physical drive that you want to make optimal is connected and press the Enter key.

`_Target:` type the target id of the physical drive that you want to make optimal and press the Enter key.

View the monitor process for the SWXCR controller that you specified for notification of the result of this command.

## 2.5 Viewing Help

To get online help using the StorageWorks RAID Array 200 Online Management Utility for OpenVMS, type:

**`swxcr help`**

# 3

# *Messages*

*This chapter presents information about the messages that appear from the command line and extended command line utilities.*

## 3.1 Introduction

This chapter contains messages that appear as a result of conditions that occur when running either the monitor or the parity check utilities; or the extended command line work-with-drives command.

Tables 3–1 and 3–2 list and describe the messages that appear from the Monitor process.

Table 3–3 lists and describes the messages that result from the parity utility.

Table 3–4 lists and describes the messages that result from the working -with-drive commands.

## 3.2 Messages from the Monitor Process

Table 3-1 lists messages that result from the monitor utility.

**Table 3-1 Monitor Utility 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 a SS$_CTRLERR error. The <ctlr-name> is the controller name (for example, "DRA"). |

**Table 3-1 Monitor Utility Messages (Cont.)**

| Message | Meaning |
|---|---|
| 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 EEPROM configuration was transferred, but you should run the standalone utility **???** to resolve this problem. (It is not clear whether the EEPROM configuration or the NVRAM configuration is the correct one). The <ctlr-name> is the controller name (for example, "DRA"). |
| NVRAM configuration does not match. | The monitor process determined, when reading the configuration from EEPROM on the controller, that the NVRAM configuration does not match (possibly due to a device not responding, and so on). The NVRAM configuration is transferred and the monitor process can continue. You should run the standalone utility **???** to resolve this problem. The <ctlr-name> is the controller name (for example, "DRA"). |

**Table 3-1 Monitor Utility Messages (cont.)**

| Message | Meaning |
|---|---|
| Replacement table full. Monitor can no longer report on rebuilds. Monitor process is terminating!!! | 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.<br><br>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 with an SS$_ABORT status. |

The following messages (Table 3-2) are sent by the monitor process through the reporting mechanisms that you set up (mail, log file, or OPCOM).

**Table 3-2 Monitor Utility Messages Reported**

| Message | Meaning |
|---|---|
| Internal log structures getting full.<br>OR<br>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.<br>If you see either message, shutdown 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 target ID. The <drive-list> is a list of logical RAID drives which will be rebuilt as part of the rebuild of the specified hard disk. |

**Table 3-2 Monitor Utility Messages Reported (cont.)**

| Message | Meaning |
|---|---|
| The hard disk at channel <old-chn>, target <old-tgt> 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>. | When a hot spare rebuild is in progress, the monitor reports this message. It indicates the channel and target of both the hard disk which failed and the hot spare hard disk which is replacing it.<br><br>In addition, a list of affected system drive(s) displays. The <old-chn> is the channel number of the hard disk being replaced; <old-tgt> is the target id of the hard disk being replaced; <new-chn> is the channel number of the hot spare device; <new-tgt> is the target id of the hot spare device.<br><br>The <drive-list> is a list of logical RAID drive numbers, that are the drives which have blocks contained on the disk which is being replaced. Each of these logical RAID drives will become critical and will be rebuilt as part of the rebuild of the hard disk. The monitor will report on the progress of the rebuild, including the rebuild of each logical RAID drive.<br><br>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 will report that the logical RAID drives are still degraded. |

**Table 3-2 Monitor Utility Messages Reported (cont.)**

| Message | Meaning |
|---|---|
| Logical RAID drive <drive-num> is <drive-status> | Monitor is reporting on a logical RAID drive status. If /NOREPEAT is specified, only changes in logical RAID drive status are reported. If /REPEAT is specified, any changes in status are reported, and any logical RAID drive status which is not 'on line' is reported. <drive-num> is the logical RAID drive unit number (0 through 7) and <drive-status> is one of the following: "Optimal," "Degraded," "Failed," or "Unknown." |
| The hard disk at channel <chn>, target <tgt> had a <err-type> error.<br>OR<br>The hard disk at channel <chn>, target <tgt> had <count> <err-type> errors | A 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. Note that the number of errors is the increment since the last report of the specified error type **not the total accumulation of that type of error on the hard disk.** <chn> is the channel number and <tgt> is the target id. <count> is the number of errors that occurred; <err-type> is one of the following: "parity," "soft," "hard," or "miscellaneous." |
| The hard disk at channel <chn>, target <tgt> is <dev-sts>. | Monitor is reporting on the status of the hard disk at the specified channel and target. Disk status changes are reported only once, regardless of the /[NO]REPEAT switch. <chn> is the channel number and <tgt> is the target id. <dev-sts> is one of the following: "failed," a" hot spare," "write only," or "optimal." |

**Table 3-2 Monitor Utility Messages Reported (cont.)**

| Message | Meaning |
|---|---|
| A new hard disk has been inserted at channel <chn>, target <tgt>.<br>OR<br>The hard disk at channel <chn>, target <tgt> has been removed. | These two messages notify you of insertion and removal of drives. <chn> is the channel number and <tgt> is the target id. These messages are reported only once, regardless of the /[NO]REPEAT switch. |
| The hard disk at channel <chn>, target <tgt> is not being automatically replaced by a hot spare. Please replace the disk and rebuild it. | 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. Your action is required to replace the drive. <chn> is the channel number and <tgt> is the target id. |
| Logical RAID drive <drive-num> is rebuilding.<br>OR<br>Logical RAID drive <drive-num> is rebuilding (<percent>% complete). | When a logical RAID drive starts rebuilding, The monitor displays this message. If /NOREPEAT is specified, the first message displays. If /REPEAT is specified, then the second message displays, and includes the percent of that disk that has been rebuilt. <drive-num> is the logical RAID drive unit number (0 through 7). <percent> is the percentage complete. |

**Table 3-2 Monitor Utility Messages Reported (cont.)**

| Message | Meaning |
|---|---|
| The rebuild (on hard disk at channel <chn>, target <tgt>) has successfully completed.<br>OR<br>The rebuild (on hard disk at channel <chn>, target <tgt>) completed with some errors.<br>OR<br>The rebuild (on hard disk at channel <chn>, target <tgt>) failed because the new drive failed.<br>OR<br>The rebuild (on hard disk at channel <chn>, target <tgt>) failed because the source drive failed.<br>OR<br>The rebuild (on hard disk at channel <chn>, target <tgt>) failed because of bad blocks in the source drive(s). | 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 target id. |
| 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. |

**Table 3-2 Monitor Utility Messages Reported (cont.)**

| Message | Meaning |
|---|---|
| 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 has failed. |
| Monitor processing terminating. Fatal error encountered: <err_text>. | The monitor process encountered the specified fatal error. Your action depends on the error specified in <err_text>. |

## 3.3 Messages from the Parity Check Utility

Table 3–3 lists the messages that result from the SWXCR CHECK command.

**Table 3-3 Parity Check Utility 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 blocks 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 unit number specified in the command line device name does not correspond to a logical RAID drive. The following DCL command shows the valid SWXCR devices: "$ SHOW DEVICE DR " Select one of the DR devices shown in this output as input for the SWXCR CHECK command. |
| Error in check-consistency command occurred. | Controller returned unknown status from the Parity Check command. Status is unknown. |
| Parity check cannot be performed on RAID 0 (nonredundant) logical RAID drive.<br>OR<br>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. |
| 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.<br>The disk should be replaced and the logical RAID drive rebuilt. |

## 3.4 Messages from the Working with Drive Commands

The working -with-drive commands described in Chapter 2, Section 2.4, may display error messages that appear in the following format:

%SWXCR–E–BADCHAN, channel out of range

where the *BADCHAN, channel out of range*, portion of the message varies based on the error, and provides unique identification of the error, *BADCHAN*, and a description of the error, *channel out of range*.

Table 3–4 lists the error messages that result from the working with drive commands (spare, fail, rebuild mkopt):

**Table 3–4 Work with Drive Command Error Messages**

| Message | Meaning |
|---|---|
| %SWXCR-E-BADCHAN,<br>channel out of range | The command you entered included a channel which does not exist for the controller specified. |
| %SWXCR-E-BADTARG,<br>target out of range | The command you entered included a target which does not exist for the controller specified. |
| %SWXCR-E-BADDEV,<br>invalid device name | The command you entered included a device which does not exist for the controller specified. |
| %SWXCR-E-BADDEV,<br>device is not an SWXCR | The command you entered specified a device which cannot be effected by this command, because it is not a SWXCR device. |
| %SWXCR-E-NOSUCHDEV,<br>no device at channel %d target %d | The command you entered specified a valid target x and valid channel x but no device exists there. |
| SWXCR-E-DRVOPTIMAL,<br>can't rebuild an OPTIMAL drive | You entered the command to rebuild a drive with optimal status. |
| %SWXCR-E-RBLDFAIL,<br>drive failed during rebuild | There were too many bad blocks causing the drive to fail during the rebuild. |

**Table 3–4 Work with Drive Command Error Messages** (cont.)

| Message | Meaning |
|---|---|
| %SWXCR-E-RBLDINPROG, rebuild or parity check already in progress | You entered the rebuild or parity check command when either of these were already in progress for the specified channel/target. |
| %SWXCR-E-DNR, device not ready/not responding | The drive is not responding to a FAIL, MKOPT, or SPARE command. |
| Spinning up channel %d target %d | Indicates the specified device is spinning up. |

# 4

# *Using the RAID Array 200 Online Management Utility for OpenVMS 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 OpenVMS V6.2.*

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

The RAID Array 200 Online Management Utility for OpenVMS Graphical User Interface (GUI) provides the following functionality to maintain a RAID Array 200 Subsystem in an OpenVMS V6.2 environment only (the GUI portion of the utility is not available for systems running OpenVMS V6.1):

- View the status of the disk drives in the array
- Match on-screen drive representation to physical drives
- View drive vendor ID and code revision
- View the logical drives
- Check parity of logical drives
- Create hot spares
- Set the rebuild rate
- 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)

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 OpenVMS GUI

To start the Online Management GUI, type the following:

**`RUN DEVICE:[SYS0.SYSCOMMON.SYSEXE]SWXCRMGR.EXE`**
where DEVICE specifies the name of the drive on which you loaded this GUI

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 OpenVMS GUI, Displaying a 1–Channel Configuration**

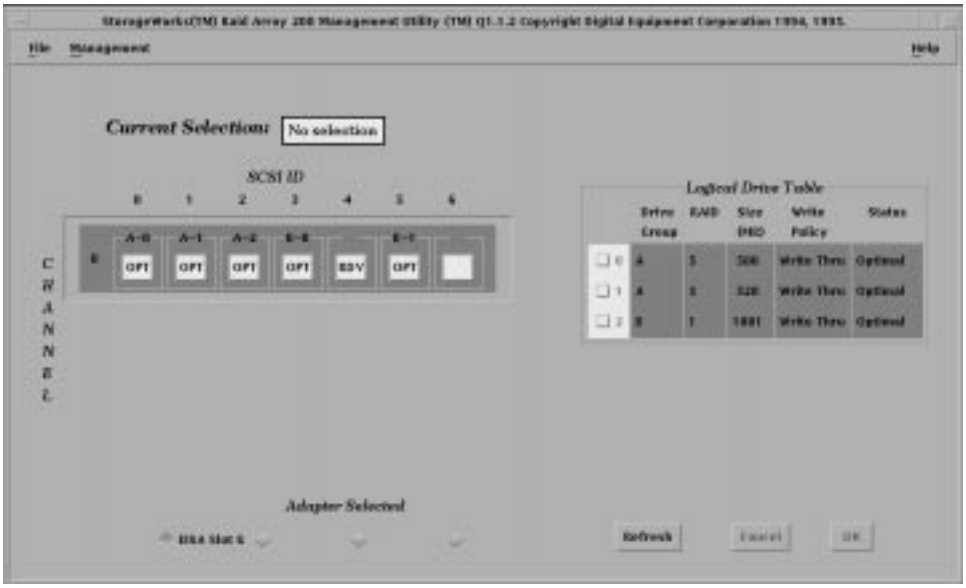

**Figure 4–2 The Base Window of the Online Management Utility for OpenVMS GUI, Displaying a 3–Channel Configuration**

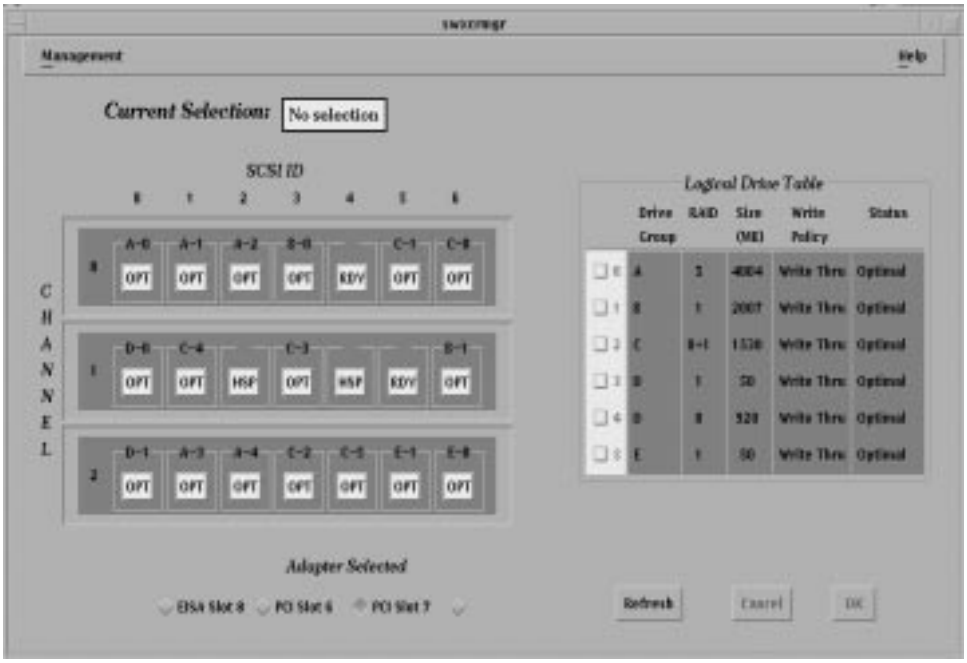


The base window displays:

- 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<br><br>For example, double click on a drive icon to display drive parameters |
| Access a menu | Left mouse button<br><br>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 OpenVMS(TM) V1.0.0 Copyright Digital Equipment Corporation 1994 |
| Resize corner | Allows you to change window size from any corner of window<br>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:

- *File*
- *Management*
- *Help*

The *File* menu provides the following items in the order shown:

*Print Configuration*

*Exit*

The *Management* menu 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*

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 OpenVMS 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 *File* 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**

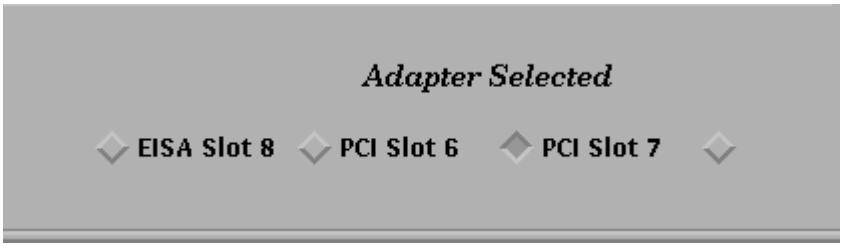

**Figure 4–4 The Subsystem Displayed after Selecting EISA Slot 8**

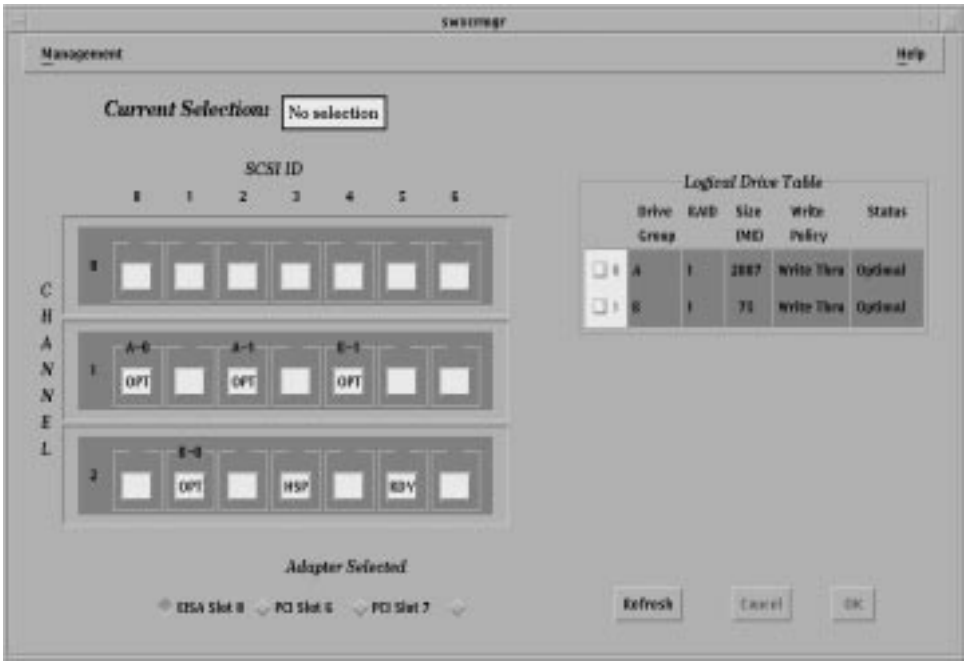

**Figure 4–5 The Subsystem Displayed after Selecting PCI Slot 6**

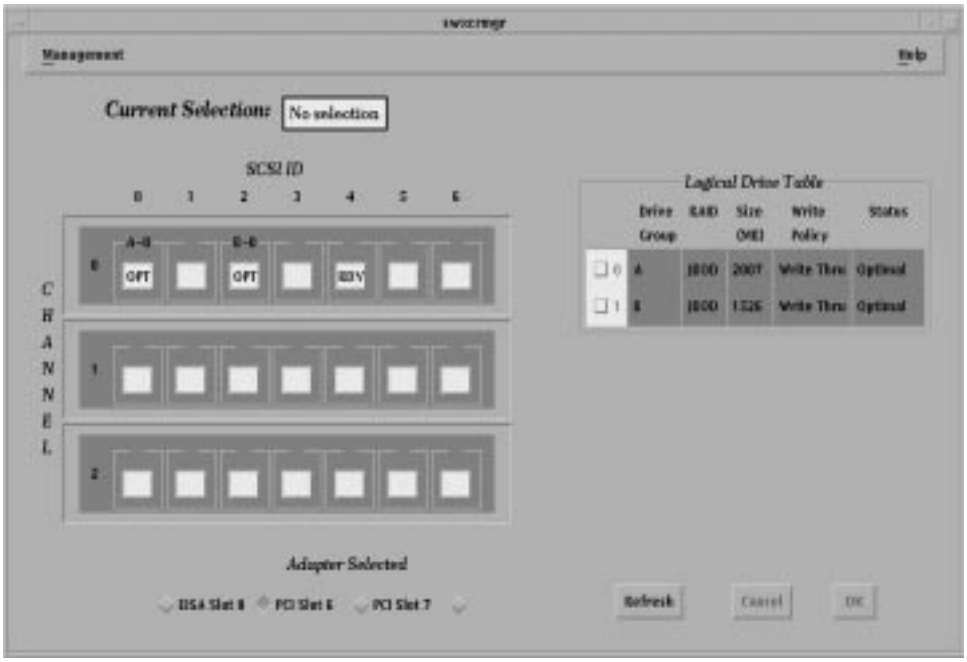

### 4.4.2 Viewing Disk Drives

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 controller and Figure 4–7 for a 3-channel controller.

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 Controller**

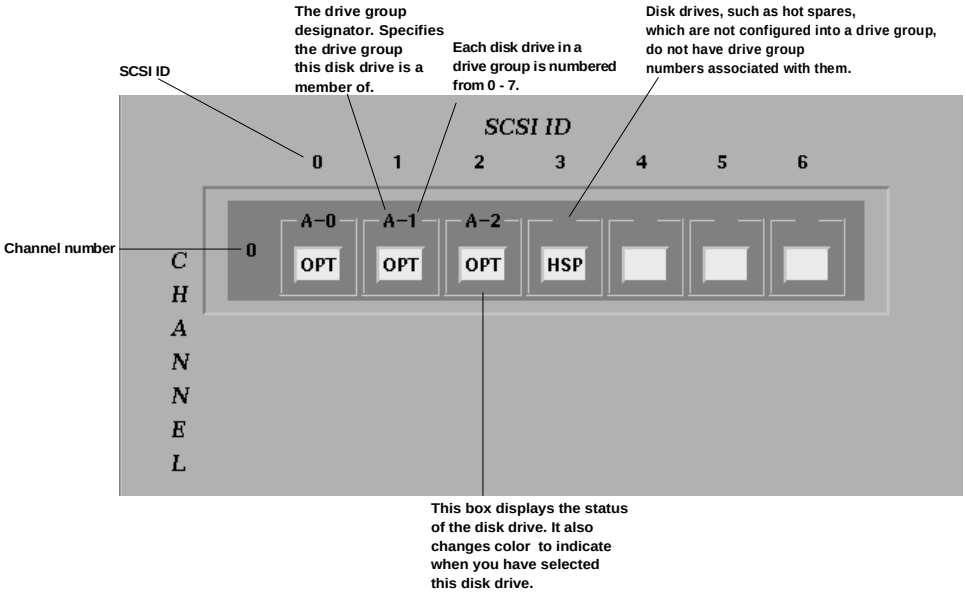

**Table 4–3 Drive Matrix Information**

| Element | Description |
|---|---|
| Channel | 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).<br><br>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.<br><br>You can configure drive groups on one channel of a controller, or across channels of a controller. |
| SCSI ID | 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. |
| Disk Drive Icons | 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. |

**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.<br>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.<br>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 Controller**

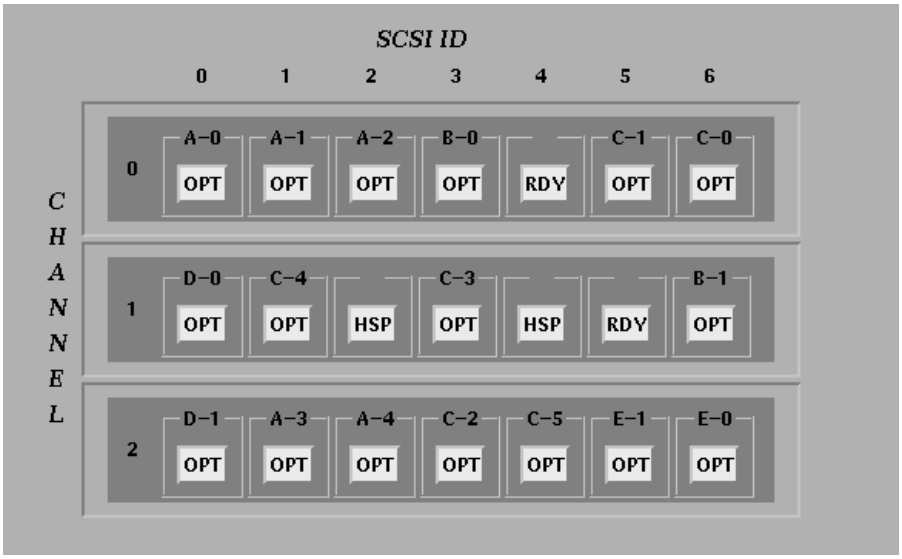

**Table 4–4 Disk Drive Status**

| 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<br>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<br>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<br>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<br>You need to replace a drive with this status<br>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<br>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 |

In addition to the drive information displayed by the drive matrix, you can view the following information by clicking on a drive icon:

- Click once on a drive icon to view the vendor ID and code revision of the physical drive that the icon represents
- Match a drive icon with the physical drive that it represents; click repeatedly on a drive icon displayed in the drive matrix, and the green LED of the physical drive it represents will flash each time you click.

### 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**

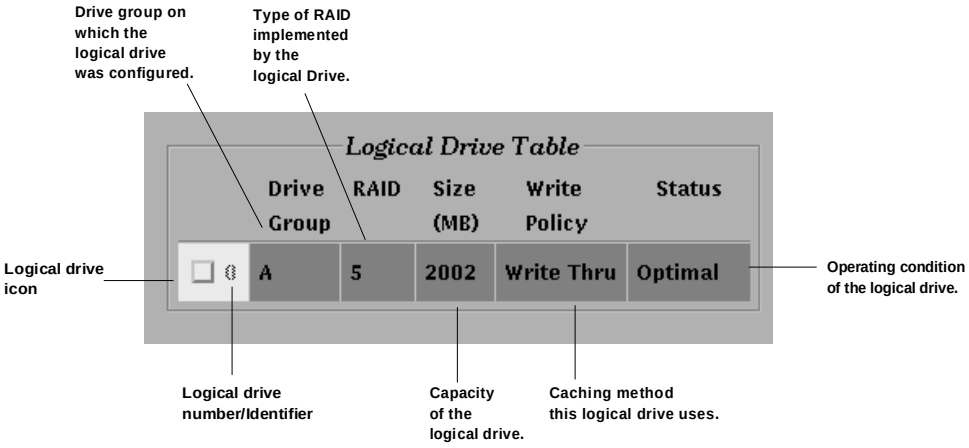

**Table 4–5 Logical Drive Table Information**

| Element | Description |
|---|---|
| Logical Drive Icon | The GUI represents each logical drive configured on your subsystem by a logical drive icon in the *Logical Drive Table*. |
| 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 *Logical Drive Table* 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.<br>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 *StorageWorks RAID Array 200 Subsystem Controller Installation and Standalone Configuration Utility User's Guide* 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.<br>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.<br>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.<br>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.<br>or<br>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.<br>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 *File* 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**

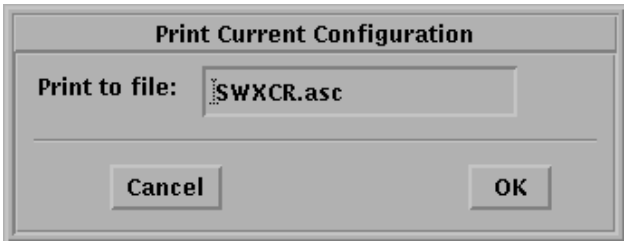


**Note**

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 maintaining a StorageWorks RAID Array 200 Subsystem with the online management utility for OpenVMS GUI under a system running OpenVMS V6.2.*

## 5.1 Maintaining the Subsystem

The RAID Array 200 Online Management Utility for OpenVMS GUI allows you to do the following to maintain a RAID Array 200 subsystem on a system running OpenVMS V6.2 (the GUI portion of the utility is not available for systems running OpenVMS V6.1):

- 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 OpenVMS 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 a hot spare drive

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 if 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**

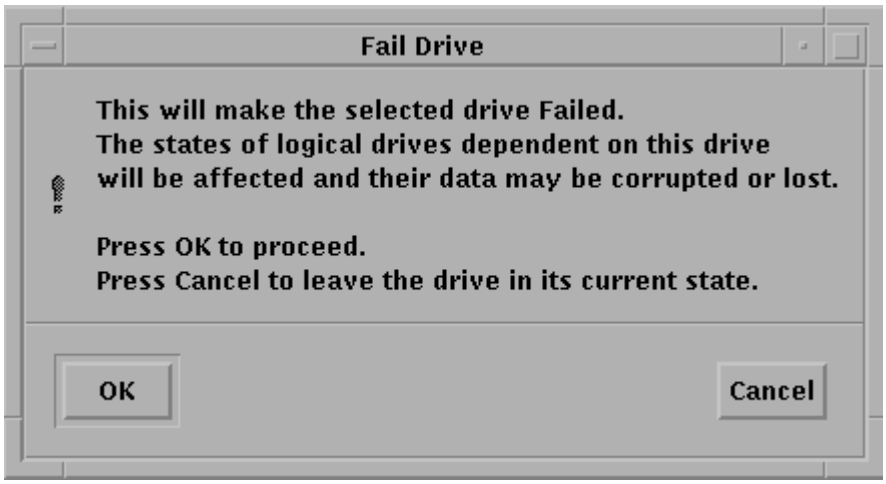


4. Locate the disk with the solid amber light in the storage enclosure and remove it.
5. Insert a replacement disk into the same slot of the storage enclosure from which you removed the disk that you failed.

   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**

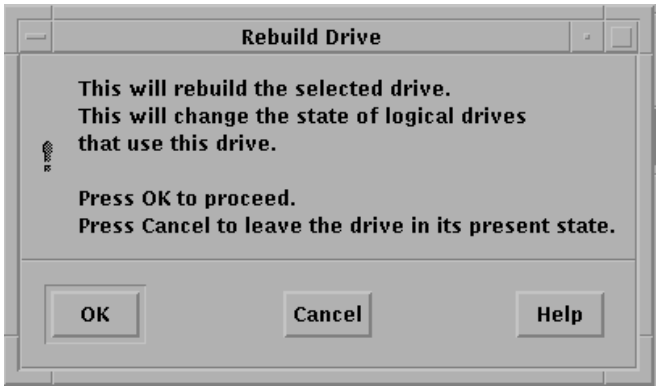


3. A window appears, similar to the one shown in Figure 5–3, with a slider bar showing the progress of the rebuild on the current logical drive. The controller rebuilds each logical RAID drive associated with this disk drive. A message appears, as shown in Figure 5–4 indicating when the controller finishes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

**Figure 5–3 The Rebuild Progress Window**

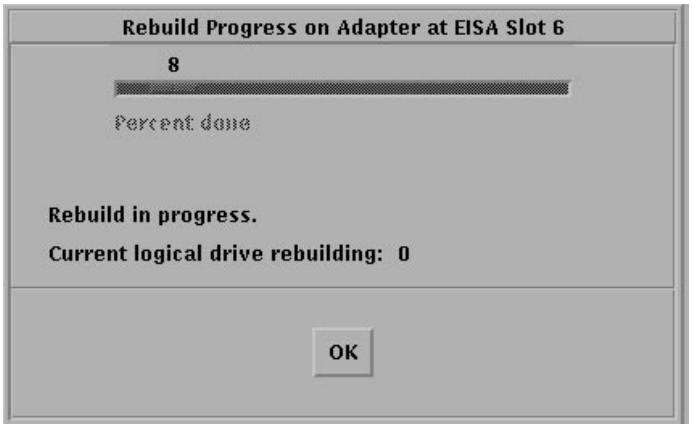


**Figure 5–4 The Rebuild Completed Message**

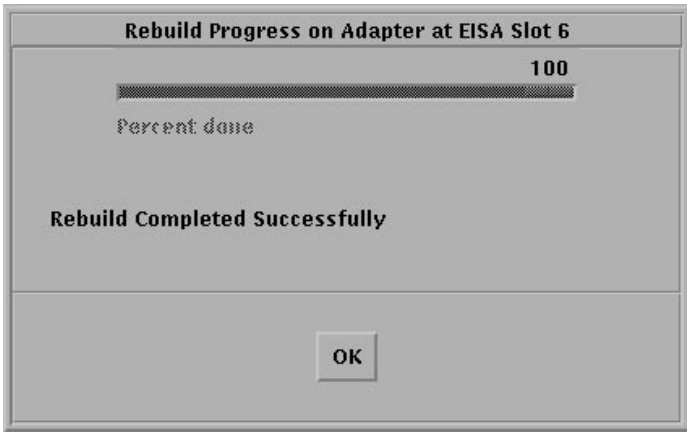


**NOTE**

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 Hot Spares

Make 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 hot spares out of any disk drive, which is not in a drive group. Use disk drives with capacities 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, 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**

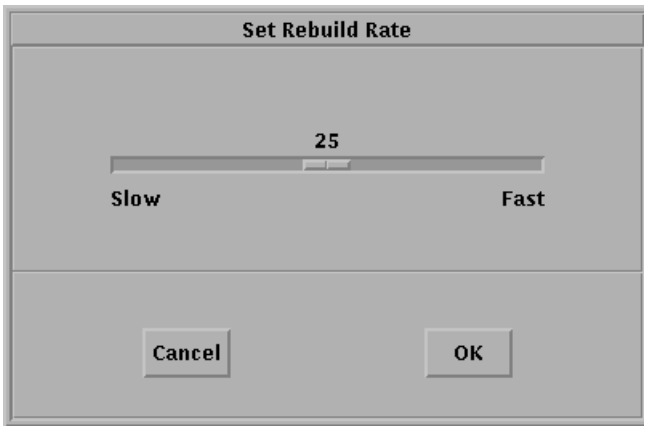


2. To exit without changing the rate, click the Cancel button in the pop-up window. To change the rate, drag the slider bar with the mouse to the value desired, and then click OK. The base window appears.

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.

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**

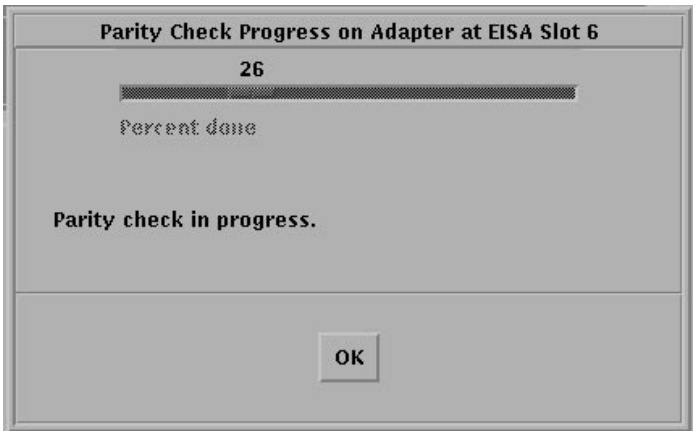


**Figure 5–9 Parity Check Completed Message**

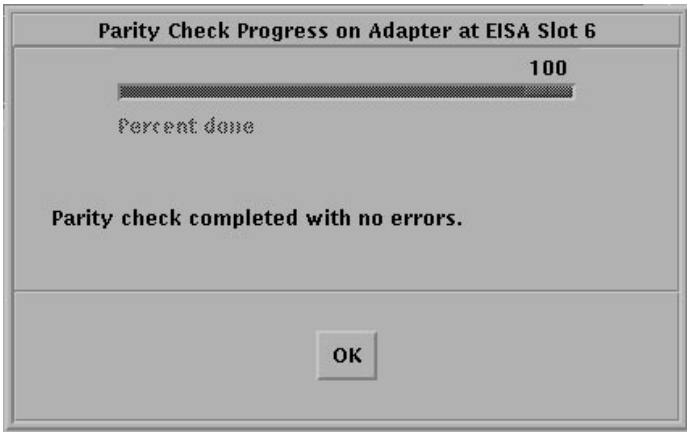

**Figure 5–10 Parity Status Report Appears When Parity Check Completes**

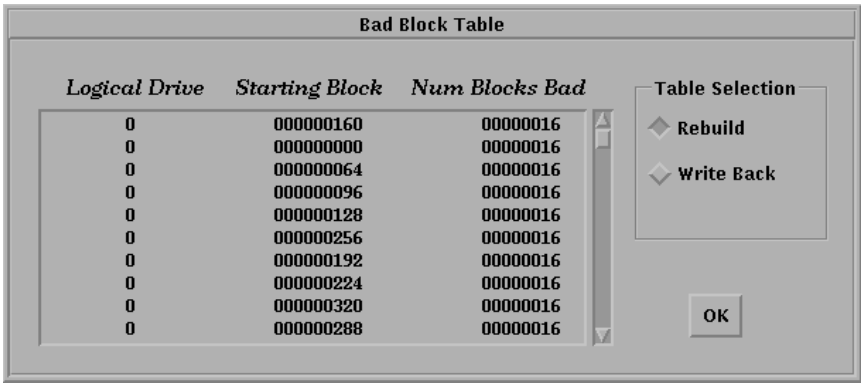


5. Click *OK* to close the bad block table after viewing it. *A Parity Check Progress* window appears, as shown in Figure 5–11.

**Figure 5–11 Parity Errors Detected Message**

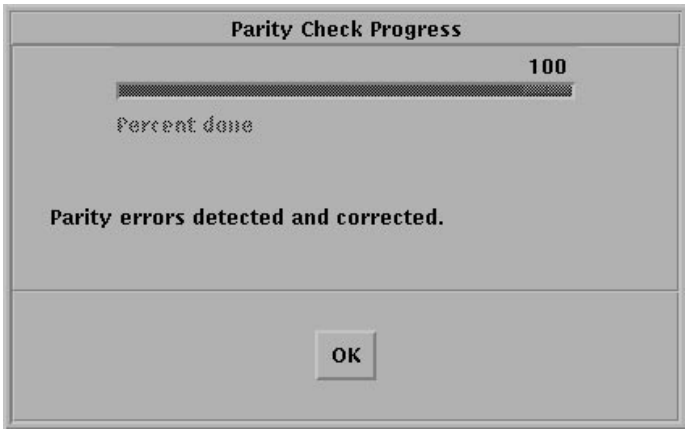


6. Click OK to close the *Parity Check Progress* window.

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