

IRIS Series 3000 Owner's Guide

Version 1.1

**Silicon Graphics, Inc.
2011 Stierlin Road
Mountain View, CA 94043**

Technical Publications:

Marcia Allen
Robin Florentine
Susan Luttner
Celia Szente
Diane M. Wilford

Illustrations:

Anna Szabados

Special thanks to
Herb Kuta, Donl Mathis, and Rick McLeod.

© Copyright 1986, Silicon Graphics, Inc.

All rights reserved.

This document contains proprietary information of Silicon Graphics, Inc., and is protected by Federal copyright law. The information may not be disclosed to third parties or copied or duplicated in any form, in whole or in part, without prior written consent of Silicon Graphics, Inc.

The information in this document is subject to change without notice.

IRIS Series 3000 Owner's Guide

Document numbers: 007-5220-010, 007-5220-011

UNIX is a trademark of AT&T Bell laboratories.
Ethernet and XNS are trademarks of Xerox Corporation.
Geometry Engine is a trademark of Silicon Graphics, Inc.
VAX is a trademark of Digital Equipment Corporation.

User Information

Please read the following information on product safety and interference potential.

Computer and CRT:

WARNING - For continued protection against risk of fire, replace fuses only with the same type and rating of fuse.

No operator-serviceable parts inside unit.

Class A Computing Device

WARNING: This equipment generates, uses, and can radiate radio frequency energy, and if not installed and used in accordance with the instruction manual, may cause interference to radio communications. It has been tested and found to comply with the limits for a Class A computing device pursuant to Subpart J of Part 15 of FCC Rules, which are designed to provide reasonable protection against such interference when operated in a commercial environment. Operation of this equipment in a residential area is likely to cause interference in which case the user at his own expense will be required to take whatever measures may be required to correct the interference (FCC, Part 15.818 a).

Instructions to User

This equipment generates and uses radio frequency energy and if not installed and used properly, ie., in strict accordance with the operating instructions, reference manuals, and the service manual, may cause interference to radio or television reception. It has been tested and found to comply with the limits for a Class A computing device pursuant to Subpart J of Part 15 of FCC Rules, which are designed to provide reasonable protection against such interference when operated in a commercial installation.

If this equipment does cause interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

1. Reorient the receiving antenna.
2. Relocate the equipment with respect to the receiver.
3. Move the equipment away from the receiver.
4. Plug the equipment into a different outlet so that equipment and receiver are on different branch circuits.

If necessary, consult a Silicon Graphics field engineer for additional suggestions.

Silicon Graphics is not responsible for any radio or TV interference caused by unauthorized modifications to this equipment. It is the responsibility of the user to correct such interference.

List of Figures

Figure	Page
IRIS Series 3000 System	2-3
IRIS 3030 Cabinet Front Panel (IRIS 3020 Similar)	2-4
IRIS 3010 Cabinet Front Panel	2-5
IRIS Cabinet Connections - Nineteen-inch 60 Hz Monitor	2-7
IRIS Cabinet Connections - Fifteen-inch 60 Hz Monitor	2-9
IRIS Cabinet Connections - 33 Hz Monitor	2-10
Monitor Control Panel for Nineteen-inch 60 Hz Monitor	2-12
Monitor Control Panel for Fifteen-inch 60 Hz Monitor	2-13
Monitor Control Panel for 33 Hz Monitor	2-14
IRIS Configuration Switches	3-2
Sample IRIS Boot Information	3-6
Sample <i>/etc/passwd</i> File	4-4
Sample <i>/etc/group</i> File	4-6
Standard I/O Panel and Auxiliary I/O Panel	7-2
Sample <i>/etc/ttytype</i> File	7-3
Sample <i>/etc/inittab</i> File	7-4
Sample <i>/etc/gettydefs</i> File	7-5
Monitor Back Panel for Nineteen-inch 60 Hz Monitor	A-6
Monitor Back Panel for Fifteen-inch 60 Hz Monitor	A-7
Monitor Back Panel for 33 Hz Interlaced Monitor	A-8
IRIS Cabinet Back Panel	A-11

List of Tables

Table	Page
IRIS Configuration Switches	3-9
PROM Monitor Commands	3-11
Boot Devices	3-11
Special Device Files in <i>/dev</i>	4-11, 4-12
Approximate Tape Capacities for Backups	4-21
Sample Tape Drive Commands	4-22
Files Critical to UNIX Operation	4-27
Disk Partition Sizes	6-2
Disk Storage Requirements	6-3
Two Ways to Configure a Second Disk	6-9
Serial Device Names	7-2
Cartridge Tape Drive Specifications	7-18
Pin Definitions for Modem Cable	7-20
Pin Definitions for a Null Modem Cable	7-20
Simplified Null Modem Pin Definitions for Printers	7-21
Three-wire Null Modem for Terminals	7-22
Summary of Cable Types	7-22
Primary and Secondary Video Options	8-1
IRIS Component Specifications	A-2
IRIS Environmental Specifications	A-12
Device Abbreviations	B-1

CONTENTS

1.	Road Map	1-1
1.1	Installation Overview	1-1
1.2	IRIS Series 2000 Turbo Products	1-2
1.3	Product Support	1-2
1.4	Conventions	1-2
2.	Hardware Installation	2-1
2.1	Hardware Components	2-1
2.2	Keyboard and Mouse Connection	2-6
2.3	Monitor to Cabinet Video Connections	2-6
2.3.1	Nineteen-inch 60 Hz Non-interlaced Monitor	2-6
2.3.2	Fifteen-inch 60 Hz Non-interlaced Monitor	2-8
2.3.3	33 Hz Interlaced Monitor	2-8
2.4	Monitor Adjustment	2-11
2.5	Monitor AC Power Cable Connection	2-15
2.6	Cabinet AC Power Connection	2-15
2.7	IRIS to Ethernet Connection	2-15
3.	Booting the IRIS	3-1
3.1	Workstation Booting Instructions	3-1
3.2	Terminal Booting Instructions	3-4
3.3	The Startup Environment	3-5
3.4	Boot Options	3-8
3.4.1	PROM Monitor Boot	3-9
3.4.2	Automatic Boot	3-12
4.	Workstation System Administration	4-1
4.1	Basic Operation and Shutdown	4-1
4.2	UNIX Configuration	4-2
4.2.1	Naming an IRIS Workstation	4-3
4.2.2	Adding a New Account	4-3
4.2.3	Entering Multi-user Mode Automatically	4-8
4.2.4	UNIX Configuration Files	4-8
4.2.5	Device Files	4-11
4.3	Network Communication	4-13
4.3.1	XNS Network Commands	4-13
4.3.2	TCP/IP Network Commands	4-14
4.3.3	Using Multiple Kernels	4-19
4.4	Backups	4-19
4.4.1	Making a Bootable Backup Tape	4-20
4.4.2	Making Periodic Backups	4-21

4.4.3	Storing Key UNIX Files on a Second Disk	4-25
4.5	Crash Recovery	4-26
4.5.1	Recovery Strategies	4-27
4.5.2	Recovery with <i>init</i>	4-28
4.5.3	Recovery by Booting from Tape	4-29
4.5.4	Recovery by Booting from the Second Disk	4-33
5.	IRIS 3010 System Administration	5-1
5.1	Host Software Installation	5-1
5.2	Operation	5-3
5.2.1	Connecting to the Host	5-3
5.2.2	Issuing Commands to the Terminal Emulator	5-4
5.2.3	Running a Demonstration Program	5-6
5.2.3	Shutting Down the IRIS Terminal	5-7
5.2.4	Boot Failure Recovery	5-7
5.3	Software Configuration	5-8
5.3.1	Configuration Tools	5-8
5.3.2	Configuring the Startup Environment	5-10
5.3.3	XNS Configuration	5-14
5.3.4	TCP/IP Configuration	5-14
5.3.5	Miscellaneous UNIX Configuration	5-15
6.	Disk Configuration	6-1
6.1	Disk Specifications	6-2
6.2	IRIS Workstation Disk Configuration	6-3
6.2.1	Changing the Swap Partition Size on an IRIS 3020	6-4
6.2.3	Changing the Swap Partition Size on an IRIS 3030	6-6
6.2.4	Configuring a Second Disk on an IRIS Workstation	6-9
6.3	IRIS Terminal Disk Configuration	6-13
6.3.1	Changing the Swap Partition Size on an IRIS Terminal	6-13
6.3.2	Creating and Using a Backup Floppy Set	6-17
7.	Optional Peripherals	7-1
7.1	ASCII Terminal	7-1
7.1.1	Connecting an ASCII Terminal to the IRIS	7-1
7.1.2	Configuring Software for an ASCII Terminal Connection	7-3
7.2	Dial and Button Box	7-5

7.2.1	Connecting the Dial and Button Box to the IRIS	7-5
7.2.2	Configuring Software for the Dial and Button Box Connection	7-6
7.3	Modem (IRIS Workstation Only)	7-7
7.3.1	Connecting a Modem to the Workstation	7-7
7.3.2	Configuring Software for a Modem Connection	7-7
7.4	Serial Printer	7-10
7.4.1	Connecting a Printer to the IRIS	7-10
7.4.2	Configuring the Software for a Printer Connection	7-10
7.5	Color Graphics Printer	7-13
7.5.1	Connecting the Color Graphics Printer to the IRIS	7-13
7.5.2	Using the Color Graphics Printer	7-13
7.6	Digitizer Tablet	7-14
7.6.1	Connecting the Digitizer Tablet to the IRIS	7-15
7.6.2	Configuring Software for a Digitizer Tablet	7-16
7.6.3	Using a Digitizer Tablet	7-17
7.7	Tape Drives (IRIS Workstation Only)	7-17
7.7.1	Half-inch Tape Drive	7-17
7.7.2	Quarter-inch Cartridge Tape Drive	7-18
7.8	Floppy Disk Drive	7-18
7.9	Stereo Optic Viewer	7-19
7.9.1	Connecting the Stereo Optic Viewer to the IRIS	7-19
7.10	Serial Port Cabling	7-19
8.	Video Options	8-1
8.1	Supported Options	8-2
8.1.1	60 Hz Non-interlaced Monitor	8-2
8.1.2	33 Hz Interlaced Monitor	8-2
8.1.3	RS-170A	8-2
8.1.4	European Video Standard	8-2
8.1.5	Genlockable RS-170A or European Video Standard	8-3
8.2	Booting with the Secondary Video Driver	8-4
8.3	Using the Graphics Library with the Video Options	8-4
9.	Demonstration Programs and Gifts	9-1
9.1	Demonstration Programs	9-1
9.1.1	Running Demonstration Programs	9-1

9.1.2	Using the Mouse with the Window Manager	9-2
9.1.3	Window Manager Demonstration Programs	9-2
9.2	Gifts	9-9
Appendix A:	IRIS Specifications	A-1
A.1	Cables	A-2
A.2	Monitor	A-2
A.2.1	60 Hz Non-interlaced Monitors	A-2
A.2.2	33 Hz Interlaced Monitor	A-2
A.2.3	Monitor Control Panel	A-2
A.2.4	Monitor Back Panel	A-2
A.3	IRIS Electronics Cabinet	A-9
A.3.1	Power Switch	A-9
A.3.2	Optional Tape Drives	A-9
A.3.3	Standard I/O Panel	A-9
A.3.4	Cabinet Power Panel	A-10
A.4	Site Selection	A-12
Appendix B:	System Messages	B-1
B.1	Hardware Error Messages	B-2
B.1.1	General Hardware Error Messages	B-2
B.1.2	Disk Controller Messages	B-2
B.1.3	Tape Drive Messages	B-3
B.1.4	Floppy Controller Messages	B-4
B.2	System Software Messages	B-4

1. Road Map

This document explains how to install and test IRIS series 3000 products, and how to perform system administration. Read this document carefully before installing your IRIS.

1.1 Installation Overview

The table below summarizes the procedures for installing IRIS series 3000 products.

To complete this procedure:	See Chapter/Section:
Install hardware	Chapter 2
Install optional peripherals, including ASCII terminals	Chapter 7
Install non-standard video monitors	Chapter 8
Boot the IRIS workstation	Section 3.1
Boot the IRIS terminal	Section 3.2
Configure software (IRIS workstation)	Chapter 4
Configure software (IRIS terminal)	Chapter 5
Configure disk	Chapter 6

For basic information about UNIX commands, the UNIX text editor, and navigating the file system, see the pamphlet entitled *Getting Started with Your IRIS Workstation*.

1.2 IRIS Series 2000 Turbo Products

This manual is applicable to the IRIS 2300T (3010 equivalent), 2400T (3020 equivalent), and 2500T. Early IRIS series 2000 models with the Turbo upgrade may not have the same standard peripherals described here. The 60 Hz monitor may be different from the one described in this guide. The mouse may be mechanical rather than optical. The early IRIS 2400 Turbo system required a separate junction box. For a description of these features, see the *IRIS Workstation Guide Series 2000, Version 1.0*.

1.3 Product Support

Silicon Graphics provides a comprehensive product support and maintenance program for IRIS series 3000 products. For further information, contact Customer Service through the Geometry Hotline

Silicon Graphics Geometry Hotline	
(800) 252-0222	U.S. except California (toll-free)
(800) 345-0222	California (toll-free)
(800) 443-0222	Canada

1.4 Conventions

This document uses the standard UNIX convention for referring to entries in the UNIX documentation. The entry name is followed by a section number in parentheses. For example, *cc(1)* refers to the *cc* manual entry in Section 1 of the *UNIX Programmer's Manual, Volume I*.

In command syntax descriptions and examples, square brackets surrounding an argument indicate that the argument is optional. Words in *italics* represent variable parameters, which should be replaced with the string or value appropriate for the application.

In text descriptions, filenames and UNIX commands are written in *italics*. IRIS Graphics Library and PROM monitor commands are written in **typewriter font**.

2. Hardware Installation

This chapter describes how to install and connect the standard components of the IRIS 3000 series systems (see Figure 2-1).

WARNING
Do <i>not</i> connect the IRIS to an external power source until each cable has been connected and checked.

2.1 Hardware Components

The typical configuration for IRIS series 3000 systems includes these components:

- Electronics cabinet with one disk drive
- Nineteen-inch tilt-and-swivel 60 Hz non-interlaced monitor (standard) or fifteen-inch tilt-and-swivel 60 Hz non-interlaced monitor
- Keyboard with right-hand and left-hand connectors for the mouse
- Optical mouse with grid pad
- Ethernet transceiver

For a more detailed description of the components, see Appendix A.

The IRIS 3000 series products can also be ordered with these hardware options:

- Dial and button box
- 33 Hz interlaced monitor
- Digitizer tablet
- Tektronix 4692 color printer
- Half-inch tape drive
- Quarter-inch cartridge tape drive
- Stereo optic viewer cable
- Floppy disk drive (standard on the IRIS 3010)
- Additional hard disk drive

This manual is applicable to the IRIS 2300T, 2400T, and 2500T. Early IRIS series 2000 models with the Turbo upgrade may not have the same standard peripherals described here. The 60 Hz monitor may be different from the one described in this guide. The mouse may be mechanical rather than optical. The early IRIS 2400 Turbo system required a separate junction box. For a description of these features, see the *IRIS Workstation Guide Series 2000, Version 1.0*.

To install optional peripherals, see Chapter 7.

The major components have the following functions and features:

- The *electronics cabinet* for the standard configuration of the IRIS series 3000 is a floor-standing unit with a 20-slot backplane, a disk drive, and a power supply (see Figure 2-1). The cabinet uses forced-air cooling and is mounted on casters. A quarter-inch cartridge tape drive (internal) or a half-inch streaming tape drive (external) is optional on the IRIS workstation (see Figure 2-2). A floppy disk drive is standard on the IRIS terminal (see Figure 2-3). See Chapter 6 for disk specifications and Chapter 7 for tape specifications.
- Three types of *monitor* are available from Silicon Graphics for the IRIS series 3000. The fifteen-inch and nineteen-inch tilt-and-swivel monitors are 60 Hz high-resolution RGB color monitors. Silicon Graphics also offers a 33.43 Hz nineteen-inch interlaced monitor. (This monitor will be referred to as 33 Hz rather than 33.43 Hz.)

The IRIS series 3000 can also be ordered with drivers for other types of monitors (see Chapter 8, Video Options).

- The *keyboard* is an 83-key, up-down encoded keyboard, with right-hand and left-hand connectors for the optical mouse.
- The *mouse* is a three-button optical mouse with a grid pad.




Figure 2-1: IRIS Series 3000 System




Figure 2-2: IRIS 3030 Cabinet Front Panel (IRIS 3020 Similar)




Figure 2-3: IRIS 3010 Cabinet Front Panel

2.2 Keyboard and Mouse Connection

Connect the mouse to the keyboard by inserting the mouse cable in either the right-hand or left-hand nine-pin D connector in the keyboard. Secure the mouse connector to the keyboard with the two retaining screws on the connector. Connect the keyboard to the IRIS by inserting the keyboard cable into **Port 1** of the standard I/O panel on the IRIS.

2.3 Monitor to Cabinet Video Connections

The procedure for connecting the monitor to the IRIS is different for each type of monitor.

2.3.1 Nineteen-inch 60 Hz Non-interlaced Monitor

Use the color-coded coaxial video cables to connect the monitor to the cabinet (see Figure 2-4).

1. If you are using only one monitor, plug the terminators into the *unused* video connectors. The terminators are attached to the monitor by chains.

To connect several monitors in a series (daisy chain), remove the terminators from all but the last monitor. Plug the terminators on the last monitor into the unused video input sockets.

Attach the outgoing video cables to the unused BNC connectors on the back of the monitor.

2. Connect each color-coded cable end to the corresponding input socket on the monitor back panel.

Push each cable into its connector and rotate to lock into place.

3. Connect the other end of each color-coded cable to the corresponding output socket on the standard I/O panel.

Push each cable into the connector and rotate to lock into place.




Figure 2-4: IRIS Cabinet Connections
Nineteen-inch 60 Hz Non-interlaced Monitor

2.3.2 Fifteen-inch 60 Hz Non-interlaced Monitor

Use the color-coded coaxial video cables to connect the monitor to the cabinet (see Figure 2-5).

1. If you are using only one monitor, set *all* of the input impedance switches on the back of the monitor to the **75 Ω** position. The input impedance switches are push-button switches located between the input sockets. **IN** is the **75 Ω** position. **OUT** is the **High** position.

To connect several monitors in a series (daisy chain), set the input impedance switches to the **High** position for all but the last monitor, which should be set to the **75 Ω** position.

Attach the outgoing video cables to the unused BNC connectors on the back of the monitor.

2. Connect each color-coded cable end to the corresponding input socket on the monitor back panel.

Push each cable into its connector and rotate to lock into place.

3. Connect the other end of each color-coded cable to the corresponding output socket on the standard I/O panel.

Push each cable into the connector and rotate to lock into place.

2.3.3 33 Hz Interlaced Monitor

Use the color-coded coaxial video cables to connect the monitor to the cabinet (see Figure 2-6).

1. If you are using only one monitor, set *all* of the input impedance switches on the back of the monitor to the **75 Ω** position.

To connect several monitors in a series (daisy chain), set the input impedance switches to the **High** position for all but the last monitor, which should be set to the **75 Ω** position.

Attach the outgoing video cables to the unused BNC connectors on the back of the monitor.

2. Connect each color-coded cable end to either one of the corresponding input sockets on the monitor back panel.

Push each cable into its connector and rotate to lock into place.

3. Connect the other end of each color-coded cable to the corresponding output socket on the standard I/O panel.

Push the cable into the connector and rotate to lock into place.




Figure 2-5: IRIS Cabinet Connections
Fifteen-inch 60 Hz Non-interlaced Monitor




Figure 2-6: IRIS Cabinet Connections
33 Hz Interlaced Monitor

2.4 Monitor Adjustment

The monitor control panel on the right side of the monitor allows the characteristics of the monitor to be adjusted.

The nineteen-inch monitors (60 Hz and 33 Hz) have two knobs labeled **Brightness** and **Contrast**, and a button labeled **Degauss** (see Figures 2-7 and 2-9).

The fifteen-inch monitor has two recessed dials labeled **Brightness** and **Contrast**, and a button labeled **Degauss** (see Figure 2-8).

NOTE: Color rendering and stability may drift for the first 45 minutes after powerup.

1. After the monitor has warmed up, turn the **Brightness** and **Contrast** controls to the maximum (clockwise) setting.
2. Turn the **Brightness** control down until the gray raster just disappears in relation to the black area at the edge of the screen. Lighter brightness settings impair image sharpness and color fidelity.
3. After the **Brightness** control has been set, adjustment of the **Contrast** control is a matter of personal taste.
4. If the color purity or convergence appear to be out of adjustment, hold down the **Degauss** button on the monitor control panel for about five seconds and then release it. The image should improve noticeably.

Once the nineteen-inch 60 Hz monitor has been adjusted, the **Brightness** and **Contrast** knobs can be pushed in: so that they are flush with the surface of the monitor.





Figure 2-7: Monitor Control Panel for Nineteen-inch 60 Hz Non-interlaced Monitor




Figure 2-8: Monitor Control Panel for Fifteen-inch 60 Hz Non-interlaced Monitor




Figure 2-9: Monitor Control Panel for 33 Hz Interlaced Monitor

2.5 Monitor AC Power Cable Connection

Plug the monitor into a 120/240 VAC outlet.

1. Connect the female end of the AC power cable to the **Input** power socket on the monitor back panel.
2. Connect the male end of the monitor power cable to a wall outlet.

2.6 Cabinet AC Power Connection

The cabinet power socket is located on the cabinet power panel (see Figure A-4).

CAUTION
Do not connect the IRIS to a switched power outlet.

1. Connect the female end of the AC power cable to the power socket labeled **Power** on the cabinet power panel.
2. Connect the male end of the cabinet power cable to an appropriate outlet. See Appendix A, IRIS Specifications, for the IRIS power requirements.

2.7 IRIS to Ethernet Connection

The IRIS can communicate with other hosts and terminals (that have the necessary communication software) through an Ethernet local area network. The IRIS can be connected to an Ethernet local area network while the network is operating.

To connect the IRIS to an Ethernet local area network, you need:

- An *Ethernet transceiver* to attach to the Ethernet.
- One 75-foot, 15-conductor drop cable to connect the IRIS cabinet to the Ethernet transceiver.

The Ethernet port is located on the standard I/O panel, on the back of the cabinet.

To connect the Ethernet cable to the cabinet, follow these steps:

1. Connect the male end of the drop cable to the Ethernet port, labeled **Ethernet** on the standard I/O panel. Use the slide lock to secure the cable in the connector.
2. Select an appropriate tap point on the Ethernet coaxial cable.

NOTE: Approved Ethernet coaxial cable is marked with rings at 2.5-meter intervals (minimum distance). Transceivers should be placed at these rings to minimize transceiver reflection, which can induce transmission errors.

Tap into the Ethernet cable (instructions are included with each transceiver).

3. Connect the transceiver to the Ethernet cable.
4. Connect the female end of the drop cable to the transceiver.

3. Booting the IRIS

This chapter contains step-by-step procedures for booting and checking out a new IRIS system. This chapter also describes the bootstrap options and the startup environment for the IRIS. Before following the procedures in this chapter, you must have installed your system as described in Chapter 2. After completing the procedures in this chapter, see Chapter 4 for workstation system administration procedures, or Chapter 5 for terminal system administration procedures.

This chapter contains the following sections:

<i>Section</i>	<i>Topic</i>
3.1	Workstation Booting Instructions
3.2	Terminal Booting Instructions
3.3	The Startup Environment
3.4	Boot Options

3.1 Workstation Booting Instructions

To boot the default software shipped on the disk and run the flight simulator demonstration program, follow these steps:

1. Set the **Power** switch on the display monitor to **On**.
2. Set configuration switches 1 through 9 on the cabinet back panel to **Closed** (see Figure 3-1).
3. Set the **Power** switch on the cabinet front panel to **On**. If the power for the workstation is already on, press the **Reset** button. The workstation displays the PROM monitor prompt:

```
iris>
```

Wait 20 seconds for the disk to spin up to speed. (The boot will fail if you don't wait that long.)




Figure 3-1: IRIS Configuration Switches

4. Enter the letter `b` and press `Return`.

```
b
```

The workstation displays system information similar to that shown in Figure 3-2.

5. At this point, the IRIS is running in UNIX *single-user mode*, which is used only for system maintenance. For normal operation of the IRIS workstation, start *multi-user mode* by entering the `multi` command.
6. When the workstation prompts for the correct date, enter the date in the requested format, as illustrated in the example below. The characters `mmddhhmm[yy].[ss]` represent digits for month-day-hour-minute-year-seconds. The year and the seconds are optional.

```
Is the date Wed July 21 08:10:33 PST 1989 correct? (y or n) n
Enter the correct date (mmddhhmm[yy][.ss]):0721083089.33
```

7. The workstation prompts for a file system check:

```
Do you want to check filesystem consistency? (y or n)
```

If you answer `yes`, the file systems are checked with the `fsck` program. Only answer `no` if the system has been shut down with the `reboot` command, as described in Section 4.1. Even then, you are advised to run `fsck` occasionally. (See the *UNIX Programmer's Manual* for more about `fsck`.) The system now performs initializations required for multi-user mode. Once multi-user mode has been started, a login prompt appears:

```
IRIS login:
```

8. To test a newly installed system, log in to the guest account:

```
IRIS login: guest
```

9. To use the Flight simulator, type:

```
~/demos/flight
```

10. To leave the *flight* program, press the `ESC` key.

11. To log out, type:

```
logout
```

Important

<p>Now that your IRIS workstation is up and running, make sure you have a bootable backup tape containing the stand-alone software, the root file system, and the user file system. You should receive a bootable backup tape with your workstation. This tape can be used to boot the workstation and rebuild the disk in case the file system is damaged.</p>

For more information about system administration, such as adding new accounts, changing the time zone, and enabling network communication, see Chapter 4.

Section 3.4 contains more detailed information about the boot procedure, including procedures for booting from various devices. Section 3-4 also describes the PROM monitor, a command interpreter that controls the boot environment. Table 3-2 shows the commands understood by the PROM monitor. Table 3-3 shows the boot devices.

3.2 Terminal Booting Instructions

To boot the default software shipped on the disk and run the flight simulator demonstration program, follow these steps:

1. Set the **Power** switch on the display monitor to **On**.
2. Set configuration switches 1 through 9 on the cabinet back panel to **Closed** (see Figure 3-1).
3. Set the **Power** switch on the cabinet front panel to **On**. If the power is already on, press the **Reset** button. The IRIS terminal displays the PROM monitor prompt:

```
iris>
```

Wait 20 seconds for the disk to spin up to speed. (The boot will fail if you don't wait that long.)

4. Enter the letter **b** and press **Return**.

```
b
```


- When the terminal prompts for the correct date, enter the date in the requested format, as illustrated in the example below. The characters *mmddhhmm[yy].ss* represent digits for month-day-hour-minute-year-seconds. The year and the seconds are optional.

```
Is the date Wed July 21 08:10:33 PST 1989 correct? (y or n) n
Enter the correct date (mmddhhmm[yy][.ss]):0721083089.33
```

- After the startup sequence is complete, the IRIS 3010 automatically runs the terminal emulator program, *wsiris*. Enter the sequence `~!` to the first prompt to enter the shell:

XNS or IEEE-488 option:

```
Connect to what host? ~!
```

TCP/IP option:

```
Enter IRIS IP address : ~!
```

If your IRIS has the IBM option, enter `~!` to return from the *t3279* menu to the shell.

The shell prompt appears:

```
%
```

- To use the flight simulator, type:

```
~demos/flight
```

- To leave the *flight* program, press the **ESC** key.

To return to the terminal emulator program, enter the `exit` command.

See Chapter 5 for more information on using the terminal emulator, connecting to a host, and configuring the user environment.

3.3 The Startup Environment

During the startup process, the IRIS displays a set of system information, including memory size, hardware configuration, and file system identification (see Figure 3-2).

Kernel Number:

SYSTEM 5 UNIX #6: [Fri May 4 11:15:09 PST 1985]

Copyright:

(c) Copyright 1983 - UniSoft Corporation
(c) Copyright 1983 - Silicon Graphics Inc.

Actual Memory:

real = 4194304

Kernel Size:

kmem = 491520

Available User Memory:

user = 3702784
bufs = 819200 (max=8k)

Hardware Configuration:

dsd0 at mbio 0x7f00 ipl1
qic0 slave 0
md0 (Priam V170 Name: IP-2) slave 1
md1 (Priam V170 Name: IP-2) slave 2
mf0 not installed
st0 not installed
nx0 (FW2.6 HW 2.0) (0800.1400.0120) at mbio 0x7ffc ipl2
fpa installed

Root File System Device Name:

root on md0a

Swap Space Device Name and Size:

swap on md0b [8865K]

Figure 3-2: Sample IRIS Boot Information

The IRIS terminal boots into the terminal emulator program. For more information on IRIS terminal operation, see Chapter 5.

The IRIS workstation boots into single-user mode, which should be used only for system maintenance. For normal operation, the workstation must be put into multi-user mode. The *multi* command initiates the process of putting the workstation into multi-user mode. The details of the procedure are controlled by the file */etc/rc*, which contains commands for starting daemons and mounting file systems. For more information, see *brc(1)*.

1. To enter multi-user mode, enter the *multi* command:

```
multi
```

2. If the */etc/model* file becomes corrupted, the system prompts for the model number, writing the response to */etc/model*.
3. The workstation asks to run the UNIX file system check program *fsck*.

```
Do you want to check filesystem consistency? (y or n)
```

If you answer *yes*, the file systems are checked with the *fsck* program. Only answer *no* if the system has been shut down with the *reboot* command, as described in Section 4.1. Even then, you are advised to run *fsck* occasionally.

If *fsck* is run and finds no problems, it displays one line for each test. If *fsck* finds any problems, it prompts for permission to repair the file system. See the *UNIX Programmer's Manual* for more information on *fsck*.

4. Enter the correct date using the specified format:

```
Enter the correct date (mmdhmm[.ss][yy]):07210830.3389
Is the date Wed July 21 08:30:33 PST 1989? (y or n) y
```

In the example above, the characters *mmdhmm.ssy* represent month-day-hour-minute-seconds-year. The seconds and the year are optional. The entry above sets the date to 33 seconds past 8:30 a.m. on July 21, 1989.

The system now performs initializations required for multi-user mode. Once multi-user mode has been started, a login prompt appears:

```
IRIS login:
```

5. Section 4.2 explains how to configure the workstation to boot into multi-user mode automatically.

3.4 Boot Options

The IRIS boot environment is determined by the setting of configuration switches 1 through 4 on the cabinet back panel. The two basic boot options are automatic boot or PROM monitor boot:

- | | |
|-------------------|---|
| Automatic boot | At powerup, the IRIS tries to boot from a file called <i>defaultboot</i> on the device specified by the configuration switches (see Table 3-1 and Section 3.4.2). |
| PROM monitor boot | At powerup, the IRIS enters the PROM monitor and waits for further boot instructions (see Section 3.4.1). |

To boot from a secondary video driver, see Chapter 8.

Since the IRIS can be booted from different devices (hard disks, tape drives, etc.), the PROM monitor provides the *ls* command, a version of the UNIX *ls(1)* command, for displaying the names of the files on the attached devices. You can specify a device name and a pathname in the *ls.* command (see Table 3-2).

For example, `ls /` searches the root file system (/) on the default device and lists its contents. If you type `ls` without an argument, the IRIS displays a list of files on the default device.

The IRIS is shipped with the configuration switches set for booting from the PROM monitor. To select automatic booting, change the settings of the configuration switches on the standard I/O panel. Table 3-1 lists the defined configuration switch settings.

If you attempt to boot a nonexistent device or file from the PROM monitor, the system prints an error message and returns you to the PROM monitor. If an automatic boot fails, the system waits a few seconds and tries again. Automatically booting an IRIS 3030 always fails on the first attempt after power-up, but the automatic retry should succeed.

Configuration Switches			
Switch	Switch Name	Position ¹	Meaning
1 - 4	Boot environment	CCCC OCCC COCC OCCC OCOC	Hard disk boot Cartridge tape boot Floppy disk boot Network boot PROM monitor
5	Autoboot	C O	PROM Monitor boot Automatic boot
6	Quiet mode	C O	Display system information Don't display system information
7	Monitor select	C O	Display on primary monitor Display on secondary monitor
8-9	Reserved	C	

Table 3-1: IRIS Configuration Switches

As Table 3-1 shows, switches 1 through 4 select the device from which the IRIS is to be booted. Switch 5 specifies whether the IRIS should perform an automatic boot or a PROM monitor boot. Switch 6 determines whether or not system information is displayed on the screen after the IRIS is reset. Switch 7 selects the display monitor type. (For an explanation of selecting a display type, see Chapter 8, Video Options.)

3.4.1 PROM Monitor Boot

The PROM monitor is a command interpreter designed primarily to control the IRIS boot environment. The IRIS enters the PROM monitor when the power is turned on or the system is reset in either of two cases:

- Configuration switches 1 through 4 are set for PROM monitor boot (“OCOC”).
- Switch 5 is set to “C”, regardless of the setting of switches 1 through 4.

¹ C means **C**losed and O means **O**pen.

Table 3-2 summarizes the PROM monitor commands. Table 3-3 lists the device names recognized by the PROM monitor.

To boot through the PROM monitor, follow these steps:

1. Set the **Power** switch on the display monitor to **On**.
2. Set configuration switches 1 through 4 to the appropriate boot device setting (see Table 3-1). The default boot device is the hard disk, which has the setting "CCCC".
3. Set switch 5 to **Closed**.
4. Set the **Power** switch on the cabinet front panel to **On**.

If the clock battery has run down, the system displays this ominous message:

```
Power has been lost to the machine.
```

This message refers only to the clock, however.

The IRIS displays the PROM monitor prompt:

```
iris>
```

If the prompt does not appear, press the **Reset** button.

Wait 20 seconds for the disk to spin up to speed. (The boot will fail if you don't wait that long.)

5. To boot from the device specified by configuration switches 1 through 4, enter the letter **b** and press **RETURN**.

The IRIS reads the file *defaultboot* from the default device.

Command	Description
b	Boot <i>defaultboot</i> from the default boot device.
[b] <i>file</i>	Boot <i>file</i> on the default boot device. The default boot device is determined by the configuration switches (see Table 3-1).
[b] <i>dev</i> :	Boot file <i>defaultboot</i> from specified device.
[b] <i>dev</i> : <i>file</i>	Boot <i>file</i> on device <i>dev</i> .
h	Display a list of PROM monitor commands.
ls [<i>dev</i> :] [<i>pathname</i>]	List the contents of directory <i>pathname</i> on device <i>dev</i> . With no arguments, <i>ls</i> provides a list of files on the default device (see Table 3-3).
xns[<i>.hostname</i>]: [<i>file</i>]	Boot <i>file</i> over the network using XNS.
exit	Restart the PROM monitor.

Table 3-2: PROM Monitor Commands

Device	Description
hd0	Disk drive 0 (generic)
md0	Disk drive 0 (IRIS 3010/3020)
si0	Disk drive 0 (IRIS 3030)
hd1	Disk drive 1 (generic)
md1	Disk drive 1 (IRIS 3010/3020)
si1	Disk drive 1 (IRIS 3030)
ct0	Cartridge tape drive 0 (generic)
mq0	Cartridge tape drive (IRIS 3020/3030)
sq0	Cartridge tape drive (IRIS 3030)
fd0	Floppy disk 0 (generic)
mf0	Floppy disk 0 (IRIS 3020/3030)
sf0	Floppy disk 0 (IRIS 3030)

Table 3-3: Boot Devices

You can boot the IRIS from any attached device through the PROM monitor, regardless of the settings of switches 1 through 4. See Table 3-2 for a listing of the PROM monitor commands.

The examples below show how to boot from various devices using the PROM monitor. Table 3-3 lists the available device names. In all of the examples, if you enter a device name but no filename, the IRIS looks for a file with the name *defaultboot* on the specified device. If you enter only a device or filename, or both, with no specific PROM monitor command, the IRIS tries to boot from the specified device and file.

Booting from the Tape Drive

To boot from the tape drive, specify the tape drive device name and filename. For example, to boot the file *mdfex* on the tape drive, enter this command:

```
ct0:mdfex
```

To be able to boot from a tape, it must be in *cpio* format. See Section 4.4 for information on making a bootable backup tape.

Booting from the Disk

To boot from the disk drive, specify the disk device name and filename. For example, to boot the file *vmunix1* on the **a** partition of drive number one, **hd1**, enter this command:

```
hd1:vmunix1
```

Booting over the Ethernet

To boot the IRIS over an XNS Ethernet, use this command syntax:

```
xns[.hostname]:filename
```

If the hostname is omitted, the first host to recognize *filename* responds.

For example, to boot the IRIS from the file */usr/local/boot/goboot* on the host *cruncher*, enter this command:

```
xns.cruncher:/usr/local/boot/goboot
```

3.4.2 Automatic Boot

When configuration switch 5 is set for automatic boot, the IRIS attempts to boot the file *defaultboot* on the device specified by configuration switches 1 through 4 (see Table 3-1).

To configure the IRIS to boot automatically:

1. Set the **Power** switch on the display monitor to **On**.
2. Set configuration switch 5 to **Open**.
3. Set configuration switches 1 through 4 to the appropriate settings for the boot device (see Table 3-1).
4. Set the **Power** switch on the cabinet front panel to **On**.

NOTE: If the power for the IRIS is already on, press the **Reset** button. Do not press the **Reset** button while the IRIS is running UNIX without first using the */etc/reboot* command.

4. Workstation System Administration

The system administrator is responsible for configuring the IRIS workstation to meet local requirements and for maintaining the installed system. This chapter covers the following procedures:

<i>Section</i>	<i>Topic</i>
4.1 Basic Operation and Shutdown	
4.2 UNIX Configuration	Naming the workstation Adding user accounts Booting into multi-user mode automatically UNIX configuration files Special device files in <i>/dev</i> XNS network commands
4.3 Network Communication	Using multiple kernels
4.4 Backups	Making a bootable tape Making periodic backups Storing key UNIX files on a second disk
4.5 Crash Recovery	Recovering from a crash by: <ul style="list-style-type: none">• using <i>init</i>• booting from tape• booting from the second disk

4.1 Basic Operation and Shutdown

If you have not yet installed and booted the workstation, see Chapters 2 and 3. For a description of the startup environment, see Section 3.3.

The IRIS workstation boots into single-user mode, which should be used only for system maintenance. For normal operation, the workstation must be put in multi-user mode, as described in Section 3.3. Section 4.2 explains how to configure the workstation to boot into multi-user mode automatically.

CAUTION

Do not press the **Reset** button while the IRIS workstation is running UNIX. For information on rebooting the system, see Chapter 3. If the IRIS workstation is not running UNIX and is under control of the PROM monitor, then the **Reset** button or the **Power** switch may be used.

Before turning off power to the workstation, you should perform a soft reboot of the system to ensure that all in-memory disk-blocks have been written to the disk and that all directory information on the disk is current.

To shut down the workstation, follow these steps:

1. Make sure there are no other users on the system:

```
who
```

2. Become the super-user.

```
su
```

3. Reboot the system:

```
reboot
```

4. Set the **Power** switches on the cabinet and monitor to **Off**.

4.2 UNIX Configuration

One of the strengths of the UNIX operating system is its flexibility. The system administrator configures the workstation to the needs of the local user community, and the users can fine-tune their own environments.

The following subsections contain instructions for performing common system administration tasks on an IRIS workstation. To perform these tasks, become the super-user by logging in as **root**:

```
IRIS login: root
```

Once you are logged in as **root**, you can return to a login shell by entering the **exit** command.

For more information on UNIX configuration, see the *UNIX Programmer's Manual* and the list of configuration files in Section 4.2.4.

4.2.1 Naming an IRIS Workstation

The default name of a new IRIS workstation is IRIS. If you have more than one workstation on a network, you must assign each workstation a unique name. The name can be up to eight characters long and must contain no blanks.

To change the name of a workstation, edit the file `/etc/sys_id` and reboot the system. To display the system name at login, change the name in the file `/etc/gettydefs`. To display a message at login, edit the file `/etc/motd` (message of the day). To determine the current name of the system, use the `hostname(1)` command.

4.2.2 Adding a New Account

The IRIS workstation is shipped with six user accounts: *root*, *rootcsh*, *rootsh*, *guest*, *demos*, and *mexdemos*. The *root* and *rootcsh* accounts are super-user accounts with a C shell environment. The *rootsh* account is a super-user account with a Bourne shell environment. The *guest* account is a sample user account with a C shell environment. The *demos* and *mexdemos* accounts are for running demonstrations (see Chapter 9).

The accounts shipped on a new IRIS workstation have no passwords. To establish passwords for these accounts, follow the procedure below. See the *passwd* manual page for more information.

Typically, each user on an IRIS workstation is given his or her own account. The procedure for creating a new account on the IRIS workstation consists of adding the new user to the file `/etc/passwd`, creating a home directory for the user, and setting up startup files to define the user's environment. The details of this procedure are presented below. You should become the super-user to create a new account (see above).

1. Add a line for the new account to the file `/etc/passwd`. Figure 4-1 contains a sample `/etc/passwd` file.

The file `/etc/passwd` contains a line for each account on a UNIX system. Each line has seven fields separated by colons (:).

Field	Contents
1	Account name
2	Encrypted user password
3	User number
4	Group number
5	User's real name
6	Home directory (default /)
7	Login shell (default <code>/bin/sh</code>)

```

root::0:0:Superuser:/:/bin/csh
rootcsh::0:0:Superuser:/:/bin/csh
rootsh::0:0:Superuser:/:/bin/sh
daemon:*:1:1:/:
bin:*:2:2:Binary Files:/:
uucp:*:3:5:UUCP Login Account:/usr/spool/uucppublic:/usr/lib/uucp/uucico
adm:*:5:3:Administration:/usr/adm:
iris::7:0:/:/bin/tesh
uucpadm:*:8:8:UUCP Administration:/usr/lib/uucp:
lp:*:9:9:Line Printer:/:
mexdemos::996:997:/:usr/people/mexdemos:/bin/csh
demos::997:997:/:usr/people/demos:/bin/csh
guest::998:998:/:usr/people/guest:/bin/csh
games:*:999:999:Games:/usr/games:/bin/sh
steve::11:20:Steve Brown:/usr/people/steve:/bin/csh

```

Figure 4-1: Sample */etc/passwd* File

Be sure that the new entry contains a unique account name, a unique user number, and a group number. The user number is used to assign ownership of files to that user.

The group number is the group identification given to the user at login. When the user logs in, he or she has access to a pool of files shared by all users in that group.

The group number corresponds to the groups in the file */etc/group* (see *group(4)*). The */etc/group* file shipped on the workstation has group number 20 set to *user*; this is a default group number to use for new users.

Leave the password field empty. The system encrypts the password and adds it to the file when the user establishes a password (see step 4). In Figure 4-1, the asterisks (*) in the password fields for some accounts prevent those names from being used as login accounts.

Ordinarily, the home directory in field 6 is the user's own directory (see step 2). The login shell in field 7 can be either the C shell (*/bin/csh*) or the Bourne shell (*/bin/sh*). The default is the Bourne shell.

For example, to add a new account for a user named *jose* to the sample file in Figure 4-1, you might enter this line:

```
jose::10:20:Jose Aranda:/usr/people/jose:/bin/csh
```

See *passwd(4)* for more information.

2. Make a home directory for the new user, with the same name specified in field 6 of the */etc/passwd* entry, and set the protections and ownership of the new directory. Typically, all users' home directories are made subdirectories of the directory */usr/people*.

The *mkdir(1)* command creates the new directory.

```
mkdir /usr/people/jose
```

The *chgrp(1)* command changes the group to which the directory belongs. The example below puts the directory */usr/people/jose* in the group *user*.

```
chgrp user /usr/people/jose
```

The *chmod(1)* command sets the protection parameters on a file or directory. The example below sets protection codes to allow only the owner to create files in the directory. All users may read the file and search the directory. Protection parameters can be set only by the super-user or by the owner of the file or directory.

```
chmod 755 /usr/people/jose
```

The *chown(1)* command changes the ownership of the directory. The ownership of a directory can be set only by the superuser. The example above establishes *jose* as the owner of the directory */usr/people/jose*.

```
chown jose /usr/people/jose
```

3. Jose is a member of group *user* by virtue of his */etc/passwd* entry. He can be in other groups as well, defined by */etc/group*. To include the new user in groups other than *user*, edit the file */etc/group*. Entries to this file are optional. See *group(4)* for more information.

```

sys:*:0:
system:*:0:
daemon:*:1:
bin:*:2:
adm:*:3:
sgi_use:*:4:
uucp:*:5:uucp
sgi_use:*:6:
sgi_use:*:7:
uucpadm:*:8:uucp
lp:*:9:
user:*:20:
vlsi:*:21:steve

```

Figure 4-2: Sample */etc/group* File

Figure 4-2 contains an example */etc/group* file. The file contains a line for each group on the system. Each line consists of four fields separated by colons (:).

Field	Contents
1	Group Name
2	Encrypted Group Password
3	Group Number
4	Group Members

To add a user to a group, append the user's name to the list of group members. Separate group members with commas. For example, to add *jose* to the group *vlsi*, append the name *jose* to the line for the group *vlsi*:

```
vlsi:*:21:steve,jose
```

4. The new user should log in and use the *passwd* command to establish a password for the account. Issue the *exit* command to return to a login shell if you are now the super-user.

```

passwd
New password: password
Re-enter new password: password

```

The password is encrypted and added to the file */etc/passwd*.

If the user forgets his or her password, the encrypted password can be removed from the user's entry in */etc/passwd*.

The *passwd* command can also be used to change passwords. The *passwd* program queries the user for the old password before establishing a new one.

5. Add startup files like *.cshrc*, *.login*, and *.profile* to the new user's home directory. These files control the details of the user's environment. You can copy the startup files from */usr/people/guest*, modifying them if necessary.

The choice of the C shell or the Bourne shell determines which files are necessary. If you choose the C shell, the user's home directory should contain the files *.login* and *.cshrc*. The C shell reads the *.cshrc* file in the user's home directory each time a shell is started. If the shell is a login shell, the C shell then reads the *.login* file in the user's home directory.

If you choose the Bourne shell, the user's home directory should contain the file *.profile*. The Bourne shell reads the *.profile* file in the user's home directory when a login shell is started. For more information on the C shell and the Bourne shell, see *cs(1)* and *sh(1)*.

Each user's shell startup file should use the *tset* command to read */etc/ttytype* and set the terminal type, so that screen editors and other programs know how to communicate with the terminal. The usual way to use *tset* with the C shell is to add these lines to your *.login* file:

```
set noglob
set tmp = (`tset -s -Q`)
setenv TERM = tmp[1]
setenv TERMCAP = tmp[2]
unset tmp noglob
```

Add this line to your *.profile* file to use *tset* with the Bourne shell:

```
eval `tset -s -Q`
```

For more information, see *tset(1)* and *ttytype(4)*.

4.2.3 Entering Multi-user Mode Automatically

By default, the IRIS workstation boots into single-user mode, which should be used only for system maintenance. When the system is running in single-user mode, only the root file system is mounted, none of the system daemons are started, and the terminal lines do not allow users to log in. To enter multi-user mode, the super-user issues the *multi* command.

The super-user can also edit the file */etc/inittab* so that the system automatically enters multi-user mode each time it is booted. The first line of */etc/inittab* tells *init* which mode to enter after booting UNIX:

```
is:s:initdefault:
```

Replace the *s* (for single) in the second field with 2 or 3:

XNS protocol:

```
is:2:initdefault:
```

TCP/IP protocol:

```
is:3:initdefault:
```

The system will now enter multi-user mode each time it is rebooted. For more information on the *init* process, see *init(1M)* and *inittab(4)*.

4.2.4 UNIX Configuration Files

The directories */etc* and */usr/lib* contain the major system configuration files. They can be edited by the system administrator as needed to serve the needs of the local user community.

<i>/etc/bcheckrc</i>	This file contains commands to perform a file system check, and to set the date using <i>/etc/rc.getdate</i> if possible, or to prompt for the date if not.
<i>/etc/checklist</i>	This file lists the file systems processed by default by <i>fsck</i> .
<i>/etc/cshrc</i>	This file is read at login by accounts that specify the C shell as the login shell. It contains commands that define the user's environment. See <i>cs(1)</i> and <i>cshrc(4)</i> .
<i>/etc/gettydefs</i>	This file contains entries for line speeds and terminal settings used by <i>getty(1M)</i> for initializing devices. See <i>gettydefs(4)</i> .
<i>/etc/group</i>	This file contains the list of groups that are used to control access to files. See Section 4.2 and <i>group(4)</i> .

<i>/etc/inittab</i>	This file contains the table that <i>init</i> uses to dispatch processes for initializing devices, starting daemons, and starting the shell. See <i>inittab(4)</i> , and <i>init(1M)</i> .
<i>/etc/model</i>	This file contains the model number of the IRIS.
<i>/etc/motd</i>	This file contains the message of the day, which is displayed each time a user logs in to an IRIS workstation.
<i>/etc/passwd</i>	This file contains information about the users with accounts on the workstation. It includes each user's name, user number, group number, home directory, and login shell. See Section 4.2.2 and <i>passwd(1)</i> .
<i>/etc/profile</i>	This file is read at login by accounts that specify the Bourne shell as the login shell. It contains commands that define the user's environment. See <i>sh(1)</i> or <i>profile(4)</i> .
<i>/etc/rc</i>	This command file is executed by <i>init(1M)</i> at the start of multi-user mode. Typically, it is used to start daemons and mount file systems. See <i>brc(1M)</i> .
<i>/etc/rc.fs</i>	This shell script mounts file systems on appropriate devices, according to the model number of the workstation. It is called by the program <i>/etc/rc</i> when multi-user mode is entered.
<i>/etc/sys_id</i>	This file contains the name of the system, which is always <i>IRIS</i> on a new workstation. See Section 4.2, <i>hostname(1)</i> , and <i>sys_id(4)</i> .
<i>/etc/rc.getdate</i>	If a file with this name exists and is executable, the standard output of its execution is used as the argument to <i>date(1)</i> . Otherwise, the user is prompted for the date.
<i>/etc/termcap</i>	This file contains definitions of different terminal types. See <i>termcap(4)</i> .
<i>/etc/ttytype</i>	This file contains a mapping of the terminal types attached to each tty port on the IRIS workstation. This file is read by <i>tset(1)</i> and by <i>login</i> to initialize the terminal (TERM) environment variable at login. See <i>tset(1)</i> and <i>ttytype(4)</i> .
<i>/etc/TZ</i>	This file contains an entry for the time zone. Several different utilities use this file. <i>/etc/TZ</i> contains three fields:

1. Standard heading for time zone
2. Offset from Greenwich Mean Time (in hours)
3. Optional daylight savings time zone

The IRIS workstation is shipped with the time zone set for Pacific Standard Time:

```
PST8PDT
```

These are examples for other time zones:

```
EST5EDT
CST6CDT
MST7MDT
```

For more information, see *TZ(4)*.

/usr/lib/crontab

This file contains entries for commands to be run at fixed intervals by the *cron* daemon (the clock daemon). See *cron(1M)*.

/usr/lib/acct/holidays

This file contains a list of holidays. It is used by the *calendar(1)* program.

/usr/lib/Mail.rc

This file contains set commands used for initializing the mail system.

The next four files configure the UNIX-to-UNIX communication system UUCP. Further information on UUCP can be found in the *UNIX Programmer's Manual, Volume II*.

/usr/lib/uucp/L-devices

This file sets the line speeds for the ports used by *uucp(1C)*. The file contains a series of one-line entries that follow the format of this example:

```
DIR tty06 0 4800
```

This entry specifies that the device */dev/tty06* is a directly connected computer that can be used at 4800 baud.

/usr/lib/uucp/L-dialcodes

This file contains the dial-code abbreviations used in the */usr/lib/uucp/L.sys* file. The entries are in the form:

```
abb dial-sequence
```

abb is the abbreviation of a location, and *dial-sequence* is the dial sequence associated with the location. For example:

sf 415

This entry sends the sequence 415 to the dial unit.

/usr/lib/uucp/L.sys

This file contains information about sites that *uucp(1)* can communicate with.

/usr/lib/uucp/USERFILE

This file contains information that limits user accessibility. It specifies which files a user on the local machine can access, which files can be accessed from a remote computer, which login name is used by a particular remote computer, and whether a remote computer should be called back to confirm its identity.

4.2.5 Device Files

File	Description
drum	Paging device (used for virtual memory)
EXOS/*	Directory of TCP/IP devices
floppy	Optional floppy disk drive (generic blocked device)
rfloppy	Optional floppy disk drive (generic raw device)
ib[0-9]	IEEE-488 devices
kmem	Kernel memory; see <i>mem(7)</i>
md0[a-h]	Disk zero partitions (IRIS 3010/3020 block devices) md0a contains the root (/) file system; md0c contains the user (/usr) file system (IRIS 3020) or a copy of the root file system (IRIS 3010). See <i>efs(4)</i> .
rmd0[a-h]	Disk zero partitions (IRIS 3010/3020 raw devices)
md1[a-h]	Optional disk partitions (IRIS 3010/3020 block devices)
rmd1[a-h]	Optional disk partitions (IRIS 3010/3020 raw devices)
mem	Memory; see <i>mem(7)</i>
mf0a	Optional floppy disk drive (IRIS 3010/3020 blocked device)
rmf0a	Optional floppy disk drive (IRIS 3010/3020 raw device)

Table 4-1: Special Device Files in */dev*

File	Description
null	Null device (zero length on input, data sink on output) See <i>null(7)</i> .
pxd	IBM driver
rmt1	Cartridge magnetic tape (generic raw blocked device)
rmt2	Cartridge magnetic tape (generic raw blocked device and no rewind on open or close)
rmt3	Half-inch magnetic tape (raw blocked device)
rmt4	Half-inch magnetic tape (raw blocked device with no rewind on open and close)
sf0a	Optional floppy disk drive (IRIS 3030 blocked device)
rsf0a	Optional floppy disk drive (IRIS 3030 raw device)
si0[a-h]	Disk zero partitions (IRIS 3030 block devices) si0a contains the root (/) file system; si0c contains the user (/usr) file system. See <i>efs(4)</i> .
rsi0[a-h]	Disk zero partitions (IRIS 3030 raw devices)
si1[a-h]	Optional disk partitions (IRIS 3030 block devices)
rsi1[a-h]	Optional disk partitions (IRIS 3030 raw devices)
sq0	Cartridge magnetic tape (IRIS 3030 raw blocked device)
nrsq0	Cartridge magnetic tape (IRIS 3030 raw blocked device and no rewind on open or close)
syscon	System console (linked to <i>/dev/console</i>)
sys tty	System console (linked to <i>/dev/console</i>)
tek	Color graphics printer device
tty	A synonym for the terminal device associated with a process See <i>termio(7)</i> and <i>tty(7)</i> .
ttyd1	Serial Port 2 on standard I/O panel
ttyd2	Serial Port 3 on standard I/O panel
ttyd3	Serial Port 4 on standard I/O panel
ttyn*	XNS Ethernet network terminal devices
ttyw[0-9]	Window ttys; <i>ttyw0</i> is the console.
ttyT?	Pseudo tty devices

Table 4-1: Special Device Files in */dev* (continued)

4.3 Network Communication

An IRIS workstation can communicate with other devices over an Ethernet local area network using the XNS and TCP/IP protocols. Only one communication protocol can be used at one time.

The following sections present a brief overview of the commands used for communicating over the Ethernet. For more information about the network software for the IRIS workstation, see the manual pages for these commands in the *UNIX Programmer's Manual, Volume IA*. To make the workstation emulate a terminal, see Chapter 5.

The following two sections use these terms:

- A *network* is a collection of computers and terminals connected together.
- A *host* is a computer on a network.
- A *local* host is the host you are using.
- A *remote* host is another computer on the same network.

4.3.1 XNS Network Commands

To communicate over an XNS Ethernet local area network, the IRIS workstation uses three commands:

<i>xcp</i>	Copies a file from one computer to another.
<i>xx</i>	Runs a command on a remote host.
<i>xlogin</i>	Initiates a login on a remote host.

xcp copies a file from one host to another. The command line specifies both the source host and file first, followed by the destination host and file:

```
xcp [sourcehost:]pathname [destinationhost:]pathname
```

If no host name is specified, the local host is assumed. In these examples, the user must have an account with the same user name on both hosts.

```
xcp sqiral.c olympus:/oh4/doc/install/sqiral.c
```

This example copies the file *sqiral.c* in the current directory of the local machine to the file *sqiral.c* in the directory */oh4/doc/install* on a host named *olympus*.

```
xcp sting:/usr/include/stdio.h test
```

This example copies the file */usr/include/stdio.h* from a remote machine *sting* to the file *test* on the local machine.

```
xcp puppy:temp_vi sting:temp_vi
```

This example copies the file *temp_vi* from the remote machine named *puppy* to a file with the same name on a remote machine named *sting* in the user's home directory on each machine.

```
xcp -r -v sting:/usr/include/* localinclude
```

This command is an example of recursive copying. It copies all files and directories from */usr/include* on the remote machine *sting* to the directory *localinclude* on the local host. The *-v* option prints the names of the files it creates.

xx runs commands on a remote host. Like *xcp*, this command assumes that the user has accounts under the same user name on both the remote and local host.

The *xx* command takes the command to be executed as an argument. For example, to find the load average on another machine, enter this command:

```
xx hostname uptime
```

xlogin initiates a login on a remote host across the network. The command takes the remote host name as an argument. For example, to log in remotely to a host with the name *olympus*, enter this command:

```
xlogin olympus
login:
...
```

For more information about the network software for the IRIS workstation, see the manual pages for these commands in the *UNIX Programmer's Manual, Volume IA*.

4.3.2 TCP/IP Network Commands

This section describes the commands for communicating over an Ethernet local area network using TCP/IP. You must be running the TCP/IP kernel to use TCP/IP communications (See Section 4.3.3 below).

Each computer to be accessed with the network communication tools must have a unique name. The file */etc/hosts* contains these names. (See the documentation that accompanies the network software for more information.)

To communicate over an TCP/IP Ethernet local area network, the IRIS workstation uses five commands:

rcp Copies a file from one computer running UNIX to another computer running UNIX.

<i>rsh</i>	Starts a shell on a remote host running UNIX.
<i>rlogin</i>	Initiates a login on a remote host running UNIX.
<i>ftp</i>	Transfers files to and from a remote host using FTP protocol.
<i>telnet</i>	Communicates with a host using TELNET protocol.

rcp copies a file from one host to another. The command line specifies the source host and file, followed by the destination host and file:

```
rcp [sourcehost:]pathname [destinationhost:]pathname
```

If no host name is specified, the local host is assumed. In these examples, the user must have an account with the same user name on both hosts.

```
rcp sqiral.c olympus:/oh4/doc/install/sqiral.c
```

This example copies the file *sqiral.c* in the current directory of the local machine to the file *sqiral.c* in the directory */oh4/doc/install* on a host named *olympus*.

```
rcp sting:/usr/include/stdio.h test
```

This example copies the file */usr/include/stdio.h* from a remote machine *sting* to the file *test* on the local machine.

```
rcp puppy:temp_vi sting:temp_vi
```

This example copies the file *temp_vi* from the remote machine named *puppy* to a file with the same name on a remote machine named *sting* in the user's home directory on each machine.

```
rcp -r sting:/usr/include/* localinclude
```

This command is an example of recursive copying. It copies all files and directories from */usr/include* on the remote machine *sting* to the directory *localinclude* on the local host.

rsh connects your terminal to a remote host and executes the commands you specify. Like *rcp*, this network utility assumes that the user has accounts under the same user name on both the remote and local host.

If a command is specified, the *rsh* program takes the command as an argument. For example, to find the load average on another machine, enter this command:

```
rsh hostname uptime
```

rlogin initiates a login on a remote host across the network. The command takes the remote host name as an argument. For example, to log in remotely to a host named *olympus*, enter this command:

```
rlogin olympus
login:
...
```

If you do not specify a command to be executed, *rlogin* starts a login process on the remote host.

Like *rlogin*, *telnet* also initiates a login on a remote host across the network. The program takes the name of the host as an argument. For example, to log in remotely to a host named *puppy*, enter this command:

```
telnet puppy
```

The screen shows:

```
Trying...
connected to puppy.
Escape character is ^].
```

```
login:
...
```

To disconnect from the host, first log out from the host. Then, exit the *telnet* program. To do this, type **[CTRL-J]**. Typing **[CTRL-J]** gets the attention of the *telnet* program. To the *telnet>* prompt, type:

```
quit
```

For more information on *telnet*, see *telnet(1C)*.

Using the File Transfer Program

ftp is a program for transferring files using the ARPANET File Transfer Protocol. Below is a brief description of some of the *ftp* commands. For more information, see *ftp(1C)*.

1. To start the *ftp* program, type:

```
ftp hostname
```

hostname is the host to which you wish to transfer files. For example, suppose your *hostname* is *olympus*. Then, you type:

```
ftp olympus
```

Assuming your *hostname* is *olympus*, the screen shows:

```
Connected to olympus.  
220 olympus FTP server (Version 4.81 Mon Sep 26 08:36:28 PDT 1983) ready.
```

The host then requests your name and password. The following example assumes the user's name is Peter.

```
Name (olympus:peter):peter  
Password (olympus:peter):  
331 Password required for peter.  
230 User peter logged in.
```

After you've logged in, the *ftp* prompts appears:

```
ftp>
```

2. Now you're ready to send and receive files using the commands below. Each bulleted item shows both the command syntax and an example for that command.
 - To send a file to the remote host, type a command using the following syntax:

```
send localfile remotefile
```

For example, this command sends the file *myfile* to a file of the same name on the host machine:

```
send myfile
```

The screen indicates that the file has been sent to the remote machine.

- To get a file on the remote machine and bring it to the local machine, type:

```
get localfile remotefile
```

For example, this command gets the file *yourfile* from the remote host and stores it on the local machine in the current directory as *myfile*:

```
get yourfile myfile
```

The screen indicates that the file has been received by the local machine. You can see if the file has been received by typing:

```
!ls
```

- To transfer multiple files from the local machine to a directory on the remote machine, type:

```
cd remotedirectory  
mput localfile1 localfile2 ...
```

For example, to transfer the files *thisfile* and *thatfile* to the directory */usr/john/filexfer*, type:

```
cd /usr/john/filexfer  
mput thisfile thatfile
```

The screen indicates that the files have been sent.

- To transfer multiple files from the remote machine to the local machine into the current directory, type:

```
mget localfile1 localfile2 ...
```

For example, to transfer the files *localinclude* and *graph.c*, type:

```
mget localinclude graph.c
```

The screen indicates that the files have been received.

- You can transfer either binary or ASCII files with *ftp*. By default, *ftp* transfers ASCII files. To transfer binary files, type:

```
binary
```

To transfer ASCII files, type:

```
ascii
```

3. To exit *ftp*, type:

```
quit
```

4.3.3 Using Multiple Kernels

The UNIX operating system is read into memory from a disk file at boot time. The file can have any name, but there are two names that are recognized in special ways: */defaultboot* and */vmunix*. The PROM monitor recognizes the name *defaultboot* in the root directory of the disk, and boots that file when no name is explicitly given. Certain UNIX commands recognize the name *vmunix* and expect it to contain a copy of the kernel currently running in memory.

New kernels are shipped in a directory called */kernels*, with names such as *3000.std* (the standard kernel), *3000.tcp* (the TCP/IP optional kernel), and so on. A copy of one of the kernels is installed as */defaultboot* and linked to */vmunix*. (A link is a method of giving the same file two different names - see *cp(1)*).

When a software option includes a new kernel (for example the TCP/IP option), the kernel will automatically be installed as */defaultboot* and */vmunix*. In most cases, this is the proper thing to do. The PROM monitor will boot that kernel by default, and the various kernel-dependent UNIX commands will find the image of the kernel in the */vmunix* file. However, some sites switch between more than one kernel, to have access to more than one communications option. In such an environment, it will be necessary to copy the kernel being used to */vmunix*, and link it to */defaultboot* (This is done as the last step before taking the system down with the old kernel to reboot it with the new choice.) The *kernel(1M)* program takes care of these details. The argument is the short name of the kernel you want to install, e.g. *std*, *tcp*, *ibm*, and so on. It can only be run by the superuser.

4.4 Backups

- When you install your IRIS workstation, make sure you have a *bootable backup tape* containing the stand-alone software, the root file system, and the user file system. You should receive a bootable backup tape with your workstation. Make a new bootable backup tape after you install new software or change the size of the swap partition on your disk (see Section 4.4.1).
- After the workstation is in use, make *periodic backup tapes* of either the entire disk or the user data stored on the disk (see Section 4.4.2).
- If your workstation has a second disk, store a copy of the *key UNIX files* on the second disk (see Section 4.4.3).

The IRIS workstation supports two tape drives: a cartridge tape drive and a half-inch tape drive. The cartridge tape drive can be used both for making

bootable backup tapes and periodic backup tapes, and for installing software updates. The IRIS workstation can be booted from the cartridge tape drive in case the root file system on the disk is damaged (see Section 4.5.3).

The half-inch tape drive can be used for making periodic backup tapes and for transferring files, but it is not a bootable device.

The tape drives can be used with either *tar(1)* or *cpio(1)*, the standard UNIX archiving tools. Use the *mkboot(1M)* command to make a bootable backup cartridge tape.

4.4.1 Making a Bootable Backup Tape

The bootable backup tape can be used to boot the workstation and rebuild the disk in case the file system on the disk is damaged beyond use. Make a bootable backup tape as soon as you install a new workstation. Make a new bootable backup tape whenever you install a new software distribution or change the size of the swap partition on your disk.

To make a bootable backup tape, follow these steps:

1. Become the super-user and reboot the system into single-user mode:

```
su
reboot
b
fsck
```

Correct any errors reported by *fsck*.

2. Change your working directory to the root file system:

```
cd /
```

3. Mount the */usr* file system:

IRIS 3020:

```
mount /dev/md0c /usr
```

IRIS 3030:

```
mount /dev/si0f /usr
```

If */usr* is too large, back it up onto a separate tape, using the procedure described in Section 4.4.2.

4. Put a tape in the tape drive.

5. Run the *mkboot* program. (If you have already backed up */usr* onto a separate tape, omit the */usr* argument.)

```
mkboot /usr
```

See Section 4.5.3 for the procedures for rebuilding a damaged file system from the bootable backup tape.

4.4.2 Making Periodic Backups

This section describes the procedure for making routine backups of the IRIS workstation disk(s). The disk(s) can be backed up to a tape on a local tape drive or across a network to a tape on another workstation.

Each file system is backed up onto a separate tape. Table 4-2 below gives the capacity of various tapes. The *df* command tells you how much of the file systems are used. Use this information to determine the size of tape required to back up a file system.

CAUTION	
Do not use multiple tapes to copy a file system that is too large to fit onto one tape. Inserting a second tape, even when you are prompted to do so, could cause a system crash. If more than one tape is required, back up subdirectories of the file system onto separate tapes.	

Tape Width	Tape Length	Approx. Capacity
Quarter-inch	450 ft.	40 Mb
Quarter-inch	600 ft.	60 Mb
Half-inch (PE format)	2400 ft.	44.6 Mb

Table 4-2: Approximate Tape Capacities for Backups

Table 4-3 lists samples of the UNIX commands for making periodic backup tapes. See *find(1)* for more information on the arguments to *find*.

Cartridge Tape	Half-inch Tape
Complete Backup tar c . cpio -oha1 .	tar c3 . cpio -oBa3 .
Incremental Backup find . -mtime -7 -print tar c - find . -mtime -7 -print cpio -oha1	find . -mtime -7 -print cpio -oBa3
Read Tape tar x cpio -ihumd1	tar x3 cpio -iBumd3
List Tape Contents tar xv cpio -ihtv1	tar tv3 cpio -iBtv3

Table 4-3: Sample Tape Drive Commands

In general, you should make periodic backup tapes of everything on the disk. You may want to back up dynamic parts of the file system—such as the user files—more frequently than the more stable parts of the file system. To simplify this process, you can set up the file system so that the files that change most frequently are stored in one directory, which can be backed up independently. If you follow the recommendations in Section 4.2.2 for establishing new accounts, then most user files will be stored in the directory */usr/people*.

Each *tar* or *cpio* command generates a single output file, which is written to tape. To put several of these output files on one tape, use the tape devices without rewind. When these devices are used, output files can be written one after another onto the tape; the tape does not rewind to the beginning after each write, but rather remains where it was after the last write.

NOTE: If you interrupt the non-rewind cartridge tape device on an IRIS 3020 or 2400T, you must reboot the system to use the tape drive again. To avoid this problem, remove the cartridge from the drive to stop the device.

To use the no-rewind tape devices, replace the *tar* and *cpio* commands in steps 4 and 8 below with these lines:

```
tar c2 .      # non-rewind cartridge tape device
tar c4 .      # non-rewind half-inch tape device
cpio -oh2     # non-rewind cartridge tape device
cpio -oh4     # non-rewind half-inch tape device
```


Backing Up a System with a Local Tape Drive

To back up the disk on a workstation with a local tape drive, follow these steps:

1. Become the super-user:

```
su
```

If you are only backing up the */usr* file system, go to step 6.

2. Reboot the system into single-user mode:

```
reboot
b
```

Correct any errors reported by *fsck*.

3. Install a tape in the tape drive.
4. Back up the root file system using one of these commands:

Using *tar*:

```
tar c .          # rewind cartridge tape device
tar c3 .        # rewind half-inch tape device
```

Using *cpio*:

```
cpio -oh1 .     # rewind cartridge tape device
cpio -oh3 .     # rewind half-inch tape device
```

5. Mount the */usr* file system:

```
mount /dev/md0c /usr
```

6. Change directories to */usr*:

```
cd /usr
```

7. Install another tape in the tape drive.
8. Back up the */usr* file system using one of these commands:

Using *tar*:

```
tar c .          # rewind cartridge tape device
tar c3 .        # rewind half-inch tape device
```

Using *cpio*:

```
cpio -oh1 . # rewind cartridge tape device
cpio -oh3 . # rewind half-inch tape device
```

9. If you have any additional file systems, mount them and back them up following the same procedure.

Backing Up a System without a Local Tape Drive

An IRIS workstation with no tape drive is backed up across an XNS Ethernet onto a tape loaded on another workstation. Both systems must have network communication capabilities. The system with a tape drive must have a host name, which is referred to in this description as *hostA*, and must be in multi-user mode.

To back up a disk onto a tape installed in a remote workstation, follow the procedure for backing up a system with a local tape drive, but use the command lines below for steps 4 and 8. Replace *rmt1* with *rmt2*, *rmt3*, or *rmt4* depending on your backup device (see Table 4-1).

4. Back up the root file system:

Using *tar*:

```
tar caBf - . | xx hostA dd ibs=10k obs=200k of=/dev/rmt1
```

Using *cpio*:

```
cpio -oha . | xx hostA dd ibs=10k obs=250k of=/dev/rmt1
```

8. Back up the *usr* file system:

Using *tar*:

```
tar caBf - . | xx hostA dd ibs=10k obs=200k of=/dev/rmt1
```

Using *cpio*:

```
cpio -oha . | xx hostA dd ibs=10k obs=250k of=/dev/rmt1
```

4.4.3 Storing Key UNIX Files on a Second Disk

If you have two disks on your workstation, store a copy of the key UNIX files on the second disk. You could boot from the second disk if the root file system on the first disk became corrupted (see Section 4.5.4).

See Chapter 6 for a discussion of disk configuration. Note that it is not possible to boot from a second disk that has been configured as a single file system.

The procedure below assumes that the file system on the a partition of the second disk is named */n*. To store a copy of key UNIX files on the second disk, follow these steps:

1. Change to the root directory:

```
cd /
```

2. Create a file that contains the key UNIX files. (The backslash in this command is typed only when immediately followed by a **RETURN** as is the case here):

```
find .cshrc .login .profile vmunix* bin etc dev \  
-print -type f | sort > /keyfiles
```

The key UNIX files consist of:

- five files from the root directory,
 - all of the files in the */dev* directory, and
 - some of the files from the */bin* and */etc* directories.
3. Edit */keyfiles* to remove any files prefixed with *bin* and *dev* that are not listed below.

<i>bin/cat</i>	<i>bin/chgrp</i>	<i>bin/chmod</i>	<i>bin/chown</i>	<i>bin/cp</i>
<i>bin/cpio</i>	<i>bin/csh</i>	<i>bin/dd</i>	<i>bin/echo</i>	<i>bin/ln</i>
<i>bin/ls</i>	<i>bin/mkdir</i>	<i>bin/more</i>	<i>bin/mv</i>	<i>bin/ps</i>
<i>bin/pwd</i>	<i>bin/rm</i>	<i>bin/rmdir</i>	<i>bin/sh</i>	<i>bin/sleep</i>
<i>bin/stty</i>	<i>bin/su</i>	<i>bin/sync</i>	<i>bin/tar</i>	<i>bin/tset</i>
<i>bin/vi</i>				

<i>etc/checklist</i>	<i>etc/clri</i>	<i>etc/cshrc</i>	<i>etc/fsck</i>	<i>etc/fsdb</i>
<i>etc/group</i>	<i>etc/init</i>	<i>etc/inittab</i>	<i>etc/link</i>	<i>etc/mkfs</i>
<i>etc/mknod</i>	<i>etc/mnttab</i>	<i>etc/mount</i>	<i>etc/passwd</i>	<i>etc/profile</i>
<i>etc/reboot</i>	<i>etc/sys_id</i>	<i>etc/tar</i>	<i>etc/telinit</i>	<i>etc/termcap</i>
<i>etc/ttytype</i>	<i>etc/umount</i>	<i>etc/unlink</i>		

4. Copy the files listed in */keyfiles* into a file system on the second disk. This procedure assumes the file system is named */n*.

```
cpio -pdaumv /n < /keyfiles
```

5. Move */keyfiles* to the second disk.

```
mv /keyfiles /n
```

You now have a copy of the key UNIX files on the second disk. To find out how to recover from a crash by booting from this copy, see Section 4.5.4.

4.5 Crash Recovery

Depending on the cause, a system crash can leave an IRIS workstation in a state that requires some software repair. This section describes the steps you should take immediately after a workstation crash and the procedures for repairing the disk if the software is damaged.

See Appendix B for a list of system error messages and their probable causes. Also, see the *crash(8)* manual entry.

If the IRIS workstation stops running for some reason, follow this recovery sequence:

1. Try to reboot the workstation by entering the *reboot* command:

```
reboot
```

2. If the system does not respond, wait one minute, and press the **Reset** button.

The system should attempt to follow the normal boot sequence (see the description of *init*, below).

3. Boot the system in single-user mode.
4. Run the file system check program, *fsck*. This program may identify problems with the file system (see Section 3.3 and *fsck* in the *UNIX Programmer's Manual, Volume II*).
5. If *fsck* runs without any problems, enter multi user mode:

```
multi
```

If rebooting and running *fsck* does not solve the problem, see the procedures below for identifying the problem and recovering the system.

4.5.1 Recovery Strategies

Accidental changes to the files critical to UNIX operation can prevent the UNIX operating system from booting. Table 4-4 lists some of the files that are most likely to cause serious problems.

Critical Files	
<i>File</i>	<i>Purpose</i>
/	Root directory
/vmunix	Kernel
/vmunix1	Kernel that runs on the second disk
/defaultboot	Kernel used by default in booting
/dev	Device directory
/dev/console	Console tty device
/dev/syscon	System console device
/dev/systty	System tty device
/dev/md0a	Root file system device
/dev/rmd0a	Raw root file system device
/dev/drum	Paging device for virtual memory
/etc	Miscellaneous file directory
/etc/init	Program that starts other user processes
/etc/inittab	Table for <i>init</i>
/etc/passwd	Table of user login names, passwords, and shells
/bin	Commands directory
/bin/su	Used to set environment before <i>execing</i> shell
/bin/sh	Bourne shell
/bin/csh	C shell

Table 4-4: Files Critical to UNIX Operation

The following subsections describe three procedures for diagnosing the problem and recovering a system that does not boot:

- Using the initialization program, *init*, to identify the cause of the problem.
- Using a bootable backup tape to load a new root file system onto the disk. This procedure deletes any data that was on the root file system of the disk.
- Booting the UNIX kernel from a second disk and reloading the key files. This approach requires a second disk with a copy of the files needed to boot the system.

4.5.2 Recovery with *init*

When UNIX software is booted, it starts the initialization program, *init*, which controls the startup sequence.

The *init* program provided with an IRIS workstation has been modified so that it attempts to recover from the most common file system errors. For most errors, *init* recovers automatically but does not repair the cause of the problem. Therefore, once the system comes up, you should repair the source of the problem.

Repairing the problem usually entails editing the damaged files or reloading the files from a backup tape.

When it starts running, the *init* program invokes a single-user shell by calling the switch user program, *su*. Upon startup, *su* looks in the file */etc/passwd* for an account called *root*. The *su* program then starts a shell of the type specified in the file */etc/passwd* for the account *root*. (The default shell is the C shell, */bin/csh*.)

During this process, the system does not boot correctly if it encounters any of these conditions:

1. The */etc/passwd* file is missing.
2. The entry for *root* is missing or in error in the */etc/passwd* file.
3. Any other files critical to starting UNIX are missing or corrupted.

If any of these problems exists, *init* may be unable to execute the *su* program, or *su* may execute too quickly. The first time an error occurs, *su* should take no more than 15 seconds to execute; on subsequent errors, *su* should take no more than 30 seconds to execute.

If *init* cannot execute *su*, it displays this message:

```
exec1p of /bin/sh failed; errno = n
```

Refer to the error numbers listed in the introduction to Section 2 in the UNIX Programmer's Manual for the meaning of *n*.

If *su* executes too quickly, it displays a message explaining the problem. For example:

```
SINGLE USER MODE  
SU is broken
```

If *su* does not execute correctly, *init* tries to execute the Bourne shell, */bin/sh*, bypassing *su*. If *sh* fails to execute or executes too quickly, *init* prints the error message:

```
execvp of /bin/sh failed; errno = n
```

or the message:

```
SINGLE USER MODE
/bin/sh is broken
```

If the Bourne shell fails to execute, *init* tries to execute the C shell, */bin/csh*. If *init* is able to execute one of these shells, you may be able to identify the problem and repair it yourself.

If *init* is unable to execute either shell, it tries to execute this command:

```
/bin/tar xv /vmunix /vmunix1 /bin /dev/etc vmunix vmunix1 bin dev etc
```

With this command, *init* attempts to extract key files from a cartridge tape in any possible way (both with and without leading slashes).

If this fails, */etc/tar xvpU* is tried instead of */bin/tar xv*. This new attempt extracts everything on tape, unlinking existing files to overwrite them, and giving precedence to the modes on tape.

If this fails, *init* attempts to extract whatever it can from tape (hoping for a bootable kernel), and reads it into a file called */lastditch*. It then attempts to execute this file.

Whether or not *init* is able to recover the system, you may want to rebuild the file system from a backup copy. See the procedures below for recovering from a backup tape and recovering from a backup disk.

4.5.3 Recovery by Booting from Tape

If the *root* or *usr* file system on the workstation's disk is damaged, you can rebuild them system from a bootable backup tape made for this purpose (see Section 4.4.1). This procedure writes over the contents of the disk. If the bootable backup tape does not contain the most recent user file system, you should update the user file system from the periodic backup tapes after rebuilding the disk. When recovering files from a periodic backup tape, you need to change directories to the directory from which the tape was made.

To restore a disk from a backup tape, follow these steps:

1. Set configuration switches 1 through 5 for booting from the PROM monitor (see Table 3-1).

2. Reset the workstation by pressing the **Reset** button.

The IRIS displays the PROM monitor prompt:

```
iris>
```

3. Insert the bootable backup tape in the tape drive.
4. Boot the *mdfex* or *sifex* program.

IRIS 3020:

```
b ct0: mdfex
```

IRIS 3030:

```
b ct0:sifex
```

CAUTION
You must follow all instructions exactly when you use the <i>flex</i> programs. Failure to follow instructions could result in serious damage to your file system.

The *mdfex* or *sifex* initial screen and prompt appear.

```
SGI Formatter/Exerciser Version :x Date
Drive: 0, Name: name, Serial: serial number
```

IRIS 3020:

```
mdfex x>
```

IRIS 3030:

```
sifex x>
```

NOTE: If an error message appears instead, call your local System Engineer or the Geometry Hotline (see Chapter 1).

5. Instruct *mdfex* or *sifex* to create the root file system by entering the *t* command:

```
t
```


6. Accept the *mdfex/sififex* defaults for tape file and unit numbers by pressing **RETURN** after each of these prompts:

```
Tape file (2)?  
Unit (0)?
```

7. Confirm file system *a* by pressing **RETURN** in response to the file system prompt:

```
File System (a)?
```

mdfex/sififex displays a message about the block sizes it will use for the copy. When it is ready it displays the message:

```
Type 'go <return>' to begin ...
```

8. Type:

```
go
```

followed by **RETURN**.

The system displays these messages:

```
Copy started...  
1 2 3 4....  
Tape to Disk Copy complete
```

IRIS 3020:

```
mdfex x>
```

IRIS 3030:

```
sifex x>
```

9. Re-enter the PROM monitor by entering the quit command:

```
q
```

Confirm your intention to quit by entering *y*.

10. Boot the disk:

```
b
```

UNIX should come up in single-user mode. If the user file system was not damaged, go to step 14.

If the */usr* file system suffered major damage, perform the remaining steps in this section to restore it.

If you made a separate backup tape for */usr*, remove the bootable backup tape and put the */usr* tape in the tape drive.

11. Make a clean user file system:

IRIS 3020:

```
mkfs /dev/rmd0c
labelit /dev/rmd0c usr sgi
```

IRIS 3030:

```
mkfs /dev/rsi0f
labelit /dev/rsi0f usr sgi
```

12. Mount the new */usr* file system and change to that directory:

IRIS 3020:

```
mount /dev/md0c /usr
cd /usr
```

IRIS 3030:

```
mount /dev/si0f /usr
cd /usr
```

13. Rewind the tape, position the tape to the third file on it, and copy */usr* from the tape to the disk:

```
mt rewind
mt fsf 2
cpio -ivhmd2
```

14. Any periodic backup tapes more recent than the bootable backup tape should now be used to recover the root and, if necessary, user file systems as fully as possible.

4.5.4 Recovery by Booting from the Second Disk

A small number of key files stored on a second disk is sufficient to reboot the system if the original root file system becomes corrupted. If you have copies of the key UNIX files on a second disk, you can use them to boot the workstation. It is not possible to boot from a second disk drive that has been configured as only one file system (see Chapter 6). You can then try to rebuild the root file system by copying these key files from the second disk to the first.

The procedure in this section assumes that the key UNIX files have been stored on the second disk using the procedure in Section 4.4.3.

To boot from the second disk, follow these steps:

1. Set configuration switches 1 through 5 for booting from the PROM monitor (see Table 3-1).
2. Reset the workstation by pressing **Reset** button.

The IRIS displays the PROM monitor prompt:

```
iris>
```

3. Boot the file *vmunix1* on the second disk:

```
hd1:vmunix1
```

The system displays the superuser shell prompt:

```
#
```

4. Run the file system check program, *fsck*, on the first disk (the disk you are recovering). List the file systems explicitly on the command line:

IRIS 3020:

```
fsck /dev/rmd0a /dev/rmd0c
```

IRIS 3030:

```
fsck /dev/rsi0a /dev/rsi0f
```

Repair any problems reported by *fsck*.

5. Make a */root* directory:

```
mkdir /root
```

6. Mount the root file system:

IRIS 3020:

```
mount /dev/md0a /root
```

IRIS 3030:

```
mount /dev/si0a /root
```

7. Change to the root directory and copy the key system files from the second disk to the first disk using *cpio(1)*:

```
cd /  
cpio -pdaumv /root < /keyfiles
```

8. Shut down the system with the *reboot* command. Try to boot UNIX from the first drive. If this is successful, recreate the file system on the *b* partition of the second disk. (This is required because it was used as the swap partition while (*vmunix1* was running.)

IRIS 3020:

```
mkfs /dev/md1b  
labelit /dev/rmd1b tmp sgi
```

IRIS 3030:

```
mkfs /dev/si1b  
labelit /dev/rsi1b tmp sgi
```

If problems persist, call your local System Engineer or the Geometry Hotline.

5. IRIS 3010 System Administration

This chapter contains operation and configuration procedures for the IRIS 3010. Topics covered are shown in the table below. At this point, it is assumed that you have installed and booted the IRIS (see Chapters 2 and 3).

5.1 Host Software Installation
5.2 Operation
5.2.1 Connecting to the Host
5.2.2 Issuing Commands to the Terminal Emulator
5.2.3 Running a Demonstration Program
5.2.4 Shutting Down the IRIS Terminal
5.2.5 Boot Failure Recovery
5.3 Software Configuration
5.3.1 Configuration Tools
5.3.2 Configuring the Startup Environment
5.3.3 XNS Configuration
5.3.4 TCP/IP Configuration
5.3.5 Miscellaneous UNIX Configuration

5.1 Host Software Installation

This section contains instructions for installing the host-resident IRIS 3010 support software in a VAX UNIX 4.2BSD environment. For an IRIS workstation System V environment, install the software according to Section 4 of *Installing Software Updates*. The entry name for the workstation host-resident software is *host*. Instructions for installing host software in other environments are shipped with the software.

The software is delivered on a 1600 bpi half-inch magnetic tape in *tar(1)* format.

The software can be stored anywhere on the host file system. The procedure below assumes that the distribution software is installed in a directory called */usr/iris*.

To install the host software, follow these steps:

1. Log in to UNIX from a terminal that is already attached to the host.

2. Create a directory to hold the distribution software:

```
mkdir /usr/iris
```

3. Change to the new directory:

```
cd /usr/iris
```

4. Load the distribution tape onto a tape drive and read the distribution software into the new directory, using the *tar* command:

```
tar xv
```

5. Add the *TERMCAP* entry for the IRIS Terminal in */usr/iris/c/iris.termcap* to the terminal description data base in */etc/termcap*.

6. Change directories to */usr/iris/c*:

```
cd /usr/iris/c
```

7. Compile the C language remote host Graphics Library (*librgl2.a*), the IRIS 3010 configuration tools, and the demonstration programs, using the */usr/iris/c* makefile:

```
make all
```

8. To install the library in */usr/lib* execute:

```
make install
```

9. If a FORTRAN language remote host Graphics Library (*libfrgl2.a*) is desired, change directories to */usr/iris/f77* and compile it:

```
cd /usr/iris/f77  
make all
```

The FORTRAN demonstration programs are also compiled. To install the library in */usr/lib*, execute:

```
make install
```

For a terminal with network communication capabilities, the host must also be equipped with the network communication support software and hardware. See *VAX 4.2BSD UNIX XNS Release N1.0 Installation Guide* for the hardware and software installation procedures for a host running 4.2BSD UNIX.

5.2 Operation

The IRIS 3010 is configured to run a terminal emulation program automatically after booting. The program is named *wsiris*(1). It causes the IRIS to behave like a regular character terminal, with two additional capabilities:

- Setting up communication with the host.
- Interpreting graphics command sequences from the host.

The terminal emulator program can be told to execute an interactive subshell from which UNIX commands can be issued. A complete UNIX operating system is not provided on the IRIS 3010. Only those commands necessary to boot and perform basic system administration are included.

With the IBM option installed, the IRIS 3010 runs a different terminal emulator, *t3279*(1). Its user interface is completely different from that described below.

5.2.1 Connecting to the Host

This section describes the terminal emulator interface used for connecting to the host. See Section 3.2 for the procedures for booting the terminal emulator program. See Section 5.3.2 for the procedure for configuring the terminal emulator program for the needs of a specific host. See Section 5.2.2 for the commands recognized by *wsiris*.

When *wsiris* first comes up on the IRIS, it displays the startup message:

```
IRIS GL2 Terminal Emulator
```

The terminal then prompts for instructions regarding the connection to the host. The prompt varies, depending on the configuration of the system.

- An IRIS 3010 with no Ethernet controller, with only XNS Ethernet communication software, or with previously configured TCP/IP communication software prompts for the name of the host computer:

```
Connect to what host?
```

- To establish communication over the serial line, enter `serial`.
- To establish communication over the IEEE-488 bus, enter either `serial` or `ieee488`. If the host does not support logins on the IEEE-488 bus, enter `serial`. Once you are logged in, set the environment variable `IRISFILE`

to the name of the IEEE-488 device. The serial line will be used for the login, and the IEEE-488 bus will be used for graphics. Environment variables are described in *csH(1)* and *sh(1)*.

- To establish communication over an Ethernet, enter the name of the host.

For Ethernet communication, the host must have network support software installed. For XNS communication, the host must be running the *xnsd(1M)* daemon. For TCP/IP communication, the host must be running the *telnetd(1C)* daemon. If you do not know the name of a UNIX host, enter the *hostname(1)* command on a terminal that is already logged onto such a host.

The IRIS establishes a connection with the host, which should display its usual login prompt.

- An unconfigured IRIS 3010 with TCP/IP communication capabilities prompts for the Internet addresses of the IRIS and the host:

Enter IRIS IP address:

Enter host IP address:

Only Internet addresses consisting of four numbers separated by periods (*n.n.n.n*) are accepted.

To convert a two-field Internet address to the correct format, insert two fields of zeros in the middle of the address. For example, if the representation of a host name in the */etc/hosts* table on a 4.2BSD UNIX system is "42.2", enter the address as "42.0.0.2".

The IRIS establishes a connection with the host, which should display its usual login prompt.

Once communication with the TCP/IP host has been established, you can copy an Internet host table onto the terminal's hard disk (see Section 5.3.4). You can then connect to an TCP/IP host by simply entering the name of the host when prompted. See the documentation that accompanies the communications software.

5.2.2 Issuing Commands to the Terminal Emulator

After a connection with the host has been established, the terminal emulator program recognizes some local commands, which are stripped from the data stream and processed locally instead of being transmitted to the host. The commands begin with ~.

To issue a command to the terminal emulator, enter the escape character at the beginning of a line, followed by the rest of the command. The character prior

to the '~' must have been a **RETURN** or a **LINEFEED**. *wsiris* examines the next few characters to see if they are any of the following commands:

- ~~ Send a single '~'.
- ~. Force *wsiris* to exit. Ethernet connections are closed, but serial connections may remain open.
- ~! Escape to a C shell on the IRIS 3010. The shell allows the user to run programs from the local disk, such as the flight simulator:

```
~demos/flight
```

If you are familiar with the UNIX editor *vi*, you can edit the configuration files from the local C shell.

To return to communication with the remote host, exit the C shell by entering the `exit` command.

~! can also be entered in response to the hostname prompts before a connection is established.

- ~**DEL** Reboot the IRIS with */etc/reboot*. You are asked for confirmation. This command succeeds only for the super-user.
- ~**BREAK** If using serial communications, send a break character.
- ~%**D***n* Toggle the value of the `-d n` debugging option (see Section 5.3.2). ~%**D1** interactively turns logging on and off. When you turn logging on, log output is appended to the current log file. If *n* is 3, the new logfile *name* consists of all characters following the 3 until terminated by a **RETURN**. If logging is on, the current log file is closed and the newly named one is opened. Subsequent log output is appended to it. If you specify a null file name, the name of the current log file is displayed. ~%**D2** is ignored.
- ~%**M***n* Set the monitor type to *n*. *n* can have one of five values:

0	30Hz
1	60Hz
2	NTSC
3	50Hz
9	PAL
- ~%**P** Toggle the value of the `-p` option (see Section 5.3.2). The `-p` option controls whether textport output is printed when the textport is off.
- ~%**R** Reset the display by forcing a `ginit` and a `tpon` to be executed. This will load the default color map and turn on the textport. If the window manager is running, the terminal also executes any `mapcolor` commands stored in the window manager configuration file (see Section 5.3.2).

~%**S***speed* Change the serial communications baud rate to *speed*. The parameter *speed* can have these values:

75 110 134 150 300 600 1200 2400 4800 9600 19200

The default is 9600 baud.

~%**T** Toggle the textport. That is, turn the textport on if it is off, and off if it is on.

The textport does not automatically disappear when it is turned off. If the remote system seems not to be echoing input, especially after a graphics program has been interrupted, the problem may be that the textport has been left off. Issue this command to turn it back on.

~%**U** Unlock the keyboard. Whenever the IRIS is accepting graphics commands, the keyboard is locked and any input from it is thrown away. Use this command to force characters through to the remote system, such as the interrupt character to kill the remote program. The keyboard will be locked again at the next switch back into graphics mode.

~%**X** Toggle the value of the `-x` option (see Section 5.3.2). This option determines whether the handling of `XON` (`CTRL-g`) and `XOFF` (`CTRL-s`) is done locally.

~%**Z***n* Toggle the value of the `-z n` option (see Section 5.3.2). Currently, the only valid value for *n* is 1.

The escape character can be changed from `'~'` with the `-e` option specified in the *wsiris* configuration file (see Section 5.3.2). If `-e none` is specified, there is no escape character and none of the above interpretations occur. In this case, if the window manager is not running, there is no way to force *wsiris* to exit.

If the remote program queues the keyboard, all keystrokes will normally accumulate in the event queue. To divert them to *wsiris*—in order to, for example, issue one of the `'~'` commands described above—hold the `SETUP` key down while typing. The events from individually queued keyboard keys are also not placed in the queue while `SETUP` is down. See the *IRIS User's Guide* for further information on queuing the keyboard and the event queue.

5.2.3 Running a Demonstration Program

After booting a new IRIS 3010 and connecting to a host, you should test the system by running a demonstration program. The hard disk that comes with the IRIS 3010 contains a non-bootable copy of the *flight* demonstration program. To run *flight* from the hard disk, you must first boot UNIX on the terminal, then escape into the C shell and run the program from the shell.

1. Boot the terminal from the hard disk, as described in Section 3.2.
2. When the *wsiris* terminal emulator is running, enter the sequence `~!` to enter the C shell:

```
Connect to what host? ~!
```

The terminal displays the C shell prompt:

```
%
```

3. To run the demonstration program, enter the pathname:

```
~demos/flight
```

The first screen of the *flight* program appears.

4. To leave the *flight* program, press the **ESC** key.

The IRIS again displays the C shell prompt.

5. To return to the terminal emulator program, enter the *exit* command:

```
exit
```

5.2.4 Shutting Down the IRIS Terminal

To shut down an IRIS 3010 running *wsiris*:

1. Issue the **DEL** command.
2. Set the **Power** switch on the front of the cabinet to the **Off** position.

5.2.5 Boot Failure Recovery

If a critical file on the IRIS 3010 disk becomes corrupted, the boot procedure fails and the terminal emulator does not run. If this happens, use the procedure in this section to restore the file system from the copy stored on the backup partition of the disk.

1. Turn control of the IRIS over to the PROM monitor by setting configuration switches 1 through 5 to "OCOCC" and resetting the terminal (see Section 3.4).

2. Boot the recovery system:

```
iris> b md0c:recover
```

This will boot a copy of the UNIX kernel on the backup partition, and run a program to copy the backup partition onto the normal operating partition. After this program is finished, it returns control of the IRIS to the PROM monitor, which displays its prompt:

```
iris>
```

3. Boot the terminal normally from the local disk (see Section 3.2).
4. Repeat the configuration sequence described in Section 5.3 .

If this procedure fails to return the IRIS to its former working state, the problem may be more serious than a corrupted file system. Contact Silicon Graphics field service personnel through the Geometry Hotline number listed in Chapter 1.

If you are familiar with UNIX, you can use a 9600 baud ASCII terminal connected to **Port 2** of the IRIS 3010 to log in as *root*. This may be useful for diagnostic purposes in case the standard console textport is not functioning properly. To connect an ASCII terminal to the IRIS, see Section 7.1.

5.3 Software Configuration

The IRIS 3010 software is based on the same System V UNIX kernel used on the IRIS 3020 and 3030. IRIS 3010 users do not necessarily interact with UNIX, however, because the terminal emulator program running on top of UNIX provides an interface almost identical to the interface on IRIS terminals. The terminal emulator interface is described in Section 5.2.

Because the IRIS 3010 is UNIX-based, the IRIS window manager runs on it. The window manager is described in the *IRIS User's Guide, Version 2.1*.

An IRIS 3010 can be ordered with both XNS and TCP/IP Ethernet communication capabilities. Configuration of the IRIS 3010 for network communication is not necessary for normal operation, but some configuration can simplify the process of communicating over the network and can increase the terminal's communication capabilities.

5.3.1 Configuration Tools

A number of files stored on the IRIS 3010's hard disk allow for three kinds of configuration:

- Startup environment configuration (see Section 5.3.2).

- Network communication configuration (see Sections 5.3.3 and 5.3.4).
- Miscellaneous UNIX configuration (see Section 5.3.5).

The following subsections describe these files in detail. Samples of the configuration files are duplicated in the host software distribution.

If you are familiar with UNIX and its utilities, you can edit these files locally on the IRIS 3010, using the *wsiris* shell escape command (*~!*) and the *vi* editor. If you prefer, you can edit the files on the host, using whatever text editor is most familiar to you, then transfer the files to the disk on the IRIS 3010. This strategy has the advantage of leaving backup copies of the customized configuration files on the host.

To support configuration of the IRIS 3010 from a host, the host software distribution includes two utility programs: *iftp* and *irsh*. Both programs are available when you are logged on to the remote host from the IRIS 3010. These commands are run from the host while the IRIS 3010 is running *wsiris*.

File Transfer Program

The *iftp* program transfers files between the host and the IRIS 3010. The command takes three arguments:

```
iftp mode hostfile 3010file
```

Mode can have one of three values:

- w** Write *hostfile* to *3010file*. A temporary file on the IRIS 3010 is written to and then renamed.
- o** Overwrite. As **w**, but the temporary file is not made.
- r** Read from *3010file*, write to *hostfile*.

3010file must be an absolute pathname (i.e., it must start with '/').

Local Execution Program

The *irsh* program issues commands to the UNIX operating system running on the IRIS 3010 from the host. The program takes the IRIS 3010 UNIX command and arguments as its arguments:

```
irsh 3010cmd arg ...
```

The main uses of *irsh* are renaming and removing files on the IRIS 3010 disk.

To rename the file *oldname* to *newname*, enter this sequence:

```
irsh mv oldname newname
```

To remove *file*, enter:

```
irsh rm file
```

Configuration from a Unix Host

After the distribution software has been installed on the UNIX host, you can use this procedure to configure the IRIS 3010 from the host:

1. Boot the IRIS 3010 from its local disk and log in to the host (see Section 3.2).
2. Change directories to */usr/iris/c*.

```
$ cd /usr/iris/c
```

3. Edit the configuration files as described in the remainder of this chapter.
4. Transfer the files to the IRIS 3010 by invoking the *iftp* program for each file. To transfer all files in a single step, use the makefile argument *install-config-files*.

```
$ make install-config-files
```

5.3.2 Configuring the Startup Environment

You can configure the startup environment on the IRIS 3010 in these ways:

1. Custom-configuring the terminal emulator program, *wsiris*.
2. Causing the window manager to start automatically during bootstrap.
3. Installing your own startup program.

The next three sections describe these procedures.

Terminal Emulator Configuration

The terminal emulator program, *wsiris*, is customized through the file *.wsirisrc*. The working copy of *.wsirisrc* is stored on the IRIS 3010's hard disk. The file contains any number of the option flags listed below, separated by whitespace (spaces, tabs, and newlines), and possibly followed by a hostname. Arguments to individual option nags are separated from their options by whitespace.

The file may contain any of these options:

- d n** Debug. You may use multiple **-d** options. The legal values for *n* are:
- 1 Generate a logfile, named */LOGFILE.0*, of the communication between the IRIS 3010 and the host. Graphics commands and data are interpreted and written to the logfile as a command name and a series of arguments. The arguments are shown as *type=value*, where *type* is a single character representing the argument type:

a, A	array
b, B	byte (8 bits)
c, C	character string
f, F	float (32 bits)
l, L	long (32 bits)
s, S	short (16 bits)
o	logical (8 bits)
O	logical (32 bits)

If the character is capitalized, then the the argument was sent by the IRIS, otherwise it was sent by the remote system. The *value* of array types is always 'array'.

- 3 Use the next argument for the name of the logfile created by the **-d 1** option in place of */LOGFILE.0*.
- e c** Set the escape character from the default '~' to *c*. Lines beginning with the escape character can be commands to *wsiris*. If *c* is **none**, there will be no escape character.
- h** Use half-duplex serial communications. The default is full-duplex.
- i** Try the TCP/IP communication protocol first. When using Ethernet communications in the absence of the **-i** option, the system first attempts a connection to *hostname* using XNS protocols. If that fails, a connection using TCP/IP is then attempted.
- l line** Use device *line* for the serial connection. The device names of the three available serial ports of the IRIS 3010 are:

Port 2	ttyd1
Port 3	ttyd2
Port 4	ttyd3

The default *line* is 'ttyd2'.

- p** Print textport output even when the textport has been turned off. By default, after the receipt of a **tpoff** command from the remote system, all textport output is thrown away until a **tpon** is received. (See the *IRIS User's Guide* for information on **tpoff** and **tpon**.) The **-p** option inhibits this, and characters written to the textport will appear when

it is turned on again.

-s *speed* Use *speed* baud for serial communications. The legal values for *speed* are:

75 110 134 150 300 600 1200 2400 4800 9600 19200

The default is 9600.

-x Enable local `XON/XOFF` flow control. Normally, when read from the standard input, the characters `XON` (`CTRL-q`) and `XOFF` (`CTRL-s`) are sent to the remote system. Specifying this option causes them to be used locally to stop and restart the display of text. When using serial communications, this option also enables the use of these characters by the serial driver of the IRIS for automatic flow control of data being received from the host. This automatic flow control is normally disabled as it will not work correctly if there are two independent sources of these characters.

-y If using serial communications, interpret `XON` and `XOFF` characters from the remote host as output flow control commands. By default, they are treated as ordinary data. You must not give this flag if you are using eight-bit (fast mode) communications.

-z *n* Special instructions. Multiple **-z** options can be given. The legal values for *n* are:

- 1 Interpret `ginit` commands as `gbegin` commands. (See the *IRIS User's Guide* for information on `ginit` and `gbegin`.)
- 2 Execute a `ginit` immediately upon startup. This is useful when *wsiris* has had to be restarted and the remote program is still sending graphics commands.
- 3 When attempting a connection using Ethernet communications, try only one network protocol instead of both. TCP/IP is tried if the **-i** flag is given, and XNS if it isn't.
- 4 Emulate a Tektronics 4010.

Following the option flags, a hostname can be specified. Doing this causes *wsiris* always to connect to this host and not to give the prompt:

```
Connect to what host?
```

The IRIS 3010 is shipped with a default */.wsirisrc* file that contains these options:

```
-z 3 -x
```

The host distribution contains a copy of this file named *wsirisrc*. To configure the IRIS 3010 from the host, edit the host copy of the file, then transfer it to the IRIS with this command:

```
iftp w wsirisrc /.wsirisrc
```

The default configuration file may not work for some serial communications environments, e.g., if half-duplex is required. In this case, */.wsirisrc* will have to be edited locally on the IRIS 3010.

If an invalid */.wsirisrc* is put on the IRIS 3010 disk, it will not boot correctly. If this happens, use the boot failure recovery procedure described in Section 5.2.5.

Window Manager Configuration

The IRIS 3010 can be configured to start the window manager, *mex*, automatically. The window manager options are controlled through the file */.mexrc*. The *IRIS User's Guide* contains a description of what *mex* expects to find in this file. The default configuration file on the IRIS 3010 is different from the one described in the *IRIS User's Guide*.

The boot sequence checks for the existence of */.mexrc*. Only if it finds one will *mex* be started. The IRIS 3010 is shipped with the default window manager configuration file with the name */.mexrc*, which means that *mex* is not started up during bootstrap. To have *mex* start automatically, change the filename to */.mexrc* by entering this command on the host:

```
irsh mv /mexrc /.mexrc
```

To disable *mex* once it's been enabled, enter:

```
irsh mv /.mexrc /mexrc
```

The host distribution contains a copy of *mexrc*. After editing this file, you can copy it to the IRIS with this command:

```
iftp mexrc /.mexrc
```

Installing Your Own Startup Program

The program that the IRIS 3010 runs after booting can be changed from *wsiris* to a custom startup program. To do this, edit the file */bin/tesh* and find the line:

```
set termemul = /bin/wsiris
```

Replace */bin/wsiris* with the pathname of the program you want to have run.

If the IBM option is installed, *t3279* is run instead of *wsiris*, and this line in */bin/tesh* must be changed:

```
set termemul = /bin/t3279
```

Replace */bin/t3279* with the pathname of the program you want to have run.

If you add custom files to the root partition, you should duplicate them on the backup partition so that they can be restored in case of boot failure. Before copying these files, make sure that the IRIS boots and correctly runs your startup program. The IRIS 3010 disk contains a script, */etc/rootcopy*, that copies the entire root partition to the backup partition. Run *rootcopy*, specifying the argument *ok*:

```
rootcopy ok
```

5.3.3 XNS Configuration

Operation of the IRIS 3010 using XNS requires no configuration. The IRIS does not need a hostname, since it initiates all connections, and the XNS protocol does not require that both ends of a circuit have a name.

The terminal emulation program can be configured to initiate a connection to a single host automatically every time it is started up. This procedure is documented in Section 5.2.1.

5.3.4 TCP/IP Configuration

The TCP/IP option to the IRIS 3010 terminal is configured with the hostname file */etc/sys_id*, and the host tables */etc/hosts* and */etc/hosts.equiv*. The hostname

file contains the hostname of the IRIS 3010. The information that goes into the host tables is described in the documentation that accompanies the communications software.

A sample hostname file, *sys_id*, is included in the host distribution. To give the IRIS 3010 a hostname, edit this file to contain the desired name, and then enter the command:

```
iftp w sys_id /etc/sys_id
```

Hostnames must not be more than eight characters.

Next, create the two host table files on the host. Then transfer them to the IRIS 3010 terminal disk. If, for example, the two files on the host are named *hosts* and *hosts.equiv*, enter these commands:

```
iftp w hosts /etc/hosts
iftp w hosts.equiv /etc/hosts.equiv
```

Until these files exist on the IRIS 3010, the terminal emulator always prompts for the Internet addresses of the IRIS and the host. Once the terminal emulator knows the local address, it prompts only for the host address.

The terminal emulator can be configured to initiate a connection with a single host automatically each time it is started up. This procedure is documented in Section 5.2.1.

If the IRIS 3010 disk contains an invalid */etc/hosts* file, it will probably not be able to communicate with the host. To recover from this situation, use one of these strategies:

- Edit */etc/hosts* locally on the IRIS 3010.
- Transfer a valid copy from the host over the serial communications line.
- Use the boot failure recovery procedure described in Section 5.2.5 to start over.

5.3.5 Miscellaneous Unix Configuration

Time Zone

Several UNIX utilities need to know the time zone in order to display the time correctly. The time zone entry is contained in the file */etc/TZ*. The file has three fields:

1. A three-character abbreviation for the time zone.

2. The offset in hours of the time zone from Greenwich Mean Time.
3. A three-character abbreviation for the time zone when in daylight savings time. If daylight savings time is not used, this field should be omitted.

The IRIS 3010 is shipped with the time zone set for Pacific Standard Time:

```
PST8PDT
```

The file *TZ* in the host distribution contains a sample time zone file. After you edit it, to transfer it to the IRIS 3010 disk, execute:

```
iftp w TZ /etc/TZ
```

Automatic Date Setting

By default, the IRIS 3010 prompts the user to check the date each time it is booted. If an executable file */etc/rc.getdate* is found, the user is not prompted, and the result of its execution is used in place of what the user would have typed to the prompt:

```
Enter the correct date (mmddhhmm[yy][.ss]):
```

A sample version of this file named *rc.getdate* is included in the host distribution. It is a shell script that gets the date from an IRIS workstation designated to be the network date server. To use it, edit it and change "host" to be the name of the date server machine. Then execute:

```
iftp w rc.getdate /etc/rc.getdate
irsh chmod a+w /etc/rc.getdate
```

To disable the automatic date setting, remove */etc/rc.getdate*:

```
irsh rm /etc/rc.getdate
```

6. Disk Configuration

This chapter covers these topics:

6.1 Disk Specifications
6.2 IRIS Workstation Disk Configuration
6.3 IRIS Terminal Disk Configuration

The tables in Section 6.1 should be used as reference material for the procedures described later in the chapter. Table 6-1 shows the default division of all the available disks into partitions. This chapter explains how to change partition size.

The disks on the IRIS 3010 and 3020 are called *md*; the disks on the IRIS 3030 are *si*. To find out what brand of disk you have on your workstation, issue the *sgilabel* command:

IRIS 3010/3020:

```
sgilabel md0
```

IRIS 3030:

```
sgilabel si0
```

A first disk is disk zero; a second disk is disk 1. Disks are divided into a number of partitions depending on the brand and size; the partitions are labeled with lower-case letters *a* through *g*. The two formatter/exerciser programs, *mdfex* and *sifex*, correspond to the two types of disk, *md* and *si*. These programs enable you to change the size of the swap partition on your disk.

The commands *mkfs(1M)* and *labelit(1M)* enable you to configure a second disk. The basic steps in this procedure are making file systems, labeling them, and modifying configuration files so that the file systems are mounted and checked by *fsck(1)*.

6.1 Disk Specifications

Table 6-1 shows the disks available on IRIS series 3000 models, the number of megabytes per cylinder on each disk, and the size of the partitions.

NOTE
The disk on your IRIS 3020 may appear to have more space than is shown below, but this space is not supported by Silicon Graphics' software.

On all disks, partition **g** represents the entire disk and can be used in place of, not in addition to, the other partitions. For the IRIS 3030 only, the **f** partition is a combination of the **c** and **d** partitions.

Table 6-2 contains estimates of the disk storage requirements for UNIX software, the IRIS Graphics Library, and options that are available for the IRIS. These estimates fluctuate slightly from one release to the next as new utilities and other software are added to the system.

IRIS Model/ Disk Drive	Mb/Cyl.	Partitions	Cylinders	Mb
IRIS 3010 <i>Tandon</i>	.0348	md0a	123	4.28
		md0b	360	12.53
		md0c	123	4.28
		md0g	606	21.1
IRIS 3020 <i>Vertex</i> <i>(Priam)</i>	.0609	md0a	150	9.14
		md0b	149	9.08
		md0c	670	40.82
		md0g	969	59.04
<i>Maxtor</i>	.0696	md0a	150	10.4
		md0b	149	10.38
		md0c	586	40.8
		md0g	885	61.58
IRIS 3030	.163	si0a	56	9.18
		si0b	103	16.88
		si0c	200	32.77
		si0d	446	73.07
		si0f	646	105.84
		si0g	805	131.89

Table 6-1: Disk Partition Sizes

Software Category	Space Required (Megabytes)
Total, minimum system:	16.00
On-line manual pages (<i>/usr/man/*</i>)	2.50
Demonstration programs (<i>/usr/people/demos/*</i>)	4.25
Games (<i>/usr/games/*</i>)	1.25
Compiler Options:	
FORTRAN compiler option	1.00
Pascal compiler option	0.50
Other Options:	
IEEE 488 communications option	0.50
TCP/IP communications option	1.00
IBM option	0.75
Emacs text editor	2.00
Total with all options installed:	30.00

Table 6-2: Disk Storage Requirements

6.2 IRIS Workstation Disk Configuration

The three subsections below describe these procedures:

- Changing the swap partition size on an IRIS 3020
- Changing the swap partition size on an IRIS 3030
- Configuring a second disk on an IRIS 3020/3030

If you have an IRIS 3010, see Section 6.3.

On an IRIS 3020 or 3030, disk zero is divided into three partitions: the root partition, which contains the root file system; the swap partition, which is used by UNIX for paging; and the user partition, which contains the */usr* file system. On an IRIS 3020, the root partition is device **md0a**, the swap partition is **md0b**, and the user partition is **md0c**. On an IRIS 3030 the root partition is device **si0a**, the swap partition is **si0b**, and the user partition is **si0f**.

The size of the swap partition determines the maximum amount of virtual memory that can be used. At any instant in time, the sum of the sizes of all the processes known to UNIX cannot exceed the size of the swap partition.

If the simultaneous processes running on your workstation require more space on the swap partition, you can make it larger at the expense of some of the free space in the user file system. You can increase the space available for user files by reducing the swap partition size, but this reduces the total of the process sizes that can exist simultaneously on the system.

To find out the size of the swap partition on your workstation, issue the command:

IRIS 3020:

```
sgilabel md0
```

IRIS 3030:

```
sgilabel si0
```

The *sgilabel* command prints out the disk drive label. The label includes the brand of disk drive and the size of each partition. Table 6-1 also gives the size of each partition for the IRIS 3020/3030.

6.2.1 Changing the Swap Partition Size on an IRIS 3020

To change the size of the swap partition on an IRIS 3020, follow the steps below. For simplicity, the procedure describes a specific, although arbitrarily chosen, change: an increase of 100 cylinders.

1. Log in as root:

```
login: root.
```

2. Back up the */usr* file system and make sure you have a bootable backup tape, in case the root file system is unintentionally corrupted (see Section 4.4).
3. Reboot the workstation:

```
reboot
```

4. Boot the disk formatter program:

```
b stand/mdfex
```

CAUTION
You must follow all instructions exactly when you use the <i>fex</i> disk formatter programs. Failure to follow instructions could result in serious damage to your file system.

5. When the *mdfex* prompt appears, enter the letter *s*, then *l*, to display the drive label on your screen.

6. Press **RETURN**. The entry below appears.

```
a: base; (1)      size; (150)
```

Press **RETURN** to accept this entry. This entry indicates that partition **md0a** begins at cylinder 1 and has a size of 150 cylinders.

Partition **md0a** is the root partition and does not change.

7. To add 100 cylinders to the swap partition (**md0b**), 100 cylinders must be subtracted from the user partition (**md0c**). The base, or beginning, of partition c changes by 100 cylinders. Type the revised entries after the numbers in parentheses:

Vertex:

```
b: base; (151)    size; (149)249
c: base; (300)400 size; (586)486
```

Maxtor:

```
b: base; (151)    size; (149)249
c: base; (300)400 size; (670)570
```

8. Continue to press **RETURN** until the workstation displays the **set?** prompt.
9. Enter **q** twice to quit both the **set** function and **mdfex**. The workstation displays the **mdfex** prompt and the message:

```
Quit
Label on drive 0 needs updating... do it?
```

10. Confirm your intention to update the drive label by entering **y**. The workstation displays the message:

```
--confirm quit with 'y':
```

11. Type **y** to confirm your intention to quit **mdfex** and return to the PROM monitor. The PROM monitor prompt appears.
12. Boot the system in single-user mode:

```
b
```

13. Make the new, smaller user file system:

```
mkfs /dev/rmd0c
```

14. Name the file system and give it a volume name with the *labelit* command:

```
labelit /dev/rmd0c usr sgi
```

15. Mount the */usr* file system:

```
mount /dev/md0c /usr
```

16. Restore the */usr* files from tape:

```
cd /usr  
tar xv      OR      cpio -ihmudlv
```

17. Unmount */usr*.

```
umount /dev/md0c
```

18. Check the file system:

```
fscck /dev/rmd0c
```

19. Make a new bootable backup tape (see Section 4.4.1).

20. Begin multi-user mode:

```
multi
```

6.2.2 Changing the Swap Partition Size on an IRIS 3030

To change the size of the swap partition on an IRIS 3030, follow the steps below. For simplicity, the procedure describes a specific, although arbitrarily chosen, change: an increase of 100 cylinders.

1. Log in as root:

```
login: root
```

2. Back up the */usr* file system on partition **si0f** and make sure you have a bootable backup tape, in case the root file system is unintentionally corrupted (see Section 4.4).
3. Reboot the workstation:

```
reboot
```

4. Boot the disk formatter program:

```
b stand/sifex
```

CAUTION

<p>You must follow all instructions exactly when you use the <i>flex</i> disk formatter programs. Failure to follow instructions could result in serious damage to your file system.</p>
--

5. When the *sifex* prompt appears, enter the letter **s**, then **1**, to display the drive label on your screen.
6. Press **RETURN**. The entry below appears.

```
a: base (1)      size (56)
```

This entry indicates that partition **si0a** begins at cylinder 1 and has a size of 56 cylinders.

Partition **si0a** is the root partition and does not change. Press **RETURN**.

7. To add 100 cylinders to the swap partition, subtract 100 cylinders from the */usr* partition (partition **f**). Type the revised entries after the numbers in parentheses:

```
b: base; (57)      size; (103)203
c: base; (160)     size; (200)
b: base; (360)     size; (446)
c: base; (160)260 size; (646)546
```

8. Continue to press **RETURN** until the workstation displays the **set?** prompt.

9. Enter `q` twice to quit both the `set` function and `sifex`. The workstation displays the `sifex` prompt and the message:

```
Quit
Label on drive 0 needs updating... do it?
```

10. Confirm your intention to update the drive label by entering `y`. The workstation displays the message:

```
--confirm quit with 'y':
```

11. Type `y` to confirm your intention to quit `sifex` and return to the PROM monitor. The PROM monitor prompt appears.
12. Boot the system in single-user mode:

```
b
```

13. Make the new, smaller user file system:

```
mkfs /dev/rsi0f
```

14. Name the file system and give it a volume name with the `labelit` command:

```
labelit /dev/rsi0f usr sgi
```

15. Mount the `/usr` file system:

```
mount /dev/si0f /usr
```

16. Restore the `/usr` files from tape:

```
cd /usr
tar xv    Or    cpio -ihmud1v
```

17. Unmount `/usr`.

```
umount /dev/si0f
```

18. Check the file system:

```
fsck /dev/rsi0f
```

19. Make a new bootable backup tape (see Section 4.4.1).
20. Begin multi-user mode:

```
multi
```

6.2.3 Configuring a Second Disk on an IRIS Workstation

On workstations with two disk drives, the second disk is configured by the user. Users may choose several strategies for configuring the second disk, depending on individual needs. Table 6-3 summarizes the configuration options.

Number of File Systems	System	Partition	File System Usage
One	IRIS 3020	md1g	user files
	IRIS 3030	si1g	user files
Three	IRIS 3020	md1a	copy of key files and user files
		md1b	<i>/tmp</i>
md1c		user files	
IRIS 3030	si1a	copy of key files and user files	
	si1b	<i>/tmp</i>	
	si1f	user files	

Table 6-3: Two Ways to Configure a Second Disk

If you want as much space as possible for an extra file system, it is best to configure the second disk as one large file system. On the other hand, if you would like to keep a copy of the key UNIX files on your second disk, it is better to configure the second disk as several file systems.

NOTE: You are advised to put copies of the key UNIX files on the second disk to make it easier to recover from a crash. You are able to boot from a second disk only if it is configured as three file systems.

The following two sections give the procedures for configuring a second disk.

Configuring a Second Disk as Three File Systems

To configure a second disk as several file systems, follow these steps:

1. Become the super-user and reboot the system in single-user mode:

```
su
reboot
b
```

2. Create directories in / (root) for the new file systems to be mounted on the disk partitions shown in the table below. The names must be six characters or less. The procedure below uses the names *n1* and *n2*.

System	Partition	File System Name
IRIS 3020	md1a	<i>n1</i>
	md1b	<i>tmp</i>
	md1c	<i>n2</i>
IRIS 3030	si1a	<i>n1</i>
	si1b	<i>tmp</i>
	si1f	<i>n2</i>

Note that */tmp* already exists and will be mounted on the *b* partition.

```
cd /
mkdir n1 n2
```

3. Make the new file systems:

IRIS 3020:

```
mkfs /dev/rmd1a
mkfs /dev/rmd1b
mkfs /dev/rmd1c
```

IRIS 3030:

```
mkfs /dev/rsi1a
mkfs /dev/rsi1b
mkfs /dev/rsi1f
```

- Name the file systems and give them volume names with the *labelit* command:

IRIS 3020:

```
labelit /dev/rmd1a n1 sgi
labelit /dev/rmd1b tmp sgi
labelit /dev/rmd1c n2 sgi
```

IRIS 3030:

```
labelit /dev/rsi1a n1 sgi
labelit /dev/rsi1b tmp sgi
labelit /dev/rsi1f n2 sgi
```

- Edit */etc/rc.fs* to contain the following lines:

IRIS 3020:

```
mount /dev/md1a /n1 && echo ": /n1\c":
mount /dev/md1b /tmp && echo ": /tmp\c":
mount /dev/md1c /n2 && echo ": /n2\c":
```

IRIS 3030:

```
mount /dev/si1a /n1 && echo ": /n1\c":
mount /dev/si1b /tmp && echo ": /tmp\c":
mount /dev/si1f /n2 && echo ": /n2\c":
```

These lines automatically mount the file systems on the second disk each time multi-user mode is started.

- /etc/checklist* is the default list of file systems to be checked by *fsck*. Edit */etc/checklist* and add the new file systems:

IRIS 3020:

```
/dev/rmd1a
/dev/rmd1b
/dev/rmd1c
```

IRIS 3030:

```
/dev/rsi1a
/dev/rsi1b
/dev/rsi1f
```

7. Enter the *sync* command to flush the new information to the disk.

```
sync
```

Configuring a Second Disk as One File System

To configure a second disk as one large file system, follow these steps:

1. Become the super-user and reboot the system into single-user mode:

```
su
reboot
b
```

2. Create a directory in / (root) for the file system to be mounted on **md1g** (for an IRIS 3020) or **si1g** (for an IRIS 3030). Its name must be six characters or less. In this procedure, the name *n* is used.

```
cd /
mkdir n
```

3. Make the user file system:

IRIS 3020:

```
mkfs /dev/rmd1g
```

IRIS 3030:

```
mkfs /dev/rsi1g
```

4. Edit */etc/rc.fs* to contain the following lines:

IRIS 3020:

```
mount /dev/md1g /n && echo ": /n\c":
```

IRIS 3030:

```
mount /dev/si1g /n && echo ": /n\c":
```

This line automatically mounts the new file system each time multi-user mode is started.

5. `/etc/checklist` is the default list of file systems to be checked by `fsck`. Edit `/etc/checklist` and add the new file systems:

IRIS 3020:

```
/dev/rmd1g
```

IRIS 3030:

```
/dev/rsi1g
```

6. Name the file system and give it a volume name with the `labelit` command:

IRIS 3020:

```
labelit /dev/rmd1g n sgi
```

IRIS 3030:

```
labelit /dev/rsi1g n sgi
```

7. Enter the `sync` command to flush the new information to the disk.

```
sync
```

6.3 IRIS Terminal Disk Configuration

This subsection describes three procedures:

- Changing the swap partition size on an IRIS 3010
- Configuring a second disk on an IRIS 3010
- Creating and using a backup floppy set

6.3.1 Changing the Swap Partition Size on an IRIS Terminal

The disk on the IRIS 3010 is divided into three partitions: the root partition, which contains the file system; the swap partition, which is used by the kernel for paging; and the backup partition, which is a copy of the root partition. The root partition is device `md0a`, the swap partition is `md0b`, and the backup partition is `md0c`.

The size of the swap partition determines the maximum amount of virtual memory that can be used. At any instant in time, the sum of the sizes of all the processes known to UNIX cannot exceed the size of the swap partition. To

find out the size of the swap partition on your terminal, escape from the *wsiris* terminal emulator with `~!`. If you are using the IBM terminal emulator, *t3279*, escape to the shell by following the instructions in *IBM Terminal Emulation*, Section 3.4, Escaping to the Shell.

In response to the shell prompt, enter:

```
sgilabel md0
```

The `sgilabel` command prints out the disk drive label. The label includes the brand of disk drive and the size of each partition. Table 6-1 also gives the size of each partition for the IRIS 3010.

If the simultaneous processes running on your IRIS require more space on the swap partition, you can make it larger. If your IRIS 3010 uses less virtual memory than is currently available, you can create more storage room on the disk by reducing the size of the swap partition.

The backup partition is an exact copy of the root partition. To keep the two partitions the same size, divide the space reclaimed from the swap partition equally between the root and backup partitions.

To change the sizes of the partitions, follow these steps:

1. Configure the terminal for PROM monitor boot (i.e., set configuration switches 1 through 5 to **Closed**).
2. Boot the system:

```
b
```

3. Escape from the terminal emulator by entering:

```
~!
```

4. Become the super-user:

```
su
```

5. Ensure that the backup partition is up to date:

```
rootcopy ok
```

6. Mount the backup partition:

```
mount /dev/md0c /mnt
```

7. Edit `/mnt/etc/inittab` to suppress automatic recovery when booting from the backup partition. Find this line:

```
rr::sysinit:/etc/rootcopy >/dev/console 2>&1
```

Put an `x` between the two colons:

```
rr:x:sysinit:/etc/rootcopy >/dev/console 2>&1
```

8. Verify that the backup partition is bootable:

```
reboot
b md0c:recover
```

9. Reboot the terminal into the disk formatter/exerciser, `mdfex`:

```
b stand/mdfex
```

CAUTION

You must follow all instructions exactly when you use the *fex* disk formatter programs. Failure to follow instructions could result in serious damage to your file system.

The system displays the *mdfex* opening screen and prompt.

10. Type `s`. The system displays the *set* prompt:

```
Set ?
```

Respond with the letter `1` to set the disk's label. Press **RETURN**. The system displays the size of the root partition (partition `a`):

```
a: base; ( 1) size; ( 123)
```

The numbers in parentheses are the current values. The size is expressed in cylinders; each cylinder is 34Kb. Before pressing **RETURN**, enter the new size for partition `a`. For example, if you are reducing the swap partition by 2 Mb, you will increase the size of partition `a` by 1 Mb, which represents 30 cylinders. Respond to the size prompt with the new figure:

```
a: base; ( 1) size; ( 123) 153
```

Press **RETURN** to enter the new size and to display the base cylinder for partition **b**:

```
b: base; ( 124)
```

Change the base for partition **b** to the cylinder immediately following the last cylinder in partition **a**. For example, if you have added 30 cylinders to partition **a**, the new base for partition **b** is 154.

```
b: base; ( 124) 154
```

Press **RETURN** to enter the new base and to display the size of partition **b**. Enter the new size:

```
b: base; ( 124) 154 size; ( 360) 300
```

Press **RETURN** several times, until this prompt appears:

```
Set ?
```

Enter **q** twice to quit both the **set** function and **mdfex**. The workstation displays the **mdfex** prompt and the message:

```
Quit
Label on drive 0 needs updating... do it?
```

Confirm your intention to update the drive label by entering **y**. The workstation displays the message:

```
--confirm quit with 'y':
```

Type **y** to confirm your intention to quit **mdfex** and return to the PROM monitor. The PROM monitor prompt appears.

11. Boot the *recover* program, create a filesystem on the newly sized root partition, and copy the backup partition to it:

```
md0c:recover
mkfs /dev/rmd0a
labelit /dev/md0a root sgi
cd /
mount /dev/md0a /mnt
sh # you must use the Bourne shell for the tar command
tar cfB - [!m][!n][!t]* .??* | ( cd /mnt ; tar xvfpB - )
exit
```

The copy takes about eight minutes.

12. Reboot and change the base and size of the backup partition. Follow the procedure described in step 10, but change the values for partition **c**. Increase the size of partition **c** by the same amount as partition **a**. For the example in step 10, you would enter these figures:

```
c: base; ( 484) 424 size; ( 123) 153
```

Exit from the *set* routine, update the disk label, and exit from *mdfex*.

13. Boot from the root partition and undo the change to */etc/inittab* made in step 7.
14. Copy the root partition to the backup partition:

```
rootcopy ok
```

6.3.2 Creating and Using a Backup Floppy Set

The IRIS 3010 includes several tools for making a set of floppies for backup and using them to build a file system on the hard disk. The two procedures below describe how to create the floppies and how to build the disk from them.

The floppy set is to be used only in extraordinary circumstances. There is no need for the typical user to create the floppy set, as the backup partition on the disk can be used if there is problem with the file system.

Creating the Floppy Set

1. Become the super-user on the IRIS 3010.
2. Format eight floppies using *flp*:

```
flp format
```

3. Put *mdfex* on one of the formatted floppies:

```
mkfs -G 1 1 /dev/rfloppy 1280
mount /dev/floppy /mnt
cp /stand/mdfex /mnt/mdfex
umount /dev/floppy
```

4. Insert a new floppy and issue the *mkbootflp* command to make it a bootable UNIX floppy:

```
mkbootflp
```

5. Create a set of *tar* floppies containing the complete system:

```
cd /
tar cvfb /dev/rfloppy 80
```

Insert a new (formatted) floppy each time the system prompts for a new tape.

Using the Floppy Set

1. Turn on the power or reboot the terminal to enter the PROM monitor.
2. Boot *mdfex* off the *mdfex* floppy:

```
b mf:mdfex
```

CAUTION

<p>You must follow all instructions exactly when you use the <i>fex</i> disk formatter programs. Failure to follow instructions could result in serious damage to your file system.</p>

3. Format the hard disk by entering *f*. The system displays a warning and the message:

```
Type 'go <return>' to start ...
```

4. To set the label, enter the letters *s*, then *1*. The system prompts for the name, serial number, and type of drive. Type 3010 after the *Name* prompt. Press **RETURN** to continue to the next item.

The serial number is shown on the hard disk drive. Enter the type of drive after the *#?* prompt. Type *21*, the number corresponding to Tandon TM262. Continue to press the **RETURN** key until the *set?* prompt appears.

5. Enter *q* twice to quit both the *set* function and *mdfex*. The workstation displays the *mdfex* prompt and the message:

```
Quit
Label on drive 0 needs updating... do it?
```

Confirm your intention to update the drive label by entering *y*. The workstation displays the message: `--confirm quit with 'y':`

7. Type *y* to confirm your intention to quit *mdfex* and return to the PROM monitor. The PROM monitor prompt appears.
8. Boot *defaultboot* from the bootable UNIX floppy:

```
b mf:
```

The system prompts for run level. Enter *s*. Ignore the messages:

```
INIT: cannot open /etc/inittab errno: 2
INIT: execlp of /bin/sh failed; errno = 2
```

9. Transfer the contents of the floppy to partition *a* of the hard disk:

```
/etc/flpuxfer md0a
```

10. Shut down the system and boot *defaultboot* off the hard disk:

```
/etc/reboot
b
```

The system prompts for run level. Answer *s*.

11. Move a few files:

```
/bin/mv /bin/tar /bin/sh /etc/init /tmp
```

12. Insert the first of the six *tar* floppies. Read the complete system onto the disk:

```
cd /
/tmp/tar xvfb /dev/rfloppy 80
```

Insert the subsequent floppies as the system prompts for additional tapes. Insert the floppies in order; otherwise *tar* fails.

13. Copy the system to the backup partition:

```
/etc/rootcopy ok
```

14. Reboot the system:

```
/etc/reboot
```


8. Video Options

The IRIS workstation supports seven video options:

- 60 Hz non-interlaced (fifteen-inch or nineteen-inch)
- 33 Hz interlaced monitor
- RS-170A video (genlockable and non-genlockable)
- European video standard (genlockable and non-genlockable)

NOTE: Early IRIS series 2000 models with the Turbo upgrade may not have the same monitors described here.

Each IRIS workstation contains drivers for two video options, one *primary* video option and one *secondary* video option. Table 7-1 lists the possible combinations of primary and secondary video options.

Primary Video	Secondary Video
60 Hz Non-interlaced	33 Hz Interlaced
60 Hz Non-interlaced	RS-170A
60 Hz Non-interlaced	European Video Standard
33 Hz Interlaced	60 Hz Non-interlaced
33 Hz Interlaced	RS-170A
33 Hz Interlaced	European Video Standard

Table 8-1: Primary and Secondary Video Options

NOTE: The RS-170A and European video standard monitors are not available from Silicon Graphics.

The configuration for the primary and secondary video drivers is specified when an IRIS workstation is ordered. See below for the procedures for using the IRIS Graphics Library with different video options.

8.1 Supported Options

8.1.1 60 Hz Non-interlaced Monitor

The 60 Hz non-interlaced high-resolution monitor has a visible resolution of 1024 pixels by 768 lines. These monitors are available in two sizes: fifteen-inch and nineteen-inch. For more information on these monitors, see Chapter 2.

8.1.2 33 Hz Interlaced Monitor

The 33 Hz interlaced high-resolution monitor has a visible resolution of 1024 pixels by 768 lines. For more information, see Chapter 2.

8.1.3 RS-170A

The RS-170A option has a visible resolution of 636 pixels by 485 lines and a frame rate of 30 Hz.

The IRIS workstation provides red, green, blue, and sync video outputs (RGBS). The RGB outputs from the IRIS are noncomposite 0.7 V p-p into 75 ohms, positive bright. The sync output is TTL, low active, and is capable of driving a 75-ohm load. It meets RS-170A timing with one exception: the oscillator that determines horizontal scanning frequency stability in the IRIS differs from the frequency stability of the RS-170A color television timing standard ($\pm 0.005\%$ versus $\pm 0.000279\%$). This stability is adequate for most applications but not for broadcasting. For broadcast applications, use either the genlockable RS-170A option or a timebase corrector.

The IRIS RGBS output can be encoded into an NTSC signal with a stand-alone NTSC color encoder or an NTSC color encoder connected to a sync generator genlocked to the IRIS sync output. This NTSC signal can be connected to video tape recorders or an NTSC-type monitor.

8.1.4 European Video Standard

The European video standard has a visible resolution of 780 pixels by 575 lines, a frame rate of 25 Hz, and a total vertical resolution of 625 lines.

The IRIS workstation provides red, green, blue, and sync video outputs (RGBS). The RGB outputs from the IRIS are noncomposite 0.7 V p-p into 75 ohms, positive bright. The sync output is TTL, low active, and is capable of driving a 75-ohm load. It meets European low-resolution video timing, with one exception: the oscillator that determines horizontal scanning frequency stability in the IRIS differs from the frequency stability of the European color television timing standards ($\pm 0.005\%$ versus $\pm 0.0001\%$). This stability option is adequate for most applications but not for broadcasting. For broadcast applications, use the genlockable European low-resolution video or a timebase corrector.

The IRIS RGBS output can be encoded into a PAL signal by a stand-alone PAL color encoder or a PAL color encoder connected to a sync generator genlocked to the IRIS sync output. This PAL signal can be connected to a video tape recorder or a PAL-type monitor.

8.1.5 Genlockable RS-170A or European Video Standard

The genlockable RS-170A or European video option synchronizes the IRIS RGB outputs with those of an external sync master.

The IRIS RGBS output can be encoded into an NTSC or PAL signal with a stand-alone color encoder or a PAL color encoder connected to a genlockable sync generator which is also genlocked to the sync master.

The genlockable RS-170A or European video option provides operation with commercially available video tape recorders and mixing/keying equipment (when color encoded with external equipment). If a genlockable video option is ordered, two BNC connectors are provided with the IRIS workstation system. These connectors are located on the auxiliary I/O panel and are labeled **Genlock Input** (see Figure 7-1). If the IRIS is not used in a daisy chain, a 75-ohm terminator is required on one of the two BNC connectors.

Operation of the genlockable RS-170A or genlockable European Video Standard options are identical. The IRIS must be operating in its supported low-resolution mode (RS-170A or European Video Standard) to genlock. Genlock provides for vertical and horizontal alignment of the IRIS RGB and Sync outputs to an external composite sync reference input presented to the **Genlock Input** on the auxiliary I/O panel (see Figure 7-1). This input is a loop-through connection requiring external 75-ohm termination.

The genlocking process takes approximately two seconds. When a genlock reference input is connected, the genlocking process is initiated automatically. If no composite sync reference input is connected, the appropriate output (RS-170A or European Video Standard) runs referenced to an internal timing source. Genlocking takes place whenever the IRIS is reset, the genlock reference input is disconnected and reconnected, or an "out-of-lock" condition is detected. An "out-of-lock" condition typically results from a discontinuity on the genlock reference input, such as a power interruption or a source switch.

When genlocked, the IRIS output may be adjusted horizontally to compensate for external equipment delays. To do this, use the horizontal phase adjust potentiometer, labeled **HPhase Adjust**, which is reached through an access hole on the auxiliary I/O panel (see Figure 7-1). This potentiometer provides an adjustment range of ± 1.5 microseconds on the IRIS RGBS output with respect to the genlock reference input.

8.2 Booting with the Secondary Video Driver

Configuration switch 7 on the cabinet I/O panel controls which video driver is initialized when the workstation is booted in PROM monitor boot configuration.

To boot using the secondary display:

1. Set switch 7 to **Open** to initialize the video hardware to drive the secondary display. (Setting switch 7 to **Closed** causes the PROM monitor to initialize the hardware to drive the primary display.)
2. Set the **Power** switch on the cabinet front panel to **On**.

The PROM monitor prompt appears on the screen:

```
iris>
```

If the prompt does not appear, press the **Reset** button.

3. Boot UNIX. To boot from the default device, type:

```
b
```

See Table 3-2 for a list of PROM monitor commands.

Instead of booting using the PROM monitor, the secondary display can be selected using the `setmonitor` command (see below).

NOTE: Although it can be used as the secondary display, the RS-170A monitor option cannot be used as the boot display.

8.3 Using the Graphics Library with the Video Options

The `setmonitor` command in the IRIS Graphics Library specifies the video compatibility of the IRIS workstation display output. To change the active video driver through the Graphics Library, issue the command:

```
setmonitor(type)
```

Specify one of these types: HZ60 (60 Hz non-interlaced monitor), HZ30 (33 Hz interlaced monitor), NTSC (RS-170A monitor), or PAL (European video standard).

These monitor types are defined from the include file `set.h`. The command `getmonitor`, which returns the current display controller mode, returns one of these types. (For more information, see the *IRIS User's Guide*.)

With a low-resolution monitor, you should also reduce the size of the viewport, using the `viewport` command:

```
viewport(left,right,bottom,top)
```

To map the graphics output to the entire screen, specify the exact available pixel range:

NTSC monitor

```
viewport(0,635,0,484)
```

PAL (European video standard) monitor

```
viewport(0,779,0,574)
```

The `ortho` commands control the mapping of graphics coordinates to screen coordinates. Used in conjunction with the `viewport` command, the `ortho2` command lets you program in screen coordinates on low-resolution monitors. To establish a one-to-one mapping of graphics coordinates to screen coordinates, use parallel ranges in the two commands:

NTSC monitor

```
ortho2(-0.5,635.5,-0.5,484.5)
```

PAL monitor

```
ortho2(-0.5,779.5,-0.5,574.5)
```

If the graphics programs are run on a low-resolution monitor with the default viewport specifications, only the lower left portion of the image will be visible on the screen. UNIX will not work in NTSC or PAL mode because it requires more pixel area. To run UNIX, switch to a high-resolution monitor with `setmonitor`.

If you design a program to run on both high-resolution and low-resolution monitors, keep in mind that fonts have a fixed pixel width. This means that text on one monitor will appear different on another monitor.

9. Demonstration Programs and Gifts

The IRIS workstation is shipped with a set of demonstration programs, programming examples, and other miscellaneous tools. This chapter tells how to operate the demonstration programs. The demonstration programs that run under the window manager are described in detail in Section 9.1.3.

9.1 Demonstration Programs

Demonstration programs can be run any time after the system is booted. They are located in the directories `/usr/people/demos` and `/usr/people/mexdemos`. The directory `mexdemos` contains demonstration programs that run with the window manager, while the programs in `demos` do not work correctly if the window manager is running.

9.1.1 Running Demonstration Programs

1. Boot the workstation in multi-user mode.
2. Log in to the workstation as `demos` or `mexdemos`.

```
IRIS login: demos
IRIS login: mexdemos
```

3. If you are running the window manager demonstration programs (`mexdemos`), run `startup`:

```
startup
```

The textport is labeled with a banner, which indicates it is the console window.

4. Display a listing of the available programs with the `ls` command:

```
ls
```

5. Enter the name of the program you want to run. For example, to use the flight simulator, enter the word *flight*, followed by **RETURN**.

`flight`

Some of the demonstration programs exit automatically when they finish running. To exit the open-ended demonstration programs or to interrupt a program, press **ESC**. To clear the screen of graphics "leftovers" if you are not running the window manager demonstrations, enter the command `gclear(1G)`.

More information on the demonstration programs is included in Section 1 of the on-line *UNIX Programmer's Manual*. Manual pages for the window manager commands are labeled 1W; other demos are labeled 1D.

9.1.2 Using the Mouse with the Window Manager

There are three buttons on the mouse. They are named LEFTMOUSE, MIDDLEMOUSE, and RIGHTMOUSE. Almost all the window manager functions can be controlled by the mouse x-y valuator and the RIGHTMOUSE.

Before continuing, make sure you understand how to use the mouse to communicate with windows. This information is located in the *IRIS User's Guide, Version 2.1* in the section called "The Default User Interface to the Window Manager". Pay special attention to the subsection called "Creating new textports and selecting textports for input", which tells how to select a window for interaction ("attach" to a window), and exit a window.

9.1.3 Window Manager Demonstration Programs

The demonstration programs described in this section fall into three main groups.

- *2D demonstration programs* illustrate how the window manager works and provide instruction on the color-handling capabilities of the IRIS.
- *3D object manipulation* demonstration programs include *arch*, *flow*, *heme*, *jet*, and *shuttle*. These programs move 3D objects around. The IRIS can perform simultaneous manipulation of different 3D objects in independent windows.
- *Shaded object* and *curved surface* demonstration programs illustrate shading and surface treatment on the IRIS.

Window manager demonstration programs require a large number of bitplanes. *flow* requires 12; *arch*, *heme*, *jet*, *shuttle*, and surface editors require 24; *zshade* requires 32.

2D Window Manager Demonstration Programs

A set of 2D programs is included in the collection of demonstration programs. These tools include the following programs: *bckgrnd*, *showmap*, *demomakemap*, *showramp*, *gamma*, *cedit*, *interp*, and *mag*. These tools can be classified into three categories: backgrounds, color manipulation programs, and pixel access programs. To start any of these tools, type the name of the tool in the console window.

Backgrounds

bckgrnd is a non-interactive, textured “wallpaper” with art deco designs. A particular background fills all available space which is not covered with another window.

Color Manipulation Programs

showmap is a window filled with up to 1024 little squares. It displays the current colors in the color map, counting from the color 0 at the lower left corner up to the largest color index at the top right. (For more information on color maps, see the *IRIS User's Guide* and the *makemap(1W)* command in the *UNIX Programmer's Manual*.)

demomakemap makes and remakes the color map for the demonstration programs. Before other graphics programs are executed, run *demomakemap* to build the default color map for the demonstration package. To see the effect of changing the color map using *demomakemap*, run *showmap* in a separate window. (At any time during the demonstration session, the original color map can be restored by executing *demomakemap*.)

demomakemap also works in conjunction with the *gamma* program. *demomakemap* reads the value from the file *.gamma* in your home directory and uses it to calculate its shade scales (see below).

showramp is a window with three Gouraud-shaded rectangles. The rectangles are smoothly colored using physically and chromatically adjacent color indices that form a color ramp in the color map.

The *gamma* program uses a technique called gamma-correction to change the “smoothness” of the transitions between the color intensities in a color map. The physical characteristics of the CRT phosphor, the ambient lighting, and the vision and preference of the viewer affect the perception of the color map. A strictly linear color ramp often appears too dark, especially towards the middle of the ramp. *gamma* allows the increments from color to color to be exponential, and thus, the middle of the color ramp can be brighter or darker than in a linear color ramp.

The *gamma* command accepts a single floating-point argument. A value of 1.0 creates a strictly linear color ramp. Values above 1.0 increase the brightness in the middle of a color ramp; values below 1.0 decrease the brightness. Values between 1.0 and 3.0 are typical.

The *gamma* program does not immediately affect the color map. It stores its argument in a file called *.gamma* in your home directory. Whenever a color ramp is made, such as in the program *demomakemap*, this file is referenced for a value for *gamma*. The *gamma* program itself does not open a window.

Below is an exercise in running the previously described programs. To run these programs, use the following procedure:

1. Bring up the windows *showmap* and *showramp*.
2. Set the gamma correction value and use it to change the color map:

```
gamma 0.5
demomakemap
```

Wait a moment for the colors to change. Observe the effects on the object and the brightness and smoothness of the color ramps in the *showmap* and *showramp* windows. (*demomakemap* does not open a window.)

3. Repeat the experiment with a different *gamma* value:

```
gamma 3.0
demomakemap
```

Remember that the value you select for *gamma* is saved until you change it.

cedit is an interactive color editor in a graphics window. To use *cedit*:

1. Bring up the *cedit* window and select it for input.
2. To select a color index for editing, move the cursor to a pixel outside the *cedit* window itself.
3. To choose the color index of the pixel, press the LEFTMOUSE button and release it.

The selected color index is painted on the "sample chip" on the right of the *cedit* window, and the "color sliders" are set to the red, green, and blue values of that color.

4. To change the intensity of the color value of the sample chip, press the LEFTMOUSE button down while the cursor is located on one of the three color sliders and move the sliders up and down.

Try picking a color from *showmap* and editing it with *cedit*.

interp displays and alters a region of the color map. Use *interp* to make color ramps for depth-cued or smooth-shaded objects. *interp* makes a color ramp of any size between any two color indices, interpolating the red, green and blue intensities of all indices between the two prescribed extremes. Two squares show the extreme colors of a color ramp and a long rectangle shows the intermediate color range.

The following sequence describes how to use *interp* with *showmap* and *cedit* to create a user-defined color ramp.

1. Attach to the *cedit* window. Use *cedit* to edit any two color indices displayed by *showmap*.
2. Detach from *cedit* and attach to *interp*. Point to one of the extreme colors from *showmap* that you have edited and press the LEFTMOUSE button.
4. Point to the other extreme color and again press the LEFTMOUSE button. The colors appear as two large squares in the *interp* window. (If you make a mistake while choosing extreme colors, for example, the cursor is over an incorrect color index, then continue picking colors with the LEFTMOUSE button until both extreme colors are visible in the *interp* window.)
4. Press MIDDLEMOUSE; the system interpolates between the colors you have chosen and places the new values into the color map.

Pixel Access Program

mag is an interactive program that magnifies or enlarges a pixel region on the screen. To use *mag*:

1. Attach to the *mag* window.
2. Place the cursor on some image on the screen, then press and release the LEFTMOUSE button. A magnified image of the area around the cursor fills the *mag* window. Pressing the MIDDLEMOUSE button toggles the outlining of the magnified pixels.

3D Object Manipulation Demonstration Programs

This section describes the fast 3D demonstration programs. They illustrate 3D transformations, polygon fill, depth-cuing, hidden surface removal, and overlapping windows. These demonstration programs are double-buffered for smooth motion.

Arch, Flow, Heme, Jet, and Shuttle

arch is a simulated architectural model in which the eye perceives a 3D environment. One building (the white one) is always displayed. Several other buildings can be displayed in wireframe or as solid objects.

Once you select a window for input, you interact with the program through both the keyboard and the mouse. Each mouse button or combination of mouse buttons causes a different motion of the eye within the environment. The **F** key toggles the other buildings; **S** toggles the polygon fill on the buildings. The **M** key starts a constant motion that continues even after the window is no longer selected for input (detached).

flow is an array of nearly 20,000 data points in a frame. The data was calculated at NASA-Ames Research Center on the Illiac-IV computer. *flow* presents some solutions to the Navier-Stokes equations for flow in a turbulent fluid. LEFTMOUSE controls viewer altitude and azimuth. MIDDLEMOUSE controls distance and twist (rotation on the z-axis). RIGHTMOUSE starts a sequence of small rotations. The sequence continues even if the window is detached.

heme is a depth-cued molecular model comprised of 1022 points and 549 vectors. It is controlled by a popup menu. To see the menu, hold down MIDDLEMOUSE. To select or de-select an entry, position the cursor on the entry and press and release LEFTMOUSE. Two menus appear. The left column of menu entries independently turns off and on each of the five parts of the molecule. The right column of menu entries controls the transformations of the molecule. The menu entries in the right column are graphical icons.

A LEFTMOUSE press changes the color of a menu entry and starts that motion. Another LEFTMOUSE press restores the initial state. Simultaneously specified opposite motions cancel each other out, leaving no motion in the specified direction. The **ENTER** or **M** key causes current motion to continue even after the program's window is detached.

The bungalow icon (shaped like a house) resets the program back to the center of the window (especially helpful if the molecule has been dragged out of sight). The stop sign icon exits the program and cancels the window.

jet is a wireframe model of a fighter plane. The image includes 660 vectors and can be displayed with constant intensity or depth-cuing. Pressing one of the three mouse buttons enables rotation about one of the three axes. The amount of rotation is determined by the left-to-right motion of the mouse. Other combinations of mouse buttons allow *x-y* translation controlled by the mouse and scaling at fixed positive and negative rates. The **ENTER** button enables and disables depth-cuing. Press **M** to start or stop a continuous motion that continues after the window is detached.

The commands to manipulate the *shuttle* and control depth-cuing are the same as those for *jet*. Any alphanumeric key starts or stops the doors flopping open and closed. The motion does not continue after the program is detached.

Shaded Object and Curved Surface Demonstration Programs

The IRIS uses special microcode and hardware to draw Gouraud-shaded, z-buffered images rapidly. The Geometry Engines provide the capability to manipulate splines in real-time. The surface design and rendering demonstration package shows off these features. It allows you to modify the basis points that determine a curved, wireframe surface in one window and pass them to another window where a smooth-shaded, z-buffered rendering of the surface is made. The package consists of two programs: a surface editor and a rendering program.

Surface Editor

Start the surface editor window by using any of these commands: *surfcar*, *surfegg*, *surfabstract* or *surfjet*. Any of these commands runs the same program with different sets of basis points. Before you can interact with the surface editor, select the window for input.

The surface editor has two parts. The upper part of the display is the wireframe surface and the control points that define it. The bottom of the window is a menu of seven screen iconic buttons.

The simplest use of the surface editor allows one to edit the displayed surface by manipulating its basis points. To do this:

1. Move the cursor (now shaped like the mouse) clear of the window.
2. Press the mouse buttons one at a time and observe the glyphs. LEFTMOUSE gives a stylized arrow in the x direction, MIDDLEMOUSE gives one in the y direction, and RIGHTMOUSE gives one in the z direction.
3. Move the cursor into the top half of the display. Pick a basis point with the mouse, select a direction to move it, push the corresponding mouse button, and move the mouse from left to right to move the point. As the basis point is moved, the portion of the surface directly affected by the point is modified in real-time.
4. Now look at the menu which is lower in the display. The leftmost two screen buttons are for viewing the surface. The rightmost five screen buttons are for editing the surface. The leftmost screen button, the one showing three mutually perpendicular axes, is for rotating the surface 90 degrees around a particular axis. Move the cursor over the leftmost screen button and press and release LEFTMOUSE. The surface rotates 90 degrees around the x -axis.

5. Press MIDDLEMOUSE while the cursor is over the same button. The surface rotates around the y -axis. Pressing the RIGHTMOUSE causes the same motion around the z -axis. Do not try to start a second rotation until the first one is complete and the hourglass cursor has been replaced by the arrow.
6. Now select the second screen button, the bloodshot eye. Press and release LEFTMOUSE. Mouse motions now rotate the surface. Press RIGHTMOUSE to exit this mode and return the surface to its original spot.
7. Pick the eye icon button again and press and release RIGHTMOUSE. Now the surface, along with the array of index planes, rotates as guided by the mouse. RIGHTMOUSE gets you out of the rotate mode and leaves the display at its new position.
8. Use the third and fourth buttons (counting from left to right) to further subdivide the basis points that comprise the surface into twice as many basis points as before. More basis points allow for more precise editing of the surface, but also increase the complexity of the model and therefore slow it down.
9. Use the two rightmost screen buttons to make more modifications to the surface. The fifth screen button is unused. Each of the two rightmost buttons, the cubes with slices down the middle, selects a plane of symmetry (in a different axis) across the surface. When a plane is chosen (by pressing a mouse button as the cursor lies over one of the iconic buttons), half of the wireframe surface is made invisible. After a plane of symmetry has been chosen, any changes to a basis point on the visible side is also made to a corresponding point on the invisible side.

To make both sides visible again, move the cursor over the button corresponding to the present axis of symmetry and press a mouse button.

Rendering Program

After following the procedures for manipulating the basis points of an object as outlined above, you can use the following procedure to render the surface as a Gouraud-shaded, z -buffered image in another window.

1. Move the cursor over the second menu button, the bloodshot eye. Press and release RIGHTMOUSE. Previously you have used this mode to rotate the surface and the index planes with the mouse.
2. Press and release LEFTMOUSE. The cursor changes shape to prompt you to anchor this newly created window to the screen.

A wireframe outlining the basis points appears. The database of basis points is sent to the new window. The surface is subdivided into small quadrilaterals and an illumination model is calculated. The surface is displayed as an array of dots.

3. Detach from the surface editor window and select the rendering window for input. Rotate the object by pressing the mouse button corresponding to the axis of rotation and moving the mouse from left to right. Because the operation is double-buffered, the motion seems smooth.
4. To switch the display to splines, type **[s]**. To return to the point array, type **[p]**. To turn depth-cuing off or on, type **[d]**. Use **[z]** to Gouraud-shade the object.

When you switch from the point or spline array to the shaded picture, the system switches itself from double to single buffering to accommodate z-buffering. Z-buffering is used for general-purpose hidden-surface removal.

5. Type **[c]** to select a new color ramp to smooth shade all objects on the screen. Another way to change the surface color is to bring up the *cedit*, *showmap* and *interp* graphics windows and use them to alter the color ramp that shades the surface.

The rendering program can also be executed directly without going through the surface editor stage to modify the database. To do this, type: *zshadecar*, *zshadeegg*, *zshadeabstract* or *zshadejet*.

NOTES: Hardware z-buffering is not available on machines that do not have 32 bits of image memory per pixel. Do not try to use the rendering window without a 32 bitplane system.

If a double-buffered program not set up to work with the rendering demonstration programs is currently open, the rendering program does not shade the object. It prints a warning message in the console window. Kill the outstanding double-buffer windows and try again to shade the object.

9.2 Gifts

The contents of the directory */usr/people/gifts* varies with the current software distribution. Gifts are provided as interesting and possibly useful examples, but they are not supported and are often not documented.

At the time of this printing, */usr/people/gifts* contained programming examples, window manager demonstration programs, and sample image files. Some of the programming examples are discussed in the "Programming Examples" section of the *IRIS User's Guide*. Some of the window manager tools are documented in the *UNIX Programmer's Manual, Volume IA*. Manual pages for window manager commands are labeled 1W.

To determine what is currently available in */usr/people/gifts*:

1. Change your working directory to the directory containing the programs:

```
cd /usr/people/gifts
```

2. Display a listing of the directories contained in this directory:

```
ls
```


Appendix A: IRIS Specifications

The standard configuration for IRIS series 3000 systems includes these components:

- Electronics cabinet with one disk drive
- Fifteen-inch or nineteen-inch tilt-and-swivel 60 Hz non-interlaced monitor
- Keyboard with right-hand and left-hand connectors for the mouse
- Optical mouse with grid pad
- Ethernet transceiver

The size specifications for each of these components is listed in Table A-1.

IRIS series 3000 systems can also be ordered with these hardware options:

- Dial and button box
- Nineteen-inch 33 Hz interlaced monitor
- Digitizer tablet
- Tektronix 4692 printer
- Half-inch tape drive
- Quarter-inch cartridge tape drive
- Stereo optic viewer cable
- Floppy disk (standard on IRIS terminal)

Early IRIS series 2000 models with the Turbo upgrade may not have the same standard peripherals described here. The 60 Hz monitor may be different from the one described in this guide. The mouse may be mechanical rather than optical. The early IRIS 2400 Turbo system required a separate junction box. For a description of these features, see the *IRIS Workstation Guide Series 2000, Version 1.0*.

See Chapter 7 for more information on optional peripherals.

Component	Height	Width	Depth
Cabinet	30.0"	18.0"	27.0"
15" Monitor	14.2"	14.4"	15.9"
19" Monitor	21.9"	19.9"	21.3"
Keyboard	1.5"	19.0"	8.5"
Mouse	1.0"	2.0"	3.0"
Transceiver	2.0"	7.0"	4.0"

Table A-1: IRIS Component Specifications

A.1 Cables

The IRIS is supplied with the following cables for connecting components:

- Four bundled, color-coded, coaxial video cables connect the video output of the cabinet to the monitor.
- Two 10-foot, 3-wire, grounded AC power cables provide power for the IRIS monitor and cabinet.

A.2 Monitor

The IRIS can drive several types of monitors. See Chapter 7, "Video Options", for a discussion of all possible video options. The monitors available from Silicon Graphics are high-resolution 60 Hz non-interlaced or 33 Hz interlaced. A single system can support two video options. The video display configuration for each system is specified when the system is ordered.

A.2.1 60 Hz Non-interlaced Monitors

The 60 Hz non-interlaced monitors have a monitor control panel on the lower right front. On the back of the monitors are several ports for receiving video signals and a power socket. These monitors are available in two sizes, nineteen-inch and fifteen-inch. (See Figures A-1 and A-2.)

A.2.2 33 Hz Interlaced Monitor

A 33 Hz interlaced monitor is an optional monitor for the IRIS system. This monitor has a monitor control panel on the right front. On the back of the monitor are several ports for receiving video signals and a power socket. (See Figure A-3).

A.2.3 Monitor Control Panel

The monitor control panel has adjustment controls and a **Power** switch .

Nineteen-inch Monitor (60 Hz Non-interlaced)

The monitor control panel on this monitor is shown in Figure 2-7.

- The button labeled **Brightness** adjusts the DC levels of the red, green, and blue signals equally. Turning this button clockwise increases the monitor's brightness. You can push the button into the monitor once the brightness is adjusted.
- The button labeled **Contrast** adjusts the AC gain levels of the red, green, and blue signals. Turning this dial clockwise increases the monitor's contrast. Once the contrast is adjusted, you can push the button in so that it is flush with the surface of the monitor.
- The button labeled **Degauss** demagnetizes the monitor screen.
- The rocker switch labeled **Power** controls power to the monitor.
- The power light indicates that the monitor power is on.

Fifteen-inch Monitor (60 Hz Non-interlaced)

The monitor control panel on this monitor is shown in Figure 2-8.

- The recessed dial labeled **Brightness** adjusts the DC levels of the red, green, and blue signals equally. Turning this knob clockwise increases the monitor's brightness.
- The recessed dial labeled **Contrast** adjusts the AC gain levels of the red, green, and blue signals. Turning this dial clockwise increases the monitor's contrast.
- The button labeled **Degauss** demagnetizes the monitor screen.
- The button labeled **Power** controls power to the monitor.
- The power light indicates that the monitor power is on.

33 Hz Interlaced Monitor

The monitor control panel on the 33 Hz interlaced monitor is shown in Figure 2-9.

- The knob labeled **Brightness** adjusts the DC levels of the red, green, and blue signals equally. Turning this knob clockwise increases the monitor's brightness.
- The knob labeled **Contrast** adjusts the AC gain levels of the red, green, and blue signals. Turning this knob clockwise increases the monitor's contrast.
- The button labeled **Degauss** demagnetizes the monitor screen.
- The switch labeled **Power** controls power to the monitor.

- The light labeled **Health** lights when power to the monitor is switched on.

A.2.4 Monitor Back Panel

The monitor back panel has several connectors for the cables that connect the monitor to the cabinet (see Figures A-1, A-2, and A-3).

Nineteen-inch Monitor (60 Hz Interlaced)

- Either of the two BNC sockets labeled **Sync** receives the video sync signal from the cabinet.
- Either of the two BNC sockets labeled **R**, **G**, and **B** receives the red, green, or blue video signal from the cabinet.
- The fuse protects the circuit. See the information printed next to the fuse for type and rating requirements. (See Figure A-1.)
- The power selector switch selects the power voltage.
- The power input receptacle attaches to the power cable. See the information printed next to the receptacle for power requirements. (See Figure A-1.)
- The four 75 Ω terminators attached to the monitor by chains control the impedance of the video signals. When a single monitor is used, the terminators plug into the unused BNC sockets. When more than one monitor is used in a daisy-chain, the terminators are removed from the sockets on all but the last monitor.

Fifteen-inch Monitor (60 Hz Non-interlaced)

- Either of the two BNC sockets labeled **Sync** receives the video sync signal from the cabinet.
- Either of the two BNC sockets labeled **R**, **G**, and **B** receives the red, green, or blue video signal from the cabinet.
- The fuse protects the circuit. See the information printed next to the fuse for type and rating requirements. (See Figure A-2.)
- The power selector switch selects the power voltage. See the information printed on the switch for power requirements. (See Figure A-2.)
- The power receptacle attaches to the power cable.
- The four impedance switches control the impedance of the video signals. When the switch is **IN**, it is in the **75 Ω** position. When the switch is **OUT**, it is in the **High** position.

33 Hz Interlaced Monitor

- Either of the two BNC sockets labeled **Ext Sync** receives the video sync signal from the cabinet.
- The two BNC sockets labeled **V D** are not used.
- Either of the two BNC sockets labeled **R**, **G**, and **B** receives the red, green, or blue video signal from the cabinet.
- The circuit protector is either a 5-amp fuse (100-125 VAC) or a 2-amp fuse. (200-240 VAC).
- The 120/240 volt power receptacle connects to the power cable.
- The 25-pin plug connects the control cable from the cabinet to the monitor.
- The four impedance switches control the impedance of the video signals.




Figure A-1: Monitor Back Panel for Nineteen-inch 60 Hz Monitor




Figure A-2: Monitor Back Panel for Fifteen-inch 60 Hz Monitor




Figure A-3: Monitor Back Panel for 33 Hz Interlaced Monitor

A.3 IRIS Electronics Cabinet

There are three panels on the back of the IRIS cabinet: a standard I/O panel, an auxiliary I/O panel, and a power panel (see Figure A-4). On the front of the IRIS is a cartridge tape drive (optional) and a **Power** switch that controls power for the IRIS system.

A.3.1 Power Switch

The **Power** switch for the IRIS is located on the front upper left corner of the cabinet. This switch controls power for the cabinet. It does not control the power for any auxiliary equipment connected to the cabinet through the convenience outlet located on the cabinet power panel.

A.3.2 Optional Tape Drives

Both a 60 Mb quarter-inch cartridge tape drive and a half-inch tape drive are offered as optional components on the IRIS.

Quarter-inch Cartridge Tape Drive

The cartridge tape drive is located on the front upper left of the cabinet. See Chapter 7, "Optional Peripherals", for information on using the tape drive.

Half-inch Tape Drive

The half-inch tape drive is housed in an enclosure that is separate from the electronics cabinet.

A.3.3 Standard I/O Panel

The standard I/O panel is located on the upper right rear of the cabinet (see Figure A-4). This panel has ports for connecting the monitor, a network drop cable, a floppy disk drive, and various RS-232 devices.

- The RS-232 connector labeled **Port 1** attaches to the keyboard.
- The RS-232 connectors labeled **Port 2**, **Port 3**, and **Port 4** are available for RS-232 or RS-423 serial lines.
- The 15-pin D socket labeled **Ethernet** connects the IRIS to an Ethernet drop cable.
- The BNC socket labeled **Sync** provides the video sync signal for the monitor.
- The three BNC sockets labeled **Red**, **Green**, and **Blue** provide the red, green, and blue video signals for the monitor.

- The **Reset** button resets the processor, which in turn resets the rest of the system. After the **Reset** button has been pressed, the IRIS either reboots automatically or waits for boot instructions.

CAUTION

Do not press the **Reset** button while the IRIS is running UNIX. For information on rebooting the system, see Chapter 3. If the IRIS is not running UNIX and is under control of the PROM monitor, then the **Reset** button or the **Power** switch may be used. See the discussion on Crash Recovery in Chapter 4.

- The LED labeled **Halt** lights when the processor is in a halt state.
- The alphanumeric diagnostic LED labeled **Status** indicates system status and displays startup diagnostics.
- The nine-element DIP switch labeled **Configuration** controls the IRIS' startup diagnostics, and the boot environment.

A.3.4 Cabinet Power Panel

The cabinet power panel has a power inlet receptacle (see Figure A-4).

- The male 3-pin input power receptacle labeled **Power** accepts power for the IRIS system.
- For operation in the US, a 20 amp, 250V, 3AB normal-blow ceramic tube fuse protects the circuit. For European operation, a 6.3 amp, 250V, type T slow-blow fuse protects the circuit.




Figure A-4: IRIS Cabinet Back Panel

A.4 Site Selection

Table A-2 contains a list of guidelines for site selection for your IRIS. Although site selection is the customer's responsibility, Silicon Graphics' representatives will provide consulting services upon request.

Temperature	50 — 86°F (operating) 32 — 122°F (non-operating)
Relative humidity	20 — 80%
Minimum clearance	3" all sides
Monitor desk space	
Nineteen-inch Hz	30" width 30" depth
Fifteen-inch	22" width 22" depth
Cabinet space	24" width 33" depth 33" height
System Power Requirement	single phase line, neutral, ground
115 VAC Configuration	
Operating Range	93-132 V AC 47-63 Hz
220 VAC Configuration	
Operating Range	186-264 V AC 47-63 Hz
System Power Consumption	Cabinet 1250 VA / 1000 W Nineteen-inch 60 Hz Monitor 225 VA / 150 W Fifteen-inch Monitor 185 VA / 100 W Nineteen-inch 33 Hz Monitor 170 VA / 116 W
Heat dissipation	Cabinet 3412 BTU/hr Nineteen-inch 60 Hz Monitor 512 BTU/hr Fifteen-inch Monitor 341 BTU/hr Nineteen-inch 33 Hz Monitor 395 BTU/hr
Card Slots	20

Table A-2: IRIS Environmental Specifications

Appendix B: System Messages

When the UNIX kernel on an IRIS workstation reaches an unrecoverable error condition, it displays an error message preceded by the word `panic`:

If the error message includes one of the abbreviations listed in Table B-1, the condition is probably caused by a hardware problem. Table B-1 lists the devices associated with each abbreviation. The abbreviated device names are usually followed by a digit indicating a physical unit or letters referencing the function that lead to the problem.

Abbreviation	Device
si	IRIS 3030 disk drive
mt	tape drive (any type)
qic	quarter-inch cartridge tape
md	IRIS 3010/3020 hard disk drive
mf	floppy disk drive
dsd	tape, small disk, and floppy disk controller
nx	Ethernet

Table B-1: Device Abbreviations

If the panic message does not indicate one of the devices listed in Table B-1, there may be a hardware problem with another part of the system (processor board, memory board, or backplane, for example) or a software problem in the kernel.

Anytime you get a panic message, try to run `sync` and reboot (see Section 3.1 or 3.2). Run `fsck` to check the file system. If the message reappears, call the Geometry Hotline (see Chapter 1).

NOTE: Do not interrupt (e.g., by pressing `Ctrl-c`) processes that are accessing the no-rewind tape device. Attempts to do this always result in a kernel panic.

B.1 Hardware Error Messages

B.1.1 General Hardware Error Messages

- **IO err in swap**

While swapping a user process, a hard error occurred on the swap disk.

- **parity error**

A parity error occurred in the system memory. Another message preceding this error message indicates where in physical memory the error occurred. Unix cannot diagnose the memory failure.

- **init**

The system was unable to read the root file system. Either the disk drive or the root file system is damaged.

B.1.2 Disk Controller Messages

IRIS 3010/3020 Disk Controller Messages

A. The following messages result from a bad disk controller.

- **md[0-1][a-h]: couldn't start!**
- **dsd: zero status on command**
- **dsd: ccb timeout, "command" dev=0 unit=[0-1]**
- **dsd: ccb timeout during init**
- **dsd: cib timeout, "command" dev=0 unit=[0-1]**

B. This message indicates that no label is installed on the disk.

- **(**No label**)**

C. The following error probably results from a bad connection from the disk to the controller, a bad disk, or hard disk errors encountered when one reads or writes to the disk.

- **dsd0: hard error, "command" dev=0 unit=[0-1]**

IRIS 3030 Disk Controller Messages

A. This message indicates that no label is installed on the disk.

- **(**No label**)**

B. The following messages indicate that a hard error has occurred and recovery may not be possible.

- **"Read" or "Write" on si[0-4], slice [a-h]**

- **siintr: sitab.b_active == 0**
 - **siintr hard error NN: "error" block: NN cmd: command**
- C. The following messages occur in non-interrupt mode and could mean a faulty controller, a bad set of disk cables, or incorrect jumpering of the controller.
- **sicmd: timeout wait for status**
 - **sicmd: timeout waiting for cmd "command" to complete**
 - **sicmd: status error: "error"**
 - **sicmd: error stat "status" err "error" iostat "iostat"**
- D. The following error indicates a recoverable problem with the disk. This message may appear when there is a bad track on the disk. If this message appears repeatedly, contact Customer Service through the Geometry Hotline (see Chapter 1).
- **soft error "error" on block NN**

B.1.3 Tape Drive Messages

- A. The following are general messages for the tape drive.
- **qic0: no cartridge in drive**
Cartridge isn't inserted properly.
 - **qic0: unit not ready**
 - **qic0: write protected**
Dial on the cartridge needs to be twisted to enable writing on the tape.
- B. The following error message indicates that the tape drive has no power or that the control cable for the tape drive is disconnected or reversed.
- **qic0: can't get status during init**
- C. The following error messages indicate that there is a problem with the tape drive that inhibits its operation. The operation using the tape drive must be restarted.
- **qic0: couldn't start**
 - **qic0: hard error during "Read or Write" status= XX**
 - **qic0: "command" failed**

8.1.4 Floppy Controller Messages

- A. The following are general messages.
- **mf0: write protected**
 - **mf0: disk formatted**
 - **Reading on mf0**
 - **Writing on mf0**
- B. The following message indicates a problem with the floppy disk controller.
- **mf0a: hard error, cmd='command',error='error'**

B.2 System Software Messages

- A. The following error messages indicate a problem with the Ethernet hardware or software.
- **nxpresent cleared**
 - **xns_ttstart**
- B. The following message occurs while the system is booting. This message indicates that the system attempted to configure a disk drive with no entry in the *bdevsw[]* array (the array of block devices). The system was incorrectly configured.
- **getmajor**
- C. The following messages indicate that there is a problem with the buffer cache/inode tables.
- **devtab**
 - **no fs**
 - **no imt**
- D. The following message indicates that the *init* process was killed. If the process was killed through a user program or the shell *kill* command, then this message may not reflect a problem. If the *init* process was killed during the boot procedure, then this message indicates a problem with the root file system.
- **init died!**

- E. The following message occurs when the system attempts to put a time-driven event on its queue, and there is no room remaining. If this occurs frequently, then the system has been incorrectly configured.
- **timeout table overflow.**
- F. The following messages indicate problems with the kernel.
- **trap**
 - **kernel address error**
 - **kernel bus error**
 - **ksyscall**
- G. The following messages are related to graphics consistency checks.
- **duplicate gr_getshmem**
 - **gr_unlockmem**
- H. The following messages indicate problems occurring while trying to swap some portion of a user process to disk.
- **IO err in push**
 - **hard IO err in swap**
- I. The following message occurs during the autoconfiguration process if no disk was found that supports a swap area.
- **no swap device.**
- J. The following message indicates that all available swap space has been used.
- **swkill.**
- K. Error messages not listed above are related to internal kernel consistency checks.